

STATISTICAL SCALING ANALYSIS OF TCP/IP DATA USING CASCADES

S. Roux¹, D. Veitch², P. Abry¹,

(1) CNRS, UMR 5672, ENS Lyon, France
{sroux,pabry,flandrin}@ens-lyon.fr
{d.veitch,l.huang}@emulab.ee.mu.au

L. Huang², J. Micheel³, P. Flandrin¹

(2) EMUlab, Dept. of E&EE, Univ. Melbourne, Australia
(3) WAND, Univ., Waikato, New Zealand
joerg@cs.waikato.ac.nz

ABSTRACT

The scaling properties of Internet data are analysed in detail through the unifying viewpoint of *Infinitely Divisible Cascades*. From exceptionally precise TCP/IP traffic traces are extracted time series including arrival rate, durations, and interarrival times of TCP connections. We show that IDC's offer a pertinent description of these series. Relations between them are investigated, yielding insights on the sources of the scaling and possible modelling approaches.

1. MOTIVATION

The performance of packet networks, ultimately the losses and delays experienced by packets, is strongly dependent on the nature of the traffic carried. Not only mean rates, but the detailed structure of the traffic flow, its *burstiness*, has a major impact. This structure is very rich, and it is well accepted [6] that *scaling* properties and models are the natural language to describe it: traffic has fractal features.

In this paper we consider TCP/IP traffic, the dominant traffic protocol in many networks including the Internet. The existence of Long-Range Dependence (LRD) in such traffic is now well established [6], and recent work has shown the relevance of multifractal models [5, 12]. A preliminary study [13] showed that the more general *Infinitely Divisible Cascade* (IDC) framework can lead to further insights. Multiscaling, multifractality and hence exact self-similarity are all special cases of Infinitely Divisible Cascades. We expand on this work by looking at new, more accurate and longer data sets, a greater range of time series derived from them, and use a richer set of statistical methods. From the wavelet coefficients of the time series, we use tools we have developed to detect the presence of scaling and the corresponding range(s) of scales over which it exists, and second, to estimate the corresponding cascade parameters. Particular attention is paid to discriminating non stationarity and scaling phenomena, and related estimation issues.

The nature of traffic is evolving rapidly, however we believe that detailed scaling behaviour cannot arise by chance but reflects the presence of robust natural underlying mechanisms. On large time scales the fact of *heavy tailed* file sizes is one such mechanism for the generation of LRD. Over

Partially supported by the CNRS grant TL99035, Programme Télécommunications and French MENRT grant ACI jeune Chercheur 2329. The support of Ericsson Australia is gratefully acknowledged.

small time scales, recent work [4] suggests several contributing factors for the source of multifractal behaviour, however many questions remain unanswered. It is beyond the scope of this paper to enter deeply into the network origins of complex scaling, however we go beyond the aims of [13] by investigating relationships between the scaling found in different time series, in order to isolate the possible sources at the statistical level, and inform parsimonious modelling.

2. DATA

We use exceptionally precise TCP/IP data made available by the WAND group at the University of Waikato. This archive, the 'Auckland II' traces, are taken from both directions of the access link of the University of Auckland to the external Internet. The capture hardware developed at WAND (measuring ATM technology at 155 Mbits/s for Auckland II) is capable of loss-less measurement with synchronized timestamps accurate to below 1 μ s. Refer to <http://wand.cs.waikato.ac.nz/wand/wits/index.html> and [9] for full details. The two traces analysed here are described briefly in the table below. Packets consist of a *header*, where

| Trace | length used | begins: (h:m:s) | # IP pkts |
|-------|-------------|-----------------|--------------|
| Nov29 | 20:44:24 | 13:42:58 1999 | 31.3 million |
| Feb3 | 05:15:45 | 11:29:52 2000 | 14 million |

addressing and control information is kept, and a *payload*, the actual data to be transmitted. Internet Protocol (IP) packets carry payloads across heterogeneous networks with no mechanism for reliable delivery, and no knowledge of the whole of which they are a part. The User Datagram Protocol (UDP) uses this basic service to transmit blocks of data in one direction to an end process. Transmission Control Protocol (TCP) packets are also carried in IP payloads. Information in the TCP headers, together with state information in the TCP/IP program stack running on end computers, constitutes a two-way TCP *connection* (TCPC) between end points which provides higher level services such as retransmission of lost packets, and flow control. Web sessions involve many TCPCs, ensuring the error free delivery of text, images and other data.

