On the role of flows and sessions in Internet traffic modeling: an explorative toy-model

Fabio Ricciato, Angelo Coluccia, Alessandro D'Alconzo, Darryl Veitch, Pierre Borgnat, Patrice Abry

On the role of flows and sessions in Internet traffic modeling: an explorative toy-model

Fabio Ricciato[*†], Angelo Coluccia[*], Alessandro D'Alconzo[†], Darryl Veitch[‡], Pierre Borgnat[§], and Patrice Abry[§]

[†] Forschungszentrum Telekommunikation Wien (ftw.)
[*] University of Salento, Lecce, Italy.
[‡] CUBIN (ARC Centre for UltraBroadband Information Networks) — University of Melbourne, Australia
[§] CNRS, Laboratoire de Physique ENS Lyon, Université de Lyon
Corresponding Author's Email: ricciato@ftw.at.

Abstract—In this work we present a simple toy-model that is able to explain certain empirical observations reported in a set of previous papers by Hohn *et al.* [1]–[3] about the wavelet spectrum of real traffic traces. Therein, the authors found that the wavelet spectrum is substantially invariant to flow scrambling and truncation, suggesting that super-flow structures above the transport layer — i.e., sessions — can be ignored for modeling the packet arrival process. Based on the proposed toy-model, we offer an interpretation framework that goes in the opposite direction, wherein sessions, not transport-layer flows, should be taken as the main structural entities in simplified on/off models.

I. INTRODUCTION

The problem of modeling Internet traffic has attracted considerable interest in the research community since more than a decade. Although to date there is no commonly accepted model, two fundamental properties of Internet traffic have found extensive empirical evidence: Long-Range Dependence (LRD) of the aggregate traffic, and the ubiquituous presence of heavy-tailed (HT) distributions at different levels of granularity (origin-destination pairs, Layer-4 flows, sessions, etc.). Moreover, HT is widely accepted as the principal cause of LRD [4, p. 225], following the fundamental result by Taqqu [5, p. 26]. Still the problem of identifying a parsimonious *structural model* for Internet traffic remains open. The difficulty lies in the intrinsic complexity, heterogeneity and variability of modern Internet traffic.

Since the introduction of the wavelet analysis [6] it has become common practice to consider the wavelet spectrum as a key characteristic of the traffic process. The motivation for such approach lies in the fact that the temporal correlation structure of the traffic process plays a role in some engineering tasks, e.g. buffer dimensioning — at least for a limited horizon of timescales [7]. A popular way to represent the wavelet spectrum is by means of the so-called Log-Diagram plot ("LD plot") [6]. As a matter of fact, the adherence of a model to the empirical LD plot of a real traffic trace has become a popular criterion to demonstrate the *goodness* of the model[1].

[1]Furthermore, the model ingredients that can emulate a *realistic* LD plot are sometimes taken as explanatory of the dynamics at play in real traffic, with some confusion between the evocative and explanatory nature of the model — see e.g. discussion in [8].

In [2] the Cluster Point Process (CPP) was proposed as a simple parsimonious model for Internet traffic at the packet level. The CPP model is based on a two-layer hierarchy, where packet arrivals are grouped into *clusters*. The packets-in-cluster pattern — a renewal process with Gamma distributed interarrivals was assumed in [2] — dictates the correlation structure at small timescales, while the distribution of the cluster size shapes the wavelet spectrum at large scales (typically LRD). In this way the CPP model can emulate the whole LD observed in real traces. The central aspect of CPP is that the clusters are assumed fully independent — they arrive according to a Poisson process.

The same authors devoted a series of papers to investigate the structure of Internet traffic with particular focus on the correlation between packets, flows and *sessions* [1]–[3]. It is important to remark that in all those works the flow classification follows strictly the canonical definition based on the IP 5-tuple ([source,destination] × [address,port], protocol). As a result, the notion of *flow* coincides with transport-layer entities, i.e. TCP connections and UDP streams (hereafter referred to as "L4-flows"). The methodology adopted in those works is based on the notion of "semi-experiments", described in the following. First, the packets are grouped into L4-flows based on the 5-tupla (classification), thus identifying a two-layer hierarchy — packets and L4-flows. Then, transformations are imposed to L4-flows such as truncation, arrival time reshuffling, and packet scrambling. Each type of transformation destroys specific components of the (unknown) structure present in the real trace. In this way it is possible to identify which structural components are responsible for the characteristic LD plot observed in the original process, and which ones instead are superfluous — at least with respect to the goal of reproducing a "realistic" LD.

