

How heavy-tailed distributions affect simulation-generated time averages

Citation for published version (APA):

Fishman, G. S., & Adan, I. J. B. F. (2003). *How heavy-tailed distributions affect simulation-generated time averages*. (SPOR-Report : reports in statistics, probability and operations research; Vol. 200316). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/2003

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

How Heavy-Tailed Distributions Affect Simulation-Generated Time Averages¹

George S. Fishman² and Ivo Adan³

Abstract

This paper describes how a heavy-tailed file-size distribution with $1 < \alpha < 2$ affects statistical inference based on sample-path data from a discrete-event simulation of a telecommunications network. Representing the session initiation and file transmission processes as an M/G/ ∞ queueing system, we derive the conditional distribution of $X(t)$, the number of active sessions (number in system) at time t , and its equilibrium covariance function. Assuming that file size (service time) follows the Lomax heavy-tailed distribution, the paper shows that the conditional distribution converges to its equilibrium distribution at rate $t^{1-\alpha}$. Also, for a sample-path of length τ , the variances of sample averages converges as $\tau^{1-\alpha}$. Moreover, the central limit theorem does not hold.

To meet a specified accuracy $\nu \geq$ variance-to-mean for the sample average number of active sessions, a discrete-event simulation requires impracticably long sample paths for $\nu = 10^{-4}$ and $1.2 \leq \alpha \leq 1.8$, for $\nu = 10^{-3}$ and $1.2 \leq \alpha \leq 1.4$, and $\nu = 10^{-2}$ and $\alpha = 1.2$. This motivated a study of alternative file-size distributions, namely, the lognormal and Weibull which have also been suggested in the literature. The study approximates the Lomax by these distributions, matching their means and $1 - \zeta$ quantiles to retain its major features. For $\zeta = 10^{-3}$, both alternatives require considerably shorter paths to achieve the same accuracy. For $\zeta = 10^{-5}$, required paths are also shorter, except for $(\nu = 10^{-3}, \alpha = 1.8)$ and $(\nu = 10^{-2}, \alpha \geq 1.6)$. Although the lognormal requires longer paths than the Weibull, it better approximates the Lomax. For $\zeta = 10^{-3}$, the required lognormal sample sizes are practicable for all ν and α . For $\zeta = 10^{-5}$, they are practicable for almost all scenarios.

Keywords – Lomax distribution, M/G/ ∞ queue, network modeling, sample-path average, discrete-event simulation.

¹The authors gratefully acknowledge the assistance of Marko Boon in checking the numerical calculations.

²Department of Operations Research, CB 3180, University of North Carolina, Chapel Hill, NC 27599. E-mail: gfish@unc.edu.

³Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands. E-mail: iadan@win.tue.nl.

Analysis of telecommunications network data reveals that traffic shows substantial similarity in burstiness on time scales of hours, minutes, and milliseconds. As examples, Beran et al. (1995) demonstrate this behavior for variable-bit-rate (VBR) video, Paxson and Floyd (1995) show it for Wide Area Networks, and Leland et al. (1994) and Willinger et al. (1995) show it for Ethernet Local Area Networks. This *self similarity* implies that the correlation between the numbers of arrivals in fixed-width non-overlapping time intervals decreases slowly as the the time distance between the intervals increases. One representation consistent with self similarity assumes that the autocorrelation function for these intervals follows a *power law*, a representation that Crovella and Bestavros (1996) observe sometimes characterizes the right tail of the file-size distribution as well. The power-law characterization has encouraged the use of *heavy-tailed distributions* to describe the behavior of session initiation times and file sizes.

By a heavy-tailed distribution, we mean one for which

$$\text{pr}(Z > z) \sim cz^{-\alpha} \quad \text{as } z \rightarrow \infty, \quad (1)$$

where Z denotes a random variable, $c > 0$ a location parameter, α a shape parameter taking values in $(1, 2)$, and \sim indicates that the ratio of the right and left sides in (1) converge to unity as z increases without bound. This heavy-tailed behavior has stimulated considerable research on probabilistic characterizations of the *number of network residents* and the *network residence time*. Resnick and Samorodnitsky (2000) contains references to much of this work.

However, the heavy-tailed representation is far from a universally accepted tool for characterizing all network data. Downey (2001,2002) argues that the lognormal distribution provides a good fit to file-size data in most systems. Cao et al. (2003) show good fits for interarrival times using the Weibull distribution, with probability density function (p.d.f.) $\{f(x) = \phi\beta^{-\phi}x^{\phi-1}e^{-(x/\beta)^\phi}, \phi, \beta > 0, x \geq 0\}$ for $0 < \phi < 1$. As the aggregate connection rate increases, they also show empirically that the arrival process tends to Poisson, consistent with Khinchin's theorem (e.g., Gnedenko and Kovalenko 1987). Also, Feldman and Whitt (1997) show how approximating a heavy-tailed distribution by a mixture of exponential distributions makes the resulting model more amenable to analysis.

Although a heavy-tailed characterization may allow better fidelity in network modeling, its introduction has aroused concern about how it affects statistical inferences drawn from sample-path data collected on discrete-event

network simulations (Crovella and Lipsky 1997). The present study examines this issue using an M/G/ ∞ queueing model to characterize the number of *active sessions* (number in system). In particular, we focus on how a heavy-tailed file-size (service-time) distribution F affects estimation, from simulation-generated data, of the mean number of active sessions and on its equilibrium probability mass function (p.m.f.). In particular, for the number of active sessions process, we show that convergence from an arbitrarily selected initial state to equilibrium and convergence of the variance of sample averages are considerably slower than for more conventional file-size distributions. The paper also describes how approximating the heavy-tailed Lomax distribution by a lognormal or a Weibull distribution generally leads to considerably shorter sample paths to obtain sample averages within a specified accuracy. Lastly, for the method of approximation in Sec. 5.1, we show how the lognormal distribution approximates the Lomax distribution better than the Weibull. Although the model under consideration is very simple, the conclusion, that heavy tailed distributions preclude simulation as a means of estimating quantities with a moderate to high degree of accuracy, of course carries over to more complex and realistic network models.

The M/G/ ∞ model has an infinite number of servers, i.i.d. exponentially distributed interarrival times, and i.i.d. service times from a general distribution. In the present context, an infinite number of servers implies that no session initiation waits to begin transmitting its file or message. The specification for interarrival times implies that arrivals or session initiations form a Poisson process, a not untenable assumption in a network in which many independent sources each contribute a relatively small proportion of the total traffic. Most importantly, assuming a fixed transmission rate, we can treat service times as properly scaled file sizes. In what follows, we refer interchangeably to arrival process and session-initiation process. Likewise, we refer interchangeably to service time and file size, with a higher frequency for the former to emphasize that our results apply more widely than just in telecommunications networks.

Heavy-tailed behavior causes the correlation within a simulation-generated sample path to decrease slowly, thereby requiring excessively long warm-up intervals. Also, the variances of sample averages, computed from sample-path data, may not converge to zero as the sample-path length, τ , increases without bound. This makes the sample-path data useless. If the variances do converge, they frequently do so at rates less than τ^{-1} and the sample averages do not obey a normal central limit theorem. The slower than τ^{-1}

convergence indicates that extremely long simulations need to be run to estimate the means to within a meaningful degree of relative accuracy and the absence of a normal limit law makes it difficult to gauge how good the estimates are, as for example, using LABATCH.2 (Fishman 2001). Moreover, these impediments for variances and normality carry over to the estimation of the d.f.s.

