

A LONG-RANGE DEPENDENT MODEL FOR NETWORK TRAFFIC WITH FLOW-SCALE CORRELATIONS

PATRICK LOISEAU, PASCALE VICAT-BLANC PRIMET, AND PAULO GONÇALVES

ABSTRACT. For more than a decade, it has been observed that network traffic exhibits long-range dependence and many models have been proposed relating this property to heavy-tailed flow durations. However, none of these models consider correlations at flow scale. Such correlations exist and will become more prominent in the future Internet with the emergence of flow-aware control mechanisms correlating a flow's transmission to its characteristics (size, duration, etc.).

In this work, we study the impact of the correlation between flow rates and durations on the long-range dependence of aggregate traffic. Our results extend those of existing models by showing that two possible regimes of long-range dependence exist at different time scales. The long-range dependence in each regime can be stronger or weaker than standard predictions, depending on the conditional statistics between the flow rates and durations. In the independent case, our proposed model consistently reduces to former approaches.

The pertinence of our model is validated on real web traffic traces, and its ability to accurately explain the Hurst parameter is validated on both web traces and numerical simulations.

Keywords: Network traffic, Long-range dependence, Heavy-tailed distributions, Poisson process

1. MOTIVATION

In the last years, a considerable research effort has been devoted to the mathematical modeling of network traffic, with particular emphasis on statistical approaches. However, comprehensive modeling of all the characteristics of the traffic is a very arduous problem for it encompasses several difficulties of different natures, such as transport protocols, control mechanisms, or complexity due to the users' behavior. The design of identifiable models, while sufficiently versatile to account for essential characteristics of real traffic, is therefore an important issue for many applications such as the simulation of realistic traces for numerical studies, or the detection of abnormal traffic behaviors.

An important step ahead in this direction was the discovery in 1993 of the self-similar nature of aggregate traffic time series at large time scales [26, 40]. Following up, several models appeared, proposing the heavy-tailed nature of the activity periods (the flow durations) as a plausible origin of self-similarity, and formalizing the relation between α_{ON} , the heavy-tail exponent of the flow-duration distribution and H , the Hurst self-similarity parameter:

$$(1.1a) \quad H = \frac{3 - \alpha_H}{2},$$

$$(1.1b) \quad \text{where } \alpha_H = \min(\alpha_{ON}, 2).$$

These seminal works opened up a vast debate regarding the importance of the long-range dependence property in network-traffic time series. Yet, it is now commonly reckoned that due to its large-scale nature, long-range dependence only impacts the queueing performance of large buffers (see *e.g.* [33]) which are not so many in modern networks. Nonetheless, the generation of realistic traffic traces remains highly demanding in terms of reliable sources at the flow-level [4]. As a result, the clear comprehension of the long-term correlations produced by these models is not only interesting on its own, but it can also serve other purposes, like characterizing the performance of load estimators, or designing new procedures for load change detection [34]. Moreover, as we will see, the model we propose takes into account certain traffic features that the emergence of flow-aware control mechanisms in the future Internet may significantly accentuate. We present our contribution as a first step towards a better understanding of the impact of those mechanisms on the traffic properties.

So far though, all existing models that lead to relation (1.1), rely on the simplifying assumption that all flows have the same (mean) throughput no matter what their size (or duration). In particular, this implies that the tail indices α_{SI} and α_{ON} of the flow-size and of the flow-duration distributions respectively, are equal and thus, interchangeable in relation (1.1). For real world traffic, however, this independence assumption rarely holds, and the tail indices α_{SI} and α_{ON} are different in general (as it has been observed on web traffic since 1997 [12, 35]).

Several sensible causes can explain a possible statistical dependence between the flows' rates and durations (or sizes), among which the transient phase of the TCP protocol or the existence of mirror sites to speed up large files transfer over the net, to cite only two. More importantly, this *flow-scale* correlation is likely to become an inherent feature of the future Internet as flow-aware control procedures will advisedly condition the flows' treatment (*e.g.* the transmission rate) to their characteristics (*e.g.* their duration or their size).

In these conditions, relation (1.1) is not ensured to be reliably usable to predict the Hurst parameter of the aggregate traffic. Hence, a model that explicitly includes the correlation between flow rates and durations is lacking to understand the impact of these flow-scale dependencies on aggregate traffic properties, and to accurately predict the resulting Hurst parameter.

In this paper, we address this challenge and we propose an extension of prior models including the correlation between flows' rates and durations, that makes explicit the individual effects of α_{ON} and α_{SI} on self-similarity. Flows are represented as a marked planar Poisson process. This non-classical viewpoint allows us to simply calculate the autocovariance function of the aggregate traffic's instantaneous bandwidth and to deduce the resulting Hurst parameter.

The rest of the article is organized as follows. We first briefly review in Section 2 existing models that reproduce the long-range dependence property of network traffic. In Section 3, we propose an extension of these models which includes the correlation between flows' rates and durations, and we expose the theoretical derivations leading to the aggregate traffic's Hurst-parameter prediction. In Section 4, based on numerical simulations and on recent web traffic traces, we experimentally demonstrate the pertinence of our model and its ability to correctly predict the Hurst parameter. Finally, we conclude in Section 5.

2. RELATED WORK

There exists a large number of models able to reproduce the long-range dependence property [7, 17]. We describe here only those which explicitly ground the origin of long-range dependence in the notion of *flows*. Following up Mandelbrot's idea, these models rely on the introduction of a heavy-tailed distribution of infinite variance. We distinguish two categories of such models, which mainly differ in the flow arrival process: the *renewal models* and the *infinite source Poisson models*.

2.1. Renewal models. Renewal models are based on the same general setting firstly introduced by Mandelbrot in 1969 [32] in an economical context. Each source, emitting only one flow at a time, is modeled as a renewal reward process, where the inter-arrival time (*i.e.* the interval between two consecutive flows) is heavy-tailed; and the reward is the rate of the flow. Many variants of such models have been considered, mainly differing in the distribution of the reward [43, 27, 28] (see also [17]). We present here only briefly the particular variant where the reward strictly alternates between 0 and 1: the classical ON/OFF model [44]. This model allows including the notion of idle time between the transmission of two flows through the OFF periods. Moreover, ON- and OFF-periods distributions can have different tail indices α_{ON} and α_{OFF} . Both distributions are assumed to have a finite mean (*i.e.*, $\alpha_{ON}, \alpha_{OFF} > 1$), and at least one of these distributions is assumed to have an infinite variance (*i.e.*, $\alpha_{ON} < 2$ or $\alpha_{OFF} < 2$). We denote by $W_t^{(n)}$ the instantaneous rate of source n at time t ($W_t^{(n)} = 0$ or 1), and define the cumulative input process from M aggregated sources:

$$W_{Tt} = \int_0^{Tt} \sum_{n=1}^M W_u^{(n)} du.$$

Then, the two following limit theorems hold [42]:

$$(2.1) \quad \lim_{T \rightarrow \infty} \lim_{M \rightarrow \infty} \frac{W_{Tt} - \mathbb{E}\{W_{Tt}\}}{T^H M^{1/2}} = \sigma B_H(t)$$

and

$$(2.2) \quad \lim_{M \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{W_{Tt} - \mathbb{E}\{W_{Tt}\}}{T^{1/\alpha_H} M^{1/\alpha_H}} = c\Lambda_{\alpha_H}(t),$$

where H is defined as in relation (1.1a) with $\alpha_H = \min(\alpha_{\text{ON}}, \alpha_{\text{OFF}}) < 2$, and σ and c are constants. If both the ON- and OFF-periods distributions have finite variance, the limit process is an ordinary Brownian motion (exhibiting no long-range dependence), no matter what the order of the two limits. Here, B_H denotes a fractional Brownian motion of Hurst parameter H , that is a self-similar Gaussian process, whose increments are stationary. Its increment process (called fractional Gaussian noise) is said to be long-range dependent of Hurst parameter H , meaning that its autocovariance function decays as a power law of index $(2H - 2)$. The result (2.1), which establishes the long-range dependence of the aggregate traffic's bandwidth in the Gaussian limit has been widely used in the last decade to model network traffic, since taking first the limit on the number of sources (*i.e.* aggregating the sources) shows more natural. Conversely however, when time scale goes first to infinity, Equation (2.2) shows that the limit process Λ_{α_H} is a Lévy stable motion: a self-similar process with stable marginals (*i.e.* non-gaussian and heavy-tailed) and independent stationary increments. In the case where M and T tend simultaneously to infinity, [37] gives conditions on the ratio between the growth rate of both quantities that ensure one, or the other limit regimes (see also [18] where an intermediate case with a ratio between these two conditions is studied). In practice, when the number of sources and the time scale are finite, [23] provides sufficient conditions on the aggregation levels in both the horizontal (time) and the vertical (flows) directions, for the Gaussian approximation to hold. One particularity of the ON/OFF model is that the flow arrival process is not Poisson in general [31]. However, if Equation (2.1) or (2.2) is rescaled in such a way that the idle OFF periods grow with the number of sources, then the flow arrival process tends towards a Poisson process [14]. In that case, the OFF durations (hence the index α_{OFF}) have no effect, and the model asymptotically becomes an infinite source Poisson model. Before entering into details, let us mention that, whenever $\alpha_{\text{ON}} < \alpha_{\text{OFF}}$, the predicted Hurst parameter does not depend on α_{OFF} . This particular situation was observed in early Internet traffic measurement [12] and more recent measurement campaigns (see *e.g.* [6] in the context of social networks) also demonstrate that the distributions of inter-session times have a lighter tail than that of session times.