The main finding of [1]–[3] is that inter-flow correlation can be eliminated — by flow reshuffling — without impacting the "important" correlation properties of the traffic aggregate, i.e. its LD profile. This suggests that super-flow structures above the transport layer — which we call *sessions* hereafter — are superfluous and can be ignored in the final model. *Notably, this does not mean that such higher structures are not present in real traffic, but only that they are ininfluential*

for modeling the packet arrival process.. In other words, it is not the existence of sessions that is questioned, but only their role in a parsimonious model.

In this work we critically revise this conclusion. More specifically, we do not contrast the two-layer model approach — packets and clusters, as in CPP [2] — but argue against the identification of the main structural entity (cluster) with L4-flows. Our arguments contain also an implicit criticism to the modelling methodology based on LD emulation.

Curiously, the key of our arguments lies in some empirical observations reported in [1]–[3] that the authors did not completely explain therein — we review them in detail in the next section. In this work we build an ultra-simplified "toy-model" that is able to consistently reproduce and explain all those observations. Based on such model, we achieve the opposite conclusion as in [1]–[3]: we find that the natural definition of cluster object shifts from L4-flows to a higher *session* level. In other words, our toy-model provides an alternative interpretation framework where *sessions, not L4-flows are the key structural entities* in a two-layer hierarchical model. We believe that such result — although still in a preliminary stage — should motivate a re-orientation of the traffic modeling research agenda towards a better understanding of super-flow structures above the transport layer.

The rest of this work is organized as follows. In §II we review some earlier works on wavelet analysis of real traces, highlighting a set of empirical observations left unexplained therein. In §III we present preliminarily some aspects of LD plots for on/off processes. In §IV we introduce formally our model and demonstrate analytically that it can consistently explain all the previous observations. We also validate our results via simulations. In §V we justify the key model ingredients. In §VI we discuss the consequences of the model, taken as a virtual interpretation framework for real traffic. Finally, in §VII we draw the conclusions and identify some points for future research.

II. REVIEW OF PREVIOUS WORKS

Here we review some key observations made in [1]–[3]. The flow definition therein is based on the canonical IP 5-tupla plus a timeout T: a sequence of IP packets with a common 5-tupla value and interarrival times smaller than T (e.g. 60 sec) are classified as belonging to the same *flow*, or more precisely *L4-flow*. The count process of packet and flow arrivals in fixed-length time bins are denoted by X and Y respectively (the time index is omitted for a more compact notation). The LD of a generic process Z will be denoted by $\mathfrak{L}(Z)$. The authors observe empirically that:

(i) $\mathfrak{L}(X)$ is *biscaling*: it scales linearly with slope h_x at large time-scales (LRD) while is approximately flat (i.e., it scales linearly with a very small slope) at small time scales, with an intermediate knee at scale q_x.

(ii) $\mathfrak{L}(Y)$ is *similar* to $\mathfrak{L}(X)$: biscaling as well, with same slope at large scales but different knee point, i.e. $h_y = h_x$ and $q_y > q_x$.

(iii) Invariance to scrambling: the qualitative behaviour of $\mathfrak{L}(X)$ does NOT change significantly if the flow arrivals are randomly scrambled, i.e. flows are relocated according to Poisson while keeping the same flow size.

(iv) Invariance to truncation: the qualitative behaviour of $\mathfrak{L}(X)$ does NOT change significantly if the flow size is truncated while keeping the same flow arrival times — note that the process Y can be considered as an extreme case of truncation, where the flow size is set to one packet.

(v) Effect of joint truncation and scrambling: $\mathfrak{L}(X)$ changes significantly and becomes almost flat if the flows are scrambled *and* truncated

In particular, observation (iii) is used in [1]–[3] to argue that the *L4-flow* arrival structure is NOT necessary to reproduce the characteristic LD plot of $X(t)$, and therefore for the purpose of modeling the packet arrival process there is no need to consider super-flow structures, i.e. *sessions*. This does not mean that sessions are absent from real traffic, but only that they can be ignored for modelling the traffic at the packet level.

In the following we present a simple synthetic toy-model that offers an explanation for all the above observations, but leads to an opposite conclusion with respect to the roles of flows and sessions. As our model is based on a simple extension of the basic packetized $M/G/\infty$ process, in the next section we provide a preliminary view of such process and of its LD spectum.