Using the Lomax distribution with $1 < \alpha < 2$, we show that:

- i. The influence of initial conditions on the mean number of active sessions at time t decreases at a rate $\frac{1}{t^{\alpha-1}}$.
- ii. The variance of the sample arithmetic average of the mean number of active sessions, based on a sample-path length, τ , decreases at rate $\frac{1}{\tau^{\alpha-1}}$.
- iii. Asymptotic normality does not hold for the sample average.
- iv. If $\alpha > 2$, the rates become t^{-1} and τ^{-1} for the warm-up interval and variance respectively.
- v. In general, the variance of a sample average decreases as τ^{-1} if and only if the second moment of the service-time distribution is finite.
- vi. Similar results hold for estimates of the d.f. as well as the mean.

These observations prompted us to investigate how statistical inference based on simulation-generated sample-path data would fare when using a lognormal (Sec. 5.2) or a Weibull (Sec. 5.3) distribution to generate service times. To make these distributions approximate the Lomax, their parametric values were assigned by matching means and $1 - \zeta$ quantiles. For $\zeta = 10^{-3}$ and 10^{-5} , our results indicate that:

- vii. The Weibull and lognormal distributions require considerably shorter sample paths than the Lomax to estimate the mean number of active sessions to within a high or moderate specified accuracy. The Weibull requires a shorter sample path than the lognormal for $\zeta = 10^{-3}$ and longer paths for $\zeta = 10^{-5}$.
- viii. The ratio of Lomax to either the Weibull or lognormal sample-path lengths grows rapidly as the specified accuracy for the sample average increases. In some cases, the ratio exceeds 10^{10} .

- ix. In some cases, the lognormal and the Weibull require longer sample paths than the Lomax. These occur for a gross accuracy specification for the sample average.

Some qualification for this method is needed. The first step in fitting a distribution to empirical data is to estimate the values of its parameters, for example, by the maximum likelihood method. If several distributions are under consideration, this is to be done for each. The next step is to perform formal goodness-of-fit tests or graphical comparisons of each fitted d.f. with the empirical d.f. and with each other. If, in the present context, statistical analysis points to the lognormal or Weibull d.f. as fitting the empirical file-size d.f. significantly better, or at least as well as, a heavy-tailed d.f., then the choice for file-size distribution is clear. If the analysis demonstrates that a heavy-tailed distribution fits the data significantly better, then our results indicate that its adoption for simulation inevitably leads to excessive computing time.

However, there is an alternative. Suppose one assigns values to the parameters of the Lomax distribution by conventional methods, but assigns values to the parameters of the lognormal or Weibull distribution according to the scheme mentioned earlier and described in Sec. 5.1. That is, the resulting lognormal or Weibull distribution is an approximation to the statistically fitted Lomax distribution. Then it is of interest to determine how well this alternatively fitted d.f. compares to the Lomax d.f. Section 5.4 examines this issue and concludes that:

- x. The lognormal approximates the Lomax distribution better than the Weibull does.

While encouraging consideration of this approximating scheme to save simulation computing time, we emphasize that an appropriate statistical analysis should precede its usage in each study.

1 Heavy-Tailed Distributions

Heavy-tailed distributions satisfying (1) include the:

Pareto Distribution $\mathcal{Pa}(\alpha, b)$ with p.d.f.

$$f(x) = \frac{\alpha b^\alpha}{x^{\alpha+1}} \quad \alpha > 1, x \geq b > 0 \quad (2)$$

with

$$\text{mean} = \frac{\alpha b}{\alpha - 1}$$

and mode at $x = b$,

Lomax Distribution $\mathcal{L}(\alpha, k)$

$$f(x) = \frac{\alpha k^\alpha}{(x + k)^{\alpha+1}} \quad \alpha > 1, k > 0, x \geq 0 \quad (3)$$

with

$$\text{mean} = \frac{k}{\alpha - 1}$$

and mode at $x = 0$,

Inverted Gamma Distribution $\mathcal{IG}(\alpha, \beta)$

$$f(x) = \frac{e^{-1/\beta x}}{\Gamma(\alpha)\beta^\alpha x^{\alpha+1}} \quad \alpha > 1, \beta > 0, x \geq 0 \quad (4)$$

with

$$\text{mean} = \frac{1}{(\alpha - 1)\beta}$$

and mode at $x = \frac{1}{(\alpha+1)\beta}$, and

Inverse Weibull Distribution $\mathcal{IW}(\alpha, \beta)$

$$f(x) = \frac{\alpha e^{-(1/\beta x)^\alpha}}{\beta^\alpha x^{\alpha+1}} \quad \alpha > 1, \beta > 0, x \geq 0 \quad (5)$$

with

$$\text{mean} = \frac{\Gamma(1 - 1/\alpha)}{\beta}$$

and mode at $x = \frac{1}{\beta} \left(\frac{\alpha}{1+\alpha} \right)^{1/\alpha}$.

These distributions are bounded above and, for $\alpha < 2$, have infinite variance and infinite higher moments. Although they can also be defined for $0 < \alpha < 1$, this would induce an infinite mean. The Pareto, sometimes called the Pareto distribution of the first kind, is monotone decreasing with

restricted domain $[b, \infty)$. The Lomax, sometimes called the Pareto distribution of the second kind or the Johnson Type VI distribution, is also monotone decreasing but has the larger domain $[0, \infty)$. If X is from $\mathcal{Pa}(\alpha, b)$, then $Y = k \left(\frac{X}{b} - 1 \right)$ is from $\mathcal{L}(\alpha, k)$. The inverted Gamma and inverse Weibull p.d.f.s are also unimodal, but allow more flexibility for the locations of their modes. If Z is a Gamma distributed, then $X = 1/Z$ is inverted Gamma distributed. If Z is Weibull distributed, then $X = 1/Z$ is inverse Weibull distributed. For a general reference, see Johnson et al. (1994).

2 Sample-Path Average

Let

$X(t) :=$ the number in the system at time t or, equivalently,
the number of active sessions

with unknown equilibrium probability mass function (p.m.f.)

$$\pi_j := \lim_{t \rightarrow \infty} \text{pr}[X(t) = j] \quad j = 0, 1, \dots$$

and mean

$$\mu := \sum_{j=0}^{\infty} j \pi_j < \infty.$$

For the present, assume μ is finite. Simulations customarily start in some arbitrary system state \mathcal{S}_0 . In order for simulation-generated sample-path data to provide a basis for a meaningful estimate of μ , we require that

$$\text{pr}[X(t) = j | \mathcal{S}_0] \rightarrow \pi_j \quad \text{as } t \rightarrow \infty. \quad (6)$$

To estimate μ from sample-path data, $\{X(s), 0 \leq s \leq t + \tau\}$ for some $t \geq 0$ and $\tau > 0$, one customarily computes

$$\hat{\mu}(t, \tau) := \frac{1}{\tau} \int_t^{t+\tau} X(s) ds, \quad (7)$$

$[0, t)$ being a *warm-up interval* selected to reduce the influence of the initial conditions on the segment of the sample path $\{X(s), t \leq s \leq t + \tau\}$ used for estimation. We then require the stronger conditions that

$$\text{E}[\hat{\mu}(t, \tau) | \mathcal{S}_0] \rightarrow \mu \quad \text{as } \tau \rightarrow \infty \quad (8)$$

and for any fixed $\tau > 0$

$$\mathbb{E}[\hat{\mu}(t, \tau) | \mathcal{S}_0] \rightarrow \mu \quad \text{as } t \rightarrow \infty. \quad (9)$$

Limit (9) implies that a sufficiently long warm-up interval $[0, t)$ eliminates bias in $\hat{\mu}(t, \tau)$ as an estimate of μ , whereas limit (8) implies that using a sufficiently long sample path, $[t, t + \tau)$, eliminates bias regardless of the length of the warm-up interval. It is important to note that (6) does not imply (8) and (9).

