BUFFER PROBLEMS IN TELECOMMUNICATIONS NETWORKS

.

Lester R. Lipsky and John E. Hatem

Department of Computer Science and Engineering
University of Connecticut
Storrs, CT 06269-3155
lester@brc.uconn.edu .
and
.
The Taylor L. Booth Center for Computer Applications and Research
233 Glenbrook Road, U-31
University of Connecticut
Storrs, Connecticut 06269-3031
johnh@brc.uconn.edu .

March 14, 1996

ABSTRACT

We present new exact formulas for calculating buffer overflow probabilities for various single server queues. We then present the results of a parametric study of the buffer size needed to prevent overflow or loss in systems where data arrivals are "bursty", "self-similar", or "fractal". Such erratic behavior can be caused (or adequately described) by renewal processes whose interarrival distributions are power-tail (or Pareto, or Lèvy, or "long-tail") with infinite variance. We compare the behavior of such systems with that of better behaved systems, namely where the interarrival times, or service times, have hyperexponential-2 or Erlangian distributions. We show that power tails can cause problems for intermediate values of the utilization parameter, ρ, and become very serious (beyond the usual $1/(1-\rho)$ factor) when ρ is close to 1, and/or when α approaches 1. Since various researchers have reported data with α values between 1.1 and 1.5, this may prove to be serious indeed.

Buffer Problems in Telecommunications Networks

Lester R. Lipsky and John E. Hatem

.

Department of Computer Science and Engineering

University of Connecticut

Storrs, CT 06269-3155

.

and

.

The Taylor L. Booth Center for Computer Applications and Research

233 Glenbrook Road, U-31

University of Connecticut

Storrs, Connecticut 06269-3031

1 Introduction

In recent years there has been an ever increasing interest in the development of systems which will be able to process incoming traffic from various communications networks. In the last three years, numerous papers have appeared indicating that the traffic to be expected in the future will have an extraordinary character. For instance, Leland, et.al., [LELA94] have measured and analyzed the arrival of many millions of packets on ETHERNET networks at Bellcore, while Beran, et.al., [BERA95] have measured and analyzed several millions of frames from Variable-Bit-Rate (*VBR*) video services. The data displayed in both papers (which are typical of the many appearing lately) show enormous instability of arrival rates. No matter how large the measurement interval, the number of arrivals per unit time varies widely. This has been described as "self-similar", "bursty", and "fractal" behavior. These, and other papers, have argued that $r(k)$, the auto-correlation function lag-k of the number of arrivals per time interval (not to be confused with the autocorrelation function for the inter-arrival times themselves), must go to zero so slowly that $\sum_{k=1}^{\infty} r(k) = \infty$. They imply that any realistic model of such traffic must include very long-range correlation effects.

The data can be explained more simply, however. Lipsky and Fiorini [LIPS95] have shown that a renewal process (no correlation of interarrival times) where the interarrival times have a power-tail (*PT*) distribution (i.e., distribution functions which behave as $1 - F(x) \Rightarrow 1/x^\alpha$ for large x) can have such autocorrelations. Furthermore, Van de Liefvoort and Weng have generated "self-similar" data of the kind described in [LELA94] by simulating a renewal process where the interarrival times come from a single PT distribution with a finite mean (but *infinite variance*) [LIEF94]. Many attempts to do the same using various compound Poisson processes (without power tails, but with built-in burstiness) have *not* been successful. Greiner, Jobmann and Lipsky [GREI95] describe in detail the properties of PT distributions, defining an analytic class of well-behaved distributions (a sub-class of which are *Phase Distributions* which can be used in Markov Chain models) that have truncated power tails (*TPT*), and in the limit become PT distributions. This class was first used in [LIPS86] to explain the long-tail behavior of measured CPU times at Bellcore in 1986 [LELA86]. It was also used to show what might happen in data-retrieval systems which have power-tail file sizes [GARG92], and even to explain the distribution of medical insurance claims [LOWR93]. Greiner, et.al. [GREI95] then used these distributions to study the behavior of steady-state GI/M/1 queues, as a model for telecommunications networks. They described the effects of different α's on the *geometric parameter, s* [LIPS92] as a function of the *utilization parameter*, ρ, where

$$\rho := [arrival\ rate] \cdot [mean\ service\ time]. \tag{1}$$

The variance of a PT distribution is infinite if $\alpha \leq 2$. Their calculations showed that s stays close to 1 as ρ decreases from 1, and drops off ever more slowly for smaller α. Note that the mean queue length for GI/M/1 queues is proportional to $1/(1-s)$. Thus steady-state performance of these queues becomes worse gradually as α drops below 2 with ρ fixed, becoming disastrous as α approaches 1 from above (i.e., the mean still exists).

We maintain that the affects of noisy/bursty traffic can be modelled adequately with PT distributions, and most if not all of the pathological behavior which might occur in real systems will be reflected, at least qualitatively, in an appropriate queueing model where the arrival process is a renewal process. In this paper, we present new exact formulas for calculating the buffer size needed to prevent excessive overflow, or loss in steady-state GI/M/1, GI/M/1/N, M/G/1, and M/G/1/N queueing systems. We also show how PT distributions can be incorporated into these models. We then present the results of a parametric study of the effects of PT and other distributions on buffer overflow.

We expected to find that pathological behavior caused by PT (and other large variance) distributions would be widespread. Instead, we found that buffer overflow [aside from the $1/(1-\rho)$ factor], though much larger than the M/M/1 queue, can be kept under control unless α is less than 1.5. The data of [LELA94] indicated an α of approximately 1.4, so serious buffer problems can be expected in the future.

We also found that two-phase hyperexponential distributions with large coefficients of variation, C^2, comparable to an equivalent truncated tail distribution, yielded wildly inconsistent results. (H_2 distributions are 3-parameter functions, two fixed by the mean and variance. Performance varies drastically on choice of the third parameter.)

Of course, all this was done assuming steady-state behavior. But this may require inordinately many arrivals before such large queue lengths could be seen in reality. Discrete event simulation models must necessarilly suffer from the same problem. [GREI95] presented an argument showing that the closer α is to 1, the more arrivals must occur before any system's steady state can be approached. It remains for the future to yield appropriate descriptions of transient behavior.

