Tail Probabilities for the Multiplexing of Fractional α-Stable Broadband Traffic

Fotios C. Harmantzis, Dimitrios Hatzinakos, Irene Katzela

Abstract— **In this paper, we investigate the tail probabilities of a multiplexer driven by α-Stable self-similar traffic. We consider a parsimonious 4-parameter traffic model, which best captures the long range dependence and heavy-tails of aggregate packet traffic in broadband networks. Input traffic with these characteristics, induces buffer dynamics which are qualitatively different from those which arise in traditional traffic management. Using the effective bandwidths theory, we extend the recent results on α-Stable self-similar-driven queues with infinite buffer to the finite buffer case that model routers/switches more accurately. Queuing simulations with real broadband traffic emphasise the improvements in engineering considerations (e.g., connection admission control, buffer management, statistical multiplexing gains), with respect to the existing results so far. Our experiments involve a large set of real broadband network traffic which consists of Ethernet LAN, Internet WAN and MPEG-1 compressed video traces.**

Keywords— **Broadband traffic modeling, self-similarity, heavy tails, queuing analysis, multiplexing, effective bandwidths.**

I. Introduction

THE term "self-similar" was first used by Mandelbrot [19] during the 1960's to designate those processes that are scalable over time (or space) without losing their statistical properties. In other words, a continuous-time process $X = \{X(t), t \geq 0\}$ is *self-similar* (H*-ss*), with self-similarity parameter H, if it satisfies the condition:

$$X(t) \stackrel{d}{=} c^{-H} X(ct),\ \forall t \geq 0,\ \forall c > 0,\ 0 < H < 1, \tag{1}$$

where the equality is in the sense of finite-dimensional distributions. There are many different self-similar processes [25]. We typically consider those that have stationary increments, and call them H*-sssi* processes. The *Hurst parameter* H is the scaling parameter of self-similarity. Self-similarity manifests itself in a number of equivalent ways, the most relevant of which is the display of *long range dependence* (LRD) or *long memory* for H*-sssi* processes. Mandelbrot refers to this phenomenon as the *Joseph effect*.

In the 1990's extensive measurable studies from a wide range of data networks and services/applications, have convincingly demonstrated the self-similarity or fractal nature of data traffic [5], [10], [17], [23], [24]. As a consequence, a large number of traffic models have been proposed in order to successfully characterize the real statistical behaviour of the traffic met in networks today. The reason for that is that self-similarity has serious implications for the analysis, design, and control of broadband networks, where traditional schemes, typically Markovian in nature, are currently used.

One of the best known stochastic models that is able to account LRD if $1/2 < H < 1$, in a flexible and parsimonious manner, is the *Fractonal Brownian Motion* (FBM) model, initially introduced by Norros [22], and lately proposed for modeling of aggregate packet network traffic, e.g. ATM, Frame-Relay, Ethernet [10]. The process is Gaussian, and basically is the single Gaussian H-sssi process, for a given $H \in (0, 1)$.

Although Gaussian processes can display long-range dependence, their marginal distributions, being normal, concentrate their mass around the mean. α-Stable distributions with $0 < \alpha < 2$, on the other hand allow for much greater variability (burstiness) [25]. Mandelbrot refers to the variability of stable random variables as the *Noah effect*. By considering in this paper α-Stable self-similar processes with $0 < \alpha < 2$, we are, in fact, dealing with traffic models that can exhibit both the Joseph and the Noah effects.

The model we consider in this paper, is based on the *Linear Fractional Stable Motion* (LFSM) [25], an α-stable H-sssi process with $0 < \alpha < 2$, capable of capturing both the LRD and the heavy-tails effect. Actually, in modeling, it is the increments of H-sssi processes that are of interest; so we use the stationary sequence induced by the LFSM process, called the *linear fractional stable noise*. The model, referred to as the *Skewed Stable and Self-Similar (S4) Model*, is introduced in [11] for heavy-traffic modeling and it is validated against real Ethernet, VBR video, and WWW traffic. The S4 Model is a generalization of the FBM model, since it encompasses Gaussianity. In [9], the authors highlight the higher accuracy of the S4 Model in capturing the queuing performance of real traffic compared to the FBM model.

