Proceedings of the 2004 Winter Simulation Conference
R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

CONFIDENCE INTERVAL ESTIMATION IN HEAVY-TAILED QUEUES USING CONTROL VARIATES AND BOOTSTRAP

Pablo Jesús Argibay-Losada
Andrés Suárez-González
Cándido López-García
Raúl Fernando Rodríguez-Rubio
José Carlos López-Ardao

Departmento de Enxenería Telemática
ETSE de Telecomunicación
Universidade de Vigo
36200 Vigo, SPAIN

ABSTRACT

The heavy-tailed condition of a random variable can cause difficulties in the estimation of parameters and their confidence intervals from simulations, specially if the variance of the random variable we are studying is infinite. If we use a standard method to obtain confidence intervals under such circumstances we shall typically get inaccurate results. To face up this problem, and trying to contribute to find accurate confidence interval estimation methods for such cases, in this paper we propose the use of a control variate method combined with a bootstrap based confidence interval computation. The control variate approach is doubly interesting to address the problem of infinite variance. We tested this approach in a M/P/1 queue system with infinite variance in the queue waiting time and got quite accurate results.

1 INTRODUCTION

Heavy-tailed distributions and distributions with infinite variance play an important role in the modeling of several variables in communication networks. In the literature we can find good references relating these special characteristics [2] to several magnitudes like the size of the files downloaded from HTTP or FTP servers [3] [4], the duration of sessions [5], or even to certain characteristics exhibited by human-computer interactions [6] [7]. In fact, in [8] Paxson shows that the presence of heavy-tailed distributions is an invariant in the internet.

Moreover, network engineering is not the only important field where heavy-tailed distributions have a considerable practical relevance: many financial tasks also use them in models regarding financial and insurance risks [9].

So it should be quite clear how important is to consider that kind of random variables in simulation, as simulation is one of the most powerful tools at time to make performance studies within such engineering and economic areas. But the use of random variables with those characteristics leads to important problems when trying to analyze the results of simulations.

In an M/P/1 queue system, Gross et al. [10] describe problems regarding the estimation of the mean queue waiting time. Fischer et al. [11], Chen [12] and Sees and Shortle [13] study the estimation of quantiles in the presence of the heavy-tail condition.

Argibay et al. [14] study the use of a control variate (CV) to help in the estimation of the mean queue waiting time of the M/P/1, improving both the estimated mean and its confidence intervals (CIs) when the coefficient of the CV method is calculated beforehand from the classical queueing theory.

Our objective is to find an accurate method to estimate the confidence intervals for the mean queue waiting time when affected by the heavy-tailed behavior of the service time but thinking in its usefulness in a more generic scenario (G/P/1). In this paper we extend the work in [14] but now calculating the coefficient of the CV method from the simulation data itself combined with some bootstrap-based confidence interval estimation techniques. Nevertheless we use the M/P/1 queue system in order to validate our results and show that our method achieves accurate confidence intervals for the mean queue waiting time.

In Section 2, we describe mathematically the system queue under study, the M/P/1. In Section 3 we describe the problems related to the accuracy of estimators of CIs for the mean queue waiting time in that queue. In Section 4 we present the approach we propose to try to obtain better results by means of a control variate that help us with the problem of infinite variance of the estimator. In Section 5 we describe the method we are going to use for the construction of confidence intervals, based on bootstrap percentiles, and in Section 6 we describe the results of applying that method to the M/P/1 queue, achieving confidence intervals with good coverage. Finally, in Section 7 we describe some conclusions and further work.

2 THE M/P/1 QUEUE

The M/P/1 is the queue system we are going to work with. Customers arrive according to a Poisson process, and demand independent and identically distributed (iid) service times which follow a Pareto distribution. The queue discipline is "first come first served" and the queue capacity is infinite. We will work with the stochastic process of the consecutive customers' waiting times, $W = \{W_j; j = 1, 2, \dots\}$.

Since the M/P/1 is a special case of the M/G/1, we can use the Pollaczek-Khinchin formula, that gives us the mean queue waiting time:

$$\overline{W} = \frac{\lambda \cdot \overline{S^2}}{2 \cdot (1 - \rho)} \tag{1}$$

where S is the demanded service time random variable, λ is the mean arrival rate of the Poissonian arrival process, and ρ the utilization factor of the system [18].