III. PRELIMINARIES: LD PLOTS OF $M/G/\infty$ PROCESS

We start considering a classical $M/G/\infty$ process where *sessions* arrive according to a Poisson process of intensity λ_s. In the packetized version each session generates quanta of information (i.e., packets) at fixed rate $\lambda_p = 1/T_p$. In the fluid version, sessions have a constant continuous transmission rate equal to r (without loss of generality we set $r = 1$). We are interested in the LD behaviour of the aggregate process, denoted by $X(t)$. Let us consider a discrete version where the time axis is divided into bins of length b (typically a few milliseconds), so that $X(t)$ counts the total amount of information sent in the t-th time bin by all active flows.

Taqqu's Theorem [5, p. 26] predicts that when the session sizes are heavy-tailed (HT) distributed the *asymptotic* behaviour of $X(t)$ at large time-scales is Long Range Dependent (LRD) and asymptotically Self-Similar (SS). The Hurst parameter H and the session tail index a ($a \leq 2$) are linked by a simple equation:

$$H = (3 - a)/2. \tag{1}$$

Remark that this result is valid only in the limit of large time-scales. Taqqu's theorem does not give any indication about the boundary of the large scales region, i.e. the onset of LRD — this issue is addressed in a separate recent contribution [9].

In the following we plot the LD of several realizations of $M/G/\infty$ process. LD are computed by the routine `LDestimate` available from [10] with three vanishing moments. In a first set of simulations we consider the fluid version

Fig. 1: LD for fluid $M/G/\infty$ with different session size distribution. The mean session size is constant for all four experiments. Sampling interval ($j = 1$) $b = 100$ ms. Number of flows $N = 10^5$. Reference slopes were added manually. The vertical line indicates the scale of mean flow duration.


Fig. 2: LD for fluid $M/G/\infty$ with session size distribution pareto, tail index $\alpha = 1.25$, for different values of location parameter k. Sampling interval ($j = 1$) $b = 100$ ms. Number of flows $N = 10^5$. Reference slopes were added manually. The vertical line indicates the scale of mean session duration.

of the process for different session size distributions. Note that in the following simulations the length of the total observation period D is several orders of magnitude larger than the mean session duration, which ensures that the LRD region is observable (see [9]). For HT we have used Pareto distributions with CDF $F(x) = 1 - (k/x)^\alpha$, mean $E(x) = k\alpha/(\alpha - 1)$, with tail index $1 < \alpha < 2$ and location parameter $k > 0$.

In Fig. 1 we plot the LD of $M/G/\infty$ for different session distributions with the same mean value: pareto $\alpha = 1.25$, pareto $\alpha = 1.6$, exp-neg and uniform. In Fig. 2 we consider pareto distributions with same tail index ($\alpha = 1.25$) and different mass location k — hence differen mean size. In all plots two different regions are clearly visible. At large time-scales (**L**-region) the LD spectrum is consistent with the theory: flat for light-tail distributions (Fig. 1c,1d) and close to a straight line in the heavy-tail cases (Fig. 1a,1b), with slope in agreement with eq. (1). At lower time-scales (**M**-region) the LD is straight with slope $\beta = 2$ in all cases. The existence of this region was already noted in [11]. According to [12] this is due to the discontinuities of the transitions between the on and off states of individual sessions, which leads to a trivial form of self-similarity at those timescales.

Fig. 2 reports the LD plots associated to session size distributions with the same tail-index but different mean, i.e. different mass location parameter k. It is evident that the change-point between the **M**- and **L**-region (denote by q_{ML}),

i.e. the onset of the LRD domain, depends on the location of the distribution mass: the longer are the sessions, the larger is the timescale q_{ML} at which LRD starts to be visible. In a separate recent work it is shown that the onset of LRD must be larger than the mean session duration (see [9, eq. (18)]), a result that in consistent with the plots in Fig. 1 and Fig. 2. Consequently, if the length of the observation period D is too short, one could not observe the **L**-region. This observation has some fundamental implications for the LRD analysis of real traffic traces (see discussion in [9]).

