Seeing the difference in IP traffic: Wireless versus Wireline

Julien Ridoux
LIP6 - UPMC
8, rue du Capitaine Scott.
75015 Paris
FRANCE
julien.ridoux@lip6.fr

Antonio Nucci †
Narus, Inc.
500 Logue Avenue
Mountain View, CA 94043
USA
anucci@narus.com

Darryl Veitch ‡
ARC Special Research Centre on
Ultra-Broadband Information Networks
The University of Melbourne
Victoria 3010, AUSTRALIA
d.veitch@ee.mu.oz.au

Abstract— **With the explosive growth of the Internet over the last 10 years, a lot of work has been dedicated to understanding the underlying mechanisms of wired IP traffic. Recently, the rapid deployment of large-scale wireless infrastructures in various environments and the interesting mixture of traffic carried coupled with the large diversity of devices accessing the medium (Cellphones, Laptops, PDAs) have triggered the attention and curiosity of the research community. This paper analyzes in depth the properties of several large traces of packet data collected between the wireless access point and the IP cloud from an operational wireless service provider. We determine unambiguously the influence of network variables such as the arrival patterns of packet and flows, flow durations and flow interactions, on the aggregate statistics of TCP traffic. In doing so, we highlight the main differences and similarities between wireless and wired IP traffic, and between the two directions (from wireless devices to IP cloud and vice-versa), and show how they can be distinguished. The resulting insights provide a foundation for models of such traffic, necessary for improved resource allocation schemes as well as for the effectiveness of future services and applications.**

Index Terms— **wireless traffic characterization, Internet traffic, wavelets, semi-experiments**

I. INTRODUCTION

For over ten years now, researchers have sought to characterize the traffic carried on the Internet as an essential first step prior to modeling, traffic engineering or application design. A series of studies focusing on the core of the Internet provided insights into the properties and characteristics of IP traffic. Properties such as self-similarity [1], long-range dependency (LRD) [2] and scaling behavior at small time-scale [3] have been discovered, highlighted and discussed. While studies on backbone IP traffic provide indications as to the nature of wireline traffic, which as a result is now better understood, the deployment and growth of wireless access networks of different kinds now requires the examination of wireless IP traffic.

In the last few years the demand for wireless Internet access has increased exponentially. Mobility seems to be transforming the communication industry, shifting the momentum from Internet access to broadband wireless access. Carriers have slowed expansion of their fiber networks in anticipation of new wireless technologies, and engineers have refocused development towards the products and services that will enable broadband wireless communications to become ubiquitous. Meanwhile, Wi-Fi networks and hot-spots are widely implemented in homes, schools, businesses, cafes and other public areas, focusing mainly on supporting nomadic user behavior. These networks provide people with high-speed access to the Internet, without being tethered by a cord or cable (in a range of around 300 feet from the access point). WiMAX, which is still in its infancy, will deliver last-mile broadband connectivity over a larger geographical area than Wi-FI (one to six miles) enabling greater mobility for high-speed data applications, while 3G (GPRS/EDGE, UMTS) technology provides long-range wireless access for a larger range of devices supporting voice and data.

Unfortunately, not much is known about the nature of the traffic carried by each of these infrastructures, and still less on how they might combine. Researchers have only recently focused on this problem. Data from Wi-Fi networks has been collected and analyzed to answer specific questions related to mobility and related user behaviors [4], as well as to measure the characteristics of the associated traffic [5], [6], [7], [8]. On the other hand, some studies have emerged for large-scale wireless access networks set up by Internet Service Providers (ISPs), aiming to quantify the performance of GPRS and to understand its traffic characteristics, either on testbeds [9] or in a real environment [10]. While the number of excellent studies on wireless traffic is growing, this area of research is still fairly immature.

Compared to prior work realized in wireless environments, the present study brings a certain level of novelty in the raw material studied and tries to bridge the knowledge gap between wireline and wireless IP traffic. We have the opportunity to study traffic traces captured in the operational environment of a nation-wide ISP wireless access network. Based on CDMA-1xRTT technology, this wireless access network gathers together the large scale and mobility properties one can find in the GPRS/EDGE world (seamless mobility and limited set of applications), and the diversity of devices and applications as found in Wi-Fi networks (laptops/PDAs and any application). The scope of this paper is to understand the similarities and differences between this traffic and the traditional IP traffic, and iden-

† This work has been conducted while Antonio Nucci was at Sprint Labs.

‡ Darryl Veitch is with the ARC Special Research Centre on Ultra-Broadband Information Networks (CUBIN). CUBIN is an affiliated program of National ICT Australia (NICTA). Darryl Veitch was partially supported by the ARC

1-4244-0222-0/06/$20.00 (c)2006 IEEE

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the Proceedings IEEE Infocom.

tify unambiguously the underlying mechanisms involved. Our methodology relies on the comparison of "typical" wired Internet traces (captured on backbone links) to wireless CDMA-1xRTT ones. A strong emphasis is given to the structure of packet and flow arrival processes, and the importance of flow interactions. The present study adopts then a different point of view compared to [10]: we are not currently interested in analyzing the performance of protocols on the wireless environment but we try to understand the intrinsic properties of the wireless traffic. To be able to answer the above questions and conjecture the causes of key statistical findings, we introduce a large set of virtual experiments which enable the exploration of 'what if' scenarios, and we interpret the results using wavelet based spectral analysis. We find that in some important respects the wireless traffic is similar to Internet traffic, and also that in other important respects, it is not.

The rest of the paper is organized as follows. In section II we introduce the architecture of the wireless access network from which we collect the packet traces, and discuss interesting statistics of the time series extracted from the raw data. In section III we present our methodology based on 'semi-experiments' and wavelet analysis to enable us to dig deeper towards an understanding of the underlying mechanisms, while in section IV we offer the results and interpretations of our main findings. Section V concludes the paper and discusses possible directions for future research.

II. Packet Traces

The packet traces we use in this study are from the Sprint IPMON project [11]. On the Sprint IP backbone network, about 60 monitoring systems are deployed all across the backbone. Using optical splitters, the monitoring systems capture the first 44 bytes of all IP packets and time-stamp each of them. As the monitoring systems use the Global Positioning System (GPS) for their clock synchronization, the error in time-stamp accuracy is bounded to less than 5 microseconds. We used three of the monitoring systems to collect data from a lightly loaded OC-12 (622 Mbps) backbone link and from two OC-3 (155 Mbps) links between the Customer Data Network (CDN) and the core network, that is one for each direction. As TCP "flavored" IP traffic makes up over 90% of all packets and bytes, in the following we focus on this component only.

We begin with a brief discussion of the wireless data, its collection, and then give a description of basic statistics of the time series extracted from the raw data. We then describe the backbone link that will be used to compare the differences and similarities between the traditional and well-studied wired IP traffic, and wireless traffic. As our proposed methodology compares wired to wireless traffic, figures 4 to 7 present results for both kinds of traffic traces.

A. Wireless Data

Monitoring the links between the Customer Data Network and the core network represents a challenging experiment. From this measurement point, a diverse mixture of traffic coming from different end-mobile users, geographically spread across a very wide area, is visible. Monitoring the traffic between the CDN and the core network allows then to observe traffic originated from and intended to numerous wireless cells. Different traffic profiles originated by a variety of mobile devices are also aggregated together. For example, 3G-phones and PDAs each feature data roaming, whereas they are expected to generate different traffic profiles because of their different data processing capacities. On the other hand, laptop user mobility will be nomadic, but similar to PDAs in terms of applications used. In this paper we focus our attention on the aggregated traffic with no distinction among these three mobile device categories. Studying the traffic profiles associated to each of these is out of the scope of this paper and we reserve the problem for future work.

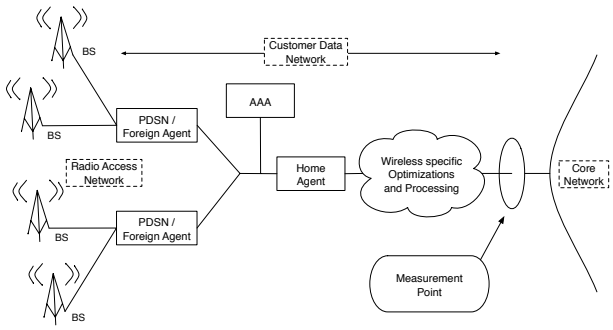

Fig. 1. Wireless and Measurement Architecture.