We are interested in those systems where the service time is a Pareto RV.

The cumulative distribution function (cdf) of the Pareto is given by:

$$\mathrm{F}(x) = 1 - \left(\frac{\mathsf{m}}{x}\right)^{\mathsf{a}} \qquad \forall x \geq \mathsf{m} > 0$$

In [10] a Pareto distribution —with $\mathsf{m} = 1$ and shifted to 0— is used in an M/P/1 to illustrate the problems of simulating such system when a is near 2. Specifically, when a is in $(2, 3)$ the variance of W is infinite. In this paper we also fix m to 1 to show the benefits of our proposed method in a similar scenario. We will also fix ρ to 0.5 to minimize the effects of the transient state in the simulations.

In not heavy-tailed distributions —like the exponential or the normal ones— the probability that the random variable takes a great value is so negligible that if we do not consider those values to calculate some moments of the distribution, we will still get a pretty good estimation of them. This can be the case of the mean, the second moment, and as a consequence, the variance.

The Pareto distribution is a particular example of heavy-tailed distributions. Nevertheless, in the case of heavy-tailed distributions, the probability of such large values, although still being small, is enough to make them have a great influence in some important parameters of the distribution. If those parameters are being estimated through simulation, the problem arises because such not negligible probability is paralelly not so significant to be likely for such large values to appear even in a long simulation; and so the lack of those unlikely samples could finally affect drastically the results.

This effect, and its implications in the simulation of queueing systems, will be discussed in the next section.

3 PROBLEMS WHEN SIMULATING THE M/P/1

We want to estimate a confidence interval for the mean queue waiting time of the M/P/1.

The classical theory of construction of CIs assumes independent and identically distributed samples from a distribution with finite mean and variance.

But if in the M/P/1 the shape parameter of the Pareto, a, is smaller than 3, the variance of W will be infinite. This will imply that we cannot use the central limit theorem to give a CI for $\overline{W}$. Instead, the infinite variance will imply that the sample mean, appropriately normalized, will tend to a stable distribution. To show it, we note that the two conditions that W must achieve to be in the domain of attraction of a stable law are [23]:

1. $$\frac{1-\mathrm{F}_W(x)}{1-\mathrm{F}_W(x)+\mathrm{F}_W(-x)} \to \mathsf{p}$$
   $$\frac{\mathrm{F}_W(-x)}{1-\mathrm{F}_W(x)+\mathrm{F}_W(-x)} \to \mathsf{q}$$
   where $\mathrm{F}_W(x)$ denotes the cdf of W. In our case, W is nonnegative, so we met this condition with $\mathsf{p} = 1, \mathsf{q} = 0$.
2. $$1 - \mathrm{F}_W(x) + \mathrm{F}_W(-x) \sim \frac{2-\alpha}{\alpha} x^{-\alpha} \mathrm{L}(x)$$

where $\mathrm{f}(x) \sim \mathrm{g}(x)$ means $\lim_{x\to\infty} \frac{\mathrm{f}(x)}{\mathrm{g}(x)} = 1$, and $\mathrm{L}(x)$ is a slowly varying function at infinite. In the case of W, it suffices to see that $1-\mathrm{F}_W(x) \sim \mathrm{K}\cdot x^{-\alpha}\cdot \mathrm{L}(x)$, with K a constant value. To show that, we use the fact that the tail asymptotics of a G/G/1 queue follows [25]:

$$\Pr(W > x) \sim \frac{\rho}{1-\rho}\Pr(S_{\mathrm{e}} > x)$$

with S_{e} being the RV associated to the equilibrium distribution (residual life) of the service time, which in our case is the residual life of the customers:

$$\mathrm{F}_{S_{\mathrm{e}}}(x) = \frac{1}{\overline{S}}\int_0^x (1-\mathrm{F}_S(y))\,\mathrm{d}y$$

For the Pareto,

$$\mathrm{F}_{S_{\mathrm{e}}}(x) = \begin{cases} \frac{\mathsf{a}-1}{\mathsf{am}}\cdot x & x \in [0,\mathsf{m}] \\ 1 - \frac{\mathsf{m}^{\mathsf{a}-1}}{\mathsf{a}}\frac{1}{x^{\mathsf{a}-1}} & x \geq \mathsf{m} \end{cases}$$

so

$$\Pr(W > x) \sim \frac{\rho}{1-\rho}\frac{\mathsf{m}^{\mathsf{a}-1}}{\mathsf{a}}\frac{1}{x^{\mathsf{a}-1}}$$

and the last expression is of the form $\mathrm{K}\cdot x^{-\alpha}\cdot \mathrm{L}(x)$, with $\alpha = \mathsf{a}-1$