Whereas FBM is the only H-sssi process with $0 < H < 1$, there are many different H-sssi processes when $0 < \alpha < 2$. Consequently, after [14], similar models based on other α-stable H-sssi processes, e.g., α-stable Levy motion, log-fractional stable motion, have tried to address the problem of self-similarity and heavy-tails at the same time [2], [8], [16].

This established presence of self-similarity and heavy-tails over a wide range of time scales in packet traffic processes is expected to have an impact on queuing performance and buffer engineering. In this paper, we seek to gain some insights into these issues by investigating the effects of LRD and heavy-tails on the behaviour of the buffer queue at a multiplexer. More precisely, we consider a single server queue with infinite capacity and constant rate C bits/slot, as a surrogate for a multiplexer, and feed it with a α-stable self-similar input, generated by the S4 Model. It has been reported in [12], [8] that LRD and heavy-tails in the input traffic will indeed induce buffer dynamics which are qualitatively very different from those which arise in Markovian models and in the Gaussian case. If $Q(t)$ denotes the buffer content (in number of bits) at the multiplexer, we are interested in tail probabilities $P[Q(t) > B]$ as a means for estimating the buffer overflow probabilities (BOP) for the corresponding finite

F. C. Harmantzis and D. Hatzinakos are with the Communications Group, The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada. E-mail: fotios, dimitris@comm.toronto.edu.

I.Katzela is with Lucent Technologies, Canada.

buffer system of size B. Such asymptotics are often the first guiding step to size up the buffer at the multiplexer in order to guarantee quality of service (QoS) requirements.

In this paper we review the *large buffer asymptotic* obtained for the continuous-time storage model driven by LFSM [13]. As far as we know this is the only asymptotic result reported in the literature. We use $A_i(t)$ to model a stochastic flow as a LFSM. Suppose the aggregate $A(t)$ of N independent copies A_i feeds the buffer; we show, using the *many sources asymptotic* [6], that the tail probability for $Q(t)$ for fixed t, obeys a better asymptotic than the large buffer one. This aims to develop connection admission control (CAC) algorithms which can make better use of systems resources, estimating for example the refined notion of *effective bandwidths* [15].

The paper is organized as follows: In Sec. 2 we review basic properties of the α-Stable distributions, the LFSM process, and we describe the S4 Model. In Sec. 3 we describe the buffer overflow asymptotics obtained in a multiplexer driven by LFSM traffic. Then, we apply the many sources asymptotic, estimating the effective bandwidth of the source directly from the traffic trace. In Sec. 4 we present simulation experiments and numerical results, and provide comparisons to the queuing behaviour of real broadband traffic. In the Conclusions we summarize the main observations and refer to future work.

II. The Model and Preliminaries

A. α-Stable Distributions

The most convenient way to describe a stable distribution is by its characteristic function, because few stable density functions are known in closed form. This lack of closed formulas has been a major drawback to the use of stable distributions by practitioners.

α-Stable distributions have been proposed as a model for many types of physical and economic systems [1]. There are three main reasons for using a stable distribution to describe a system. Firstly, there are solid theoretical reasons for expecting a non-Gaussian stable model. Secondly, the Generalized Central Limit Theorem, which states that the only possible non-trivial limit of normalized sums of independent identically distributed terms is stable. Finally, empirical observations show that many large data sets exhibit heavy tails and skewness. Such data sets are poorly described by a Gaussian model, but can be well described by a stable distribution.

