Network Multiplexer with Truncated Heavy-Tailed Arrival Streams

Predrag R. Jelenković
Department of Electrical Engineering
Columbia University
New York, NY, 10027
predrag@comet.columbia.edu

Abstract

This paper investigates the asymptotic behavior of a single server queue with truncated heavy-tailed arrival sequences. We have discovered and explicitly asymptotically characterized a unique asymptotic behavior of the queue length distribution. Informally, this distribution on the log scale resembles a *stair-wave* function that has steep drops at specific buffer sizes. This has important design implications suggesting that negligible increases of the buffer size in certain buffer regions can decrease the overflow probabilities by order of magnitudes.

A problem of this type arises quite frequently in practice when the arrival process distribution has a bounded support and inside that support it is nicely matched with a heavy-tailed distribution (e.g. Pareto). However, our primary interest in this scenario is in its possible application to controlling heavy-tailed traffic flows. More precisely, one can imagine a network control procedure in which short network flows are separated from long ones. If the distribution of flows is heavy-tailed this procedure will yield truncated heavy-tailed distribution for the short network flows. Intuitively, it can be expected that with short flows one can obtain much better multiplexing gains than with the original ones (before the separation). Indeed, our analysis confirms this expectation.

Keywords: *Truncated Heavy-Tailed Distributions; Regular Variation; Network Multiplexer; Single Server Queue; GI/GI/1 Queue; Long Range Dependence.*

1 Introduction

Increasing empirical evidence has demonstrated the presence of heavy-tailed (subexponential) characteristics in communication network traffic streams. Early discoveries on the heavy-tailed nature of Ethernet traffic was reported in [26]. More recently, in [12] the long-range dependency of Ethernet traffic was attributed to the heavy-tailed file sizes that are transfered over the network. Heavy-tailed characteristics of the scene length distribution of MPEG video streams were explored in [17, 24].

These empirical findings have encouraged theoretical developments in the modeling and analysis of heavy-tailed phenomena. In this area there have been two basic approaches: self-similar processes and fluid renewal models with heavy-tailed renewal distributions. The investigation of queueing systems with self-similar long-range dependent arrival processes can be found in [32, 14, 27, 35, 39, 13].

Basic tools for the analysis of fluid renewal models with a single heavy-tailed arrival stream are the classical results on subexponential (heavy-tailed) asymptotic behavior of the waiting time distribution in a GI/GI/1 queue [10, 34, 41] (these results were used in [3, 16]). Asymptotic expansion refinements of these results can be found in [42, 1, 11]. Generalizations to queueing processes (random walks) with dependent increments were investigated in [4, 22, 5].

Queueing models with multiple long range dependent arrival streams are of particular interest for engineering communication networks. Unfortunately, the analysis of these models is much more difficult due to the complex dependency structure in the aggregate arrival process [16]. An intermediate case of multiplexing a single heavy-tailed stream with exponential streams was investigated in [7, 23, 38, 2].

For the case of multiplexing more than two heavy-tailed arrival processes general bounds were obtained in [9, 29]. In [7] a limiting process obtained by multiplexing an infinite number of On-Off sources with regularly varying on periods was analyzed. This limiting arrival process, so called $M/G/\infty$ process [36, 40], appears to be quite promising for the analysis of a corresponding fluid queue. In [23], under specific stability conditions, an explicit asymptotic formula for the behavior of the infinite buffer queue length distribution with $M/G/\infty$ arrivals was derived. In the same paper it was shown with simulation experiments that the derived asymptotic formula yields good approximation for multiplexing finitely many heavy-tailed On-Off sources. Asymptotic expression for the expected value of the first passage time in a fluid queue with M/G/∞ arrivals was derived in [37]. New results on a finite buffer queue can be found in [21, 18]. A recent survey of results on fluid queues with heavy-tailed arrivals can be found in [8]; for the latest comprehensive list of references see [20].

In this paper we examine the idea of separating the transmission of long and short network flows. This idea of flow separation was presented in a sequence of papers by P. Newman et al. [30, 31]. In these papers authors mostly discuss system and implementation issues. They explain how ATM switching technology can be combined with IP packet (connectionless) routing to implement this idea. ATM virtual circuits (VC) are suggested for transmitting long flows, while IP packet routing is used for transmitting short flows. One of the potential benefits of having this network architecture is to avoid costly router table lookups (usually done in software) for packets in long flows and use fast ATM hardware forwarding for its transimission.

Here, we focus on some performance aspects that can arise in implementing the flow separation idea. We use the increasing expected residual life time of heavy-tailed distributions to guarantee the extraction of long flows with a simple threshold based mechanism. This is important because establishing a VC involves complex signaling and "pays off" only if the flow for which a VC is established is indeed long. This is presented in Section 2.

The main focus of our investigation is in characterizing the queueing behavior with short (truncated heavy-tailed) arrivals. Our main result in this direction is presented in Theorem 2 in Section 4; preliminary results needed for this theorem are contained in Section 3. Informally, the queue length distribution (characterized by Theorem 2) on the log scale resembles a *stair-wave* function that has steep drops at specific buffer sizes. This has important design implications suggesting that negligible increases of the buffer size in certain buffer regions can decrease the overflow probabilities by order of magnitudes. These steep drops of the queue overflow probability represent the increased multiplexing gain of the short flows. Numerical illustrations of our results is presented in Section 5 (see Figures 2 and 4). The paper is concluded in Section 6.