So far we have considered fluid sessions (continuous transmission). We now turn to packetized sessions (discrete transmission) where the packets-in-session arrivals are taken as Poisson for the sake of simplicity. Note that the packetized $M/G/\infty$ process is equivalent to a special case of the CPP model. The impact of packetization on the wavelet spectrum is limited to small time-scales, where the LD becomes flat. This leads to a the emergence of a third region, namely the **S**-region (small scales) below the **M**-region. Fig. 3 shows the complete LD for a packetized $M/G/\infty$ processs with exponential session size. The three regions are clearly visible:

- **S**mall scales: LD is flat due to the Poisson structure of packets-in-session.
- **M**edium scales: LD is linear with slope $\beta = 2$; the SS property in this region is due to the discontinuity of the
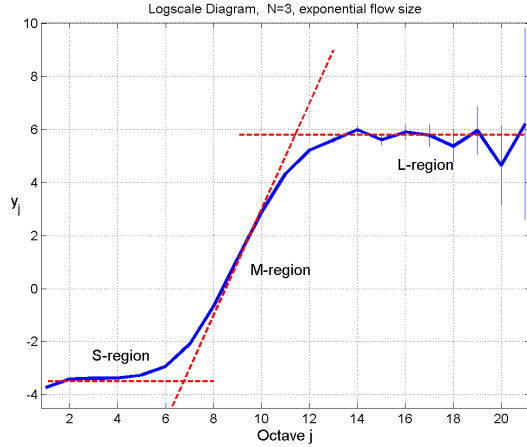
Fig. 3: The three timescale regions for a packetized $M/G/\infty$ process with exp-neg activity periods.

Fig. 4: LD plot of a biscaling process W and W+P.

on/off transitions (see [12]).

- **Large scales**: LD depends on the tail behaviour of the session size: LRD for heavy-tailed pareto, flat for light-tail exponential (as in Fig. 3).

IV. A TOY-MODEL

We consider a CPP process (or equivalently packetized $M/G/\infty$) similar to that introduced in the previous section. The sessions (clusters) identify periods of activities by individual users, therefore they are assumed independent and arrive as Poisson. Within each session, packets arrive as Poisson with rate λ_p. The flows are derived from the set of sessions as explained below.

We randomly classify the sessions into two classes, namely $\mathcal{S}_0$ and $\mathcal{S}_1$, with probablity ϕ. For simplicity we assume an equal probability classification, i.e. $\phi = 0.5$. By taking the packet count process for each class separately we identify the two sub-processes X_0 and X_1, with $X = X_0 + X_1$. Note that the child processes X_0 and X_1 are also CPP, and inherit the same session tail-index of the mother process X. In other words we have (randomly) decomposed the original process into two *equivalent* sub-processes. All three processes have *exactly* the same temporal correlation structure at all time-scales, hence the same LD.

We now move to build the flows. The count process for the flows will be denoted by Y. For the class $\mathcal{S}_0$ we simply map each session into a single flow. For the class $\mathcal{S}_1$ we fragment each session into a number of distinct flows. We consider here the extreme case where *each individual packet is mapped into a distinct flow* (maximal fragmentation). It is important to remark that we are considering here only the two extreme cases of session-to-flow mapping — one flow per session and one flow per packet — for the sake of simplicity and ease of explanation. The central features of the model should not change if more articulated, intermediate fragmentation schemes are considered. Later in §V we will motivate such

construction and discuss how it relates to concrete aspects of real traffic, protocols and applications in the modern Internet.

For each class $\mathcal{S}_0$ and $\mathcal{S}_1$ we define the associated flow count process Y_0 and Y_1, respectively. By construction $Y_1 = X_1$, while Y_0 is Poisson. Clearly, it holds that $Y = Y_0 + Y_1$.

In the rest of this section we show that such simple model is consistent with the set of observations listed previously in §II. In other words, if one applies the same analysis procedure of [1]–[3] — semi-experiment transformations and LD computation — to our toy-model, all those observations would hold.

A. Notation

It is convenient to introduce a few elements of notation for LD plots. We denote by P the Poisson process. Given a generic count process W (e.g. $W = X$ or $W = Y$) we denote by $\mathfrak{L}(W)$ its LD spectrum. If two processes W_i and W_j yield the same LD modulo a vertical offset we say that they are *LD-equivalent*. We will adopt the symbol "$\cong$" to indicate this condition. Hence:

$$\mathfrak{L}(W_i) = \mathfrak{L}(W_j) \Leftrightarrow W_i \cong W_j \qquad (2)$$

If two processes W_i and W_j are combined into the sum process W, we write: $W = W_i + W_j$ (the symbol "$+$" indicates superposition in this context). If the process W can be constructed as the superposition of k independent realizations of the process W_i, we write $W = kW_i$. Conversely, if the process W_i is derived from a mother process W by means of some form of *random decimation*, e.g. per-flow sampling with factor $\psi < 1$, we write $W_i = \psi W$. In both cases — composition and random decimation — the processes W and W_i yield the same LD behaviour except for a vertical offset of $\log_2(k)$ or $\log_2(\psi)$ respectively, therefore $W_i \cong W$.