1) Wireless Network Architecture: Figure 1 shows the wireless access network with mobile devices (3G-phones, PDAs, laptops) communicating with the Base Stations (BSs) over wireless links. The BSs interact with the rest of the voice or data network through the Base Station Controllers (BSCs). Usually each BS covers tens of mobile devices belonging to the same cell site (usually covering tens of miles), while each BSC interacts with tens of BSs covering a wide geographical area (hundreds of miles). This part of the network is common to both wireless voice and data traffic. The network separates only beyond the BSCs where voice frames are forwarded to the Public Switched Telephone Network (PSTN) while data frames are forwarded to Packet Switched Data Network (PSDN). In this paper we are interested in the data traffic only and we deliberately ignore all the traffic delivered to the PSTN. All network devices along the path described above are part of the *Radio Access Network* (RAN) that provides the basic transmission, radio control, and management functions needed for the mobile user to access the resources of the core network and the end-user services network. It is this access network that terminates the air interface to the Mobile Station (MS) and converts the air frames to packet format for transporting traffic between mobile users on other RANs or to the core network. In the network architecture at our disposal the RAN is based on a CDMA-1xRTT technology.

Between the Radio Access Network and the core network, the traffic goes through what we call the *Customer Data Network* (CDN) whose entry points are represented by Packet Data Switch Node/Foreign Agent (PDSN/FA). This part of the network is in charge of several specific functions like Authoriza-

tion, Authentication and Accounting (AAA) and mobility management through a Mobile IP infrastructure. When a user appears on the RAN, the FA relays its authorization request to the AAA server for authentication. Once authenticated, the request packet is forwarded to the Home Agent (HA) that replies with an IP address from its pool. From this moment, an association is maintained allowing the new user to access the network resources while being free to move to other cells. The CDN is also responsible for transcoding, bandwidth optimization and other proxy functions specific to the type of traffic processed.

2) Data sets: As it is evident from the architecture of the network, all traffic from the wireless access network to the core network and vice-versa must traverse a set of links and "boxes" between the CDN and the core network. This wired architecture is typical of wireless access networks. In the network studied we can capture all packets in both directions by monitoring two OC-3 links. Two wireless traces were captured on April 1st 2004 during a continuous 24 hour period, allowing to capture the users behavior over an entire day. Average values of their link utilization and packet rate are indicated in table I. In the following, by "Wless-OUT", we refer to packet traces capturing traffic originated from wireless devices to the core network, while "Wless-IN" refers to traffic from the core network to the wireless devices. We start our discussion by showing some basic statistics and properties.

TABLE I
AVERAGE CHARACTERISTICS OF THE WIRELESS TRACES.

	Wless-OUT	Wless-IN
Duration	23h59m59s	23h59m59s
Link Capacity	OC-3	OC-3
Util. (Mbit/s)	2.249	7.447
Util. (Pkt/s)	1893	2121
Total # Packets	25,476,058,647	84,339,787,498
Average Packet IAT	404.995 μs	341.590 μs
Total # Flows	29,970,428	35,202,266

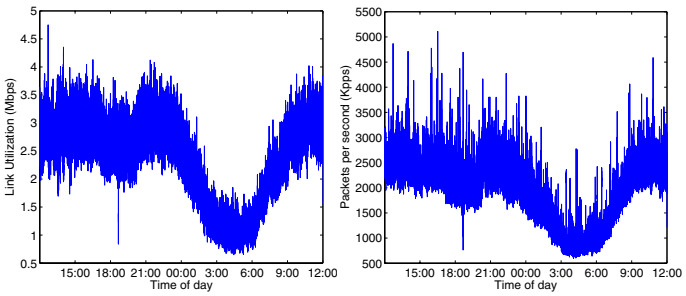

(a) Link Utilization (Mbps) (b) Link Utilization (Kpps)

Fig. 2. Link Utilization of Wless-OUT.

Figures 2 and 3 show the link utilization in Mbps (on the left) and the corresponding packet arrival time series in Kpps with 1-second bin durations (on the right) respectively for Wless-OUT and Wless-IN. Both directions exhibit a strong diurnal behavior with a decrease of the traffic in the early hours of the day and an increase later in the evening. The link in the Wless-IN direction has a utilization more than twice the one of Wless-OUT. Even with this proportion in mind, Wless-OUT appears to be less

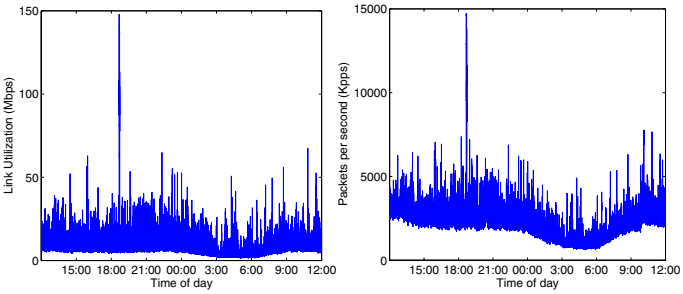

(a) Link Utilization (Mbps) (b) Link Utilization (Kpps)

Fig. 3. Link Utilization of Wless-IN.

bursty than Wless-IN. In the latter, spikes in traffic up to 10 to 80 times larger than the average are present across the entire trace.

By comparing the link utilization and packet arrival time series, we can see an interesting difference between the two directions. In Wless-OUT the spikes in the packet arrival time series are not apparent in the link utilization, suggesting they consist of highly localized groups of low total byte count, while in Wless-IN spikes of similar magnitude are common to both time series.

In figures 4, 5 and 6 we plot respectively the distributions of the number of packets, the associated packet rate, and the packet inter-arrival time (IAT) in TCP flows. Not surprisingly, the curves for Wless-IN and Wless-OUT in figure 4 are very similar, since they share common context information about the two directions of the same flows. The same comment can be made about Wless-IN and Wless-OUT in figures 5 and 6.

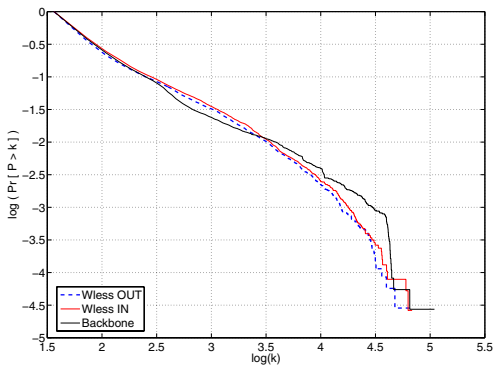

Fig. 4. Empirical Distribution of number of packets P in TCP flows.

In figure 7 we plot the empirical distribution of the TCP flow arrival process, which we call $Y(t)$ below. It is interesting to note a remarkable difference in the shape of the distributions for the two directions. While the distribution of the Wless-IN looks similar to an exponential distribution with a large exponent, the distribution of Wless-OUT exhibits an interesting "wavy" behavior with period around 0.6ms. Unfortunately we were not able to identify the reasons of this deterministic periodicity, as it would require identification of the type of network device and details of the data path followed by the traffic. We leave it as part of future work but discuss and conjecture possible causes in section IV-C. Finally, we notice the large portion of very small flow inter-arrivals in Wless-IN. We believe that this char-

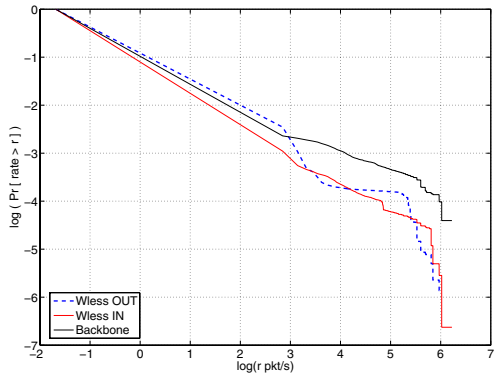

Fig. 5. Empirical Distribution of packet rates in TCP flows.

