Discrete-time heavy-tailed chains, and their properties in modelling network traffic

JOSÉ ALBERTO HERNÁNDEZ and IAIN W. PHILLIPS
Loughborough University, United Kingdom
JAVIER ARACIL
Universidad Autónoma de Madrid, Spain

The particular statistical properties found in network measurements, namely self-similarity and long-range dependence, cannot be ignored in modelling network and Internet traffic. Thus, despite their mathematical tractability, traditional Markov models are not appropriate for this purpose, since their memoryless nature contradicts the burstiness of transmitted packets. However, it is desirable to find a similarly tractable model which is, at the same time, rigorous at capturing the features of network traffic.

This work presents the *discrete-time heavy-tailed chains*, a tractable approach to characterise network traffic as a superposition of discrete-time "on/off" sources. This is a particular case of the generic "on/off" heavy-tailed model, thus showing the same statistical features as the former; particularly, self-similarity and long-range dependence, when the number of aggregated sources approaches infinity.

The model is then applicable to characterise a number of discrete-time communication systems, for instance ATM and Optical Packet Switching, and further derive meaningful performance metrics, such as the average burst duration and the number of active sources in a random instant.

1. INTRODUCTION

The analysis of network measurements is crucial in understanding network behaviour and evaluating the actual performance levels of networks [Frost and Melamed 1994]. In the first instance, such analysis must be followed by an attempt to extract the properties and features of measurements, and secondly must set the grounds for proposing models based on such empirical evidence, which are at the same time capable of generalisation to other scenarios [Floyd and Kohler 2003].

In the past years, network modellers have undertaken research attempting to find

the particular statistical characteristics of network traffic, and have reported their highly-variable nature, and the difficulties in capturing their properties [Feldmann et al. 1998; Karagiannis et al. 2002]. However, this challenge has been partially solved with the identification of long-range dependence, which persists in the network traces at different scenarios [Leland et al. 1994; Beran et al. 1995; Paxson and Floyd 1995; Crovella and Bestavros 1997].

Several monofractal models have been proposed to characterise network traffic, see the so-called "black-box" models like the Fractional Gaussian Noise [Norros 1995] and fractional ARIMA time-series [López-Ardao et al. 2000], and the structural "on/off" source model [Willinger et al. 1997]. The latter stands out since it also gives an explanation of the causes inducing self-similarity and long-range dependence in the traces, rather than just fitting the statistical properties without a justification. Finally, Feldmann et al. [1998] propose a cascade model, motivated by the layered structure of the TCP/IP protocol stack, which explains the multifractal behaviour of network traffic.

Essentially, the traffic volume traversing a certain monitored point in the network can be considered as a combination of multiple two-state flows, each of them switching from active periods of packet transmission (the "on" state) to idle or silent periods (the "off" state). This approach has motivated a large number of studies on Markov processes, see [Adas 1997] for a summary of them. However, Markov chains are short-memory processes, thus failing to capture the burstiness of packet streams, hence unsuitable in modelling long-range dependent processes such as network traffic [Paxson and Floyd 1995; Willinger and Paxson 1998]. Instead, according to Willinger et al. [1997], long-range dependence can be induced if either the active or both the active and the idle periods of transmission are heavy-tailed distributed. Furthermore, when the number of sources and the time-scale are large, the aggregated traffic approximates the Fractional Gaussian Noise [Taqqu et al. 1997].

This work introduces the *discrete-time heavy-tail chains*, a discrete-time model in the light of the "on/off" heavy-tailed source model defined by Willinger et al. [1997], with a tractable methodology such as that of classical Markov chains. Its main characteristics and properties are also studied, along with a brief discussion of its relationship with long-range dependent processes and, particularly, Fractional Brownian Motion.

This discrete-time nature of this model makes it applicable to a number of synchronous TDM (aka, Time-Division Multiplexed) communication systems, for example, slotted wireless communications channels (GSM/UMTS), Asynchronous Transfer Mode (ATM) and possibly synchronous Optical Packet Switching [O'Mahony M., Simeonidou D., Hunter D. and Tzanakaki A. 2001; Guillemot C. et al 1998]. In the light of this, the discrete-time version of the classical "on/off" heavy-tailed source model, proposed in the following, is of interest in the simulation, modelling and performance evaluation of such discrete-time communication scenarios.