A stable distribution is characterized by four parameters: The *characteristic exponent*, or *index of stability*, $\alpha \in (0,2]$; the *scale*, or *spread parameter*, $\sigma \geq 0$; the *skewness*, or *symmetry parameter*, $\beta \in [-1,1]$; and the *shift*, or *location parameter*, $\mu \in \Re$. A random variable X is said to have a stable distribution if and only if its characteristic function has the form:

$$E(exp^{iXt}) = \begin{cases} exp\{-\sigma^\alpha |t|^\alpha (1 - i\beta \text{sign}(t) \tan\frac{\pi\alpha}{2}) + i\mu t\} & ; \quad \alpha \neq 1 \\ exp\{-\sigma|t|(1 + \frac{2i\beta}{\pi}\text{sign}(t) ln|t|) + i\mu t\} & ; \quad \alpha = 1 \end{cases}$$

We shall call such a distribution α-Stable and denote it by writing $X \sim S_\alpha(\sigma,\beta,\mu)$. The characteristic exponent α determines the rate of decay, i.e., the heaviness of the tails of the distribution, while the parameter β is an indication of the skewness of the distribution, with $\beta = 0$ corresponding to the symmetric case; α and β together determine the shape of the distribution. The parameter μ shifts the distribution to the left or right, whereas the parameter σ merely expands or contracts it around μ; they both have no effect on its shape. We may therefore represent the *standard* stable distribution with $\sigma = 1$ and $\mu = 0$, by $X \sim S_\alpha(1,\beta,0)$. We also write $X \sim S\alpha S$ when X is *symmetric* α-stable, i.e., when $\beta = 0$. Finally, $X \sim S_\alpha(\sigma,1,\mu)$ when X is a *totally (positively) skewed* α-Stable.

The tail probabilities in the non-Gaussian cases ($\alpha < 2$) are known asymptotically. The following property states for the tail approximation of an α-stable distribution:

Tail Approximation [25]. Let $X \sim S_\alpha(\sigma,\beta,\mu)$ with $0 < \alpha < 2$. Then, as $x \to \infty$

$$P(X > x) \sim \sigma^\alpha \frac{1+\beta}{2} C_\alpha x^{-\alpha}, \tag{2}$$

where $C_\alpha = \frac{1-\alpha}{\Gamma(2-\alpha)cos(\pi\alpha/2)}$, if $\alpha \neq 1$. Thus, stable laws have heavy tails; in contrast, normal distributions have exponential tails. This implies that the tails of stable laws are significantly thicker than those of the normal distribution; in fact the smaller the value of α, the thicker the tails.

Stability [25]: Let X_1 and X_2 be independent random variables, with $X_i \sim S_\alpha(\sigma_i,\beta_i,\mu_i)$, $i = 1,2$. Then $X_1 + X_2 \sim S_\alpha(\sigma,\beta,\mu)$, with

$$\sigma = (\sigma_1^\alpha + \sigma_2^\alpha)^{1/\alpha},\ \beta = \frac{\beta_1\sigma_1^\alpha + \beta_2\sigma_2^\alpha}{\sigma_1^\alpha + \sigma_2^\alpha},\ \mu = \mu_1 + \mu_2. \tag{3}$$

B. The LFSM Process

There are different extensions of FBM to the α-stable case; the one which is most commonly used is the *Linear Fractional Stable Motion* (LFSM). The *well-balanced LFSM* is a continuous-time stochastic process $\{L_{\alpha,H}, t \geq 0\}$ defined as follows [25]

$$L_{S,H}(t) = \int_{-\infty}^{\infty} (|t-x|^{H-1/\alpha} - |x|^{H-1/\alpha}) M_S(dx) \tag{4}$$

where $0 < \alpha < 2$, $0 < H < 1$, $H \neq 1/\alpha$, and M_S is an α-stable random measure on $\Re$ with Lebesque control measure. The LFSM is a H-sssi process, and if one sets $\alpha = 2$, LFSM reduces to FBM.

Consider the increments $\{N_{S,H}(k) = L_{S,H}(k+1) - L_{S,H}(k),\ k = \ldots,-1,0,1,\ldots\}$ of the H-sssi LFSM process introduced by (4). The stationary sequence (defined in discrete time now) $\{N_{S,H}(k), k \geq 0\}$ is the *Linear Fractional Stable Noise* (LFSN)

$$N_{S,H}(k) = \int_{-\infty}^{\infty} (|k+1-x|^{H-1/\alpha} - |k-x|^{H-1/\alpha}) M_S(dx) \tag{5}$$

We say that the increments of LFSM have *long-range dependence* when $H > 1/\alpha$ and *short-range dependence* when $H < 1/\alpha$ [25]. There is no LRD when $0 < \alpha \leq 1$ because H is constrained to lie in the interval $(0, 1)$. The sequence $\{N_{S,H}(k), k \geq 0\}$ is of a particular interest because it displays both the "Joseph" and the "Noah" effects.