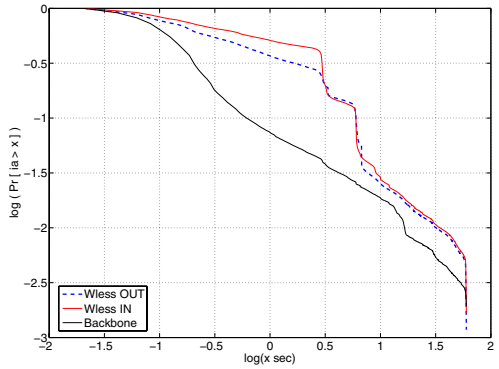

Fig. 6. Empirical Distribution of packet inter-arrivals times in TCP flows.

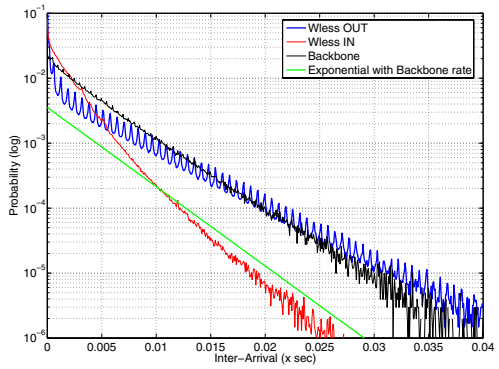

Fig. 7. Empirical Distribution of TCP Flows inter-arrivals times.

acteristic is a key factor in explaining the different burstiness observed in the two directions. The deeper analysis provided in section IV adds weight to this intuition.

B. Wireline Data

We selected a backbone wireline trace based on its equivalent duration (a continuous 24 hours capture) and its comparatively low link utilization (for an Internet backbone link), to facilitate comparison. In the following, we will refer to this trace as "Backbone". It was collected on a lightly loaded OC-12 (655 Mbps) link on January 28th 2004. This Backbone trace is not related to the wireless traffic described above, but has been chosen as a good representative of the typical characteristics observed in previous studies characterizing Internet backbone traffic. Average values of the link utilization and packet rate are indicated in table II.

TABLE II
AVERAGE CHARACTERISTICS OF THE WIRELINE TRACE.

	Backbone
Duration	23h59m59s
Link Capacity	OC-12
Util. (Mbit/s)	40.999
Util. (Pkt/s)	10656
Total # Packets	478,304,079,349
Average Packet IAT	81.328 μs
Total # Flows	51,219,857

Figure 8 shows the traffic rate in Mbps (on the left) and the corresponding packet arrival time series in Kpps with 1-second bin size (on the right). Visually, the burstiness characteristic of this trace seems to be different from either Wless-OUT and Wless-IN.

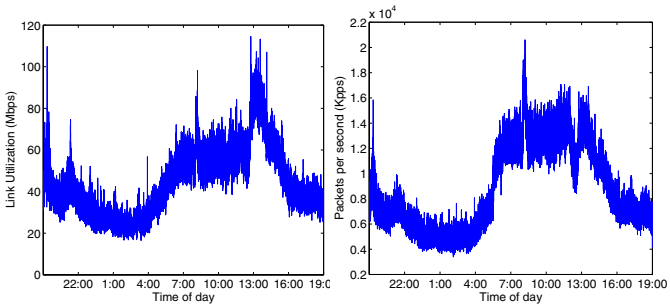

(a) Link Utilization (Mbps) (b) Link Utilization (Kpps)

Fig. 8. Link Utilization of Backbone.

In figure 4 we see that the tail of the number of packets per flow is heavier for the Backbone link, which is what we would expect given that it does not have the same bandwidth limitations and does not have the same billing scheme. All traces exhibit heavy tails however, indicating the wide range of download sizes attempted in each case. Figure 5 also shows that higher rates are accessible, and used, by Internet users more often than wireless users, although they are similar at lower rates. Figure 6 and figure 7 are more interesting. In the former, the two wireless traces are similar, and clearly different from the backbone, where the higher available rates allow closer spacing of packets on average. In the latter however, all three traces are different. Nevertheless, Wless-OUT is distinctly the most different one and Backbone and Wless-IN are closer to each other. The presence of the exponential distribution of same rate as the Backbone distribution shows it.

In this section we presented and discussed briefly some characteristics of the time series extracted from the raw data collected. We highlighted some interesting properties of the data but we could not determine unambiguously the influence of arrival patterns of packets and flows, or that of flow interactions on the aggregate TCP traffic. In the next section we propose another methodology based on wavelet analysis to enable us to dig deeper towards an understanding of the underlying mechanisms.

III. METHODOLOGY

Our methodology is based on using a wavelet based spectrum as the metric to observe and interpret the effect of *semi-experiments*.

The term *semi-experiments* was coined in [12] to describe a methodology of virtual experimentation which enables, based on a single data set, the exploration of 'what if' scenarios aiming to determine the causes of statistical properties of the data. It is a systematic extension of the idea of block-wise *shuffling* introduced in [13] to explore the presence of long-range dependence (LRD) in time series of byte counts. Typically, a semi-experiment involves replacing a single specific aspect of the real data with a simple, neutral model substitute. One then compares the statistics before and after, drawing conclusions on the role played by the structure removed by the 'manipulation'.

The metric we use as the basis of comparison is the wavelet spectrum, defined below. Although this is a second order characterization only, it is comprehensive in that all time scales are examined, and is reliable in practice as it offers a view which is quasi-independent across scales. If the semi-experiment has changed the process significantly, this will typically make its presence felt at second order over some scales at least.

We will first describe the wavelet analysis in more detail, then give an example introducing the use of semi-experiments, and then introduce some important modeling ideas.

A. Wavelet Analysis

Wavelets have become a tool of choice in the analysis of traffic data because they are well suited to studying scale invariant properties. Thus, they are capable of 'dealing' with the known long-range dependent (LRD) properties of packet counts (and other time series), whose difficult statistical properties can cause many other statistical tools to perform poorly both in the measurement of scaling parameters such as the Hurst exponent, and more generally, for example, in terms of robustness to non-stationarity. However, the properties of wavelets which make it effective for LRD are also useful for other reasons. In particular, their ability to de-correlate data means that they provide a way of isolating and examining behavior (be it LRD or not) separately at different time scales. Thus, they are an invaluable investigative tool to help see 'what is happening' in a time series at different scales. We use them in this sense in this paper. LRD is observed, but is not a focus of the work. We use the wavelet analysis code freely available at [14].

Performing the Discrete Wavelet Transform (DWT) of a process X consists in computing coefficients that compare, by means of inner products, X against a family of functions, that is

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

The wavelets $\psi_{j,k}(t) = 2^{-j/2}\psi(2^{-j}t - k)$ derive from an elementary function ψ, called the mother wavelet, dilated by a scale factor $a = 2^j$ and translated by $2^j k$. They are required to have excellent localization properties jointly in time and frequency. A key practical advantage 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 (note that the means of wavelet coefficients are identically zero) of its wavelet coefficients satisfies:

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

where $\Psi(\nu)$ denotes the Fourier transform of ψ. In fact, equation (2) can be viewed as defining a kind of wavelet energy spectrum, analogous to a Fourier spectrum, but much better suited to the study of long-range dependent processes. We will also use the term *wavelet spectrum* to refer to equation 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 is sometimes called the *Logscale Diagram* (LD):

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

The thick gray curve in figure 9 represents the LD of the measured packet arrival process. The vertical lines mark 95% confidence intervals on the estimation of $\mathbf{E}|d_X(j,k)|^2$. The horizontal axis is calibrated both in scale a (top edge of plot, in seconds), and octave $j = \log_2 a$.

It is important to note the following three facts about the wavelet spectrum:

- the spectrum of a Poisson process of intensity λ is flat: $\mathbf{E}|d_X(j,k)|^2 = \lambda$,
- for a *simple* point process (one whose points are isolated), in the limit $j \to -\infty$ of small scales, $\mathbf{E}|d_X(j,k)|^2 \to \lambda$, where λ is the average arrival intensity,
- if the LD of a process X is $L(j)$, then that of a superposition of N i.i.d. copies is simply $\log_2(N) + L(j)$: the (log) spectrum simply moves up.