2 The Basic Models

Our system is made up of a single processor receiving data from an arrival stream of variable-sized packets. The arrival stream may be considered a renewal process. The receiving processor has a finite primary memory buffer which can hold at most N packets. If a no-loss system is required, then we assume there exists an unbounded secondary or *backup-buffer* that will store the overflow (e.g., a *disc-array sub-system*), and transfer the data to the primary buffer when space becomes available. The assumption of (almost) "infinite buffer" is not unreasonable, given the emerging technologies for fairly highspeed massive storage. We will show presently that for GI/M/1 queues under heavy load, if the primary buffer is large enough so that only 1% of arriving packets will have to be placed in the backup, then a backup buffer that is k times the size of primary will overflow only one time in 10^{2k}. Although there may be many problems associated with the transfer (e.g., loss of first-come-first-served sequencing, extra processors needed), we assume that the transfer can be made at least as quickly as it takes to drain the primary buffer, so there is no change in effective service rate. This is, then, a GI/G/1 open queueing system. If there is no backup buffer, then there must be losses, and we have a GI/G/1/N system.

We will assume that either "GI" or "G" is exponential, yielding a total of four different types of queues. First we will assume that packet arrivals can be considered a general renewal process, where each packet must be serviced in a time taken from an exponential distribution with mean time $1/\lambda$ (a GI/M/1 queue). If no backup buffer is provided, then we have a GI/M/1/N queue. In an alternate view (see [LIKH95]) a Poisson process with a "disbursed" batch of packets whose number is distributed by a power tail, can also generate self-similar data. If the packets can be "reassembled" at the receiving node and counted as one customer whose service time is taken from a PT distribution then we have an M/G/1 queue, or an M/G/1/N queue if there is no backup buffer.

2.1 Properties of Power-tail Distributions

These distributions are thoroughly described in [GREI95]. A summary is given here. A *Probability Distribution Function* (*PFD*), for some random variable, X, is defined as:

$$F(x) := \mathrm{Prob}(X \leq x),$$

while its *Reliability Function* is given as

$$R(x) := \mathrm{Prob}(X > x) = 1 - F(x).$$

Also, if it exists, the *probability density function* (*pdf*) is defined as

$$f(x) := \frac{dF(x)}{dx} = -\frac{R(x)}{dx}.$$

A *Power-tail Distribution* can be defined by its behavior for very large x. That is, if

$$R(x) \longrightarrow \frac{c}{x^\alpha}, \tag{2}$$

then $R(x)$ [or $F(x)$] is a PT distribution with power α, where α is a positive, real number. From elementary calculus it is easy to show that if $\alpha \leq 1$ then the distribution has an infinite mean. If $1 < \alpha \leq 2$ then $F(x)$ has a finite mean, but an infinite variance. In general,

$$\mathrm{E}(X^\ell) := \int_0^\infty x^\ell f(x)\, dx = \infty \quad \forall \quad \ell \geq \alpha. \tag{3}$$

Such distributions have been known to exist for a very long time. Pareto used them in describing the distribution of wealth in economics. Lèvy showed that all *stable* distributions with infinite variance have power tails. Thus they are also referred to as Pareto-, or Lèvy-Pareto-, or simply Lèvy Distributions in the literature for various disciplines. For a more complete discussion, the reader is referred to William Feller's Book [FELL71], or [GREI95]. These distributions have been ignored in computer science and related fields in the past because an extremely large number of events (a number of the order of 10^7 would not be very large) must occur for the tail to be felt. What does it mean for a model to predict a steady-state mean queue length of, say 10,000 customers, when there are hardly that many customers in the user community? Only now, with the imminent arrival of the *information super-highway* (and the world-wide web) can we expect to see so many events (customers - packets) in a relatively short time.

In general, simple PT distributions (the one used by Pareto was of the form: $f(x) = c\,x^{\mu-1}/(1+x)^{\alpha+\mu}$) are difficult to use for Laplace transforms, and do not have direct matrix representations. But a most useful sub-class of them is given in [GREI95]. The particular one we use here is defined as:

$$R_M(x) = \frac{1-\theta}{1-\theta^M} \sum_{n=0}^{M-1} \theta^n \, \exp(\mu x/\gamma^n), \tag{4}$$

where θ and γ are parameters satisfying the inequalities: $0 < \theta < 1$ and $\gamma > 1$. It is not hard to show that the ℓ^{th} moments are given by

$$\mathrm{E}(X_M^\ell) = \frac{1-\theta}{1-\theta^M} \cdot \frac{1-(\theta\gamma^\ell)^M}{1-\theta\gamma^\ell} \cdot \frac{1}{\mu}. \tag{5}$$

Next define their limit function

$$R(x) := \lim_{M\to\infty} R_M(x) = (1-\theta) \sum_{n=0}^{\infty} \theta^n \, \exp(\mu x/\gamma^n). \tag{6}$$

Then it can be shown that $R(x)$ satisfies (2), and α is related to θ and γ by

$$\theta\gamma^\alpha = 1\,, \quad \text{or} \quad \alpha := -\frac{\log(\theta)}{\log(\gamma)}\,. \tag{7}$$

It then follows that

$$\mathrm{E}(X) := \lim_{M\to\infty} \mathrm{E}(X_M^\ell) = \infty \ \text{ for } \ \ell \geq \alpha.$$

That is, Equation (3) is satisfied. We refer to the functions, $R_M(x)$ as *truncated power-tail (TPT) distributions*, because, depending on the size of M, they look like their limit function, the true power-tail, $R(x)$. But for some large x, depending upon M, they drop off exponentially. These functions are easy to manipulate algebraically. Furthermore, they are $M-$dimensional *phase* distributions whose vector-matrix representations , $< \mathbf{p}_\mathrm{M}, \mathbf{B}_\mathrm{M} >$ are given by (using the notation of [LIPS92]):

$$\mathbf{B}_\mathrm{M} = \mu \begin{bmatrix} 1 & 0 & 0 & \vdots & 0 \\ 0 & \delta & 0 & \vdots & 0 \\ 0 & 0 & \delta^2 & \vdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \vdots & \delta^{M-1} \end{bmatrix} \quad \text{and} \quad \mathbf{p}_\mathrm{M} = \frac{1-\theta}{1-\theta^M}[1 \ \theta \ \theta^2 \ \cdots \ \theta^{M-1}]. \tag{8}$$

where $\delta := 1/\gamma$. We need these matrices to calculate the properties of finite-buffer queues. They generate the functions given above by the relations

$$R_M(x) = \Psi_M[\exp(-x\mathbf{B}_\mathrm{M})]$$

where for any square matrix, $\mathbf{Z}$

$$\Psi_M[\mathbf{Z}] := \mathbf{p}_\mathrm{M}\, \mathbf{Z}\, \epsilon'. \tag{9}$$

