Proceedings of the 2001 Winter Simulation Conference
B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, eds.

USING QUANTILE ESTIMATES IN SIMULATING INTERNET QUEUES WITH PARETO SERVICE TIMES

Martin J. Fischer
Denise M. Bevilacqua Masi

Center for Telecommunications and Advanced Technology
Mitretek Systems
McLean, VA 22102, U.S.A.

Donald Gross
John Shortle

Dept. of Systems Engineering and Operations Research
George Mason University
Fairfax, VA 22030, U.S.A

Percy H. Brill

Department of Management Science
University of Windsor
Windsor, Ontario N9B 3P4, CANADA

ABSTRACT

It is readily apparent how important the Internet is to modern life. The exponential growth in its use requires good tools for analyzing congestion. Much has been written recently asserting that classical queueing models assuming Poisson arrivals or exponential service cannot be used for the accurate study of congestion in major portions of the Internet. Internet traffic data indicate that heavy-tailed distributions (e.g., Pareto) serve as better models in many situations for packet service lengths. But these distributions may not possess closed-form analytic Laplace transforms; hence, much of standard queueing theory cannot be used. Simulating such queues becomes essential; however, previous research pointed out difficulties in obtaining the usual moment performance measures such as mean wait in queue. In this paper, we investigate using quantile estimates of waiting times (e.g., median instead of mean), which appear to be considerably more efficient when service times are Pareto.

1 INTRODUCTION

In this study, we look at simulating M/G/1 queues with Pareto service (M/P/1). Fowler (1999) gives an extensive table showing heavy-tailed distributions found in various places within the Internet, including five of the seven OSI protocol levels. For example, FTP transfers (application level) and session durations (session level) are among those mentioned as having Pareto distributions. Requested file sizes, reading time per page (both at the application level), and packet interarrival times (network level) are listed as having heavy-tailed distributions. Naldi (1999) does an empirical study of dial-up to an ISP and shows that while arrival times appear Poisson, the holding times are heavy-tailed. Paxson (2000) discussed the tremendous diversity in the Internet, but mentioned that one invariant is the presence of heavy-tailed distributions. Thus, we study here an M/P/1 model. Fischer, et al (2001) showed that problems arose in obtaining mean wait in queue (Wq) in simulating M/P/1 queues. Point estimates of Wq, when compared to the Pollaczek-Khintchine (PK) formula were as much as 50% off for run sizes as large as 20,000,000 transactions, and confidence interval estimates using batch means did not come close to covering the PK value. A preliminary study using the median (see Fischer et al, op cit) appeared to be considerably more accurate when compared to a numerical technique called the Transform Approximation Method or TAM (see Brill, et al 2000), even for run sizes as low as 5,000,000. Here, we further investigate the median and look at estimating the .8, .9, and .99 quantiles as well.

2 THE PARETO DISTRIBUTION

The Pareto distribution is a power-tailed distribution. A cumulative distribution function, $F(x)$, has a *power tail* if there exists positive constants c and a such that for $\overline{F}(x) = 1 - F(x)$

$$\lim_{x \to \infty}\left[\overline{F}(x)\right] = cx^{-a} .$$

That is, the tail decays geometrically in the limit (as opposed to the more familiar exponential decay of the exponential, and gamma, for example). Power-tail distributions are a subset of a broader class of distributions whose tails decay slower than exponential, i.e.,

$$\lim_{x \to \infty} \left\lfloor e^{ax} \overline{F}(x) \right\rfloor = \infty \ .$$

This broader class is referred to as heavy-, fat- or long-tailed distributions, and include the lognormal and the Weibull (with shape parameter <1). These latter distributions have tails that decay slower than any exponential, but not as slow as the Pareto, so that a power-tailed distribution is also a heavy-tailed distribution, but not necessarily the reverse.