2.2. Infinite source Poisson models. In the simplest version of the infinite source Poisson models, flows arrive at the link as a Poisson process of rate λ . For each flow, data is transmitted at a fixed rate (arbitrarily set to 1) during a heavy-tailed distributed random time with tail index α_{ON} . This model is also known as the M/G/ ∞ model that was originally considered in [11]. Amazingly then, the same limit equations (2.1) and (2.2) as for the ON/OFF model are obtained, when replacing the number of sources M by the flow-arrival rate λ , and posing $\alpha_H = \alpha_{\text{ON}}$ (see [17, 22]). Furthermore, it is shown in [37] that the same conditions on the arrival rate and on the time scale growths determine the two limit regimes. There exists many variants of this elementary M/G/ ∞ model, which all rely on the same mechanism: a Poisson arrival process and a heavy-tailed distributed duration of flows. Differences mainly lie in the way data is transmitted within a flow (see also a survey of infinite Poisson models in [20], and the references in [37]). In [25], the authors consider an infinite source Poisson model with a general form of “workload function” within a flow and establish the Gaussian limit result. In [36], a similar model is considered (where the “workload function” is called “transmission schedule”), and the authors establish the alternative convergence towards a Lévy motion. The Poisson shot-noise model developed in [3], is morally very similar to the two previously introduced models. The “workload functions” or “transmission schedules” are here termed “shots”, which still arrive according to a Poisson process. An application of this model is proposed in [38] where the shape of the shot is representative of the AIMD (Additive Increase Multiplicative Decrease) mechanism with Poisson losses. In [15], the authors consider a model where the flow rate is a random variable that remains constant over the full duration of the flow. The flow sizes are drawn at random according to a heavy-tailed distribution of tail index α_{SI} and the rate variable, drawn independently of the size, follows a finite-mean distribution. The flow duration is then also a heavy-tailed random variable with the same index $\alpha_{\text{ON}} = \alpha_{\text{SI}}$, and the model leads to a long-range dependent traffic with Hurst parameter H defined as in relation (1.1). A similar model where the durations and the rates are drawn independently is proposed in [16]. The last model that we

mention in the infinite source Poisson model category is the Cluster Point Process (CPP) model [21], based on a discrete point process approach. Flows are “clusters” of points (packets) and a flow-arrival time is determined by the instant of its first packet. The flow-arrival process is Poisson, the number of points in the cluster (the flow size) is heavy-tailed distributed with index α_{SI} , and points (packets) within a cluster (flows) follow a renewal process with some inter-arrival distribution determining the mean flow rate. From point-processes theory [10, 13], the resulting aggregate traffic’s spectrum is shown to correspond to long-range dependence with the same Hurst exponent H as in relation (1.1a) with $\alpha_H = \alpha_{SI}$. As in the model of [15], the flow duration here is heavy-tailed distributed with tail index $\alpha_{ON} = \alpha_{SI}$, an equality that was disallowed in [12] by measurements on real Internet traffic. This can be caused for example by the slow-start mechanism which, allowing larger flows to reach a higher mean rate, naturally yields a lighter tail for the flow-duration distribution. The authors of [21] suggest the use of a multi-class CPP, where small flows (mice class) would be allotted small rates and larger flows (elephant class) higher rates. However, they do not develop further the implications on the aggregate traffic, and the central question that we address in the present work remains open: When flow-duration and flow-size distributions have different tail indices, which specific role do those play in the origin and intensity of long-range dependence of the aggregate traffic?

2.3. Planar point process setting. Our proposed model is an infinite source Poisson model, where flows are represented as a marked planar Poisson process. This setting was originally introduced in the context of multifractal analysis in [5] to extend binomial multiplicative cascades, and reused in [9], also to deal with multiplicative processes. To the best of our knowledge, in the context of additive processes, this approach was only used in the aforementioned articles [15, 16] to derive long-range dependence. Besides being very intuitive, this formalism allows readily computing the autocovariance function in the case of non-independent marks (rates).

3. A MODEL INCLUDING CORRELATION BETWEEN FLOW RATES AND DURATIONS

3.1. Definitions and notations. We consider the point process $\{(T_i, D_i), i \geq 0\}$ depicted on Figure 1, where T_i represents the arrival time of flow i and D_i its duration, and we assume that it is a (planar) Poisson process of intensity Λ [11, 24]. The intensity Λ is a measure that fixes the mean number of points in any region of the plane; for instance $\Lambda(C(t_1))$ is the mean value of the random variable formed by the number of points lying in the cone $C(t_1)$ (*i.e.* the number of flows active at time t_1 , see Figure 1). If the density of Λ is constant over the plane, the point process is called homogeneous. Here, to take into account the heavy-tailed distribution of the flow size, the density of an infinitesimal square of size $dt \times dd$ centered on (t, d) , takes on the particular form:

$$(3.1) \quad \Lambda(dt, dd) = \begin{cases} \frac{C dt dd}{d^{\alpha_{ON}+1}} & \text{if } d \geq d_{\min} \\ 0 & \text{otherwise} \end{cases},$$

where $d_{\min} > 0$ is the minimal flow duration, C is a positive constant, and $\alpha_{ON} > 1$ is the tail index of the flow-duration distribution. As Λ depends on d , the point process is clearly non-homogenous (the smaller d is, the higher the density). However, since Λ does not vary with t , the arrival times T_i are independent of the durations D_i and the resulting traffic is stationary. With this particular form of the intensity measure Λ , the planar Poisson process model is equivalent to an infinite source Poisson model, with a finite flow arrival rate ($\lambda = \int_{d=0}^{\infty} \int_{t=0}^1 \Lambda(dt, dd) < \infty$) and independent durations drawn from a Pareto distribution of tail index α_{ON} . Each flow i emits data at a constant rate $R_i \geq 0$, drawn at random, over its full duration D_i . The rates $(R_i)_{i \geq 0}$ do form a sequence of independent random variables, but in contrast to previous models (such as the ones proposed in [15, 16]), they are not assumed independent of the durations $(D_i)_{i \geq 0}$. We will formalize the explicit statistical bond between these two quantities later in this section, when it becomes necessary to pursue the calculations. We designate by S_i the size of flow i : $S_i = D_i R_i$.

Let us consider two time instants $t_1 < t_2$ and let us introduce the following notations (see Figure 1):

$$(3.2) \quad W_{t_1 \setminus t_2} = \sum_{(T_i, D_i) \in C(t_1) \setminus C(t_2)} R_i,$$

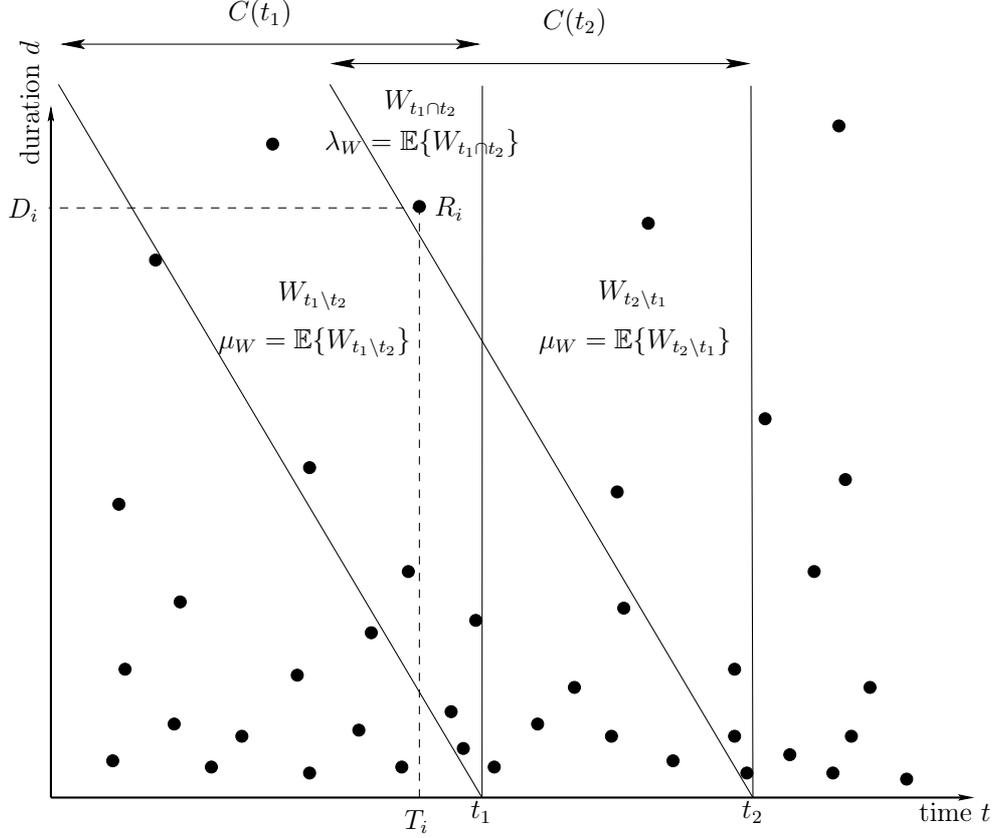


FIGURE 1. Setting of the model. Each point represents a flow. The x -coordinate represents the start time T_i and the y -coordinate the duration D_i . R_i is the rate of the flow. At time t_1 , the active flows are those whose corresponding point on the graph lies in cone $C(t_1)$ (the left border of the cones have a slope of -1).

$$(3.3) \quad W_{t_1 \cap t_2} = \sum_{(T_i, D_i) \in C(t_1) \cap C(t_2)} R_i,$$

$$(3.4) \quad W_{t_1} = W_{t_1 \setminus t_2} + W_{t_1 \cap t_2} = \sum_{(T_i, D_i) \in C(t_1)} R_i,$$

and similarly for the variables $W_{t_2 \setminus t_1}$ and W_{t_2} . The random variable W_{t_1} is the instantaneous throughput at time t_1 (rates summation of all active flows at time t_1). $W_{t_1 \setminus t_2}$ represents the traffic generated by flows active at time t_1 but no longer active at time t_2 , while $W_{t_1 \cap t_2}$ denotes traffic coming from flows active at time t_1 and still alive at time t_2 . We will use these two intermediate variables to facilitate our derivations, and we define the corresponding means:

$$(3.5) \quad \lambda_W = \mathbb{E}\{W_{t_1 \cap t_2}\} = \Lambda(C(t_1) \cap C(t_2)),$$

$$(3.6) \quad \mu_W = \mathbb{E}\{W_{t_1 \setminus t_2}\} = \mathbb{E}\{W_{t_2 \setminus t_1}\} = \Lambda(C(t_1) \setminus C(t_2)) = \Lambda(C(t_2) \setminus C(t_1)).$$

Last equality $\Lambda(C(t_1) \setminus C(t_2)) = \Lambda(C(t_2) \setminus C(t_1))$ simply comes from time-shift invariance of measure Λ . Finally, with these notations, we have:

$$(3.7) \quad \mathbb{E}\{W_{t_1}\} = \mathbb{E}\{W_{t_2}\} = \lambda_W + \mu_W.$$

With this setting, we now compute the autocovariance of the instantaneous bandwidth W_t , in order to evaluate the Hurst parameter. We recall that a stationary process is said long-range dependent of

Hurst parameter H if its autocovariance function decreases like τ^{2H-2} when τ goes to infinity and where $1/2 < H < 1$.

3.2. Computation of the instantaneous bandwidth's autocovariance. With the notations introduced in the preceding section, we now calculate the autocovariance function $\mathbb{E}\{W_{t_1}W_{t_2}\} - \mathbb{E}\{W_{t_1}\}\mathbb{E}\{W_{t_2}\}$ of the instantaneous throughput W_t .

