Long-Tail Buffer-Content Distributions in Broadband Networks

Gagan L. Choudhury

AT&T Labs, Room 1L-238, Holmdel, NJ 07733-3030;
email: gagan@buckaroo.att.com

Ward Whitt

AT&T Labs, Room 2C-178, Murray Hill, NJ 07974-0636;
email: wow@research.att.com

We identify conditions under which relatively large buffers will be required in broadband communication networks. For this purpose, we analyze an infinite-capacity stochastic fluid model with a general stationary environment process (without the usual independence or Markov assumptions). With that level of generality, we are unable to establish asymptotic results, but by a very simple argument we are able to obtain a revealing lower bound on the steady-state buffer-content tail probability. The bounding argument shows that the steady-state buffer content will have a long-tail distribution when the sojourn time in a set of states with positive net input rate itself has a long-tail distribution. If a set of independent sources, each with a general stationary environment process, produces a positive net flow when all are in high activity states, and if each of these sources has a high-activity sojourn-time distribution with a long tail, then the steady-state buffer-content distribution will have a long tail, but possibly one that decays faster than the tail for any single component source. The full buffer-content distribution can be derived in the special case of a two-state fluid model with general high and low activity time distributions, assuming that successive high and low activity times come from independent sequences of i.i.d. random variables. In that case the buffer-content distribution will have a long tail when the high-activity-time distribution has a long tail. We illustrate by giving numerical examples of the two-state model based on numerical transform inversion.

1 Introduction

It has now become common to use infinite-capacity stochastic fluid models to represent the buffer content in broadband communication networks, e.g., see Anick, Mitra and Sondhi [7] and Roberts [30]. Typically the buffer receives input from several access lines and is emptied by a single high-speed backbone communication link. A realistic fluid model can be based on a countable-state environment process $\{Z(t) : t \geq 0\}$. When $Z(t) = i$, there is constant net fluid flow into or out of the buffer at rate r_i, with there being no decrease in buffer content when the buffer is empty. Each environment state can correspond to a set of activity levels for the sources. In a simple model for one source there may be just two activity levels: high and low (which may be on and off, but need not be). However, the general fluid model framework allows any number of levels for each source. Even though the fluid model is an idealization, the possibility of many levels for each source makes it a fairly believable model for a wide variety of sources, including video; e.g., see Sen, Maglaris, Rikli and Anastassiou [31].

There is growing concern that traffic burstiness may make large buffer contents relatively likely. In particular, it is thought that, with realistic models of traffic, the distribution of the buffer content in the infinite-capacity model may have a long tail, e.g., it may decay as a power instead of exponentially; i.e., if B is the steady-state buffer content, then we may have

$$P(B > x) \sim \alpha x^{-\eta} \text{ as } x \to \infty \tag{1}$$

instead of

$$P(B > x) \sim \alpha e^{-\eta x} \text{ as } x \to \infty , \tag{2}$$

where α and η are positive constants and $f(x) \sim g(x)$ means that $f(x)/g(x) \to 1$ as $x \to \infty$. Alternatively, the buffer-content distribution may have a Weibull-like tail (which we also regard as a long tail)

$$P(B > x) \sim \alpha e^{-\eta x^b} \text{ as } x \to \infty \tag{3}$$

for $0 < b < 1$, as suggested by analysis of fractional Brownian motion input by Norros [27] and Duffield and O'Connell [16].

The main difficulty with the long-tail asymptotic forms (1) and (3) is that larger buffers will be required in order to provide the desired quality of service. In addition, the less familiar asymptotic forms (1) and (3) make network management more difficult. For example, an intuitively appealing notion of

effective bandwidths has been developed based on the exponential asymptotic form (2), e.g., see Chang and Thomas [13], deVeciana, Kesidis and Walrand [32] and Whitt [33]. Unfortunately, the asymptotic theory supporting the nice additive effective bandwidths breaks down when (1) or (3) holds instead of (2); e.g., see Duffield and O'Connell [16] and Duffield [15].

The possibility of long-tail buffer-content distributions such as in (1) and (3) is indicated by extensive traffic measurements in recent years; e.g. see Meier-Hellstern, Wirth, Yan and Hoeflin [26], Leland, Taqqu, Willinger and Wilson [25], Paxson and Floyd [29], Willinger [34] and Willinger, Taqqu, Leland and Wilson [35]. In traffic data, the observed long-tail distributions are for high and low activity periods. The observed long-range dependence is the observation that the variance of the number of packets or bytes arriving in an interval of length t grows faster than $O(t)$ as $t \to \infty$.

As indicated in the papers discussing traffic measurements, the data tend to be inconsistent with the familiar traffic models. This has led to the development of new traffic models, e.g., involving fractals; see [18]. Within the context of existing models, an important question is: How can the models better reflect the data? Toward that end, we ask: When will a stochastic fluid model have a long-tail buffer-content distribution as in (1) or (3)?

The purpose of this paper is to show that long-tail distributions such as in (1) and (3) arise quite naturally in stochastic fluid models when a sojourn time of the environment process $Z(t)$ in a set of high-activity states itself has a long-tail distribution. We do not contend that long-tail buffer-content distributions necessarily will prevail. Instead, we present conditions under which they will occur, and other conditions under which they will not occur. Hopefully these results will help to interpret traffic data and prospective network controls.

