Proceedings of the 1985 Winter Simulation Conference
D. Gantz, G. Blais, S. Solomon (eds.)

A RELAXATION TIME HEURISTIC FOR REDUCING INITIALIZATION BIAS IN SIMULATION OF M/E_k/1 QUEUEING SYSTEMS

Emily Roth

GTE Laboratories Incorporated
40 Sylvan Road
Waltham, Massachusetts 02254

ABSTRACT

In a previous paper, a heuristic technique was introduced for reducing initialization bias in Monte-Carlo simulation of stationary, infinite-capacity, Markovian queueing systems. The technique differs from many existing methods in that, rather than relying on a statistical analysis of output from preliminary simulation runs, it is based on characteristics of the underlying theoretical model. Experimental results presented earlier provide convincing evidence that use of the heuristic yields an unbiased estimate of the steady-state expected queue length in M/M/1 systems with virtually no cost added to the simulation. We now present results of experiments designed to validate the technique for estimating the steady-state expected number of customers in M/E_k/1 queueing systems.

INTRODUCTION

Most initial transient problem heuristics suggested in the literature are based on some type of statistical analysis of output from preliminary simulation runs. In Roth and Rutan [1], we introduced a heuristic derived through a markedly different approach: exploiting characteristics of the underlying theoretical queueing model in an attempt to locate a suitable truncation point without need for preliminary replications. We have provided validation results for simulating the equilibrium expected queue length of M/M/1 queueing systems. We now apply this technique, the relaxation time heuristic, to a different behavior measure and a more general class of queueing models. Specifically, we examine the performance of the relaxation time heuristic in simulation of the steady-state expected number of customers in M/E_k/1 systems.

RELAXATION TIME HEURISTIC

The relaxation time heuristic is a straightforward application of approximations described in Odoni and Roth [2]. In that paper, numerical solutions to stationary, infinite-capacity, Markovian queueing systems were examined to identify general characteristics of the manner in which these systems approach equilibrium. In particular, the authors found that transients of aggregate measures such as the expected queue length decay over time according to one of four patterns, dependent on initial conditions. Practically speaking, for systems which begin at rest, decay of transient effects is initially faster than exponential. Then, for larger values of time, it is approximated well by pure exponential decay. Thus, if this dominant exponential function can be specified, it may be used to provide an upper bound on the length of time until the system has effectively reached steady state.

Odoni and Roth [2] suggest an approximation for the exponential time constant dominant when 99% of the initial transients have dissipated. They provide a closed-form expression for this approximate time constant, τ_R, which depends only on the system arrival rate λ, service rate μ, and coefficients of variation of the interarrival and service times, C_A and C_S, respectively. For M/E_k/1 queueing systems, this expression reduces to:

$$\tau_R = \frac{1 + 1/k}{2.8\,\mu\,(1 - \sqrt{\varrho})^2} \tag{1}$$

As decay is initially faster than exponential, the time required for one of the aggregate behavior measures of interest to reach 99% of its equilibrium value may be bounded from above by pure exponential decay with rate τ_R. This result implies that after four time constants the observed response variable should be within $\exp(-4) \cong 0.02$ of its equilibrium value.

Our initial transient problem heuristic is based on the assumption that once the response variable is within 2% of its equilibrium value, remaining transients may be ignored. Using the approximation (1), we delete data representative of the first $4\tau_R$ units of model time.

Note the simplicity of implementing this rule. Its use requires only values of four queueing system parameters; there are no ambiguous rule parameters for the user to define. Since the heuristic is based on characteristics of the theoretical queueing model, performance should be unaffected by design parameters of the simulation experiment. We must, however, validate that bias remaining after time $4\tau_R$ is, in fact, insignificant, and that this deletion interval is not too long to be excessively wasteful of data (resulting, perhaps, in overly large confidence intervals).

COMPARISON RULES

As in the M/M/1 study, we compare the practicality and performance of the relaxation time heuristic with those of four other rules: (i) the Kelton-Law algorithm, (ii) a rule attributable to Gordon, (iii) a modification of the Gordon rule based on the truncated sample mean, and (iv) no truncation. These rules form a representative sample of the types of methods currently available to the practitioner (for surveys, see Gafarian, Ancker, and Morisaku [3] and Wilson and Pritsker [4]).

