

Cluster Processes: A Natural Language for Network Traffic

Nicolas Hohn, Darryl Veitch, *Senior Member, IEEE*, and Patrice Abry

Abstract—We introduce a new approach to the modeling of network traffic, consisting of a semi-experimental methodology combining models with data and a class of point processes (cluster models) to represent the process of packet arrivals in a physically meaningful way. Wavelets are used to examine second-order statistics, and particular attention is paid to the modeling of long-range dependence and to the question of scale invariance at small scales. We analyze in depth the properties of several large traces of packet data and determine unambiguously the influence of network variables such as the arrival patterns, durations, and volumes of transport control protocol (TCP) flows and internal flow structure. We show that session-level modeling is not relevant at the packet level. Our findings naturally suggest the use of cluster models. We define a class where TCP flows are directly modeled, and each model parameter has a direct meaning in network terms, allowing the model to be used to predict traffic properties as networks and traffic evolve. The class has the key advantage of being mathematically tractable, in particular, its spectrum is known and can be readily calculated, its wavelet spectrum deduced, interarrival distributions can be obtained, and it can be simulated in a straightforward way. The model reproduces the main second-order features, and results are compared against a simple black box point process alternative. Discrepancies with the model are discussed and explained, and enhancements are outlined. The elephant and mice view of traffic flows is revisited in the light of our findings.

Index Terms—Internet data, long-range dependence, multifractals, point processes, scaling, time series analysis, traffic modeling, wavelets.

I. INTRODUCTION

WE seek to model, and understand, the statistical nature of the flow of data packets passing through telecommunications links, such as high-speed links in the Internet “backbone.” By data packets, we mean *Internet protocol* (IP) packets, which are the universal medium of transport in the present-day Internet. For our purposes, the effect of the highly complex, layered structure of the network on data can be abstracted to the concept of *flow*. A flow is a set of packets that are part of an identifiable exchange between two end points; for example, they may carry the bytes of a file transfer between two computers (see

Section III-A for a technical definition). At a given measurement point in the interior of the network, packets from many thousands of intermingled flows pass, and individual flows are seen to begin, pass through bursty and idle phases, and end. Flows are highly variable, with durations ranging from less than a second to many hours, from just a single packet to billions [see Fig. 2(b) and (c)].

The set of arrival times of packets can be viewed as a point process on the real line. A central aim of traffic modeling is to be able to describe key features of this process, using parameters with direct and verifiable physical meaning in terms of the nature of traffic sources and the network’s transformations of them. This is important for network engineering because the degree and nature of traffic burstiness determines the properties of queuing delays (and losses) in switching devices and, thereby, the quality of the services delivered over the network.

Although many traffic models have been proposed to date (for point process examples, see [1] and [2]), none have been accepted as definitive. The complexity required to adequately describe the statistics of traffic is potentially very high. First, the structure of packet arrivals within flows could in itself be rich. Then, packet arrivals could be correlated across flows through interactions in queues and through reactive flow control such as the transport control protocol (TCP) that is active in the Internet. This feedback mechanism attempts to control the rate of most flows to avoid packet loss and maximize link utilization, effectively linking different flows dynamically. At another level, the statistics of “sessions,” which are groups of flows correlated through a higher level protocol or computer application, could be essential to take into account (this approach is adopted in [3]). For example, the downloading of a webpage results in the generation of multiple correlated TCP file transfers corresponding to the text, data, and images constituting the page. In this paper, we propose the use of a particular class of point processes: *Poisson cluster models* [4]. They are relatively simple, yet strongly motivated by empirical features of traffic, in particular, the role of flows, and their tractability allows the quantitative investigation of key properties as a function of meaningful network parameters. They are also easily synthesized and have marginals that are intrinsically positive. Through these models, we are able to give strong answers to several outstanding questions and clarify many issues. Although cluster models have been used in various fields such as meteorology, we are not aware of prior applications to IP packet traffic modeling. Very recent applications of cluster processes in networking have concerned the Web’s hypertext transfer protocol (HTTP) request arrivals [5] and TCP packet losses [6].

Manuscript received October 7, 2002; revised March 14, 2003. This work was supported in part by the French MENRT under Grant ACI Jeune Chercheur 2329, 1999. The associate editor coordinating the review of this paper and approving it for publication was Dr. Rolf Riedi.

N. Hohn and D. Veitch are with the Australian Research Council Special Research Center for Ultra-Broadband Information Networks, Department of Electrical and Electronic Engineering, The University of Melbourne, Victoria, Australia (e-mail: n.hohn@ee.mu.oz.au; d.veitch@ee.mu.oz.au).

P. Abry is with the CNRS, UMR 5672, Laboratoire de Physique, Ecole Normale Supérieure de Lyon, Lyon, France (e-mail: pabry@ens-lyon.fr).

Digital Object Identifier 10.1109/TSP.2003.814460

Our primary statistical tool is wavelet analysis. Apart from the high computational efficiency of the discrete wavelet transform that is necessary for the examination of the huge data sets typical in telecommunications, this is motivated by their natural suitability for signals with scale invariance. The discovery of scale invariance in packet data—the so called “fractal traffic”—was the most significant development in tele-traffic in the 1990s. On the whole, it refers to the near universal presence of long-range dependence (LRD), or persistent memory over “large” time scales, in time series extracted from raw traffic data such as byte or packet counts in successive time intervals [7]. The accepted physical explanation for this phenomenon lies in the heavy-tailed (finite mean, infinite variance) nature of source characteristics including session durations and file sizes. Long memory, however, is not the only issue concerning scaling. An equally remarkable feature, but one receiving far less attention, is the ubiquity and distinctiveness of the characteristic onset scale of LRD, which is found at around 1 s. One unresolved issue is what features of traffic determine this scale? Evidence for other kinds of scaling behavior have also been reported. Multifractal scaling [8], [9] has been suggested as a model of the extreme burstiness often observed at small scales (below 1 s) and sometimes above it [10], and infinitely divisible cascades [11] have been put forward as a means of unifying the scaling behavior across all scales. For a recent survey of wavelet methods and their application to scaling behavior in traffic, see [12].

One of our main goals was to explain all forms of scaling present in both statistical and networking terms. The importance of this arises from the fact that scaling typically implies high variability, which, in the case of traffic entering switches, implies worse queuing performance, as explored, for example, in [13]. Furthermore, its presence implies an underlying mechanism or mechanisms that need to be understood. Unless the source of such behavior is known, it will not be possible to predict how it, and its impact, will evolve over time. We contribute substantially to this issue. Through a model with a firm physical basis, we show that there are good reasons to believe that there is in fact no true scaling behavior at second order over small scales, which in turn implies no true multifractal behavior over those scales. We also provide explicit formulae capable of predicting the onset scale of LRD as a function of meaningful parameters.

Another goal is to contribute to a clarification of the meaning and role of the elephant (large but rare) and mice (small but numerous) flow concept, which has become popular in describing packet traffic. Rather than proposing fixed definitions of these categories, we let the data speak for itself and point out the orthogonal roles of “volume” versus “rate”-based approaches and the importance of time-scale.

This paper builds on the recent work described in [14]. The starting point of that paper was the surprising observation that the scaling seen in the point process of packet arrivals is broadly similar to that found in the arrival process of flow arrival points only, namely, clear LRD at large scales, evidence for a second, though less clear, scaling regime at small scales, and a transition scale at around 1 s separating them. This similarity led to the following question: In what way are the twin scaling regimes at

the IP level due to or influenced by the corresponding features at the flow level? Of the conclusions, the following, based on a second-order wavelet analysis, directly inspires the models we investigate here.

- The scaling in the flow arrival process is not responsible for that at the IP level, and further, it does not influence it significantly at either small or large scales.
- Dependencies between packet arrival processes across different flows are very weak.
- The structure at small scales has its origin in the packet patterns within flows.
- The LRD has its origins in the heavy-tailed nature of flow volumes (a known result) and does not have a component due to packet processes within flows (new result).

These findings (which are both discussed more fully and considerably extended in Section III and are consistent with recent work of [15]) have two very strong implications for traffic modeling. They suggest that, *for the purpose of modeling the overall process of IP packets*, flows can be treated as statistically independent. Thus, the point process of packet arrivals is seen as the superposition of independent point processes: one for each flow. Second, the lack of impact of the detailed nature of the flow arrival statistics suggests that they can be effectively modeled as a Poisson process. Finally, the isolation of the LRD as a property of the number of packets per flow allows them to be modeled using simple and intuitive heavy-tailed ingredients. Cluster models are ideally suited to modeling the above features.

We point out that although the arrival process of flows is not important for the overall packet process, it is of great interest in other contexts, such as the performance of web servers and proxies. Flow arrivals themselves have a rich structure, and there are many open questions. Some recent results can be found in [16] and [17].

The traces studied here and in [14] are of lightly loaded links. The central observation of independent flows underlying our model is likely to break down on heavily loaded links; however, exactly when this will occur is not clear. Low utilization notwithstanding, it is likely that a backbone link transports groups of flows that share bottleneck links elsewhere in the network, resulting in in-group dependencies. Nonetheless, such interactions were found to be negligible for the traces considered here, suggesting that the model could still apply at quite high utilizations and be a useful dimensioning tool for core networks.

The paper is structured as follows. Section II reviews the wavelet transform and gives examples of its use for scaling processes. In Section III, the technical details of the data and its processing are given, followed by the body of data analysis underlying the choice of the models. Section IV is the main part of the paper, where the cluster models are introduced, their properties given, and the fit to the data examined. Further analyses on the data are then performed, leading to suggested refinements to the model in Section IV-D, and a discussion on elephants and mice. Section V uses the model to examine in a well defined context the question “does traffic become more bursty or more Poisson as link rates increase?” and related issues. We conclude in Section VI.




Fig. 1. LD examples. (a) Poisson and fGn. (b) Poisson and Gamma-renewal. (c) GR and fGn. The upper dashed curves are the LDs of the superpositions. The * mark a characteristic upper saturation scale j_{GR}^* for Gamma renewal.

II. 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 [18]; in addition, see [19] for theoretical and practical details of their use in the spirit of this paper. Here, we briefly describe the key features, address some issues of particular importance at small scales, and give a short guide to interpretation.