From raw data the number of TCP payload bytes and packets in bins of size $\delta = 40$ ms were extracted to form the IP-level time series IPbytTCP and IPpktTCP, and similarly IPpktUDP and IPpktUDP count UDP payloads. TCP-level time series require the tracking of packets belonging to indi-

vidual TCPCs, a non-trivial task. For example many connections begin but never terminate. At both IP and TCP levels we consider only ‘complete’ connections which begin and end correctly, and which transmit at least *some* data. In our traces these constitute $\sim 86\%$ both of total connections and bytes, and $\sim 90\%$ of packets. Other connections were examined but found not to have interesting scaling properties. At the TCP level we study the following series indexed by bin number: Arr, Dep, Act are respectively the numbers per bin of TCPCs arriving, departing and active. Other series are indexed by arrival order: Iar, Dur, Pkt, Byt are respectively inter-arrival times, durations in ms, and the number of packets and bytes for each TCPC.

3. ANALYSIS TOOLS

The scaling properties of the data are studied through the analysis of the scale dependencies of their wavelet coefficients. Let X denote any of the time series introduced above and let $\{T_X(a, t) = \langle X, \psi_{a,t} \rangle\}$ denote its wavelet coefficients, where the $\psi_{a,t}(u) = a^{-1}\psi(a^{-1}(u-t))$ are dilated and translated templates of a reference pattern ψ . For details on wavelet transforms, see e.g., [7].

The analysis tool used to study scaling is the so-called infinitely divisible cascade model. An IDC relates the probability density function (pdf) of the log of the wavelet coefficients $h_X(a, t) = \ln |T_X(a, t)|$ at scale a to that at scale a' by a convolution kernel, called the *propagator*:

$$p_a(h) = \int G_{a,a'}(h-h')p_{a'}(h')dh', \quad (1)$$

with the key assumption that the Laplace transform $\widehat{G}_{a,a'}(q)$ of $G_{a,a'}(h)$ is separable in the variables a and q :

$$\ln |\widehat{G}_{a,a'}(q)| = H(q)(n(a) - n(a')). \quad (2)$$

This definition, together with the fact that

$$\ln \mathbb{E}|T_X(a, t)|^q = \ln \mathbb{E} \exp(q \ln |T_X(a, t)|) = \sum_n c_{a,n} q^n / n!,$$

where the $c_{a,n}$ are the cumulants of $h_X(a, t)$, yields

$$\begin{aligned} \ln \mathbb{E}|T_X(a, t)|^q &= H(q)(n(a) - n(a')) + \ln \mathbb{E}|T_X(a', t)|^q \\ c_{a,n} - c_{a',n} &= C_n(n(a) - n(a')) \\ \ln \mathbb{E}|T_X(a, t)|^q &= H(q)/H(p) \ln \mathbb{E}|T_X(a, t)|^p + K_{q,p}(a') \\ c_{a,n} &= C_n/C_m c_{a,m} + \beta_{n,m}(a') \\ H(q) &= \sum_n C_n q^n / n! \quad , \end{aligned} \quad (3)$$

the central consequences of the propagator, expressed in equivalent moment or cumulant form. Note that $H(q)$ and $n(a)$ are defined up to a multiplicative constant and a multiplicative and additive constant, respectively. For further details on IDCs the reader is referred to e.g., [2, 13].

If $n(a) \equiv \ln a$, (3) implies that the $\mathbb{E}|T_X(a, t)|^q$ behave as power laws of scale. If such behaviour holds in the limit of small scales, the IDC reduces to a multifractal analysis [11, 13], where the function $H(q)$ is deeply related to the (ir)regularity or high variability of the sample paths of X . Rewriting $H(q) = C_1(q + C_2/(2C_1)q^2 + \dots)$, one sees

$C_2/(2C_1)$ as a measure of the departure of $H(q)$ from linearity and therefore of multi- versus mono-fractality. In the even more specific case where $n(a) = \ln(a)$ is valid for all scales, $H(q) = qH$, and X is a H -self-similar process [6, 1], and $C_1 \equiv H$, $C_2 = 0$. LRD is another special case, specific to $q = 2$ and $H(2) + 1 \in (1/2, 1]$, where a power-law behaviour exists at large scales [6, 1], implying that the past history of X has a non-negligible influence. Note that this means that the traditional Hurst exponent is $H = H(2) + 1 \simeq C_1 + 1$ (since C_2 will be small).