ϵ' is the column vector with all 1's. The general method of representing processes by matrix operators is called *Linear Algebraic Queueing Theory* (LAQT) by [LIPS92]. We define

$$\mathbf{V}_\mathrm{M} := \mathbf{B}_\mathrm{M}^{-1},$$

then we get

$$\mathrm{E}(X_M^\ell) = \ell!\, \Psi_M[\mathbf{V}_\mathrm{M}^\ell].$$

Furthermore, the Laplace transform of $F_M(\cdot)$ is

$$B_M^*(s) := \int_0^\infty e^{-sx}\, f_M(x)\, dx = \Psi_M[(\mathbf{I} + s\mathbf{V}_\mathrm{M})^{-1}]$$

Note that the matrix $\mathbf{B}$, representing $R(x)$, is infinite dimensional, and has an infinite set of eigenvalues, $\{\beta_n\}$, with an accumulation point at 0. So, in principle, its inverse does not exist. But with judicious use of limits, all calculations can be carried out.

2.2 GI/M/1 Queues - No Packets Lost

Here we assume that the time to process a packet is exponentially distributed, with mean $1/\lambda$, and the buffer is unbounded in size (bigger than we'll ever need). The packet arrivals constitute a renewal process, with interarrival-time distribution $F(x)$. As already discussed, this constitutes an open GI/M/1 queue. It is well known that the steady-state probability for finding k customers in the queue, $\pi(k)$, is given by [LIPS92]

$$\begin{aligned} \pi(0) &= 1 - \rho \\ \pi(k) &= (1-s)\cdot\rho\cdot s^{k-1}, \quad k > 0 \end{aligned}$$

where s is the *geometric parameter* satisfying the equation

$$s = B^*[\lambda(1-s)]. \tag{10}$$

$B^*(z)$ is the Laplace transform of the interarrival distribution. Alternatively, in the LAQT representation, s is the smallest eigenvalue of the matrix

$$\mathbf{A} := \mathbf{I} + \frac{1}{\lambda}\mathbf{B} - \mathbf{Q}. \tag{11}$$

Let $\bar{x}$ be the mean interarrival time. Then

$$\rho = \frac{1}{\lambda \bar{x}}$$

and the mean queue length (including the one being served) for the process is

$$\bar{q} := \sum_{k=0}^{\infty} k \cdot \pi(k) = \frac{\rho}{1-s} \cdot$$

What is needed here is the probability that an arriving packet will find exactly $k > 0$ other packets already there. The arrival probabilities are given by

$$a(k) = (1-s) \cdot s^k = \pi(k) \cdot \frac{s}{\rho} \cdot$$

Then the probability that an arriving packet will have to be stored in the backup buffer is

$$\Pr(N) = \sum_{k=N}^{\infty} a(k) = (1-s) \sum_{k=N}^{\infty} s^k = s^N.$$

We see that the smaller s is, the less likely overflow will occur. Equivalently, the closer s is to 1, the bigger $\bar{q}$ and $\Pr(N)$ will be, giving less desirable performance.

There are some general statements one can make. For instance, when $\rho = 1$, so does s. If $R(0) = 1$ (a *non-defective* distribution) then $s = 0$ when $\rho = 0$. Also, only for the M/M/1 queue does $s = \rho$ for all ρ. We say that if $s > \rho$ then system performance is worse, and if $s < \rho$ then system performance is better than one could ask for. It has been shown [LIPS92] that the slope of the curve, s versus ρ at $\rho = 0$ is $\bar{x} f(0)$. So if this is less than (greater than) 1, then for small ρ, performance is better (worse) than the equivalent M/M/1 queue. At the other end, at $\rho = 1$, the slope is $2/(C^2+1)$. If $C^2 > 1$ ($C^2 < 1$) then performance is worse (better). It is also known [LIPS92] that near $\rho = 1$ performance depends only on the moments of the interarrival time distribution, and thus on α and θ. In particular, if $\alpha \leq 2$ then $C^2 = \infty$ and the slope is 0. This means that s will remain close to 1 even as ρ decreases.

In general, for small ρ performance depends only on the behavior of $R(x)$ when x is small. For instance, in modelling PT distributions using (6) with $\mu = (1-\theta)/(1-\theta\gamma)$, it follows that $f_{\mathrm{M}}(0) > 1$ for all θ and all α. A different function (other than $e^{-\mu x}$) could have been chosen which would have yielded a smaller s for small ρ (e.g., $\mu^2 x e^{-\mu x}$ instead of $e^{-\mu x}$). But the performance for $\rho \to 1$ would be the same. This shows the difficulty in selecting test functions in exploring the general performance of systems. Without more knowledge of a particular system, no model can be relied upon to give an accurate picture of the performance for small or intermediate ρ. However, qualitative behavior can be surmised.

When ρ is close to 1, it is better to look at $\Pr(N)$ as a function of $t := 1 - s$, for then

$$P := \Pr(N) = s^N = e^{N\log(s)} = e^{N\log(1-t)} \approx e^{-tN} \quad \text{for} \quad t \ll 1.$$

(Remember that t is a function of ρ.) This formula, and all examples we will examine in the rest of this paper, can be looked at from a different point of view. For a *fixed* probability of overflow, say $P = 1/10^k$, the buffer size must be:

$$N(\rho) = \frac{\log_e P}{(-t)} = \frac{k\log_e(10)}{t}.$$

Clearly, a doubling of the buffer size (for fixed ρ and thus fixed t) will reduce the probability of overflow to $1/10^{2k}$. Therefore we see that inexpensive backup buffers can reduce packet loss to arbitrarily small values if t is not too small.