2 Flow Separation

Let us first define a source (session) model. We assume that a source (session) can be represented as an alternating sequence of active (on) and silence (off) periods; during an activity period (at least on a fluid level) the source produces a constant bit rate traffic of rate r; each on period will be referred as a *flow*. Let $\{\tau_n^{on}, n \geq 0\}$, $\{\tau_n^{off}, n \geq 0\}$, be two independent sequences of i.i.d. random variables representing the durations of successive on and off periods, respectively; e.g., a sample path realization of this on-off source model is shown on the right-hand side of Figure 1.

We model on periods as being regularly varying. The class of functions of regular variation $\mathcal{R}_{-\alpha}, \alpha \in \mathbb{R}$ was invented by Karamata [25]. These functions provide a framework for the asymptotic inversion of the Laplace Transform (see the main reference book [6]). Pareto distributions are the best known example from this family. Formally, we say that a distribution function F is regularly varying, $1 - F \in \mathcal{R}_{-\alpha}$ if it is given by

$$F(x) = 1 - \frac{l(x)}{x^\alpha} \qquad \alpha \geq 0,$$

where $l(x) : \mathbb{R}_+ \to \mathbb{R}_+$ is a function of slow variation, i.e., $\lim_{x\to\infty} l(\delta x)/l(x) = 1, \delta > 1$. Functions of regular variation have played a role in the queueing theory since the classical result of Cohen [10] on the asymptotic property of a workload distribution in a $GI/GI/1$ queue.

Next, we propose a simple, threshold based algorithm for separating the flows (see Figure 1). We assume that at the beginning of each flow (on period) we start a clock that measures the duration τ_e of that flow. When τ_e exceeds a threshold B we term the remainder of that flow as being long, and from that time on we handle that flow separately in a connection oriented fashion.

2.1 Optimizing the Choice of B

Regularly varying distributions are characterized by increasing expected residual life time. In other words, something that had lasted for a significant amount of time B is expected to continue to exist for an amount of time proportional to B. This is formally expressed in the following lemma. Throughout the paper we will use the customary notation $f(x) \sim g(x)$ as $x \to \infty$ to denote $\lim_{x\to\infty} f(x)/g(x) = 1$

Lemma 1 *If* $\mathbb{P}[\tau^{on} > x] = l(x)/x^\alpha (\in \mathcal{R}_{-\alpha}), \alpha > 1$, *then*

$$\mathbb{E}[\tau^{on} - B | \tau^{on} > B] \sim \frac{B}{\alpha - 1} \quad \text{as} \quad B \to \infty.$$

Proof: Follows from Karamata's theorem. $\diamond$

Thus, by going back to our threshold based flow separation algorithm we see that the long flows are indeed going to be long, i.e. their expectation is roughly going to be $B/(\alpha - 1)$. This is important because establishing a virtual connection for a long flow is more

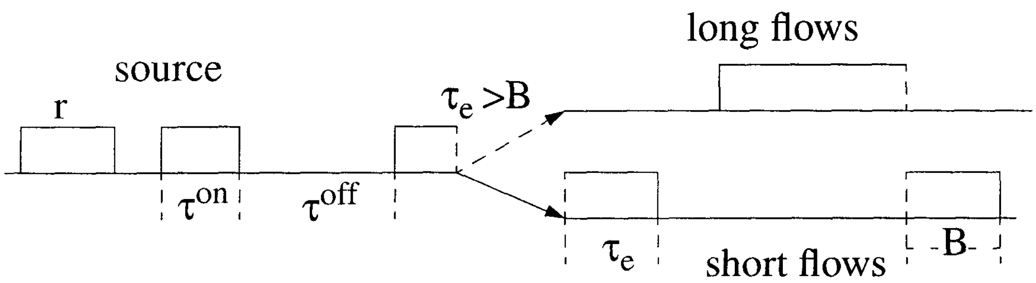

Figure 1: Threshold based flow separation.

complex (than connectionless transmission) and in general "pays off" only if the duration of the flow is long. Using Lemma 1 we can optimize the choice of B such that the remainder of the flow which is transmitted via a VC has a desired expected value.

In passing, we would like to mention that an important issue is to decide when to close down a VC; e.g., this can be done by having an appropriate silence detector.

2.2 Queueing Short Flows

It is reasonable to expect that a sequence of short flows, since long flows are extracted, will yield much better queueing performance. Because of that, it will be quite inefficient to assign to the sequence of short flows a capacity that is equal to its peak rate. Hence, if we assign a capacity $c < r$ to each session, it is important to be able to estimate buffer overflow probabilities and optimize the buffer design.

Also, in the interest of increased utilization, many short sessions can share the same link capacity and buffer space. For a large number of sources it is often good to approximate the arrival process as an M/G/∞ arrival process (see [23]), i.e., to assume that the beginnings of short flows arrive according to a Poisson arrival process of rate Λ = (sum of the rates of all multiplexed sessions). At this point, the analysis of a fluid queue with truncated heavy-tailed M/G/∞ arrivals appears to be a very difficult problem. In order to make the analysis feasible, we assume that short flows are arriving instantaneously with Poisson rate Λ. This queueing system will provide an upper bound on the performance of a corresponding fluid queue. Then, the queue length distribution observed at the flow arrival times evolves as

$$Q_{n+1}^B = (Q_n^B + r\tau_n^{on,B} - c\tau_n)^+, \quad n \geq 0, \tag{1}$$

where τ_n are Poisson inter-arrival times and

$$\tau_n^{on,B} \stackrel{\text{def}}{=} \tau_n^{on} 1(\tau_n^{on} < B) + B1(\tau_n^{on} \geq B).$$

Characterizing the stationary asymptotic behavior of $\mathbb{P}[Q_n^B > x]$ as both x and B go to ∞ is the main question that we will explore in the rest of the paper. To answer this question, we first investigate large deviations of truncated heavy-tailed sums in the next section. Then, by using this large deviation result, in the subsequent section we characterize the behavior of $\mathbb{P}[Q_n^B > x]$.