The standard form for the two-parameter (location and shape) Pareto distribution function defined over the nonnegative real numbers can be written as

$$F(x) = 1 - \frac{\gamma^{\alpha}}{\left[(\gamma + x)^{\alpha}\right]} \quad (\gamma, \alpha > 0; x \geq \gamma).$$

We note that the Pareto distribution function can be directly derived as a gamma mixture of ordinary exponential densities. With no loss in generality, henceforth we shall use the one-parameter (shape only) version of the Pareto given by

$$F(x) = 1 - \frac{1}{\left[(1 + x)^{\alpha}\right]} \qquad (x \geq 0)\,,$$

where α is the shape parameter. The corresponding density function is

$$f(x) = \frac{\alpha}{\left[(1 + x)^{\alpha+1}\right]} \qquad (x \geq 0),$$

and it is straightforward to show that the Pareto is indeed a power-tailed distribution.

A major consequence of power-tailed behavior is the disappearance of moments. It is easy to see that for a Pareto to have its k^{th} moment, $E[X^k]$, we need $\alpha > k$. If $\alpha > 1$, then the mean, $E[X]$, is

$$E[X] = 1/(\alpha - 1)$$

and if $\alpha > 2$, it follows that

$$E[X^2] = \frac{2}{\left[(\alpha - 2)(\alpha - 1)\right]} \ .$$

Thus, no matter what the value of the parameter α, a Pareto random variable cannot have all its moments and hence does not have an analytic Laplace transform, which renders standard queueing analyis of M/G/1 impossible. Hence the requirement for simulation, or other numerical techniques (such as TAM).

In this study, we consider Pareto distributions with α values of 3.5, 2.25, 2.083333, 2.05 and 2.020202, yielding CVs (coefficient of variation = standard deviation/ mean) of 1.53, 3, 5, 6.4, and 10 respectively. (Naldi (1999) in his study, found that holding time CVs generally ranged between 1.5 and 2.5). The traffic intensity was set to 0.8 for all runs. Arena and GPSS/H were the major simulation packages used. Resulting mean estimates of queue wait were compared to PK values and quantile estimates compared to values obtained from TAM.

3 PROBLEMS WITH ESTIMATING MEAN QUEUE WAIT

Figure 1 illustrates the difficulty in estimating mean wait in queue for the M/P/1 case, with Pareto shape parameter 2.083333 (CV=5). The estimates for mean queue wait, Wq, for Arena and GPSSH simulations versus the run length for run sizes from 500,000 to 20,000,000 are shown in the graph, along with the PK value. The graph shows that even at the longest run length, simulation estimates are about half the PK value.

Figure 2 shows percent errors from PK in estimating mean queue wait for all CV cases run. Except when CV=1.53 (Pareto shape parameter α = 3.5), the errors are very sizable, with the errors increasing for increasing CV (decreasing shape parameter α). We believe the problem for the $\alpha < 3$ cases stems from the fact that the Monte Carlo sampling from the Pareto greatly underestimates the variance.

Table 1 presents some Monte Carlo samples from various Pareto distributions, illustrating this. In the table, percent errors from the theoretical values for mean and variance are shown, and we see that for α = 2.05 (CV= 6.4), while the mean is fairly well estimated, the variance is over 50% underestimated. Further, it appears from Table 1 that in order to get decent estimates of the nth moment in a Monte Carlo sample, the distribution must possess at least one higher moment. For example, when α is 4.05, the first four moments of the Pareto exist and the estimates for both mean and standard deviation are quite good. For α = 3.05, the first three moments exist and estimates of the standard deviation are fairly good (except for the oddity of the one million run-length case). For α = 2.05, the standard deviation cannot be estimated sufficiently. When α = 1.05, even the mean cannot be estimated to any decent degree of accuracy. And, of course, for α = .8, the mean does not exist.

But the table also shows that good estimates of the median are obtainable, regardless of the existence of moments (errors far less than 1%). Further, good estimates of the .8, .9 and even the .99 quantiles are obtained, even for quite small run lengths. This leads us to believe that using quantile estimates for queue waits from a simulation run for M/P/1 queues might be preferable to moment estimates such as the mean.

Another reason for using quantile estimates for queue wait in an M/P/1 is that from theory (see Cohen (1969), page 256 or Gross and Harris (1998), page 228), in order for the kth moment of the waiting time distribution to exist, one must have the (k+1)st moment of the service time existing. For Pareto shape parameter $\alpha < 3$, the third moment does not exist and hence the variance of the queue wait does not exist yielding a poor estimate of the mean (see Figure 2, CV = 3,5,6.4 and 10 cases [α = 2.25 ,2.08. 2.05, 2.02 respectively]). But for α = 3.5 (CV=1.53), the third moment does exist and the mean is well estimated.