A major goal here is to treat quite general models. In particular, since considerable dependence has been found in traffic data, we want to avoid the usual Markov and independence assumptions. Hence, here the underlying environment process $\{Z(t) : t \geq 0\}$ is assumed to be a general stationary process, by which we mean (as usual) that the joint distribution of the vector $(Z(t_1 + s), \ldots, Z(t_n + s))$ is independent of s for all n and all n-tuples $(t_1, \ldots, t_n)$ with $0 < t_1 < \ldots < t_n$.

With that generality (without extra conditions), we are unable to establish asymptotic results of the form (1)–(3), but we are able to establish a lower bound for the buffer-content tail probabilities, which enables us to provide conditions under which the steady-state buffer-content distribution necessarily has a long tail. The proofs here are remarkably short, but their brevity is deceptive. The proofs rely on the (deep) theory of stationary point processes, as in Baccelli and Brémaud [9], Brandt, Franken and Lisek [12] and Franken,

König, Arndt and Schmidt [20].

Moreover, we show that in one special case (a two-state fluid model with high and low activity times that are independent sequences of i.i.d. random variables) the buffer-content distribution can be fully analyzed; i.e., we can compute the exact buffer-content distribution and asymptotic approximations. We also indicate that this special case can arise in practice.

However, the cases of primary interest are more complicated, involving the superposition of input from many different sources (statistical multiplexing). In such more complicated settings, we establish bounds on the buffer-content tail probabilities. We show that multiplexing may soften or even remove the long-tail effect. However, more work is needed here, since our results are only in terms of lower bounds.

Our bounding results complement recent more detailed limit theorems by Asmussen, Henriksen and Kluppelberg [8], Boxma [11], Duffield [15], Duffield and O'Connell [16] and Jeleković and Lazar [21], [22] (some of which have been established at the time or after the results here.) (Earlier related results are by Cohen [14] and Pakes [28]; see [2].) As illustrated by our results for the model with a two-state environment process in Section 3, stronger results of the form (1) and (3) tend to require additional independence and Markov assumptions. Establishing full convergence is clearly much more challenging, but our analysis is appealing because, we are able to quickly address an important engineering problem.

Before leaving the introduction, we point out that even if there is long-tailed behavior in individual uncontrolled sources, that does not necessarily mean that it will show up in buffer-content distributions. It may disappear either due to multiplexing (as we will discuss in Section 4) or due to individual connection controls (e.g., finite window sizes allowed by end-to-end protocols), or congestion controls imposed by the network (e.g., leaky-bucket-type controls limiting short-term and long-term data rates). On the other hand, long-tail buffer-content distributions may occur even in the presence of controls if each source can choose its high activity rates such that the controls allow them to be sustained for arbitrarily long durations.

2   A Lower Bound for Fluid Models

We initially only assume that the environment process $\{Z(t) : t \geq 0\}$ is a stationary stochastic process. We thus initially do not impose any independence conditions. As mentioned in the introduction, our model thus allows the superposition of many sources, each with many different activity levels. Moreover,

4

the activity durations may have general (including long-tail) distributions.

Divide the environment states into two subsets: $U \equiv \{i : r_i \geq \lambda > 0\}$ and $D \equiv \{i : r_i < \lambda\}$. Let $-\mu = \min\{r_i\}$. We assume that U and D are nonempty and that $\mu > 0$. It is elementary that the fluid content process is bounded below, for each sample path, by the fluid content process with the two-state environment process with net flow rate λ in U and net flow rate $-\mu$ in D.

Let X_n and Y_n be the successive holding times in sets U and D, respectively. (We assume that these holding times are well defined, which will always be the case if $Z(t)$ is countably valued with only finitely many jumps in any bounded interval w.p.1.) Since $\{Z(t) : t \geq 0\}$ is a general stationary process, these random variable are in general not independent. However, in the Palm version (which corresponds to conditioning upon the time $Z(t)$ enters U and making that the origin), $\{(X_n, Y_n) : n \geq 1\}$ is itself a stationary sequence; see Franken, König, Arndt and Schmidt [20], Brandt, Franken and Lisek [12] or Baccelli and Brémaud [9]. Let (X, Y) be generic random variables distributed as (X_n, Y_n) in the Palm version. Assume that $EX < \infty$ and $EY < \infty$. Let $F^c(x) = 1 - F(x)$ for any cumulative distribution function (cdf) F.

Theorem 1. *Assume that a proper steady-state distribution exists for the buffer content in the general stationary model above. Then*

$$P(B > x) \geq F^c(x) \equiv \frac{1}{(EX + EY)} \int_{x/\lambda}^{\infty} P(X > u)du \ . \tag{4}$$

Proof. A lower bound on the buffer content is obtained by assuming that the buffer content is 0 whenever $Z(t) \in D$. Let $T(t)$ be the time interval since the last transition to or from the set U before time t. This leads to the bound

$$P(B > x) \geq P(Z(0) \in U, \ T(0) > x/\lambda) \tag{5}$$

Let $\{(X_n, Y_n)\}$ be the stationary sequence in the Palm version. Stationary point process theory implies that

$$P(Z(t) \in U) = \frac{EX}{EX + EY} \tag{6}$$

and

$$P(Z(t) \in U, \ T(t) > x) = \frac{1}{EX + EY} \int_{x}^{\infty} P(X > u)du \ ; \tag{7}$$

e.g., see Section 4.3.3 of Brandt et al. [12]. Combining (5) and (7) yields (4). □

To apply Theorem 1 to establish asymptotic relations, we exploit the following elementary lemma; see p. 17 of Erdélyi [17].