A. Definitions and Properties

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

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

The wavelets $\psi_{j,k}(t) = 2^{-j/2}\psi(2^{-j}t - k)$ derive from an elementary function ψ , which is called the mother wavelet, dilated by a scale factor $a = 2^j$ and translated by $2^j k$. They are required to have excellent localization properties jointly in time and frequency. The function ψ is, moreover, characterized by its number of vanishing moments, which are defined as the largest integer N such that $\int t^k \psi(t) dt \equiv 0$ for $k = 0, 1, \dots, N-1$. Wavelets with higher N are smoother and are capable of analyzing signals with higher order divergences. A key practical advantage of the DWT 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

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

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

$$\Gamma_X(\nu) \sim c_f |\nu|^{-\alpha}, \quad |\nu| \rightarrow 0, \quad \text{with } \alpha \in (0, 1) \quad (3)$$

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

$$\mathbf{E} |d_X(j, k)|^2 \sim c_f C(\alpha) 2^{\alpha j}, \quad j \rightarrow +\infty \quad (4)$$

where $C(\alpha) = \int |\nu|^{-\alpha} |\Psi(\nu)|^2 d\nu$ is close to a constant. In fact, (2) can be viewed as defining a kind of wavelet energy spectrum, which is analogous to a Fourier spectrum but much better suited to the study of fractal processes. Just as in the Fourier case, the wavelet spectrum of a sum of independent processes is just the sum of the individual wavelet spectra, and multiplication by a constant c results in scaling the spectrum by c^2 .

To estimate the wavelet spectrum from data, the 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 octave j (scale $a = 2^j$), perform very well because of the short-range dependence in the wavelet domain. A plot of the logarithm of these estimates against j we call the *logscale diagram* (LD):

$$\text{LD : } \log_2 S_2(j) \text{ versus } \log_2 a = j.$$

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 Fig. 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 two relevant issues that are frequently ignored in applied work, in particular, in network traffic analysis.

- 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 the short-range dependence property of the sequences $d_X(j, \cdot)$ for each j to estimate them from data in a robust way.

- 2) The $O(n)$ pyramidal algorithm that calculates the $d_X(j, k)$ requires initialization by projecting $X(t)$ into some initial *approximation space* at an initial scale $a = \delta$. If this step is omitted, initialization errors result, which can be very significant for the smallest scales: $j = j_\delta + 1$ and $j_\delta + 2$, where $j_\delta = \log_2 \delta$. Furthermore, frequently $X(t)$ is only available via a discretised version $X_\tau(k)$: the result of a nonoverlapping averaging filter being applied to $X(t)$ about the points $t = k/\tau$, where τ is the sampling period. This limits the available scales to those above $j_\tau = \log_2 \tau$ and again results in errors over the first two available octaves $j = j_\tau + 1$ and $j_\tau + 2$. This is important as three fourths of the data is concentrated at these scales! For point processes, however, the initialization can be performed exactly. For simplicity, we use the Haar wavelet, where the initialization amounts simply to taking normalized counts, and use the higher order ($N > 1$) Daubechies wavelets to check the robustness of the conclusions.

C. Examples

In Fig. 1, LDs are given of some continuous time processes. The Fourier spectrum of each of these is known analytically, and so, we can evaluate the exact wavelet spectrum through (2). Here and below, the horizontal axis is calibrated both in scale a (top edge of plot, in “microseconds” (mus), “seconds,” or “hours,” as appropriate) and octave $j = \log_2 a$.

In plot (a), the horizontal line is for a Poisson process with $\lambda = 1$, viewed as a continuous-time process with delta functions at each arrival point, with spectrum $\Gamma_X(\nu) = \lambda$ (in this paper, we exclude the $\lambda^2 \delta(\nu)$ term corresponding to the “mean”). Equation (2) predicts $\mathbb{E}|d_X(j, k)|^2 = \lambda$, which is a flat wavelet spectrum corresponding to perfect but trivial second-order scaling ($\alpha = 0$). It is important to understand that this level corresponds to variance and not to rate: Means are eliminated by the wavelet analysis. The other straight line with slope $\alpha = 0.6$ is a continuous time *fractional Gaussian noise* (fGn): a generalized process with perfect scaling given by $\Gamma_X(\nu) = c_f |\nu|^\alpha$. The dashed curve in plot (a) is the LD of a superposition of the above two processes. Its form is a reminder that to add two curves in a log-log plot, one is really adding the underlying quantities and then taking the logarithm of the total. This same point is illustrated further in plot (b), where a Poisson process with $\lambda = 0.01$ and a Gamma renewal process with shape parameter of 0.2 are plotted; the dashed line representing their superposition is visually much closer to the Gamma renewal curve at large scale. Finally, plot (c) combines a Gamma renewal process with a fGn. The spectrum of Gamma renewal processes will be explored in detail in Section IV, where plot (c) will be particularly instructive in relation both to data and to cluster models.

III. EXPERIMENTAL RESULTS

In this section, we describe the main experimental findings underlying the models we subsequently select. We begin with some details on the data itself and then summarize and extend prior work.

TABLE I
TIME PERIODS

Traces	Date	Local Time	Rate (Mbps)	Link
MelbISP-1	20000425	19:00 to 22:00	0.03	Unknown
MelbISP-2	20000427	19:00 to 22:00	0.03	Unknown
MelbISP-3	20000428	19:00 to 22:00	0.03	Unknown
AUCK-a0	19991201	13:00 to 16:00	1.5	OC3
AUCK-b0	20010330	13:00 to 16:00	3.5	OC3
AUCK-c0	20010402	02:00 to 05:00	0.4	OC3
AUCK-c1	20010402	02:00 to 05:00	0.5	OC3
AUCK-d0	20010402	13:00 to 16:00	3.7	OC3
AUCK-d1	20010402	13:00 to 16:00	2.5	OC3
UNC-a0	20000927	19:30 to 20:30	178	OC12
UNC-a1	20000927	19:30 to 20:12	46	OC12
Abilene	20020814	10:00 to 10:10	418	OC48c

A. Raw Data

We analyze Internet traffic traces taken from lightly loaded links in a variety of geographical regions, with a wide range of average bit rates. The main body of traces we study—a selection from the Auckland II and Auckland IV data sets [20]—were recorded from the Internet access link of the University of Auckland. High precision *DAG* hardware allowed loss-less measurement of the OC3 ATM (155 Mb/s) link with timestamp accuracy of 100 ns [21]. The traces gather the timestamp of each IP packet, the packet size, and whether it is transporting TCP data or data from other protocols such as the *user datagram protocol* (UDP). UDP offers a simple transfer service with no flow control and is used for example for video streaming. As TCP “flavored” IP traffic makes up over 80% of all packets and bytes, we extract and focus on this component. As summarized in Table I, we focus on two 3-h periods during weekdays: 2:00 to 5:00 and 13:00 to 16:00, corresponding to apparently stationary traffic at “low” and “high” rate, respectively. The Precision *DAG* measurement was also used for the very recent high-rate Abilene trace collected at an OC48 (2.45 Gb/s) Internet backbone in Indianapolis, made available by NLANR [22].

We also study some traces with less accurate timestamps. UNC-a0 and UNC-a1, which were recorded by the DiRT group [23] at the University of North Carolina, are noteworthy for their high bit rate. The three traces included from a small Internet provider based in Melbourne (MelbISP) provide diversity in the packet rate within individual flows, owing to the speed limitations of modems.

For time and space reasons, not all analyses were performed, or reported on, over all traces; however, conclusions were always based on consistent results over multiple traces.

B. Time Series Extraction

The raw traces are processed with the CAIDA Coralreef tool suite [24] and our own C programs, allowing the extraction of each IP packet header and timestamp (for further details of TCP/IP protocols, see [25]). The information therein allows IP packets to be categorized into different *flows*. A flow is defined as a set of time-ordered packets with the same 5-tuple: IP protocol carried, source address, destination address, source port and destination port, where no packet inter-arrival exceeds a given time interval, fixed here at 64 s [24].