2.2.1 Some Examples

Our purpose in this paper is to examine the affect which PT distributions and their truncated cousins have upon buffer sizes. For our "base" function we have chosen $\alpha = 1.4$ and $\theta = 0.5$. [GREI95] has shown that for fixed α, system behavior is quite insensitive to changes in θ, so any intermediate value will do. On the other hand, performance is *very* sensitive to α. The value $\alpha = 1.4$ fits the data given in [LELA94]. For comparison, we have included the Erlangian-2 function, $E_2(x) = \mu^2 x e^{-\mu x}$, and the hyperexponential-2 function, $H_2(x) = p\mu_1 e^{-\mu_1 x} + (1-p)\mu_2 e^{-\mu_2 x}$. In all cases, the μ's have been chosen to give a mean interarrival time of 1. As always, the Erlangian-2 has a coefficient of variation of $C^2 = 1/2$. The H_2 function, however, is a 3-parameter function, and even after choosing an appropriate C^2, one arbitrary parameter remains. For our first set of calculations we chose $C^2 = 4.75$, the same as the TPT for $M = 8$, and $p = .0001$. The results for $P = .01$ (1 percent primary buffer overflow probability) are given in Figure 1. As would be expected, the necessary buffer size grows unboundedly as ρ approaches 1. Therefore, in order to control the variation along the y-axis, N was multiplied by $1 - \rho$. Even so, for large M, C^2 becomes unboundedly large, so we plotted $\log[(1-\rho)\,N]$ versus ρ instead.

It is true that for all distributions with finite variance (see [LIPS92]),

$$\lim_{\rho\to 1} \frac{t}{(1-\rho)} = \frac{2}{C^2+1}.$$

Therefore,

$$\lim_{\rho\to 1}[(1-\rho)\cdot N(\rho)] = \log_e(1/P)\left(\frac{C^2+1}{2}\right).$$

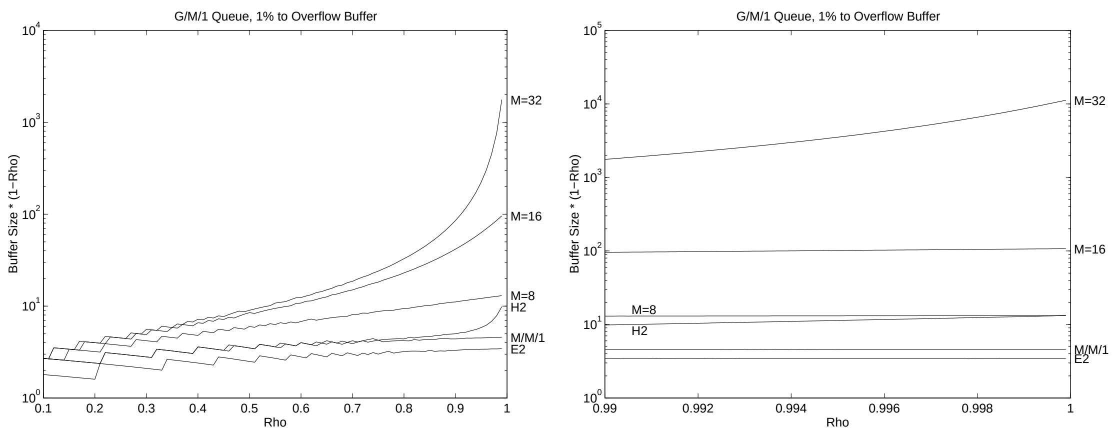

Figure 1: **Primary Buffer Size Needed in a GI/M/1 Queue for Overflow To Be Less Than 1%, as a Function of** $\rho = 1/(\lambda \cdot \bar{x})$. Because the buffer size can become very large as ρ approaches 1, the function actually plotted is $\log[(1-\rho)N]$. All curves (except that for $M = \infty$) are finite at $\rho = 1$, as shown by the inset figure. The curves are discontinuous because N is an integer function, and have negative slopes for small ρ because of the factor $1 - \rho$.

Except for very large M, the inset of Figure 1 clearly shows that a limit exists. But since C^2 goes to infinity as M does, the PT distribution itself must yield an infinite limit. This is certainly true. [GREI95] has shown that as ρ approaches 1

$$t(\rho) \longrightarrow (1-\rho)^{\frac{1}{\alpha-1}}.$$

Therefore,

$$(1-\rho) \cdot N(\rho) \longrightarrow \frac{1}{(1-\rho)^{\frac{2-\alpha}{\alpha-1}}} \longrightarrow \infty.$$

Put differently, for PT distributions with $1 < \alpha \leq 2$,

$$\lim_{\rho \to 1} (1-\rho)^{\frac{1}{\alpha-1}} \cdot N(\rho)$$

is finite and greater than 0.

While these limits are interesting in their own right, and show that the calculations are consistent with theory, they are not of much use for practical performance analysis. Most of the extreme behavior occurs for $\rho > .9$, where almost any system would be expected to behave badly. Of more interest is the range $0.5 < \rho < 0.9$. In this range it is clear that E_2 and Poisson renewal processes cause no great buffer problems. Even the H_2 distribution gives results very similar to that of the M/M/1 queue except above $\rho = .9$ where it finally rises abruptly to approach the same limit as the curve for the TPT with $M = 8$, as it must, since they have the same C^2. The