Lemma 1. *If $f(x) \sim g(x)$ as $x \to \infty$, then $\int_x^\infty f(y)dy \sim \int_x^\infty g(y)dy$ as $x \to \infty$.*

We apply Lemma 1 to obtain the following corollary to Theorem 1 for the case of a power law.

Corollary 1. *In the setting of Theorem 1, if $P(X > x) \sim \beta x^{-(\eta+1)}$ as $x \to \infty$, where $\eta > 0$, then the lower bound tail probability in (4) satisfies*

$$F^c(x) \sim \frac{\beta \lambda^\eta}{\eta(EX + EY)} x^{-\eta} \text{ as } x \to \infty \ . \tag{8}$$

Theorem 1 indicates that we can detect conditions in source activity leading to a long-tail buffer-content distribution by estimating the tail behavior of the random variable X. We can estimate the distribution of X in the usual way by forming a histogram. It is significant that the successive variables X_n need *not* be mutually independent in such estimation and in the setting of Theorem 1. It is also significant that the distribution of Y plays no role.

In applying Theorem 1 there are many ways to choose the sets U and D. To make it most likely that the high-activity time X has a long tail, we should let λ be as small as possible, while still being positive, so that U includes all states with positive rates.

Even though Corollary 1 to Theorem 1 only yields a lower bound on the asymptotic tail behavior of $P(B > x)$ when $P(X > x) \sim \beta x^{-(\eta+1)}$ as $x \to \infty$, it lends support to the notion that (1) may hold, with the asymptotic decay rate η determined but the asymptotic constant α yet to be specified (with $\alpha > \beta \lambda^\eta / \eta(EX + EY)$). Additional support for this notion is provided by Section 3.

Corollary 1 extends to the setting of distributions with regularly varying tails, see p. 275 of Feller [19] and Bingham, Goldie and Teugels [10]. A real-valued function $L(x)$ is said to be *slowly varying* if it is a positive function such that $L(cx)/L(x) \to 1$ as $x \to \infty$ for each $c > 0$. A simple example is $\log x$. The following corollary can be proved by applying Theorem 1 on p. 281 of Feller [19].

Corollary 2. *In the setting of Theorem 1, if $P(X > x) = x^{-(\eta+1)}L(x)$ as $x \to \infty$, where $\eta > 0$ and $L(x)$ is slowly varying, then the lower bound tail probability in (8) satisfies.*

$$F^c(x) \sim \frac{\lambda^\eta}{\eta(EX + EY)} x^{-\eta} L(x) \quad as \ \ x \to \infty \ .$$

We can further extend Corollary 2 to the setting of subexponential distri-

6

butions. Let G_2 be the convolution of a cdf G with itself. A cdf G on the nonnegative real line is said to be *subexponential*, denoted by $G \in \mathcal{S}$ if

$$\frac{G_2^c(x)}{G^c(x)} \sim 2 \quad \text{as} \quad x \to \infty ; \tag{9}$$

e.g., see Appendix 4 of [10] and references there. Condition (9) can be understood as relating the distribution of the sum of two random variables to the distribution of the maximum; i.e., if X_1 and X_2 are independent random variables with common cdf G, then (9) is equivalent to

$$P(X_1 + X_2 > x) \sim P(\max\{X_1, X_2\} > x) \quad \text{as} \quad x \to \infty . \tag{10}$$

In order to treat integrals of tail probabilities, it is useful to consider another class of cdf's. The cdf G is said to belong to the class $\mathcal{S}^*$ (a subset of $\mathcal{S}$) if G has finite expectation m and

$$\int_0^x \frac{F^c(x-y)}{F^c(x)} F^c(y) dy \sim 2m \quad \text{as} \quad x \to \infty ; \tag{11}$$

see Klüppelberg [24]. The cdf's in $\mathcal{S}$ can be thought of as long tailed, because if $G \in \mathcal{S}$, then

$$e^{\eta x} G^c(x) \to \infty \quad \text{as} \quad x \to \infty \quad \text{for each} \quad \eta > 0 . \tag{12}$$

Klüppelberg [24] showed that if $G \in \mathcal{S}^*$, then $G_e \in \mathcal{S}$, where G_e is the associated stationary-excess cdf associated with G, i.e.,

$$G_e(x) = \frac{1}{m} \int_0^x G^c(u) du , \quad x \geq 0 . \tag{13}$$

Hence we have the following corollary, further showing how the long-tail character of the cdf of X is inherited by the buffer-content cdf.

Corollary 3. *In the setting of Theorem 1, if the cdf of X belongs to $\mathcal{S}^*$, where $\mathcal{S}^*$ is defined in (11), then $F \in \mathcal{S}$, where is $\mathcal{S}$ defined in (9) and F is the lower bound cdf in (4).*

We conclude this section by pointing out that Theorem 1 does not cover all the cases. There are other ways in which the steady-state buffer content B could have a long-tail distribution. Since we have made no independence assumptions, it is possible to have many successive low-activity times be very short. To take an extreme case, suppose that they are so short that they can be

regarded as having zero length. If the random number of these successive short low-activity periods has a long-tail, then even if the associated high-activity times are of constant length, the whole period will be like one high-activity period with a long-tail distribution. In summary, a long-tail buffer-content distribution can also be caused by periods with a long-tail distribution in which there are many high-activity subintervals, none of which is long, but nevertheless almost all of the time during that period the environment is in the high-activity state.

Consistent with this observation, we note that it can be difficult to fit an environment process to traffic data. If activity is measured in a very small time scale, the environment process can fluctuate very rapidly, producing very short high and low activity times, masking the effect Theorem 1 is intended to capture. Nevertheless, we believe that with judicious definitions Theorem 1 and its corollaries can help detect the need for large buffers.