Combining the first two, we learn that when comparing the LDs of different traces, in the limit of small scales they will asymptotically reach values, which depend on their average arrival intensities.

B. Introduction to Semi-Experiments

In [12], [15], [16] this method was used to explore the role of the flow arrival process $Y(t)$ in the structure of the packet arrival process $X(t)$. A long list of manipulations were performed, modifying aspects such as the flow arrival process, the internal dynamics of flows, and the number of packets per flow. For this introduction, we restrict ourselves to two of the most important used in the works above, and illustrate them with results from an Auckland IV data set [17] like the one used in [15]. The thick gray line in figure 9 shows the LD of the data. The constant slope (relative to confidence intervals) above scales of 2[sec] corresponds to LRD.

The first semi-experiment employs a manipulation of Y:

- **[A-Pois]:** Re-position flow **A**rrival times according to a **Pois**son process with the same rate and randomly permute the flow order. Flows are translated to their new starting points without having their internal packet structure altered.

For the Auckland trace of figure 9, the **[A-Pois]** manipulation completely erases the original flow arrival process Y **and** removes inter-flow dependencies. Despite this radical removal of structure, the resulting LD is barely distinguishable from the original. It follows that not only can Y be taken as Poisson, but also that flows can be treated as *independent*, eliminating the need to consider session level structure to explain or model X. Note that this does not contradict as such the presence of closed loop effects such as TCP flow control. It simply means that we have *observed* that the dependencies due to any such feedback may be ignored, for the purpose of describing and understanding the aggregate statistics of $X(t)$.

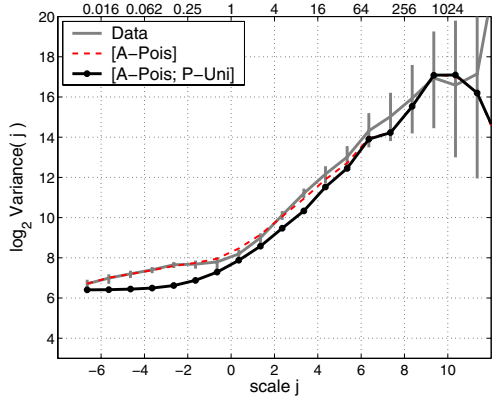

Fig. 9. Semi-experiments **[A-Pois]** and **[A-Pois;P-Uni]** on a typical Auckland IV trace

The second semi-experiment manipulates the structure of packets within flows:

- **[A-Pois; P-Uni]:** In addition to **[A-Pois]**, within each flow separately, **P**acket arrival times are **Uni**formly distributed between the original arrivals of the first and last packet of the flow. Flow durations and packet counts are preserved.

Looking at figure 9, at small scales this manipulation flattens the spectrum to its small scale Poisson limit (recall the properties discussed in section III-A). Compared with **[A-Pois]**, the removal of in-flow burstiness has reduced energy (variance) over small scales without significantly affecting large scale behavior. This indicates that the energy in the data above a neutral Poisson model, at small scales, is due to the burstiness within flows, and again not to dependencies between flows or details of flow arrivals, and also that in-flow burstiness is not the cause of LRD.

As we see in section IV, the choice of semi-experiment is guided by observation in a truly experimental procedure where one tried to pinpoint as precisely as possible the structural elements most responsible for the main statistical features. This requires the constant invention of new semi-experiments.

C. Cluster Processes as Models

Based on the semi-experiments just described and a number of others which reinforced and extended the conclusions above, in [15] a Bartlett-Lewis point process (BLPP) was proposed as a model for $X(t)$. A BLPP is a Poisson cluster process [18], that is it consists of a Poisson process defining the locations of 'seeds', about which independent and identically distributed (i.i.d.) clusters of points are placed. Let the arrival times $\{t_F(i)\}$ of flows (the seeds) follow a Poisson process of rate λ_F. The packet arrival process can be written as

$$X(t) = \sum_i \mathcal{G}_i(t - t_F(i)), \tag{3}$$

where $\mathcal{G}_i(t)$ represents the arrival process of packets within flow i. In the particular case of the BLPP, a cluster is a finite renewal process consisting of a random number $P \geq 1$ of points (including the seed) with inter-arrival time variable A. $\mathcal{G}_i(t)$ then reads

$$\mathcal{G}_i(t) = \sum_{j=1}^{P(i)} \delta\left(t - \sum_{l=1}^{j-1} A(i,l)\right), \tag{4}$$

where $A(i,l)$ denotes the $l-$th inter-arrival for flow i (the inner sum is zero if $j = 1$) and $P(i)$ is the number of packets in flow i. In [15] a choice of gamma distributed inter-arrivals A, with mean μ and shape parameter $c > 1$, was found to account in a simple way for the observations made on in-flow burstiness. A heavy tailed (infinite variance) choice of P accounts for the long-range dependence.

The BLPP model is useful to us here as a reference from which to understand traditional Internet traffic, and to compare against our findings from wireless data. One can go beyond BLPP models in several ways. One way is to allow greater structure for the seed process. In [16] that observation was made that Y is **not** Poisson but in fact is long-range dependent! however, for the purposes of modeling X, this source of LRD could be neglected compared to the main source, namely the heavy tailed nature of P. We will show in this paper that the role of Y is far from negligible in the wireless context.

IV. Data Analysis

In this section we offer results and interpretations based on the methodology of semi-experiments defined above. It is useful to structure the results into two sets of comparisons, first Wless-OUT and Wless-IN, then Wless-IN and Backbone, in order to structure the expected differences and similarities.

Due to the architecture of the wireless access network and the placement of the measurement point at a concentration point of traffic, the same flows are seen by both Wless-OUT and Wless-IN. They are therefore guaranteed to be alike in many respects, enabling us to interpret differences as being in large part due to features of the wireless network, and the subsequent processing before the measurement point.

In contrast, the Backbone trace has nothing to do with the measurement point, and its flows will be completely different. However, since Wless-IN consists of traffic which traversed the Internet, it will share, in a general sense, a networking environment with Backbone. We might then expect differences to be mainly due to application mix, and the effect of slower flow control acting through from the wireless side, but not to details of the wireless link layer.

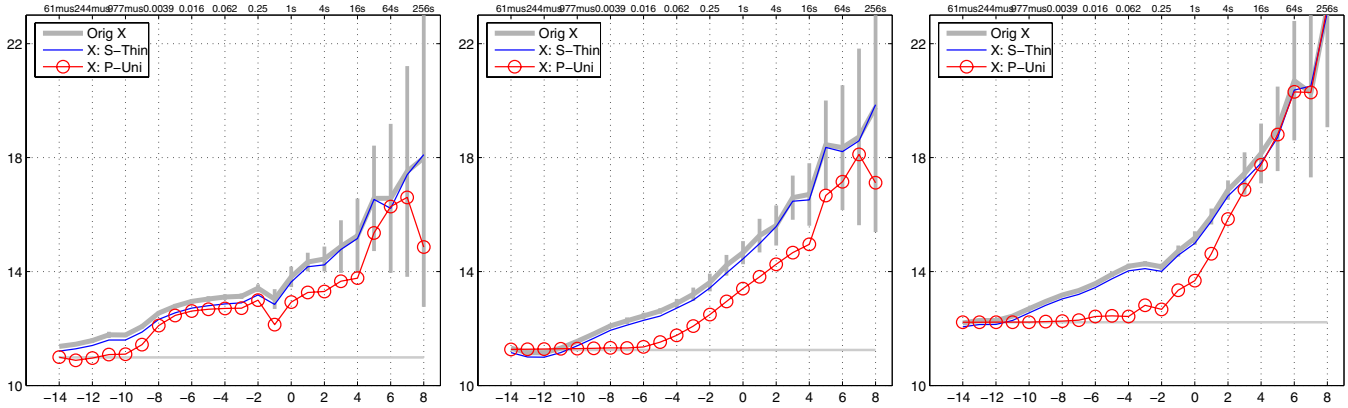

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 10. Flow and packet level manipulations on X. Examining flow independence and impact of in-flow burstiness.