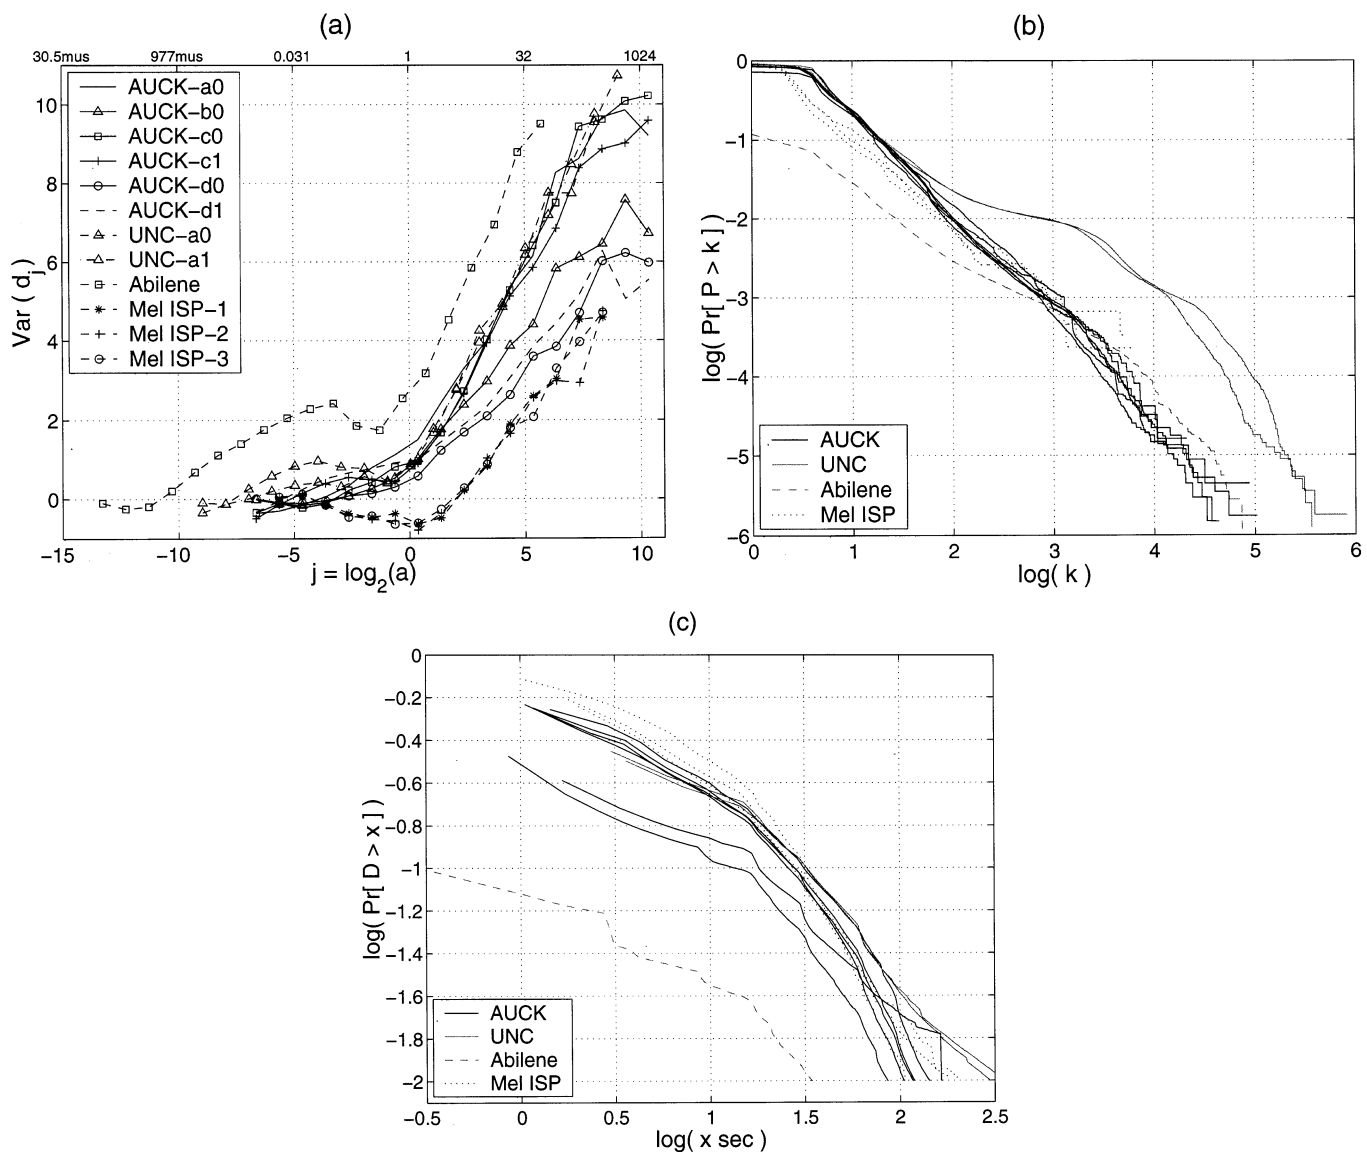


Fig. 2. TCP packet arrivals. (a) Ubiquity of biscaling behavior. (b) Heavy-tailed body and tail of P (number packets in flows). (c) Heavy-tailed flow durations D .

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 $t_P(k)$ of packets indexed in arrival order $k = 1, 2, \dots, K$. This time series defines the continuous time point process $X(t)$ of packet arrivals we wish to model or, equivalently, the interarrival sequence $A(k) = t_P(k) - t_P(k-1)$. At the “flow level,” statistics of individual flows are collected, beginning with the ordered arrival instants $t_F(i)$, $i = 1, 2, \dots, I$ of flows. The intrinsically discrete series $P(i)$ and $D(i)$, $i = 1, 2, \dots, I$ 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.

Considerable computation is required to perform the packet and flow level analyses here. The UNC-a0 trace, for example, consists of 2 GB compressed and contains 800 000 flows and 77 million packets, all individually tracked. To run our C and Matlab programs, we used a dedicated file server delivering compressed data off a RAID over Gigabit Ethernet to a dual

processor 900-MHz Dell workstation running Linux with 1 GB of fast memory.

C. Central Observations

The founding observation underlying our approach is the prevalence of “biscaling,” that is the observation of dual scaling regimes separated by a distinct “knee” in the packet arrival process $X(t)$. This is shown in Fig. 2(a) for the traces of Table I, where for ease of comparison the plot ordinates have been normalized (for more details, though on different traces, see [14]). At large scales, the LRD is clearly seen in each trace, and the “knees” in the curves are distinctive and all located in a narrow band at about 1 s. At smaller scales evidence for scaling is also present, which, although much noisier, recurs consistently across traces. Fig. 2(b) shows the remarkable power-law form of the distribution of P across traces and similarly for D in plot (c). In Section IV, we discuss the consequences of the fact that P , in addition to a power-law tail that contains only around 1% (depending on the exact definition of “tail”) of the mass, also has a distribution body which is close to power-law

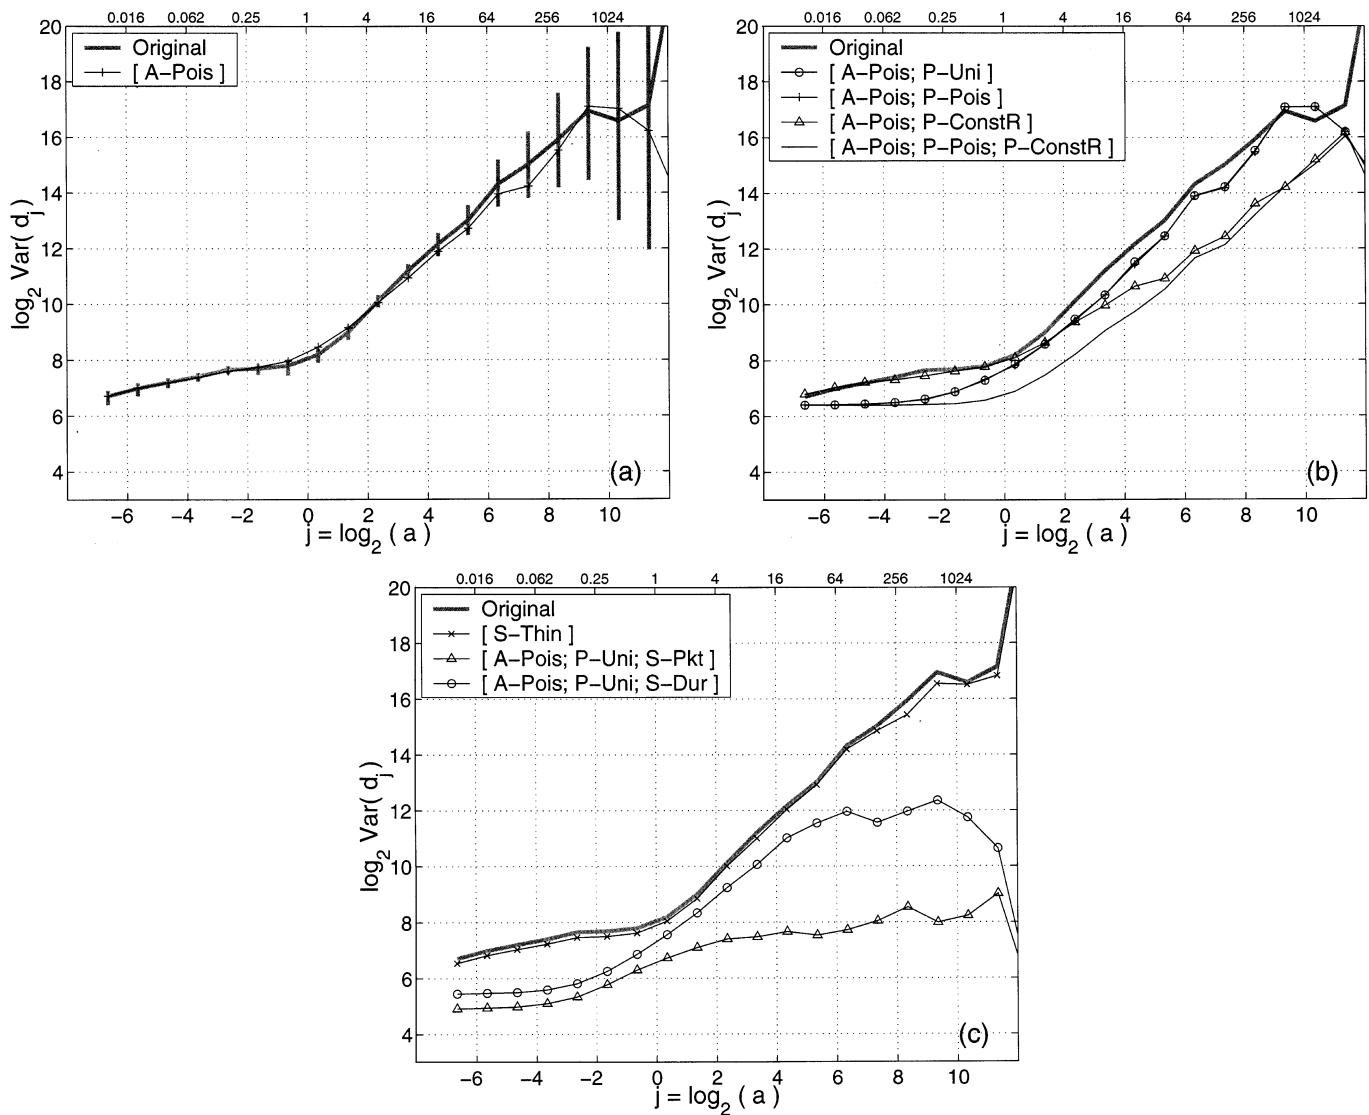


Fig. 3. Dissecting AUCK-c1 with the semi-experimental method. (a) Flow arrivals have negligible impact. (b) Small scales determined by in-flow structure, and D can be taken as proportional to $1/P$ (note that [A-Pois; P-Uni] and [A-Pois; P-Pois] are almost indistinguishable), and flow rate changes translate large scale behavior. (c) Thinning has no structural effect, and LRD is carried by heavy tailed P and/or D .

but with different parameters. In all cases, results from the same group (AUCK, UNC, MelbISP) are very consistent.

We employ a technique we call the *semi-experimental method*, which is invaluable as a means to track down the origins of, the connections between, and to selectively test models of, portions of the traffic structure, without having to postulate a full model from the outset. It involves transforming the original packet process in selective ways. Three categories of such “manipulation” will be used.

- A** Flow Arrival manipulation.
- P** Packet-in-flow manipulation.
- S** Flow Selection manipulation.

Our presentation is similar to but different from that of [14], and we examine the data in more depth both here and later in Section IV.

The thick grey curve in Fig. 3(a) is the LD of the trace AUCK-c1. The other curve ([A-Pois]) is constructed from the data by completely randomising the arrival process of flows, while maintaining in full the integrity of the packet arrival

patterns within each flow. More precisely, the flow arrival times are replaced by a sample path of a homogeneous **Poisson** process (conditional on the observed number of flows), the flow order is randomly permuted, and the flows themselves are then translated to the corresponding new arrival times. Despite this radical erasure of the flow arrival structure, *and* interflow dependencies, the resulting LD is barely altered. The result for other traces is just as striking (in Fig. 3, confidence intervals are placed on only one curve for readability). These results contradict modeling approaches which postulate the need for “session level” structure linking flows, at least for lightly loaded links.

