Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

DETERMINING A WARM-UP PERIOD FOR A TELEPHONE NETWORK ROUTING SIMULATION

Christopher W. Zobel

Department of Management Science
and Information Technology
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061, U.S.A.

K. Preston White, Jr.

Department of Systems Engineering
University of Virginia
Charlottesville, VA 22903, U.S.A.

ABSTRACT

We present a new approach to determining the warm-up period for steady-state simulation of telephone traffic. The underlying simulation model captures the sophisticated interactions that determine the acceptance and routing of calls between origin and destination nodes across the telephone network. Recognizing that both the arrival and duration of calls are Markovian, approximate satisfaction of the *equivalence property* of Jackson networks signifies a stochastic steady state. We are able to determine the onset of steady-state behavior, therefore, by monitoring arrival and departure rates observed during the simulation and testing for equivalence. Application of the rule is illustrated using a simple three-node network.

1 INTRODUCTION

It is well known that the initial output of a stochastic simulation may not be representative of steady state behavior. Consider, for example, an M/M/1/c queuing model and suppose that we are interested in determining the average waiting time for customers in the system. If the arrival rate exceeds the service rate, then we would expect the average number in queue to be near capacity and the steady-state waiting times of customers in the system to be appreciable. If the simulation is initialized empty and idle, however, the first few arriving customers will have little or no waiting time. Including these initial observations in the steady state average will artificially lower the resulting estimate of the expected waiting time.

To mitigate the bias introduced by unrepresentative initial conditions, steady-state simulation runs typically include a *warm-up or start-up period.* This period allows the system to evolve to a condition more representative of steady-state, before observations are retained for output analysis. Methods for determining an appropriate length for the warm-up period are efficient if these truncate only the biasing observations; effective if these truncate all such biasing observations; and robust if these apply across a wide range of models. The search for efficient, effective, and robust truncation rules has a long history and continues to be an area of active research. Cash, *et al.* (1992), Goldsman, *et al.* (1994), Law and Kelton (1991), Pawlikowski (1990), Stamoulis and Tsitsiklis (1990), and White (1997) provide ample entrée into the literature on the initialization problem.

In this paper we describe a special-purpose truncation rule developed and applied to determine the warm-up period for a telephone network routing (TNR) problem. We first provide some background for the problem in order to motivate the development of the approach. We then discuss our approach and present some sample experimental results based on the simulation of several specific network configurations.

2 BACKGROUND - TNR PROBLEM

Consider a network consisting of a set of nodes connected together by a series of links or paths, and suppose that these links represent cables through which telephone calls can travel from node to node within the system. At any given time, a series of telephone calls may be in progress within this network, each one associated with an origin node from which the call was initiated, and also with a destination node which ultimately receives the call. While a call is in progress it occupies a continuous pathway, or route, through the network, connecting these origin and destination nodes. Every origin/destination node pair, or OD pair, is associated with a unique set of such connecting routes. Figure 1 illustrates a sample five-node telephone network, and shows each of the one- and two-link routes that connect a chosen OD pair.

The arrival of calls to the network may be modeled as a set of independent Poisson processes, each one associated with a different OD pair. The length of time that each call remains within the network, once it has been assigned to a route, is also stochastic, and we assume that these holding times are IID exponential random variables.

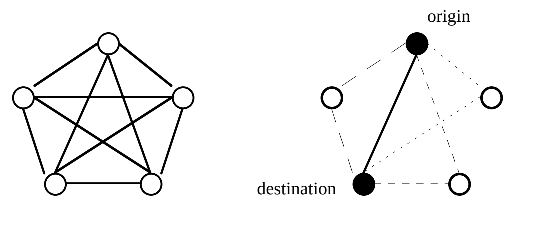

Figure 1: Sample Telephone Network

We assume that each link in the network can carry only a finite number of calls at a given time. By extension, each of the routes associated with a given OD pair also has finite capacity, and we may determine the current available capacity on any given route by the minimum available capacity among its component links. Since each network link may be shared by routes associated with several different OD pairs, a given route with no calls currently occupying it may theoretically have no available capacity for adding a call. Because different OD pairs are often associated with varying call arrival rates, the calls which come from OD pairs with higher arrival rates will tend to take up more of the link capacity on any shared links.

If there is a non-trivial number of nodes in a given network, the overall number of routes in the network can be quite large. Even by restricting consideration to only one- and two-link routes, there will still be $o(N^3)$ total routes in a fully-connected N-node network. It is this complexity, as well as the interaction between different, distinct OD pairs and their associated routes, which makes analyzing the telephone network routing problem so challenging.

There currently is significant interest in developing optimal decision rules for routing calls within such a telephone network. Given the complexity of the problem domain, this is widely recognized as an extremely difficult task for which no exact solutions are currently available (Kelly 1995; Kawashima 1995). One of the more recent approaches to developing such decision rules uses simulation of the telephone network to empirically generate the parameters of a large-scale Markov decision process (MDP); the resulting MDP is then solved to produce a network routing policy (Zobel 1998). Because this approach, known as *Simulation for Model Generation* (SMG), is computationally intensive, it is important to have an efficient procedure for determining an appropriate start-up period for the steady-state simulation of the underlying network. This paper is the result of attempting to address this need.