In the following, we describe briefly the manner in which each of the comparison rules attempts to locate a suitable truncation point. For greater detail, see Roth and Rutan [1].

The Kelton-Law technique makes use of generalized least squares regression to fit a line to output values averaged over several epochs and over multiple replications. Kelton and Law [5] reason that if these batch means are representative of steady-state behavior, they should be fit well by a line with zero slope. By fitting successively earlier epochs, the truncation point is reduced until a test for zero slope fails. The smallest acceptable epoch is then used as a truncation point for the remaining replications in the experiment.

The Gordon heuristic (rule TR6 in Wilson and Pritsker [4]) selects as a truncation point the smallest epoch number such that the observed mean calculated from untruncated data achieves a prespecified level of stability. The truncated means rule also seeks stability, deleting successively more data representative of early epochs until the truncated sample mean satisfies a stability requirement.

A serious failing of most existing truncation rules is the necessity to define, frequently with little or no guidance, one or more critical rule parameter(s). In addition to considerations such as the number and length of initial replication runs, the Kelton-Law method requires specification of seven input parameters while the Gordon and truncated means rules each require at least one. In contrast, no truncation and the relaxation time heuristic require neither preliminary runs nor specification of input parameters apart from those needed to define the underlying queueing model.

PERFORMANCE EVALUATION

The primary motivation for using a truncation algorithm is to reduce bias caused by initial conditions. Thus, we first remove from consideration any rule which does not yield data truncated in such a way that it consistently passes a bias test. On the other hand, excessive truncation may result in confidence intervals large enough that interpretation of results is difficult. Therefore, in addition to bias, we monitor estimator variance by calculating the mean squared error.

Other properties to assess when evaluating an initial transient problem heuristic are ease of use and computation cost. Excessive truncation is, therefore, a negative attribute to watch for; depending on whether the experimental design calls for single or multiple replications, generating sufficient data to compensate for a large truncation point (which implies a large deletion interval) could add significantly to the cost of simulation.

We have designed the experiments as multiple sets of independent replications to enable us to exploit Central Limit Theorem properties. Specifically, each experiment comprises R sets of K replications each. We selected values of R and K, along with design parameters such as run length, to provide unquestionable differentiation between untruncated data and data truncated according to one of the heuristics. The Kelton-Law, Gordon, and truncated means rules use the first set of K replications to determine a truncation point which is then used for each remaining replication in the experiment. The validation experiments were designed in this manner in an attempt to duplicate, as closely as possible, conditions under which the rules would be used in practice.

We compare performance under each of the five rules through experiments with varying system parameters of Erlang order k, traffic intensity ϱ, and service rate μ. Design parameter specifications such as replication length N (in epochs), sample frequency Δt, number of replications per set K, and number of sets of replications R were based on results of preliminary experiments. Additional details of the simulation programs are presented in Appendix I.

The system being tested is the stationary, single-queue, $M/E_k/1$ queueing system which begins at rest. The behavior measure of interest is the steady-state expected number of customers in the system.

Performance Criteria

We quantify bias present in the estimate of the mean of the response variable by comparing the simulation grand mean $\bar{\bar{x}}^t$ under rule t with the theoretical steady-state mean $\overline{X}$. (Throughout, we will use uppercase letters to represent random variables and corresponding lowercase letters for observed values; see Appendix II for a summary of notation.) Therefore, if B^t is defined to be the absolute bias of the estimate of the behavior measure $\overline{X}$, observed bias may be expressed as

$$b^t = \sum_{r=1}^{R} \bar{x}_r^t/R - \overline{X}$$

$$= \bar{\bar{x}}^t - \overline{X}$$

where $\bar{x}_r^t$ is the observed mean number in the system over the K replications in set r truncated according to rule t.