In Fig. 3(b), we turn our attention to the packet statistics within flows. The curve [A-Pois; P-Uni] retains the flow placement of [A-Pois], as well as the original $P(i)$ and $D(i)$, but smooths out the packet arrivals within each flow. More precisely, if $P(i) = 1$ for flow i , then the sole packet is simply placed at its surrogate arrival point $t'_F(i)$. If $P(i) = 2$, then the second point is placed at $t = t'_F(i) + D(i)$. If $P(i) \geq 3$, then




Fig. 4. Examining flow variability (AUCK-d1). (a) Flow density plot over $(R(i), P(i))$ showing high mass over a distribution of rates. (b) Packet density plot (flow density weighted by number of packets). (c) Coefficient of variation per flow. In the main high mass region, flows are overdispersed.

the $P(i) - 2$ internal points are independently placed according to a uniform distribution over the duration of the flow. A clear difference is apparent at small scales. The wavelet spectrum has become flat, and the level in the LD is consistent with a Poisson process with the same average rate as $X(t)$. We conclude that the richness at small scales, and the (possible) scaling behavior, is due to the internal structure of flows and that conversely, the LRD is not due to this structure.

After performing [A-Pois; P-Uni], the only original features of the traffic left, where the origin of the LRD must lie, are the flow durations $D(i)$ and the flow packet counts $P(i)$. To narrow down this statistical origin more precisely, we select flow subsets according to different criteria. In Fig. 3(c), we first examine the effect of random thinning in the manipulation [S-Thin], where the flow and packet structure is fully retained, flows being randomly selected with probability 0.9. The resulting LD has the same shape as the original, with a variance which is approximately 90% of it, which is consistent with an independent and identically distributed (i.i.d.) superposition model. In contrast, in [A-Pois; P-Uni; S-Dur], we select only those flows with du-

urations below the 90% percentile. The result is the removal of the LRD. A similar result is obtained with [A-Pois; P-Uni; S-Pkt], when a selection is made based on the 90% percentile of P .

The result of [A-Pois; P-Uni; S-Pkt] is in keeping with the findings of [26] that show how the LRD at the IP level can be explained by the heavy-tailed distribution of file sizes. To explain that of [A-Pois; P-Uni; S-Dur], we are led to examine the relationship between D and P . However, although duration is a natural descriptor of a flow, it is a highly derivative one in that it is a dependent function of both the traffic source and the effect of the network. On the other hand, $P(i)$ acts like an independent variable describing the source, and the average rate $R(i) = P(i)/D(i)$, $i \geq 2$ combines source and link characteristics, since the average (and peak) rate of a flow is conditioned by the bandwidths of links it traversed before reaching the measurement point. Focussing, therefore, on rate rather than duration suggests that one might extend the in-flow packet manipulation so that $D(i)$ is no longer preserved but made a linear function of $R(i)$. A simple way to do this (in an average sense) is to reposition the packets in a flow according to a Poisson process,

which is a manipulation we call **[P-Pois]**. As seen in Fig. 3(b), the two curves **[A-Pois; P-Uni]** and **[A-Pois; P-Pois]** are almost indistinguishable. This shows that flows for which it would not be appropriate to slave $D(i)$ to rate (effectively to $1/P(i)$), such as those with very large gaps, have a negligible impact. Making $D(i)$ a dependent variable in this way opens up the possibility of renewal models for packets in flows and explains the observations of **[A-Pois; P-Uni; S-Dur]** as a simple consequence of those of **[A-Pois; P-Uni; S-Pkt]**.

We now consider flow behavior as a function of the “quasi independent” variables: average rate and flow volume. Because P is discrete, a scatter plot of $(R(i), P(i))$ hides mass along discrete lines and is very misleading. We therefore discretise the scatterplot to form the density plot [see Fig. 4(a)], where each square in the (R, P) plane is shaded according to the number of points within it. The mass is highly concentrated (most flows have a small number of packets), and therefore, a logarithmic scale is used to greatly enhance the outer regions. For a fixed packet volume, the average rates cover a wide range and, similarly, a flow with a given rate may contain many packets or as few as the minimum of 2. Furthermore, although the spread of values indicates high variability across flows, we do not see any bimodality that would suggest a need to classify flows into two or more classes. Simplifying things somewhat, the picture that emerges is that, in the range of rate values where the density is highest, the packet volume distribution is approximately independent of rate (and is heavy tailed). In Fig. 4(b), we give *packet density* rather than flow density, in effect weighting plot (a) by the packet impact of each underlying flow. The dark elements at large $P(i)$ correspond to volume-elephant flows, which have an appreciable packet impact despite arising from a very small percentage of flows—they were invisible in plot (a). Our conclusions are not altered however the epicentre of activity is still located at the dark region of plot (a). We return to the question of elephants in Section IV-D. We next look more deeply inside flows in two orthogonal ways.

Fig. 4(c) gives the value of the index of dispersion σ/μ of the inter-arrivals within a flow, calculated individually for each flow with at least three packets and then averaged over squares in a log-log plot. We see that they cover a wide variety of values, but the most extreme of these are not in the main region of high mass, as revealed in Fig. 4(b), which is of the same trace and shares the same scale. (At very small rate, we have a small number of very regular flows. These, due to TCP *keepalive* packets, have little impact.) On the contrary, the values in the main high mass region are reasonably uniform, with a weighted average value around 1.4; this is overdispersed compared to Poisson but not extremely so.

Fig. 3(b) includes two experiments designed to reveal the role of rate. The manipulation **[A-Pois; P-ConstR]** rescales the packet inter-arrivals within each flow by a constant $s(i)$ such that the average flow rates are moved to a common value: $R^* = s(i)R(i)$, which is chosen here to be the median rate. Despite preserving $P(i)$ as well as the individuality of packet structures within flows, the impact is notable; the entire large scale behavior is translated by a significant amount. This is more clearly seen in **[A-Pois; P-Pois; P-ConstR]**, where the small scale structure is simplified. In a third experiment, **[A-Pois;**




Fig. 5. Examining the inter-arrival process (AUCK-d0). (a) Inter-arrival distribution with fitted Gamma (shape = 0.6, mean = 1.2 ms). (b) Autocorrelation of (detrended) inter-arrivals.

P-ScaledR] (not shown), a uniform rescaling $R'(i) = sR(i)$ resulted in a simple horizontal translation of the LD by $-\log_2(s)$ (this is not surprising given Poisson flow arrivals). We conclude that rate acts as a scale parameter but that the wide distribution of rates noted from Fig. 4(b) cannot be easily represented by a single value at medium to large scales, unless it is tuned very carefully. The question of how to determine and interpret such a value is investigated below.

We now disregard flows and examine the inter-arrival series $A(k)$. Fig. 5 shows its histogram for AUCK-d0, which fits well to a Gamma random variable with $\sigma/\mu = 1.29$. The autocorrelation in plot (b) is negligible over small lags (small scale). Similar results apply for other traces, but it should be remembered that the time scale corresponding to a lag varies inversely as the packet rate. Whilst these results are true as such, they are in fact misleading. This can be revealed using a multiscale analysis and explained using a cluster model.

IV. CLUSTER MODELS

In this section, we define and evaluate two models for the point process $X(t)$ of packet arrivals, inspired by the observations above.

A. Black Box Model: Gamma Renewal (GR)

A renewal process is a simple point process, where the inter-arrival variables $\{A(k)\}$, $k \in \mathbf{Z}$ are i.i.d. We will examine its utility as a direct model for the inter-packet times. Although we seek meaningful constructive models rather than those of black box type, there are good reasons to first examine a renewal model. First, Fig. 5(b) directly suggests it. The second reason is the observation from Fig. 1(c) that a renewal process has the potential to generate scaling (or apparent scaling) behavior at small scales. The possibility of gaining a statistical understanding of this effect in a very simple context is worth pursuing. Finally, the spectrum of a renewal process plays a direct role in the cluster models introduced in Section IV-B.

The spectrum of the continuous time renewal process $X(t)$ is [4]

$$\Gamma_X(\nu) = \lambda_A \left[(1 - \Phi_A(\omega))^{-1} + (1 - \Phi_A(-\omega))^{-1} - 1 \right] \quad (5)$$

where $\Phi_A(\omega) = E[\exp(i\omega A)]$ is the characteristic function of the inter-arrival distribution, and $\omega = 2\pi\nu$ is the unnormalized frequency. Fig. 5(a) justifies a Gamma distribution for A , with characteristic function $\Phi_A(\omega; b, c) = (1 - i b \omega)^{-c}$, where $c > 0$ is the shape parameter. The exponential case is $c = 1$, corresponding to the Poisson process. As b is a scale parameter, $\Phi_A(\omega; b, c) = \Phi_A(b\omega; 1, c)$. The mean and standard deviation are given by $\mu_A = bc$ and $\sigma_A = b\sqrt{c}$ the coefficient of variation by $\sigma_A/\mu_A = 1/\sqrt{c}$, and the arrival intensity $\lambda_A = 1/\mu_A$, respectively. The following properties of the GR spectrum hold:

$$\Gamma_{GR}(\nu) \stackrel{\nu \rightarrow 0}{\approx} \lambda_A \left[\frac{1}{c} + \frac{(c^2 - 1)}{12c} (b\omega)^2 + O(\omega^4) \right] \rightarrow \frac{\lambda_A}{c} \quad (6)$$

$$\stackrel{\nu \rightarrow \infty}{\approx} \lambda_A \left[1 + \frac{2 \cos\left(\frac{c\pi}{2}\right)}{(b\omega)^c} + o(\omega^{-c}) \right] \rightarrow \lambda_A. \quad (7)$$

One can show that in the over-dispersed case ($c < 1$) of interest here, $\text{Re}(\Phi_A(\omega))$ is monotonic decreasing, from which it follows that the spectrum is as well. Since a monotonic spectrum implies a monotonic wavelet spectrum, the LD of GR with $c < 1$ monotonically increases from the asymptotic level $\log_2(\lambda_A)$ up to $\log_2(\lambda_A/c)$, as in Fig. 1(b). The small-scale asymptotic level is that of a Poisson process as well as of rate λ_A . However, this limit is not specific to Poisson but is due to the *general point process property that points do not coincide*.