3 Large Deviations of Truncated Heavy-Tailed Sums

Let $\{Y, Y_i, i \geq 1\}$ be a sequence of non-negative i.i.d. random variables with density $f(x) = l(x)/x^\alpha \in \mathcal{R}_{-\alpha}, \alpha > 1$. Next, for each $B > 0$, construct an i.i.d. sequence $\{Y^B, Y_i^B, i \geq 1\}$ with density $f^B(x) = f(x)/\mathbb{P}[0 \leq Y \leq B], 0 \leq x \leq B$.

Theorem 1 *Let $S_n^B = \sum_{i=1}^n Y_i^B, n \geq 1$. If $f(x) = l(x)/x^\alpha$, $(\in \mathcal{R}_{-\alpha})$, $\alpha > 1$, then for any constant $K > 0$, fixed $k = 0, 1, \ldots$, fixed $0 < \delta < 1$, and uniformly for all $k + 1 \leq n \leq K \log B$,*

$$\mathbb{P}[S_n^B \geq (k+\delta)B] \sim \binom{n}{k+1} h_k(\delta) \frac{l(B)^{k+1}}{B^{(k+1)(\alpha-1)}}, \tag{2}$$

as $B \to \infty$, where

$$h_k(\delta) \stackrel{\text{def}}{=} \int_{\substack{0 \leq x_i \leq 1, 1 \leq i \leq k+1 \\ x_1 + \cdots + x_{k+1} \geq k+\delta}} x_1^{-\alpha} \cdots x_{k+1}^{-\alpha} dx_1 \cdots dx_{k+1}. \tag{3}$$

Remarks: **(i)** Note that the interval $k + 1 \leq n \leq K \log B$ is not the largest one for which (2) holds, but it suffices our needs. **(ii)** Also, note that $h_k(\delta), k = 0, 1$, are explicitly given by

$$h_0(\delta) = \frac{1}{(\alpha-1)\delta^{\alpha-1}}(1 - \delta^{\alpha-1}), \tag{4}$$

$$h_1(\delta) = \left((\alpha-2)(\alpha-1)^2\delta^\alpha\right)^{-1}\{(\alpha-2)(\delta^\alpha-\delta) \quad (5)$$
$$-(\alpha-2)\delta^\alpha(1+\delta)^{-\alpha+1}{}_2F_1\left(1-\alpha,\alpha,2-\alpha,(1+\delta)^{-1}\right)$$
$$+(\alpha-2)\delta(1+\delta)^{-\alpha+1}{}_2F_1\left(1-\alpha,\alpha,2-\alpha,\delta(1+\delta)^{-1}\right)$$
$$+(\alpha-1)\delta^\alpha(1+\delta)^{-\alpha}{}_2F_1\left(2-\alpha,\alpha,3-\alpha,(1+\delta)^{-1}\right)$$
$$-(\alpha-1)\delta^2(1+\delta)^{-\alpha}{}_2F_1\left(2-\alpha,\alpha,3-\alpha,\delta(1+\delta)^{-1}\right)\},$$

where ${}_2F_1$ is the Hypergeometric function.

Proof: Observe that for any $y > 0$, $\mathbb{P}[S_n^B \geq (k+\delta)B]$ can be decomposed as follows

$$\begin{aligned}
&\mathbb{P}[S_n^B \geq (k+\delta)B] \\
&= \mathbb{P}[S_n^B \geq (k+\delta)B, \max_{1\leq i\leq n} Y_i^B \leq y] \\
&\quad + n\mathbb{P}[S_n^B \geq (k+\delta)B, Y_1^B > y, \max_{2\leq i\leq n} Y_i^B \leq y] \\
&\quad + \binom{n}{2}\mathbb{P}[S_n^B \geq (k+\delta)B, Y_1^B > y, Y_2^B > y, \\
&\qquad \max_{3\leq i\leq n} Y_i^B \leq y] \\
&\qquad \vdots \quad \vdots \\
&\quad + \binom{n}{k+1}\mathbb{P}[S_n^B \geq (k+\delta)B, Y_1^B > y, \dots, \\
&\qquad Y_{k+1}^B > y, \max_{k+2\leq i\leq n} Y_i^B \leq y] \\
&\quad + \mathbb{P}[S_n^B \geq (k+\delta)B, \cup_{1\leq i_1<i_2<\cdots<i_{k+2}\leq n} \\
&\qquad \{Y_{i_1}^B > y, Y_{i_2}^B > y, \dots, Y_{i_{k+2}}^B > y\}] \\
&\stackrel{\text{def}}{=} P_{0n} + nP_{1n} + \binom{n}{2}P_{2n} + \cdots \\
&\qquad + \binom{n}{k+1}P_{(k+1)n} + P_{(k+2)n} \qquad (6)
\end{aligned}$$