After the preliminaries outlined in section 2, the remainder of this work is organised as follows: sections 3 and 4 introduce the *discrete-time heavy-tailed chains* model, and analyze its main properties and features. Section 5 provides the link from the discrete-time heavy-tailed chains to long-range dependence. Section 6

shows how to estimate the model parameters following the Hill's estimator. Finally, section 7 shows how to infer the average number of active sources in a given link, and section 8 summarises the main contributions of this work.

2. PRELIMINARIES

Let $Y = \{Y_n, n = 1, 2, \ldots\}$ be a stationary random process which takes values 0 (idle or "off" state) or 1 (active or "on" state). Then, such sequence Y is said to be a Markov chain if it satisfies the Markov or memoriless property given by:

$$
\begin{aligned}
\mathbb{P}(Y_{n+1} &= y_{n+1}|Y_n = y_n, Y_{n-1} = y_{n-1}, \ldots, Y_{n-k} = y_{n-k}) \\
&= \mathbb{P}(Y_{n+1} = y_{n+1}|Y_n = y_n), \qquad y_j \in \{0,1\}, \quad j \in \mathbb{Z}^+, \quad k = 1, \ldots, n-1
\end{aligned}
\tag{1}
$$

Under such "memoriless" property and starting from an "on" state, the probability to have a sequence of d consecutive "on" states, starting in an "on" state, is given by a geometric distribution, i.e:

$$
p(d) = (1 - p_{11})p_{11}^{d-1}, \quad d \in \mathbb{Z}^+
\tag{2}
$$

where p_{11} refers to the probability of jumping from an "on" state to an "on" state in one step, that is, of remaining in the same "on" state: $p_{11} = \mathbb{P}(Y_2 = 1|Y_1 = 1)$. According to this, traditional Markov chains exhibit the following drawbacks concerning the modelling of network traffic:

(1) The probability of remaining in an "on" state is a constant. This contradicts the empirically observed burstiness phenomenon. Namely, in real networks, when a packet arrives, it is very likely that more packets follow its arrival [Jain and Routhier 1986; Paxson and Floyd 1995]. The conditional probability of remaining in an "on" state should then increase.

(2) The probability of d consecutive "on" states decays exponentially following a geometric distribution (eq. 2), instead of hyperbolically or Pareto-like as a condition to induce long-range dependence [Willinger et al. 1997].

The following section proposes a modification of the memoriless nature of the Markov condition to encourage the creation of bursts with Pareto-like duration.

3. MODEL DEFINITION

Definition 3.1 *Discrete-time heavy-tailed chain.* The random process $X = \{X_n, n = 1, 2, \ldots\}$, which takes values 0 (idle or "off" state) or 1 (active or "on" state), is a discrete-time heavy-tailed chain if it satisfies:

$$
\begin{aligned}
\mathbb{P}(X_{n+1} &= 1|X_n = 1, \ldots, X_{n-k+1} = 1, X_{n-k} = 0) = \\
&= \mathbb{P}(X_{n+1} = 1|X_n = 1, \ldots, X_{n-k+1} = 1, X_{n-k} = 0, X_{n-k-1} = i_1 \ldots, X_{n-k-m} = i_m) = \\
&= \left(\frac{k}{k+1}\right)^{\alpha_{\text{on}}}, \quad k \geq 1
\end{aligned}
\tag{3}
$$

and

$$\mathbb{P}(X_{n+1} = 0 | X_n = 1, \ldots, X_{n-k+1} = 1, X_{n-k} = 0) =$$
$$= \mathbb{P}(X_{n+1} = 0 | X_n = 1, \ldots, X_{n-k+1} = 1, X_{n-k} = 0, X_{n-k-1} = i_1 \ldots, X_{n-k-m} = i_m) =$$
$$= 1 - \left(\frac{k}{k+1}\right)^{\alpha_{\mathrm{on}}}, \quad k \geq 1 \tag{4}$$

for any m and any $(i_1, \ldots, i_m), i_l \in \{0, 1\}, l = 1, \ldots, m$ and for $n \geq 1$. Similarly:

$$\mathbb{P}(X_{n+1} = 0 | X_n = 0, \ldots, X_{n-k+1} = 0, X_{n-k} = 1) =$$
$$= \mathbb{P}(X_{n+1} = 0 | X_n = 0, \ldots, X_{n-k+1} = 0, X_{n-k} = 1, X_{n-k-1} = i_1 \ldots, X_{n-k-m} = i_m) =$$
$$= \left(\frac{k}{k+1}\right)^{\alpha_{\mathrm{off}}}, \quad k \geq 1 \tag{5}$$