Fig. 1(c) illustrates how, for a range of scales close to the upper asymptotic level, the LD of a GR process can appear to follow a straight line: a “pseudo scaling.” To quantify this, we define a lower cutoff frequency ν^* , where the spectrum can be said to “first” deviate from its asymptotic value. Fix a deviation parameter $\epsilon \in (0, 1)$. Define ν^* as the smallest ν such that the second term of (6) deviates from the first by ϵ times the distance

$\lambda_A|(1/c - 1)|$ between the asymptotic levels. The result, which respects the role of the scale parameter b , is

$$\nu^* = \frac{1}{2\pi b} \left(\frac{12\epsilon}{c+1} \right)^{\frac{1}{2}}, \quad c > 0. \quad (8)$$

The LD equivalent $j_{GR}^* = -\log_2 \nu^*$ is marked by asterisks in Fig. 1 ($\epsilon = 0.1$). Expressions for the center of the zone where such a pseudo scaling exists, and its slope, can also be derived, allowing predictive tests of the model. Approximate expressions for $c \in (0.2, 1]$ are given by $j_{GR}^c(b, c) \approx 1/b \cdot 1$, and $\alpha_{GR}(c) \approx (1 - c)/4$.

The model is easily calibrated through the sample mean and variance of the inter-arrivals. Comparing the resulting GR wavelet spectrum against the AUCK-c1 trace in Fig. 8(a), we see reasonable agreement at low scales and up to the onset of LRD. In general, however, the predictive ability of the GR model fails badly. The reasons for this become clear when one moves to the cluster model and result in useful insights, as we presently show.

Our final but important comment relates to the pitfalls in interpretation that “pseudo slopes” can cause. Since, for realistic values of c , $j_{GR}^*(b, c)$ is the same order of magnitude as μ_A , for both practical and physical reasons one is led to focus analysis on scales above it. This is standard practice in traffic analysis, as it seems inefficient to study time series that are mostly zeros. We have verified that if one does so, pseudo-slopes exist not only at second order but also more generally. Consequently, if one performs, for example, a wavelet multiscaling analysis of the type described in [12] over a range of moment orders, one finds empirical indications of multiscaling (possibly multifractal) behavior. This can lead to an erroneous belief that the data is much richer than a mere renewal process when, in fact, in this respect, it is entirely consistent with it. Indeed, it is likely that in many cases, the evidence for multifractal behavior over time scales below 1 s has been misinterpreted. In this paper, we do not present results beyond second order. Nonetheless, it is clear that if scaling (over some given scale range) is only *apparent* at second order, then the process cannot be multifractal as multifractality would imply true scaling over a range of orders, including second order. Exploring this issue in detail is the subject of ongoing work (see also [15]).

B. Flow Based Model: Poisson-Gamma Renewal (P-GR)

The main observations of Section III, that flows can be seen as independent entities arriving according to a Poisson process, fits naturally into a cluster process framework. A stationary Poisson cluster process on the real line [4], [27] consists of a Poisson process defining the locations of “seeds” about which a group of points are placed according to i.i.d. copies of another process. In a harmless abuse of notation, symbols already defined for the data will be reused.

Let the arrival times $\{t_F(i)\}$ of flows (the seeds) follow a Poisson process, also of rate λ_F . The packet arrival process can be written as

$$X(t) = \sum_i \mathcal{G}_i(t - t_F(i)) \quad (9)$$

where $\mathcal{G}_i(t)$ represents the arrival process of packets within flow i . Let the $\{\mathcal{G}_i\}$ be i.i.d., and consider a representative $\mathcal{G}(t)$. It is a point process containing a finite number $P \geq 1$ of points (packets) with the first located at $t = 0$.

Determining an appropriate process for $\mathcal{G}_i(t)$, given the complexity of TCP dynamics and network heterogeneity, is a challenge (see [28] for an interesting fluid model approach). Recall, however, from Section III-C that the manipulations **[P-Uni]** and **[P-Pois]** showed that simple “constant rate” models accounted for most of the second-order properties seen at the packet level. A (finite) renewal process model is a simple way to obey this finding, which has the advantage of falling within the theoretical framework of *Bartlett–Lewis* cluster processes. We choose the inter-arrival random variable A to be Gamma distributed [with c.f. $\Phi_A(\omega; b, c)$] for several reasons. First, it has a scale parameter, making it consistent (see below) with the observations on rate dependence of Fig. 3(b). Second, we have seen that **[P-Pois]** failed to reproduce important qualitative behavior at small scales. We will see below that incorporating burstiness through the variance to mean ratio is, in many cases, sufficient to reinstate this structure. This is easily and naturally achieved in the Gamma family, as the second parameter c is equivalent to this ratio, and $c = 1$ corresponds to **[P-Pois]**. Thus, finally, although the parameters λ_A, c of Gamma are not derived from network “first principles,” they do have physical meaning taken directly from data, and two is clearly the minimum number necessary.

The number of packets in a flow is a random variable P with density $p_j = \Pr(P = j)$, probability generating function $G_P(z) = \sum_{j=0}^{\infty} p_j z^j$, $|z| \leq 1$, and distribution function F_P (we take $p_0 = 0$). From Fig. 2(b), it is taken to be heavy tailed, that is, $1 - F_P(j) \stackrel{j \rightarrow \infty}{\sim} Lj^{-\beta}$, $\beta \in (1, 2)$, implying $\mathbf{E}[P] \equiv \mu_P < \infty$ but infinite variance.

Assembling these components, the flow model can be written as

$$\mathcal{G}_i(t) = \sum_{j=1}^{P(i)} \delta \left(t - \sum_{l=1}^{j-1} A_i(l) \right) \quad (10)$$

where $\delta(t)$ is a delta function centered at $t = 0$, $A_i(l)$ denotes the l th inter-arrival for flow i , and the inner sum is defined to be zero if $j = 1$. The average arrival intensity is given by $\lambda_X = \lambda_F \mu_P$. The parameters of the model are $\lambda_F; \lambda_A, c$; and μ_P, β . This is the smallest number allowing a packet level description of traffic with physical meaning: one parameter for flow arrivals, two for in-flow packet arrivals, and two for flow volume.

Apart from its physical motivation, one of the main advantages of the model is that its second-order properties are tractable. The expressions from [4, p. 417] and [27, p. 79] can be coerced into the following instructive real form:

$$\Gamma_X(\nu) = \lambda_F \left(\frac{\mu_P}{\lambda_A} \Gamma_{\mathcal{G}}(\nu) + (S_{\mathcal{G}}(\omega) + S_{\mathcal{G}}(-\omega)) \right) \quad (11)$$

where $\Gamma_{\mathcal{G}}(\nu)$ is the spectrum of the *stationary* renewal process with the same parameters as the finite flow renewal process, here $\Gamma_{\mathcal{G}}(\nu) = \Gamma_{\text{GR}}(\nu; b, c)$, and

$$\text{Re}(S_{\mathcal{G}}(\omega)) = \frac{\Phi_A(\omega)}{(1 - \Phi_A(\omega))^2} (G_P(\Phi_A(\omega)) - 1). \quad (12)$$

One sees immediately that the flow arrivals enter only through λ_F , which is a simple variance prefactor with an interpretation that one independently superposes “ λ_F ” of the same thing. Furthermore, the parameter $1/\lambda_A$ acts as a scale parameter: $\Gamma_X(\omega; \lambda_F, \lambda_A, c, F_P) = \Gamma_X(\omega/\lambda_A; \lambda_F, 1, c, F_P)$. This is a direct consequence of choosing Φ_A with a scale parameter obeying $b \propto \lambda_A^{-1}$. The third striking feature is that the expression consists of two terms of which the first $(\lambda_X/\lambda_A)\Gamma_{\text{GR}}(\nu)$ is familiar from Section IV-A. To understand the second, we note that

$$\text{Re}(S_{\mathcal{G}}(\omega)) \stackrel{\omega \rightarrow 0}{\sim} LB(\beta)(2\pi\lambda_A)^{2-\beta} \omega^{-(2-\beta)} \rightarrow \infty \quad (13)$$

$$\stackrel{\omega \rightarrow \infty}{\sim} -\frac{\cos\left(\frac{c\pi}{2}\right)}{(b\omega)^c} \rightarrow 0 \quad (14)$$

where $B(\beta) = \psi(1 - \beta) \cos(\pi\beta/2)/(2\pi)^{(2-\beta)} > 0$, $\psi(\cdot)$ denoting Euler’s Gamma function ((13) can be derived using a Taylor expansion of $\Phi_A(\omega)$ and employing a standard Tauberian theorem [29, p. 333]). Thus, at high frequency, the spectrum is dominated by the scaled GR term and, at low frequency, by the divergent second term. Comparing with (3), we see that the model is LRD with parameters $(c_f, \alpha) = (2\lambda_F LB(\beta)\lambda_A^{2-\beta}, 2 - \beta)$. It is significant that (13) depends only on the intensity λ_A of the GR flow processes and not on the second-order statistics: At large scale, the finer details of the flows cease to matter. This remains true if the standard deviation σ_P of P exists, in which case

$$\Gamma_X(\nu) \stackrel{\nu \rightarrow 0}{\sim} \lambda_F (\sigma_P^2 + \mu_P^2). \quad (15)$$

Recall that the GR term is monotonic decreasing when $c < 1$. The second term shares this property as $\nu \rightarrow 0$ and was observed to obey it for all ν where it is nonnegligible. Carrying over these observations to the wavelet spectrum, the generic shape of the LD for the model is similar to that of the dashed curve in Fig. 1(c): a monotonic function with the form of a (scaled) GR process, saturating at medium scales before crossing over to a LRD behavior at large scale.