Since each $\bar{x}_r^t$ is a sum of means from independent replications, a Central Limit Theorem argument can be used to assert that $\{\bar{x}_1^t, \bar{x}_2^t, \ldots, \bar{x}_R^t\}$ is a random sample from an approximately normal distribution with unknown mean γ. Our initial check for the presence of significant bias is the following hypothesis test:

$$H_0: \gamma = \overline{X}$$
$$H_1: \gamma \neq \overline{X} \qquad (2)$$

with test statistic

$$Z^t = \frac{\overline{\overline{X}}^t - \overline{X}}{\sqrt{\mathrm{Var}(\overline{\overline{X}}^t)}}.$$

If R is large (R > 30 should suffice), the t-statistic Z^t is approximately normally distributed. Under these conditions, we reject H_0 if $|z^t| > 1.96$ where

$$z^t = \frac{\bar{\bar{x}}^t - \overline{X}}{\sqrt{\mathrm{Var}(\bar{\bar{x}}^t)}}$$

with

$$\mathrm{Var}(\bar{\bar{x}}^t) = \left[\sum_{r=1}^{R} (\bar{x}_r^t - \bar{\bar{x}}^t)^2/(R-1)\right]/R .$$

An alternative way to evaluate the presence of bias is to compute the observed probability of coverage, P_{cov}, the likelihood that $\overline{X}$ is contained in the confidence interval about the observed mean generated during the simulation. To estimate P_{cov}, we examine confidence interval coverage in each of the R sets of replications; P_{cov} is equal to the fraction of sets for which a 95% confidence interval about the observed mean for the set, $\bar{x}_r^t$, covers the true mean $\overline{X}$.

To be specific, we define an indicator random variable, I_r, as

$$I_r = \begin{cases} 1 \text{ if } |\bar{x}_r^t - \overline{X}| \le \Delta_r \\ 0 \text{ otherwise} \end{cases}$$

The confidence interval half-width, Δ_r, is given by

$$\Delta_r = t_{K-1,0.975} \sqrt{\mathrm{Var}\,(\overline{X}_r^t)} \;,$$

in which $t_{K-1,0.975}$ is the critical value of the two-tailed t-statistic with 95% confidence and K−1 degrees of freedom. The variance for the mean number in the system derived from a set of K replications, $\mathrm{Var}(\overline{X}_r^t)$, is estimated by

$$\mathrm{Var}\,(\bar{x}_r^t) = \left[\sum_{k=1}^{K} (\bar{x}_{kr}^t - \bar{x}_r^t)^2/(K-1) \right] / K \,,$$

with $\bar{x}_{kr}^t$ defined to be the observed mean of the data generated in replication k of set r and truncated according to rule t.

Finally, the coverage probability may be expressed as

$$P_{cov} = \sum_{r=1}^{R} I_r/R \,.$$

We can also compute the half-width for a 95% confidence interval for the true coverage, Δ_{cov}, using

$$\Delta_{cov} = 1.96\,[P_{cov}(1 - P_{cov})/R]^{1/2}$$

(see Law and Kelton [6]).

By design, with unbiased data and $R = \infty$ the coverage probability should equal 0.95. Given the finite nature of each experiment, the observed values may be somewhat less than this design value. Successful bias removal should, however, result in confidence intervals for P_{cov} that consistently cover much of the region about 0.95.

As mentioned above, there are two primary statistical attributes of concern: bias and variance. The two tests discussed previously indicate presence of bias. To quantify the spread of the estimates, we examine the mean squared error under each rule t, MSE^t, using

$$MSE^t = \sum_{r=1}^{R} (\bar{x}_r^t - \overline{X})^2/R = [R-1]\,\mathrm{Var}\,(\bar{x}^t) + (b^t)^2 \,.$$

Rules that do not reject the null hypothesis in (2) are compared by means of the MSE^t values; the closer the mean squared error is to zero, the more precise the estimate.

Parameter Selection

Specifying input parameters for the Kelton-Law, Gordon, and truncated means heuristics is not a simple task. After preliminary testing, we chose the following values. For the Kelton-Law algorithm we used all parameter values suggested in Kelton and Law [5]. The deviation ϵ required in application of the Gordon and truncated means rules was defined to be proportional to the untruncated sample mean. Specifically, we measured relative rather than absolute error through use of

$$\epsilon = \alpha \sum_{k=1}^{K} \sum_{n=1}^{N} x_{nk1}/NK \,,$$

where x_{nk1} is the observed number of customers in the system at epoch n in replication k of set 1.

Preliminary experiments revealed that the consistency of bias removal using these latter two heuristics is extremely sensitive to the ϵ value selected. Although we are unable to quantify the relationship in any meaningful way, the variance of the truncation point distribution seems to be highly correlated with ϵ. The multiplier values of $\alpha = 1.0$ and $\alpha = 0.04$ for the Gordon and truncated means heuristics, respectively, were found to yield estimates which pass the bias test in a reasonably consistent manner and, thus, were selected as adequate for this work. It is important to note that we make no claim to having selected rule parameters that are in any sense "optimal"; rather, we selected values that performed well in preliminary experiments.