Now consider the particular case that W has the typical biscaling LD observed in the traces: linear with slope $l_w << 1$ (i.e., almost flat) at scales below a certain knee-point q_w, and linearly scaling with slope h_w above the knee-point. If another process S has the same biscaling behaviour with the same slopes $l_h = l_w$ and $h_s = h_w$ but a different location of the knee, i.e. $q_s \neq q_w$, then we say that W and S are *LD-similar* and write $W \sim S$. Note the difference between the notion of *LD-equivalence* (symbol "$\cong$") and *LD-similarity* (symbol "$\sim$"): LD-equivalence implies that the spectrum of the two

process is exactly the same at all scales (modulo a vertical shift), while with LD-similarity this holds true only at the extremes of the scale range, with an intermediate region of possibly different behaviour.

It is easy to show (see geometrical construction in Fig. 4) that if W is biscaling with almost flat spectrum at low timescales, then adding a Poisson component leads to a similar LD spectrum, formally:

$$W \text{ is biscaling with } l_w << 1 \Rightarrow W \sim W + P. \qquad (3)$$

Such simple notation is sufficient to demonstrate formally the effects of the semi-experiments on our toy-model, which is the goal of this work, and we do not elaborate it further[2].

B. Derivation of LD-similarity

By means of the notation introduced above we now proceed to demonstrate the observations (i)-(iv) stated above.

LD similarity for X and Y. By construction, it holds that:

$$Y = Y_0 + Y_1 = Y_0 + X_1 = P + X_1 \qquad (4)$$

where we have used $Y_1 = X_1$ (by the maximal fragmentation in class $\mathcal{S}_1$) and $Y_0 = P$. Recall that the child process X_1 is LD-equivalent to X, hence we derive:

$$Y = \frac{1}{2}X + P \sim \frac{1}{2}X \cong X. \qquad (5)$$

This proves the LD-similarity $X \sim Y$, i.e. observation (i). Moeover, since in eq. (5) the energy of P (only one point per session) is much lower than the energy of $\frac{1}{2}X$ (one point per packet), the latter dominates in the sum and therefore Y is (almost) equivalent to $\frac{1}{2}X$, hence to X.

Invariance to flow scrambling. The scrambling transformation relocates each flow according to a Poisson process. We will use the notation $\tilde{X}$ to denote the process resulting from the application of flow scrambling to X. Scrambling the mother process is equivalent to scrambling each of the two child components:

$$\tilde{X} = \tilde{X}_0 + \tilde{X}_1. \qquad (6)$$

By construction the flows in X_0 already Poisson distributed, therefore the scrambling transformation does not have any statistical impact, i.e. $\tilde{X}_0 = X_0$. Instead the flow scrambling turns the process X_1 into Possion: each individual packet is marked as a flow and therefore relocated, hence $\tilde{X}_1 = P$. With these positions eq. (6) leads to:

$$\tilde{X} = X_0 + P = \frac{1}{2}X + P \sim \frac{1}{2}X. \qquad (7)$$

Since $\frac{1}{2}X \cong X$ (decimation), eq. (7) leads immediately to the LD-similariy $\tilde{X} \sim X$, which proves the observation (ii).

[2]In principle such notation could be refined and extended, in combination with elements of geometric construction of LD, so as to develop a sort of "LD calculus".

Note that the energy of both components in the last term of eq. (7) is the same as there was no decimation.

Invariance to flow truncation. This transformation truncates the size of individual flows: only the first m packets of each flow are kept, while the remaining ones (if any) are removed. We will use the notation $\overline{X}^{(m)}$ to denote a process resulting from the m-truncation of X. The extraction of the Y process can be seen as an extreme case of flow truncation, where all but the first packets of each flow are removed, i.e. $Y = \overline{X}^{(1)}$. As done above, we separate the effects of flow truncation for the two child components:

$$\overline{X}^{(m)} = \overline{X_0}^{(m)} + \overline{X_1}^{(m)}. \qquad (8)$$

By construction X_1 contains only single-packet flows, hence the truncation is ineffective and we obtain:

$$\overline{X}^{(m)} = \overline{X_0}^{(m)} + X_1 = \overline{X_0}^{(m)} + \frac{1}{2}X. \qquad (9)$$

The effect of flow truncation onto X_0 is twofold. First, it subtracts energy to $\overline{X_0}^{(m)}$ (less points) thus letting the other component prevail in eq. (9). Second, it removes the scaling at very large timescales, as it truncates the tail of the session size distribution. This again results in X_1 dominating the LD at large scales. When $m = 1$, then $\overline{X_0}^{(1)}$ becomes Poisson, and eq. (9) leads to

$$\overline{X}^{(1)} = P + \frac{1}{2}X \sim \frac{1}{2}X \cong X \qquad (10)$$

thus we end up with the LD-similarity $\overline{X}^{(1)} \sim X$ which proves observation (iii).