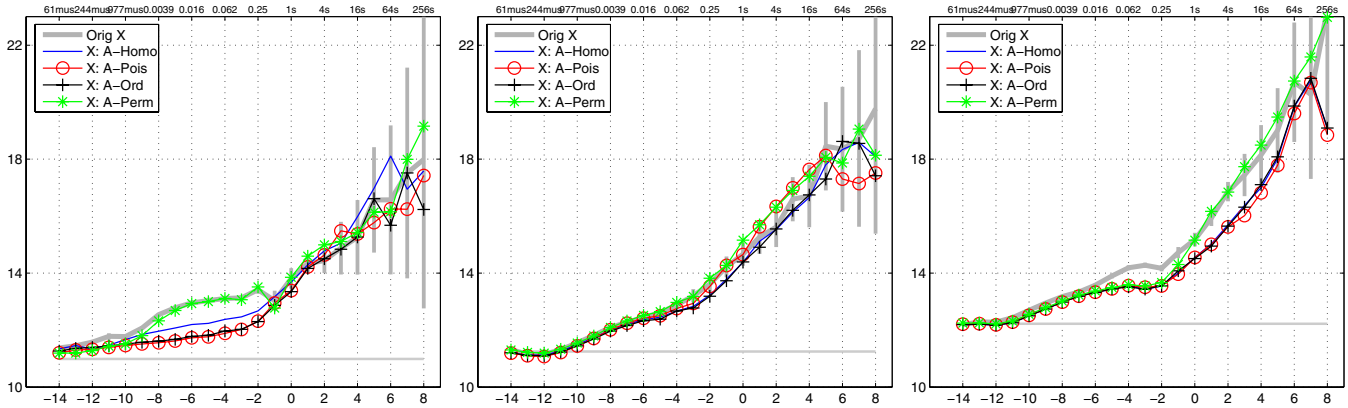

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 11. Flow level manipulations on X. Examining interaction between flows, and the structural importance of Y.

In the following, figures 10 to 17 present the results produced by the semi-experiments methodology on the 3 traces studied, respectively from left to right Wless-OUT, Wless-IN and Backbone. To improve the readability of those results we start the analysis by comparing Wless-OUT to Wless-IN.

A. Comparing Wless-IN and Wless-OUT

We begin the analysis of the wireless traces by the examination of the results of semi-experiments acting at packet and/or flow level on $X(t)$. We then focus on the examination of the consequences of the semi-experiments on $Y(t)$.

1) The Packet Arrival Process X: The first manipulation, **[S-Thin]**, is designed to directly test for the independence of flows. It consists of a thinning of the flow process by selecting flows randomly with some probability, here equal to 0.9. The results are similar for both Wless-OUT in figure 10(a) and Wless-IN in figure 10(b). The energy plot has the same shape, but has dropped by a small amount close to $\log_2(0.9) = -0.15$, just as one would expect if flows were independent. Looking ahead briefly to the Backbone trace in figure 10(c), we see that the same conclusion also holds there, just as it did for the Auckland trace of the previous section. At least in this important respect then, we see that *wireless traffic is similar to wireline traffic*.

The second manipulation is **[P-Uni]**, which as we know from the previous section, uniformises the packet arrivals within flows (since it is not combined with **[A-Pois]** here, Y is not affected), leaving intact both the duration and the number of packets in the flow. At scales below $j = -10$ we again see the same result for each of Wless-OUT and Wless-IN: the excess energy above the asymptotic Poisson level corresponding to the average arrival intensity (marked as the horizontal gray line) has been eliminated, revealing that it was due to the organization of packets within flows. For Wless-IN however, this zone continues up to $j = -6$, whereas an energy 'bump' appears clearly in Wless-OUT over scales $j \in [-9, -2]$. This unexpected bump of energy in the wavelet spectrum of Wless-OUT is only slightly affected. We conclude that *the main source of energy behind this bump is **not** in-flow burstiness*.

We now move to figures 11(a) and 11(b) where we use a set of manipulations affecting **flow arrivals**. First consider **[A-Pois]** from the previous section. Although it has little effect for Wless-IN, it effectively removes the energy bump in Wless-OUT (over scales $j \in [-9, -2]$). This is another strong indication that in-flow burstiness is not responsible for it, and points

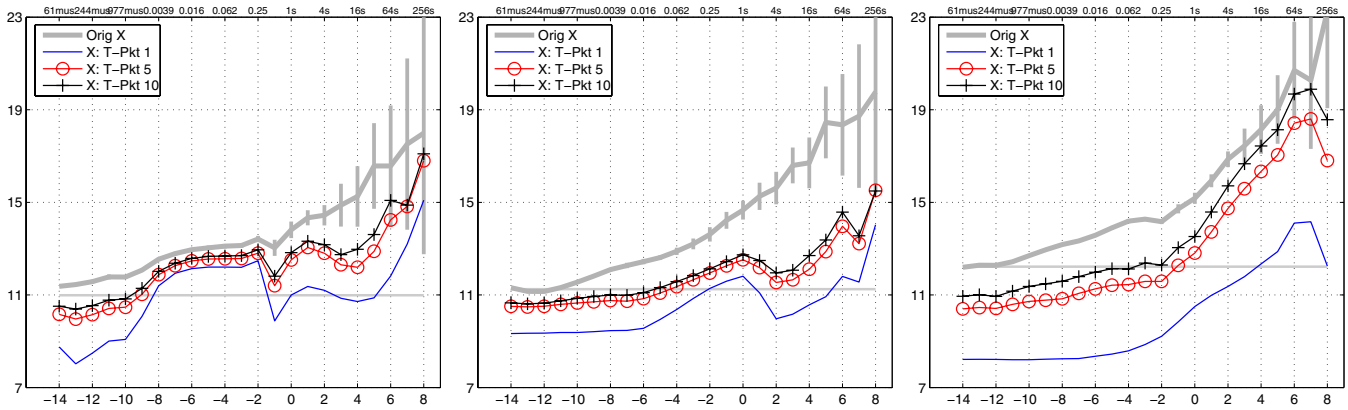

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 12. Packet level manipulations on X. Examining the effect of packet weighting of the Y 'skeleton'.

the finger at Y. However, since **[A-Pois]** is a very drastic manipulation, it is not yet clear what exactly does cause it.

[A-Ord] re-positions the points of Y according to a Poisson process, just as **[A-Pois]** does, but unlike it does not permute the flow bodies. Dependencies between flows are therefore perturbed at small scales, but not destroyed. Despite this much gentler modification of the original process, the effect on Wless-OUT is indistinguishable compared to **[A-Pois]**. Reducing the severity of the manipulation still further, **[A-Homo]** replaces Y with a renewal process rather than a Poisson process, based on the empirical inter-arrival distribution of Y (in fact the actual inter-arrivals are used, but permuted). Despite retaining the correct inter-arrival marginal in Y, much of bump still vanished, indicating that flow inter-arrivals cannot be taken as independent. In other words, the distribution of flows inter-arrivals is not sufficient to model accurately Y. Finally, moving in the opposite direction to **[A-Ord]**, **[A-Perm]** retains Y, but like **[A-Pois]**, permutes the flow bodies about these points. The result is almost indistinguishable from the original plot, a result which again argues for the independence of flows. As expected, each of **[A-Ord]**, **[A-Homo]** and **[A-Perm]** has little effect on Wless-IN, since they are all lesser 'sub-manipulations' of **[A-Pois]**.

The results above lead to two conclusions:
(i) *flows are independent in both Wless-IN and Wless-OUT*
(ii) *the structure of the flow process Y has a strong influence on Wless-OUT, but negligible influence on Wless-IN*

Following our observation of the importance of Y, we introduce a new family of semi-experiments that focus on it whilst still studying X. The **T**runcation manipulation **[T-Pkt** n**]** consists in keeping only the first n packets in each flow. With this definition, **[T-PKt 1]** collapses simply to Y. Note that this manipulation preserves the total number of flows, but removes packets, and therefore energy, from the process. The resulting wavelet spectra accordingly drop to lower values. LRD is also affected, since its main source, the heavy tail of the number of packets per flow, has been cut off.