3   A Two-State Fluid Model

The fluid model in Section 2 is very general, but for it we are only able to compute a lower bound on the tail of the steady-state buffer-content distribution. In this section we consider a restricted fluid model with a two-state environment process for which we can exactly compute the buffer-content distribution and derive asymptotics for its tail.

We assume that there is constant positive flow in at rate λ in the high-activity state and constant positive flow out at rate μ in the low-activity state. In addition, we assume that the successive high and low-activity times X_n and Y_n come from independent sequence of i.i.d. random variables. To have stability, we also assume that $\lambda EX < \mu EY$. As a technical regularity condition for getting a proper steady-state buffer-content distribution, we assume that X and Y have nonlattice distributions. (A lattice distribution concentrates on a countable periodic subset, e.g., $\{k\delta : k \geq 0\}$ for some δ.)

A fluid model with just two states may seem simplistic, but it can appear in practice. One example is a single access line with high and low rates, where the high access line rate is higher than the data rate on the backbone. This may happen in the context of a PC and modem, where the maximum data rate from the PC to the modem is higher than the modem output data rate.

A second example is a superposition of several constant-bit-rate (CBR) sources and one two-state variable-bit-rate (VBR) source, such that the combined data rates of the CBR sources and the low-activity data rate of the VBR source is below the backbone data rate, while the combined data rates of

the CBR sources and the high-activity data rate of the VBR source is above the backbone data rate. More generally, the two-state fluid model may serve as a rough approximation, which is worth studying because it can be fully analyzed.

With the extra conditions on the fluid model, we are able to calculate the steady-state buffer content distribution and its asymptote exactly. We draw on Kella and Whitt [23], where it is shown that the steady-state buffer-content distribution in this fluid model can be simply related to the steady-state virtual waiting-time distribution in the GI/G/1 queue with service times λX_n and interarrival times μY_n. (Also see [23] for previous related work.) As can be seen from Section 2, we will be interested in the case in which X has a long-tail distribution, but Y could be anything. For example, if Y is exponential, then the buffer-content distribution can be obtained by solving the M/G/1 queue. However, we are also able to solve the general GI/G/1 case.

By Corollary 3 of [23],

$$P(B > 0) = \left(\frac{EY}{EY + EX} \right) \left(\frac{\lambda + \mu}{\lambda} \right) \left(\frac{\lambda EX}{\mu EY} \right) \tag{14}$$

and the conditional buffer content given that it is nonempty satisfies

$$(B|B > 0) \stackrel{\mathrm{d}}{=} (V|V > 0) \stackrel{\mathrm{d}}{=} W + S , \tag{15}$$

where $\stackrel{\mathrm{d}}{=}$ denotes equality in distribution, V is the steady-state virtual waiting time and W is the steady-state actual waiting time in the GI/G/1 queue, and S is independent of W with the stationary-excess distribution of the service-time distribution, i.e.,

$$P(S > x) = \frac{1}{\lambda EX} \int_x^\infty P(\lambda X > u) du$$

$$= \frac{1}{EX} \int_{x/\lambda}^\infty P(X > u) du . \tag{16}$$

By (15), the Laplace transform of $(B|B > 0)$ is the product of the Laplace transforms of W and S, i.e.,

$$Ee^{-s(B|B>0)} = (Ee^{-sW})(Ee^{-sS}) . \tag{17}$$

9

By (11),

$$Ee^{-sS} = (1 - Ee^{-s\lambda X})/s\lambda EX . \tag{18}$$

Since the Laplace transform of W can be calculated and exactly inverted numerically [1], [2], so can the Laplace transform of $(B|B > 0)$, to yield $P(B > x|B > 0)$ for any desired x. Combining this with (14) yields the tail probabilities $P(B > x)$ themselves.

Abate, Choudhury and Whitt [4] already gave sufficient conditions for $P(V > x)$ to have an exponential tail and determined the asymptotic parameters, which applies to $P(B > x)$ by (17). (This extends the classical Cramér-Lundberg approximation for W discussed in [3].) Now we establish the precise power tail asymptotics for B assuming power tail asymptotics for the high activity times X_n.

Theorem 2. *For the two-state fluid model, if $P(X > x) \sim \beta x^{-(\eta+1)}$ as $x \to \infty$, then*

$$P(B > x) \sim \left(\frac{EY}{EY + EX}\right)\left(\frac{\lambda + \mu}{\lambda}\right)\left(\frac{\rho}{1 - \rho}\right)\frac{\beta\lambda^\eta x^{-\eta}}{\eta EX} \text{ as } x \to \infty ,$$

where $\rho \equiv \lambda EX/\mu EY$.

Proof. First, from (11), Lemma 1 and the assumed asymptotics for X,

$$P(S > x) \sim \frac{\beta\lambda^\eta x^{-\eta}}{\eta EX} \text{ as } x \to \infty .$$

Second, by the Corollary to Theorem 1 of [2] (due to Cohen [14] and Pakes [28]),

$$P(W > x) \sim \left(\frac{\rho}{1 - \rho}\right)\frac{\beta\lambda^\eta}{\eta EX}x^{-\eta} \text{ as } x \to \infty .$$

Hence, by (15) and Proposition (8.14) on p. 278 of Feller [19], which gives the asymptotic behavior of the convolution of two distributions with power tails,

$$P(B > x|B > 0) \sim \frac{\beta\lambda^\eta x^{-\eta}}{(1 - \rho)\eta EX} \text{ as } x \to \infty .$$

Combining this with (14) completes the proof. □.