Lemma 3.1. *If the planar point process $\{(T_i, D_i), i \geq 0\}$ is Poisson, then*

$$(3.8) \quad \mathbb{E}\{W_{t_1}W_{t_2}\} - \mathbb{E}\{W_{t_1}\}\mathbb{E}\{W_{t_2}\} = \mathbb{E}\{W_{t_1 \cap t_2}^2\} - \mathbb{E}\{W_{t_1 \cap t_2}\}^2 = \text{Var}\{W_{t_1 \cap t_2}\}.$$

Proof. With our notations:

$$\mathbb{E}\{W_{t_1}W_{t_2}\} = \mathbb{E}\{W_{t_1 \setminus t_2}W_{t_2 \setminus t_1}\} + \mathbb{E}\{W_{t_1 \setminus t_2}W_{t_1 \cap t_2}\} + \mathbb{E}\{W_{t_1 \cap t_2}W_{t_2 \setminus t_1}\} + \mathbb{E}\{W_{t_1 \cap t_2}^2\}.$$

Due to the Poisson assumption, the random variables $W_{t_1 \setminus t_2}$, $W_{t_2 \setminus t_1}$, $W_{t_1 \cap t_2}$ are mutually independent, so that

$$\begin{aligned} \mathbb{E}\{W_{t_1}W_{t_2}\} &= \mathbb{E}\{W_{t_1 \setminus t_2}\}\mathbb{E}\{W_{t_2 \setminus t_1}\} + \mathbb{E}\{W_{t_1 \setminus t_2}\}\mathbb{E}\{W_{t_1 \cap t_2}\} + \mathbb{E}\{W_{t_1 \cap t_2}\}\mathbb{E}\{W_{t_2 \setminus t_1}\} + \mathbb{E}\{W_{t_1 \cap t_2}^2\} \\ &= \mu_W^2 + 2\lambda_W\mu_W + \mathbb{E}\{W_{t_1 \cap t_2}^2\} \\ &= (\lambda_W + \mu_W)^2 - \lambda_W^2 + \mathbb{E}\{W_{t_1 \cap t_2}^2\}, \end{aligned}$$

which readily yields the result in view of the definition of λ_W (Equation (3.5)) and of Equation (3.7). \square

Lemma 3.1 shows that this autocovariance only depends on the variance of traffic generated by flows that are active at times t_1 and t_2 . It is quite natural and was noticed in [5, 9] in a different context. To complete the calculation of the autocovariance function, we then need to compute $\text{Var}\{W_{t_1 \cap t_2}\}$.

If the flows' rates were all equal to 1, $W_{t_1 \cap t_2}$ would simply be the number of points in $C(t_1) \cap C(t_2)$, and the variance $\text{Var}\{W_{t_1 \cap t_2}\}$ equal to λ_W (the mean and the variance of a Poisson process are equal). The autocovariance would then be entirely controlled by the value of λ_W . Before carrying on the calculation, let us denote by N the random variable corresponding to the number of points in $C(t_1) \cap C(t_2)$ and by

$$(3.9) \quad \lambda_N = \mathbb{E}\{N\},$$

its mean. Then, the autocovariance function takes on the general form:

Lemma 3.2.

$$(3.10) \quad \mathbb{E}\{W_{t_1}W_{t_2}\} - \mathbb{E}\{W_{t_1}\}\mathbb{E}\{W_{t_2}\} = \lambda_N \mathbb{E}\{R_i^2\},$$

where it is implicitly understood that the expectation of R_i^2 is computed in $C(t_1) \cap C(t_2)$.

Proof. To evaluate the value of $\text{Var}\{W_{t_1 \cap t_2}\}$, we successively compute the values of $\mathbb{E}\{W_{t_1 \cap t_2}\}$ and $\mathbb{E}\{W_{t_1 \cap t_2}^2\}$.

$$\begin{aligned} \mathbb{E}\{W_{t_1 \cap t_2}\} &= \mathbb{E}\{\mathbb{E}\{W_{t_1 \cap t_2} | N\}\} \\ &= \sum_{k=1}^{\infty} \mathbb{E}\left\{\sum_{i=1}^k R_i | N = k\right\} \mathbb{P}(N = k) \\ &= \sum_{k=1}^{\infty} \sum_{i=1}^k \mathbb{E}\{R_i | N = k\} \mathbb{P}(N = k) \\ &= \sum_{k=1}^{\infty} k \mathbb{E}\{R_i\} \mathbb{P}(N = k) \\ &= \lambda_N \mathbb{E}\{R_i\}. \end{aligned}$$

$$\begin{aligned} \mathbb{E}\{W_{t_1 \cap t_2}^2\} &= \mathbb{E}\{\mathbb{E}\{W_{t_1 \cap t_2}^2 | N\}\} \\ &= \sum_{k=1}^{\infty} \mathbb{E}\left\{\left(\sum_{i=1}^k R_i\right)^2 | N = k\right\} \mathbb{P}(N = k). \end{aligned}$$

By independence of the sequence $(R_i)_{i \geq 0}$,

$$\mathbb{E}\left\{\left(\sum_{i=1}^k R_i\right)^2 \mid N = k\right\} = \mathbb{E}\left\{\left(\sum_{i=1}^k R_i\right)^2\right\} = k\mathbb{E}\{R_i^2\} + (k^2 - k)\mathbb{E}\{R_i\}^2,$$

so that

$$\begin{aligned} \mathbb{E}\{W_{t_1 \cap t_2}^2\} &= \sum_{k=1}^{\infty} (k\mathbb{E}\{R_i^2\} + (k^2 - k)\mathbb{E}\{R_i\}^2) \mathbb{P}(N = k) \\ &= \sum_{k=1}^{\infty} k\mathbb{E}\{R_i^2\} \mathbb{P}(N = k) + \sum_{k=1}^{\infty} k^2 \mathbb{E}\{R_i\}^2 \mathbb{P}(N = k) - \sum_{k=1}^{\infty} k\mathbb{E}\{R_i\}^2 \mathbb{P}(N = k) \\ &= \lambda_N \mathbb{E}\{R_i^2\} + (\lambda_N + \lambda_N^2) \mathbb{E}\{R_i\}^2 - \lambda_N \mathbb{E}\{R_i\}^2 \\ &= \lambda_N \mathbb{E}\{R_i^2\} + \lambda_N^2 \mathbb{E}\{R_i\}^2. \end{aligned}$$

□

Until now, we have not used the specific form of the measure Λ (Equation (3.1)), nor we have considered an explicit form for the correlation between R_i and D_i . The result of Lemma 3.2 depends only on the Poisson hypothesis and on the independence assumption of the sequence $(R_i)_{i \geq 0}$. It shows that the autocovariance function is the product of two terms: λ_N and $\mathbb{E}\{R_i^2\}$. The first term λ_N (the mean number of points in $C(t_1) \cap C(t_2)$) depends only on the measure Λ (*i.e.* on the points repartition). To compute the second term $\mathbb{E}\{R_i^2\}$, we need a functional model to describe the correlation between R_i and D_i .

Our goal here is to introduce a correlation which leads to different tail indices α_{SI} and α_{ON} for the flow-size and for the flow-duration distributions. As we already mentioned, independence between the two random variables R_i and D_i leads to identical tail indices, no matter what finite mean distribution is chosen for R_i (see *e.g.* [15, 16]). Therefore, different tail indices $\alpha_{\text{SI}} \neq \alpha_{\text{ON}}$ can only come from a statistical correlation between R_i and D_i . A naive choice could be to deterministically set each flow rate to $R_i = K D_i^{\beta-1}$, where $\beta = \frac{\alpha_{\text{ON}}}{\alpha_{\text{SI}}}$. In that case, we have $S_i = K D_i^\beta$, which effectively leads to a heavy-tailed flow-size distribution with tail index α_{SI} different from α_{ON} . However, assuming that a flow's rate is deterministically imposed by its duration is not realistic. That is why, in our model, R_i is a random variable but whose mean and variance are statistically conditioned to the flow duration D_i . The conditional expectation is set to $\mathbb{E}\{R_i | D_i\} = K D_i^{\beta-1}$ ($\beta > 0$), as this is the sole algebraic relation compatible with the desired heavy-tailed distributions with different tail indices. As for the conditional variance, we also choose to model its dependence on D_i with a power law of the type: $\text{Var}\{R_i | D_i\} = V D_i^\gamma$. Experimentally, this choice fits a larger number of real traffic data than a constant conditional variance; from an analytic viewpoint, it permits a simple analysis of long-range dependence in the aggregate traffic. Next proposition shows that, provided a simple condition on β and γ , the prescription of these conditional moments indeed leads to a heavy-tailed flow-size distribution, with tail parameter $\frac{\alpha_{\text{ON}}}{\beta}$.

Proposition 3.1 (Flow-size distribution). *Suppose that D_i follows a Pareto distribution with tail index α_{ON} . Suppose that $\mathbb{E}\{R_i | D_i\} = K D_i^{\beta-1}$ and $\text{Var}\{R_i | D_i\} = V D_i^\gamma$, with K and V two positive constants. If $\gamma < 2(\beta - 1)$, then the distribution of the flow size S_i is heavy-tailed distributed with index $\frac{\alpha_{\text{ON}}}{\beta}$:*

$$(3.11) \quad \mathbb{P}(S_i > s) \underset{s \rightarrow \infty}{\sim} \left(\frac{s}{K}\right)^{-\frac{\alpha_{\text{ON}}}{\beta}}.$$

Proof. The proof is deferred to appendix A. □

Remark 3.1. *When $\gamma \geq 2(\beta - 1)$, there is no guarantee that the flow-size distribution is heavy-tailed with index $\frac{\alpha_{\text{ON}}}{\beta}$.*

Remark 3.2. *Because we clearly have $\mathbb{E}\{R_i^2\} = \mathbb{E}\{\mathbb{E}\{R_i^2 | D_i\}\} = K^2 \mathbb{E}\{D_i^{2(\beta-1)}\} + V \mathbb{E}\{D_i^\gamma\}$, the rate variable has finite variance if $\alpha_{\text{ON}} > \gamma$ and $\alpha_{\text{ON}} > 2(\beta - 1)$. Moreover, it ensures that the autocovariance function of the instantaneous bandwidth $(W_t)_{t \in \mathbb{R}}$ exists, and also justifies the Gaussian approximation*

for the traffic. Similarly, since $\mathbb{E}\{S_i\} = \mathbb{E}\{D_i\mathbb{E}\{R_i|D_i\}\} = K\mathbb{E}\{D_i^\beta\}$, the condition $\alpha_{\text{ON}} > \beta$ warrants a finite-mean distribution for the flow-size variable. This is consistent with a tail index larger than one in Proposition 3.1 when $\gamma < 2(\beta - 1)$.

We now have all the ingredients of our model to establish the algebraic decay of the resulting autocovariance function.