and

$$\mathbb{P}(X_{n+1} = 1 | X_n = 0, \ldots, X_{n-k+1} = 0, X_{n-k} = 1) =$$
$$= \mathbb{P}(X_{n+1} = 1 | X_n = 0, \ldots, X_{n-k+1} = 0, X_{n-k} = 1, X_{n-k-1} = i_1 \ldots, X_{n-k-m} = i_m) =$$
$$= 1 - \left(\frac{k}{k+1}\right)^{\alpha_{\mathrm{off}}}, \quad k \geq 1 \tag{6}$$

for any m and any $(i_1, \ldots, i_m), i_l \in \{0, 1\}, l = 1, \ldots m$ and for $n \geq 1$. Here, the α_{on} and α_{off} parameters take positive values in the range $1 < \alpha_{\mathrm{on}}, \alpha_{\mathrm{off}} < 2$. On the other hand, $\mathbb{P}(X_1 = 0) = 1/2$.

In the light of the previous equations, one can easily see that the two consistency conditions stated by the Kolmogorov Consistency Theorem [Billingsley 1995, Page 482] are met. Accordingly, there is a stochastic process $\{X_n, n \in \mathbb{Z}^+\}$ on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with the conditional distributions defined by eq. 3, eq. 4, eq. 5 and eq. 6.

For clarity, let us consider the following example: Let us assume that the chain is initially in the "on" state, i.e. $X_1 = 1$. Thus, according to eq. 3, the conditional probability of remaining in the "on" state in the following time-slot is given by $\mathbb{P}(X_2 = 1 | X_1 = 1) = \left(\frac{1}{2}\right)^{\alpha_{\mathrm{on}}}$. Obviously, the conditional probability of jumping to state "off" would be: $1 - \left(\frac{1}{2}\right)^{\alpha_{\mathrm{on}}}$ (see eq. 4).

Following this, let us assume the chain remains in state "on" at time $n = 2$. In this case, the conditional probability of remaining in state "on" for a further time is increased to $\mathbb{P}(X_3 = 1 | X_2 = 1, X_1 = 1) = \left(\frac{2}{3}\right)^{\alpha_{\mathrm{on}}}$. Clearly, the size of this change depends on the "on"-state tail index α_{on}. Hence, the conditional probability of jumping to the "off" state has decreased to $\mathbb{P}(X_3 = 0 | X_2 = 1, X_1 = 1) = 1 - \left(\frac{2}{3}\right)^{\alpha_{\mathrm{on}}}$ (see eq. 4), and so forth.

On the other hand, let us assume the chain eventually jumps to the "off" state at time $n = k$. In this case, the conditional probability of remaining in the "off" state at time $n = k + 1$ would be $\mathbb{P}(X_{k+1} = 0 | X_k = 0, X_{k-1} = 1, \ldots) = \mathbb{P}(X_{k+1} = 0 | X_k = 0) = \left(\frac{1}{2}\right)^{\alpha_{\mathrm{off}}}$. Again, this conditional probability would increase to $\left(\frac{2}{3}\right)^{\alpha_{\mathrm{off}}}$

if the system jumps again to the "off" state, and so forth.

As shown in this example, the conditional probability of the process remaining in the same state increases the longer the chain has already stayed in it, since:

$$\left(\frac{k}{k+1}\right)^{\alpha} > \left(\frac{k-1}{k}\right)^{\alpha}, \qquad \alpha > 0$$

Such state persistence effect is stronger the smaller the value of the tail index α, which controls the speed at which the conditional probability approaches one. Figure 1 shows how the conditional probability of remaining in the same state increases with history, for several values of α, with $1 < \alpha < 2$.
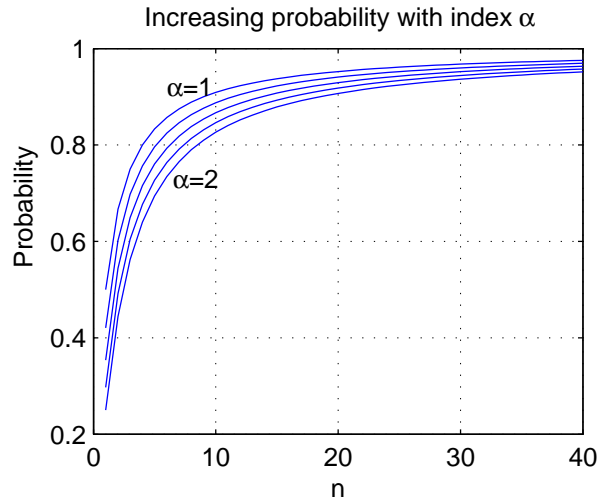

Fig. 1. Conditional probability of remaining in the same state ("on" or "off") during n timeslots for various tail-indices α.