In figure 12 we show the results of the new semi-experiments. The plots are necessarily nested as n decreases. The fact that the energy in the range of scales associated with the bump remains high even for **[T-PKt 1]** confirms our suspicion that it is dominated by Y.

It is interesting to note that, with increasing n and therefore a reduced 'masking effect' due to packets in flows, we are able to see an energy bump even in Wless-IN, although it is smaller in magnitude and occurs at larger scale (in fig. 12(b) over scales $j \in [-5, 1]$). We speculate that this bump, which is due to link layer protocol or shaping effects on the wireless side, impacts the traffic returning through the Internet, but that this 'reflection' is sufficiently smoothed out that it no longer has an appreciable effect. The truncation manipulation however was able to reveal it. This could potentially be used as a means of identifying traffic originating from wireless access networks.

2) Flow Arrival Process Y: In this section we try to determine more precisely which aspect of Y is responsible for the energy bump we observed. This knowledge will help in the search for, and in the future validation of, network level causes, and ultimately allow a more precise model to be developed. For the same reason we also wish to understand more clearly how Y impacts on X, as it is only of interest to model those aspects with high impact.

We know from **[A-Homo]** that although the marginal inter-arrival distribution of Y is very important, so are correlations between different interarrivals. We now design manipulations to break the latter down in different ways.

The first manipulation type, **[A-LocUni** t**]**, provides a **Loc**al smoothing by **Uni**formising points of Y on a time scale of t. More precisely, we segment the time axis into adjacent windows of width t, and for each window uniformly re-position the points of Y which fall within it. Figure 13 shows the results for 3 values of t. We also show X and Y separately, each with its respective asymptotic Poisson lines. By comparing X and Y (upper and lower gray curves respectively), we can see directly the impact of Y on X in terms of pure energy, that is viewing Y as a, generally small, subset of the points of X.

The results shown in figures 13(a) and 13(b), as expected, show the removal of energy above the Poisson level due to the low pass filtering, thereby progressively eroding the bump as t is increased. However, the energy is removed only up to the timescale t of the manipulation, which is to be expected, but

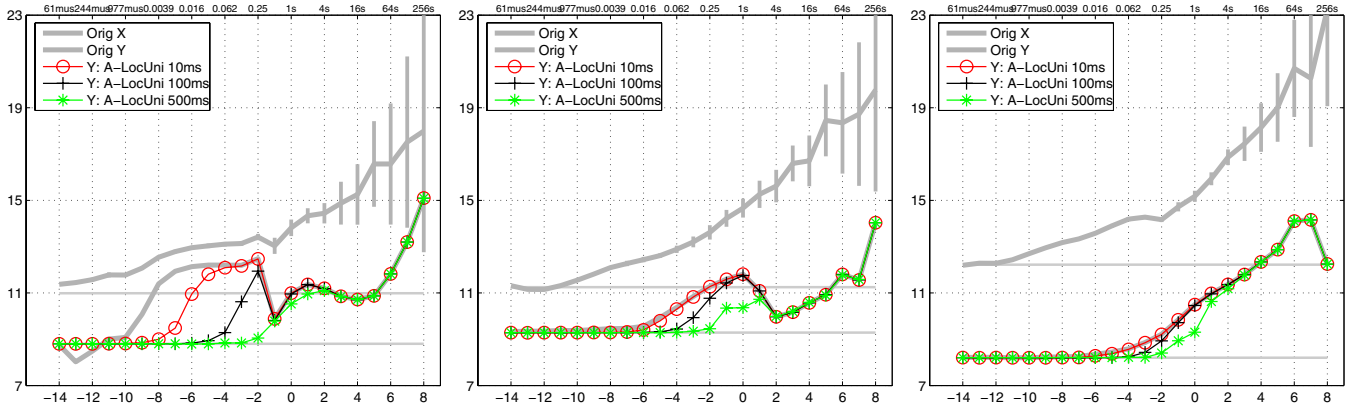

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 13. Time-based smoothing manipulations on Y only, with X as a comparison.

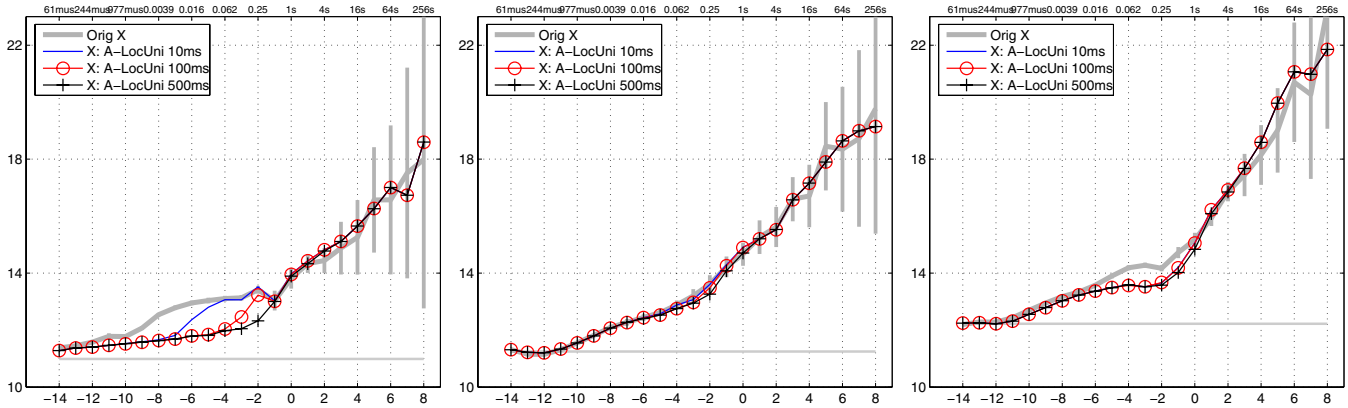

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 14. Time-based smoothing manipulations on X, showing the effect of adding in flow bodies.

not beyond. Thus the manipulation fails to find a particular structure which underlies the energy bump and which could be 'surgically' removed. The corresponding plots in figures 14(a) and 14(b), where the same manipulations on Y are performed but we observe the resulting impact on X, is nonetheless informative. It reinforces the conclusions of the previous subsection, that since the flow bodies act independently, we can understand X in terms of a simple packet weighting of the 'skeleton' Y. If Y itself has too little energy, such as in figure 14(b), details of its structure are lost in X. If it does have enough, such as in figure 14(a), the impact is nonetheless smoothed and reduced by the addition of the flow bodies.

We now move from a time based local randomization of Y to one based on packet index. The manipulation type **[A-IAPerm** n**]** acts on **I**nter**A**rrivals of Y, **Perm**uting them randomly within non-overlapping groups of n consecutive flows. Note that once n is large enough to include all of Y, then this manipulation reduces to **[A-Homo]** defined previously.

Figure 15 shows the results of **[A-IAPerm** n**]** on Y for various n. In contrast to the time based randomization, here we see the energy bump in figure 15(a) targeted much more precisely. For each n, energy is lost across the entire range of scales defining the bump. This suggests that a clustering of Y is responsible for the energy bump observed, which we are succeeding in breaking up. From the values of n, it seems that most clustering of Y occurs in groups of around 25, but that such involve as many as 100 flows. Interestingly, these values are different, in fact higher, for Wless-IN (figure 15(b)). This may be related to the mechanism which pushes the bump to higher scales, and requires further investigation. Figure 16 shows the results of the **[A-IAPerm** n**]** applied on Y but observed on X. In the same way as for figure 14, we observe that the structural characteristic of Y is smoothed by the addition of the flow bodies.