As both conditions are met, the average of iid steady state samples of size 1 of W, when appropriately normalized, tends to an α-stable distribution with mean zero. It might be possible to use the convergence to a stable distribution to make inferences for $\overline{W}$, but that stable distribution depends on unknown parameters (the index of stability, α, the scale parameter, σ, and the skewness parameter, β). Since most stable laws are not available in closed form, one has to resort to tables —limiting the accuracy of the inferences— or perform complex simulations to estimate the distribution of the stable RV.

We can identify two main problems here. The first is the obtention of samples of W from which we can make inferences; the second is how can we make those inferences taking into account the actual existence of moments with infinite value.

3.1 Obtaining samples of W

We are working with a stochastic process, W, and we want to estimate its mean. Most inference procedures work over iid samples of RVs. In general, when working with the waiting time of the customers in a queue system, we will have to resort to simulation to obtain samples. For those samples to be independent, we can make several simulation runs —each one of size m— and average the waiting times of the customers in each simulation run:

$$\overline{W}[m] = \sum_{j=1}^{m}\frac{W_j}{m}$$

Each average tends to $\overline{W}$ when m tends to infinite.

If we are working with an M/G/1 queue, it is possible to use another approach to obtain iid samples from the marginal probability density function (pdf) of W. This can be a valuable tool to compare methods that make inferences about W —as can be the case of a confidence interval for $\overline{W}$— since this one would be equivalent to an "ideal" simulation scenario obtaining independent samples in steady state. The method uses the well known expression

$$\mathrm{f}_W(w) = \sum_{k=0}^{\infty}(1-\rho)\cdot\rho^k\cdot \mathrm{f}_{\mathrm{r},k}(w) \qquad (2)$$

[17] where

$$\mathrm{f}_{\mathrm{r},k}(w) = \mathrm{f}_{\mathrm{r}}(w) * \mathrm{f}_{\mathrm{r}}(w) * \cdots * \mathrm{f}_{\mathrm{r}}(w) \qquad (k\text{-times})$$

where $*$ is the convolution operator, and $\mathrm{f}_{\mathrm{r}}(w)$ the pdf of the service time residual life [19].

3.2 Making inferences for W

Once we have iid samples of some RV to work with, we must choose some method for making inferences about that RV. In our case we want to construct a CI for $\overline{W}$. If we apply the normal approximation to the samples of an infinite variance RV, the results are bad behaved. In Figure 1 we can see what happens if we apply the normal approximation to the averages obtained from several simulations of an M/P/1 queue with $\rho = 0.5$ and Pareto demanded service time with $\mathsf{a} = 2.1$ and $\mathsf{m} = 1$. We have obtained 1000 CIs, each one constructed from $n = 50$ averages individually calculated —after the respective simulation— over a $m = 10000$ length sample of W. For the sake of clarity, we have represented only 100 randomly chosen intervals. We can see that, although we have run a considerable number of simulations, the results have poor accuracy. The empirical coverage is only 0.717. Notice the logarithmic scale in the Y axis. Many intervals have enormous values and many others do not include the theoretical value.

The root of our problems is the fact that the sum of iid samples of $\overline{W}[m]$, appropriately normalized, does not tend to a normal distribution due to the infinite variance

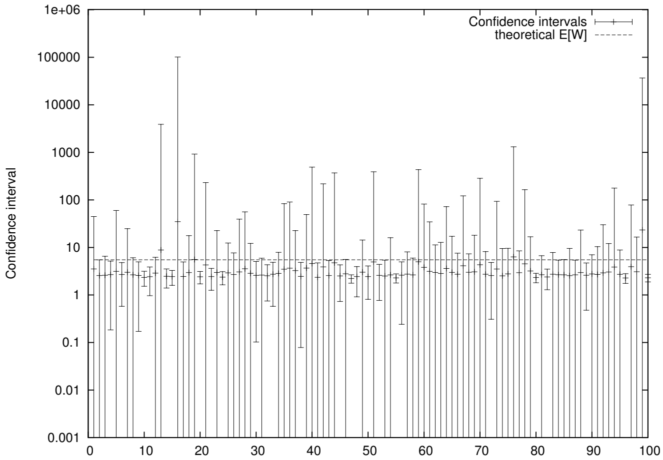

Figure 1: 95% confidence intervals using normal approximation.

of this RV. Instead, it tends to a stable distribution. In this case, few methods to construct CIs are available.