Clearly, this property encourages the creation of bursts, and agrees with the empirical fact that packets are typically transmitted back-to-back [Jain and Routhier 1986].

The following sections discusses the probability distribution of such bursts and other related properties of this model.

4. MODEL PROPERTIES

Let us assume that $X_1 = 1$, that is, the chain is originaly in the active state, and let $U_{1,\text{on}}$ denote the number of time-slots spent in the "on" state until the first jump to the "off" state occurs. Similarly, let $U_{1,\text{off}}$ denote the number of time-slots spent in the "off" state. In other words, $U_{i,\text{on}}$ and $U_{i,\text{off}}$ denote the durations in time-slots of the i-th sequence $(i = 1, 2, \ldots)$ of "on" and "off" periods respectively.

Remark 4.1 *Mutual independence of "on" and "off" durations.* The random variables $U_{i,\text{on}}$ and $U_{i,\text{off}}$ are mutually independent for $i = 1, 2, \ldots$. In fact, let $d, k \in \mathbb{Z}^+$, then:

Table I.	A few values of $\zeta(\alpha)$	
$\zeta(1.00)$	$=$	∞
$\zeta(1.10)$	$\approx$	10.58
$\zeta(1.15)$	$\approx$	7.25
$\zeta(1.25)$	$\approx$	4.60
$\zeta(1.40)$	$\approx$	3.11
$\zeta(1.50)$	$\approx$	2.61
$\zeta(1.65)$	$\approx$	2.16
$\zeta(1.75)$	$\approx$	1.96
$\zeta(1.90)$	$\approx$	1.75
$\zeta(2.00)$	$\approx$	1.65

$$\mathbb{P}(U_{i,\text{on}} = d | U_{i,\text{off}} = k) = \mathbb{P}(U_{i,\text{on}} = d | U_{i,\text{off}} \geq 1) = \mathbb{P}(U_{i,\text{on}} = d) \tag{7}$$

by using eq. 3 and eq. 5 and taking into account that $\mathbb{P}(U_{i,\text{off}} \geq 1) = 1$. Note that both eq. 3 and eq. 5 indicate that the duration of any actual "on" or "off" period only depends on the time elapsed from the last change in the chain.

THEOREM 4.2 DISTRIBUTION OF "ON" AND "OFF" DURATIONS. *The probability distribution of the random variables $U_{i,on}$ and $U_{i,off}$, $i = 1, 2, \ldots$ is Pareto-like with tail indices α_{on} and α_{off} respectively. That is,*

$$\mathbb{P}(U_{i,on} < d) = 1 - \left(\frac{1}{d}\right)^{\alpha_{on}} \text{ and } \qquad \mathbb{P}(U_{i,off} < d) = 1 - \left(\frac{1}{d}\right)^{\alpha_{off}} \tag{8}$$

$$\tag{9}$$

for $d = 1, 2, \ldots$.

PROOF. *Without loss of generality let us assume that $X_1 = 1$. Thus,*

$$\mathbb{P}(U_{1,on} \geq d) = \left(\frac{d-1}{d}\right)^{\alpha_{on}} \left(\frac{d-2}{d-1}\right)^{\alpha_{on}} \cdots \left(\frac{1}{2}\right)^{\alpha_{on}} = \frac{1}{d^{\alpha_{on}}} \tag{10}$$

for $d = 1, 2, \ldots$ and, similarly, $\mathbb{P}(U_{i,off} \geq d) = 1/d^{\alpha_{off}}$, which completes the proof of the theorem. □