Now, let us choose θ such that $(k+1)/(k+2) < \theta < 1$, and $y = B^\theta$. Then, for all sufficiently large B such that $K\log B < \delta B^{1-\theta}$, we have that $nB^\theta < \delta B$, which implies

$$P_{in} \equiv 0 \qquad 0 \leq i \leq k. \qquad (7)$$

Next, the estimate of $P_{(k+2)n}$ follows from

$$\begin{aligned}
P_{(k+2)n} &\leq \binom{n}{k+2}(\mathbb{P}[Y_1 > B^\theta])^{k+2} \\
&= o((l(B)/B^{\alpha-1})^{k+1}) \quad \text{as} \quad B \to \infty. \quad (8)
\end{aligned}$$

To finish the proof we need to estimate $P_{(k+1)n}$. Observe that for any $0 < \epsilon < \delta$ there exists sufficiently large B_ϵ, such that for all $B > B_\epsilon$,

$$\sum_{i=k+2}^{n} Y_i^B \leq \epsilon B \qquad (9)$$

on the set $\{\max_{k+2\leq i\leq n} Y_i^B \leq y\}$. Hence, for $B > B_\epsilon$

$$\begin{aligned}
P_{(k+1)n}(B) &= \mathbb{P}\left[S_n^B \geq (k+\delta)B, Y_1^B > y, \dots,\right. \\
&\qquad \left. Y_{k+1}^B > y, \max_{k+2\leq i\leq n} Y_i^B \leq y\right] \\
&\leq \mathbb{P}\left[S_{k+1}^B \geq (k+\delta-\epsilon)B\right] \\
&= \int_{B(\delta-\epsilon)}^{B} f^B(y_1)dy_1 \int_{B(1+\delta-\epsilon)-y_1}^{B} f^B(y_2)dy_2 \\
&\qquad \cdots \int_{B(k+\delta-\epsilon)-y_1\cdots-y_k}^{B} f^B(y_{k+1})dy_{k+1}. \qquad (10)
\end{aligned}$$

By upper bounding $l(x)$ with its maximum value in the interval $[B(\delta-\epsilon), B]$, and by changing the variables $y_i = x_i B$, $1 \leq i \leq k+1$ we obtain

$$\begin{aligned}
P_{(k+1)n}(B) &\leq \frac{(\max_{B(\delta-\epsilon)\leq x\leq B} l(x))^{k+1}}{\mathbb{P}[Y \leq B]^{k+1}B^{(k+1)(\alpha-1)}}h_k(\delta-\epsilon) \\
&\sim \frac{l(B)^{k+1}}{B^{(k+1)(\alpha-1)}}h_k(\delta-\epsilon), \qquad (11)
\end{aligned}$$

as $B \to \infty$, where for the last asymptotic relation we have used $\max_{B(\delta-\epsilon)\leq x\leq B} l(x) \sim l(B)$ as $B \to \infty$, which follows from Theorem 1.2.1, p. 6, [6]. Next, by observing that $h_k(\delta)$ is continuous in δ, by passing $\epsilon \to 0$ in (11) we obtain the asymptotic upper bound

$$\limsup_{B\to\infty} \frac{P_{(k+1)n}(B)B^{(k+1)(\alpha-1)}}{l(B)^{k+1}} \leq h_k(\delta). \qquad (12)$$

For the lower bound we observe that for sufficiently large B ($> B_\epsilon$)

$$\{S_{k+1}^B \geq (k+\delta)B\} \subset \{Y_1^B > y, \dots, Y_{k+1}^B > y\},$$

which implies

$$\begin{aligned}
P_{(k+1)n} &\geq \mathbb{P}[Y^B \leq B^\theta]^{n-k-1}\mathbb{P}\left[S_{k+1}^B \geq (k+\delta)B,\right. \\
&\qquad \left. Y_1^B > y, \dots, Y_{k+1}^B > y\right] \\
&= \mathbb{P}[Y^B \leq B^\theta]^{n-k-1}\mathbb{P}\left[S_{k+1}^B \geq (k+\delta)B\right] \\
&\geq \mathbb{P}[Y^B \leq B^\theta]^{n-k-1}\frac{(\inf_{B\delta\leq x\leq B} l(x))^{k+1}h_k(\delta)}{\mathbb{P}[Y \leq B]^{k+1}B^{(k+1)(\alpha-1)}} \\
&\sim \frac{l(B)^{k+1}h_k(\delta)}{B^{(k+1)(\alpha-1)}} \quad \text{as} \quad B \to \infty, \qquad (13)
\end{aligned}$$

where in the last relation we have used $\inf_{B\delta\leq x\leq B} l(x) \sim l(B)$, as $B \to \infty$ (Theorem 1.2.1, p. 6, [6]). Finally, by combining (12) and (13) we finish the proof of the theorem. ◇

4 Asymptotic Queueing Behavior

Observe that queueing recursion (1) can be generically described as

$$Q_{n+1}^B = (Q_n^B + A_{n+1}^B - C_{n+1})^+, \qquad (14)$$

where $\{A^B, A_n^B, n \geq 1\}$ and $\{C, C_n, n \geq 0\}$ are two independent i.i.d. sequences, and A_n^B has a bounded (truncated) support. Note that (14) represents the customer waiting times in a GI/GI/1 queue.

Without loss of generality we can assume that for any $B > 0$

$$A_n^B := A_n 1(A_n \leq B) + B1(A_n > B),$$

where $\{A, A_n, n \geq 1\}$ is an i.i.d. process. According to the classical result of Loynes' [28], under the stability condition $\mathbb{E}A_n < \mathbb{E}C_n$, recursion (14) admits a unique stationary solution, and for all initial conditions $\mathbb{P}[Q_n^B \leq x]$ converges to the stationary distribution $\mathbb{P}[Q^B \leq x]$. For the rest of this paper we will assume that all the queueing systems under consideration are in their stationary regimes. When $B = \infty$ we denote Q^∞ simply as Q.