Design Parameters. The number of replications in a set, K, and the number of sets of replications in each experiment, R, must be large enough that the Central Limit Theorem arguments used in measuring bias are valid. We chose K = 40 and R = 400 for the final experiments.

To ensure meaningful comparisons among rules, the replication length, $N\Delta t$, must be small enough that the transient period accounts for a significant portion of each replication. If not, the biased data representing transient behavior may be dominated by unbiased data representative of the steady state. This can result in negligible differences between estimates derived from truncated and untruncated data. We found, through preliminary testing, that $N\Delta t = 7\tau_R$ yields output which allows for clear differentiation. Specifying the replication length to be proportional to the approximate relaxation time, τ_R, has the added advantage of maintaining consistency in the run length as system parameters vary.

System Parameters. We examine the performance of the truncation rules over a range of $M/E_k/1$ queueing systems. The traffic intensity, ϱ, is varied from 0.25 to 0.925 by holding the arrival rate fixed and altering the service rate μ. The Erlang order, k, of the service time distribution ranges from 1 (i.e., an M/M/1 system) to 20.

Table 1: Ranking of Mean Squared Errors

ϱ	k	Δt	REL	K–L	GOR	TRNC MEAN
0.25	1	15	2*	1	4	3
	5		2	1	4	3
	10		2	1	4	3
	15		2	1	4	3
	20		2	1	4	3
0.50	1	100	1	—	3	2
	5		1	—	3	2
	10		2	1	4	3
	15		2	1	4	3
	20		2	1	4	3
0.75	1	500	1	—	2	3
	5		1	—	3	2
	10		2	1	—	3
	15		2	1	3	4
	20		3	1	4	2
0.85	1	2000	1	—	2	3
	5		2	1	4	3
	10		2	1	4	3
	15		1	—	3	2
	20		2	—	1	3
0.90	1	5000	3	2	1	4
	5		3	1	2	4
	10		2	—	1	3
	15		3	2	1	4
	20		3	1	2	4
0.925	1	10000	3	—	1	2
	5		3	1	2	4
	10		2	—	—	1
	15		3	2	1	4
	20		4	2	1	3

* "—" indicates rule failed bias test. Entry numbers represent relative MSE^t values, i.e., "1" = smallest MSE^t.

Table 2: Square Root Mean Squared Error

ϱ	k	Δt	REL	K–L	GOR	TRNC MEAN
0.25	1	15	0.0670*	0.0534	0.0958	0.0863
	5		0.0531	0.0499	0.0800	0.0658
	10		0.0509	0.0450	0.0735	0.0618
	15		0.0512	0.0460	0.0726	0.0620
	20		0.0505	0.0446	0.0731	0.0609
0.50	1	100	0.1359	—	0.1860	0.1548
	5		0.0972	—	0.1358	0.1039
	10		0.0841	0.0785	0.1340	0.1018
	15		0.0852	0.0631	0.1333	0.1071
	20		0.0814	0.0672	0.1253	0.1030
0.75	1	500	0.3229	—	0.3635	0.3747
	5		0.1993	—	0.2508	0.2456
	10		0.1751	0.1412	—	0.2051
	15		0.1759	0.1598	0.2013	0.2112
	20		0.1886	0.1419	0.2034	0.1847
0.85	1	2000	0.5653	—	0.6576	0.6706
	5		0.3346	0.2743	0.4233	0.3537
	10		0.3114	0.2586	0.4337	0.3906
	15		0.3134	—	0.3740	0.3298
	20		0.2751	—	0.2717	0.3011
0.90	1	5000	0.9141	0.8229	0.8017	1.0421
	5		0.5072	0.4249	0.4319	0.6279
	10		0.4888	—	0.4290	0.5952
	15		0.4762	0.4061	0.3852	0.5035
	20		0.4502	0.4095	0.4333	0.5539
0.925	1	10000	1.1739	—	0.9554	1.1200
	5		0.7374	0.6339	0.6408	0.7822
	10		0.6130	—	—	0.5760
	15		0.6382	0.5804	0.5233	0.7439
	20		0.6055	0.5548	0.4822	0.5840

* "—" indicates rule failed bias test.