A final reason for using quantile estimates is that they provide more meaningful information as to congestion in a queueing situation, for often of more interest than mean waiting time is the tail of the waiting time distribution, about which, quantile estimates provide information.

4 M/P/1 QUANTILE ESTIMATION

Figure 3 shows results of simulation runs producing quantile estimates for an M/P/1 queue, with Pareto α = 2.083333 (CV=5). The simulated values are compared to values obtained from TAM (see section 1). Note how quickly the simulation quantile estimates settle down as run size increases and that for quite small run sizes, quantile estimates become very close to those from TAM. Both GPSSH and ARENA gave quite similar results, so we show only the ARENA plot as the simulated values.

The ARENA simulation run of 20 million transactions took a little over 30 minutes, with about another 15 minutes of manual-computer interaction to get the quantiles using the output analyzer histogram iteratively to zero in on the quantile values shown in Figure 3. GPSSH ran faster, taking about 8 minutes for the 20 million run length, while TAM took less than half a minute. Figure 3 indicates that a simulation run-length of one million appears sufficient, and for that size run length, ARENA takes under 1.5 minutes and GPSSH under 0.5 minutes. While TAM is certainly more efficient than simulation, it is not as versatile and cannot be used in more complex situations, such as networks. But the TAM and simulation results do appear to validate each other.

Table 2 shows results for TAM, ARENA and GPSSH runs for all the CV cases. Only the 20 million run-length simulation cases are shown. Note how close the values are for TAM and the two simulation packages, especially for the .5, .8, and .9 quantiles.

Table 2 also shows quantile values using an approximation due to Brichet et al. (1996). They show that the tail of the waiting time of an M/P/1 queue asymptotically follows

$$P(T > x) \sim \frac{\lambda}{(1-\rho)(\alpha-1)} \frac{1}{x^{\alpha-1}},$$

where $a(x) \sim b(x)$ means that $\lim_{x\to\infty} a(x)/b(x) = 1$, λ is the arrival rate, ρ is the traffic intensity, and α is the Pareto parameter. The table shows that the approximation for the .5, .8, and .9 is not very close to either simulation or TAM. In other words, the .9 quantile is not far enough out in the tail for the approximation to be accurate. The approximation begins to give better estimates for the .99 quantile. Preliminary numerical experiments with TAM show that the approximation becomes much more accurate past the .995 quantile. For example, for the CV = 3 case, TAM estimates the .995 quantile to be 229.5; simulation gives a value of 227 (ARENA, 20 million transactions). These both correspond to about the .9955 quantile in Brichet's approximation. Further out in the tail, TAM estimates the .999 quantile to be 722, simulation gives a value of 845, and inverting Brichet's approximation gives a value of 761. Plugging these values back into the approximation gives quantiles of .9989, .9991, and .9990 respectively. Thus, the tail is very flat, and a wide range of times give roughly the same quantile. Further out in the tail, the approximation may be a better option, since computing times become large for TAM and simulation. But, for lower quantiles like .9 or .95, numerical methods like simulation or TAM are needed.

5 CONFIDENCE INTERVALS FOR QUANTILES

Section 3 and 4 results show the difficulty in estimating moment-based measures for the M/P/1 queue, and the improved convergence in obtaining quantile estimates for the M/P/1. However, further work in output analysis for quantile-based estimators must be done. Wood and Schmeiser (1995) have done some work with overlapping batch quantiles, and conclude that very large sample sizes and batch sizes are needed to obtain reliable standard error estimators when using overlapping batch quantiles, even for i.i.d. data. Heidelberger and Lewis (1984) describe two extensions of the method of batch means to quantile estimation. Chen and Kelton (1999) discuss algorithms for quantile-based confidence interval estimation, but nothing has been done on queues with power- or heavy-tailed service. This section discusses the use of some of these estimation procedures with power-tailed service.