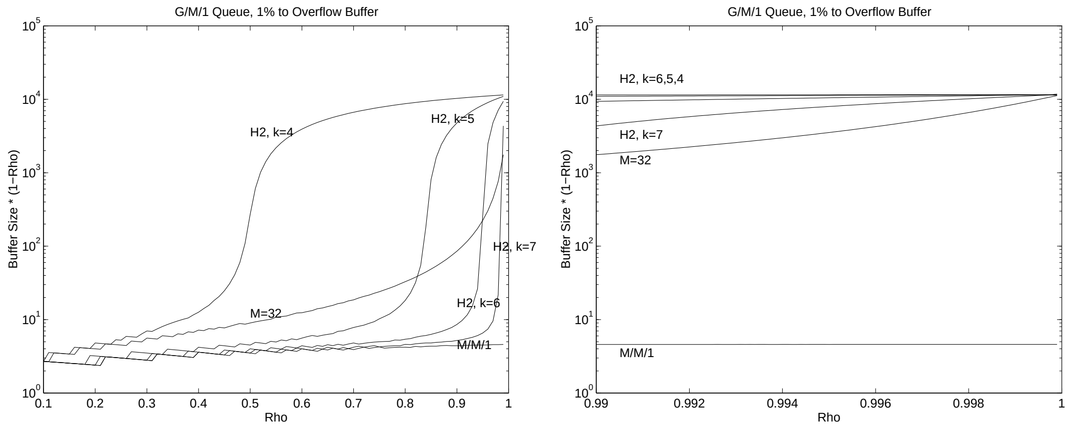

Figure 2: **Comparison of Primary Buffer Sizes between a TPT with $M = 32$ and various H_2 distributions with the same C^2, as a Function of $\rho = 1/(\lambda \cdot \bar{x})$.** Buffer size*$(1 - \rho)$ is plotted on a logarithmic scale. For the H_2's, $p = 10^{-k}$, for $k = 4$, 5, 6, 7. The M/M/1 queue is included for reference.

implication is strong here that C^2 is *not* as significant as it is, say in the mean queue length of the steady-state M/G/1 queue. To explore this further we compared the TPT with $M = 32$, and $C^2 = 5033.44\cdots$ with various H_2 distributions with approximately the same mean and coefficient of variation ($C^2 = 5033$). For the different H_2's we selected $p = .0001$, .00001 .000001, and .0000001 (p must be smaller than .0004 order to get such a big C^2). The results are given in Figure 2.

Recall that all the buffer sizes grow unboundedly as ρ approaches 1, so we once again plot $(1 - \rho) \cdot N(\rho)$ on a logarithmic scale. It is clear from the figure that none of the curves have anything in common, except near $\rho = 0$ and at $\rho = 1$. The inset shows they have the same asymptotic value as ρ approaches 1. The TPT increases smoothly throughout the range, but the others behave as would an M/M/1 queue for small ρ, and at different values of ρ jump rapidly to a higher level. We expect that this is an artifact of the H_2 distributions. They can each be thought of as generating a Poisson stream of packets, interspersed infrequently (p) with an extremely long pause ($1/\mu_1$). When ρ is small, the pause is long enough for the queue to drain. As ρ increases, enough packets arrive during the busy times to back up the queue sufficiently so that it cannot drain during the quiet time. The smaller p is, the closer μ_2 is to 1, so the queue cannot build up even during the busy times unless ρ is very close to 1. Clearly, this describes such specialized behavior that H_2 functions cannot be used to describe a general behavior pattern, at least not for buffer problems. We see as a general rule that in the range of interest, C^2 does

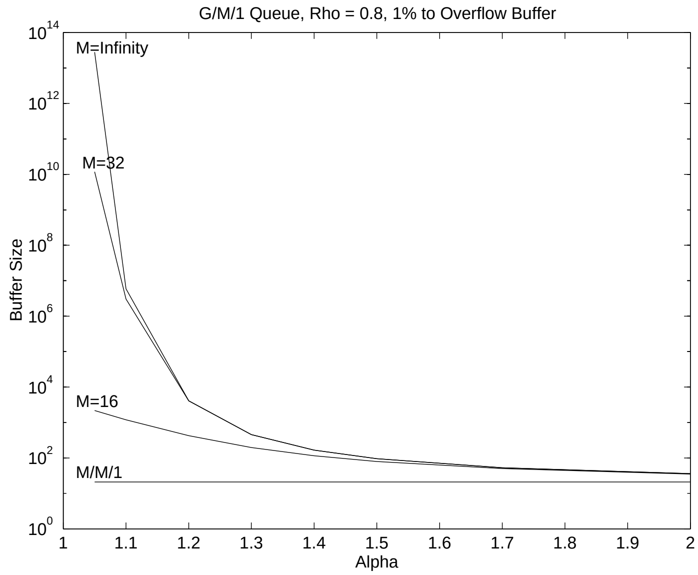

Figure 3: **Primary Buffer Size of a GI/M/1 Queue for the PT and Various TPT distributions as a Function of α, With $\rho = 1/(\lambda \cdot \bar{x}) = 0.8$.** The M/M/1 queue is included for reference.

not tell the whole story (or even a good part of it).

In the previous figures we chose the power parameter to be $\alpha = 1.4$, matching the experimental value that appeared in [LELA94]. We now describe how performance varies over the critical range of $1 < \alpha \leq 2$ for intermediate $0.7 \leq \rho \leq 0.9$. In Figure 3 we see how the degree of truncation of the power tail affects performance when $\rho = 0.8$. The M/M/1 queue ($M = 1$) is again included as reference. We see that even for an M as little as 16, the TPT and PT distributions yield comparable results for $\alpha \geq 1.4$. But below that value, the buffer sizes become extraordinately large, and below $\alpha = 1.1$ different truncations yield very different results. Even $M = 64$ does not come close to the full PT distribution. In this region, the buffer sizes are so big that they become meaningless for a real-world situation. In the near future, at least, can we expect a host to process 10^{12} packets in a single hour, let alone store them? Thus we must conclude that systems experiencing PT arrivals with $\alpha < 1.1$ never reach a steady state. Another modelling procedure must be found.

Figure 4 Shows that for $\alpha \geq 2$ PT distributions and their truncated cousins behave like other distributions. Their unusual behavior only becomes significant when α goes below 1.4. System behavior seems to vary smoothly with increasing ρ. But keep in mind that buffer size is given as a log scale. There is a factor of 2 difference between the $\rho = 0.7$ and the $\rho = 0.9$ curves.