When analysing data, the initial questions are i) can we identify a range of scales where the IDC model applies, ii) can we estimate the corresponding $H(q)$ (or equivalently the C_n) and $n(a)$ functions? The favourable statistical properties of the wavelet coefficients of scaling processes [1] indicate that the ensemble averages $\mathbb{E}|T_X(a, t)|^q$ can be efficiently estimated by the time averages $S_q(j) = 1/n_j \sum_k |T_X(a = 2^j, 2^j k)|^q$ (n_j is the number of coefficients available at octave j). This technique has the advantage of being insensitive to non-stationarities in signal variance. Because the wavelet coefficients are by nature centered about 0, the alternative of estimating the cumulants of their logarithm is a delicate task, but desirable due to their close relationship with the IDC structure. To overcome the estimation difficulties, we use the so-called wavelet transform modulus maxima (WTMM) technique [8] which performs time averages on the local maxima only. The fit to the data is tested by checking the affine behaviours given in (3), mainly those of $\log_2 S_q(j)$ vs $\log_2 S_p(j)$. These same relations yield estimators for $H(q)$ or the C_n 's, and the $n(a)$. See [13, 2] for complete definitions, discussions and algorithms. Finally, note that since the C_n are small, their reliable estimation becomes increasingly difficult as n grows. We therefore restrict ourselves to the estimation of the approximation $H(q)n(a) = C_1(q + C_2/(2C_1)q^2)(\ln a)$.

4. DATA ANALYSIS

At the TCP level, for the Iar, Arr, Dep and Act time series, the results of the IDC analysis can be summarized as follows (for space reasons only the Iar series appears in figure 1). The $\log_2 S_q(j)$ vs $\log_2(2^j)$ plots (top left) show two scaling ranges: coarse scales (CS) and fine scales (FS), about a change point j_* . In each range linear behaviour is observed, indicating power law evolution of the $S_q(j)$ with scale $a = 2^j$. The $\log_2 S_q(j)$ vs $\log_2 S_p(j)$ plots, for $p, q \in [0, 5]$, (top right) have close to affine form over almost the entire range of available octaves and the estimation procedures show (middle left) that $n(a)$ is a piece-wise log function with a knee around j_* , in agreement with the $\log_2 S_q(j)$ vs j plots. Figure 2 shows the Gaussian like pdf of the $h_X(j, k)$ for Iar over various scales (left plot) and the effect of normalising them through the estimated propagator $H(q)n(a)$ (right). The fact that the pdf's collapse onto each other is a clear visual validation of the IDC. Estimation procedures for $H(q)$ (middle right in figure 1) allow us to claim that a small yet significant departure from the simplest linear behaviour is observed. The table below gives estimated C_1 , C_2 and j_* values for each times series for both the coarse (CS) and fine (FS) scales.

For the coarse scales, the estimated C_1 values together

| | FS | j_* | CS | FS | j_* | CS |
|------------|----------|-------|-------|----------|-------|-------|
| | Iar | | | Act | | |
| C_1 | -0.32 | 6 | -0.20 | 0.48 | 8 | 0.43 |
| $C_2/2C_1$ | 0.03 | | 0.02 | -0.02 | | -0.02 |
| | Dep | | | Arr | | |
| C_1 | -0.54 | 8 | -0.25 | -0.54 | 8 | -0.25 |
| $C_2/2C_1$ | -0.02 | | -0.04 | -0.02 | | -0.04 |
| | IPpktTCP | | | IPpktUDP | | |
| C_1 | -0.32 | 8 | -0.08 | -0.42 | 8 | -0.20 |
| $C_2/2C_1$ | -0.05 | | -0.19 | 0.06 | | -0.12 |

with the fact that $n(a) \propto \ln(a)$ indicates LRD with parameter $H \simeq 0.8 \pm 0.05$ for Iar, Arr and Dep while Act is H -self similar: $H \simeq 0.45 \pm 0.05$. The C_1 and C_2 estimates for the fine scales together with the fact that $n(a) \propto \ln(a)$ can be interpreted as multifractal behaviour, as concluded in [13]. Rather than trying to infer multifractal spectra, a difficult task given the poor precision of C_2 , we emphasise that it is mainly because the C_1 's at fine scales differ and are smaller than those at large scales that accounts for the high variability of the sample paths, rather than mono vs multi-fractality as such.

Of particular interest is that all the measured scaling behaviour, as well as first order statistics and marginal distributions, yield astonishingly similar results for Arr and Dep, indicating strong dependence. Moreover, the C_1 's at fine and coarse scales for Act are much closer to each other than for the other time series, indicating that the IDC is almost scale invariant (i.e., $n(a) = \ln a, \forall a$). Since Act results from combining Arr and Dep, this again suggests strong dependence between the two.