Table 3: Truncation Point

ϱ	k	Δt	REL	K–L	GOR	TRNC MEAN
0.25	1	15	1905*	1000	3004+	2701
	5		1143	1000	1801+	1574
	10		1048	750+	1657+	1465
	15		1016	750+	1616+	1395
	20		1000	750+	1582+	1294
0.50	1	100	1665	(250)	2640+	2141
	5		999	(230)+	1582+	1180
	10		916	750	1457+	1183
	15		888	230	1411+	1211
	20		874	420	1381+	1196
0.75	1	500	2388	(420)	2953	3110
	5		1433	(190)	1976	1933
	10		1313	690	(1959)	1593
	15		1273	1000+	1673	1774
	20		1254	500	1534	1190
0.85	1	2000	1994	(210)	2517	2565
	5		1196	460	1681	1346
	10		1096	460	1635	1480
	15		1063	(210)	1441	1192
	20		1047	(250)	991	1294
0.90	1	5000	1953	1380	1225	2584
	5		1172	680	821	1614
	10		1074	(250)	667	1478
	15		1042	690	523	1145
	20		1025	750+	919	1380
0.925	1	10000	1808	(500)	1055	1644
	5		1085	750+	775	1219
	10		995	(250)	(273)	835
	15		964	750+	513	1314
	20		949	750+	433	875

* Entry numbers represent number of epochs. Parentheses indicate rule failed bias test.

"+" indicates truncation point was increased to upper limit.

EXPERIMENTAL RESULTS

We look first at the consistency with which each rule yields estimates that pass the bias test (2). With no truncation, the null hypothesis was rejected in every case; therefore, we exclude this rule from further consideration. In Table 1, results under each of the remaining heuristics are illustrated. The relaxation time and truncated means heuristics satisfy the bias criterion in each experiment. The Kelton-Law and Gordon techniques cause rejection of the null hypothesis for one or more sets of system parameters.

The confidence interval coverage probabilities are quite consistent for each truncation rule. While the design value of 0.95 is not always contained in the 95% confidence interval for P_{COV}, the confidence interval does typically include much of the 0.90–0.95 range.

Also in Table 1, we rank the mean squared error values for each case that satisfies the bias criterion (2). The Kelton-Law heuristic appears to be particularly successful in keeping the MSE small. In Table 2, actual MSE values are indicated. Here we note that differences in magnitude across rules are not striking.

For those situations in which bias removal is adequate, the major component of the MSE is the variance of the estimator. Given a fixed replication length, this variance is affected directly by the size of the truncation interval. Thus, we expect that for our experiments application of the Kelton-Law heuristic yielded truncation points that are smaller than those of the other rules. This hypothesis is supported by the data in Table 3, the truncation points generated during the

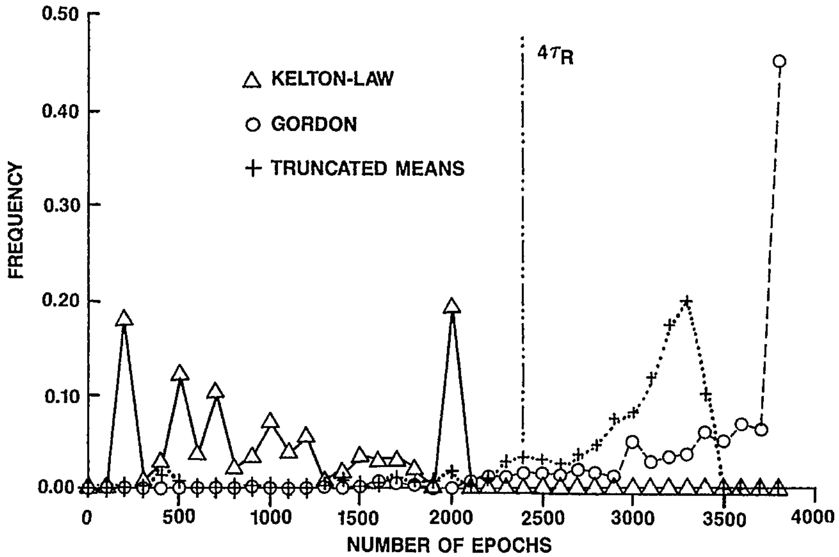

Figure 1: Truncation Point Distribution (M/M/1, ϱ = 0.75, Δt = 500, N = 4179)