C. The S4 Model

[11] proposes a parsimonious model, called the *Skewed Stable and Self-Similar (S4) Model*, for network heavy-traffic approximation, based on the LFSN. We intend to adopt it for the purposes of our research, and investigate its queuing performance, extending the work reported in [13].

The S4 Model is used to model a heavy-traffic packet-count process, it is linear and is defined by

$$M(i) = c_1 \cdot N_{S,H}(i) + m, \; i = 0, 1, \ldots, \tag{6}$$

where $M(i)$ denotes the volume of traffic carried by the network element in the time unit i, c_1 and m are positive real constants and $N_{\alpha,H}$ is the well-balanced LFSN defined in (5) with $\beta = 1$, $\sigma = 1$, $\mu = 0$, and $H > 1/\alpha$ to ensure LRD, i.e., $N_{\alpha,H} \sim S_\alpha(1, 1, 0)$ is a *totally skewed standard* α-Stable distribution.

The model process (6) is a totally skewed α-Stable process with the same parameter α as that of the LFSN used to construct it (see property 1.2.3 p.11 in [25]).

The four parameters α, H, c_1 and m have direct physical meaning:

- $\alpha \in (1, 2]$ is the index of stability
- $H \in (0, 1)$ is the Hurst parameter (index of self-similarity)
- $c_1 > 0$ is the dispersion around the mean of the traffic
- $m > 0$ is the mean of the traffic

Details for the parameter estimation can be found in [12].

III. Asymptotics of the BOP for Aggregate Homogeneous Traffic

Suppose that the arrival process $A(t)$, i.e., the amount of total load in the time interval $[0, t), t \in (-\infty, \infty)$, at a multiplexer is the superposition of N independent traffic flows with the same set of parameters α, H, c_1 and m (homogeneous traffic sources). Their distributions are according to

$$A_i(t) = m \cdot t + c_1 \cdot L^i_{\alpha,H}(t), t \in \Re$$

i.e., $A(t) = \sum_{i=1}^{N} A_i(t)$. The link has a shared buffer of size B, capacity C and utilization $\rho = Nm/C$.

A. The Large Buffer Asymptotic

We are interested in the residual of the queuing distribution at the multiplexer, whereas the arrival process is a superposed α-Stable self-similar one. Following [13]

$$\begin{aligned} P(Q(t) > B) &= P(\sup_{t \geq 0}(A(t) - Ct) > B) \\ &\cong \sup_{t \geq 0} P(L_{\alpha,H}(1) > \frac{B - (Nm - C)t}{c_1 t^H}), \end{aligned}$$

using (1), i.e., the fact that $L_{\alpha,H}(t)$ is an H-sssi process. The optimum point in the above inequality is obtained at

$$t = \frac{HB}{(1 - H)(C - Nm)} \tag{7}$$

(please note that (7) is independent of α, e.g., it is the same for an FBM-driven queue).

But, if $\{X(t), t \in \Re\}$ is LFSM, then for each fixed $t \in \Re$, $X(t)$ has an α-Stable distribution [25]. In our case, $L_i(t)$ are independent well-balanced LFSM; for $t = 1$, they are independent α-Stable variables, $L^i_{\alpha,H}(1) \sim S_\alpha(\sigma_1, \beta_1, 0)$, where the parameters σ_1 and β_1 are given by the following expressions [25]:

$${\sigma_1}^\alpha = 2\int_0^\infty ((1 - x)^d - x^d)^\alpha dx + \int_0^1 |(1 - x)^d - x^d|^\alpha dx$$

and

$$\beta_1 = \int_0^1 ((1 - x)^d - x^d)^{<\alpha>} dx$$

where $d = H - 1/\alpha$ and $x^{<a>} = sign(x)|x|^a$ is the signed power.