COROLLARY 4.3. *The expectation of the random variables $U_{i,on}$ and $U_{i,off}$, $i = 1, 2, \ldots$ is the Riemann function evaluated at the tail-indices α_{on} and α_{off} respectively. That is,*

$$E[U_{i,on}] = \zeta(\alpha_{on}) = \sum_{d=1}^{\infty} \frac{1}{d^{\alpha_{on}}} \tag{11}$$

and $E[U_{i,off}] = \zeta(\alpha_{off})$. Note that $U_{i,on}$ and $U_{i,off}$ are both positive a. s. and, thus, $E[U_{i,on}] = \sum_{d=1}^{\infty} \mathbb{P}(U_{1,on} \geq d)$ and $E[U_{i,off}] = \sum_{d=1}^{\infty} \mathbb{P}(U_{1,off} \geq d)$

Obviously, the burst length variance is infinity.

The Riemann function $\zeta(x)$ has no closed form. A few example values are shown in table I.

5. LONG-RANGE DEPENDENCE

The previous theorems indicate that the discrete-time heavy-tailed chain is a special case of the "on/off" process considered in [Taqqu et al. 1997]. In fact, the model in [Taqqu et al. 1997] encompasses any on/off process with independent "on" and "off" durations, such that, when $x \to \infty$, $\mathbb{P}(U_{i,\mathrm{on}} > d) \sim l_{\mathrm{on}} x^{-\alpha_{\mathrm{on}}} L_{\mathrm{on}}(x)$ and $\mathbb{P}(U_{i,\mathrm{off}} > d) \sim l_{\mathrm{off}} x^{-\alpha_{\mathrm{off}}} L_{\mathrm{off}}(x)$ for $d > 0$ where $l_{\mathrm{on}} > 0$, $l_{\mathrm{off}} > 0$ are constants and $L_{\mathrm{on}}(x) > 0$, $L_{\mathrm{on}}(x) > 0$ are *slowly varying* functions, namely $\lim_{x \to \infty} L(tx)/L(x) = 1$ (see [Resnick 1997] for further details concerning heavy tails). It is worth noticing that the discrete-time heavy-tailed chains meet these two requirements, as can be readily seen from eq. 8. According to this, it can be concluded that the *discrete-time heavy-tailed chains* basically provide a *discrete-time* version of the "on/off" model considered in [Taqqu et al. 1997].

Indeed, the superposition of an infinite number of "on/off" heavy-tailed sources leads to a random process, say $Z = \{Z_n, n = 1, 2, \ldots\}$, which shows *long-range dependence*. More precisely, there exist a real number $\beta \in (0, 1)$ and a constant value c_ρ such that the process' autocorrelation function $\rho(k)$, $k = 1, 2, \ldots$ fulfills:

$$\lim_{k \to \infty} \frac{\rho(k)}{c_\rho k^{-\beta}} = 1 \qquad (12)$$

as pointed out in [Beran 1994, Definition 2.1]. As a consequence, the variance of the aggregated process: $\mathrm{Var}(\sum_{i=1}^{n} Z_i) \approx K n^{2H}$ as $n \to \infty$, where $K > 0$ is a constant and H is the *Hurst parameter*, with $1/2 < H < 1$, according to [Beran 1994, Theorem 2.2]. Note that the process Z shows *slowly-decaying* autocorrelation and variance of the aggregated process, $\sum_{i=1}^{n} Z_i$. Intuitively, such slowly-decaying behaviour implies that the traffic burstiness smooths out very slowly with the aggregation level, in contrast to processes with independent increments.

Moreover, the process Z, made up by aggregating a large number of such "on/off" sources converges to *Fractional Gaussian Noise* [Taqqu et al. 1997], a long-range dependent process with Gaussian marginal distribution, very often used in the so-called "black-box" modelling of network traffic [Norros 1995]. Finally, it is worth pointing out that the Hurst parameter is related to the tail-indices α_{on} and α_{off} as follows:

$$H = \frac{3 - \min(\alpha_{\mathrm{on}}, \alpha_{\mathrm{off}})}{2} \qquad (13)$$

The reader is referred to [Beran 1994] and to [Resnick 1997] for a detailed study of long-range dependent models and heavy tail modelling. The following section shows how to obtain the tail-indices of the "on" and "off" states, based on measured data.