Impact of joint scrambling AND truncation. We have seen that our toy-model embeds two distinct sub-process with equal LD behaviour. Each of the considered transformations (scrambling and truncation) impacts only one of the two components, transforming it into Poisson. Therefore both truncation (with $m = 1$) and scrambling, if indivudally applied, lead to a similar LD profile as the original process. It is easy to see that if both are applied jointly, then the process degenerates into Poisson, and LD becomes flat. This is consistent with observation (iv).

C. Simulation results

In Fig. 5 we show the LD plots obtained via simulations with our toy-model. We have considered two different session size distributions: exp-neg and pareto. In each case, we report the LD of the original process X along with the LD resulting after truncation and scrambling. It can be seen that all curves yield the same biscaling behaviour of the original process X.

V. MAPPING TO REALITY

In order to give credibility to the toy-model presented above, we need to identify the real-world phenomena that could be responsible for the model ingredients, and particularly the session-into-flows fragmentation. It turns out that several

(a) session size exp-neg ($\lambda_s = 1/5$).

(b) session size pareto ($k = 1, a = 1.25$).

Fig. 5: LD plots of the original process (X_p) and its derivations after flow scrambling ($X_{p,pois}$) and truncation ($X_{p,trun}$).

different kinds of activities in the modern Internet generate sessions split into different flows.

Scanning and unwanted traffic. A scanning source produces a "session" of packets towards different addresses and/or ports, therefore each packet is classified in a separate L4-flow. As a matter of fact, in the Internet there is a lot of scanning traffic - and unwanted traffic in general. Notably, the traces used in [1]–[3] date after the "explosion" of scanning traffic that is documented in [13]. In another work, we have found that such traffic has an impact on global rate (and delay) statistics, introducing large spikes, some of which appear regularly and periodically (see [14]). More in general, the impact of unwanted traffic on the wavelet spectrum of real traces is a point that deserves further study.

Peer-to-peer. Many P2P applications produce sessions rich in small flows. For instance, in file-sharing P2P applications the same content, typically a large file, is split into many small chunks — often of fixed size — and transferred via smaller connections from/to different hosts. In Skype, voice calls between two parties are routed over one or more middle-nodes. The middle-node can be dynamically changed during the call, which leads a single voice stream to be split into different L4 connections.

WEB. In general a complex web page consists of several elements that can be transferred simultaneously in different TCP connections, either due to the use of HTTP 1.0 or because the objects are located at different specialized servers (advertisements, images, etc.). Note that in [15] it has been reported that also the *number of flows* within WEB sessions is HT, a point that is fully consistent with our toy model.

All such applications are central to the modern Internet. There is ample evidence that P2P and WEB are responsible for a major share of the global Internet traffic. Therefore session splitting into small flows should be regarded nowadays as the norm, rather than the exception.

Recall that we have illustrated our toy-model with respect to only the most extreme form of fragmentation, i.e. flows of a single packet. This approach greatly simplifies the model because avoids the introduction of a probabilistic fragmentation scheme — note that maximal fragmentation is deterministic – and therefore help in focusing on the fundamental aspects of the model. In the real Internet single-packet flow fragmentation is produced only by scanning, while web and P2P obviously produce flows of non-trivial size. It is easy to see that the central properties of our toy-model — most prominently LD invariance to flow scrambling and truncation — would still hold if more articulated probabilistic fragmentation schemes are considered, for example with flows size distributed according to some probabilistic law, to better represent the behavior of real applications.

VI. IMPLICATIONS

The power of our toy-model resides in the following aspects. First, it can consistently explain the full set of empirical observations that were presented in [1]–[3] — to the best of our knowledge no other model exists that does so. Second, it is extremely simple and parsimonious, which accounts to its generality. Third, its generative components are reminiscent of real phenomena found in modern Internet traffic, as discussed in §V.