Note that Theorem 2 is consistent with Theorem 1, yielding the *same asymptotic decay rate η* but a larger asymptotic constant. Theorems 1 and 2 together support the power-tail approximation $P(B > x) \approx \alpha x^{-\eta}$ when a high activity time X has power tail approximation $P(X > x) \approx \beta x^{-(\eta+1)}$. Theorems 1 and 2 yield the asymptotic decay rate η and provide guidance concerning approximations of the asymptotic constant α.

The power tail asymptote may well yield a satisfactory approximation, but we would suggest caution. Experience in [2] indicates that power tail asymptotics often do not provide especially good approximations. As support for this conclusion, note that the asymptote for $P(B > x)$ in Theorem 2 depends on the off-time distribution only through its mean, but the tail probability depends on the entire distribution. Thus, when we actually have the two-state fluid model, it is desirable to calculate the steady-state buffer content distribution exactly using numerical transform inversion, given the transforms of the high-activity and low-activity distributions, using the methods of [1], [2].

We conclude this section by giving some illustrative numerical examples. Without loss of generality, we assume that $EX = \lambda = 1$, which implies that the mean service time in the associated GI/G/1 queue is 1. We let X have a *Pareto mixture of exponential* (PME) distribution, as in [2]. (For other ways to construct long-tail distributions, see Abate and Whitt [6].) In particular, we consider the two cases

$$G_2^c(x) \equiv P(X > x) = \frac{1}{2x^2}(1 - (1 + 2x)e^{-2x}), \ x > 0 \ , \qquad (19)$$

and

$$G_4^c(x) \equiv P(X > x) = \frac{243}{32x^4}\left(1 - \left(1 + \frac{4x}{3} + \frac{8x^2}{9} + \frac{32x^3}{81}\right)e^{-4x/3}\right) \ , \quad x > 0 \ , (20)$$

with associated Laplace-Stieltjes transforms

$$\hat{g}_2(s) \equiv \int\limits_0^\infty e^{-st}dG_2(t) = 1 - s + \frac{s^2}{2}ln\left(1 + \frac{2}{s}\right) \qquad (21)$$

and

$$\hat{g}_4(s) = 1 - s + \frac{4}{3}s^2 - \frac{27}{16}s^3 + \frac{81}{64}s^4ln\left(1 + \frac{4}{3s}\right) \ . \qquad (22)$$

From (19) and (20), we see that

$$G_2^c(x) \sim \frac{1}{2x^2} \quad \text{and} \quad G_4^c(x) \sim \frac{243}{32x^4} \text{ as } x \to \infty \ . \qquad (23)$$

From (22) and Theorem 2, we see that $P(B > x)$ is asymptotic to $\alpha_2 x^{-1}$ and $\alpha_4 x^{-3}$ as $x \to \infty$ for appropriate constants α_2 and α_4 in the two cases.

For the rest of the model, we let the distribution of the low activity time Y be exponential. This leaves two parameters: (1) the mean low activity time, EY,

and (2) $\rho = 1/\mu EY$. Since Y has an exponential distribution, the associated queueing model is M/G/1. Hence we can apply the numerical transform inversion algorithm in Abate and Whitt [6] to the transform in (12), using the classical Pollaczek-Khintchine expression for the transform Ee^{-sW}. However, we could also allow the distribution of Y to be nonexponential (e.g., long-tail) by using an expression for Ee^{-sW} in the GI/G/1 setting, as in [1], [2].

We give three examples for $G_2(x)$ in Table 1: ($\rho = 0.8$, $EY = 1$), ($\rho = 0.8$, $EY = 1000$) and ($\rho = 0.3$, $EY = 1$). We give two examples for $G_4(x)$ in Table 2: ($\rho = 0.8$, $EY = 1$) and ($\rho = 0.3$, $EY = 1$). In all cases we give the power tail asymptotic approximation provided by Theorem 2 as well as the exact value computed by numerical transform inversion. In both cases, $P(B > x) \sim \alpha x^{-\eta}$ as $x \to \infty$, but $\eta = 1$ for $G_2^c(x)$ while $\eta = 3$ for $G_4^c(x)$. We display the asymptotic constant α in each case in Tables 1 and 2 too. The exact probability $P(B > 0)$ is given for $x = 0$; the asymptotic approximation for $x = 0$ is infinity. As in [2], we see that the asymptote underestimates the exact tail probabilities for x suitably large, i.e., for $x \geq 10$. The tail probabilities clearly decay more rapidly with $G_4^c(x)$ than with $G_2^c(x)$. The asymptote is not a good approximation until x becomes large, but if interest is in very small tail probabilities, then the asymptotic approximation can be very useful.

For the current model, it is easy to compute both the asymptotic parameters α and η. However, in a more general setting, η is easier to determine than α (η is just one less than the η for the associated on-time distribution). Hence, a more elementary approximation would be to use the true value of η and set $\alpha = 1$. For our examples, the quality of this cruder approximation can be ascertained by looking at how the actual value of α differs from 1.

4 The Effect of Multiplexing

In Section 2 we saw that the holding time X in the set U having a long-tail distribution implies that the steady-state buffer-content distribution is bounded below by a long-tail distribution. We now want to identify sufficient conditions in terms of individual sources for the buffer content distribution to have a long-tail distribution when there is statistical multiplexing.