6. INFERENCE OF THE TAIL-INDEX α

Clearly from the previous sections, the estimation of each state's tail-indices α_{on} and α_{off} is key in determining essential aspects of the model, such as the burst length and their expectation. Taking into account the fact that both the "on" and the "off" sequences are heavy-tailed, the Hill's estimator provides a good way to

estimate the chain's index parameters α_{on} and α_{off}. Following Hill [1975], the Hill's estimator proceeds as follows:

Remark 6.1 *Hill's estimator.* Let $X_{i,\text{on}}$ and $X_{i,\text{off}}$ $i = 1, \ldots, n$ refer to the i-th sample obtained from $U_{i,\text{on}}$ and $U_{i,\text{off}}$ respectively, that is, samples of sequences of "on" and "off" periods. Let $X_{(1),\text{on}} \geq X_{(2),\text{on}} \geq \ldots X_{(k+1),\text{on}}$ denote the $k + 1 < n$ largest "on" sequences, sorted from largest to smallest, and let $X_{(1),\text{off}} \geq X_{(2),\text{off}} \geq \ldots X_{(k+1),\text{off}}$ refer to the $k + 1 < n$ largest "off" sequences, again sorted from largest to smallest. Then the Hill's estimator for α_{on} and α_{off}, based on such $k < n$ subsample, is given by:

$$\hat{\alpha}_{\text{on},k} = \left(\frac{1}{k+1} \sum_{i=1}^{k} \log \frac{X_{(i),\text{on}}}{X_{(k+1),\text{on}}} \right)^{-1} \tag{14}$$

$$\hat{\alpha}_{\text{off},k} = \left(\frac{1}{k+1} \sum_{i=1}^{k} \log \frac{X_{(i),\text{off}}}{X_{(k+1),\text{off}}} \right)^{-1} \tag{15}$$

When k is significantly large, the estimate $\hat{\alpha}_k$ approaches the real index α. Figure 2 shows the performance of the Hill's estimator for a discrete-time heavy-tailed chain with tail indices: $\alpha_{\text{on}} = 1.5$ and $\alpha_{\text{off}} = 1.75$. The chain was simulated with $T = 10^7$ states. Out of them, only the $k = 1000$ largest sequences of active and idle periods have been considered to estimate α_{on} and α_{off}, as shown.
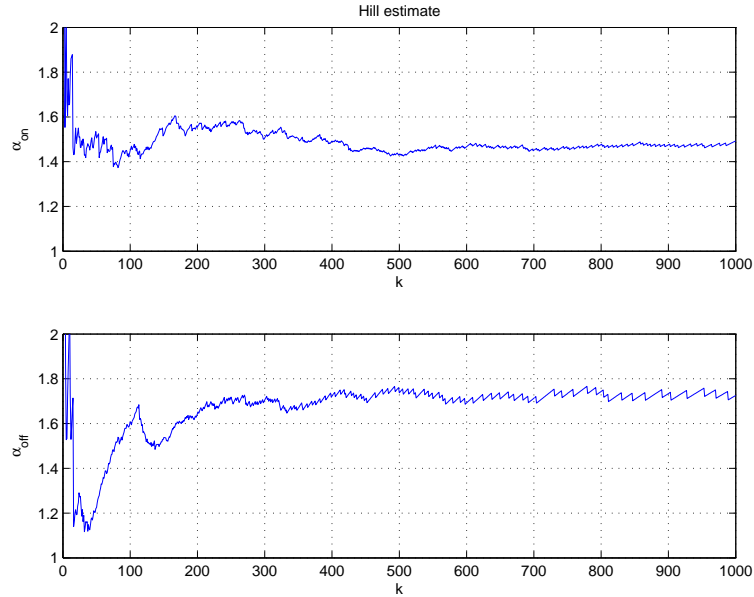

Fig. 2. Example of the performance of the Hill's estimator for a discrete-time heavy-tailed chain with $\alpha_{\text{on}} = 1.5$ and $\alpha_{\text{off}} = 1.75$.

In conclusion, the Hill's estimator is a very good procedure to obtain the discrete-time heavy-tailed chain main characteristic parameters: α_{on} and α_{off}.

7. NUMBER OF ACTIVE SOURCES

This section derives the probability distribution of the number of active sources, say in a router's output link, at a random discrete-time. Clearly, this value provides a measure of the router's load as seen by a new source arriving randomly (i.e. a Poisson arrival), and is key in the design of network resources.