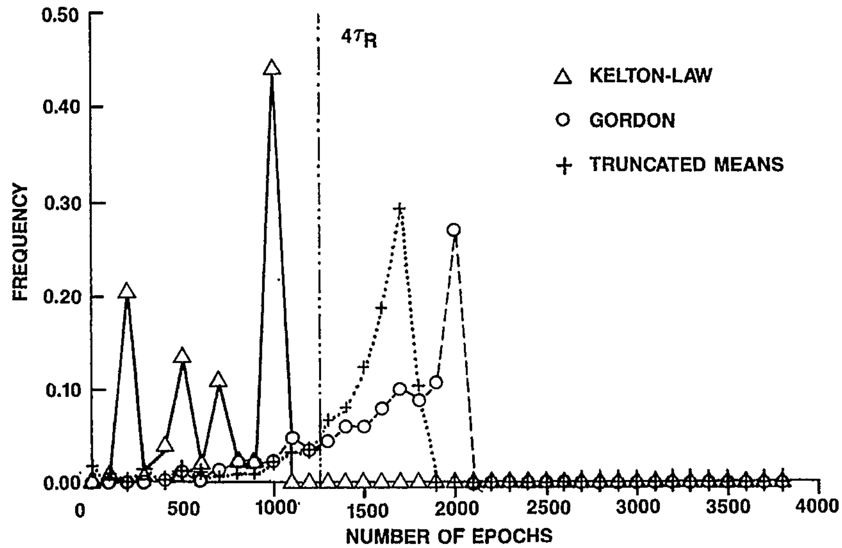

Figure 2: Truncation Point Distribution ($M/E_{20}/1$, ϱ = 0.75, Δt = 500, N = 2194)

experiments. At this point, however, it is important to recall the manner in which we applied the Kelton-Law, Gordon, and truncated means algorithms. In essence, for each of these rules, an experiment results in selecting a single point according to the (unknown) truncation point distribution. If these distributions have large variances, the truncation point selected, n^t, may vary greatly from an *average* value.

In Figures 1 and 2 we show truncation point distributions for M/M/1 and $M/E_{20}/1$ systems, respectively, with ϱ = 0.75, sample frequency Δt = 500, and run length $N\Delta t = 7\tau_R$. To generate these distributions we completed the following steps. First, we applied the rules to 400 independent sets of 40 replications each yielding four sets of 400 n^t values. The 400 n^t values for each rule were then divided into cells of width 100 epochs. Finally, the cell counts were normalized and plotted as frequency distributions.

We note that there is a distinguishable spike at the upper end of each distribution, particularly with the Kelton-Law and Gordon techniques. For these rules and the truncated means heuristic, a maximal n^t value was specified to insure that sufficient data remained for the statistical analysis. We used upper bounds of roughly 0.5N, 0.9N, and 0.8N for the Kelton-Law, Gordon, and truncated means heuristics, respectively. We may infer by the size of these spikes that, due to the particular design parameters and random number seeds, the Kelton-Law and Gordon rules defaulted to the maximal truncation point values quite frequently. We decided not to extend the run length (and, thus, the maximal truncation point values), as even when n^t equals the default value, bias removal is typically acceptable.

Although the truncation point distributions are highly dependent on the particular random number streams generated during the experiments, Figures 1 and 2 are useful in illustrating the variability of the n^t values yielded by the Kelton-Law, Gordon, and truncated means rules. A major difficulty in using such a technique, i.e., one which relies on output from preliminary replications, is that when testing we can, at best, evaluate only average performance. Even with appropriate rule parameters, due to randomness in the preliminary replications an atypical n^t value may be generated. One ramification is that, even if these rules satisfy the bias criteria most of the time in testing, in practice the user will be uncertain as to the suitability of a particular truncation point; there will be no way of determining if the n^t value is an anomaly.

CONCLUSIONS

The results presented in the previous section illustrate that the relaxation time heuristic is indeed an appropriate vehicle for removing initialization bias in simulation of the steady-state expected number of customers in $M/E_k/1$ queueing systems. Regarding ease of implementation, it has several attributes that make it clearly superior to most existing techniques.