If there exists a positive constant κ such that for fixed $t \geq 0$

$$\tau \text{var}[\hat{\mu}(t, \tau) | \mathcal{S}_0] \rightarrow \kappa < \infty \quad \text{as } \tau \rightarrow \infty, \quad (10)$$

then $\hat{\mu}(t, \tau)$ has approximate variance κ/τ for large τ . Moreover, under relatively general conditions,

$$\frac{\hat{\mu}(t, \tau) - \mu}{\sqrt{\kappa/\tau}} \xrightarrow{d} \mathcal{N}(0, 1) \quad \text{as } \tau \rightarrow \infty, \quad (11)$$

thus providing a basis for evaluating the accuracy of $\hat{\mu}(t, \tau)$ when τ is sufficiently large.

The limits (8) and (9) are essential for meaningful statistical estimation. The limits (10) and (11) are convenient, but their absence does not rule out the possibilities that $\text{var}[\hat{\mu}(t, \tau) | \mathcal{S}_0]$ decreases with τ and that a limit law may exist for $[\hat{\mu}(t, \tau) - \mu] / \sqrt{\kappa g(\tau)}$, as $\tau \rightarrow \infty$, for some properly defined decreasing function, $\{g(\tau), \tau > 0\}$. Hereafter, the notation $\hat{\mu}(\infty, \tau)$ denotes $\hat{\mu}(t, \tau)$ as $t \rightarrow \infty$; i.e., $\hat{\mu}(\infty, \tau)$ is the average over a sample-path of length τ when the system is in equilibrium.

3 The M/G/ ∞ Queue

The M/G/ ∞ queueing model provides a convenient platform for uncovering the implications of using a heavy-tailed distribution for the i.i.d. service times, S_1, S_2, \dots with p.d.f. $\{f(x), x \geq 0\}$, d.f. $\{F(x) = \int_0^x f(y) dy, x \geq 0\}$, and mean ω^{-1} . It has i.i.d. exponential interarrival times, T_1, T_2, \dots , with rate λ and infinite service capacity.

It is known that in equilibrium the number in the system, $X(t)$, has the Poisson distribution

$$\pi_j = \frac{e^{-\mu} \mu^j}{j!} \quad j = 0, 1, \dots$$

with mean $\mu = \lambda/\omega$.

Common simulation practice assigns \mathcal{S}_0 as the empty and idle state. Accordingly, we hereafter assume that $X(0) = 0$. The Appendix gives the more comprehensive results for $l > 0$.

Let

$$\mu(h) := \mu\omega \int_h^\infty [1 - F(x)]dx, \quad (12)$$

where $\{\omega[1 - F(x)], x \geq 0\}$ is the p.d.f. of the residual service time. It is this function that determines all the distributional properties of $\{X(t), t \geq 0\}$. It is known that (Riordan 1951, Beneš 1957)

Proposition 3.1 *For the $M/G/\infty$ queue, $h \geq 0$, and $r = 0, 1, \dots$:*

$$\begin{aligned} \text{pr}[X(t) = r | X(0) = 0] &= p_{0r}(t) \\ &:= \frac{1}{r!} e^{-\mu + \mu(t)} [\mu - \mu(t)]^r \\ \text{E}[X(t) | X(0) = 0] &= \mu - \mu(t) \\ \text{E}[X(t)X(t+h) | X(0) = 0] &= [\mu - \mu(t)] [\mu - \mu(t+h)] + \mu(h) - \mu(t+h) \\ \text{cov}[X(t), X(t+h) | X(0) = 0] &= \mu(h) - \mu(t+h). \end{aligned} \quad (13)$$

□

(1951).

The sample-path average $\hat{\mu}(t, \tau)$ has conditional expectation

$$\text{E}[\hat{\mu}(t, \tau) | X(0) = 0] = \frac{1}{\tau} \int_t^{t+\tau} \text{E}[X(h) | X(0) = 0] dh = \frac{1}{\tau} \int_t^{t+\tau} [\mu - \mu(h)] dh \quad (14)$$

and conditional variance

$$\begin{aligned} \text{var}[\hat{\mu}(t, \tau) | X(0) = 0] &= \frac{2}{\tau^2} \int_t^{t+\tau} \int_0^{t+\tau-s} \text{cov}[X(s), X(s+h) | X(0) = 0] dh ds \\ &= \frac{2}{\tau^2} \int_t^{t+\tau} \int_0^{t+\tau-s} [\mu(h) - \mu(s+h)] dh ds \\ &= \frac{2}{\tau^2} \int_0^\tau \int_t^{t+\tau-h} [\mu(h) - \mu(s+h)] ds dh. \end{aligned} \quad (15)$$

Hence, as $t \rightarrow \infty$,

$$\begin{aligned}
\mathbb{E}\widehat{\mu}(\infty, \tau) &= \mu \\
\text{var}\widehat{\mu}(\infty, \tau) &= \frac{2}{\tau^2} \int_0^\tau [\tau - h]\mu(h)dh \\
\mathbb{E}[\widehat{\mu}(\infty, \tau) - \mu]^3 &= \frac{6}{\tau^3} \int_0^\tau [\tau - h]h\mu(h)dh \\
\mathbb{E}[\widehat{\mu}(\infty, \tau) - \mu]^4 &= \frac{1}{\tau^4} \int_0^\tau [\tau - h] [18h\tau + 12h^2] \mu(h)dh,
\end{aligned} \tag{16}$$

where the third and fourth central moments follow from the cumulants (semi-invariants) of $\widehat{\mu}(\infty, \tau)$ in Riordan (1951). For fixed $t \geq 0$, it follows from (15) that

$$\tau \text{var}[\widehat{\mu}(t, \tau) | X(0) = 0] \rightarrow \kappa := 2 \int_0^\infty \mu(h)dh \quad \text{as } \tau \rightarrow \infty. \tag{17}$$

Proposition 3.2 *A necessary and sufficient condition for $\kappa < \infty$ is that $\mathbb{E}S_i^2 < \infty$.*

Proof. From (12),

$$\begin{aligned}
\int_0^\infty \mu(h)dh &= \mu\omega \int_0^\infty \int_h^\infty [1 - F(x)]dx dh \\
&= \lambda \int_0^\infty [1 - F(x)] \int_0^x dh dx \\
&= \lambda \int_0^\infty x[1 - F(x)]dx \\
&= \lambda \mathbb{E}S_i^2.
\end{aligned}$$

□

Proposition 3.2 immediately tells us that for heavy-tailed distributions, κ is finite if and only if $\alpha \geq 2$. Since $1 < \alpha < 2$ is often the specification to maintain a sufficiently slowly vanishing tail in f , it is of interest to study this case in detail. Section 4 does this for the Lomax distribution. In particular, it shows that, although μ is finite, a long sample path is required to estimate it.

In characterizing the number of active sessions, it also be of interest to estimate its p.m.f., $\{\pi_j, j = 0, 1, \dots\}$, or its d.f. $\{\Gamma_l := \sum_{j=0}^l \pi_j; l = 0, 1, \dots\}$. Indeed, if the mean, μ , does not exist, then estimating $\{\pi_j\}$ or $\{\Gamma_l\}$ provides

the only other bases for characterization. Although μ exists for $1 < \alpha < 2$, the long sample path required to estimate it motivates us to investigate the estimation of π_0 , noting that analogous results hold for the other probabilities, starting from $X(0) = l$ for $l = 0, 1, \dots$

Let

$$U_0(t) = \begin{cases} 1 & \text{if } X(t) = 0 \\ 0 & \text{otherwise,} \end{cases}$$

so that for $t, h \geq 0$,

$$E[U_0(t)|X(0) = 0] = p_{00}(t) = e^{-\mu+\mu(t)}$$

$$\begin{aligned} \text{cov}[U_0(t), U_0(t+h)|X(0) = 0] &= p_{00}(t)p_{00}(h) - p_{00}(t)p_{00}(t+h) \\ &= e^{-2\mu}e^{\mu(t)}[e^{\mu(h)} - e^{\mu(t+h)}]. \end{aligned}$$