Proposition 3.2 (Autocovariance function of the instantaneous bandwidth $(W_t)_{t \in \mathbb{R}}$). *Let flows be modeled as a planar Poisson process with intensity Λ defined as in expression (3.1). Assume that $\mathbb{E}\{R_i|D_i\} = KD_i^{\beta-1}$ and $\text{Var}\{R_i|D_i\} = VD_i^\gamma$, where K and V are two positive constants, and let:*

$$(3.12) \quad \begin{cases} \alpha' = \alpha_{\text{ON}} - 2(\beta - 1), \\ \alpha'' = \alpha_{\text{ON}} - \gamma. \end{cases}$$

If $\alpha' > 1$ and $\alpha'' > 1$, then

$$(3.13) \quad \begin{aligned} \mathbb{E}\{W_{t_1}W_{t_2}\} - \mathbb{E}\{W_{t_1}\}\mathbb{E}\{W_{t_2}\} &= CK^2 \frac{1}{\alpha'(\alpha'-1)}(t_2 - t_1)^{-\alpha'+1} \\ &+ CV \frac{1}{\alpha''(\alpha''-1)}(t_2 - t_1)^{-\alpha''+1}. \end{aligned}$$

Proof. First note that the conditions $\alpha' > 1$ and $\alpha'' > 1$ guarantee the existence of the autocovariance function (see Remark 3.2), and recall that $\mathbb{E}\{R_i^2\} = K^2\mathbb{E}\{D_i^{2(\beta-1)}\} + V\mathbb{E}\{D_i^\gamma\}$. Then calculating the expectations via direct integration with the explicit form of the density of measure Λ (Equation (3.1)) yields:

$$\mathbb{E}\{D_i^{2(\beta-1)}\} = \frac{1}{\lambda_N} C \frac{1}{\alpha'(\alpha'-1)}(t_2 - t_1)^{-\alpha'+1},$$

and

$$\mathbb{E}\{D_i^\gamma\} = \frac{1}{\lambda_N} C \frac{1}{\alpha''(\alpha''-1)}(t_2 - t_1)^{-\alpha''+1}.$$

The result directly follows from Equation (3.10) of Lemma 3.2. \square

Remark 3.3. *If $\gamma \geq 2(\beta - 1)$ the result remains valid although the flow sizes may not be heavy-tailed distributed with index $\frac{\alpha_{\text{ON}}}{\beta}$.*

Remark 3.4. *As its autocovariance only depends on the time difference $(t_2 - t_1)$, the instantaneous bandwidth W_t is a second-order stationary process, consistently with the time-shift invariance of the measure Λ .*

Our main result lies in the algebraic decay of the autocovariance function of Proposition 3.2, where two different regimes of long-range dependence coexist. Before presenting real traffic traces supporting our model choice, let us comment and elaborate on the specific form of this autocovariance and its interpretation.

3.3. Interpretation. The autocovariance (3.13) is the sum of two power-law terms decreasing with different exponents: $-\alpha' + 1$ and $-\alpha'' + 1$. We denote by τ^* the value of $(t_2 - t_1)$ for which these two terms coincide:

$$(3.14) \quad \tau^* = \left| \frac{\alpha'(\alpha' - 1)}{\alpha''(\alpha'' - 1)} \cdot \frac{V}{K^2} \right|^{\frac{1}{2(\beta-1)-\gamma}}.$$

Then, depending on whether $(t_2 - t_1)$ stands below or above τ^* , one of the two power laws dominates in Equation (3.13) and impose its decay to the autocovariance function. In terms of long-range dependencies, each exponent corresponds to a Hurst parameter as in Equation (1.1a) ($H = \frac{3-\alpha_H}{2}$), where α_H is either equal to α' or to α'' . Rather than plotting the autocovariance (3.13), we represent in Figure 2 the corresponding *log-diagram* [2]. A log-diagram corresponds to the logarithmic variance $\log S(j)$ of the wavelet coefficients calculated at scale j (morally equivalent to the variance of the process aggregated in consecutive time windows of size $2^j\Delta$, where Δ is the data *granularity*). For a long-range dependent process of Hurst parameter H , the corresponding log-diagram increases linearly with slope $2H$.

Figure 2 emphasizes the two domains separated by τ^* :

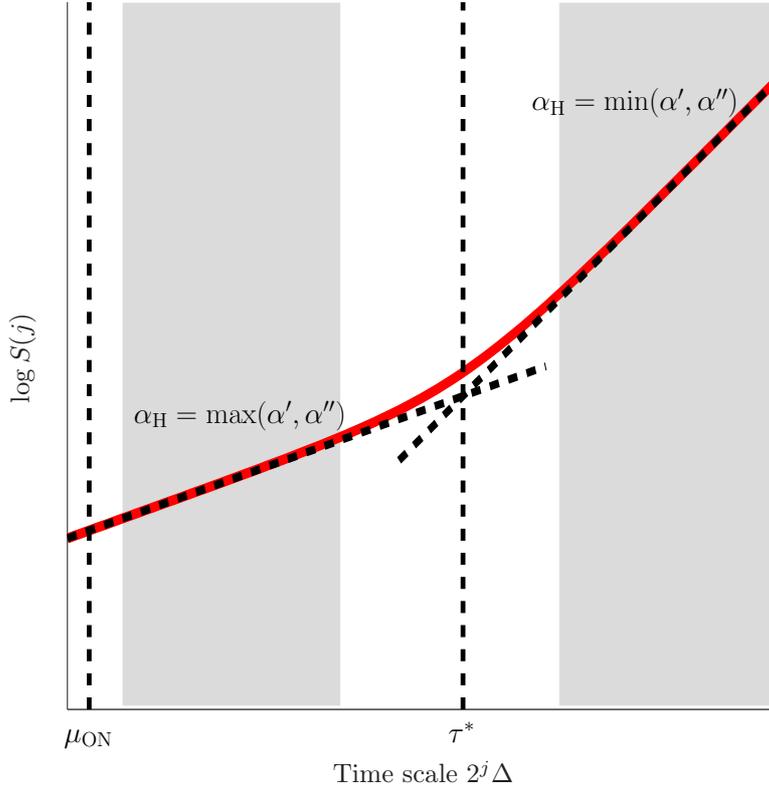


FIGURE 2. Generic log-digram of the aggregate traffic: $\log S(j)$ versus time scale j (plain line). The oblique dashed line represents power-law evolutions of the autocovariance (3.13) with indices α' and α'' defined in Equation (3.12), which correspond to the two distinguished regimes of long-range dependence (see text). Vertical dashed lines materialize the thresholds μ_{ON} and τ^* .

- at scales larger than τ^* , the smaller tail index $\alpha_{\text{H}} = \min(\alpha', \alpha'')$ governs the autocovariance decay, meaning that it is the larger Hurst exponent that characterizes the long-range dependence. This situation, referred to as asymptotic long-range dependence, holds for time lags going to infinity; it corresponds to the rigorous definition of long-range dependence.
- conversely, at scales smaller than τ^* , it is the term corresponding to the larger tail index $\alpha_{\text{H}} = \max(\alpha', \alpha'')$, *i.e.* to the smaller Hurst exponent, which dominates the autocovariance. To contrast with the asymptotic long-range dependence, this property is called pseudo long-range dependence as it holds only over a finite scale range, bounded from above by τ^* and from below by the mean flow duration μ_{ON} . Indeed, as alluded in [21] and experimentally validated in [30], μ_{ON} is a reasonable lower scale bound beyond which long-range dependence due to the heavy-tail distribution of flow duration is manifest.

The prevalence regions of these different Hurst parameters, as well as the comparison with long range dependence of usual infinite source Poisson models, depend on the parameters β and γ .

Influence of parameters β and γ . As a first remark, notice that if $\beta = 1$ ($\alpha'' = \alpha_{\text{ON}}$) or if $\gamma = 0$ ($\alpha' = \alpha_{\text{ON}}$), the classical relations (1.1) straightforwardly apply to the respective scale domains. They even extend to the entire scale axis when $\beta = 1$ and $\gamma = 0$, our model coinciding then to the usual infinite source Poisson model with (second order) mutually independent rates and durations. We denote by $H_0 = \frac{3-\alpha_{\text{ON}}}{2}$ the corresponding Hurst parameter that we take as a reference to discuss the impact of β and γ .

In the general case, τ^* delineates two long-range dependent regimes determined by α' and α'' . In terms of long-range dependence, our results are valid for all values of α' and α'' larger than one. Yet, we restrict our discussion to the case $\alpha' < \alpha''$ (i.e. $\gamma < 2(\beta - 1)$), guaranteeing the flow-size variable to be heavy-tailed distributed with index $\alpha_{\text{SI}} = \frac{\alpha_{\text{ON}}}{\beta}$ (Proposition 3.1), as it is commonly observed on real traffic traces. Then, the pseudo long-range dependence regime has a Hurst parameter $H = H_0 + \frac{\gamma}{2}$, whereas the Hurst parameter for the asymptotic long-range dependence regime reads $H = H_0 + (\beta - 1)$ and depends on the conditional mean rate exponent β :

- If $\beta > 1$ ($\alpha_{\text{SI}} < \alpha_{\text{ON}}$): This is the case that has been observed on Internet traces [12] and that we also experience with the web trace of next section. In average, the achieved rate increases with the duration of the flow, a situation due, for instance, to the transient behavior of some protocols that penalize short connections. The asymptotic long-range dependence, controlled by the tail index $\alpha_{\text{H}} = \alpha' < \alpha_{\text{ON}}$, is then stronger than the one predicted by standard models that disregard flow-scale correlations. It is worth noticing that, depending on the β value, α_{H} can possibly be smaller, equal to, or larger than α_{SI} . This means that the tail index governing the long-range dependence does not necessarily lie between α_{SI} and α_{ON} , and can be smaller than both these tail indices.
- If $\beta < 1$ ($\alpha_{\text{SI}} > \alpha_{\text{ON}}$): This situation, where the mean transmission rate of a flow decreases with its duration, could happen with scheduling policies that aim at prioritizing short flows. In this case, the asymptotic long-range dependence, controlled by the tail index $\alpha_{\text{H}} = \alpha' > \alpha_{\text{ON}}$, is weaker than the one predicted by existing models. It demonstrates that it is theoretically possible to reduce the long-range dependence strength in aggregate traffic, by activating appropriate flow-aware control mechanisms. Finally, notice that if $\gamma > 2(\beta - 1)$, the case $\beta < 1$ leads to a weakened (pseudo) long-range dependence in the intermediate scale domain.

Besides their effect on the Hurst exponents, parameters β and γ can also impact the threshold τ^* , although this latter mainly varies with the two constants K and V .

Influence of the mean K and of the variance V . The ratio $\frac{V}{K^2}$ is mostly prominent in the determination of the frontier τ^* separating the two long-range dependence regimes. Its effective role, however, is conditioned to β and γ : τ^* increases with $\frac{V}{K^2}$ when $\gamma < 2(\beta - 1)$ and it decreases otherwise. In both cases however, a large variance systematically extends the scale range where the Hurst parameter is governed by the conditional variance of the flow rate ($H = H_0 + \frac{\gamma}{2}$), whereas a small variance broadens the scale range where the Hurst parameter is determined by the conditional mean term ($H = H_0 + (\beta - 1)$). In the limit case $V = 0$, only this latter regime remains.