Theorem 2 *If $\mathbb{E}(A - C) < 0$, for all $n > 0$, $\mathbb{P}[C > x] \leq e^{-\eta x}, \eta > 0$, and A has a regularly varying distribution $\mathbb{P}[A > x] = l(x)/x^\alpha, \alpha > 1$, then*

$$\mathbb{P}[Q^B > (k + \delta)B] = \frac{h_k(\delta)(1 + o(1))}{(\mathbb{E}C - \mathbb{E}A)^{k+1}} \frac{l(B)^{k+1}}{B^{(k+1)(\alpha-1)}}, \tag{15}$$

as $B \to \infty$, where $h_k(\delta), 0 < \delta < 1, k = 0, 1, 2, \ldots$ are explicitly computable from (3).

Remark: If $k = 0$ and $\delta \lll 1$, i.e., the buffer size $b = \delta B$ is much smaller then the truncation point B then by (4), as intuitively expected, Theorem 2 gives approximately the same result as Cohen's (or Pakes') result.

Initially, we assume that A and C are integer valued, and that the distribution of C has a bounded support $\mathbb{P}[C \leq c] = 1$, $c < \infty$.

By using a well known connection between the queue length distribution and the supremum of the corresponding random walk with increments $X_n^B = A_n^B - C_n$, [15], Chapter XII, (or [33], Section 24) the probability generating function $q^B(z)$ of Q^B can be represented as

$$q^B(z) = \frac{1 - g_+^B(1)}{1 - g_+^B(z)}, \tag{16}$$

where $g_+^B(z) = \sum_{i=0}^B g_{+i}^B z^i$ is the generating function of an ascending ladder height random variable for which $g_+^B(1) < 1$ iff $\mathbb{E}X_n^B < 0$. Equation (16) can be written in its equivalent form

$$q_i^B = (1 - g_+^B(1)) \sum_{k=0}^\infty g_{+i}^{B,*k}, \tag{17}$$

where $g_{+i}^{B,*k}$ represents the k-fold convolution of g_{+i}^B.

Lemma 2 *If $\mathbb{P}[A > i]$ is regularly varying, $\mathbb{P}[C \leq c] = 1, c < \infty$, and $\mathbb{E}A < \mathbb{E}C$, then for every $\epsilon > 0$, there exist constants $K_1(\epsilon), K_2(\epsilon)$, such that for all $B > K_1(\epsilon) + K_2(\epsilon)$, $i \in [K_1(\epsilon), B]$*

$$g_{+i}^B \leq (1 + \epsilon) \frac{1 - g_+(1)}{\mathbb{E}C - \mathbb{E}A} \mathbb{P}[A > i],$$

and for $i \in [K_1(\epsilon), B - K_2(\epsilon)]$

$$g_{+i}^B \geq (1 - \epsilon) \frac{1 - g_+(1)}{\mathbb{E}C - \mathbb{E}A} \mathbb{P}[A > i],$$

Proof: Due to space limitations this proof is given in [19]. $\diamond$

Proof of Theorem 2: Let $\{Y_n^B, n \geq 1\}$ be a sequence of random variables distributed as $\mathbb{P}[Y_n^B = i] = g_{+i}^B / g_+^B(1)$, and let $S_n^B = \sum_{i=1}^n Y_i^B$. Then, from (17)

$$\begin{aligned} &\mathbb{P}[Q_n^B > (k + \delta)B] \\ &= (1 - g_+^B(1)) \sum_{n=k+1}^\infty (g_+^B(1))^n \mathbb{P}[S_n^B > (k + \delta)B]. \end{aligned} \tag{18}$$

Next, observe that due to the stability condition $\mathbb{E}A < \mathbb{E}C$

$$\lim_{B \to \infty} g_+^B(1) = g_+(1) < 1. \tag{19}$$

Hence, we can choose K sufficiently large such that for $n(B) := \lfloor K \log B \rfloor$

$$\begin{aligned} \sum_{n=n(B)}^\infty (g_+^B(1))^n \mathbb{P}[S_n^B > (k + \delta)B] \\ \leq \frac{(g_+^B(1))^{n(B)}}{(1 - g_+^B(1))} \\ = o\left(\frac{l(B)^{k+1}}{B^{(k+1)(\alpha-1)}} \right). \end{aligned} \tag{20}$$

Thus, by using stochastic dominance, Lemma 2, Theorem 1, (19) and estimate (20) we conclude that for any $\epsilon_1 > 0$ and all sufficiently large $B > B_\epsilon$, the infinite sum in (18) is bounded with

$$\begin{aligned} &\mathbb{P}[Q_n^B > (k + \delta)B] \\ &\leq (1 + \epsilon_1)(1 - g_+(1)) \left(\frac{(1 - g_+(1)) l(B)}{g_+(1)(\mathbb{E}C - \mathbb{E}A) B^{\alpha-1}} \right)^{k+1} \\ &\quad \times \sum_{n=k+1}^\infty (g_+(1))^n \binom{n}{k+1}. \end{aligned} \tag{21}$$