Now, $L_{\alpha,H}(1) = \sum_{i=1}^{N} L^i_{\alpha,H}(1)$, and using the stability property (3), $L_{\alpha,H}(1) \sim S_\alpha(N^{1/\alpha}\sigma_1, \beta_1, 0)$.

Finally, using (2) to approximate the tails of the α-Stable distribution for $1 < \alpha < 2$, we obtain the following asymptotic

$$P(Q(t) > B) \cong N \cdot K_N \cdot B^{-\alpha(1-H)}, \;\; as \; B \to \infty, \tag{8}$$

where $K_N = C_a \frac{1+\beta_1}{2}(\sigma_1 c_1 (1 - H)^\alpha (\frac{(C-Nm)(1-H)}{H})^{-\alpha H}$.

While the BOP for SRD traffic presents an exponential decay relative to the buffer size, such probability has a Weibullian tail for FBM [22]; here, the decay is at most algebraic, as shown in (8), which makes the discrepancy between that result and the previous ones even more severe (in [12] the reader can find the proof as well as the relative effect of the four parameters to the asymptotic).

The above approximation is believed to be as modestly accurate as someone would expect to find with most large deviation calculations.

B. The Many Sources Asymptotic

Kelly [15] defines the *effective bandwidth* as

$$\alpha(s, t) = \frac{1}{st}\Lambda_{A(t)}(s) \;\; 0 < s, t < \infty \tag{9}$$

where $\Lambda_{A(t)}(s) = logE[exp(sA(t))]$ is the logarithmic moment generating function of the process $A(t)$. The traditional large deviations formula is [7]

$$P(Q(t) > B) \cong exp(-\inf_{t \geq 0} \Lambda^*_{A(t)}(Ct + B)) \tag{10}$$

where $\Lambda^*_{A(t)}(a) = \sup_{\theta \in \Re}[\theta a - \Lambda_{A(t)}(\theta)]$ the *Legendre transformation* of the function $\Lambda_{A(t)}(s)$. Since the moment generating function of the stable process $A(t)$, exists only for the Gaussian case (α=2) and diverges or does not exist for all other cases [25], (10) degenerates into the not very helpful $P(Q(t) > B) \cong 1$ as was verified from our numerical calculations.

The first attempt to apply the effective bandwidth theory for α-Stable traffic is reported in [3]. The authors propose to truncate the infinite integral that defines the moment generating function of the traffic process in order to obtain a finite value. However, this approach is neither mathematically rigorous nor practically useful, according to its own authors.

Since there is no closed form formula for estimating effective bandwidths of α-Stable processes indirectly, we can only estimate $\alpha(s, t)$ empirically, e.g., directly from the trace of the aggregate traffic.

Let T be the total duration of the trace. The expectation in the statistical descriptor (9) can be approximated by the empirical average, hence the effective bandwidth can be estimated directly by

$$\hat{\alpha}(s, t) = \frac{1}{st} log[\frac{1}{T/t} \sum_{i=1}^{T/t} e^{sA[(i-1)t, it]}]$$

(for simplicity we assume that T is integer multiples of t). $A[(i-1)t, it]$ is the amount of work produced in the interval $[(i-1)t, it)$. (The measuring techniques get significant attention lately in the teletraffic community. The reader should refer to [27] for a study on measuring effective bandwidths).

Having estimated the effective bandwidth of the aggregated traffic stream, we can numerically compute[1] the *many sources asymptotic* [6] and apply (10) as follows

$$P(Q(t) > B) \cong exp(\sup_{t \geq 0} \inf_{s \in \Re} [st\hat{\alpha}(s, t) - s(Ct + B)])$$

The many sources asymptotic can be improved using the Bahadur-Rao theorem [18]; the results presented in this paper are based on the improved version of the asymptotic.