Let us mention that the time lag τ^* where the two autocovariance terms intersect is proportional to the characteristic flow duration for which mean squared rate and variance rate are equal:

$$(3.15) \quad d^* = \left| \frac{V}{K^2} \right|^{\frac{1}{2(\beta-1)-\gamma}}.$$

Both quantities differ by a multiplicative factor $\left| \frac{\alpha'(\alpha'-1)}{\alpha''(\alpha''-1)} \right|^{\frac{1}{2(\beta-1)-\gamma}}$, independent of K and V . This is coherent with the interpretation of the scale regions where the Hurst parameter is either prescribed by the conditional variance $H = H_0 + \frac{\gamma}{2}$ or by the conditional mean $H = H_0 + (\beta - 1)$. Notice though, that whenever the multiplicative factor significantly differs from one, only τ^* is meaningful to determine the frontier between the two regimes.

Finally, let us stress that we limited our study to the long-range dependence of aggregate traffic, a second-order statistical property fully determined by the autocovariance function. As a result, the first two conditional moments $\mathbb{E}\{R_i|D_i\}$ and $\text{Var}\{R_i|D_i\}$ were sufficient to identify this correlation structure unambiguously. To investigate statistical properties of the instantaneous bandwidth at higher orders, finer characterization of the conditional distribution $\mathbb{P}(R_i \in dr | D_i)$ would be needed.

Remark 3.5. *To model real traffic more accurately, one may want to use an heterogeneous mixture of a certain number of different classes k , with the same power-law conditional moments but with different indices β_k and γ_k . Each class leads to a set of new tail exponents $\alpha'_k = \alpha_{\text{ON}} - 2(\beta_k - 1)$ and $\alpha''_k = \alpha_{\text{ON}} - \gamma_k$. Assuming that traffics from the different classes sum independently, it is clear that the smallest of all these indices will asymptotically govern the autocovariance decrease and therefore imposes the*

long-range dependence parameter as in (1.1a). The correlation structure however, should behave more intricately in the intermediate scale range.

4. EXPERIMENTAL RESULTS AND DISCUSSION

This section drives an empirical study of our traffic model. Based on real web traces, we first present experimental evidences that support our choices, and demonstrate the ability of our approach to predict the Hurst parameter of aggregate traffic. Then we present a more systematic validation of our results based on numerical simulations.

4.1. Real web traffic. We use a trace acquired at the output link of the in2p3 research center (Lyon, France), with the capture tool *MetroFlux* [29]. The traffic is captured from the VLAN corresponding to RENATER (french national research and education network) web traffic, which is encapsulated in the 10 Gbps output link of in2p3. Although we captured more than one day of traffic, we restrict our trace to a 30 minutes stationary trace, corresponding to the incoming traffic between 3pm and 3:30pm on January 18, 2009. The mean throughput in this period is 127.3 Mbps, the mean flow duration is 0.12 s and the mean flow size 16 kBytes. The traffic variations are displayed on Figure 3-(a). While it is essentially composed of web activity, this traffic also encompasses an unusual large number of elephant flows, produced by the particular nature of experiments carried out at the in2p3 centre.

To analyze flow-level characteristics, we recover the flow sequence from the packet level trace using the widely used definition of a flow: a flow is a set of packets sharing the same source and destination IPs and ports, the same protocol, and with no inter-packet gap larger than some `timeout` that we set here to 100 ms. This relatively small value of the `timeout` was chosen because a thorough analysis of this particular trace revealed that many TCP connections are kept open even when idle, as they certainly use the TCP keep-alive protocol [8]. The high number of such connections is likely to be specific of the activities carried out at the in2p3 research center. On the other hand, most of the traffic is intra-continental with a corresponding RTT smaller than 100 ms. Nonetheless, we also performed the analysis with a `timeout` set to 1 s and verified that the results remain coherent and quite similar.¹ More generally though, as it was recently shown in [41], the flow definition can have a dramatic effect on model matching; and the one we propose here may fail at fitting the data if one adopts a different flow definition.

Figures 3-(b) and 3-(c) display the corresponding flow-size and flow-duration distributions. As expected, both exhibit power-law shapes, but with different tail indices: $\alpha_{SI} = 0.88$ and $\alpha_{ON} = 1.30$. These estimates were obtained with a state-of-the-art non-parametric wavelet method proposed in [19], using the 5-*th* derivative of a Gaussian wavelet. To avoid the bias introduced by partially observed flows, we retained only those whose starting and ending dates fall in the analyzed 30-minutes trace. Inevitably, this amounts to truncate the distribution tail, reducing thus the effective scale-range used for the tail exponent estimation. In accordance with our model, the empirical conditional mean and variance rate plotted in Figures 3-(d) and 3-(e) respectively, obey a power-law dependence with respect to the flow duration, over a significant scale range. We observe a small number of flows (22 in the analyzed trace) whose duration is larger than 100s and yet, that present a small rate. As these flows correspond to control traffic and they do not significantly contribute to the overall data traffic, we disregard them in the sequel. In Figure 3-(d), a superimposed straight line with slope $\beta = 1.48$ equal to the ratio between the tail indices α_{ON} and α_{SI} , closely matches the estimated slope ($\beta = 1.45$) obtained from a direct least-square fit of the data within the scale interval $[1, 90]$ s. The estimated value $\gamma = 0.43$ corresponds to the least-square fit of the conditional variance rate over the same scale interval, and, together with β , they meet the condition $\gamma < 2(\beta - 1)$ of Proposition 3.1. Let us mention that the estimation of γ , as well as that of α_{SI} , α_{ON} and β to a lesser extent, may be relatively sensitive to the chosen scale range

¹The quantities estimated when the `timeout` is set to 1 s are the as follows: $\alpha_{SI} = 0.99$, $\alpha_{ON} = 1.33$ (estimated with the method of [19]), $\beta = 1.34$ (deduced from α_{ON}/α_{SI}), to be compared to the value 1.31 (estimated from a least-square regression of Fig. 3-(d) adapted to the new `timeout`), $\gamma = 0.29$ (estimated from a least-square regression of Fig. 3-(e) adapted to the new `timeout`), $K = 10^{4.35}$ and $V = 10^{10.32}$ (estimated from the same least-square regressions). We deduce $\tau^* \sim 2000$ s and the predicted Hurst parameter in the pseudo long-range dependence regime $H = 0.98$, fully consistent with the estimated value.

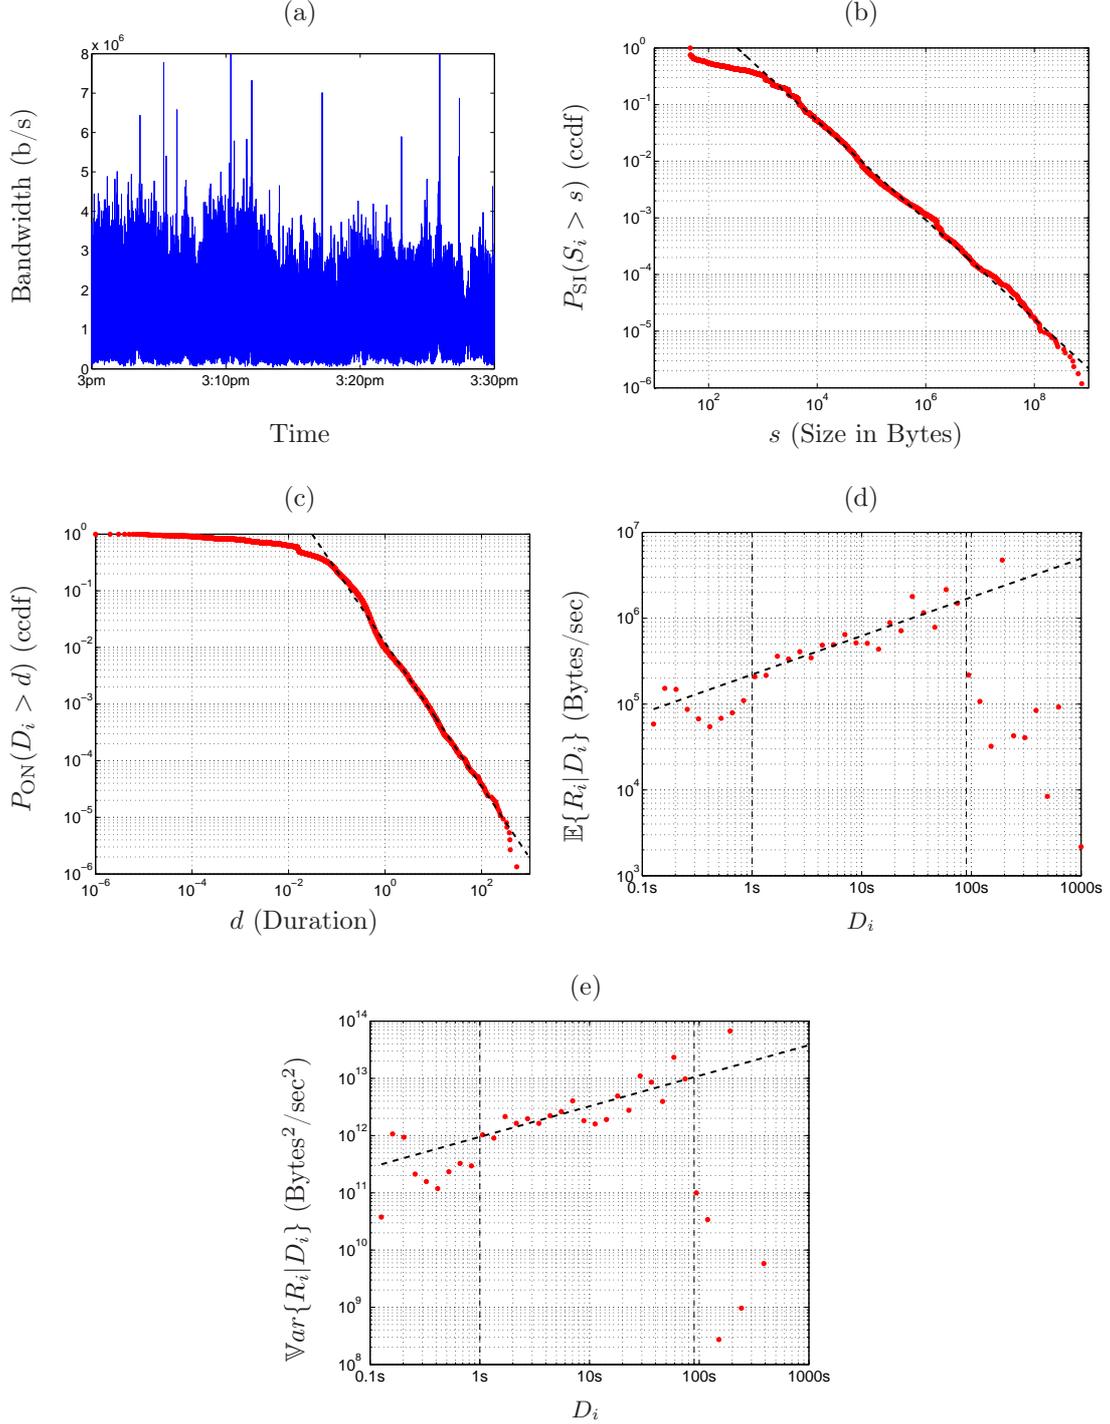


FIGURE 3. Characteristics of the in2p3 trace: (a) Aggregate bandwidth evolution (granularity 10 ms) – (b) Flow-size distribution, complementary cdf (dashed black: Pareto with tail index $\hat{\alpha}_{SI} = 0.88$ estimated with the method of [19]) – (c) Flow-duration distribution, complementary cdf (dashed black: Pareto with tail index $\hat{\alpha}_{ON} = 1.30$ estimated with the method of [19]) – (d) Conditional mean rate (dashed black: model $\mathbb{E}\{R_i | D_i\} = K D_i^{\beta-1}$ with $\beta = \hat{\alpha}_{ON} / \hat{\alpha}_{SI}$ and value $K = 10^{5.34}$ estimated by least-square fit in the range $[1, 90]$ s) – (e) Conditional variance of the rate (dashed black: model $\text{Var}\{R_i | D_i\} = V D_i^\gamma$ with values $V = 10^{11.97}$ and $\gamma = 0.43$ estimated by least-square fit in the range $[1, 90]$ s). The mean flow duration is 0.12 s and the mean flow size is 16 kBytes. In plots (d) and (e), logarithmic binning was used.