The bootstrap is a method to calculate a functional of a distribution function, as is the case of a CI. It was proposed by Efron [21] and it is based on the general idea that the relationship between a sample space and a specific group of random samples from it, closely resembles the relationship between that group of samples and a resample of them. So, analyzing that subsample and the group of samples, we can extract information about the sample space. The bootstrap, as originally proposed, can fail when the underlying distribution has infinite variance. But it was realized [15] that, when the parameter for which we want to estimate the CI is the mean, the bootstrap can be modified so that it works even in that case. If the samples are in the domain of attraction of an α-stable law with $1 < \alpha < 2$, it suffices to take a resample size, s_n, such that $\lim_{n\to\infty} \frac{s_n}{n} = 0$. In a G/G/1 queue, with service times following a Pareto with shape parameter $2 < \mathsf{a} < 3$, we are in the domain of attraction of a stable law when averaging independent samples of size 1 of W, so we can use bootstrapping as indicated.

Another approach to construct asymptotically valid inference procedures, the subsampling method, was adopted in [24]. It involves evaluating the statistic of interest at subsamples of the data and extrapolating its distribution to the actual sample size. We can apply it to distributions with infinite variance if they are in the domain of attraction of an α-stable law with $1 < \alpha < 2$, and the subsample size, s_n, and the sample size, n, are such that $\lim_{n\to\infty} \frac{s_n}{n} = 0$ [16].

4 CONTROL VARIATE TO ADDRESS INFINITE VARIANCE

Since the infinite variance of the queue waiting time appears to be an important cause of the problems related to the CI estimation, we are going to use a variance reduction technique to check its appropriateness in the M/P/1 case. Specifically, we are going to use a control variate (CV) method.

The main idea for our control variate selection is an empirical hypothesis interpretation of Equation (1): a long enough simulation run will produce a value of $\overline{W}[m]$ close to the evaluation of Equation (1) substituting the theoretical values (λ, ρ and $\overline{S^2}$) by the empirical ones (average arrival rate, frequency of the "resource busy at arrival time" event, and average square service time).

In [20] we have already evaluated positively an internal RV of mean value $1/(1-\rho)$ for polling systems with constant service times.

In the case of the M/P/1 queue with service time RV of infinite variance, our working hypothesis is that the departures of $\overline{W}[m]$ from its mean value $\overline{W}$ in long enough simulation runs will be mainly due to the departures of

$$\overline{S^2}[m] = \sum_{j=1}^{m} \frac{S_j^2}{m}$$

from its mean value $\overline{S^2}$.

To construct a CI for $\overline{W}$, we will run n simulations —each one of size m— and compute, for each one, their average queue waiting time, $\overline{W}[m]$, and their average square demanded service time, $\overline{S^2}[m]$. We will obtain samples of a new RV, T, from

$$T_i = \overline{W}[m]_i - \mathsf{c}\cdot(\overline{S^2}[m]_i - \overline{S^2}), i = 1...n$$

where $\overline{S^2}$ is the theoretical mean value of the square service time RV, which is known.

According to the theory of control variates, when both W and S^2 have finite variance, the optimum value for c is [26]

$$\mathsf{c}_{\mathsf{opt}} = \frac{\mathrm{Cov}(W, S^2)}{\mathrm{Var}(S^2)}$$

In our infinite variance case, we are going to use as c the ratio of the estimators of the covariance and variance, since it minimizes the sample variance of the new RV.

$$\mathsf{c}[n] = \frac{\sum_{i=1}^{n}(\overline{W}[m]_i - \overline{W}[n]) \cdot (\overline{S^2}[m]_i - \overline{S^2}[n])}{\sum_{i=1}^{n}(\overline{S^2}[m]_i - \overline{S^2}[n])^2} \quad (3)$$

where

$$\overline{W}[n] = \sum_{i=1}^{n} \frac{\overline{W}[m]_i}{n}$$

$$\overline{S^2}[n] = \sum_{i=1}^{n} \frac{\overline{S^2}[m]_i}{n}$$

We will obtain samples of T and process them to obtain a CI for $\overline{W}$. This will be discussed in the next section.