IV. Numerical Examples and Simulation Experiments

Data on broadband traffic is now much more easily measured and readily available. Here we consider three traces which have received considerable attention in the literature. The traces used in the experiments are: the Bellcore Ethernet LAN trace *BC-pAug89* [17], the Lawrence Berkeley Laboratory (LBL) TCP WAN trace *LBL-TCP-3* [23] and the *race* MPEG-1 video sequence [24].

The BC-pAug89 trace contains the first million packet arrivals seen on an Ethernet at Bellcore, while the LBL-TCP-3 trace contains two hours' worth of all wide-area TCP traffic between LBL and the rest of the world (they are both available from the Internet Traffic Archive[2]). The original traces were modified by breaking them into 50 ms blocks (study of fine timescales). The MPEG-1 encoded *race* video sequences consist of 40.000 frames which is equivalent to approximately half an hour[3]. In the video trace each epoch has duration 40 ms.

[1] Use of the *msa* tool at http://www.ics.forth.gr/netgroup/msa

TABLE I
S4 Model Parameters for real traffic

Trace	α	$c1$	H	m
BC-pAug89	1.89	4572.5	.87	1.1
LBL-TCP-3	1.62	1071.6	.98	0.27
race	1.49	9789.3	.99	0.76

There are many techniques available in the literature for the parameter estimation of α-Stable distributions. The authors of [4] and [21] present some of them. In this paper we used the well-known technique by McCulloch [20] which provides estimation using predefined tables and sample quantiles. For the estimation of the fractality parameter H we applied the R/S statistic, properly modified for non-Gaussian processes, as presented in [26]. The S4 Model fitted parameters for out datasets are given in Table I.

Throughout the experiments we focus only on the tail of the queue length distribution $P[Q(t) > B]$, but other performance metrics can easily be calculated, e.g., average utilization, number of admitted source, etc. In our investigations, we consider a buffer up to 250 ms for Ethernet LAN traffic, 120 ms for Internet WAN traffic and a buffer size up to 6 ms for delay sensitive MPEG-1 video traffic. We also consider link capacities 34Mbps for the three types of traffic.

The curve labelled "Simulation" in the figures 1, 2 and 3, reports the actual fraction of time the buffer exceeds the threshold B or the equivalent delay bound. To obtain the upper curve, labelled "Large Buffer", we implemented the asymptotic of Eq.

[2] http://ita.ee.lbl.gov/
[3] The sequence is available at http://nero.informatik.uni-wuerzburg.de/MPEG/

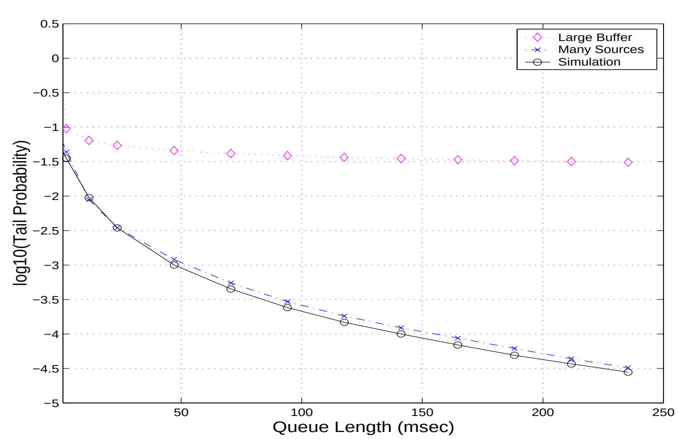

Fig. 1. Overflow probability at a link multiplexing Ethernet traffic: BC-pAug89 trace. [N=22, $\rho = 72\%$].

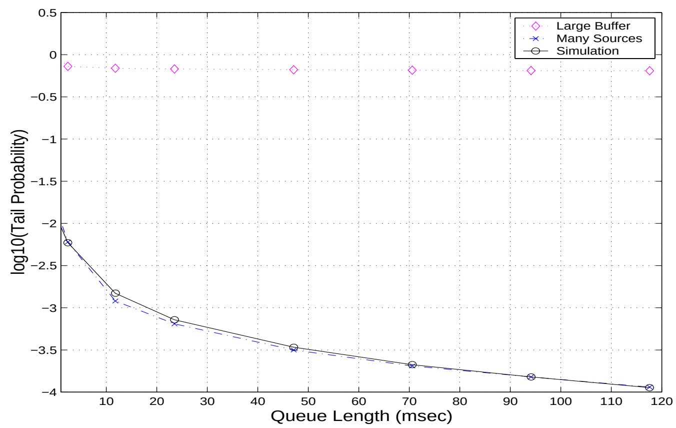

Fig. 2. Overflow probability at a link multiplexing TCP traffic: LBL-TCP-3 trace. [N=92, $\rho = 73\%$].