As an estimate of π_0 , we have

$$\hat{\pi}_0(t, \tau) = \frac{1}{\tau} \int_t^{t+\tau} U_0(h)dh,$$

with

$$E[\hat{\pi}_0(t, \tau)|X(0) = 0] = \frac{1}{\tau} \int_t^{t+\tau} e^{-\mu+\mu(h)} dh \quad (18)$$

and

$$\text{var}[\hat{\pi}_0(t, \tau)|X(0) = 0] = \frac{2}{\tau^2} \int_0^\tau \int_t^{t+\tau-h} e^{-2\mu}e^{\mu(t)}[e^{\mu(h)} - e^{\mu(s+h)}]dsdh.$$

Hence, as $t \rightarrow \infty$, we find

$$\begin{aligned} E\hat{\pi}_0(\infty, \tau) &= e^{-\mu} \\ \text{var}\hat{\pi}_0(\infty, \tau) &= \frac{2e^{-2\mu}}{\tau^2} \int_0^\tau (\tau - h) [e^{\mu(h)} - 1] dh. \end{aligned} \quad (19)$$

Since $e^{\mu(h)} - 1 = O(\mu(h))$, $h \geq 0$, we immediately see that $\text{var}\hat{\pi}_0(\infty, \tau)$ behaves the same as $\text{var}\hat{\mu}(\infty, \tau)$ in (16) as $\tau \rightarrow \infty$. Therefore, for $1 < \alpha < 2$, slow convergence holds here as well as for the mean.

4 Lomax Distribution

This section describes the case of the Lomax service-time distribution.

Proposition 4.1 *Suppose that S_1, S_2, \dots has the Lomax Distribution distribution, $\mathcal{L}(\alpha, k)$, so that $\omega^{-1} = k/(\alpha - 1)$. Then for $h \geq 0$ and $r = 0, 1, \dots$,*

$$\begin{aligned} \mu(h) &= \frac{\mu}{(h/k+1)^{\alpha-1}} \\ p_{0r}(h) &= \pi_r \left[1 - \frac{1}{(h/k+1)^{\alpha-1}}\right]^r e^{\mu/(h/k+1)^{\alpha-1}} \\ E[X(t)|X(0) = 0] &= \mu \left[1 - \frac{1}{(t/k+1)^{\alpha-1}}\right] \\ E[\hat{\mu}(t, \tau)|X(0) = 0] &= \mu \left\{1 + \frac{k}{\tau(\alpha-2)} [(t/k + \tau/k + 1)^{2-\alpha} - (t/k + 1)^{2-\alpha}]\right\} \\ &= \begin{cases} \mu \left[1 + O\left(\left(\frac{k}{\tau}\right)^{\alpha-1}\right)\right] & \text{if } 1 < \alpha < 2 \\ \mu \left[1 - \frac{k}{\tau} \ln\left(1 + \frac{\tau}{t+k}\right)\right] & \text{if } \alpha = 2 \\ \mu \left[1 + O\left(\frac{k}{\tau}\right)\right] & \text{if } \alpha > 2 \end{cases} \\ \text{var}\hat{\mu}(\infty, \tau) &= \frac{2\mu}{(\alpha-2)(\alpha-3)} \left(\frac{k}{\tau}\right)^2 \left[(\tau/k + 1)^{3-\alpha} + (\alpha - 3)\frac{\tau}{k} - 1\right] \\ &= \begin{cases} O\left(\left(\frac{k}{\tau}\right)^{\alpha-1}\right) & \text{if } 1 < \alpha < 2 \\ O\left(\frac{\ln(\tau/k+1)}{\tau/k}\right) & \text{if } \alpha = 2 \\ O\left(\frac{k}{\tau}\right) & \text{if } \alpha > 2 \end{cases} \\ E[\hat{\mu}(\infty, \tau) - \mu]^3 &= O\left(\left(\frac{k}{\tau}\right)^{\min(2, \alpha-1)}\right) \\ &\hspace{15em} \alpha \neq 2, \\ E[\hat{\mu}(\infty, \tau) - \mu]^4 &= O\left(\left(\frac{k}{\tau}\right)^{\min(2, \alpha-1)}\right) \end{aligned}$$

where all limiting behavior applies as $\tau \rightarrow \infty$ for fixed $t \geq 0$. For fixed $\tau > 0$

Table 1: Controlling bias for M/G/ ∞ : service times from $\mathcal{L}(\alpha, k)$

$$t_{\mathcal{L}}(\theta) = \min\{t \geq 0 : |\mathbb{E}[X(t)|X(0) = 0] - \mu| \leq \theta\mu\} = k[\theta^{\frac{1}{1-\alpha}} - 1]$$

$$t_{\mathcal{L}}(\theta)/k$$

α	$\theta = .01$	$\theta = .05$	$\theta = .10$
1.2	10^{10}	3.2×10^6	10^5
1.4	10^5	1,788	315
1.6	2,153	146	45
1.8	315	41	17

and $1 < \alpha < 2$,

$$|\mathbb{E}[\widehat{\mu}(t, \tau)|X(0) = 0] - \mu| \leq \frac{\mu}{(t/k + 1)^{\alpha-1}}. \quad \square$$

The proof follows by direct evaluation. The same convergence rates hold for other heavy-tailed distributions satisfying (1).

As illustration of the slow convergence rate $O\left((k/\tau)^{\alpha-1}\right)$ for $1 < \alpha < 2$, Table 1 shows the minimal required warm-up interval length, $t_{\mathcal{L}}(\theta)$, for a simulation, begun in the empty and idle state $X(0) = 0$, to observe the number in the system, $X(t)$, for $t \geq t_{\mathcal{L}}(\theta)$ with an absolute bias no greater than $\theta\mu$. The sensitivity to α is notable.

Table 2 displays the substantial sample-path lengths

$$\tau_{\mathcal{L}}(\nu) := \min\{\tau \geq 0 : \text{var}\widehat{\mu}(\infty, \tau) \leq \mu\nu\}$$

required for high ($\nu = 10^{-4}$), moderate ($\nu = 10^{-3}$), and gross ($\nu = 10^{-2}$) accuracies, where $k/t_{\mathcal{L}}(\nu)$ is the (positive) solution of

$$\left(\frac{k}{\tau}\right)^2 \left[\left(\frac{\tau}{k} + 1\right)^{3-\alpha} + (\alpha - 3)\frac{\tau}{k} - 1 \right] = \frac{\nu(\alpha - 2)(\alpha - 3)}{2}.$$

By contrast, a sample average number in system based on a single observation on each of n independent replications in equilibrium has variance $\frac{\mu}{n}$, requiring

Table 2: Accuracy for the sample mean for M/G/ ∞ : service times from $\mathcal{L}(\alpha, k)$

α	$\tau_{\mathcal{L}}(\nu)/k$		
	$\nu = 10^{-4}$	$\nu = 10^{-3}$	$\nu = 10^{-2}$
1.2	5.2×10^{20}	5.2×10^{15}	5.2×10^{10}
1.4	6.3×10^{10}	2.0×10^8	6.3×10^5
1.6	3.9×10^7	8.3×10^5	1.7×10^4
1.8	1.3×10^6	6.7×10^4	3.2×10^3

$n = 10,000$ for $\nu = 10^{-4}$ and $n = 100$ for $\nu = .01$. Since each of these replications would also require a substantial warm-up interval, as in Table 1, to achieve approximate equilibrium, it is not clear whether one long sample path or multiple independent replications gives a specified accuracy for less computing time.