5 METHOD FOR THE ESTIMATION OF CONFIDENCE INTERVALS

We have to choose some method to estimate a CI for $\overline{W}$ using the transformed samples T_i. Since bootstrap methods do not make assumptions about the theoretical distribution of the random variables —unlike the standard methods—, we are going to use it. We have tested several approaches for applying the bootstrap and chosen one that has shown to perform accurately in all tests. If we use an asterisk to say that we work over a resample, the estimator we are going to use in our method is:

$$R = \left| \frac{\overline{T}^*[n] - \overline{T}[n]}{\sigma_{T^*}[n]} \right|$$

where $\sigma_{T^*}^2[n]$ is the sample variance of the resample.

This method is going to construct a new CI over the transformed samples taking into account that the limit distribution of this estimator is unknown. We are going to suppose that this tends to a limit distribution, and therefore we are going to estimate it through bootstrapping the pairs $\left(\overline{W}[m], \overline{S^2}[m]\right)$.

Implicitly, we are using the hypothesis that the resulting controlled RV has finite variance, so we can use the standard bootstrap. The method consists of:

1. We perform n independent simulations of the M/P/1 from empty state. In each one we average the waiting times and square service times of the customers, obtaining n pairs $(\overline{W}[m],\overline{S^2}[m])$.
2. We compute (3), the average and sample variance of that transformed sample, $\overline{T}[n]$ and $\sigma_T[n]$.
3. We resample 9999 times [22] the pairs $\overline{W}[m], \overline{S^2}[m]$ —each resample of size n— and compute the coefficient $\mathsf{c}^*[\mathsf{n}]_\mathsf{i}$ which minimizes the variance of the transformed resample. Next we transform the resample —to obtain T^*— using the coefficient $\mathsf{c}^*[\mathsf{n}]_\mathsf{i}$ and compute the average of the resample, $\overline{T^*[n]_i}$, and the sample variance of the resample, $\sigma_{T^*}[n]_i^2$. We compute the estimator R over the resample and store it in an array, called M.
4. We sort the array M in growing order.
5. The 95 % CI will be $\overline{T}[n] \pm |\frac{\overline{T^*[n]} - \overline{T}[n]}{\sigma_{T^*}[n]}|_{0.95} \cdot \sigma_T[n]$, where $|\frac{\overline{T^*[n]} - \overline{T}[n]}{\sigma_{T^*}[n]}|_{0.95}$ is the 0.95 quantile estimator of the distribution of R, in this case the element number 9500 of the array.

The results are shown in next section.

6 CASE STUDY

We have tested subsampling and bootstrapping both over iid samples from the marginal pdf of W (no transient period effect) and over iid samples of $\overline{W}[m]$ in the M/P/1. We have also applied our proposed method and the normal approximation over the latter group of samples.

The utilization factor is set to 0.5. We have tested the coverage properties of 1000 computed 95% CIs and obtained their average length and the coefficient of variation of their length. We have tested the methods for several values of a for the Pareto service time; each a was tested with two sample sizes, $n = 64$ and $n = 1024$; and the relation between the resample/subsample size, s, and the sample size, n, is $s = n^{\frac{2}{3}}$ for the bootstrap and subsampling methods. As the coverages has been computed from 1000 CIs, the 95% confidence interval for the estimated coverages is$(0.939, 0.961)$ and those below this range are shown in italycs in Tables 1 and 2.

Table 1 depicts the results for subsampling and bootstrap for iid samples from the marginal pdf of W. We can see that when $\mathsf{a} < 3$ the coverage is poor for both methods. The lower the a, the poorer the coverage for the selected sample sizes. We can see that these two methods, though asymptotically correct, give results that are still not in the zone of asymptotic behavior.

Table 2 depicts the results for samples obtained from simulations. For subsampling and bootstrap, we can see that the coverage is poor when $\mathsf{a} < 3$. In this case, W will be under the domain of attraction of a distribution with $\alpha < 2$, so it will have infinite variance. The smaller the a, the poorer the coverage for the subsampling and bootstrap methods for the selected sample sizes, even lower than in the previous case —perhaps a transient period effect.