Chen and Kelton (1999) propose two methods of estimating confidence intervals for quantiles: a fixed sample size procedure, and a sequential procedure. Due to the complexity of their sequential procedure, we choose to focus on the fixed sample size procedure. The required sample size n_p for a fixed sample size procedure estimating the p quantile is the minimum n_p that satisfies

$$n_p \geq \frac{z_{1-\alpha/2}^2 p(1-p)}{(\varepsilon')^2}$$

where $z_{1-\alpha/2}$ is the 1 - $\alpha/2$ quantile of the standard normal distribution, ε' is the maximum proportion half-width of the confidence interval, and 1 - α is the confidence level. This sample size n_p ensures that the p quantile estimator is approximately between the p - ε' and p + ε' quantiles. Often in simulation studies the process $\{X_i\}$ is dependent (φ-mixing) instead of i.i.d., which is certainly the case with successive queue waiting times from simulation output. The sample size for φ-mixing processes needs to be further inflated by a factor which is proportional to $P_{x_p}(0)$, the initial point on the spectrum of a binary process $I_n(x_p)$, where $I_n(x) = 1$ if $X_n \leq x$. The quantity $P_{x_p}(0)$ is estimated using spectral methods (Heidelberger and Lewis, 1984), and uses least squares to fit a low order polynomial of degree d to the logarithm of the first K values of the averaged periodogram of $I_n(\hat{x}_p)$, where $\hat{x}_p$ is the estimated p quantile.

Here we use the Chen and Kelton fixed sample size procedure to select the run length n_p, and in addition perform independent replications of these runs to construct a confidence interval for the p quantile since we compute n_p using the assumption of i.i.d. data. Due to the complexity of estimating $P_{x_p}(0)$, we investigate the impact of omitting the inflation factor due to φ-mixing sequences. We select the M/P/1 queue, with CV=3, 5, and 10, ε' = .003, and the .5, .8, .9, and .99 quantiles for these experiments using GPSSH.

As seen in Table 3, for the .5, .8, and .9 quantiles at all CVs, the 95 percent confidence interval half-width covers the value obtained using one long run, although the half-widths are quite large, especially for the .9 quantile case. For the .99 quantile at CV=5 and CV=10, the 95 percent confidence interval half-width does not cover the value obtained using one long run, which is not surprising given the small sample size given by the Chen and Kelton formula. If the required inflation factor due to φ-mixing sequences had been incorporated, the confidence interval half-widths would likely be smaller since the number of observations per replication would have been greater. The run times were short, although manual-computer interaction was required; this time is not included. This manual-computer interaction consisted of getting the quantile value for each run from the output file using a GPSSH statement that allows one to obtain the distribution of residence time in a queue, and then copying to a spreadsheet. This takes about 15 minutes of additional time per quantile. An alternative simulator written in Visual Basic using Lindley's equation (Gross and Harris, 1998) to simulate the M/P/1 queue gave similar results. Use of Visual Basic had the advantage of eliminating the required manual-computer interaction, although run times were significantly longer.

Next we investigate use of methods of generating confidence intervals for quantiles, which do not require estimation of the inflation factor due to φ-mixing sequences. We perform independent replications of various lengths including n=100,000, based on the Chen and Kelton (1999) sample sizes computed above. In addition, n=500,000, and n=5,000,000 were run.

Table 4 indicates that using the longer run length of 500,000 observations per replication significantly tightens the 95 percent confidence intervals; half-widths are about one-fourth the magnitude of the half-widths at n=100,000. Also, at n=500,000 mean quantile values are closer to the quantile values obtained from a single long run of 20,000,000. As expected, the shorter run length of n=100,000 gives 95 percent confidence interval half-widths similar to the fixed sample size results in Table 3. Increasing the sample size to n=5,000,000 results in mean quantile values similar to those at n=500,000, and further decreases the confidence interval half-widths to about one-half the half-widths at n=500,000. Again, manual-computer interaction was required, and took about 40 minutes of additional time per CV investigated.

In summary, use of the fixed sample size procedure for estimating run lengths while omitting the inflation factor due to φ-mixing sequences does result in sample sizes that do not yield very good accuracy. Longer run lengths are required. Additionally, it would be very helpful for the simulation software packages to include in their output analysis an easy way to get quantile estimates so that manual-computer interaction between replications is not required.