for regressing the data. For different regression ranges within the interval $[1, 100]$ s, the estimate of γ varies roughly from 0.2 up to 0.5. No doubt that in our case, this sensitivity is accentuated by the quite short length of our stationary trace that leads to only a limited scale domain.

From definition (3.12), we then deduce the exponents $\alpha' = 0.34$ and $\alpha'' = 0.87$. Theoretically, $\alpha' < 1$ corresponds to a diverging autocovariance (3.13) as time lag $(t_2 - t_1)$ goes to infinity; but since we have a finite-size trace, the distributions are bounded and tail indices smaller than one (or Hurst parameter larger than 1) can be observed over a finite scale range. As for γ , we estimate $K = 10^{5.34}$ and $V = 10^{11.97}$ from the least-square fits of the plots in Figures 3-(d) and 3-(e), which leads to the threshold $\tau^* \sim 1,000$ s.

To investigate the long-range dependence of the aggregate traffic, we used the wavelet-based method described in [2], with a Daubechies wavelet with three zero moments. This method provides a robust estimation of the Hurst parameter within a confidence interval derived under normal assumption (we checked that the traffic was fairly gaussian, with a kurtosis value equal to 3.15 at an aggregation scale of 10 ms). Figure 4 displays the log-diagram of the aggregate traffic bandwidth calculated at the granularity $\Delta = 10$ ms. It follows a linear trend that is characteristic of an underlying scaling law, and whose slope reflects the Hurst parameter. By linear regression over the scale range $[1.28, 40.96]$ s (which lies far beyond the mean flow duration of 0.12 s), we estimate a Hurst exponent $\hat{H} = 0.92 \pm 0.06$. The upper bound of this regression scale interval stands far below the threshold τ^* between the two long-range dependence regimes. Under these circumstances, Proposition 3.2 predicts a dominating Hurst exponent $H = \frac{3-\alpha''}{2} = 1.06$, which is slightly higher than the estimated value. In the present case, relation (1.1) obtained with standard infinite source Poisson models, would yield a similar Hurst estimate, slightly lower than the estimated value (0.85). Nonetheless, we believe that the elaborated model of Section 3 fills a gap in traffic modeling and explains why approximate models that disregard flow-level correlations, can still apply in restricted scale ranges. Moreover, our analysis sheds a new light on the scaling observed from the real trace: it is not reminiscent of strict long-range dependence, but of *pseudo* long-range dependence. The value observed for the threshold τ^* indicates that asymptotic long-range dependence can actually not be observed on realistic-length stationary traces, at least with parameters similar to the one observed in our trace. It is also worth noticing that our study comes to the same conclusion as in [1, 41], but using a different modeling path. As discussed above though, this conclusion may vary with the considered flow definition: then, it is possible that another choice would lead to a smaller τ^* observable in practice. Future Internet traffic may also exhibit smaller τ^* due to increasing mean rates. Should this be the case, our model would certainly constitute a relevant tool to predict the long-range dependence in the asymptotic regime beyond τ^* .

As a last remark, let us point out that the flow-arrival process (*i.e.* the count process associated with the point process $\{T_i, i \geq 0\}$) is not strictly Poisson, but exhibits a slight long-range dependence with an estimated Hurst parameter of 0.65. However, after we shuffled the flow arrival times to annihilate correlations in the flow arrival process, we verified that long-range dependence observed on the aggregate traffic remained strictly unchanged. This is fully consistent with natural intuition drawn from the knowledge of ON/OFF models, and may also be explained by higher-level structure such as sessions, as it is shown in [41].

With this experimental set, we were not able to confirm the existence of a second long-range dependence regime, as the available data size did not permit to scrutinize scales beyond the critical time lag τ^* . To more thoroughly study the ins and the outs of Proposition 3.2, we now resort to numerical simulations.

4.2. Model validation via simulations. In this section, we use MATLAB simulations to generate controlled traffic traces in order to demonstrate the ability of the model proposed in Section 3 to accurately predict the Hurst parameter in different scale ranges and for different sets of parameters.

Our simulations consist in traffic from an infinite source Poisson model, where we consider sufficiently large flow arrival rates to reach the Gaussian limit case [23] (see Section 2). Moreover, we use a fluid simulation with constant bitrate over the full duration of the flow. Infinite source Poisson models with independent but intra-flow-varying rates were proposed (see Section 2.2), leading to no difference regarding large-scale correlations. It has also been experimentally verified in [30] that the protocol

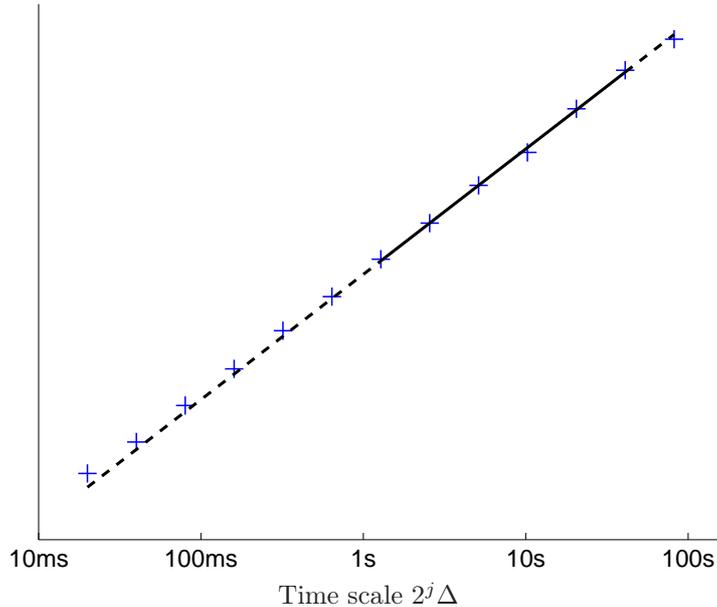


FIGURE 4. Log-diagram of the in2p3 aggregate traffic bandwidth in time windows of size $\Delta = 10$ ms. The estimated value of the Hurst parameter is $\hat{H} = 0.92 \pm 0.06$ (scale range of estimation: $[1.28, 40.96]$ s). Recall that the mean flow duration is $\mu_{\text{ON}} = 0.12$ s.

specificities only affect short time-scale properties, but leave unchanged traffic features at coarser scale, and in particular the long-range dependence property that we are interested in here. We therefore do not consider such intra-flow-level refinements that fall out of the present scope.

4.2.1. *The two long-range dependence regimes.* Due to the lack of long-term stationarity, our previous analysis of a real trace was confined to the pseudo long-range dependence regime. To overcome this limitation, we generate a stationary trace possessing a flow structure very similar to the one of the in2p3 trace: Poisson flow arrival, $\alpha_{\text{ON}} = 1.2$, $\mu_{\text{ON}} = 0.12$ s, $\beta = 1.4$, $\gamma = 0.1$, $K = 10^{5.4}$, $V = 10^{12.3}$, but with corresponding smaller values of $\tau^* \sim 300$ s and of $d^* \sim 100$ s. We also generate the trace over a much longer period (140 hours). The flows' rates are drawn at random according to a family of gamma distributions that adequately condition the rate's mean and variance to the flow duration (we observed that a gamma assumption is also reasonable for the in2p3 trace of previous section).

The resulting log-diagram of the bandwidth at granularity $\Delta = 0.05$ s is displayed on Figure 5. Like the theoretical log-diagram of Figure 2, it clearly reveals the two long-range dependence regimes discussed in Section 3.3. Moreover, the quantities τ^* and d^* seem to be reasonable separators of these two domains. In both domains, the estimated Hurst parameters are in good agreement with the theoretical values predicted by the model of Section 3: $\hat{H} = 0.98 \pm 0.01$, *versus* $H = 0.95$ for $\tau < \tau^*$ and $\hat{H} = 1.27 \pm 0.10$, *versus* $H = 1.3$ for $\tau > \tau^*$. Despite the restricted scale ranges retained for the regression to impose a good separation with the threshold τ^* , small differences observed in each regime are consistent with a reminiscent effect of the slope in the other regime.

In contrast to existing ones, the model developed in Section 3 shows that the positive correlation ($\beta > 1$) that naturally exists between flow rates and flow durations in Internet traces is able to increase long-term correlations in aggregate traffic. This effect was not observed in the previous section because of a too large threshold τ^* as compared to the short stationary trace. However, it is susceptible to become a common feature of the future Internet traffic if, for instance, the variance V remains unchanged while the mean flow rate keeps augmenting (hence threshold τ^* diminishes).