We assume that the input comes from the superposition of independent sources and that there is a fixed output rate (the bandwidth). Paralleling our assumption for Theorem 1, we assume that the sequence of successive high and low activity times for each source is well defined, so that we can focus on the stationary sequence by looking at the Palm version. Let $X^{(j)}$ and $Y^{(j)}$ be generic high activity and low activity times for source j.

Theorem 3. *In the multi-source setting above, with the sources being inde-*

pendent, suppose that there are m sources such that the environment process $Z(t) \in U$ whenever all m sources are simultaneously in high activity states. If a proper steady-state distribution exists for the buffer content, then

$$P(B > x) \geq F^c(x) \equiv \prod_{j=1}^{m} \frac{1}{(EX^{(j)} + EY^{(j)})} \int_{x/\lambda}^{\infty} P(X^{(j)} > u) du \ . \quad (24)$$

Proof. Let j index the m sources. Let U_j denote a high activity subset of states for source j; let $Z_j(t)$ be the environment process for source j; and let $T_j(t)$ be the time interval since the last transition to the set U_j before time t. Let U now denote the set of states in which all m designated sources are in high activity states and let $T(t)$ be the time before t since the last transition into a high activity state by any one of the m sources. Then, by the proof of Theorem 1,

$$P(B > x) \geq P(Z(0) \in U, T(0) > (x/\lambda)) \ , \quad (25)$$

just as in (5). However,

$$\{Z(0) \in U \ , \ T(0) > x/\lambda\} = \bigcap_{j=1}^{m} \{Z_j(0) \in U_j \ , \ T_j(0) > x/\lambda\} \quad (26)$$

and, by the assumed independence among sources,

$$P\left(\bigcap_{j=1}^{m} \{Z_j(0) \in U_j \ , \ T_j(0) > x/\lambda \right) = \prod_{j=1}^{m} P(Z_j(0) \in U_j \ , \ T_j(0) > x/\lambda) \ , (27)$$

which coincides with (24) by Theorem 1. ☐

We now state a corollary to Theorem 3 paralleling Corollary 1 to Theorem 1. (Analogs of Corollaries 2 and 3 to Theorem 1 also hold.)
Corollary. *In the setting of Theorem 3, if*

$$P(X^{(j)} > x) \sim \beta_j x^{-(\eta_j+1)} \quad as \ \ x \to \infty \ , \quad (28)$$

where $\beta_j > 0$ and $\eta_j > 0$ for $1 \leq j \leq m$, then the lower bound tail probability in (24) satisfies

$$F^c(x) \sim \lambda^{\eta} \left(\prod_{j=1}^{\infty} \frac{\beta_j}{\eta_j(EX^{(j)} + EY^{(j)})} \right) x^{-\eta} \quad as \ \ x \to \infty \ . \quad (29)$$

13

where $\eta = \sum_{j=1}^{m} \eta_j$.

Theorem 3 has some important consequences. Note that $\eta = \sum_{j=1}^{m} \eta_j$. The special case with $m = 1$ implies that if one source can yield positive net input, then a long-tail distribution for the high activity time of that one source will be inherited by the entire system; i.e., then multiplexing cannot remove the long-tail effect.

On the other hand, if a positive net input rate requires high activity from many sources simultaneously (as we expect in most applications), then the buffer-content tail probability may decay much faster than the tail probability of the high-activity period for one source; i.e., there may be significant statistical multiplexing gains, even in the setting of power tails. Indeed it is also possible that there are a few access lines (sources) with long-tailed high-activity-time distributions but their combined high-activity rates do not exceed the backbone rate, thereby not showing any long-tail effect on the buffer-content distribution. In this case, statistical multiplexing (and providing the higher backbone rate for multiple sources) would greatly improve performance.

5   Upper Bounds for Fluid Models

We can modify the procedure in Section 2 to obtain upper bounds for multi-state fluid models, under extra assumptions. We start by dividing the environment states into two subsets: $U \equiv \{i : -\mu < r_i \leq \lambda\}$ and $D \equiv \{i : r_i < -\mu < 0\}$. Again, it is elementary that the fluid content process is bounded, this time *above*, for each sample path, by the fluid content process with the two-state environment process with net flow rate λ in U and net flow $-\mu$ in D.

As in Section 2, let X_n and Y_n be the successive holding times in the sets U and D, respectively. As before, in the Palm version $\{(X_n, Y_n)\}$ is a stationary sequence. Now, however, we need to assume more. In particular, we assume that $\{X_n\}$ and $\{Y_n\}$ are independent sequences of i.i.d. random variables.

Under these strong independence conditions, the upper bound process has the structure of the GI/G/1 queue, just as in Section 3. If, in addition, $\lambda EX < \mu EY$, then the bounding buffer-content process is stable. Moreover, assumptions on X will yield upper bounds on the buffer-content-distribution tail.

Theorem 4. *Suppose that the independence conditions above hold, $\lambda EX < \mu EY$, and that X and Y have nonlattice distributions. Then the bounding buffer content has a proper steady-state distribution.*

(a) If

$$Ee^{\eta(\lambda X - \mu Y)} = 1 \quad and \quad Ee^{(\eta+\epsilon)(\lambda X - \mu Y)} < \infty \qquad (30)$$

for some $\eta > 0$ and $\epsilon > 0$, then

$$P(B > x) \le \alpha e^{-\eta x} \quad for\ all \quad x > 0$$

for positive constants α and η, with η satisfying (30).

$$(b) \quad If \quad P(X > x) \sim \beta x^{-(\eta+1)} \quad as \quad x \to \infty \ ,$$

for positive constants β and η, then

$$P(B > x) \le \alpha x^{-\eta} \quad for\ all \quad x$$

for η and positive constant α.