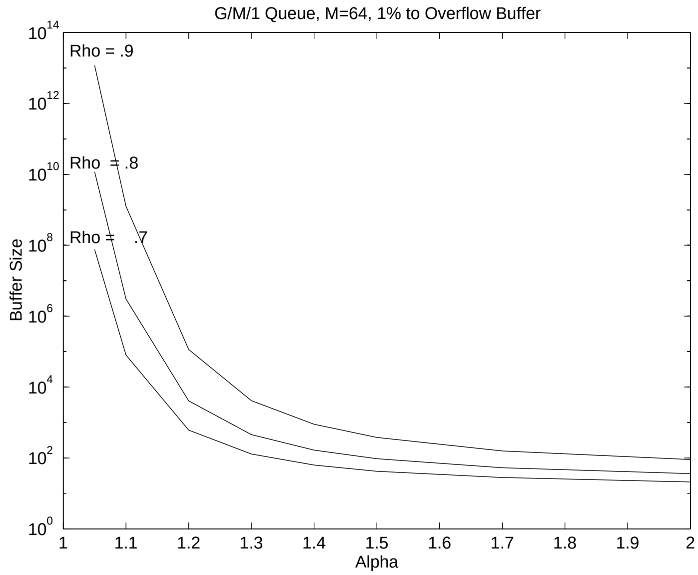

Figure 4: **Primary Buffer Size in A GI/M/1 Queue for the TPT Distribution With $M = 64$ as a function of** α, with $\rho = 1/(\lambda \cdot \bar{x}) = 0.7$, 0.8 and 0.9.

2.3 GI/M/1/N Queues - Finite Buffer

An explicit steady-state expression for the probability of finding k customers in a GI/M/1/N queue is only known in terms of LAQT and is given in [LIPS92]. They are

$$\begin{aligned} \pi(0 \mid N) &= \lambda g(N)\Psi[\mathbf{U}^N \mathbf{V}] \\ \pi(k \mid N) &= g(N)\Psi[\mathbf{U}^{N+1-k}] \end{aligned}$$

The arrival probabilities (i.e., the probability that an arriving packet will see k packets already in the buffer) are different. Let N be the size of the buffer, then

$$a(k \mid N) = K(N)\Psi[\mathbf{U}^{N-k}] \quad \text{for} \quad 0 \leq k \leq N,$$

where $\mathbf{U} := \mathbf{A}^{-1}$, $\Psi[\,\cdot\,]$ is defined in (9) and $\mathbf{A}$ is given by (11). $K(N)$ is the normallizing factor making the sum of the probabilities equal to 1. That is,

$$\sum_{k=0}^{N} a(k \mid N) = 1 \Longrightarrow \frac{1}{K(N)} = \Psi[\mathbf{U}^N + \mathbf{U}^{N-1} + \cdots + \mathbf{U}^2 + \mathbf{U} + \mathbf{I}] = \Psi[(\mathbf{I} - \mathbf{U})^{-1}(\mathbf{I} - \mathbf{U}^{N+1})].$$

Details of how to compute this are given in [LIPS92]. The probability that a packet will be lost is the same as the probability that an arriving packet will see a full buffer, and is given by $a(N \mid N)$. Thus

$$\Pr(N) = K(N).$$

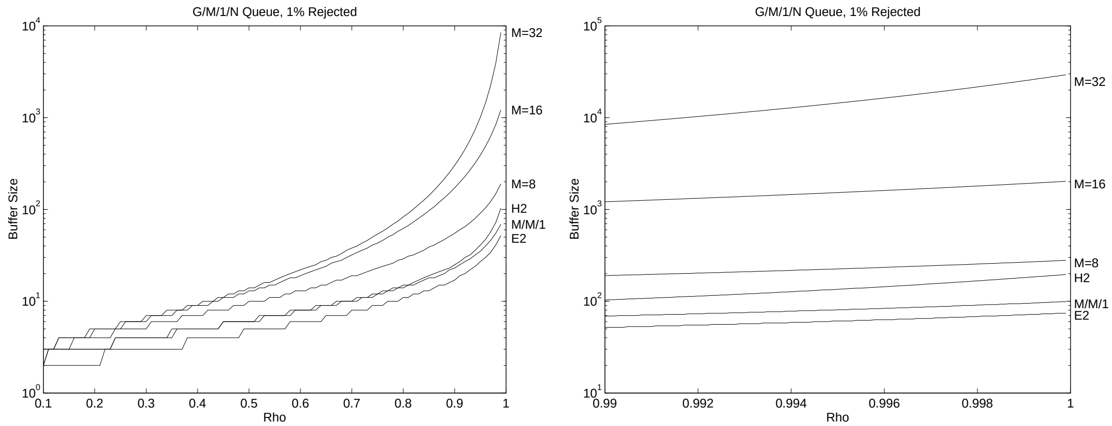

Figure 5: **Buffer Size For 1% Loss in A GI/M/1/N Queue for Various Distributions as a function of** $\rho = 1/(\lambda \cdot \bar{x})$. The distributions included are: TPT with $M = 1, 8, 16, 32$, E_2 and $H_2(C^2 = 4.75)$.

Figure 5 is similar to Figure 1 except that $N(\rho)$ is not multiplied by $(1-\rho)$, since $N(\rho = 1)$ itself is finite (the functions blow up at $\rho = 1 + P$). The same pattern occers here. The H_2/M/1/N system behaves no differently than the M/M/1/N until $\rho > 0.9$, even though it has the same C^2 value as the curve labelled $M = 8$. So we see again that C^2 does not tell the story, certainly not in the range of primary interest.