Remark 7.1 *Number of active sources in a given random time.* Let N refer to the number of connections handled by a router simultaneously, that is, discrete-time "on/off" heavy-tailed sources. The number n of flows in the "on" state at a given random time is given by the following binomial distribution:

$$\mathbb{P}(n) = \binom{N}{n} p_{\text{on}}^{n} (1 - p_{\text{on}})^{N-n} \tag{16}$$

characterised by the parameter p_{on}, which stands for the probability of a single source to be "on" at a random time. Such value is given by:

$$p_{\text{on}} = \frac{\zeta(\alpha_{\text{on}})}{\zeta(\alpha_{\text{on}}) + \zeta(\alpha_{\text{off}})} \tag{17}$$

Note that the on/off process induced by eq. 3 and eq. 5 is renewal, and, then, using [Ross 1996, Theorem 3.4.4], eq. 17 can be shown to hold in the long run. Thus, the validity of eq. 16 follows.

By the *De Moivre-Laplace theorem*, when the number N of connections is large, eq. 16 approaches a gaussian distribution $\mathcal{N}(\mu, \sigma)$ with mean $\mu = Np_{\text{on}}$ and standard deviation $\sigma = \sqrt{Np_{\text{on}}(1 - p_{\text{on}})}$. This result agrees with the fact that the superposition of infinite "on/off" sources leads to a Fraction Gaussian Noise, which is a Gaussian process [Taqqu et al. 1997].

Figure 3 shows a case example of the validity of the remark above. In this example, a router is handling $N = 100$ heavy-tailed connections with tail indices $\alpha_{\text{on}} = 1.5$ and $\alpha_{\text{off}} = 1.75$ for a number of $T = 10^{7}$ discrete time-slots. Out of them, we have chosen 10000 time-slots randomly, accounted the number of active sources in each time-slot, and plotted its histogram, along with the theoretical binomial distribution given by eq. 16. As expected, the simulation-based histogram accurately matches the theoretical probability density function.

In this example, the simulated results shows that $p_{\text{on}} = 0.5710$, thus giving the values of $\mu = 57.10$ and $\sigma = 4.95$ for the Gaussian approximation.

8. SUMMARY

This work has introduced the *discrete-time heavy-tailed chains*, a practical methodology that can be used to simulate discrete-time "on/off" sources with heavy-tailed distributed sequences of the "on" and the "off" periods, in the light of the "on/off" source model described by Willinger et al. [1997]. The literature and previous work demonstrates that the aggregation of multiple sources of such characteristics con-
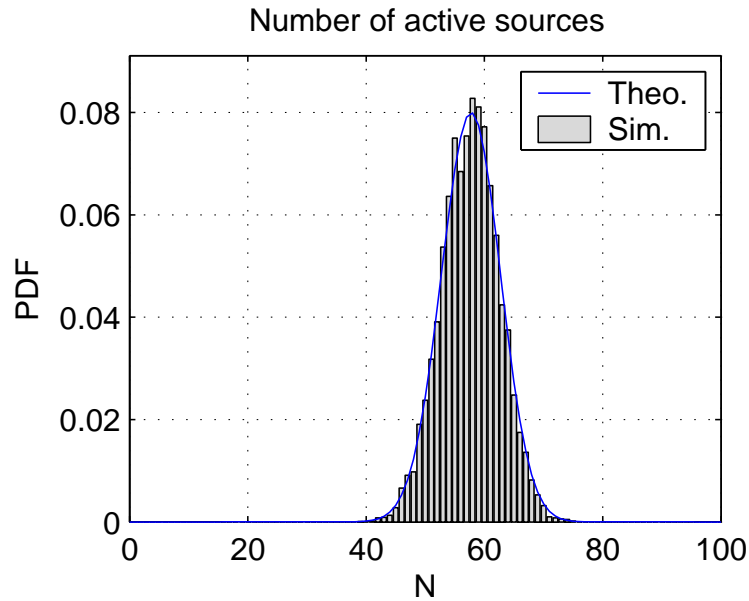
Fig. 3. Distribution of the number of active sources in a random discrete-time. Total number of sources: $N = 100$, and $\alpha_{on} = 1.5$ and $\alpha_{off} = 1.75$

verges to Fractional Gaussian Noise, a Gaussian long-range dependent process, very often used in the simulation and modelling of network traffic.