6 CONCLUSIONS

In modeling congestion on the Internet, there are many places where holding time distributions are of the heavy- or power –tail type. We have studied the M/P/1 queue, where service times are Pareto – a power tail distribution, and found that the usual moment type measures of congestion (e.g., mean wait in queue) cannot be adequately estimated, even in very long simulation runs. However, much better results occur using quantile estimates (e.g., median), even for relatively short simulation runs. Further, quanitle estimates give more and better information regarding congestion (e.g., the percent of users that have to wait more than T time units in the queue). One problem with quantile estimates, however, is the difficulty in obtaining confidence interval estimates, and this is an area where some work has been done, but more is needed.

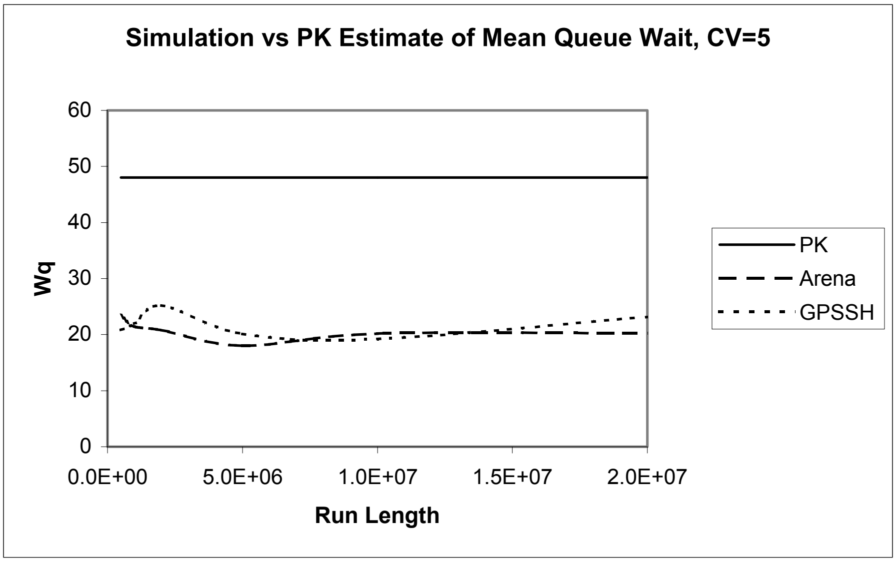

Figure 1: Simulation vs. PK Value, CV=5

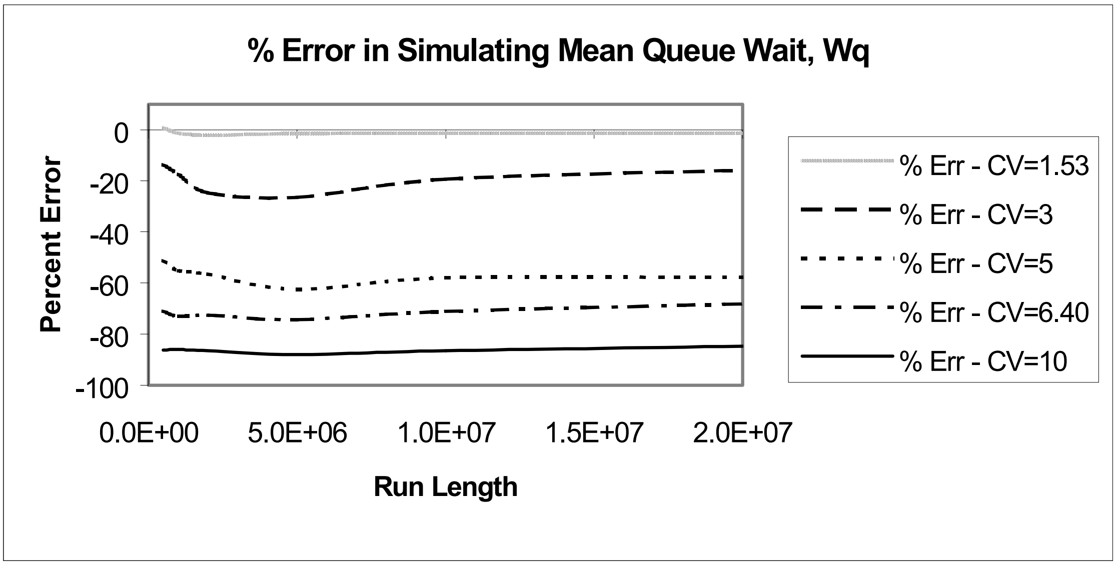

Figure 2: Errors in Simulating Mean Queue Wait vs. CV

Table 1: Monte Carlo Samples of Pareto Distributions