Despite its explicative power, we do not claim at this stage that our toy-model should be taken as representative of the real traffic process. More research and further analysis of the real traces is required to assess how much of the toy-model ingredients are found in real traffic — indeed the ultimate goal of the present contribution is to motivate further work in this direction. Nevertheless, our toy-model provides a possible explanatory framework that is original in its own, and might indicate a reasoning path towards unveiling some aspects of modern Internet traffic that would be overlooked if one accepts the idea that L4-flows should be the central entities in traffic models.

In order to present our arguments, we follow a reasoning approach based on "virtual experiments": we invite the reader to consider a "virtual Internet", wherein the traffic is generated according to our toy-model. In such "virtual Internet", the LD of the captured trace would be given by the curve X_p in Fig. 5b. Now consider a practitioner, Alice, willing to synthesize a CPP traffic model from the traces. She would then need to derive the empirical distribution of session size. Notably, the choice of the CPP model class is correct, as the toy-model is ultimately a version of CPP. The problem lies in the identification of the *clusters*. The easiest approach is to assume that clusters correspond to L4-flows, since these can be easily classified via simple 5-tupla matching. Inspired by the results of [1]–[3], Alice would probably follow this approach, and use the empirical distribution of L4-flows as a proxy for cluster distribution. The traffic model obtained in this way would be equivalent to the $X_{p,pois}$, whose LD is given in Fig. 5b. The close similarity between the LD of the captured trace and that obtained with the synthetic model — see LD of X_p and $X_{p,pois}$ in Fig. 5b — would comfort her that taking L4-flows as clusters is sufficient to capture the "important characteristics" of the traffic. Therefore she would conclude that L4-flows are the central entities of (virtual) Internet traffic, and that inter-flow correlations — or equivalently super-flow structures, i.e. sessions — do not need to be taken into account for modeling purposes. This is indeed the logical path followed in the earlier cited works [1]–[3].

In the context of our "virtual Internet" we know that such conclusion is *wrong*, because we have built the traffic generation process (i.e. the toy-model) around sessions, not L4-flows. The key question now is: What are the consequences of such modeling error ? The answer depends on the intended use of the model.

To illustrate, let us consider a very basic problem in data network engineering, namely buffer analysis. We consider a queuing system with a single server of capacity C (packets/sec) and infinite buffer. We feed the process X_p (our toy-model) and $X_{p,pois}$ (after L4-flow scrambling, or equivalently Alice's model) to this queuing system, and measure the waiting time of each packet. The server capacity C is set so as to obtain an average load $\rho = 0.5$ (recall that X_p and $X_{p,pois}$ have the same average packet rate). The distribution of the waiting time for both processes are plotted in Fig. 6. It is evident that the performance of the queuing system is dramatically different for the two processes. By eliminating the inter-flow correlation within (some of) the sessions, the aggregate packet-level process becomes less correlated, which leads to dramatically lower waiting times. In this case, a CPP model based on L4-flows as clusters instead of sessions would lead to a gross engineering error. More in general, the direct comparison between Fig. 6 and Fig. 5b provides an illustrative example of two processes — X_p and $X_{p,pois}$ in our case — that have similar LD but very different system behaviour. This should suggest some caution in using LD-similarity as a goodness criterion for traffic models.

Another potential consequence of mis-taking L4-flows for

Fig. 6: Empirical CCDF of the waiting time for the process X_p (curve $q1$) and $X_{p,pois}$ (curve $q2$) with $\rho = 0.5$. The LD for the same processes were given in Fig. 5b.

clusters relates to the empirical observability of LRD in practice. In [9] it is shown that the observability of the LRD region in the packetized M/G/∞ process — or more exactly in one realization of limited duration — depends on the duration of the observation period *relative to the mass of the cluster duration*. Now, the duration of a real trace is typically limited to 2-3 hours due to stationarity constraint. The problem here is to decide what entity must be taken as "cluster": L4-flows or sessions ? As a matter of fact, the HT property has been observed in both sessions size (see e.g. [15]) as well as L4-flow size. However this issue has an impact on the *observability* of the theoretical relationship between HT and LRD in real traces: as sessions are generally longer than L4-flow, if their mean duration exceeds a certain (unknown) value then it might be simply impossible to observe the Taqqu's Theorem at work in a real trace of a few hours. In other words, it remains to be seen whether the theoretical onset of LRD is within the stationarity limit of real traffic traces once that sessions are taken as clusters. See [9] for a discussion on this point.