Figure 17 replots selected manipulations to assist in summarizing some of the main findings of this section, and to give insight towards a model. Comparing Y to **[A-Homo]** applied to Y in figure 17(a) we see that a renewal process model for Y with the correct marginal captures the small and large scale behavior, but does not capture the energy bump, which is due to clustering of Y as we have just seen. Correspondingly, if we used such a model of Y to form X, that is if we perform **[A-Homo]** on X, we obtain a good match, with the exception of the bump. The same results hold for Wless-IN (figure 17(b)), except that there the renewal model for Y degenerates to Poisson, and the packet weighting overwhelms the bump, so that it would not in fact be necessary to model it, resulting in a Pois-

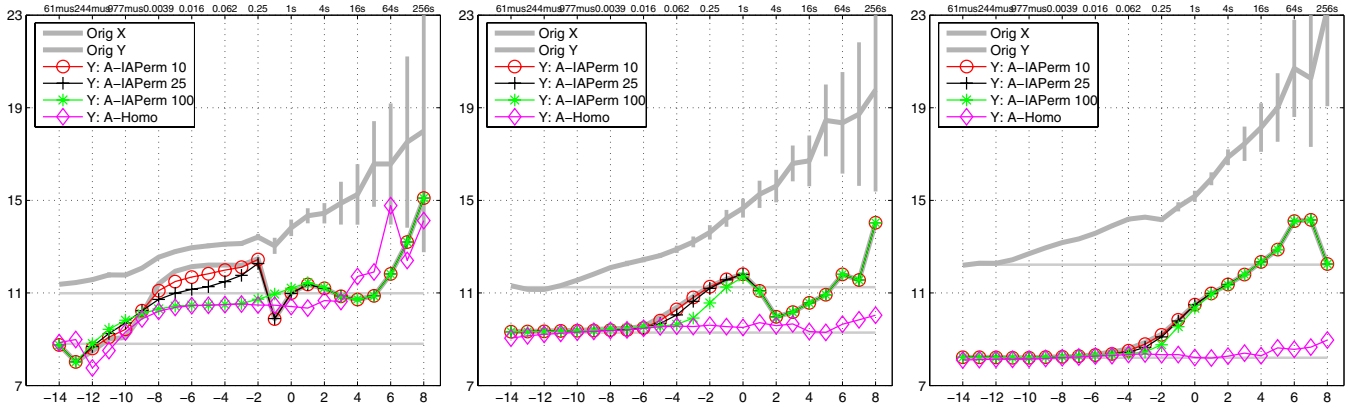

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 15. Packet index-based smoothing manipulations on Y only, with X as a comparison.

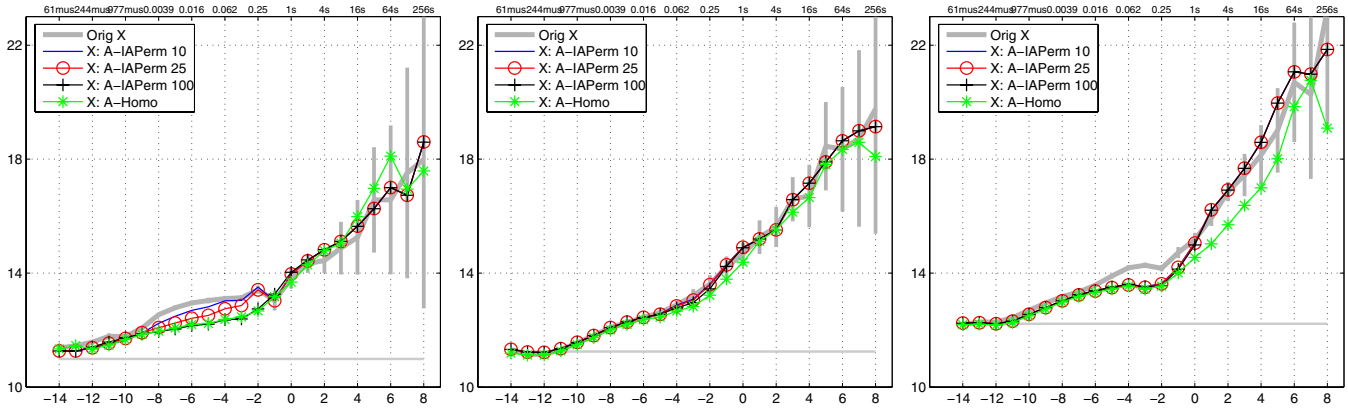

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 16. Packet index-based smoothing manipulations on X, showing the effect of adding in flow bodies.

son model for Y in this case. *The Poisson Process hypothesis remains then valid for the Wless-IN traffic but is more difficult to retain for the Wless-OUT traffic.*

B. Comparing Wless-IN and Backbone

The results of the semi-experiments for the Backbone trace appear as the third plot in the figures above (figures 10 to 17). A quick examination shows that they are much closer to the result for Wless-IN than to Wless-OUT, as expected. We accordingly focus on comparing Wless-IN and Backbone.

1) Packet Arrival Process X: As already noted above, the results for the Backbone trace in figure 10(c) are entirely consistent with those of Wless-IN (figure 10(b)) and of previous work [15]: *flows can be taken as independent, and the structure at small scales is due to in-flow burstiness*.

To first order, in figure 11 we also see a consistent picture across Wless-IN and Backbone, where flows are independent, and details of Y do not matter. There is a subtlety, however, relating to a small energy bump centered around octave $j = -3$. The result of **[A-Perm]**, which is extremely close to X except at the bump, shows that it is related to the position of flow bodies (but not to their internal structure, as the bump is visible in **[P-Uni]**) rather than to the structure of Y. On the other hand, the bump is also eliminated by the other manipulations which act directly on Y (but which *do* however perturb flow body positions). We speculate that this bump may be due to a small class of flows, with related starting times, which carry an unusually large proportion of packets. For example, they could be from a set of clients with direct access to a high capacity link.

Figure 12 shows a clear difference between wireless and wireline. There is no evidence of an energy bump at small scales related to Y. On the other hand, there is evidence that Y has long-range dependence. These findings agree with those of [16], [15].

2) Flow Arrival Process Y: The results of figure 13 show that the local time-based smoothing of Y makes very little difference, reinforcing the conclusion that for backbone traffic Y has little small-scale structure. Consequently, the impact on X in figure 14 is even lower, indeed negligible.

Figure 15 shows a similar picture under the packet-based local smoothing. The impact on Y is negligible until **[A-Homo]** radically eliminates the LRD in Y. Despite this, the impact on X as seen in figure 16 is small. These again reinforce the modeling idealization of taking Y as Poisson for backbone traffic. Figure 17 gives a direct illustration of the impact of such an

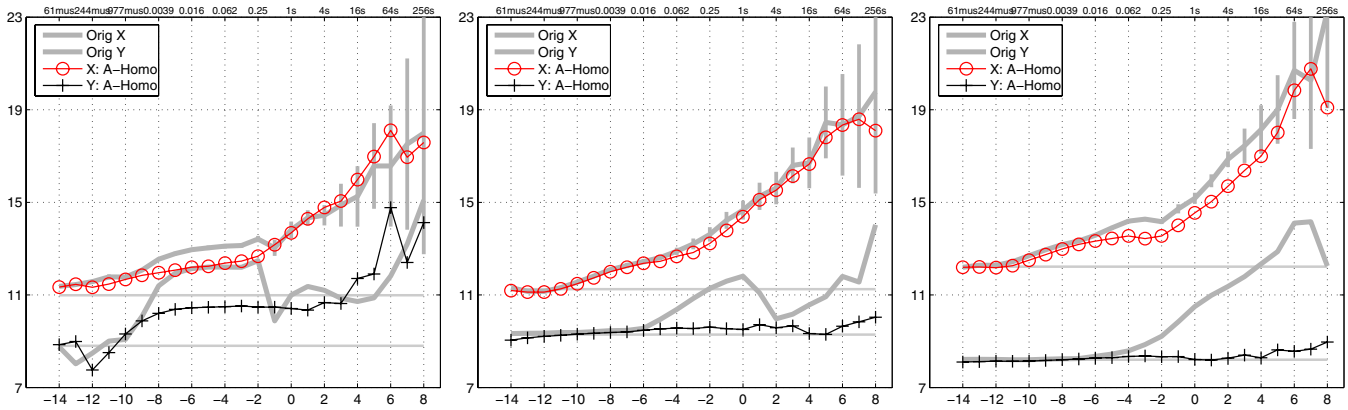

(a) Trace Wireless OUT (b) Trace Wireless IN (c) Trace Backbone

Fig. 17. Manipulations summarizing the effect of Y, and a renewal model of Y, on X.

assumption. Figure 17(c) shows indeed how the **[A-Homo]** manipulation applied on Y derived to a Poisson Process and that the overall impact when adding flow bodies is small. The same assumption may be more dangerous for wireless traffic. In the case of Wless-IN, we see in figure 17(b) that the assumption of modeling Y with a Poisson Process works even better than for the Backbone traffic. Nevertheless, in the case of Wless-OUT, because the energy bump although not having a strong influence on X in figure 17(a), is close to the spectrum of the overall traffic, it may influence it in other circumstances or for particular applications of this ideal Poisson model.