alpha	CV	No. Obs.	% error Mean	% error std. Dev	% error median	% error .8 q'tile value	% error .9 q'tile value	% error .99 q'tile value
0.8	inf	100,000	NA	NA	-0.47	-0.51	-0.33	-5.59
		500,000	NA	NA	-0.25	-0.24	-0.38	-0.58
		1,000,000	NA	NA	-0.18	-0.07	-0.38	-1.30
		5,000,000	NA	NA	-0.18	-0.14	-0.29	-0.92
1.05	inf	100,000	-59.00	NA	-0.43	-0.42	-0.26	-4.33
		500,000	-57.13	NA	-0.22	-0.20	-0.30	-0.45
		1,000,000	261.04	NA	-0.11	-0.06	-0.30	-1.01
		5,000,000	10.18	NA	-0.17	-0.11	-0.22	-0.70
2.05	6.4	100,000	-1.03	-57.44	-0.33	-0.30	-0.13	-2.47
		500,000	-0.43	-56.04	-0.20	-0.13	-0.18	-0.25
		1,000,000	0.75	68.48	-0.08	0.03	-0.18	-0.55
		5,000,000	0.00	-13.43	-0.08	-0.05	-0.13	-0.39
3.05	1.7	100,000	-0.54	-3.87	-0.38	-0.14	-0.13	-1.91
		500,000	-0.20	-2.55	-0.18	-0.14	-0.22	-0.21
		1,000,000	-0.06	14.27	-0.10	-0.03	-0.17	-0.43
		5,000,000	-0.06	1.82	-0.14	-0.07	-0.13	-0.29
4.05	1.4	100,000	-0.44	-1.56	-0.36	-0.25	-0.13	-1.64
		500,000	-0.17	-0.67	-0.20	-0.11	-0.15	-0.13
		1,000,000	-0.10	2.10	-0.12	-0.03	-0.15	-0.36
		5,000,000	-0.05	0.13	-0.14	-0.07	-0.12	-0.22