With regard to asymptotic normality, recall that a standard normal variate Z has mean 0, variance 1, skewness

$$\frac{\mathbb{E}(Z - \mathbb{E}Z)^3}{[\mathbb{E}(Z - \mathbb{E}Z)^2]^{3/2}} = 0,$$

and kurtosis

$$\frac{\mathbb{E}(Z - \mathbb{E}Z)^4}{[\mathbb{E}(Z - \mathbb{E}Z)^2]^2} = 3.$$

For $1 < \alpha < 2$, the distribution of the sample-path average $\hat{\mu}(\infty, \tau)$ has skewness $O(\tau^{(\alpha-1)/2})$ and kurtosis $O(\tau^{(\alpha-1)})$ as $\tau \rightarrow \infty$. Hence, asymptotic normality does not hold for $[\hat{\mu}(\infty, \tau) - \mu]/\sqrt{\text{var}\hat{\mu}(\infty, \tau)}$ for $1 < \alpha < 2$; the same is true for $2 < \alpha < 3$.

To evaluate the behavior of the mean and variance of $\hat{\pi}_0(t, \tau)$ for large τ , we use $e^{\mu(h)} = 1 + O(\mu(h))$, $h \geq 0$ in (18) and (19) and substitute $\mu(h) = (h/k + 1)^{1-\alpha}\mu$, yielding

$$\mathbb{E}[\hat{\pi}_0(t, \tau) | X(0) = 0] = e^{-\mu} \left[1 + O\left(\left(\frac{k}{\tau}\right)^{\min(1, \alpha-1)}\right) \right] \quad \alpha \neq 2$$

and

$$\text{var}\widehat{\pi}_0(\infty, \tau) = 2e^{-2\mu} \mathcal{O}\left(\left(\frac{k}{\tau}\right)^{\min(1, \alpha-1)}\right) \quad \alpha \neq 2.$$

Indeed, slow convergence holds here as well.

With regard to the warm-up interval, $t_{\mathcal{L}}(\theta)$, in Table 1,

$$\pi_0^{-1} \mathbb{E}[U_0(t_{\mathcal{L}}(\theta)) | X(0) = 0] = e^{\theta\mu} \geq 1 + \theta\mu.$$

5 Alternative Distributions

With the few exceptions for $\nu = 10^{-2}$, the entries in Tables 1 and 2 imply that generating file sizes from a heavy-tailed distribution in a discrete-event simulation of the M/G/ ∞ queue limits the statistical accuracy of sample averages obtainable within a practicable computing time. This realization motivates the search for an alternative distribution that fits empirical file-size data at least as well as a heavy-tailed one, but ensures considerably shorter simulated sample paths for sample averages with a specified accuracy. As mentioned in the introduction, the lognormal and Weibull distributions have been proposed as substitutes.

To compare how well two different distributions fit empirical file-size data, a statistical methodology such as maximum likelihood is first applied to estimate the distributions' parametric values. Then the fit is evaluated, sometimes, by formal hypothesis testing. More often, graphical techniques are applied. For example, a graph of a fitted d.f. versus the empirical d.f. (q-q plot) provides a convenient basis for evaluation. As to which of the two different fitted d.f.s gives the better fit, a comparison of their respective q-q plots can be helpful.

The present study calls for a distinctly different basis for comparison. Section 5.1 describes a way of choosing the parametric values for a lognormal or Weibull distribution that makes its d.f. closely resemble that of the Lomax with specified parametric values. Then we compare the sample-path lengths that the candidate and Lomax each require to meet a specified accuracy for the sample mean. Section 5.2 does this for the lognormal, and Sec. 5.3 for the Weibull. Section 5.4 assesses which of these gives the better approximation.

5.1 Approximating the Lomax Distribution

Let $\{F(x), x \geq 0\}$ denote the d.f. of file size, let $\overline{F}(x) := 1 - F(x)$, and let the subscripts \mathcal{L} and \mathcal{LN} , denote the Lomax and lognormal, distributions respectively. We require that each have the same means

$$\int_0^\infty \overline{F}_{\mathcal{L}}(x)dx = \int_0^\infty \overline{F}_{\mathcal{LN}}(x)dx. \quad (20)$$

Although

$$\frac{\overline{F}_{\mathcal{LN}}(x)}{\overline{F}_{\mathcal{L}}(x)} \rightarrow 0 \quad x \rightarrow \infty$$

for specified $\zeta \in (0, 1)$, we can control this right-tail ratio by equating the $1 - \zeta$ quantiles, $x(1 - \zeta)$, of the d.f.s

$$\overline{F}_{\mathcal{L}}(x(1 - \zeta)) = \overline{F}_{\mathcal{LN}}(x(1 - \zeta)) = \zeta, \quad (21)$$

subject to the condition

$$\zeta > \overline{F}_{\mathcal{L}}(x) > \overline{F}_{\mathcal{LN}}(x) \quad \forall x > x(1 - \zeta). \quad (22)$$

By way of interpretation, an empirical data set of $1/\zeta$ observations implies that just one observation exceeds $x(1 - \zeta)$. For example, one observation exceeds $x(.999)$ for a data set of 1,000 observations. For 100,000 observations, just one exceeds $x(.99999)$. In the present setting, we can choose ζ to fix the point beyond which $\overline{F}_{\mathcal{L}}(x(1 - \zeta))$ exceeds $\overline{F}_{\mathcal{LN}}(x(1 - \zeta))$.

5.2 Lognormal Distribution

As mentioned earlier, Downey (2001,2002) found good fits for file size using the lognormal distribution

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-[\ln(x/\eta)]^2/2\sigma^2} \quad \eta, \sigma > 0, x \geq 0$$

with

$$\text{mean} = \omega^{-1} = \eta e^{\sigma^2/2},$$

mode at $x = \eta e^{-\sigma^2}$, and variance $\eta^2 e^{\sigma^2} (e^{\sigma^2} - 1)$. Observe that the modal location tends to zero as σ increases. The corresponding equilibrium covariance function for the number in the system for the M/G/ ∞ model is

$$\mu(h) = \frac{\mu}{\sqrt{\pi}\eta e^{\sigma^2/2}} \int_h^\infty \Gamma\left(1/2, \frac{(\ln(x/\eta))^2}{2\sigma^2}\right) dx.$$

Also,

$$\tau \text{var}[\hat{\mu}(t, \tau) \mid \mathcal{S}_0] \rightarrow 2\mu\eta e^{3\sigma^2/2} \quad \tau \rightarrow \infty, \quad (23)$$

as a consequence of (17) and Proposition 3.2.

To compare sample-path lengths for the Lomax and lognormal distributions, we match the means

$$\frac{k}{\alpha - 1} = \eta e^{\sigma^2/2} \quad (24)$$

and the $1 - \zeta$ quantiles

$$x(1 - \zeta) = k(\zeta^{-1/\alpha} - 1) = \eta e^{c(1-\zeta)\sigma}, \quad (25)$$

where $c(1 - \zeta)$ denotes the $1 - \zeta$ point of the standard normal distribution. Combining (20) and (21) leads to

$$e^{[\sigma(\zeta, \alpha) - c(1-\zeta)]^2/2} = \frac{e^{c^2(1-\zeta)/2}}{(\alpha - 1)(\zeta^{-1/\alpha} - 1)},$$

which has two solutions only the smaller of which satisfies (20), (21), and (22).