2.4 M/G/1 Queues - Infinite Buffer

Irratic traffic may be caused by files whose sizes are distributed according to a power-tail law, but are broken up into numerous smaller packets which then disburse upon transmission, giving an appearance of burstiness. If we imagine that they are reassembled at the server and stored as one packet, then this could be adequately described as an M/G/1 system. But now we are faced with an obvious problem. The Pollaczek-Khinchin formula states clearly that a service distribution with infinite variance must produce an infinite mean queue length for the steady-state M/G/1 queue, for *all* ρ. If the time to process a packet is proportional to the size of the packet, then the amount of buffer space needed to hold waiting packets must be proportional to the waiting time, which over a long time must grow unboundedly if $C^2 = \infty$. However, if an entire reassembled packet can be stored in a single *slot*, then the steady-state probability that an arriving packet will find that N or more slots are already buffer taken, is *not* infinite, even though the mean queue length is infinite. We hypothesize without proof, that the steady-state

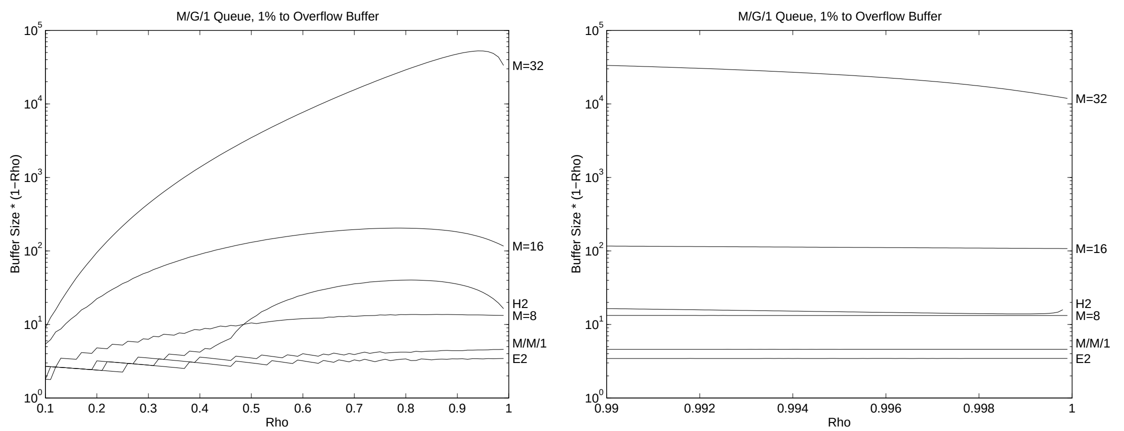

Figure 6: **Primary Buffer Size Needed for Overflow of an M/G/1 Queue To Be Less Than 1%, as a Function of $\rho = \lambda \cdot \bar{x}$.** Because buffer size can become very large as ρ approaches 1, the function actually plotted is $\log[(1-\rho)N]$. All curves appear to be finite at $\rho = 1$, as shown by the inset figure.

(and arrival) probabilities for PT distributions satisfy, for large n

$$a(n) = \pi(n) \longrightarrow (1-\rho)\frac{\text{const}}{n^\alpha}.$$

Then, if $\alpha > 1$

$$\lim_{N\to\infty} \Pr(N) = \lim_{N\to\infty} \sum_{n=N}^{\infty} a(n) = 0$$

where now (and hereafter), $\rho = \lambda\bar{x}$. λ is the Poisson arrival rate, and $\bar{x}$ is the mean service time. This says that there exists a finite N for which $\Pr(N) = \delta$ for all $1 > \delta > 0$, however small δ is. Notice that this would be true even though $\bar{q}$ is infinite if $\alpha \leq 2$, for then

$$\bar{q} := \sum_{n=1}^{\infty} n\, a(n) \approx (1-\rho)\cdot \text{const} \sum_{n=1}^{\infty} \frac{n}{n^\alpha} = (1-\rho)\cdot\text{const}\sum_{n=1}^{\infty}\frac{1}{n^{\alpha-1}} = \infty.$$

(Note that if $\alpha \leq 1$ then there can be no steady state.) We assume here, then, that each reconstituted packet takes up one slot.

The steady-state and arrival probabilities are the same for an M/G/1 queue, and are given in LAQT form by:

$$a(n) = \pi(n) = (1-\rho)\Psi[\mathbf{U}^n].$$

The probability that an arriving packet will find N or more slots full is given by:

$$\Pr(N) = (1-\rho)\sum_{n=N}^{\infty}\Psi[\mathbf{U}^n] = (1-\rho)\Psi[\mathbf{U}^N(\mathbf{I}-\mathbf{U})^{-1}].$$

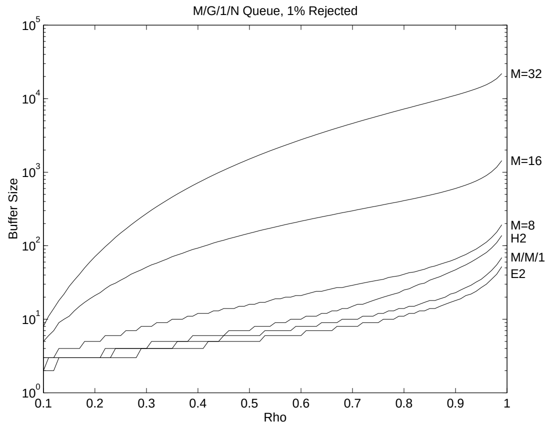

Figure 7: **Primary Buffer Size of an M/G/1/N Queue for Various TPT distributions as a Function of $\rho = \lambda \cdot \bar{x}$.** The M/M/1/N, M/$H_2$/1/N and M/$E_2$/1/N queues are also included.

We have calculated these probabilities for the usual collection of distributions, and present the results in Figure 6. The behavior is similar to that of the GI/M/1 queue (Figure 1), but the primary buffer sizes are somewhat bigger here. However we see that these curves are concave downward (that is, the buffer size needed grows somewhat more slowly than $1/(1-\rho)$). Also, the H_2 system behaves peculiarly, even for the relatively small $C^2 = 4.75$. But it does approach the same value as the TPT system for $M = 8$ as ρ approaches 1. The inset shows a slight upturn for this curve very close to 1, but this is almost surely due to numerical instability.