VII. CONCLUSIONS

We have proposed a simple toy-model that is capable to consistently explain a number of empirical observations reported in some previous works, based on the analysis of real traces. Indeed, given its simplicity, it is remarkable that it has the power to evocate consistently the whole set of LD patterns resulting from the semi-experiments. We have also presented a possible mapping of the model ingredients to real phenomena in the Internet. Despite its evocative power with respect to LD properties we do not claim at this stage that such ultra-simplified model is representative of real Internet traffic, and more work is needed to achieve a more realistic — and likely less parsimonious — model to be used in practice.

The main contribution of our toy-model is that it provides the basis of a novel interpretation framework where higher-layer sessions, rather than transport-layer flows, should be

considered as the central entities in traffic modeling — at least for the CPP model class. We hope in this way to contribute the discussion within the networking community about the role of sessions and flows, and at the same time motivate further research on characterizing the sessions, as done e.g. in [15]. This is clearly more difficult because there is not simple mechanism to classify sessions, while L4-flows can be identified directly by the IP 5-tupla. On the other hand, keeping the focus on L4-flows appears increasingly less natural in the modern Internet scenario, wherein applications and services are evolving in such a way that the coupling between application-level activities and L4-flows becomes progressively weaker.

Along the way we have touched a number of issues that still deserve further investigations, like the role of LD evocation in model assessment and the observability of HT-LRD relation in real traces.

In the progress of this work we intend to follow two complementary directions. On one hand, we are interested in studying more articulated variants of the toy-model with more realistic (probabilistic) fragmentation/composition schemes between L4-flows and sessions. On the other hand, we intend to investigate possible methods to classify sessions (advancing e.g. the work in) and leverage assess quantitatively the mapping between sessions and L4-flows in real traces. For that, we would need to advance the

REFERENCES

[1] N. Hohn, D. Veitch, P. Abry. Does Fractal Scaling at the IP Level depend on TCP Flow Arrival Processes? *Internet Measurements Workshop, Marseille, France*, November 2002.

[2] N. Hohn, D. Veitch, P. Abry. Cluster processes, a natural language for network traffic. *IEEE Transactions on Signal Processing, special issue on Signal Processing in Networking*, 51(8):2229–2244, 2003.

[3] N. Hohn, D. Veitch, P. Abry. The impact of the flow arrival process in Internet traffic. *Proc. of IEEE Int. Conf. on Acoust. Speech and Sig. Proc.*, 2003.

[4] B. Krishnamurty, M. Crovella. *Internet Measurement: Infrastructure, Traffic and Applications*. Wiley, 2006.

[5] K. Park, W. Willinger. *Self-similar Network Traffic and Performance Evaluation*. Wiley, 2000.

[6] P. Abry, D. Veitch. Wavelet analysis of long range dependent traffic. *IEEE Trans. on Information Theory*, 1(44):2–15, 1998.

[7] M. Grossglauser, J.C. Bolot. On the relevance of Long-Range Dependence in Network Traffic. *IEEE/ACM Trans. on Networking*, 7(5), 1999.

[8] W. Willinger, R. Govindan, S. Jamin, V. Paxson, S. Shenker. Scaling phenomena in the Internet: Critically examining criticality. *Proceedings of the National Academy of Sciences, vol. 99, issue 90001, pp. 2573-2580*, February 2002.

[9] P. Abry, P. Borgnat, F. Ricciato, A. Scherrer, D. Veitch. Revisiting and old friend: On the observability of the relation between long range dependence and heavy tail. *preprint*, 2008.

[10] D. Veitch, P. Abry. Matlab code for the wavelet based analysis of scaling processes. *Available from ¡http:// www.cubinlab.ee.mu.oz.au/∼darryl/secondorder_code.html¿*.

[11] F. Ricciato. Some remarks to recent papers on traffic analysis. *editorial for CCR*, July 2006.

[12] M. Roughan, D. Veitch. Some remarks on unexpected scaling exponents. *editorial for CCR*, October 2007.

[13] M. Allman, V. Pasxson, J. Terrell. A Brief History of Scanning. *ACM Internet Measurement Conference (IMC'07), San Diego*, October 2007.

[14] F. Ricciato, E. Hasenleithner, P. Romirer-Maierhofer. Traffic analysis at short time-scales: an empirical case study from a 3g cellular network. *IEEE Transactions on Network and Service Management*, March 2008.

[15] A. Bianco, G. Mardente; M. Mellia; M. Munafo; L. Muscariello. Web user session characterization via clustering techniques. *Globecom*, 2005.