4.2.2. *Relation “ H versus α_{ON} ”.* In this last section, we complete the numerical validation of the relation between the Hurst parameter(s) and the tail index of the flow-duration distribution, when the conditional

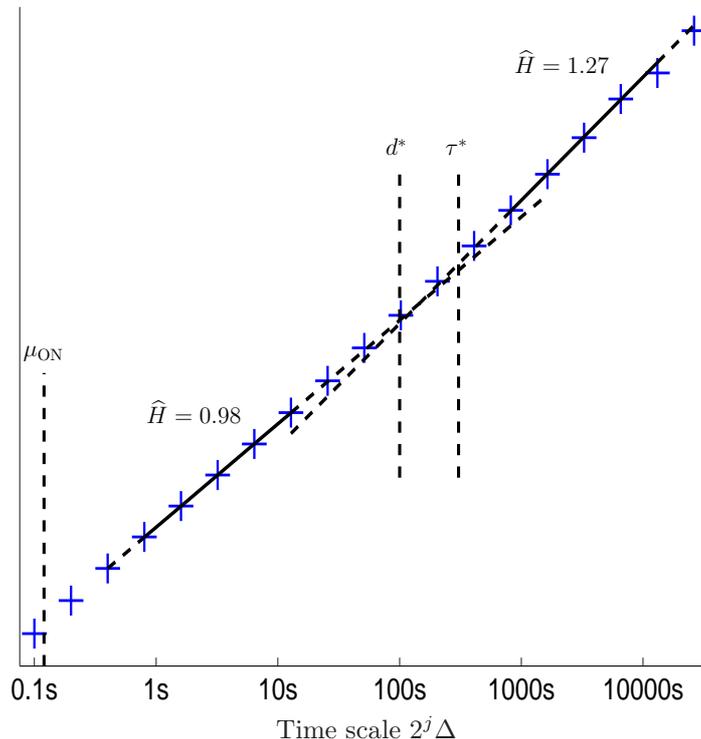


FIGURE 5. Log-diagram of the long-trace simulation’s aggregate traffic bandwidth in time windows of size $\Delta = 50$ ms. The Hurst parameter’s estimations are $\hat{H} = 0.98 \pm 0.01$ (in the scale range $[0.8, 12]$ s) and $\hat{H} = 1.27 \pm 0.10$ (in the scale range $[800, 13,000]$ s).

parameters β and γ vary. The classical relation (1.1) have already been validated on simulators [39], numerically [1] and on a real experimental platform [30]. Here, we concentrate on situations leading to the observation of a new long-range dependence regime with a Hurst exponent that usual models fail at predicting, but that our approach identified as $H = \frac{3-\alpha'}{2} = H_0 + (\beta - 1)$, or $H = \frac{3-\alpha''}{2} = H_0 + \frac{\gamma}{2}$, depending on the observed scale domain.

In the following experiments, $\mu_{ON} = 0.1$ s and $K = 10^{5.2}$ are kept constant with the same values as in the previous section. For different values of the pairs $(\alpha_{ON}, \beta; \gamma = 0)$ and $(\alpha_{ON}, \gamma; \beta = 1.4)$, we generate two collections of 25-hours traces. The first set with (α_{ON}, β) varying serves to validate the relation $H = \frac{3-\alpha'}{2} = H_0 + (\beta - 1)$, the conditional variance term is disabled ($\gamma = 0$) and the rate variance set to an arbitrary small value $V = 1$. Conversely when the parameters’ pair (α_{ON}, γ) varies and $\beta = 1.4$ remains constant, V is fixed to a sufficiently large value to ensure that the threshold τ^* lies beyond the observational scale range involved in the prediction expression $H = \frac{3-\alpha''}{2} = H_0 + \frac{\gamma}{2}$. We then confront the resulting Hurst exponents estimated in the adequate scale ranges to the theoretical predictions of our model.

Experimental results displayed in Figure 6 are in good agreement with the expected Hurst exponent values. The slight deviation observed around the knee point of each theoretical curve (where the Hurst parameter hits its critical value $1/2$ and stabilizes) is fully consistent with the arguments exposed in [1, 30] where the classical relation (1.1) is experimentally validated.

5. CONCLUSION

In this work, we extended the widely used relation (1.1) linking the aggregate traffic’s Hurst parameter to the flow-duration distribution’s tail index. We proposed a variant of the infinite source Poisson model where flow rates and flow durations are correlated (hence the flow-size and flow-duration distributions

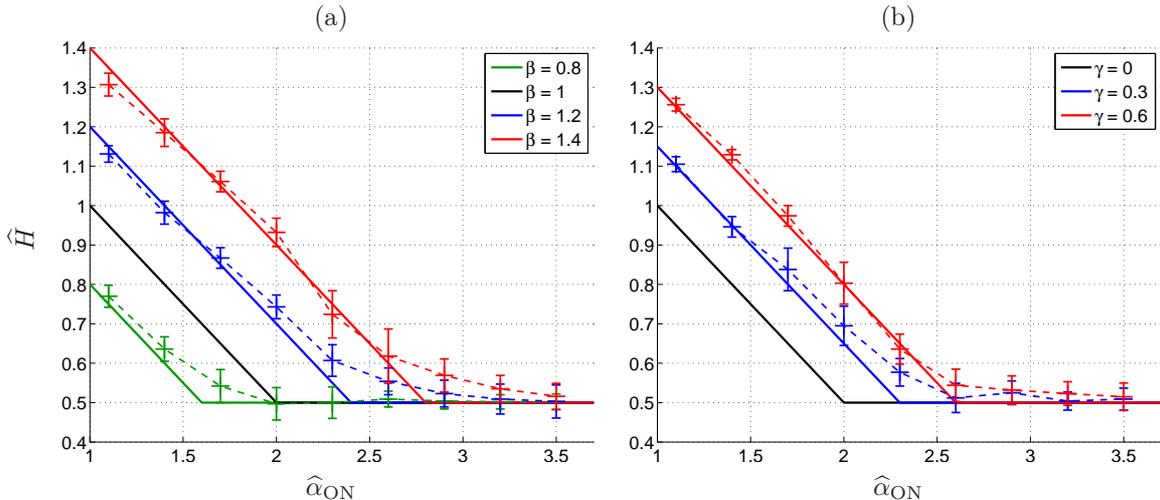


FIGURE 6. Validation of the relation “ H versus α_{ON} ” in the two different domains: (a) For $\gamma = 0$ and three different values of β in the domain where $H = \frac{3-\alpha'}{2} = H_0 + (\beta - 1)$ (for $\beta < 1$, the Hurst parameter corresponds to the pseudo long-range dependence regime; whereas it corresponds to the asymptotic long-range dependence for $\beta > 1$). – (b) For $\beta = 1.4$ and two different values of γ in the domain where $H = \frac{3-\alpha''}{2} = H_0 + \frac{\gamma}{2}$ (the Hurst parameter corresponds here to the pseudo long-range dependence regime). Solid lines represent the theoretical relations, while dashed lines are drawn from empirical estimations. Confidence intervals displayed are provided by the wavelet method of [2].

have different tail indices α_{SI} and α_{ON}). The model mainly relies on two parameters β and γ that fix the power-law evolution of the first two conditional moments of the rate given the duration. The β parameter also relates the two tail indices α_{ON} and α_{SI} .

Then, we showed that there exists two possible regimes of long-range dependence, each corresponding to a distinct heavy-tail phenomenon of intensity depending on β and γ respectively. We characterized the threshold between those regimes, as a function of the β and γ parameters, but also of the rate’s mean and variance. Our analysis of a real web-traffic trace revealed that this limit has such a value that the observed domain on reasonable-length stationary traces corresponds to pseudo long-range dependence and not to real (asymptotic) long-range dependence. On that point, our study joins the conclusions of [1, 41], using a different modeling path. We also showed that, depending on the parameters β and γ , the effect of the correlation between flow rates and durations can be either to accentuate or to weaken the long-range dependence. Our results extend previously proposed predictions like relation (1.1), with which it naturally coincides if the correlation vanishes.

The model proposed in this work is the first model, to the best of our knowledge, that includes the correlation between flow rates and flow durations. This correlation is very important. It has been observed for over a decade [12, 35] on Internet traces, and is likely to become an even more important parameter in the future Internet with the emergence of flow-aware approaches. Our results show how this correlation can modify the long-range dependence classically introduced by heavy-tailed flow durations. Then, not only does our model show the impact of natural correlations (due *e.g.* to protocol’s effects) on the aggregate traffic; but it also opens the possibility to finely control (*e.g.* reduce) the long-range dependence via flow-aware control procedures. One limitation of our model though, is that it considers an open-loop protocol (without feedback reaction); thus it may fail to model traffic from highly congested links. This certainly is an interesting direction to investigate, in particular when the correlation parameters possibly depend on the load (as it could be the case for instance with some scheduling policies).

From a mathematical viewpoint, our work introduces a new class of models based on non-independently-marked point processes, from which various processes, such as the instantaneous throughput considered in this paper, can be derived. The study of these processes, under various rescaling and limit regimes and at both large and small scales, is certainly an interesting direction to complement the present work, as they are likely to reveal new properties.

6. ACKNOWLEDGMENT

The authors would like to thank anonymous referees for their thorough review of our work and their valuable suggestions that led to improvements of this paper. The first author would also like to thank B. Loiseau for his helpful suggestions.

APPENDIX A. PROOF OF PROPOSITION 3.1

The proof relies on the fact that if $\gamma < 2(\beta - 1)$, the conditional variance of the rate is asymptotically much smaller than its square mean so that the rate becomes almost deterministically equal to $KD_i^{\beta-1}$ for long durations.

Let $s > 0$ and $\varepsilon > 0$, and write:

$$\begin{aligned} \mathbb{P}(S_i > s) &= \int_{d_{\min}}^{\infty} \mathbb{P}(R_i > \frac{s}{d} | D_i = d) \mathbb{P}(D_i \in dd) \\ &= \int_{d_{\min}}^{(\frac{s}{K(1+\varepsilon)})^{\frac{1}{\beta}}} \mathbb{P}(R_i > \frac{s}{d} | D_i = d) \mathbb{P}(D_i \in dd) \\ &\quad + \int_{(\frac{s}{K(1+\varepsilon)})^{\frac{1}{\beta}}}^{(\frac{s}{K(1-\varepsilon)})^{\frac{1}{\beta}}} \mathbb{P}(R_i > \frac{s}{d} | D_i = d) \mathbb{P}(D_i \in dd) \\ &\quad + \int_{(\frac{s}{K(1-\varepsilon)})^{\frac{1}{\beta}}}^{\infty} \mathbb{P}(R_i > \frac{s}{d} | D_i = d) \mathbb{P}(D_i \in dd). \end{aligned}$$

We denote by $A(s, \varepsilon)$, $B(s, \varepsilon)$ and $C(s, \varepsilon)$ the three terms of this sum.