8. In the video traffic case, the large buffer asymptotic did not give reasonable results, e.g., $P[Q(t) > B] > 1$, for that range of buffer sizes. The reason is that the value of α is quite small, i.e. considerable mass of probability in the tails, so the factor $N \cdot K_N$ in (8) failed to compensate for the high values of $B^{-\alpha(1-H)}$. The "Many Sources" approach refines the multiplexing result, since it is based on the definition (9) of the effective bandwidths, taking into account not only the statistical characteristics of the traffic stream, but also the link's operating point. As shown in the figures, that curve follows the measured tail probability more accurately.

V. Conclusions

The large buffer asymptotic approach, whenever applicable, is significantly more conservative compared to the many sources asymptotic result, in the case of multiplexed α-Stable self-similar traffic. The later exploits the economies of scale in the number of flows, while the former determines a flow's resource demands using only the stochastic properties of the source itself. In the future we will extend our investigations in the case of heterogeneous traffic sources.

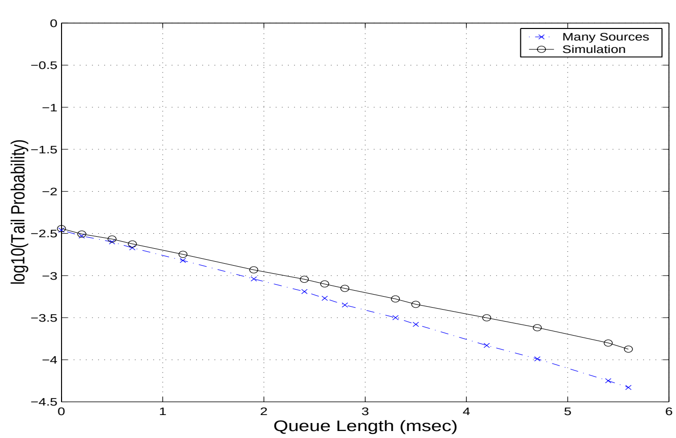

Fig. 3. Overflow probability at a link multiplexing MPEG-1 Video traffic: race sequence. [N=32, $\rho = 72\%$].

References

[1] R.J. Adler, R.E. Feldman and M.S. Taqqu. *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*. Birkhauser, Boston, 1998.

[2] S. Bates. *Traffic Characterization and Modeling for Call Admission Control Schemes on ATM networks*. PhD dissertation, University of Edinburgh, UK, 1997.

[3] S. Bates and S. McLaughlin. The Effective Bandwidth of Stable Distributions. In *Proc. ICASSP'98*, pp. 2281-2284, Seattle, WA, 1998.

[4] S. Bates and S. McLaughlin. The Estimation of Stable Distribution Parameters from Teletraffic Data. *IEEE Trans. Signal Processing*, Vol. 48, pp. 865-870, 2000.

[5] J. Beran. *Statistics for Long-Memory Processes*. Chapman & Hall, New York, 1994.

[6] C. Courcoubetis and R. Weber. Buffer Overflow Asymptotics for a Switch Handling Many traffic Sources. *J. Appl. Prob.*, Vol. 33, pp. 886-903, 1996.

[7] A. Dembo and O. Zeitouni. *Large Deviation Techniques and Applications*. Springer Verlag, 1998.

[8] J. Gallardo, D. Makrakis, and L. Orozco-Bardosa. On the Use of Alpha-Stable Self-Similar Stochastic Processes for Modeling Traffic in Broadband Networks. In *SPIE Symposium on Voice, Video and Data Communications*, pp. 281-296, Boston, MA, 1998.