Table 3 displays $\sigma(\zeta, \alpha)$ and k/η for $\alpha \in \{1.2, 1.4, 1.6, 1.8\}$ and $\zeta \in \{10^{-3}, 10^{-5}\}$. It also lists $\tau_{\mathcal{LN}}(\nu)$ to achieve $\text{var}\hat{\mu}(\infty, \tau_{\mathcal{LN}}(\nu)) \approx \nu\mu$, where

$$\tau_{\mathcal{LN}}(\nu) := 2\nu^{-1}\eta e^{3\sigma^2/2},$$

based on the asymptotic variance (23), denotes the required sample-path length that the lognormal induces. Table 4 lists the ratios of sample-path lengths

$$\begin{aligned} \text{ratio}_{\mathcal{L}, \mathcal{LN}}(\alpha, \sigma(\zeta, \alpha)) &:= \frac{\tau_{\mathcal{L}}(\nu)}{k} \times \frac{\eta}{\tau_{\mathcal{LN}}(\nu)} \times \frac{k}{\eta} \\ &= \frac{\tau_{\mathcal{L}}(\nu)}{\tau_{\mathcal{LN}}(\nu)}. \end{aligned}$$

They reveal that

Table 3: Accuracy for the sample mean for M/G/ ∞ : service times from $\mathcal{LN}(\eta, \sigma)$

ζ	α	$\sigma(\zeta, \alpha)$	$\frac{k}{\eta}$	$\tau_{\mathcal{LN}}(\nu)/\eta$		
				$\nu = 10^{-4}$	$\nu = 10^{-3}$	$\nu = 10^{-2}$
10^{-3}	1.2	1.97	1.38	6.6×10^6	6.6×10^5	6.6×10^4
	1.4	1.85	2.23	3.5×10^6	3.5×10^5	3.5×10^4
	1.6	1.69	2.50	1.4×10^6	1.4×10^5	1.4×10^4
	1.8	1.55	2.67	7.4×10^5	7.4×10^4	7.4×10^3
10^{-5}	1.2	2.77	9.40	2.1×10^9	2.1×10^8	2.1×10^7
	1.4	2.37	6.70	9.4×10^7	9.4×10^6	9.4×10^5
	1.6	2.07	5.10	1.2×10^7	1.2×10^6	1.2×10^5
	1.8	1.85	4.40	3.3×10^6	3.3×10^5	3.3×10^4

- a. With the exception of the triplets $(\zeta, \alpha, \nu) = (10^{-5}, 1.6, 10^{-2})$, $(10^{-5}, 1.8, 10^{-3})$, and $(10^{-5}, 1.8, 10^{-2})$, the lognormal requires considerably shorter sample paths than the Lomax.
- b. The exceptions induce longer sample paths.
- c. The corresponding ratios are smaller for $\zeta = 10^{-5}$ than for $\zeta = 10^{-3}$, implying that longer sample paths, relative to the Lomax, are needed when employing the lognormal, to meet the more stringent tail constraint.

5.3 Weibull Distribution

As mentioned earlier, Cao et al. (2003) found that the Weibull distribution

$$f(x) = \phi \beta^{-\phi} x^{\phi-1} e^{-(x/\beta)^\phi} \quad 0 < \phi < 1, \beta > 0, x \geq 0$$

with

$$\text{mean} = \omega^{-1} := \beta \Gamma(1 + 1/\phi)$$

Table 4: Ratio of Lomax and lognormal sample-path lengths for $\text{var}\widehat{\mu}(\infty, \tau) \approx \nu\mu$

		ratio $_{\mathcal{L}, \mathcal{LN}}(\alpha, \sigma(\zeta, \alpha))$		
ζ	α	$\nu = 10^{-4}$	$\nu = 10^{-3}$	$\nu = 10^{-2}$
10^{-3}	1.2	1.1×10^{14}	1.1×10^{10}	1.1×10^6
	1.4	4.0×10^4	1.3×10^3	40
	1.6	67	14	3.0
	1.8	4.6	2.4	1.1
10^{-5}	1.2	2.3×10^{12}	2.3×10^8	2.3×10^4
	1.4	4.5×10^3	1.4×10^2	4.5
	1.6	16	3.4	.71
	1.8	1.7	.88	.42

and infinite mode at $x = 0$ provided a good fit for file size. The small ϕ induces slow convergence in the right tail, but at an exponential rather than a polynomial rate in x . For fixed scale parameter β , the distribution tends to give weight to small and large service times. Moreover, it has finite variance, $\beta^2 [\Gamma(1 + 2/\phi) - \Gamma^2(1 + 1/\phi)]$.

The corresponding equilibrium covariance function is

$$\mu(h) = \mu \frac{\Gamma(1/\phi, (h/\beta)^\phi)}{\Gamma(1/\phi)}, \quad (26)$$

where $\mu = \lambda/\omega$, as before, and

$$\Gamma(a, z) := \int_z^\infty y^{a-1} e^{-y} dy \quad a > 0, z \geq 0.$$

Since $\Gamma(a, z) \sim z^{a-1} e^{-z}$ as $z \rightarrow \infty$, it follows from (26) that the bias $\text{E}[X(h)|X(0) = 0] - \mu = \mu(h)$ decays at an exponential rate $(h/\beta)^\phi$. Moreover, for fixed $t \geq 0$,

$$\tau \{ \mu - \text{E}[\widehat{\mu}(t, \tau) | X(0) = 0] \} \rightarrow \int_t^\infty \mu(h) dh \quad (< \kappa) \quad \text{as } \tau \rightarrow \infty$$

and

$$\tau \text{var}[\hat{\mu}(t, \tau) \mid \mathcal{S}_0] \rightarrow \kappa \quad \text{as } \tau \rightarrow \infty, \quad (27)$$

where

$$\kappa := 2\mu\beta \frac{\Gamma(2/\phi)}{\Gamma(1/\phi)} = \mu\beta \frac{2^{2/\phi}}{\sqrt{\pi}} \Gamma(1/\phi + 1/2) < \infty,$$

limits that follow from (14) and (17).

Let $\{F_{\mathcal{W}}(x), x \geq 0\}$ denote the Weibull d.f. and let $\bar{F}_{\mathcal{W}}(x) := 1 - F_{\mathcal{W}}(x)$. To approximate the Lomax by the Weibull, we again invoke (20), (21), and (22), but with $F_{\mathcal{W}}(\cdot)$ replacing $F_{\mathcal{LN}}(\cdot)$. In particular, (20) and (21) give

$$\frac{k}{\alpha - 1} = \beta \Gamma(1 + 1/\phi) \quad (28)$$

and

$$\frac{k^\alpha}{[x(1 - \zeta) + k]^\alpha} = e^{-[x(1 - \zeta)/\beta]^\phi} = \zeta. \quad (29)$$

For fixed α , one can solve

$$\frac{(-\ln \zeta)^{1/\phi}}{\Gamma(1 + 1/\phi)} = (\alpha - 1)(\zeta^{-1/\alpha} - 1) \quad (30)$$

for $\phi = \phi(\zeta, \alpha)$, and that fixes the ratio

$$\frac{k}{\beta} = (\alpha - 1) \Gamma \left(1 + \frac{1}{\phi(\zeta, \alpha)} \right). \quad (31)$$

The left side of (30) is monotone increasing in ϕ for $\psi(1 + 1/\phi) < \ln(-\ln \zeta)$ and monotone decreasing for $\psi(1 + 1/\phi) > \ln(-\ln \zeta)$, where ψ denotes the digamma function, $\{\psi(z) = \frac{d \ln \Gamma(z)}{dz}, z > 0\}$. In what follows, we choose the $\phi(\zeta, \alpha)$ for which $\psi(1 + 1/\phi) > \ln(-\ln \zeta)$. It is the only solution to satisfy (22) as well as (28) and (29).

Table 5 displays $\tau_{\mathcal{W}}(\nu)$ that achieve $\text{var} \hat{\mu}(\infty, \tau_{\mathcal{W}}(\nu)) \approx \nu \mu$ for selected ν , where

$$\tau_{\mathcal{W}}(\nu) := \frac{\beta 2^{2/\phi} \Gamma(1/\phi + 1/2)}{\nu \sqrt{\pi}},$$

based on the asymptotic variance (27), denotes the sample-path length that the Weibull induces.