An example of a wavelet spectrum for the model, evaluated using (2), appears in Fig. 6, where the magnitude of the (scaled) GR and LRD components are also plotted. The knee in the LD is now seen as the zone where these two compete. To capture its position as a function of parameters in a practically useful way, it is essential to realize that the scale at which the “road to LRD” begins may be very different from where the asymptotic LRD behavior of (13) dominates. We accordingly give two different definitions of transition scale. The first is the largest scale at which the small-scale effects, which are represented by the saturation level $\log_2 \lambda_X/c$ of the GR component, accounts for half of the wavelet spectrum. This scale, which is denoted by J_{PGR}^* , is the one we use for comparison against data, as it includes the important medium-scale effects. The second definition looks for equality between the large-scale asymptotic behaviors of the two spectral components $\Gamma_X(\nu) = \lambda_X/c$ and $\Gamma_X(\nu) = c_f \nu^{-\alpha}$, yielding

$$J_{\text{PGR}}^{**} = -\log_2 \lambda_A + \frac{1}{2 - \beta} (\log_2 \mu_P - \log_2(2LB(\beta)) - \log_2 c). \quad (16)$$

Its greater tractability encourages its use (see Section V) in describing the qualitative parameter dependence of the knee as the




Fig. 6. Comparison of LDs of AUCK-d1 and the P-GR model. The asterisk (resp. square) marks the transition scale j_{GR}^* (resp. j_{PGR}^*).

parameter dependencies of the two definitions are very similar. In order to see whether the GR component saturates before the LRD dominates, creating a plateau at medium scales as schematized in Fig. 1(c), one can compare j_{PGR}^* against j_{GR}^* , which can be rewritten as

$$j_{GR}^* = -\log_2 \lambda_A + \log_2 \left(\frac{\pi^2(c+1)}{3\epsilon c^2} \right). \quad (17)$$

In Fig. 8(a) $j_{GR}^* < j_{PGR}^*$, and so the plateau is visible, although just barely. If $j_{GR}^* \approx j_{PGR}^*$, then the plateau will have negligible width; however $j_{GR}^* \gg j_{PGR}^*$ is not possible since the departure from small scales leading to LRD can only take effect at scales where there are many packets in a flow, which is intuitively the same criterion defining GR saturation.

Another advantage of the model is that the packet inter-arrival time distribution can be calculated analytically [4], enabling comparisons against data and fitted Gamma inter-arrivals. Finally, simulation of the model is trivial and fast, apart from the long transient induced by the LRD.

C. Verification

The model works well when fitted to the packet process for the AUCK-d1 trace, as seen in Fig. 6. The use of GR flows, here with $\sigma_A/\mu_A = 1.58$ ($c = 0.4$), succeeds in modeling most of the burstiness which was not reproduced using [P-Pois] in Fig. 3(b). Here, $j_{GR}^* \approx j_{PGR}^*$, predicting that the plateau is not visible. This is the case, and j_{PGR}^* agrees visually with the onset of LRD. Furthermore, the visual agreement in the process $X(t)$ itself was found to be excellent, not only over the scales shown in Fig. 7, but over all scales. This agreement, essentially in the marginals, goes beyond second-order, even though the experiments were judged through the eyes of the wavelet spectrum, and indicates that the “physics” has been captured.

We can now explain the failure of the black box GR model. Simply, a scaled renewal process is not a renewal process. Thus, the “GR term” $\lambda_F \mu_P \Gamma_{GR}/\lambda_A$ of (11), although sharing the general form of a GR process, and the possibility of pseudo scaling with the same j_{GR}^* value, is not equivalent to one. The cluster




Fig. 7. Packet process (AUCK-d1). (a) Synthetic $X(t)$ data binned with $\tau = 50$ ms. (b) Corresponding original process $X(t)$.

model and the black box GR model can therefore never coincide at small scales unless $\lambda_F \mu_P / \lambda_A = \lambda_X / \lambda_A = 1$. Fortunately, this is almost the case in Fig. 8(a), where $\lambda_X / \lambda_A \approx 2.1$ but not in general. For the Abilene trace, $\lambda_X / \lambda_A = 278$. This result is significant since, looking solely at the results in Fig. 5, a GR process seems reasonable at small scales, but such measures cannot resolve important dependencies in the data, which are captured by the cluster model.

We now describe the parameter fitting in detail. The flow arrival parameter λ_F was estimated directly from the sample mean of flow inter-arrivals. Determining an appropriate λ_A for in-flow packet arrivals is not trivial. Simple choices such as the median of $R(i)$ [see Fig. 3(b)], or the mean, perform poorly. This is because, as we are modeling the packet arrival process, it is essential to capture the impact of each value of rate in terms of packets. We accordingly weight the average rate by $P(i) - 1$, which is the number of inter-arrivals in each flow. This results in values that are generally considerably above a simple mean, which agree well with semi-experimental comparisons.

The tail parameters (L, β) are measured via a least squares fit in a log-log plot of $1 - F_P(k)$. The fit is at logarithmically separated k and begins at small ($k = 6$) or medium (0.8 quantile) values, rather than just the far tail. The exponents of heavy tails are notoriously difficult to estimate, and the factor L is even more so. The above procedure includes more data and thus stabilizes the estimation and, in addition, is important in the present context where the distribution body is also power-law like over a range of scales. The resulting behavior in the LD is thus a mix of effects that must be appropriately captured when measuring

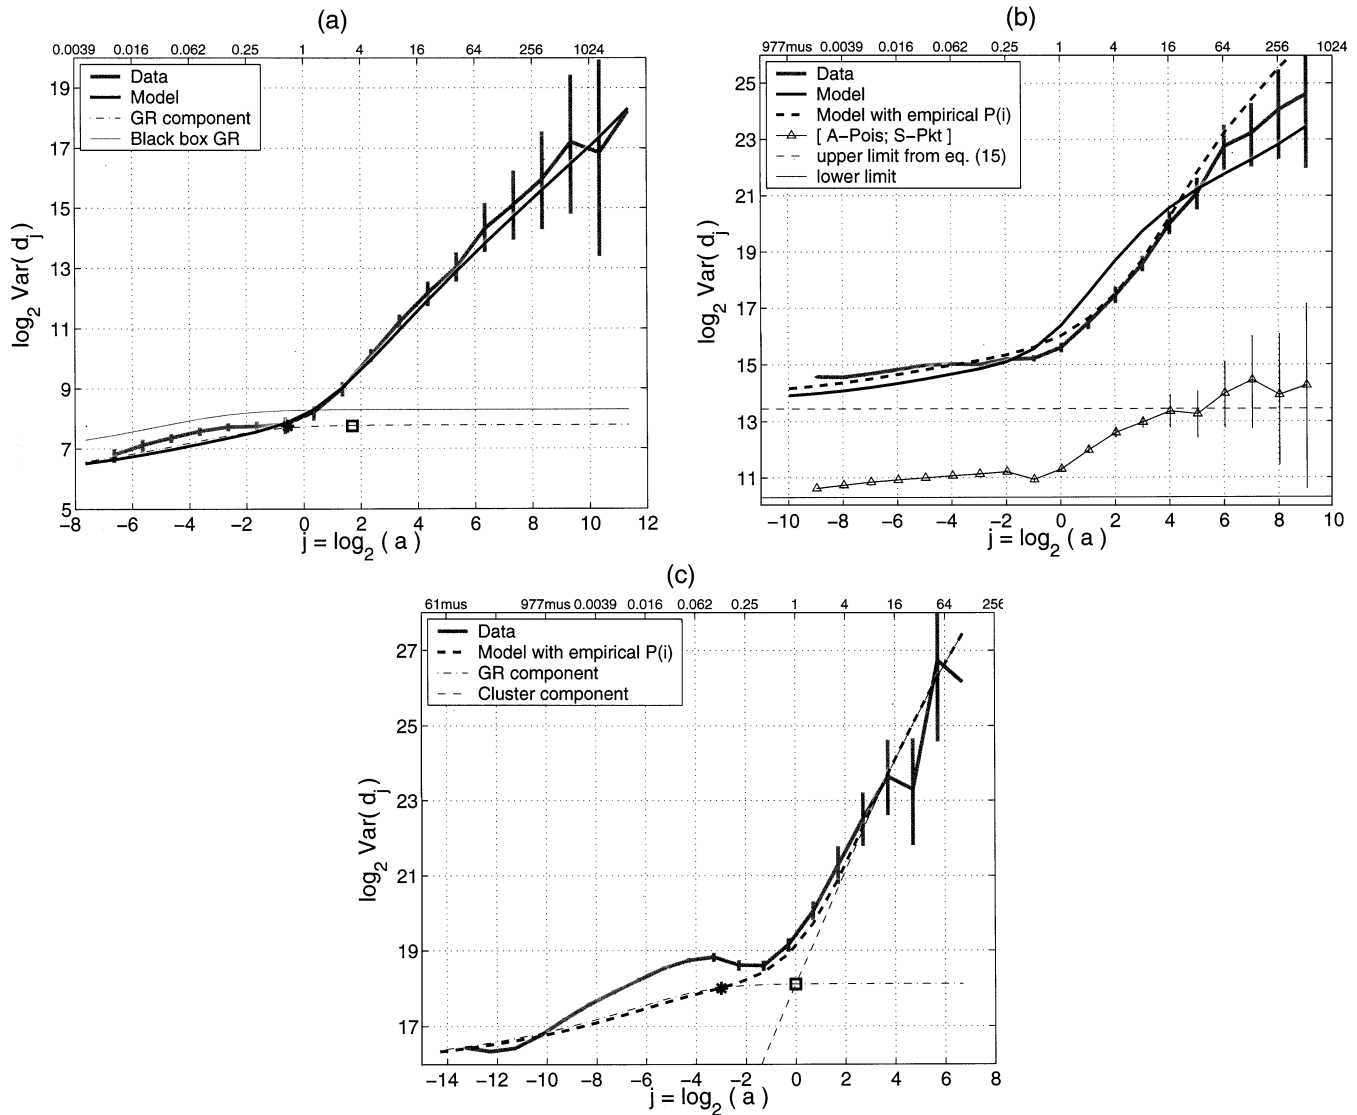


Fig. 8. Comparison of data and P-GR model. (a) Fit to AUCK-c1 is good, whereas the quality of the black box GR model is fortuitous. (b) Fit to UNC-a1 shows distortion not present when the empirical P histogram is used. A model using truncated empirical P agrees with the predicted level. (c) Abilene deviations remain even with the empirical P . The asterisk (resp. square) marks the transition scale j_{GR}^* (resp. j_{PGR}^*).

(L, β) . A measurement from the far tail only would not be consistent with (c_f, α) estimates except at extremely large scales beyond the usual observable range.

An entire distribution F_P is required for P to link its physical parameters μ_P, β, L . The discrete Pareto-like variable $H(k; a, \beta)$

$$F_H(k; a, \beta) = 1 - (ak + 1)^{-\beta} \sim 1 - Lk^{-\beta}, \quad k = 1, 2, \dots \quad (18)$$

where $a = L^{-1/\beta} > 0$ is a scale parameter, has mean $\mathbf{E}[H] = a^{-\beta} \zeta(\beta, 1/a)$ for $\beta > 1$ [the generalized Riemann Zeta function $\zeta(\cdot, \cdot)$ can be evaluated to chosen precision]. Unfortunately, $\mathbf{E}[H]$ can fail to match μ_P by a large factor. To broaden the family while respecting the power-law tail and/or bodies, we define the mixture distribution $F_P(k; p, a, \beta) = pF_H(k; a_2, \gamma) + (1 - p)F_H(k; a, \beta)$. For fixed $\gamma > 2$ (finite variance) and $a_2 > 0$, the mixture parameter $p \in [0, 1]$ allows the mean μ_P to be independently matched. For AUCK-d1, the fitting proce-

dures yields $(\lambda_F, \lambda_A, c) = (63, 8, 0.15)$ and $(p, a_2, \gamma, a, \beta) = (0.0236, 0.005, 3, 0.2711, 1.3510)$.