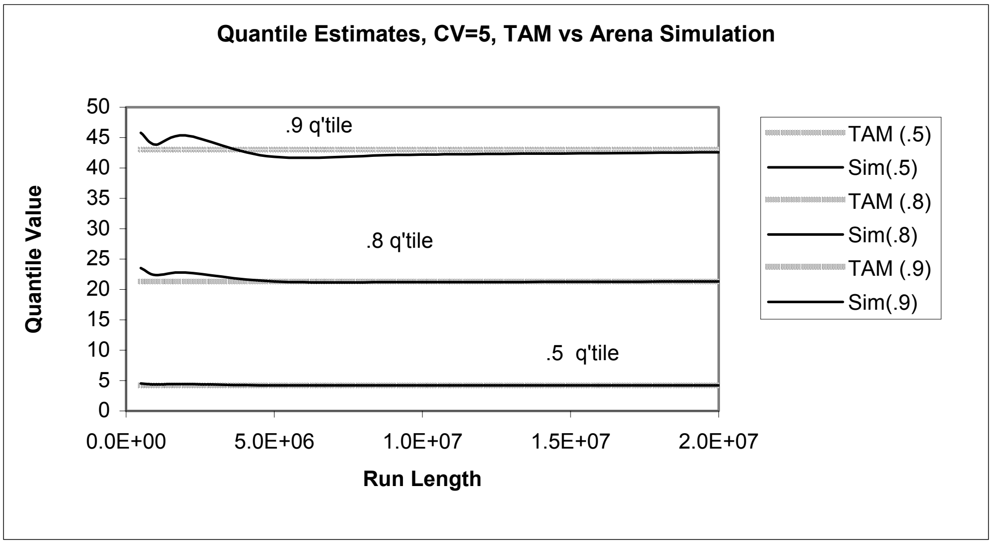

Figure 3: Simulation Runs for M/P/1, CV =5

Table 2: TAM and Simulation Results for Quantile Estimates

CV	Pareto alpha	quantile	quantile values TAM	ARENA	GPSSH	Brichet Approx.
1.528	3.5	0.5	1.24	1.23	1.25	2.30
		0.8	4.4	4.37	4.42	3.31
		0.9	7.09	7.05	7.1	4.37
		0.99	17.9	17.8	18.0	10.99
3	2.25	0.5	3.33	3.3	3.33	5.28
		0.8	15.31	15.27	15.31	10.99
		0.9	28.89	28.8	29.05	19.13
		0.99	133.13	139.00	139.5	120.68
5	2.083333	0.5	4.22	4.23	4.22	6.82
		0.8	21.33	21.34	21.31	15.88
		0.9	42.97	42.56	43.2	30.12
		0.99	246.02	259.30	269.2	252.29
6.4	2.05	0.5	4.45	4.46	4.44	7.25
		0.8	22.99	23.14	23.01	17.34
		0.9	46.96	47.43	47.4	33.56
		0.99	334.2	320.9	314.8	300.71
10	2.020202	0.5	4.68	4.66	4.67	7.68
		0.8	24.77	24.58	24.74	18.85
		0.9	51.53	51.31	51.8	37.18
		0.99	329.82	394.10	364.50	355.25

Table 3: M/P/1 Confidence Interval Estimation Using Fixed Sample Size Procedure, ε′ = .003

Fixed Sample Size Procedure w/ Replications, GPSSH							ARENA
quantile	CV	No. Obs./ Rep	No. of Reps	quantile value	quantile CIHW	Time (min) (500 Mhz)	quantile value (no reps, n=20000000)
0.5	3	106,711	20	3.43	0.13	1.56	3.30
0.8	3	68,295	20	15.61	1.24	1.29	15.27
0.9	3	38,416	20	54.31	45.70	1.11	28.80
0.99	3	4,226	20	107.29	50.93	0.88	139.00
0.5	5	106,711	20	4.18	0.25	1.98	4.23
0.8	5	68,295	20	21.61	1.91	1.55	21.34
0.9	5	38,416	20	46.65	10.98	1.21	42.56
0.99	5	4,226	20	160.94	83.57	0.90	259.30
0.5	10	106,711	20	4.78	0.27	2.09	4.65
0.8	10	68,295	20	28.65	8.16	1.92	24.63
0.9	10	38,416	20	168.65	199.63	1.36	51.17
0.99	10	4,226	20	191.60	103.25	0.88	394.10

Table 4: M/P/1 Confidence Interval Estimation Using Replications of Various Lengths, CV=5

GPSSH						ARENA
quantile	No. Obs./Rep	No. of Reps	quantile value	quantile CIHW	Time (min) (500 Mhz)	quantile value (no reps, n=20000000)
0.5	100,000	20	4.08	0.275	2.8	4.23
0.8	100,000	20	20.99	2.096		21.34
0.9	100,000	20	45.10	8.659		42.56
0.5	500,000	20	4.24	0.073	5.1	4.23
0.8	500,000	20	21.48	0.556		21.34
0.9	500,000	20	43.33	1.861		42.56
0.5	5,000,000	20	4.24	0.035	40.4	4.23
0.8	5,000,000	20	21.52	0.277		21.34
0.9	5,000,000	20	43.78	0.824		42.56

ACKNOWLEDGMENTS

This work was partially supported by the National Science Foundation Grant DMI-0071185, SGER: The Development of Procedures for Utilizing Power-tail Distributions in Queueing Models of Internet Traffic.

REFERENCES

Brichet, F., Roberts, J., Simonian, A., and Veitch, D. 1996. Heavy Traffic Analysis of a Storage Model with Long Range Dependent On/Off Sources. *Queueing Systems* **23**, 197-215.