Finally, by replacing

$$\sum_{n=k+1}^\infty (g_+(1))^n \binom{n}{k+1} = \frac{g_+(1)^{k+1}}{(1 - g_+(1))^{k+2}} \tag{22}$$

in (21) and by passing ϵ_1 to 0 we obtain

$$\limsup_{B\to\infty} \mathbb{P}[Q_n^B > (k+\delta)B] \frac{B^{(\alpha-1)(k+1)}}{l(B)^{k+1}} \le \frac{h_k(\delta)}{(\mathbb{E}C - \mathbb{E}A)^{k+1}}. \tag{23}$$

In a very similar way one can obtain a lower bound

$$\liminf_{B\to\infty} \mathbb{P}[Q_n^B > (k+\delta)B] \frac{B^{(\alpha-1)(k+1)}}{l(B)^{k+1}} \ge \frac{h_k(\delta)}{(\mathbb{E}C - \mathbb{E}A)^{k+1}}; \tag{24}$$

we skip the details. Finally, the combination of (23) and (24) concludes the proof for the case of A and C integer valued and C bounded.

Due to the space constraint the remainder of the proof is presented in the extended version of this paper [19] (all of the numerical examples presented in the following section will have A integer valued and $C \equiv 1$). $\diamond$

5 Numerical Results

In this section we illustrate with numerical examples the accuracy of our main result (Theorem 2) in approximating the queue length distribution with truncated heavy-tailed arrivals.

The following examples are based on a discretized version of the scenario described by equation (1) for the case $c = 1, r = 1$, i.e., we assume that the flows of size $\tau_n^{on,B}$ arrive instantaneously according to a Poisson process. For numerical purposes we consider a discrete (slotted) time approximation of a Poisson process in which arrivals per slot are indicated with $\{I_n, n \ge 0\}$ Bernoulli i.i.d. sequence (independent of $\tau^{on,B}$) with success probability $\mathbb{P}[I_n = 1] = 1 - \mathbb{P}[I_n = 0] = p$; $I_n = 1$ indicates that flow of size $\tau_n^{on,B}$ arrives in a particular slot. Then, the evolution of the queue per one time slot is given by

$$Q_{n+1}^B = (Q_n^B + \tau_n^{on,B} I_n - 1)^+. \tag{25}$$

Note that in this system the inter arrival time of flows is geometric with parameter p; this approximates well the exponential interarrival times of the Poisson process. For simplicity we define $A_n^B = \tau_n^{on,B} I_n$; A_n^B is a truncated heavy-tailed random variable with support $[0, B]$. Let $\rho^B = \mathbb{E}A_n^B$ and let $a_B(z) = \sum_{i=0}^B z^i a_i^B, a_i^B = \mathbb{P}[A_n^B = i]$ be a probability generating function (pgf) of A_n^B. Similarly, let $q_B(z) = \sum_{i=0}^\infty z^i q_i^B, q_i^B = \mathbb{P}[A_n^B = i]$ be a pgf of Q_n^B. Then,

$$q_B(z) = \frac{(1-\rho^B)(z-1)}{z - a_B(z)}. \tag{26}$$

Using Mathematica 2.2 we will invert (26) and compare it to the approximation suggested by equation (15), Theorem 2. We will choose $\tau^{on,B}$ to be truncated Pareto distribution, i.e.,

$$\mathbb{P}[\tau^{on} = i] = \frac{d}{i^{\alpha+1}}, \quad i \ge 1,$$

where d is a normalization constant, and

$$\tau^{on,B} = \tau^{on} 1[\tau^{on} < B] + B1[\tau^{on} \ge B].$$

Then, $a_0^B = 1 - p$ and

$$a_i^B = \frac{pd}{i^{\alpha+1}}, \quad 1 \le i \le B-1.$$

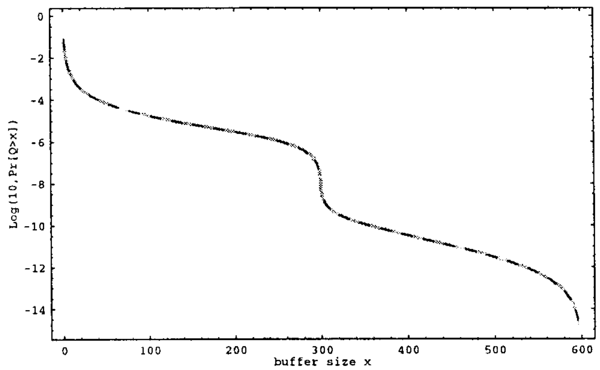

Figure 2: Illustration for Example 1.

Example 1 For the choice of arrival parameters $B = 300, \alpha = 2.8, p = 0.3$, we compute $d = 1/\zeta(\alpha+1) = 0.273345$, where $\zeta(x)$ is Zeta function. Next, $a_0^B = 0.7, a_i^B = pd/i^{\alpha+1}, 1 \le i \le B-1, a_B^B = 1 - \sum_{i=0}^{B-1} a_i, \rho^B = 0.34086$. For these values we numerically invert (26); the exact inverted values of $\mathbb{P}[Q^B > x]$ are plotted with a gray line in Figure 2. The values of approximation (15) are plotted on the same figure with dashed black lines. From the figure we can easily see that the approximation is almost identical to the exactly computed probabilities. Furthermore, if $\hat{Q}_k^B(\delta)$ denotes our approximation and $Q_k^B(\delta) := \mathbb{P}[Q^B > (k+\delta)B]$, then the relative error $e_k(\delta) := |Q_k^B(\delta) - \hat{Q}_k^B(\delta)|/Q_k^B(\delta)$ is presented in Figure 3; from the figure we observe that for buffer sizes $x \in [60, 267] \cup [309, 582]$ the relative approximation error is smaller than 1%!