Proof. As before, we apply Kella and Whitt [23] to relate the bounding process to the virtual waiting time in the GI/G/1 queue. For part (a), we then apply Theorem 2 of Abate, Choudhury and Whitt [4]. Thus, asymptotics for the bounding process corresponds to the classical Cramér-Lundberg approximation; see [3] and references therein. For part (b), we apply Theorem 2 here.
□

It should be noted that the conditions in Theorem 4 are quite restrictive. First, the independence assumption in this section is a strong requirement. It would be satisfied by many Markov-modulated fluid models, but it is a restriction. Also, the bounding model might well be unstable ($\lambda EX \ge \mu EY$) even though the original model is stable.

6   Conclusion

The results here can be summarized by saying that long-tail sojourn-time distributions of individual source high-activity regions or in system high-activity regions lead to a long-tail distribution in the steady-state buffer content. The sojourn-time distributions in low-activity regions do not play any role in determining whether there is a long-time buffer-content distribution. This conclusion is intuitively reasonable, but worth careful analysis. The result is not surprising given previous results for classical queueing models: Long-tail distributions for service times lead to long-tail distributions for delays; see [2] and references cited there. For the case of a two-state fluid model with independent high and low activity times, we exploit a connection between the queueing models and fluid models established previously in [23] to derive the

asymptotic behavior of the tail probabilities and calculate them exactly. We obtain important insights about the effect of statistical multiplexing, but more work needs to be done on that topic.

References

[1] J. Abate, G. L. Choudhury and W. Whitt, Calculation of the GI/G/1 waiting time distribution and its cumulants from Pollaczek's formulas, *Archiv für Elektronik und Übertragungstechnik* 47 (1993) 311–321.

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

[3] J. Abate, G. L. Choudhury and W. Whitt, Exponential approximations for tail probabilities in queues, I: waiting times, *Opns. Res.* 43 (1995) 885–901.

[4] J. Abate, G. L. Choudhury and W. Whitt, Exponential approximations for tail probabilities in queues, II: sojourn time and workload, *Opns. Res.*, (1996) to appear.

[5] J. Abate and W. Whitt, The Fourier-series method for inverting transforms of probability distributions, *Queueing Systems* 10 (1992) 5–88.

[6] J. Abate and W. Whitt, An operational calculus for probability distributions via Laplace transforms, *Adv. Appl. Prob.* 28 (1996), 75–113.

[7] D. Anick, D. Mitra and M. M. Sondhi, Stochastic theory of a data-handling system with multiple sources, *Bell System Tech. J.* 61 (1982) 1871–1894.

[8] S. Asmussen, L. F. Henriksen and C. Klüppelberg, Large claims approximations for risk processes in a Markovian environment, *Stoch. Proc. Appl.* 54 (1994) 29–43.

[9] F. Baccelli and P. Brémaud, *Elements of Queueing Theory*, Springer, Verlag, 1994.

[10] N. H. Bingham, C. M. Goldie and J. L. Teugels, *Regular Variation*, Cambridge University Press, Cambridge, England, 1989.

[11] J. Boxma, Fluid queues and regular variation, CWI, Amsterdam, 1996.

[12] A. Brandt, P. Franken and B. Lisek, *Stationary Stochastic Models*, Wiley, New York, 1990.

[13] C.-S. Chang and J. A. Thomas, Effective bandwidth in high-speed digital networks, *IEEE J. SAC* 13 (1995) 1091–1100.

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

[15] N. G. Duffield, Economies of scale in queues with sources having power-law large deviations scalings, *J. Appl. Prob.*, to appear.

[16] N. G. Duffield and N. O'Connell, Large deviations and overflow probabilities for the general single-server queue, with applications, *Math. Proc. Camb. Phil. Soc.*, 118 (1995) 363–374.

[17] A. Erdélyi, *Asymptotic Expansions*, Dover, New York, 1956.

[18] A. Erramilli, Application of fractals in engineering for realistic traffic processes, *Proceedings 14th Int. Teletraffic Congress*, 1994, 35–44.

[19] W. Feller, *An Introduction to Probability Theory and Its Applications*, vol. II, second ed., Wiley, New York, 1971.

[20] P. Franken, D. König, U. Arndt and V. Schmidt, *Queues and Point Processes* Akademie-Verlag, Berlin, 1981.

[21] P. R. Jelenković and A. A. Lazar, Subexponential asymptotics of a network multiplexer, Proceedings fo the thirty-third Annual Allerton Conference on Communication, Control and Computing, Urbana IL, 1995.

[22] P. R. Jelenković and A. A. Lazar, Subexponential asymptotics of a Markov-modulated G/G/1 queue, Department of Electrical Engineering, Columbia University, 1995.

[23] O. Kella and W. Whitt, A storage model with a two-state random environment, *Opns. Res.* 40 (1992) S257–S262.

[24] C. Klüppelberg, Subexponential distributions and integrated tails, *J. Appl. Prob.* 25 (1988) 132–141.

[25] W. E. Leland, M. S. Taqqu, W. Willinger and D. V. Wilson, On the self-similar nature of ethernet traffic, *ACM SIGCOMM'93*, (1993), 183–193.