Brill, P.H., Fischer, M.J., Gross, D and Masi, D.M. 2000. Internet-Type Queues with Heavy-Tailed Service Times and Their Computational Solution, submitted to the *INFORMS Journal on Computing.*

Chen, E.J., W.D. Kelton. 1999. Simulation-Based Estimation of Quantiles. *Proceedings of the 1999 Winter Simulation Conference*, pages 428-434.

Cohen, J.W. 1969. The Single Server Queue, John Wiley and Sons, New York.

Gross, D., Harris, C.M. 1998. Fundamentals of Queueing Theory, 3rd ed., John Wiley and Sons, New York.

Fowler, T.B. 1999. A Short Tutorial on Fractiles and Internet Traffic. *The Telecommunications Review* **10**, 1-14, Mitretek Systems, McLean, VA.

Fischer, M.J., Girard, A.M. Masi, D. M. and Gross, D. 2001. Simulating Internet-Type Queues with Heavy-Tailed Service or Interarrival Times. *Proceedings of the Advanced Simulation Technologies Conference 2001*, Seattle, WA.

Heidelberger, P. and Lewis, P.A.W. 1984. Quantile Estimation in Dependent Sequences. *Operations Research* **32**(1), 185-209.

Naldi, M. 1999. Measurement-based Modelling of Internet Dial-up Access Connections. *Computer Networks* **31**, 2381-2390.

Paxson, V. 2000. Why Measuring the Internet is Painfully Hard. The Fifth Informs Telecommunications Conference, Boca Raton, Florida, March 5-8.

Wood, D.C., B.W. Schmeiser. 1995. Overlapping Batch Quantiles. *Proceedings of the 1995 Winter Simulation Conference*, 303-308.

AUTHOR BIOGRAPHIES

MARTIN J. FISCHER is a Senior Fellow in Mitretek's Center for Telecommunications and Advanced Technology. He has published over 30 articles in refereed journals. He received an M.S. degree and a Ph.D. degree in Operations Research from Southern Methodist University; and a B.A. degree and an M.S. degree in Mathematics from the University of New Hampshire and Florida State University, respectively. He has over 30 years experience in the network design and performance analysis of telecommunications networks.

DENISE M. MASI is a Principal Engineer in Mitretek's Center for Telecommunications and Advanced Technology. Her research interests include queueing theory and simulation applied to telecommunication networks. She received her Ph.D. degree in Information Technology and Engineering at George Mason University in 1998, an M.S. degree in Industrial Engineering from Purdue in 1989, and a B.S. in Industrial Engineering from Texas A&M University in 1988.

DONALD GROSS is a Research Professor in the Department of Systems Engineering and Operations Research at George Mason University and Professor Emeritus of Operations Research, George Washington

University. He is the co-author of the well-known book, *Fundamentals of Queueing Theory,* has numerous publications in the field of queueing theory, and is past president of INFORMS. He was Director, Operations Research and Production Systems at the National Science Foundation, 1988-1990; 1996. He has received the INFORMS Kimball Medal for Service to the Operations Research Profession.

JOHN SHORTLE was born in Santa Barbara, CA. He received a B.S. in mathematics from Harvey Mudd College in 1992. He received a Ph.D. and M.S. in operations research at UC Berkeley in 1996. He worked for three years at U S WEST Advanced Technologies developing stochastic, queueing, and simulation models to optimize networks and operations. In 2000, he won the INFORMS Daniel H. Wagner Prize for excellence in Operations Research Practice. He is currently an assistant professor of Systems Engineering at George Mason University. His research interests include simulation and queueing applications in telecommunications.

PERCY BRILL received a B.Sc. in Math & Physics from Carleton U., an M.A. in Math Stat from Columbia U., and a Ph.D. in O.R. from U. of Toronto. He is presently Professor Emeritus in the Dept. of Management Sci. and the Dept. of Math & Stat at U. of Windsor. He has published research articles in a variety of refereed journals as well as in a number of Proceedings. He received the CanQueue 2000 Award of Distinction for long-term work in Applied Probability and Level Crossing techniques, at the Canadian queueing conference held at the U. of Western Ontario in Sept. 2000. His research interests include: level crossing theory, queueing theory, renewal theory, applied math, nonparametric statistical inference.