[9] R.G. Garroppo, S. Giordano, S. Porcarelli and G. Procissi. Testing α-Stable Processes in Modelling Broadband Teletraffic. In *Proc. ICC'00*, pp. 1615-1619, New Orleans, LA, 2000.

[10] J.L. Jerkins and J.L. Wang. A Close Look at Traffic Measurements from Packet Networks. In *Proc. GLOBECOM'98*, pp. 2405-2411, Sydney, 1998.

[11] A. Karasaridis and D. Hatzinakos. Network Heavy Traffic Modeling using α-Stable Self-Similar Processes. *IEEE Trans. on Commun.* (accepted for publication, June 2000).

[12] A. Karasaridis. *Broadband Network Traffic Modeling, Management and Fast Simulation based on α-Stable Self-Similar Processes*. PhD dissertation, University of Toronto, 1999.

[13] A. Karasaridis and D. Hatzinakos. Resource Allocation Issues for Long-Tailed LRD Internet WAN traffic. In *Proc. GLOBECOM '99*, pp. 1495-1499, Rio de Janeiro, 1999.

[14] A. Karasaridis and D. Hatzinakos. On the Modeling of Network Traffic and Fast Simulation of Rare Events using α-Stable Self Similar Processes. In *Proc. Sign. Proc. Workshop on HOS*, Banff, Alberta, 1997.

[15] F.P. Kelly. Notes on Effective Bandwidths. In *Stochastic Networks: Theory and Applications*, F.P. Kelly, S. Zachary and I.B. Ziedins, Eds. Oxford Univ. Press, pp. 141-168, 1996.

[16] T. Konstantopoulos and S.-J. Lin. Macroscopic Models for Long-Range Dependent Network Traffic. *Queuing Systems: Theory and Applications*, Vol. 28, pp. 215-243, 1998.

[17] W.E. Leland, M.S. Taqqu, W. Willinger and D.V. Wilson. On the Self-similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Trans. on Networking*, Vol. 2, pp. 1-15, 1994.

[18] N. Likhanov and R.R. Mazumdar. Cell Loss Asymptotics in Buffers Fed with a Large Number of Independent Stationary Sources. In *Proc. INFOCOMM'98*, pp. 339-346, San Francisco, CA, 1998.

[19] B.B. Mandelbrot and J.W. Van Ness. Fractional Brownian Motions, Fractional Noises and Applications. *SIAM Rev.*, Vol. 10, pp. 422-437, 1968.

[20] J.H. McCulloch. Simple Consistent Estimators of Stable Distribution Parameters. *Commun. Stat. - Simulation*, Vol. 15, pp. 1109-1136, 1986.

[21] C.L. Nikias and M. Shao. *Signal Processing with Alpha-Stable Distributions and Applications*. John Wiley & Sons, 1995.

[22] I. Norros. A Storage Model with Self-Similar Input. *Queuing Systems*, Vol. 16, pp. 387-396, 1994.

[23] V. Paxson and S. Floyd. Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Trans. on Networking*, Vol.3, pp. 226-244, 1995.

[24] O. Rose. Statistical Properties of MPEG Video Traffic and Their Impact on Traffic Modeling in ATM Systems. In *Proc. LCN'95*, pp. 397-406, Minneapolis, MN, 1995.

[25] G. Samorodnitsky and M.S. Taqqu. *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. London, UK: Chapman and Hall, 1994.

[26] M.S. Taqqu and V. Teverovsky. On Estimating the Intensity of Long-Range Dependence in Finite and Infinite Variance Time Series. In *A Practical Guide to Heavy Tails: Statistical techniques and Applications*, R.J. Adler, R.E. Feldman and M.S. Taqqu, Eds. Birkhauser, Boston, pp. 177-217, 1998.

[27] S. Tartarelli, M. Falkner, M. Devetsikiotis and I. Lambadaris. Empirical Effective Bandwidth. In *Proc. GLOBECOM'00*, pp. 672 -678, San Francisco, CA, 2000.