Regarding the method we propose, we can see in Table 2 that the coverages are good for every a, independently of the possible infinite variance of W —$\mathsf{a} \leq 3$. Even if the shape parameter is so close to 2 that $\overline{W}$ is "about to disappear", the coverage is good. The more heavy-tailed the distribution is, the bigger the average length of the CIs. We may see a strange behavior of the average length of the computed CIs as a function of the sample size n for $\mathsf{a} < 2.5$, where doing more simulation runs gets wider values. The cause may be the fixed size m of the simulations.

Regarding the normal approximation, we have to note that it can only be applied to samples with finite variance ($\mathsf{a} > 3$). We have used it on samples with the other values of a for comparison purposes. We see that the coverage is good when $\mathsf{a} = 4$. In this latter case, the coverage and the mean CI length are similar in the proposed method and the normal approximation, but the proposed method has smaller coefficients of variation.

In Figure 2 we have plotted 100 randomly chosen CIs from the proposed method for the case $\mathsf{a} = 2.1$ and $n = 64$. We can compare them with those from Figure 1 and see that the new CIs have shorter lengths and improved accuracy.

We see that the method we propose performs well compared with the other alternatives and therefore seems promising in order to tackle the problem of heavy-tailed behavior in the simulation of simple queues.

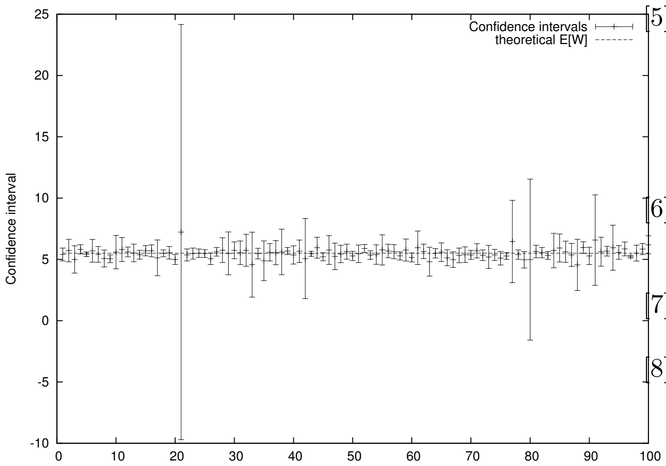

Figure 2: 100 confidence intervals with theoretical coverage of 95% calculated from simulations of the M/P/1 queue and using the proposed method with $\mathsf{a} = 2.1$ and $n = 64$. Empirical coverage is 0.94.

7 CONCLUSIONS AND FURTHER WORK

The heavy-tailed condition of a RV used as input to a queue simulator can be the cause of very low accuracy of the standard confidence interval estimation methods for parameters of the queue, as can be the case of $\overline{W}$. We have shown that the joint use of control variates and adequately chosen inference methods gives good results in a specially problematic case, the M/P/1 queue. It seems promising in order to apply it to more complicated queues, like G/G/1.

REFERENCES

[1] Willinger, W., Vern Paxson and Murad S. Taqqu. 1998. *Self-Similarity and Heavy Tails: Structural Modeling of Network Traffic.* A Practical Guide to Heavy Tails: Statistical Techniques and Applications. Springer Verlag.

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

[3] Pitkow, James E. 1999. *Summary of WWW characterizations.* World Wide Web Journal, Vol. 2, No. 1, pp. 3-13, 1999.

[4] Crovella, Mark E. and Azer Bestavros. 1997 *Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes.* IEEE/ACM Transactions on Networking, vol. 5, no. 6, December 1997. Pages 835-846.

[5] Crovella, Mark E., Murad Taqqu and Azer Bestavros *Heavy-Tailed Probability Distributions in the World Wide Web.* A Practical Guide to Heavy Tails: Statistical Techniques and Applications, by Robert J. Adler, Raisa E. Feldman and Murad S. Taqqu (Editors). Springer Verlag (December 1998).

[6] Mah, Bruce A. 1997 *An Empirical Model of HTTP Network Traffic.* Proceedings of INFOCOM 97, April 7-11, 1997.

[7] Zipf, G. K. 1949. *Human Behavior and the Principle of Least-Effort.* Addison-Wesley Press.