Perhaps the major advantage is that the rule requires specification of model characteristics **only**; there are no input parameters left to the user to define. This eliminates the need for preliminary replications and, as there is no randomness in the truncation point selection, the user may have confidence that bias is removed in a consistent manner. Implementation cost is negligible in terms of both programmer time and programming cost; all that is required is the simple calculation of τ_R [equation (1)].

An additional benefit of the relaxation time heuristic or, more specifically, τ_R, is that it can be useful in determining a suitable replication length. Defining time in terms of the approximate time constant provides consistency across systems and, in conjunction with the relaxation time heuristic, allows the user to specify, without preliminary runs, the number of data points to be used in analysis of the simulation output.

A major shortcoming of the relaxation time heuristic at this time is that the algorithm has not yet been tested on many types of queueing systems. We are currently validating the heuristic for M/M/k and $E_k/M/1$ queueing models. Of much greater interest to the practitioner is a method applicable to networks of these systems. Due to the difficulty of obtaining even approximate solutions with sufficient accuracy, we have not yet been able to validate whether or not the assumptions on which the heuristic is based are appropriate for networks of simple queues or single queues with truly general interarrival and service time distributions. We will continue working in this direction.

ACKNOWLEDGMENTS

I am extremely grateful to Nina Bohrod for her programming assistance and to W. David Kelton for supplying software for the Kelton-Law heuristic.

APPENDIX I: COMPUTER PROGRAMS

The computer programs used in this study were written in GPSSH with FORTRAN subroutines. We used the IMSL pseudo-random number generators GGEXN for exponential deviates and GGAMR for Erlang deviates. Each of these generators has a cycle length of $2^{31}-1$, adequately long to prevent cycling in our experiments. All experiments were run on an IBM 3081 computer.

APPENDIX II: NOTATION

N = number of epochs in each replication.

n^t = truncation point for rule t (all data up to and including epoch n^t are deleted).

Δt = time between epochs; sampling interval.

K = number of replications in each set.

R = number of sets of replications in each experiment.

X = steady-state number of customers in the system (random variable).

$\overline{X} = \varrho + (k+1)\varrho^2/2k(1-\varrho) = E(X)$

x_{nkr} = observed number in the system at epoch n in replication k of set r.

$\bar{x}^t_{kr} = \sum_{n=n^t+1}^{N} x_{nkr}/(N-n^t)$ = observed mean number in the system over truncated replication k of set r for rule t.

$\bar{x}^t_r = \sum_{k=1}^{K} \bar{x}^t_{kr}/K$ = observed mean number in the system over truncated replications in set r for rule t.

$\bar{x}^t = \sum_{r=1}^{R} \bar{x}^t_r/R$ = observed mean number in the system over all truncated replications for rule t

REFERENCES

1. Roth, E., and A.H. Rutan, "A Relaxation Time Approach for Reducing Initialization Bias in Simulation," *Proceedings of the 18th Annual Simulation Symposium*, Tampa, 1985, pp. 189–213.
2. Odoni, A.R., and E. Roth, "An Empirical Investigation of the Transient Behavior of Stationary Queueing Systems," *Operations Research 31*, 1983, pp. 432–455.
3. Gafarian, A.V., C.J. Ancker, Jr., and T. Morisaku, "Evaluation of Commonly Used Rules for Detecting 'Steady-State' in Computer Simulation," *Naval Research Logistics Quarterly 25*, 1978, pp. 511–529.
4. Wilson, J.R., and A.A.B. Pritsker, "A Survey of Research on the Simulation Startup Problem," *Simulation 31*, 1978, pp. 55–58.
5. Kelton, W.D., and A.W. Law, "A New Approach for Dealing with the Startup Problem in Discrete Event Simulation," *Naval Research Logistics Quarterly 30*, 1983, pp. 641–658.
6. Law, A.W., and W.D. Kelton, *Simulation Modeling and Analysis*, McGraw-Hill, New York, 1982.

DR. EMILY ROTH is a Senior Member of Technical Staff in the Economics, Operations Research and Statistics Department at GTE Laboratories in Waltham, MA. Prior to her current position, she taught at Carnegie-Mellon University. Dr. Roth received an SB in Mathematics and a PhD in Operations Research from MIT. Much of her research focuses on the development of techniques useful to practitioners solving models of stochastic systems. Of particular interest is understanding the transient behavior of queueing systems.

GTE Laboratories Incorporated
40 Sylvan Road
Waltham, MA 02254
(671) 466-2837