2.5 M/G/1/N Queues - Finite Buffer

Finally we reach our last system, and last figure. As with the open M/G/1 system the arrival probabilities and the steady-state probabilities are equal. (This is not true for GI/M/1 and GI/M/1/N queues.) Therefore, from [LIPS92]

$$a(n,\, N) = \pi(n,\, N) = G(N) \begin{cases} \Psi[\mathbf{U}^n] & \text{for } 0 \le n < N \\ \lambda \Psi[\mathbf{U}^{N-1}\mathbf{V}] & \text{for } n = N. \end{cases}$$

where

$$[G(N)]^{-1} = \Psi[(\mathbf{I}-\mathbf{U})^{-1}(\mathbf{I}-\mathbf{U}^N)] + \lambda\Psi[\mathbf{U}^{N-1}\mathbf{V}].$$

The probability that an arriving packet will be rejected is $a(N;\, N)$. Therefore,

$$\Pr(N) = \frac{\lambda\Psi[\mathbf{U}^{N-1}\mathbf{V}]}{\Psi[(\mathbf{I}-\mathbf{U})^{-1}(\mathbf{I}-\mathbf{U}^N)] + \lambda\Psi[\mathbf{U}^{N-1}\mathbf{V}]}.$$

Figure 7 shows the result of our calculations for 1.0% rejection. It is not clear how the true PT would behave, but it will be finite. As with the GI/M/1/N queue, $N(\rho = 1)$ is finite but should grow unboundedly as ρ approaches $1 + P$, which is 1.01 throughout this paper.

We have not calculated the behavior of the full PT system yet because it requires manipulating extremely large matrices, or finding a large number of the eigenvalues of $\mathbf{A}$. For PT distributions, $\mathbf{A}$ has an infinite number of eigenvalues, with an accumulation point at 1. Therefore in principle, $(\mathbf{I} - \mathbf{U})$ doesn't have an inverse. However, by using the *Spectral Decomposition Theorem*, [LIPS95] has derived an expression by which any function of $\mathbf{A}$ can be computed. It is:

$$\Psi[f(\mathbf{A})] = \sum_i \frac{f(\nu_i)}{\Psi\left[\{(1-\nu_i)\mathbf{I} + 1/\lambda\mathbf{B}\}^{-2}\right]}.$$

The sum is over all eigenvalues of $\mathbf{A}$, $\{\nu_i\}$. We have worked out a numerical procedure which allows any number of eigenvalues, together with evaluation of the $\Psi[\,\cdot\,]$ expression, to be computed automatically, but haven't yet performed the sum. The sum converges geometrically, so there should be no numerical instability.

3 Conclusion

We have shown how to integrate power-tail distributions and their truncated children into the analysis of communications networks using various GI/G/1 queues with-and-without finite buffers. Other types of test functions (e.g., $H_2(x)$) must surely be inadequate. On the other hand, more complicated processes which hueristically build in correlations, may well be unnecessary. The models given here, though relatively simple, can be used by researchers who use discrete event simulations. The reason why this could be very important is that statistics for PT distributions converge to their mean much more slowly than any other distibutions [GREI95]. Our simple analytic model and results could serve as base-line comparisons to see if an equivalent simulation is anywhere near convergence, before the researcher attempts a more complicated simulation model.

Whereas for most processes the average for a set of data converges to its mean as $1/\sqrt(n)$ (n is the number of data points), data generated by PT distributions with $1 < \alpha < 2$ converge as $1/n^{\beta}$, where $\beta = 1 - 1/\alpha$. If $\alpha \leq 1$, the data doesn't converge at all! In fact, if one were to try to test to see if a given set of data was generated by a PT distribution with $\alpha < 2$, the χ^2 test would fail even if the hypothesis were true.

If it turns out that PT distributions play a significant role in communications systems, then much more research will have to be carried out on the statistical convergence in simulations. It will probably be necessary to find new tools for analysis, based on transient systems, because the number of events needed to bring a PT system to its steady state may exceed the lifetime of that system. It is interesting to note that while there are several theories as to why PT distributions occur (all of which *probably* have some truth in them), there are no ideas or hunches or suggestions as to why a particular system should have a particular α. Since the behavior of any PT system very strongly depends upon its characteristic α, *any ideas* in this direction deserve a serious hearing.

References

[BERA95] J. Beran, R. Sherman, M. S. Taqqu and W. Willinger, "Variable-Bit-Rate Video Traffic and Long-Range Dependence", *IEEE/ACM Trans. on Networking*, 1995.

[FELL71] William Feller, *An Introduction to Probability Theory and its Applications*, Vol. II, John Wiley and Sons, New York, 1971.

[GARG92] Sharad Garg, Lester Lipsky and Maryann Robbert, "The Effect of Power-Tail Distributions on the Behavior of Time Sharing Computer Systems", *1992 ACM SIGAPP Symposium on Applied Computing*, Kansas City, MO, March, 1992.

[GREI95] Michael Greiner, Manfred Jobmann and Lester Lipsky, "The Importance of Power-tail Distributions for Telecommunication Traffic Models", Technical Report, Department of Informatics, Technical University-Munich (TUM), Submitted for publication.

[LELA86] Will E. Leland and Teunis Ott, "Load-Balancing Heuristics and Process Behavior", *Proceedings of ACM SIGMETRICS 1986*, May 27 - 30, 1986, pp. 54-69. (The proceedings appeared as vol 14, no 1, Performance Evaluation Review, May, 1986.)

[LELA94] Will E. Leland, Murad S. Taqqu, Walter Willinger and Daniel V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)", *Proc. of IEEE/ACM Trans. on Networking*, **2**,1, Feb. 1994.

[LIEF94] Appie van de Liefvoort and Hai Feng Weng, "On the Instability ('Fractal' Behavior) of Arrival Counts Generated By Power-Tail Renewal Processes", Department Report, Computer Science and Telecommunications Program, University of Missouri-Kansas City, 1994.

[LIKH95] Nikolai Likhanov, Boris Tsybakov and Nicolas D. Georganas, "Analysis of an ATM Buffer with Self-Similar ("Fractal") Input Traffic", Proc. IEEE INFOCOM'95, Boston, April 1995.

[LIPS86] Lester Lipsky, "A Heuristic Fit of an Unusual Set of Data", Bell Communications Research Report, January 1986.

[LIPS92] Lester Lipsky, *QUEUEING THEORY: A Linear Algebraic Approach*, MacMillan and Company, New York, 1992.

[LIPS95] Lester Lipsky and Pierre Fiorini, "Auto-Correlation of Counting Processes Associated with Renewal Processes", Technical Report, Booth Research Center, University of Connecticut, August 1995.

[LOWR93] Walter Lowrie and Lester Lipsky, "A Model For The Probability Distribution of Medical Expenses", Proceedings of *CONFERENCE OF ACTUARIES IN PUBLIC PRACTICE*, 1993.