[8] Paxson, Vern. 1999. *Why Analyzing the Internet is Painfully Hard.* Scaling Phenomena in Communication Networks Workshop. Institute for Mathematics and its Applications. October 22-24, 1999.

[9] Asmussen, Soren. 2000. *Ruin Probability.* World Scientific Pub Co.

[10] Gross, D., Shortle, J.F., Fischer, M.J., Masi, D.M.B., 2002. *Difficulties in simulating queues with pareto service.* Proceedings of the 2002 Winter Simulation Conference.

[11] Fischer, Martin J., Donald Gross, Denise M. Bevilacqua Masi and John Shortle. 2001. *Using Quantile Estimates in Simulating Internet Queues with Pareto*

Service Times. Proceedings of the 2001 Winter Simulation Conference.

[12] Chen, E. Jack. 2002. *Two-Phase Quantile Estimation*. Proceedings of the 2002 Winter Simulation Conference.

[13] Sees, John C. Jr. and John F. Shortle. 2002. *Simulating M/G/1 Queues with Heavy-Tailed Service*. Proceedings of the 2002 Winter Simulation Conference.

[14] Argibay-Losada, Pablo Jesús, Andrés Suárez-González, Cándido López-García, Raúl Fernando Rodríguez-Rubio, José Carlos López-Ardao and Diego Teijeiro-Ruiz. 2003. *On the Use of Control Variates in the Simulation of Queues with Heavy Tailed Service*. Proceedings of the 2003 Industrial Simulation Conference.

[15] Athreya, K. 1985. *Bootstrap of the mean in the infinite variance case, II* Technical Report 86-21. Department of Statistics, Iowa State University.

[16] Politis, Dimitris N., Joseph P. Romano and Michael Wolf. 1999. *Subsampling* Springer-Verlag New York.

[17] Benes, V.E. 1956. *On queues with Poisson arrivals*. Annals of Mathematical Statistics. **28**, 670-677.

[18] Gross, Donald and Carl M. Harris. *Fundamentals of Queueing Theory.*

[19] Kleinrock, Leonard. 1975. *Queueing systems, Vol. 1, page 201.* Wiley & Sons.

[20] Suarez Gonzalez, Andres. 2000. Surez-Gonzlez, Andrs, Cndido Lpez-Garca, Jos C. Lpez-Ardao, and Manuel Fernndez-Veiga. 2000. *On the Use of Control Variates in the Simulation of Medium Access Control Protocols* Proceedings of the 2000 Winter Simulation Conference.

[21] Efron, B. 1979. *Bootstrap methods: Another look at the jackknife.* Annals of Statistics, vol. 7. Pages 1-26.

[22] Hall, Peter. 1992. *The Bootstrap and Edgeworth Expansion*, pages 306-311. Springer-Verlag.

[23] Feller, William. 1971. *An Introduction to Probability Theory and Its Applications.* John Wiley and Sons.

[24] Politis, Dimitris N., Joseph P. Romano. 1992. *A general theory for large sample confidence regions based on subsamples under minimal assumptions.* Technical Report 299, Department of Statistics, Stanford University.

[25] Sigman, K. 1999. *A primer on heavy-tailed distributions.* Queueing Systems. 33:261-275.

[26] Bratley, Paul, Bennet L. Fox and Linus E. Schrage. 1987. *A Guide to Simulation.* Springer-Verlag, 2nd edition.

AUTHOR BIOGRAPHIES

PABLO JESÚS ARGIBAY-LOSADA is an assistant professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a telecommunication engineering degree from *Universidade de Vigo* in 2001. Nowadays, he is working toward his Ph.D. in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. His e-mail address is `Pablo.Argibay@det.uvigo.es`

ANDRÉS SUÁREZ-GONZÁLEZ is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidade de Vigo* in 2000. He is a member of ACM and SIGSIM. His current research interests include simulation methodology and analysis of stochastic systems. His e-mail address is `Andres.Suarez@det.uvigo.es`

CÁNDIDO LÓPEZ-GARCÍA is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidad Politécnica de Madrid* in 1995. His e-mail address is `Candido.Lopez@det.uvigo.es`

RAÚL FERNANDO RODRÍGUEZ-RUBIO is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidade de Vigo* in 2000. His current research interests include simulation methodology and analysis of stochastic systems. His e-mail address is `Raul.Rodriguez@det.uvigo.es`