Finally, c can be tuned to fit the LD over scales below the LRD. Alternatively, a meaningful value of c can be derived by packet weighting as for λ_A above. The flows with the largest packet volume, as they also have higher average dispersion [see Fig. 4(c)], act disproportionately to decrease the effective c value. This illustrates a more general point. The detailed parameter fitting procedures above show that meaningful values can be given to the (meaningful) parameters, thus completing the physical validation of the model. For some parameters, however, notably c and λ_A , this can be computationally intensive. However, faster methods more akin to “fitting” could be used for more routine application of the model.

In Fig. 8(b), the fit to UNC-a1 is not quite as good, although the main features are reproduced; in particular, the knee position prediction is satisfactory. Much of the discrepancy is due to the more complex form of F_P . To see this, we also plot a “semi-model” fit, where the empirical distribution of P is used

instead of the fitted model distribution. The improvement reveals that the body of the distribution of P plays an important role in the shape of the “approach” to the LRD asymptote. Indeed, we have observed that in many cases, the observed “LRD” can be dominated by the shape of F_P at “medium” scales, resulting in estimates of the LRD exponent α , which are very misleading. To illustrate the relevance of (15), in the lower part of the figure, we show a semi-experimental LD, where the empirical distribution has been truncated at the 90th percentile, rendering the data short-range dependent. The LD then saturates at a value (dashed line), which agrees well with (15).

Finally, Fig. 8(c) shows the result for the high rate Abilene trace. As the fit is poorer, we show only the semi-model fit using the empirical distribution for P . We see that despite eliminating mismatches in the shape of P , the model fails to account for some of the variability at medium scales (also reported in [15] for other OC48 traces). Understanding the reasons for this requires a return to the data as well as an enhancement to the model.

D. Elephants, Mice, and a Multiclass Cluster Model

The term “elephants and mice” has become common parlance. It refers to the fact that often a small proportion of flows—the “elephants”—have a disproportionate impact over the more numerous “mice.” Typically, this distinction is made in terms of flow volume. The heavy-tailed modeling for P respects this idea, and the results for the Auckland and UNC traces show that the P-GR model is capable of naturally modeling both elephants and mice within a single model class. However, the concept can, and should, also be applied to the orthogonal dimension of traffic rate (see [10]). An important reason for this is that what constitutes a “large impact” is scale dependent. Only a small number of packets from volume-elephant flows intersect a given small interval, so their contribution will be negligible compared with that of volume-mice. Instead, flows with very high rate—rate-elephants—would make themselves felt at such small scales. On the other hand at large scales, localized high rates are irrelevant, and the contribution of volume-elephants is significant.

Although we noted in Section III that flow rates vary widely, in the P-GR model, they share a deterministic value λ_A . This was acceptable as a single value of λ_A could be found, which represented well the range seen in the high density portions of Figs. 4(a) and 4(b). This would not be the case if rate-elephants and rate-mice were present. A cluster model incorporating two distinct classes would then be needed in order to successfully describe behavior at all scales. To calculate the spectrum of a cluster model like P-GR but where the parameters can fall into two distinct classes: (“E,” with rate λ_A^E , shape c^E , and flow volume distribution F_P^E) and “M,” with parameters λ_A^M , c^M and F_P^M), we proceed as follows. Let B be a Bernoulli random variable (independent of P etc.) taking value “E” with probability q , else “M.” Consider a cluster process where for each flow an independent copy of B determines its class. By a well-known splitting property of Poisson processes, the set of seeds of clusters of type “E” (resp. “M”) is also a Poisson process with rate $\lambda_F^E = q\lambda_F$ (resp. $\lambda_F^M = (1-q)\lambda_F$). These two new processes, which each have constant rate, shape, and flow volume distribu-

tion, are independent P-GR processes. Thus, the spectrum Γ_X of the “multiclass” cluster model is just the weighted sum of two spectra of P-GR type. This construction can easily be extended to a countable number of classes.

With these additional tools at our disposal, we return to the Abilene trace with the flow density plot of Fig. 9(a). It tells a similar story to that of Fig. 4(a), albeit with a shift to higher rate (note that the diagonal boundary across the top is an edge effect due to the short duration of the trace). However, when we move to the packet density plot of Fig. 9(b), we see a striking change in the center of mass that is not found in the AUCK traces, where the epicentres of “packet” density and flow density coincide [compare Figs. 4(a) and 4(b)]. The location in (R, P) space of this high-density region represents an empirical definition of “elephant,” which is not tied to rate or packet volume alone. It is characterized by a very small proportion of flows containing a high proportion of total packets, with a higher average rate and higher average dispersion (lower c values), as seen from Fig. 9(c). Thus, the Abilene trace contains very strong, bursty, and high rate volume-elephants, and yet, by the argument above, the volume-mice must still be important for small enough scale, suggesting that a multiclass model may be essential for a full description of this data.

In future work, we will examine the usefulness of the dual class cluster model to explain the form of the wavelet spectrum shown in Fig. 8(c) (similar spectra have been observed in OC-48 commercial backbone links [15]). Alternatives to Gamma renewal models will also be investigated to model more extreme in-flow burstiness. Although the number of parameters increases when moving to multiclass models, it may be necessary to capture important network features. Network traffic is complex and cannot be reproduced accurately, nor meaningfully understood, with just three or four parameters. As the Abilene trace is a very recent one and is from a large backbone link, these complexities are exciting to explore since in many ways, they constitute a taste of the future of traffic.

V. TOWARDS UNDERSTANDING TRAFFIC EVOLUTION

In this section, we examine in more detail the nature of the P-GR model as a function of parameters and illustrate its use as a tool to speculate on the future shape of traffic. For convenience, we recall that for large j , the LD tends to $\log_2(c_f C) + \alpha j$, or

$$\log_2(2\lambda_F \cdot LB(\beta)C(2-\beta)) + (2-\beta)(j + \log_2 \lambda_A). \quad (19)$$

A. Flow Arrival Parameter λ_F

The role of λ_F is to vary the number of flows, which, through (11), can be seen as an i.i.d. superposition leaving the form of the second-order structure *invariant*. The magnitude of second-order dependencies relative to the mean decreases as $(\lambda_F/\mu_P)^{-1/2}$; therefore, this result is not in contradiction with the well-known weak convergence of such a superposition to a Poisson process [4]. In traffic engineering, this relative decrease of variability is known as *statistical multiplexing gain* and is a standard yet powerful argument for using links with higher capacity to enable more flows to mix together, effectively lowering variability, even for LRD traffic. This argument




Fig. 9. Flow and packet density in Abilene. (a) Flow density plot over $(R(i), P(i))$. (b) Packet density plot (flow density weighted by number of packets). (c) Coefficient of variation per flow.

follows “open loop” model reasoning, where network feedback is weak. This, however, is currently valid for backbone links, as network utilizations are low and are likely to remain so.

B. Flow Structure Parameters λ_A and c

Since $1/\lambda_A$ is a scale parameter, increasing λ_A results simply in translating the wavelet spectrum toward smaller scales. This can be seen explicitly in the expressions for the transition scales j_{PGR}^{**} and j_{GR}^* and in (19) above. Increasing λ_A also obviously scales back flow durations proportionally. At a fixed scale of observation, say at the sampling rate of a particular measurement infrastructure, one would see the traffic burstiness increase and become decidedly less Poisson as both the in-flow burstiness and scaling behavior translate to smaller scale. In network terms, increased λ_A could correspond to the same traffic passing through faster access networks before reaching the measured link.

Equation (19) is independent of c . Decreasing c results mainly in an increase in burstiness at scales below LRD through the plateau height λ_X/c and an increase in the pseudo slope at octaves below j_{GR}^* . It also results in a monotonic movement of approximately the same speed of both j_{GR}^* and j_{PGR}^{**} to higher

scales. Increased flow burstiness could arise through lower utilizations on network links, resulting in less queuing and therefore less traffic smoothing, as well as through more aggressive TCP flow control.

C. Flow Volume Parameters μ_P and (β, L)

We assume that these three can be varied independently, although this can never be entirely realized in a parametric family. At scales below j_{PGR}^{**} , the tail parameters (β, L) have no impact. The plateau onset scale j_{GR}^* is entirely independent of P , and μ_P enters only as a variance factor magnifying the burstiness at scales below j_{PGR}^{**} (thus, scaling up the pseudo-slope). At the other extreme, the LRD is unaffected by μ_P but is strongly influenced by the tail parameters; the asymptotic line moves up when the tail is made heavier either by increasing L or by decreasing β . The onset scale j_{PGR}^{**} is the result of competing effects. It is pushed up when increased μ_P increases short-range burstiness and grows to a limiting value with increasing β but *decreases* with increasing L . In terms of networks, a smaller β corresponds to an increased spread of file sizes, whereas L and μ_P trade off the proportion of “small” versus “large” files.

D. Future Scenarios and Scale of Observation

The parameter dependencies above can be combined according to possible future traffic scenarios. For example, assume that increased access link rates promote a proportional increase in network usage according to $\lambda_F \mapsto \Lambda\lambda_F$, $\lambda_A \mapsto \Lambda\lambda_A$, and consider the following question: Will traffic become more or less bursty? Clearly, the answer must be time scale dependent. If observing at a scale, which is in the range $[j_{GR}^*, j_{PGR}^{**}]$ both before and after the increase, then the multiplexing effect of case I alone will apply, reducing (relative) burstiness. At scales above j_{PGR}^{**} , however, the increase in λ_A largely cancels this out, and in addition, the LRD invades lower scales. If the more generous access rates also encourage greater transfer volumes $\mu_P \mapsto \Lambda\mu_P$, then $\lambda_X \mapsto \Lambda^2\lambda_X$, and the multiplexing effect will win out.