When applied to Dur, the scaling analysis machinery leads to observations qualitatively similar to those above. However, the suspicious parameter values obtained indicated the possibility that this was an artifact generated by infinite moments in the series, as discussed in [3]. A careful analysis based both on arguments in [3] and estimates of marginals confirmed this, revealing that the structure of Dur is that of weakly dependent random variables, with heavy power law tails of infinite variance and possibly no mean. The series Byt and Pkt showed not only remarkably similar properties to Dur, both scaling and otherwise, but considerable dependence at the sample path level.

Analysis of the four IP level series led to similar observations: piecewise log IDC for the full range of scales with LRD over the coarsest scales and multifractality at the finest. Furthermore for each of the UDP and TCP protocols the IPpkt and IPbyt series were very closely related. However, the estimates of the C_1 (this includes the LRD exponent) and C_2 differed significantly between UDP ($H \simeq 0.8$) and TCP ($H \simeq 0.92$) (see table above). Examining the incoming and outgoing traffic separately yielded identical conclusions.

It is immediately obvious from low resolution plots (see the WAND web page), that the traffic is non-stationary, for example there is a clear diurnal cycle in both load (mean) and variance, and to a lesser extent there are changes in scaling properties also. The results presented here were obtained from subsets of the time series where "stationarity

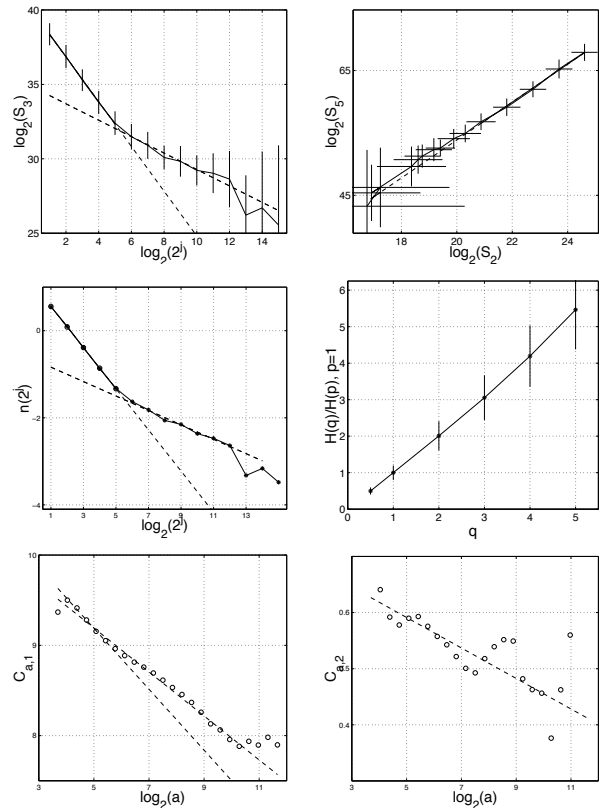


Figure 1: **Scaling for IAR time series.** Top: evidence for IDC; Middle: measurement of its propagator from the $S_q(j)$; Bottom: measurement of its propagator from the $C_{a,n}$.

of the scaling properties" was observed, that is where estimation on even finer subsets yielded compatible conclusions and parameter values. To investigate dependencies between load and scaling behaviour, the data was split into subsets of *low*, *medium* and *high* load level. The scaling properties at CS did not significantly depend on the load level, as opposed to that at FS where the C_1 and C_2 estimates over adjacent blocks (Figure 3) closely reproduced the time evolution of the average load and power. Further work is needed to see if this is a true dependency or due to high estimation variance at low load. The results presented above were for high or medium load. We also noted in several series that the variance to mean ratio was approximately constant over time, consistent with the superposition of a varying number of processes with similar characteristics.

To conclude, IDC models provide a good description of the scaling of the data over almost the entire range of scales present. The main feature of the observed IDCs is that $n(a) \neq \ln(a)$, they are not scale invariant. Instead, the $n(a)$ are piecewise log with a change point around j_* . For each series, j_* corresponds to a characteristic time of 2.5 to 3.5 seconds, in keeping with findings in [4], and of our own measurements of *round trip times* of TCPC's. As $H(q)$ was always close to linear, the approximation $H(q) = C_1q + C_2q^2/2$ was used, suggesting the cumulant based es-

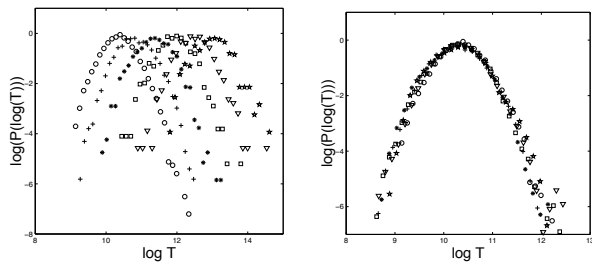


Figure 2: Pdf's of the log wavelet coefficients for Iar. Left: pdf's at octaves 6 to 11 from left to right; Right: collapsed pdf's using the estimated propagator $H(q)n(a)$.

timation method and focusing it to four parameters, j_* , the C_1 at fine and coarse scales, and $C_2/(2C_1)$. Key open questions include the origin and implications of the piecewise log behaviour. The great similarity between the Arr and Dep series, the i.i.d. heavy tailed nature of Dur and its close relation to Pkt and Byt, are also noteworthy.