JOSÉ CARLOS LÓPEZ-ARDAO is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidade de Vigo* in 1999. His e-mail address is `J.C.Ardao@det.uvigo.es`

ACKNOWLEDGEMENTS

This work was supported by the project TIC2000-1126 of the “Plan Nacional de Investigacin Cientfica, Desarrollo e Innovacin Tecnolgica” and by the “Secretara Xeral de I+D da Xunta de Galicia” through the grant PGIDT01PX132202PN

Table 1: 95% CIs behavior over iid samples from the marginal pdf of W for subsampling and bootstrap. The 95% CI for a nominal coverage of 0.95 is $(0.939, 0.961)$. Coverages below this range are shown in italycs.

a		**2.001**		**2.01**		**2.1**		**2.5**		**3**		**4**	
n		**64**	**1024**	**64**	**1024**	**64**	**1024**	**64**	**1024**	**64**	**1024**	**64**	**1024**
Subs	**Coverage**	*0.04*	*0.041*	*0.14*	*0.153*	*0.55*	*0.511*	*0.9*	*0.89*	0.96	0.95	0.96	0.955
	$\overline{L}/\overline{W}$	1.999	2.54	41	114	170	11.44	56	1.33	1.75	0.36	1.13	0.21
	CV_L	17	19	25.2	20.7	28	8	5.8	3.8	1.89	1.92	0.45	0.11
Boot	**Coverage**	*0.05*	*0.041*	*0.15*	*0.157*	*0.6*	*0.518*	*0.93*	*0.902*	0.98	0.961	0.98	0.962
	$\overline{L}/\overline{W}$	2.67	2.6	44.9	117	191	11.8	6.46	1.4	2.1	0.38	1.33	0.23
	CV_L	17	19	25.2	20.7	28	8	5.7	3.75	1.8	1.89	0.5	0.11

Table 2: 95% CIs behavior over $\overline{W}[m]$ for bootstrap, subsampling, the proposed one (labeled CV) and normal approximation. The 95% CI for a nominal coverage of 0.95 is $(0.939, 0.961)$. Coverages below this range are shown in italycs.

a		**2.001**		**2.01**		**2.1**		**2.5**		**3**		**4**	
n		**64**	**1024**	**64**	**1024**	**64**	**1024**	**64**	**1024**	**64**	**1024**	**64**	**1024**
Normal	**Cov**	*0.002*	*0*	*0.013*	*0.006*	*0.132*	*0.125*	*0.587*	*0.532*	*0.841*	*0.819*	0.948	0.956
	$\overline{L}/\overline{W}$	0.014	9e-3	0.12	0.09	0.62	0.40	0.23	0.145	0.06	0.02	0.022	5e-3
	CV_L	5.4	3.46	3.27	4.39	5.04	3.5	1.97	5.61	1.26	1.94	0.15	0.11
Subs	**Cov**	*0.025*	*0.01*	*0.119*	*0.086*	*0.328*	*0.365*	*0.687*	*0.693*	*0.845*	*0.832*	0.945	0.947
	$\overline{L}/\overline{W}$	2.19	0.46	6.5	7.2	61.2	18.4	1.79	6.75	0.14	0.06	0.021	5.3e-3
	CV_L	17.4	15.3	1	19.7	17.8	17.12	6.6	19.4	6.35	9.95	0.27	0.14
Boot	**Cov**	*0.028*	*0.01*	*0.121*	*0.088*	*0.349*	*0.372*	*0.729*	*0.715*	*0.888*	*0.847*	0.974	0.957
	$\overline{L}/\overline{W}$	2.43	0.47	7.4	7.4	66.2	18.7	1.9	6.9	0.15	0.067	0.025	5.3e-3
	CV_L	17.2	15.2	11	19.6	17.7	17	6.6	19.4	6.12	9.85	0.2	0.13
CV	**Cov**	*0.93*	0.953	0.955	0.958	0.94	0.959	0.952	0.945	0.957	0.945	0.955	0.961
	$\overline{L}/\overline{W}$	0.73	1.78	0.64	1.9	0.35	0.53	0.053	0.02	0.02	0.006	0.02	4e-3
	CV_L	18.2	3.7	2.5	3.5	1.98	3.4	0.53	1.1	0.19	0.13	0.09	0.03