Table 5: Accuracy for the sample mean for M/G/ ∞ : service times from $\mathcal{W}(\phi, \beta)$

ζ	α	$\phi(\zeta, \alpha)$	$\frac{k}{\beta}$	$\tau_{\mathcal{W}}(\nu)/\beta$		
				$\nu = 10^{-4}$	$\nu = 10^{-3}$	$\nu = 10^{-2}$
10^{-3}	1.2	0.311	1.58	2.0×10^6	2.0×10^5	2.0×10^4
	1.4	0.333	2.41	1.2×10^6	1.2×10^5	1.2×10^4
	1.6	0.370	2.51	5.8×10^5	5.8×10^4	5.8×10^3
	1.8	0.407	2.53	3.3×10^5	3.3×10^4	3.3×10^3
10^{-5}	1.2	.171	106	3.9×10^9	3.9×10^8	3.9×10^7
	1.4	.207	36.4	1.9×10^8	1.9×10^7	1.9×10^6
	1.6	.244	16.6	2.2×10^7	2.2×10^6	2.2×10^5
	1.8	.280	10.3	5.2×10^6	5.2×10^5	5.2×10^4

The entries in Tables 2 and 5 provide a basis for comparing the required lengths of the sample paths, $\tau_{\mathcal{L}}(\nu)$ and $\tau_{\mathcal{W}}(\nu)$, to achieve accuracy ν . In particular, Table 6 lists the ratios

$$\begin{aligned} \text{ratio}_{\mathcal{L}, \mathcal{W}}(\alpha, \phi(\zeta, \alpha)) &:= \frac{\tau_{\mathcal{L}}(\nu)}{k} \times \frac{\beta}{\tau_{\mathcal{W}}(\nu)} \times \frac{k}{\beta} \\ &= \frac{\tau_{\mathcal{L}}(\nu)}{\tau_{\mathcal{W}}(\nu)}. \end{aligned}$$

They reveal:

- d. With the exception of $(\zeta, \alpha, \nu) = (10^{-5}, 1.8, 10^{-2})$, the Weibull sample-path lengths are considerably shorter than those for the Lomax.
- e. The triplet $(\zeta, \alpha, \nu) = (10^{-5}, 1.8, 10^{-2})$ requires a longer sample path than the Lomax.
- f. For $\zeta = 10^{-3}$, the Weibull has longer sample paths than the lognormal. For $\zeta = 10^{-5}$, its paths are shorter.

Given the method of approximation, it is inevitable that $\tau_{\mathcal{L}}(\nu) < \tau_{\mathcal{LN}}(\nu)$ and $\tau_{\mathcal{L}}(\nu) < \tau_{\mathcal{W}}(\nu)$ for some ν . We address this issue for the Weibull, noting

Table 6: Ratio of Lomax and Weibull sample-path lengths for $\text{var}\hat{\mu}(\infty, \tau) \approx \nu\mu$

		ratio $_{\mathcal{L},\mathcal{W}}(\alpha, \phi(\zeta, \alpha))$		
ζ	α	$\nu = 10^{-4}$	$\nu = 10^{-3}$	$\nu = 10^{-2}$
10^{-3}	1.2	4.0×10^{14}	4.0×10^{10}	4.0×10^6
	1.4	1.2×10^5	4.0×10^3	1.2×10^2
	1.6	1.7×10^2	36	7.4
	1.8	10	5.2	2.5
10^{-5}	1.2	1.4×10^{13}	1.4×10^9	1.4×10^5
	1.4	1.2×10^4	3.9×10^2	12
	1.6	29	6.3	1.3
	1.8	2.5	1.3	.63

that an analogous development holds for the lognormal. Recall that all three file-size distributions induce non-negative monotone decreasing covariance functions. As Fig. 1 shows for $\zeta = 10^{-5}$, there exists an $h_0 > 0$ such that

$$\frac{\mu_{\mathcal{W}}(h)}{\mu_{\mathcal{L}}(h)} \begin{cases} \geq 1 & \text{if } 0 \leq h \leq h_0 \\ < 1 & \text{if } h > h_0. \end{cases}$$

More generally, h_0 is monotone decreasing in ζ . Therefore, it follows that there exists a $\tau_0 > h_0$ such that

$$\frac{\text{var}\hat{\mu}_{\mathcal{W}}(\infty, \tau)}{\text{var}\hat{\mu}_{\mathcal{L}}(\infty, \tau)} \begin{cases} \geq 1 & \text{if } \tau \leq \tau_0 \\ < 1 & \text{if } \tau > \tau_0. \end{cases}$$

Since $\tau_{\mathcal{L}}(\nu)$ and $\tau_{\mathcal{W}}(\nu)$ are monotone decreasing in ν , it follows that for some moderate and large ν the Lomax induces a shorter path (i.e., $\tau_{\mathcal{L}}(\nu) < \tau_0$).

One additional issue may play a role in the present study. For the Lomax, the exact variance formula in Proposition 4.1 was used, but for the lognormal,

Figure 1: Covariance ratio $\frac{\mu_{\mathcal{W}}(h)}{\mu_{\mathcal{L}}(h)}$ for $\zeta = 10^{-5}$





Figure 2: Tail-probability ratios for $\alpha = 1.2$ and $\zeta = 10^{-5}$



the asymptotic formula (23), which is larger than the corresponding exact variance, was used. Although this last distinction becomes negligible as τ increases, it could be of significance in the present study for $\nu = 10^{-2}$. Nevertheless, there is always a ν at which $\tau_{\mathcal{L}}(\nu) < \tau_{\mathcal{W}}(\nu)$.

5.4 How Good Are the Approximations?

Although the generally smaller sample sizes that lognormal and Weibull file-size distributions induce come as no surprise, their use in practice also needs a statistical justification. If either fits the data unequivocally better than the Lomax, then the choice is clear. If statistical methodology fails to detect a significant difference in fit (especially in the right tail), then the cost-saving option has an obvious attraction.

If a heavy-tailed distribution fits data better than either the lognormal or Weibull, then two options remain. Either use the heavy-tailed distribution and incur excessively long sample paths, or approximate the fitted heavy-tailed distribution by either the lognormal or Weibull. This section investigates this second option when the parameters of the approximating distribution are assigned their values by the methodology of Sec. 5.1 or by some related method. In particular, how well do these distributions approximate the Lomax and which of these gives the better approximation?

It is easily seen that

$$\overline{F}_{\mathcal{W}}(0) = \overline{F}_{\mathcal{LN}}(0) = \overline{F}_{\mathcal{L}}(0) = 1$$

$$\left. \frac{d}{dx} \left\{ \frac{\overline{F}_{\mathcal{W}}(x)}{\overline{F}_{\mathcal{L}}(x)} \right\} \right|_{x=0} \leq \left. \frac{d}{dx} \left\{ \frac{\overline{F}_{\mathcal{LN}}(x)}{\overline{F}_{\mathcal{L}}(x)} \right\} \right|_{x=0} < 0.$$

$$\left. \frac{d}{dx} \left\{ \frac{\overline{F}_{\mathcal{W}}(x)}{\overline{F}_{\mathcal{L}}(x)} \right\} \right|_{x=x(1-\zeta)} < \left. \frac{d}{dx} \left\{ \frac{\overline{F}_{\mathcal{LN}}(x)}{\overline{F}_{\mathcal{L}}(x)} \right\} \right|_{x=x(1-\zeta)} < 0$$

$$\overline{F}_{\mathcal{W}}(x) < \overline{F}_{\mathcal{LN}}(x) < \overline{F}_{\mathcal{L}}(x) \quad \forall x > x(1-\zeta).$$

However, Table 7 and the graphs in Fig. 2 for $\alpha = 1.2$ and $\zeta = 10^{-5}$ best describe the relationships among $\overline{F}_{\mathcal{W}}$, $\overline{F}_{\mathcal{LN}}$, and $\overline{F}_{\mathcal{L}}$. In particular,