Figure 3: Non stationarity. Time evolution of the mean of the data (o) and the C_1 estimates (+) (left plot) and of the variance of the data (o) and the C_2 estimates (+).

5. MODELLING OF DATA

Rather than attempting unrelated 'black box' models of each series separately, two hybrid models incorporating both empirical and modelled data are presented here which help understand underlying structural features of the data.

The great similarity between the Arr and Dep series suggests a coordination between the TCPC arrival process and connection durations which belies the latter's apparently simple i.i.d. structure. To investigate this mystery, actual arrival times were obtained from the Iar series, to which i.i.d. Pareto distributed durations were added with parameters fitted from Dur, to generate a set of surrogate departure times. These were binned to form a hybrid Dep series which shows remarkable similarities with the original. Notably, an IDC model applies with similar exponents to that of Dep, although j_* is one larger. This clearly shows that the similarities between Arr and Dep do not require a subtle interconnection. On the contrary, the structure of Dep is derived from that of Arr, and is somehow preserved from Arr even after a drastic random reordering.

The close connection between Byt and Dur suggests a very simple relationship: that often the large scale data rate of connections are independent of their duration. This

in turn suggests a very simple model for how IPbyTTCPC could be generated from Iar and Dur, namely that the total number of bytes in each connection be spread out evenly, a constant bit rate within each TCPC, and then added in bins across connections. Generating a surrogate IPbyTTCPC series in this way and comparing, we see a good correspondence at large scales, but utterly different behaviour at small scales. This strongly suggests that the rich structure of Iar is not sufficient to explain that of the final data stream, rather, the non-trivial burstiness within connections must be included, in keeping with conclusions from [4]. Similar results apply for IPpktTCP.

6. REFERENCES

- [1] P. Abry, P. Flandrin, M.S. Taqqu and D. Veitch, Wavelets for the analysis, estimation and synthesis of scaling data. A chapter in [10].
- [2] A. Arnéodo, J.F. Muzy, S.G. Roux, Experimental analysis of self-similar random cascade processes: application to fully developed turbulence, *J. Phys. II France*, 7:363–370, 1997.
- [3] P. Chainais, P. Abry, and D. Veitch, Multifractal Analysis and α Stable Processes: a Methodological Contribution. Proceedings of *ICASSP 2000*, Istanbul, June 2000.
- [4] A. Feldmann, A.C. Gilbert, Polly Huang, W. Willinger, Dynamics of IP traffic, A study of the role of variability and the impact of control, Proceedings of *ACM/Sigcomm'99*, Cambridge MA, August 1999.
- [5] A.C. Gilbert, W. Willinger, A. Feldmann, Scaling analysis of random cascades, with applications to network traffic, *IEEE Trans. Info. Theory*, Special Issue, 45(3):971–991, 1999.
- [6] 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. Networking*, 2:1–15, 1994.
- [7] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, Boston, 1997.
- [8] J.F. Muzy, E. Bacry, A. Arneodo, The multifractal formalism revisited with wavelets, *Int. Journal of Bifurcation and Chaos*, 4(2): 245–302, 1994.
- [9] Jörg Micheel, Ian Graham, N. Brownlee, *The Auckland data set: an access link observed*. to appear, Proceedings *14th ITC Specialist Seminar*, Barcelona, April 2000.
- [10] *Self-Similar Network Traffic and Performance Evaluation*. K. Park and W. Willinger, eds. Wiley, 2000.
- [11] R.H. Riedi. An improved multifractal formalism and self-similar measures, *J. Math. Anal. Appl.*, 189:462–490, 1995.
- [12] R. Riedi, M.S. Crouse, V.J. Ribeiro, R.G. Baraniuk, A Multifractal Wavelet Model with Application to Network Traffic, *IEEE Trans. on Info. Theory*, Special Issue, 45(3):992–1018, April, 1999.
- [13] D. Veitch, P. Abry, P. Flandrin and P. Chainais, Infinitely Divisible Cascade Analysis of Network Traffic Data. Proceedings of *ICASSP 2000*, Istanbul, June 2000.