Care must be taken when one moves the scale of observation as parameters vary, such as when studying packet inter-arrivals. There, the characteristic timescale $1/\lambda_X = 1/(\lambda_F\mu_P)$ shrinks with increased flow rate or volume. Since j_{GR}^* is invariant with respect to each of these, as Λ increases, the point of observation in fact moves toward the point process limit of λ_X , regardless of the actual change(s) in traffic structure. Indeed, if smaller inter-arrivals occur purely because of greater μ_P , then absolute burstiness has in fact increased at scales below j_{PGR}^{**} , whereas the change in perspective might suggest that the traffic had become more Poisson-like. At such small scales, one should also be aware of the physical limitations of the point process model, which breaks down when packet sizes are reached. At [OC48,OC3] speeds (assuming a large 1500 byte packet), the model breaks down at around $[5, 77] \mu s$ or $j = [-15, -11]$.

VI. CONCLUSION

Our analysis of the structure of TCP packet arrivals in Internet traffic led to several significant conclusions. Beginning from the concept of flows of packets, we showed (at least in the context of lightly loaded links) that both the flow arrival process and dependencies between flows have negligible impact, as do higher layer mechanisms grouping flows such as web browsing sessions. The key element was found to be the concept of independence between flows. Using wavelet analysis, the second-order statistics of packet arrivals were shown to be determined by in-flow packet arrival burstiness at small scales and heavy-tailed flow volume at large scale. The scaling-like behavior at small scales was clearly linked to the burstiness within flows.

A stationary Poisson cluster process class was proposed as an ideal model capturing these features. Poisson arrival instants with rate λ_F denote the arrival of flows. Packets within flows follow finite GR processes with rate λ_A and shape c , flow volume being given by a heavy-tailed variable P with infinite variance. The model has many advantages, including a known spectrum, positive marginals, simple synthesis, and a minimum number of parameters, each with direct physical interpretation in terms of network traffic. Its spectrum can be written as a sum of a scaled spectrum of a renewal process controlling small-scale behavior and a term controlling asymptotic large-scale behavior. A detailed description was given of the behavior of the spectrum and the wavelet spectrum, as a function of parameters and the corresponding interpretation for networks. The model offers the pos-

sibility of a new, and very simple, alternative explanation for empirical evidence of multiscaling behavior at small scales as a transitional effect over a narrow range of scales of simple in-flow burstiness, suggesting that such traffic is not truly multifractal over these time scales. An expression for the onset scale of LRD was given, analyzed as a function of network parameters, and found to be accurate. The model is highly structural, rather than black box, enabling its use as an investigative tool for the evolution of traffic properties.

The model was verified against large quantities of accurate Internet data and was found to reproduce the second-order statistics well. The parameter fitting was described in detail. It led to meaningful parameter values and visually convincing model sample paths, confirming that the model actually captures much of the network “physics.” Some departures from the model were found for a recent, very high bit rate traffic trace. Further data analysis revealed some of the underlying reasons, and a multi-class version of the model was described as a possible means to account for them.

It was shown how the model can naturally incorporate the notion of elephant and mice flows without the need to explicitly define them and treat them separately. It was also used to illustrate how a packet volume-based definition of elephants is not sufficient and how “rate-elephants” could be accounted for in the model, should they exist.

REFERENCES

- [1] B. K. Ryu and S. B. Lowen, “Point processes models for self-similar network traffic, with applications,” *Stochastic Models*, vol. 14, no. 3, pp. 735–761, 1998.
- [2] —, “Point process approaches to the modeling and analysis of self-similar traffic—Part I: Model construction,” in *Proc. Conf. Comput. Commun.*, vol. 3, San Francisco, CA, Mar. 1996, pp. 1468–1475.
- [3] C. Nuzman, I. Saniee, W. Sweldens, and A. Weiss, “A compound model for TCP connection arrivals for LAN and WAN applications,” *Comput. Networks*, vol. 40, no. 3, pp. 319–337, Oct. 2002.
- [4] D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*. New York: Springer-Verlag, 1988.
- [5] G. Latouche and M.-A. Remiche, “An MAP-based Poisson cluster model for web traffic,” *Performance Eval.*, vol. 49, no. 1–4, pp. 359–370, 2002.
- [6] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker, “On the constancy of internet path properties,” in *Proc. ACM/SIGCOMM Internet Measurement Workshop*, 2001.
- [7] 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*, vol. 2, pp. 1–15, Feb. 1994.
- [8] J.J. Lévy Véhel and R. H. Riedi, *Fractals in Engineering*, J. Lévy Véhel, E. Lutten, and C. Tricot, Eds. New York: Springer, 1997.
- [9] A. Feldmann, A. Gilbert, and W. Willinger, “Data networks as cascades: explaining the multifractal nature of internet WAN traffic,” in *Proc. ACM/Sigcomm*, Vancouver, BC, Canada, 1998.
- [10] S. Sarvotham, R. Riedi, and R. Baraniuk, “Connection-level analysis and modeling of network traffic,” in *Proc. ACM SIGCOMM Internet Measurement Workshop*, 2001.
- [11] S. Roux, D. Veitch, P. Abry, L. Huang, P. Flandrin, and J. Mischeel, “Statistical scaling analysis of TCP/IP data,” in *Proc. ICASSP Special Session, Network Inference Traffic Modeling*, Salt Lake City, UT, May 2001, pp. 7–11.
- [12] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, and D. Veitch, “The multi-scale nature of network traffic: discovery, analysis, and modeling,” *IEEE Signal Processing Mag.*, vol. 19, pp. 28–46, May 2002.
- [13] A. Erramilli, O. Narayan, A. Neidhardt, and I. Saniee, “Performance impacts of multi-scaling in wide area TCP/IP traffic,” in *Proc. IEEE Infocom*, Tel Aviv, Israel, Mar. 2000.
- [14] N. Hohn, D. Veitch, and P. Abry, “Does fractal scaling at the IP level depend on TCP flow arrival processes?,” in *Proc. ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, Nov 6–8, 2002, pp. 63–68.

- [15] Z.-L. Zhang, V. Ribeiro, S. Moon, and C. Diot, "Small-time scaling behaviors of internet backbone traffic: an empirical study," in *Proc. IEEE Infocom*, San Francisco, CA, Apr. 2003.
- [16] N. Hohn, D. Veitch, and P. Abry, "Investigating the scaling behavior of internet flow arrivals," in *Proc. Colloque, Self-Similarity Applications*, Clermont Ferrand, France, May 2002, pp. 27–30.
- [17] —, "The Impact of the Flow Arrival Process in Internet Traffic," Oct. 2002, submitted for publication.
- [18] S. Mallat, *A Wavelet Tour of Signal Processing*. New York: Academic, 1998.
- [19] 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. New York: Wiley, 2000, pp. 39–88.
- [20] <http://wand.cs.waikato.ac.nz/wand/wits/> [Online]
- [21] J.Jörg Micheel, I.Ian Graham, and N.Nevil Brownlee, "The Auckland data set: an access link observed," in *Proceedings of the 14th ITC Specialist Seminar*, 2001.
- [22] <http://www.nlanr.net/> [Online]
- [23] <http://www.cs.unc.edu/Research/dirt/> [Online]
- [24] <http://www.caida.org/tools/measurement/coralreef/> [Online]
- [25] W. Stevens, *TCP/IP Illustrated*, 9th ed. Wellesley, MA: Addison-Wesley, 1996, vol. 1, The Protocols.
- [26] 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 *Proc. ACM SIGCOMM*, 1995.
- [27] D. R. Cox and V. Isham, *Point Processes*. London, U.K.: Chapman & Hall, 1980.
- [28] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and P. Owezarski, "A flow-based model for internet backbone traffic," in *Proc. ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, Nov 6–8, 2002, pp. 35–48.
- [29] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular Variation*. Cambridge, U.K.: Cambridge Univ. Press, 1987.



Nicolas Hohn received the Ingénieur degree in electrical engineering in 1999 from Ecole Nationale Supérieure d'Electronique et de Radio-électricité, Institut National Polytechnique de Grenoble (INPG), Grenoble, France. He received the M.Sc. degree in bio-physics from the University of Melbourne, Parkville, Australia, in 2000, while working for the Bionic Ear Institute. Since 2001, he has been pursuing the Ph.D. degree with the Department of Electrical and Electronic Engineering, University of Melbourne.

His research interests include physical models of Internet traffic and theory of point processes.



Darryl Veitch (SM'98) was born in Melbourne, Australia, in 1963. He received the B.S. degree with Honours from Monash University, Melbourne, in 1985 and the mathematics Ph.D. degree in dynamical systems from the University of Cambridge, Cambridge, U.K., in 1990.

In 1991, he joined the research laboratories of Telecom Australia (Telstra), Melbourne, where he became interested in long-range dependence as a property of tele-traffic in packet networks. In 1994, he left Telstra to pursue the study of this phenomenon at the CNET, Paris, France (France Telecom). He then held visiting positions at the KTH, Stockholm, Sweden; INRIA, Sophia Antipolis and Nice, France; and Bellcore, Red Bank, NJ, before taking up a three year position as Senior Research Fellow at RMIT, Melbourne. He then joined the Electrical and Electronic Engineering Department at the University of Melbourne as a Senior Research Fellow, where, for two years, he directed the EMULab: an Ericsson-funded networking research group. He is now a member of the ARC Special Research Centre for Ultra-Broadband Information Networks (CUBIN) within the department. His research interests include scaling models of packet traffic, parameter estimation problems and queueing theory for scaling processes, the statistical and dynamic nature of Internet traffic, and the theory and practice of active measurement of packet networks.



Patrice Abry was born in Bourg-en-Bresse, France, in 1966. He received the Professeur-Agréégé de Sciences Physiques degree in 1989 from the Ecole Normale Supérieure de Cachan and the Ph.D. degree in physics and signal processing from the Ecole Normale Supérieure de Lyon and Université Claude-Bernard Lyon I, Lyon, France, in 1994.

Since October 1995, he has been a permanent CNRS researcher at the Laboratoire de Physique, Ecole Normale Supérieure de Lyon. His current research interests include wavelet-based analysis and modeling of scaling phenomena and related topics (self-similarity, stable processes, multifractal, $1/f$ processes, long-range dependence, local regularity of processes, infinitely divisible cascades, departures from exact scale invariance ...). Hydrodynamic turbulence and the analysis and modeling of computer network teletraffic are the main applications under current investigation. He is the author of the book "Ondelettes et turbulences—Multiresolution, algorithmes de décompositions, invariance d'échelle et signaux de pression (Paris, France: Diderot, éditeur des Sciences et des Arts, October 1997). He also is the coeditor of the book "Lois d'échelle, Fractales et Ondelettes" (Paris, France: Hermès, 2002).

Dr. Abry received the AFCET-MESR-CNRS prize for best Ph.D. dissertation in signal processing from 1993 to 1994.