Example 2 Here, we choose $B = 300, \alpha = 3.5, p = 0.88$. Then, $d = 0.83435$, $a_0 = 0.12$, $a_i^B = pd/i^{\alpha+1}, 1 \le i \le B-1$, $\rho^B = 0.9400$. An excellent agreement between the approximation is evident from Figure 4 (similarly, as in the preceding experiment, the gray line

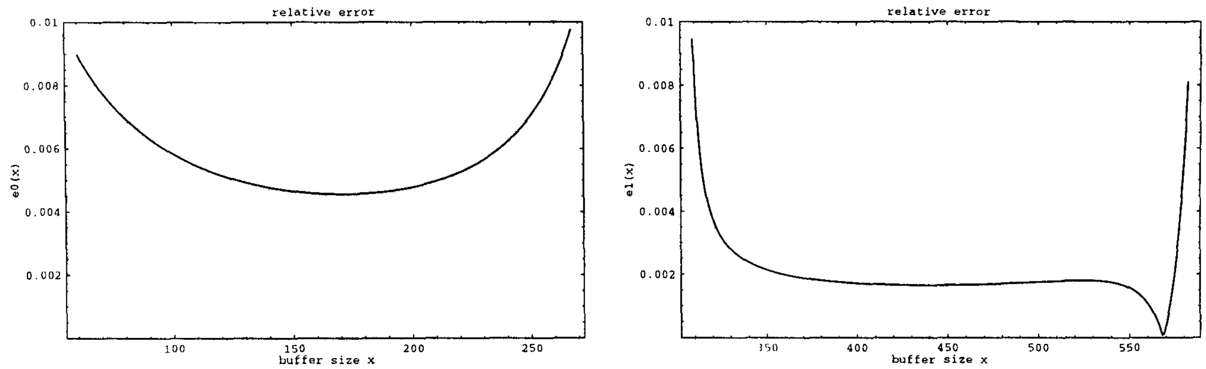

Figure 3: Illustration for Example 1.

represents the exact values and black dashed lines represent the approximation).

6 Conclusion

In this paper, we have shown that heavy-tailed distributions present a natural mathematical framework for the investigation of flow separation. The increasing expected residual life time of heavy-tailed distributions is used to optimize the design of a threshold for the flow separation.

Our main result (Theorem 2) gives an explicit asymptotic characterization of an intriguing behavior of the queue length distribution that results from queueing short (truncated heavy-tailed) flows. Informally, this distribution on the log scale resembles a *stair-wave* function that has steep drops at specific buffer sizes. This has important design implications suggesting that negligible increases of the buffer size in certain buffer regions can decrease the overflow probabilities by order of magnitudes. These steep drops of the queue overflow probability represent the increased multiplexing gain of the short flows.

Besides its applicability to the flow separation problem, the investigated framework is of independent interest for other networking scenarios where truncated heavy-tailed distributions might arise.

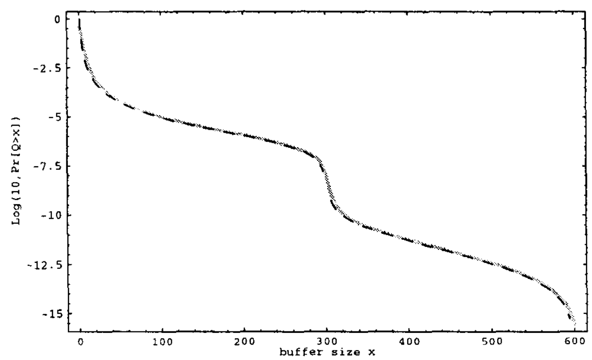

Figure 4: Illustration for Example 2.

References

[1] J. Abate, G.L. Choudhury, and W. Whitt. Waiting-time tail probabilities in queues with long-tail service-time distributions. *Queueing systems*, 16(3/4):311–338, 1994.

[2] R. Agrawal, A. M. Makowski, and P. Nain. On a reduced load equivalence for fluid queues under subexponentiality. preprint, June 1998.

[3] V. Anantharam. On the sojourn time of sessions at an ATM buffer with long-range dependent input traffic. In *Proceedings of the 34th IEEE Conference on Decision and Control*, December 1995.

[4] S. Asmussen, L. F. Henriksen, and C. Klüppelberg. Large claims approximations for risk processes in a Markovian environment. *Stochastic Processes and their Applications*, 54:29–43, 1994.

[5] S. Asmussen, H. Schmidli, and V. Schmidt. Tail probabilities for non-standard risk and queueing processes with subexponential jumps. preprint, 1997.

[6] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation.* Cambridge University Press, Cambridge, 1987.

[7] O. J. Boxma. Regular variation in a multi-source fluid queue. In *ITC 15*, pages 391–402, Washington, D.C., USA, June 1997.

[8] O. J. Boxma and V. Dumas. Fluid queues with long-tailed activity period distributions. Technical Report PNA-R9705, CWI, Amsterdam, NL, April 1997.

[9] G. L. Choudhury and W. Whitt. Long-tail buffer-content distributions in broadband networks. *Performance Evaluation*, 30:177–190, 1997.

[10] J. W. Cohen. Some results on regular variation for distributions in queueing and fluctuation theory. *J. Appl. Probability*, 10:343–353, 1973.