To handle the first term $A(s, \varepsilon)$, note that for $d \leq (\frac{s}{K(1+\varepsilon)})^{\frac{1}{\beta}}$, by Chebychev's inequality,

$$\begin{aligned} \mathbb{P}(R_i > \frac{s}{d} | D_i = d) &\leq \mathbb{P}(|R_i - Kd^{\beta-1}| > \frac{s}{d} - Kd^{\beta-1} | D_i = d), \\ &\leq \frac{Vd^\gamma}{(\frac{s}{d} - Kd^{\beta-1})^2}. \end{aligned}$$

Then, for some constants G independent of s and ε , we have:

$$\begin{aligned} A(s, \varepsilon) &\leq \int_{d_{\min}}^{(\frac{s}{K(1+\varepsilon)})^{\frac{1}{\beta}}} \frac{Vd^\gamma}{(\frac{s}{d} - Kd^{\beta-1})^2} \frac{\alpha_{\text{ON}}/d_{\min}}{(\frac{d}{d_{\min}})^{\alpha_{\text{ON}}+1}} dd \\ &= G \int_{Kd_{\min}^\beta}^{\frac{s}{1+\varepsilon}} \frac{x^{\frac{\gamma+2-\alpha_{\text{ON}}}{\beta}-1}}{(s-x)^2} dx. \end{aligned}$$

Using the series expansion:

$$\frac{1}{(s-x)^2} = \frac{1}{s^2} \sum_{n=1}^{\infty} n \left(\frac{x}{s}\right)^{n-1},$$

with uniform convergence for $x \in [Kd_{\min}^\beta, \frac{s}{1+\varepsilon}]$, we deduce:

$$A(s, \varepsilon) \leq G s^{\frac{\gamma+2-\alpha_{\text{ON}}}{\beta}-2} \sum_{n=1}^{\infty} \frac{n}{\left(\frac{\gamma+2-\alpha_{\text{ON}}}{\beta} - 1 + n\right) (1+\varepsilon)^{\frac{\gamma+2-\alpha_{\text{ON}}}{\beta}-1+n}} - \frac{G}{s} \sum_{n=1}^{\infty} \frac{n(Kd_{\min}^\beta)^{\frac{\gamma+2-\alpha_{\text{ON}}}{\beta}-1}}{\frac{\gamma+2-\alpha_{\text{ON}}}{\beta} - 1 + n} \left(\frac{Kd_{\min}^\beta}{s}\right)^n.$$

Since by hypothesis, $\gamma < 2(\beta - 1)$, we can take ε of the form $\varepsilon = \varepsilon(s) = s^{-\eta}$ with $0 < \eta < \frac{2(\beta-1)-\gamma}{\beta}$, which ensures that

$$(A.1) \quad A(s, \varepsilon) \underset{s \rightarrow \infty}{=} o\left(s^{-\frac{\alpha_{\text{ON}}}{\beta}}\right).$$

It is easy to see that we also have:

$$(A.2) \quad B(s, \varepsilon) \underset{s \rightarrow \infty}{=} o\left(s^{-\frac{\alpha_{\text{ON}}}{\beta}}\right).$$

For the last term $C(s, \varepsilon)$, application of Chebychev's inequality again shows that:

$$C(s, \varepsilon) = \int_{\left(\frac{s}{K(1-\varepsilon)}\right)^{\frac{1}{\beta}}}^{\infty} (1 - \delta(d)) \mathbb{P}(D_i \in dd),$$

with

$$|\delta(d)| \leq \frac{Vd^{\gamma-2(\beta-1)}}{\varepsilon^2 K^2}.$$

Thus, recalling that we took $\varepsilon = s^{-\eta}$ with $0 < \eta < \frac{2(\beta-1)-\gamma}{\beta}$, it is clear that

$$(A.3) \quad C(s, \varepsilon) \underset{s \rightarrow \infty}{=} \left(\frac{s}{K(1-\varepsilon)}\right)^{-\frac{\alpha_{\text{ON}}}{\beta}} + o\left(s^{-\frac{\alpha_{\text{ON}}}{\beta}}\right),$$

which completes the proof of Proposition 3.1.

REFERENCES

- [1] P. Abry, P. Borgnat, F. Ricciato, A. Scherrer, and D. Veitch. Revisiting an old friend: On the observability of the relation between long range dependence and heavy tail. *Telecommunication Systems, Special Issue on Traffic Modeling, Its Computations and Applications*, 43(3-4):147–165, 2009.
- [2] P. Abry and D. Veitch. Wavelet analysis of long-range dependent traffic. *IEEE Transactions on Information Theory*, 44(1):2–15, January 1998.
- [3] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and P. Owezarski. A flow-based model for internet backbone traffic. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement (IMW'02)*, pages 35–47, 2002.
- [4] P. Barford and M. Crovella. Generating representative web workloads for network and server performance evaluation. In *Proceedings of ACM Sigmetrics/Performance'98*, pages 151–160, 1998.
- [5] J. Barral and B. Mandelbrot. Multifractal products of cylindrical pulses. *Probability Theory and Related Fields*, 124:409–430, 2002.
- [6] F. Benevenuto, T. Rodrigues, M. Cha, and V. Almeida. Characterizing user behavior in online social networks. In *Proceedings of the Internet Measurement Conference (IMC'09)*, pages 49–62, 2009.
- [7] J. Beran. *Statistics for Long-memory processes*. Chapman & Hall, 1994.
- [8] R. Braden. Requirements for Internet Hosts – Communication Layers. RFC 1122, October 1989.
- [9] P. Chainais, R. Riedi, and P. Abry. On non-scale-invariant infinitely divisible cascades. *IEEE transactions on information theory*, 51(3):1063–1083, 2005.
- [10] D. R. Cox and V. Isham. *Point Processes*. Chapman & Hall, 1980.
- [11] D.R. Cox. Long range dependence: A review. In *Statistics: An Appraisal*, pages 55–74. Iowa State University Press, 1984.
- [12] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes (extended version). *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.
- [13] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods*. Springer, 2nd edition, 2003.
- [14] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes, Volume II: General Theory and Structure*. Springer, 2nd edition, 2007.
- [15] B. D'Auria and S.I. Resnick. Data network models of burstiness. *Advances in Applied Probability*, 38(2):373–404, 2006.

- [16] B. D’Auria and S.I. Resnick. The influence of dependence on data network models. *Advances in Applied Probability*, 40(1):60–94, 2008.
- [17] P. Doukhan, G. Oppenheim, and M.S. Taqqu. *Long-Range Dependence: Theory and Applications*. Birkhäuser, 2003.
- [18] R. Gaigalas and I. Kaj. Convergence of scaled renewal processes and a packet arrival model. *Bernoulli*, 9(4):671–703, 2003.
- [19] P. Gonçalves and R. Riedi. Diverging moments and parameter estimation. *Journal of the American Statistical Association (JASA)*, 100(472):1382–1393, December 2005.
- [20] C. A. Guerin, H. Nyberg, O. Perrin, S. Resnick, H. Rootzén, and C. Starica. Empirical testing of the infinite source poisson data traffic model. *Stochastic Models*, 19(2):151–200, 2003.
- [21] N. Hohn, D. Veitch, and P. Abry. Cluster processes, a natural language for network traffic. *IEEE Transactions on Signal Processing – Special Issue on Signal Processing in Networking*, 8(51):2229–2244, October 2003.
- [22] I. Kaj and M.S. Taqqu. Convergence to fractional brownian motion and to the telecom process: the integral representation approach. In V. Sidoravicius and M.E. Vares, editors, *In an Out of Equilibrium 2*, volume 20 of *Progress in Probability*, pages 383–427. Birkhäuser, 2008.
- [23] J. Kilpi and I. Norros. Testing the gaussian approximation of aggregate traffic. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement (IMW’02)*, pages 49–61, 2002.
- [24] J.F.C. Kingman. *Poisson processes*. Oxford university press, 1993.
- [25] T. G. Kurtz. Limit theorems for workload input models. In *Stochastic networks: theory and applications*. Clarendon Press, Oxford, 1996.
- [26] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of ethernet traffic. In *Proceedings of ACM SIGCOMM*, pages 183–193, 1993.
- [27] J.B. Levy and M.S. Taqqu. On renewal processes having stable inter-renewal intervals and stable rewards. *Les Annales des Sciences Mathématiques du Québec*, 11:95–110, 1987.
- [28] J.B. Levy and M.S. Taqqu. Renewal reward processes with heavy-tailed inter-renewal times and heavy-tailed rewards. *Bernoulli*, 6(1):23–44, 2000.
- [29] P. Loiseau, P. Gonçalves, R. Guillier, M. Imbert, Y. Kodama, and P. Vicat-Blanc Primet. Metroflux: A high performance system for analyzing flow at very fine-grain. In *Proceedings of TridentCom*, April 2009.
- [30] Patrick Loiseau, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Investigating self-similarity and heavy-tailed distributions on a large scale experimental facility. *IEEE/ACM Transactions on Networking*, 18(4):1261–1274, August 2010.
- [31] S. B. Lowen and M. C. Teich. *Fractal-Based Point Processes*. Wiley, 2005.
- [32] B. Mandelbrot. Long-run linearity, locally gaussian process, h-spectra and infinite variances. *International Economic Review*, 10(1):82–111, February 1969.
- [33] M.R.H. Mandjes and N.K. Boots. The shape of the loss curve and the impact of long-range dependence on network performance. *AEU - International Journal of Electronics and Communications*, 58(2):101–117, 2004.
- [34] M.R.H. Mandjes and P. Zuraniewski. A queuing-based approach to overload detection. In *Proceedings of NET-COOP’09*, 2009.
- [35] N.M. Markovich and J. Kilpi. Bivariate statistical analysis of tcp-flow sizes and durations. *Annals of Operations Research*, 170(1):199–216, September 2009.
- [36] K. Maulik and S.I. Resnick. Small and large time scale analysis of a network traffic model. *Queueing Systems Theory and Applications*, 43(3):221–250, 2003.
- [37] T. Mikosch, S.I. Resnick, H. Rootzén, and A. Stegeman. Is network traffic approximated by stable lévy motion or fractional brownian motion? *The Annals of Applied Probability*, 12(1):23–68, 2002.
- [38] H.X. Nguyen, P. Thiran, and C. Barakat. On the correlation of tcp traffic in backbone networks. In *Proceedings of IEEE International Symposium on Circuits and Systems (ISCAS)*, 2004.
- [39] K. Park, G. Kim, and M. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *Proceedings of the International Conference on Network Protocols*, pages 171–180, 1996.
- [40] V. Paxson and S. Floyd. Wide area traffic: The failure of Poisson modeling. In *Proceedings of ACM SIGCOMM*, pages 257–268, 1994.
- [41] F. Ricciato, A. Coluccia, A. D’Alconzo, D. Veitch, P. Borgnat, and P. Abry. On the role of flows and sessions in internet traffic modeling: an explorative toy-model. In *Proceedings of IEEE GLOBECOM’09*, pages 2880–2887, 2009.
- [42] M. S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *ACM SIGCOMM Computer Communication Review*, 27(2):5–23, 1997.
- [43] M.S. Taqqu and J.B. Levy. Using renewal processes to generate long-range dependence and high variability. In *Dependence in Probability and Statistics*, pages 73–89. Birkhäuser, 1986.
- [44] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, 5(1):71–86, 1997.

(P. Loiseau) UNIVERSITY OF CALIFORNIA, SANTA CRUZ, CA, USA
E-mail address: ploiseau@soe.ucsc.edu

(P. Vicat-Blanc Primet and P. Gonçalves) INRIA RHÔNE-ALPES/ENS LYON, UNIVERSITÉ DE LYON, LYON, FRANCE
E-mail address: {Pascale.Primet, Paulo.Goncalves}@ens-lyon.fr