[26] K. Meier-Hellstern, P. Wirth, Y.-L. Yan and D. Hoeflin, Traffic models for ISDN data users: office automation applications, *Teletraffic and Data Traffic in a Period of Change, ITC 13*, A. Jensen and B. Iversen (eds.), Elsevier, Amsterdam, 1991, 167–172.

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

[28] A. G. Pakes, On the tails of waiting-time distributions, *J. Appl. Prob.* 12 (1975) 555–564.

[29] V. Paxson and S. Floyd, Wide area traffic: the failure of Poisson modelling, IEEE/ACM *Trans. Networking* 3 (1995) 226–244.

[30] J. W. Roberts, *Performance Evaluation and Design of Multiservice Networks*, Cost 224 Project Final Report, Brussels, 1992.

[31] P. Sen, B. Maglaris, N. Rikli and D. Anastassiou, Models for packet switching of variable-bit-rate video sources, *IEEE J. Sel. Areas Commun.* 7 (1989) 865–869.

[32] G. deVeciana, G. Kesidis and J. Walrand, Resource management in wide-area ATM networks using effective bandwidths, *IEEE J. SAC* 13 (1995) 1081–1090.

[33] W. Whitt, Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues, *Telecommunication Systems* 2 (1993) 71–107.

[34] W. Willinger, Traffic modeling of high-speed networks: theory and practice, *Stochastic Networks*, F. B. Kelly and R. Williams (eds.), Springer-Verlag, New York, 1994.

[35] W. Willinger, M. S. Taqqu, W. E. Leland and D. V. Wilson, Self-similarity in high-speed packet traffic: analysis and modeling of Ethernet traffic measurements, *Statistical Science* 10 (1995) 67–85.

Gagan L. Choudhury received the B.Tech. degree in Radio Physics and Electronics from the University of Calcutta, India in 1979 and the MS and Ph.D. degrees in Electrical Engineering from the State University of New York (SUNY) at Stony Brook in 1981 and 1982, respectively. Currently he is a Technical Manager at AT&T Labs, Holmdel, New Jersey, USA. His group is responsible for the performance assessment and improvement of various AT&T products and services based on modeling and simulation. His main research interest is in stochastic modeling and its application to the performance analysis of telecommunication and computer systems.

Ward Whitt received the A.B. degree in Mathematics from Dartmouth College, Hanover, NH, USA, in 1964 and the Ph.D. degree in Operations Research from Cornell University, Ithaca, NY, USA, in 1969. He was on the faculty of Stanford University and Yale University before joining AT&T in 1977. He is currently a Technical Consultant in the Network Mathematics Research Department of AT&T Labs in Murray Hill, NJ, USA. His research has focused on probability theory, queueing models and performance analysis.

level	$\rho = 0.8$, $EY = 1$ $\alpha = 2.25$, $\eta = 1$		$\rho = 0.8$, $EY = 1000$ $\alpha = 2.001$, $\eta = 1$		$\rho = 0.3$, $EY = 1$ $\alpha = 0.464$, $\eta = 1$	
x	exact	approx.	exact	approx.	exact	approx.
0	.900	∞	.800	∞	.650	∞
10	.315	.225	.280	.200	.0513	.0464
20	.175	.113	.156	.100	.0249	.0232
30	.115	.0750	.102	.0667	.0163	.0155
40	.0825	.0563	.0734	.0500	.0121	.0116
50	.0634	.0450	.0564	.0400	.00963	.00929
100	.0278	.0225	.0247	.0200	.00474	.00464
200	.0127	.0113	.0113	.0100	.00235	.00232
300	.00816	.00750	.00725	.00667	.00156	.00155
400	.00600	.00563	.00534	.00500	.00117	.00116
500	.00475	.00450	.00422	.00400	.000934	.000929
1000	.00232	.00225	.00206	.00200	.000466	.000464
2000	.00114	.00113	.00102	.00100	.000233	.000232
3000	.000758	.000750	.000674	.000667	.000155	.000155
4000	.000567	.000563	.000504	.000500	.000116	.000116
5000	.000453	.000450	.000403	.000400	.0000929	.0000929

Table 1. A comparison of exact buffer-content tail probabilities $P(B > x)$ with the asymptotic approximations $\alpha x^{-\eta}$ for the single-source model with on-time distribution $G_2(x)$ in (19).

| level | $\rho = 0.8,\ EY = 1$ | | $\rho = 0.3,\ EY = 1$ | |
| | $\alpha = 11.39,\ \eta = 3$ | | $\alpha = 2.35,\ \eta = 3$ | |
x	exact	approx.	exact	approx.
0	.900	∞	.650	∞
10	.153	.114e-01	.405e-02	.235e-02
20	.295e-01	.142e-02	.370e-03	.294e-04
30	.616e-02	.422e-03	.992e-04	.871e-04
40	.144e-02	.178e-03	.402e-04	.367e-04
50	.405e-03	.911e-04	.201e-04	.188e-04
100	.163e-04	.114e-04	.243e-05	.235e-05
200	.166e-05	.142e-05	.298e-06	.294e-06
300	.465e-06	.422e-06	.879e-07	.871e-07
400	.191e-06	.178e-06	.370e-07	.367e-07
500	.964e-07	.911e-07	.189e-07	.188e-07
1000	.117e-07	.114e-07	.236e-08	.235e-08
2000	.144e-08	.142e-08	.294e-09	.294e-09

Table 2. A comparison of exact buffer-content tail probabilities $P(B > x)$ with the asymptotic approximation $\alpha x^{-\eta}$ for the single-source model with on-time distribution $G_4(x)$ in (20).