[11] J. W. Cohen. On the M/G/1 model with heavy-tailed service time distributions. In *9th INFORMS Applied Probability Conference*, Cambridge, Massachusetts, June 1997.

[12] M. Crovella. The relationship between heavy-tailed file sizes and self-similar network traffic. In *9th INFORMS Applied Probability Conference*, Cambridge, Massachusetts, June 1997.

[13] K. Debicki, Z. Michna, and T. Rolski. On the supremum from Gaussian processes over infinite horizon. preprint, May 1997.

[14] N. G. Duffield and N. O'Connell. Large deviations and overflow probabilities for the general single-server queue with applications. *Mathematical Proceedings of the Cambridge Philosophical Society*, 118:363–374, 1995.

[15] W. Feller. *An Introduction to Probability Theory and its Application, Volume II.* Wiley, New York, 1971.

[16] D. Heath, S. Resnick, and G. Samorodnitsky. Heavy tails and long range dependence in on/off processes and associated fluid models. preprint, 1996.

[17] D. P. Heyman and T. V. Lakshman. Source models for VBR broadcast-video traffic. *IEEE/ACM Transactions on Networking*, 4:40–48, February 1996.

[18] P. R. Jelenković. Long-tailed loss rates in a single server queue. In *Proceedings of INFOCOM'98*, San Francisco, California, April 1998.

[19] P. R. Jelenković. GI/GI/1 queue with truncated long-tailed service times. submitted for publication, 1999.

[20] P. R. Jelenković. Asymptotic results for queues with subexponential arrival processes. In K. Park and W. Willinger, editors, *Self-Similar Network Traffic and Performance Evaluation.* Wiley (Interscience Division), 1999, to appear.

[21] P. R. Jelenković. Long-tailed loss rates in a GI/GI/1 queue with applications. *Queueing Systems*, 1999, to appear.

[22] P. R. Jelenković and A. A. Lazar. Subexponential asymptotics of a Markov-modulated random walk with queueing applications. *Journal of Applied Probability*, 35(2):325–347, June 1998.

[23] P. R. Jelenković and A. A. Lazar. Asymptotic results for multiplexing subexponential on-off processes. Advances in Applied Probability, 31(2), June 1999, to appear.

[24] P. R. Jelenković, A. A. Lazar, and N. Semret. The effect of multiple time scales and subexponentiality of MPEG video streams on queueing behavior. *IEEE Journal on Selected Areas in Communications*, 15(6):1052–1071, August 1997.

[25] J. Karamata. Sur un mode de croissance régulière des fonctions. *Mathematica (Cluj)*, 4:38–53, 1930.

[26] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. In *SIGCOMM'93*, pages 183–193, 1993.

[27] N. Likhanov, B. Tsybakov, and N. D. Georganas. Analysis of an ATM buffer with self-similar ("fractal") input traffic. In *INFOCOM'95*, pages 985–991, Boston, Masssachusetts, April 1995.

[28] R. M. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proc. Cambridge Philos Soc.*, 58:497–520, 1968.

[29] P. Nain, Z. Liu, D. Towsley, and Z.-L. Zhang. Asymptotic behavior of a multiplexer fed by a long-range dependent process. In *9th INFORMS Applied Probability Conference*, Cambridge, Massachusetts, June 1997.

[30] P. Newman, G. Minishall, and T. Lyon. IP switching: ATM under IP. preprint, November 1997.

[31] P. Newman, G. Minishall, T. Lyon, and L. Huston. IP switching and gigabit routers. IEEE Communications Magazine, January 1997.

[32] I. Norros. A storage model with self-similar input. *Queueing Systems*, 16:387–396, 1994.

[33] Billingsley P. *Probability and Measure.* John Wiley & Sons, Inc., 1986.

[34] A.G. Pakes. On the tails of waiting-time distribution. *J. Appl. Probab.*, 12:555–564, 1975.

[35] M. Parulekar and A. M. Makowski. Tail probabilitites for a multiplexer with self-similar traffic. In *INFOCOM'96*, San Francisco, California, March 1996.

[36] M. Parulekar and A. M. Makowski. Tail probabilitites for $M/G/\infty$ input processes (I): preliminary asymptotices. *Queueing Systems*, 27:271–296, 1997.

[37] S. Resnick, D. Heath, and G. Samorodnitsky. Patterns of buffer overflow in a class of queues with long memory in the input stream. *Annals of Applied Probability*, 7:1021–1057, 1997.

[38] T. Rolski, S. Schlegel, and V. Schmidt. Asymptotics of palm-stationary buffer content distribution in fluid flow queues. In *9th INFORMS Applied Probability Conference*, Cambridge, Massachusetts, June 1997.

[39] B. K. Ryu and S. B. Lowen. Point process approaches to the modeling and analysis of self-similar traffic - part I: Model construction. In *INFOCOM'96*, San Francisco, California, March 1996.

[40] K. P. Tsoukatos and A. M. Makowski. Heavy traffic analysis for a multiplexer driven by $M/GI/\infty$ input processes. In *ITC 15*, pages 497–506, Washington, D.C., USA, June 1997.

[41] N. Veraverbeke. Asymptotic behavior of Wiener-Hopf factors of a random walk. *Stochastic Proc. and Appl.*, 5:27–37, 1977.

[42] E. Willekens and J. L. Teugels. Asymptotic expansion for waiting time probabilities in an M/G/1 queue with long-tailed service time. *Queueing Systems*, 10:295–312, 1992.