Table 7: Comparison of the lognormal, Weibull and Lomax distributions

Based on fit for $\alpha = 1.2$ and $\zeta = 10^{-5}$

θ	$x/k = \overline{F}_{\mathcal{L}}^{-1}(\theta)$	$\frac{\overline{F}_{\mathcal{LN}}(x)}{\overline{F}_{\mathcal{L}}(x)}$	$\frac{\overline{F}_{\mathcal{W}}(x)}{\overline{F}_{\mathcal{L}}(x)}$
.99	.008410	.8281	.3791
.95	.04367	.6588	.2869
.90	.09177	.5792	.2539
.80	.2044	.5088	.2299
.70	.3461	.4789	.2238
.60	.5306	.4688	.2268
.50	.7818	.4723	.2374
.40	1.146	.4896	.2568
.30	1.727	.5253	.2901
.20	2.824	.5935	.3511
.10	5.813	.7469	.4949
.05	11.14	.9371	.6943
.01	45.42	1.453	1.390
.001	315.2	1.983	2.579
.0001	2153	1.763	2.522

Table 7 shows that the Lomax puts less probability on small file size than either alternative, but over a substantial part of the interval $[45k, x(1 - 10^{-5}) = 14676k]$, which corresponds to $.99 < F_{\mathcal{L}}(x) < .9999$, both of these distributions put more probability in their tails than the Lomax. For the lognormal alternative, this excess is as much as $\frac{\overline{F}_{\mathcal{LN}}(x)}{\overline{F}_{\mathcal{L}}(x)} = 1.97$ at $x/k = 400$, where $\overline{F}_{\mathcal{L}}(x) = 5.8 \times 10^{-4}$. For the Weibull, the excess can be as much as $\frac{\overline{F}_{\mathcal{W}}(x)}{\overline{F}_{\mathcal{L}}(x)} = 2.81$ at $x/k = 800$, where $\overline{F}_{\mathcal{L}}(x) = 3.3 \times 10^{-4}$. These results indicate that the lognormal better approximates the Lomax than the Weibull does in the interval $[0, x(1 - 10^{-5})]$. Similar relationships were observed for $\alpha = 1.4, 1.6, 1.8$ at both $\zeta = 10^{-3}$ and 10^{-5} .

6 References

Beneš, V.E. (1957). Fluctuations of telephone traffic, *Bell System Technical Journal*, **36**, 965-973.

Beran, J., R. Sherman, M.S. Taquq, and W. willinger (1995). Long-range dependence in variable-bit-rate video, *IEEE Trans. Commun.*, **43**, 1566-1579.

Cao, J., W.S. Cleveland, D. Lin, and D.X. Sun (2003). Internet traffic tends toward Poisson and independent as the load increases, in *Nonlinear Estimation and Classification*, C. Holmes, D. Denison, M. Hansen, B. Yu, and B. Mallick, editors, Springer-Verlag, New York.

Crovella, M.E. and A. Bestavros (1996). Self-similarity in World Wide Web traffic, *Proceedings of the 1996 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 160-169.

Crovella, M.E. and L. Lipsky (1997). Long-lasting transient conditions in simulations with heavy-tailed workloads, *Proceedings of the 1997 Winter Simulation Conference*, S. Andridottir, K.J. Healy, D.H. Withers, and B.L. Nelson, editors.

Downey, A.B. (2001). The structural cause of file size distributions, *IEEE MASCOTS*, 361-370.

Downey, A.B. (2002). Evidence for long-tailed distributions in the Internet, Computer Science Department, Wellesley College, Wellesly, MA.

Feldman, A. and W. Whitt (1997) Fitting mixtures of exponentials to long-tailed distributions to analyze network performance models, *IEEE*

INFORCOM '97.

Fishman, G. S. (2001). *Discrete-Event Simulation: Modeling, Programming, and Analysis*, Springer-Verlag, New York.

Gnedenko, B.V. and I.N. Kovalenko (1987). *Introduction to Queueing Theory*, second edition, Birkhäuser, Boston.

Grossglauser, M. and J.-C. Bolot(1999). On the relevance of long-range dependence in network traffic, *IEEE/ACM Transactions on Networking*, **7**, 629-640.

Johnson, N.L., S. Kotz, and N. Balakrishnan (1994). *Continuous Univariate Distributions, Volume II*, second edition, Wiley, New York.

Leland, W.E., M.S. Taqqu, W. Willinger, and D.V. Wilson (1994). On the self-similar nature of Ethernet traffic (extended version), *IEEE/ACM Transactions on Networking*, **2**, 1-15.

Paxson, V. and S. Floyd (1995). Wide area traffic: the failure of Poisson modeling, *IEEE/ACM Trans. Networking* **3**, 226-244.

Paxson, V. and S. Floyd (1997). Why we don't know how to simulate the Internet, *Proceedings of the 1997 Winter Simulation Conference*, S. Andridottir, K.J. Healy, D.H. Withers, and B.L. Nelson, editors.

Resnick, S. and G. Samorodnitsky (2000). A heavy traffic approximation for workload processes with heavy tailed service requirements, *Management Science*, **46**, 1236-1248.

Reynolds, J.F. (1975). The covariance structure of queues and related processes—a survey of recent work, *Adv. Appl. Prob.* **7**, 383-415.

Riordan, J. (1951). Telephone traffic time averages, *Bell System Technical Journal*, **30**, 1129-1144.

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

7 Appendix

For the M/G/ ∞ queue, $h \geq 0$, and $l, r = 0, 1, \dots$,

$$\begin{aligned} \text{pr}[X(t+h) = r | X(t) = l] &= p_{lr}(h) \\ &:= \frac{e^{-\mu+\mu(h)}}{\mu^l} \sum_{j=0}^{\min(l,r)} \frac{l! [\mu-\mu(h)]^{l+r-2j}}{(l-j)!(r-j)!j!} \\ \text{pr}[X(t) = r | X(0) = l] &= p_{lr}(t) \\ &:= \frac{e^{-\mu+\mu(t)}}{\mu^l} \sum_{j=0}^{\min(l,r)} \binom{l}{j} \frac{[\mu-\mu(t)]^{l+r-2j} [\mu(t)]^j}{(r-j)!} \\ \text{E}[X(t) | X(0) = l] &= l \frac{\mu(t)}{\mu} + \mu - \mu(t) \end{aligned}$$

$$\begin{aligned} \text{E}[X(t)X(t+h) | X(0) = l] &= l(l-1) \frac{\mu(t)}{\mu} + l \left[\frac{\mu(t)}{\mu} + \frac{\mu(t+h)}{\mu} \right] \\ &\quad + [\mu - \mu(t)] [\mu - \mu(t+h)] + \mu(h) - \mu(t+h), \end{aligned} \tag{32}$$

all of which follow from the generating functions in Riordan (1951).

All means, variance, and covariances of interest are derivable from these expressions. As illustration, the Lomax distribution gives

$$\text{E}[\hat{\mu}(t, \tau) | X(0) = l] = \mu \left(\frac{l}{\mu} - 1 \right) \frac{k}{\tau(\alpha - 2)} \left[(t/k + 1)^{2-\alpha} - (t/k + \tau/k + 1)^{2-\alpha} \right] + \mu.$$

George S. Fishman George Fishman is professor emeritus of operations research at the University of North Carolina at Chapel Hill. His research interests are in the design of large-scale Monte Carlo and discrete-event simulation experiments and in the analysis of the data they generate.

Ivo Adan Ivo Adan received the M.Sc. and Ph.D. degrees in Mathematics from the Eindhoven University of Technology. Since 2000 he works at this university as associate professor in the Department of Mathematics and Computer Science. His current research interests are in the analysis of Markov processes and queueing models and in the performance evaluation of communication, production and warehousing systems.