C. Discussion and Future Directions

In this section we observed and characterized the structure of wireless traffic and compared it to prior results obtained in the wired traffic area. The set of data we had the chance to observe covers continuous 24 hours period. Even with such a limited data set, we found possible evidences of wireless traffic peculiarities. The analysis, so far, mainly adopted a statistical point of view. Here we try to offer some clues about the possible causes of the observed properties of the traffic and about the models that can be derived from the present study. As we saw, the flow arrival structure of the wireless traffic may not be negligible when designing models of traffic for the wireless access network. Obviously, some models one can design may only require a "loose" traffic estimation. A model of traffic based on a BLPP process where flows are treated as independent will be satisfying to capture most of the structure of the entire traffic (the packet arrival process). More precise models would require to capture the bump observed in the LDestimate spectrum. The "wavy" behavior we observed in the distribution of the flows inter-arrivals led us to characterize the clustering present in the flow process Y.

In order to build more accurate models, a deeper analysis in larger data sets is required. The set of data at our disposal does not permit to clearly identify network causes. Future work will focus on trying to identify the causes of the observations made. Three main directions will be investigated to determine these network causes:

- We first will study the impact of the air interface on the traffic. The channel access mechanisms, the radio channel setup delay, the rate adaptation of CDMA-1xRTT layer may be the possible causes of the particular clustering structure we observed.
- Our second intuition targets the role of the intermediate "black boxes" present in the Customer Data Network. Those boxes working at the packet, flow and application levels have proxy and transcoding functionalities that may have a specific impact on the wireless traffic.
- Finally, the diversity of devices and applications, and the particular user behaviors on a wireless access network are surely of high interest in understanding the reality of the wireless traffic.

Understanding the consequences of the three directions of research mentioned above is surely crucial from a wireless access operator perspective. As a future work we will validate (or invalidate) the two first possible causes by adding a measurement point close to the air interface (at the Base Station Controller). Having two measurement points will allow us to observe the characteristics of the traffic at the air/CDN interface and at the CDN/core network interface. We will then be able to estimate the influence of the wireless channel on the traffic characteristics and, at the same time, it will allow us to quantify the impact of the "black boxes" on the aggregated traffic. As the architecture deployed on the Sprint network is surely close to the one of any other wireless operator, this future study will help to understand better the interaction of those different networking elements.

Finally, we will observe the characteristics of specific groups of devices or applications in use on the wireless network. We will then be seeking for a possible classification of the devices or application based on their specific contribution (if it exists) to the characteristics of the traffic. Location and roaming informations could be added in order to observe the impact of users mobility on a large scale wireless access network. These analyses remain difficult to realize as they have to cope with an access to data coming from different layers of the architecture (physical layer, network layer, AAA functionality, mobileIP management) and correlate all this information. To be realized, the

setup of all those monitoring tools is highly time consuming on a commercial network, while ensuring data anonymity at the same time.

V. Conclusion

We examined precision time-stamped TCP/IP traffic at a concentration point where all traffic passing between a portion of a wide area CDMA wireless network and the Internet was visible, capturing traffic from a mix of end devices, mobility profiles, and applications. The traffic was examined in detail using wavelet energy plots, which allow a vision of the statistics of the data over all time scales. The semi-experimental method was used to observe differences between the incoming and outgoing traffic, and a benchmark backbone Internet trace.

Our first observation is that, to a large extent, the wireless and Internet data was quite similar. Clear differences were nonetheless observed. Data originating from the wireless network contained energy in a range of scales in the millisecond range. We were able to trace this to the flow arrival process, and more specifically, to clustering of neighboring arrivals in groups of the order of 50 members. We found that the packet arrival processes within flows were not a major factor, and that flows themselves could be treated as independent, confirming recent work on Internet traffic as well as the results from the backbone trace. In the trace containing data from the Internet to the wireless network, a related energy 'bump' was also found, however it appeared at larger scale, and had reduced amplitude. As a result, it no longer affected the statistics of packet arrivals significantly; however it was still clearly visible in the flow arrival process, opening up the possibility of using this as a signature of wireless traffic which could be used as a basis of detection. The signature in the flow arrival process for the backbone trace is quite different: there is no energy bump, but instead a long-range dependence which may be neglected as the entire flow arrival process has low energy compared to the weight of packets in flows.

Our findings have important implications for modeling of wireless traffic. First, in contrast with Internet traffic, the flow arrival process has significant structure which must be taken into account, and cannot be taken as a Poisson process, in contrast to Internet traffic in general. Second, in common with Internet traffic, the parsimonious choice of independent flows with simple internal burstiness structure represents aggregate statistics adequately. Thus, there is scope for an accurate model which also has few parameters and will be simple to simulate.

In future work, we hope to incorporate more contextual information allowing the nature of end devices to be identified. This will open up a rich new field for exploration where network causes will be able to be linked to the 'statistical causes found here'. We will also develop a model of the packet and flow arrival processes.

References

[1] M. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," in *Proc. ACM Sigmetrics*, may 1996.

[2] M. Grossglauser and J.C. Bolot, "On the Relevance of Long-Range Dependence in Network Traffic." *Computer Communication review*, vol. 26, no. 4, pp. 15–24, october 1996.

[3] 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*, march 2003.

[4] A. Balachandran, G.M. Voelker, P. Bahl and P. Venkat Rangan, "Characterizing user behavior and network performance in a public wireless LAN," *SIGMETRICS Perform. Eval. Rev.*, vol. 30, no. 1, pp. 195–205, 2002.

[5] D Kotz and K Essien, "Analysis of a campus-wide wireless network," in *Proc. ACM MobiCom*. ACM Press, September 2002, pp. 107–118.

[6] X. Meng, S. Wong, Y. Yuan and S. Lu, "Characterizing Flows in Large Wireless Data Networks," in *Proc. ACM Mobicom*, september 2004.

[7] D. Schwab and R. Bunt, "Characterizing the use of a Campus Wireless Network," in *Proc. IEEE Infocom*, march 2004.

[8] T. Henderson, D. Kotz and I Abyzov, "The changing usage of a mature campus-wide wireless network," in *Proc. ACM MobiCom*. ACM Press, September 2004, pp. 187–201.

[9] R. Chakravorty, J. Cartwright and I. Pratt, "Practical Experience with TCP over GPRS," in *Proc. IEEE Globecom*, november 2002.

[10] P. Benko, G. Malicsko and A. Veres, "A Large-scale, passive Analysis of end-to-end TCP Performances over GPRS," in *Proc. IEEE Infocom*, march 2004.

[11] http://ipmon.sprint.com/.

[12] 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 (IMW-2002)*, Marseille, Nov 2002, pp. 63–68.

[13] A. Erramilli, O. Narayan and W. Willinger, "Experimental Queueing Analysis with Long-Range Dependent Packet Traffic," *IEEE/ACM Transactions on Networking*, vol. 4, no. 2, pp. 209–223, April 1996.

[14] D. Veitch and P. Abry, "Matlab code for the wavelet based analysis of scaling processes," http://www.cubinlab.ee.mu.oz.au/∼darryl/.

[15] N. Hohn, D. Veitch and P. Abry, "Cluster processes, a natural language for network traffic," *IEEE Transactions on Signal Processing, special issue "Signal Processing in Networking"*, vol. 51, no. 8, pp. 2229–2244, August 2003.

[16] N. Hohn and D. Veitch and P. Abry, "The Impact of the Flow Arrival process in Internet Traffic," in *Proc. IEEE ICASSP 2003*, Hong Kong, April 2003, pp. VI 37–40.

[17] Waikato Applied Network Dynamics, http://wand.cs.waikato.ac.nz/wand/wits/.

[18] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*. Springer-Verlag, 1988.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the Proceedings IEEE Infocom.