Indeed, the key rule for constructing the heavy-tailed chains proposed in this work shows how it promotes the creation of bursts, i.e. sequences of consecutive "on" periods with infinite variance, which causes the long-range dependence observed in the aggregated traffic. The average length of such bursts depends on the α parameter that characterises the chain, which can be easily estimated via the Hill estimate. Also, it is shown how to infer the number of active sources in a given random time, if the total number of sources are known and so are the tail indices of the "on" and "off" periods.

Finally, it is worth remarking that the *discrete-time heavy-tailed chains* can be applied to the performance study and evaluation of a number of discrete-time communication systems, say ATM and Optical Packet Switching, for instance.

REFERENCES

ADAS, A. 1997. Traffic models in broadband networks. *IEEE Communication Magazine*, 82–89.

BERAN, J. 1994. *Statistics for long-memory processes*. Chapman and Hall.

BERAN, J., SHERMAN, R., TAQQU, M. S., AND WILLINGER, W. 1995. Long-rande dependence in variable-bit-rate video traffic. *IEEE/ACM Trans. Communications 43,* Feb-Mar-April, 1566–1597.

BILLINGSLEY, P. 1995. *Probability and measure*, 3 ed. Wiley.

CROVELLA, M. E. AND BESTAVROS, A. 1997. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Trans. Networking 5,* 6, 835–846.

FELDMANN, A., GILBERT, A. C., AND WILLINGER, W. 1998. Data networks as cascades: investigating the multifractal nature of Internet WAN traffic. *ACM SIGCOMM Computer Communication Review 28,* 4 (October), 42–55.

FELDMANN, A., GILBERT, A. C., WILLINGER, W., AND KURTZ, T. G. 1998. The changing nature of network traffic: Scaling phenomena. *Computer Communication Review 28,* 2, 5–29.

FLOYD, S. AND KOHLER, E. 2003. Internet research needs better models. *ACM Computer Communication Review 33,* 1, 29–34.

FROST, V. AND MELAMED, B. 1994. Traffic modeling for telecommunication networks. *IEEE Communication Magazine*, 70–81.

GUILLEMOT C. ET AL. 1998. Transparent optical packet switching: The european ACTS KEOPS project approach. *IEEE Journal on Lightwave Technology 16,* 12 (December).

HILL, B. M. 1975. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics 3,* 1163–1174.

JAIN, R. AND ROUTHIER, S. 1986. Packet trains – Measurements and a new model for computer network traffic. *IEEE JSAC SAC-4,* 6, 986–995.

KARAGIANNIS, T., FALOUTSOS, M., AND RIEDI, R. 2002. Long-Range Dependence: Now you see it now you don't! In *Global Internet Symposium (in IEEE GLOBECOM)*. Taipei, Taiwan.

LELAND, W. E., TAQQU, M. S., WILLINGER, W., AND WILSON, D. V. 1994. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Trans. Networking 2,* 1, 103–115.

LÓPEZ-ARDAO, J. C., LÓPEZ-GARCÍA, C., SUÁREZ-GONZÁLEZ, A., FERNÁNDEZ-VEIGA, M., AND RODRÍGUEZ-RUBIO, R. 2000. On the use of self-similar processes in network simulation. *ACM Trans. on Modelling and Computer Simulation 10,* 2 (April), 125–151.

NORROS, I. 1995. On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE Journal on Selected Areas in Communications 13,* 6, 953–962.

O'MAHONY M., SIMEONIDOU D., HUNTER D. AND TZANAKAKI A. 2001. The application of optical packet switching in future communication networks. *IEEE Communications Magazine 39,* 3 (March).

PAXSON, V. AND FLOYD, S. 1995. Wide Area traffic: the failure of Poisson modeling. *IEEE/ACM Trans. Networking 3,* 3, 226–244.

RESNICK, S. I. 1997. Heavy tail modeling and teletraffic data. *The Annals of Statistics 25,* 5, 1805–1869.

ROSS, S. 1996. *Stochastic Processes*, 2 ed. Wiley.

TAQQU, M. S., WILLINGER, W., AND SHERMAN, R. 1997. Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review 27,* 2, 5–23.

WILLINGER, W. AND PAXSON, V. 1998. Where mathematics meets the Internet. *Notices of the American Mathematical Society 45,* 8 (August), 961–970.

WILLINGER, W., TAQQU, M. S., SHERMAN, R., AND WILSON, D. V. 1997. Self-similarity through high variability: Statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Trans. Networking 5,* 1, 75–86.