3 APPROACH

In order to examine the steady-state behavior of the telephone routing problem, a warm-up period must be included at the beginning of each network simulation. This warm-up period must be long enough to move the system into its steady state behavior, but short enough that too much information is not lost from each simulation replication. A standard result from queuing theory known as the *equivalence property* (also referred to as “Burke’s theorem”), states that the steady state output rate of an infinite-length M/M/s queue is equal to the input rate, λ, whenever the sum of the s (identical) service rates exceeds this input rate (Burke 1956). We extend this concept of *rate in = rate out* to our more complicated network model by applying it to each of the OD pairs in the network simultaneously.

Our approach to determining the length of the start-up period involves comparing the observed arrival and departure rates of calls associated with each of the OD pairs in the network. The arrival rate of interest is the rate at which calls are actually added to each OD pair, rather than the rate at which calls arrive to the network, because of the balking which may occur due to limited link capacity. The departure rate is the overall rate of departure from a given OD pair. If, for each OD pair, the difference between these arrival and departure rates is not significantly different from zero, then we can be fairly confident that the system is close to exhibiting steady state behavior. If, however, there exist OD pairs for which there is a significant difference between these two rates, then we consider that the transition behavior of the system is still evolving towards steady state.

There are several methods by which we may make the comparison between the arrival and departure rates of the different OD pairs; our chosen approach provides good results while avoiding excessive computation. Let t represent the length of time over which the empirical data will be collected, let x_{ij} be the total number of arrivals over time t (for OD pair j and in simulation replication i) and let y_{ij} be the corresponding total number of departures over time t (for OD pair j and in simulation replication i). We may then represent the empirical arrival rate of calls to OD pair j, during simulation replication i, by $a_{ij} = x_{ij}/t$, and the corresponding empirical departure rate by $d_{ij} = y_{ij}/t$. $D_{ij} = (a_{ij}-d_{ij})$ then represents the observed difference between the arrival and departure rates.

Letting n represent the number of replications in the experiment and m the number of OD pairs in the network, we have

$$\overline{D_j} = E[D_j] = \frac{\sum_{i=1}^{n} D_{ij}}{n}, \qquad (1)$$

the mean difference between rates for OD pair j, and

$$S_j = \frac{1}{n(n-1)}\left[\left(\sum_{i=1}^{n}(D_{ij})^2 \Big/ n\right) - \left(\overline{D_j}\right)^2\right], \qquad (2)$$

the associated standard deviation. We may then define

$$j^* = \arg \min_{j=1..m} \left\{ \left| S_j \big/ \overline{D_j} \right| \right\} \qquad (3)$$

to be the index of the OD pair with the smallest coefficient of variation in absolute value.

If a 100(1-α)% confidence interval for OD pair j* contains the origin, then the corresponding 100(1-α)% confidence interval for each of the other OD pairs will also contain the origin, and we may say that the mean difference between rates is not significantly different from zero at an m·α confidence level (H_0: $E[D_{j^*}] = 0$) and thus, as suggested above, the network is in steady state with a confidence level of m·α. If, however, the confidence interval for OD pair j* does not include the origin, we claim that the network has not yet achieved steady-state (H_1: the network is not in steady-state), and additional warm-up time is needed.

4 EMPIRICAL RESULTS

Table 1 examines the behavior of four representative 3-node networks with varying arrival and departure rates and various link capacities, and under the control of different types of policies. The statistical results provided are based on simulation experiments of 100 replications apiece. For each example, we looked at start-up periods of different lengths and calculated the mean difference between the empirical arrival and departure rates for each OD pair. We collected the empirical data over different periods of time to try to determine what effect this might have on the results. In general, however, collection time appears to have little effect and very limited amounts of data are still able to identify the underlying system behavior.

Table 1: Warm-up Period Testing

routing scheme	arrival rates (/sec)	departure rates (/sec)	link capacity	start-up length (secs)	Collection length (secs)	mean difference (OD pair j*)	sigma of difference (OD pair j*)	95% LCL	95% UCL
direct	(5,5,5)	1	4	0	1	2.79	1.018	2.528	3.052
direct	(5,5,5)	1	4	0	100	0.03	0.009	0.030	0.034
direct	(5,5,5)	1	4	10	11	-0.02	1.333	-0.363	0.323
direct	(5,5,5)	1	4	10	100	0.002	0.015	-0.002	0.005
direct	(8,3,1)	1	4	0	1	3.42	0.727	3.233	3.607
direct	(8,3,1)	1	4	0	100	0.04	0.006	0.036	0.039
direct	(8,3,1)	1	4	10	11	-0.07	1.320	-0.410	0.270
direct	(8,3,1)	1	4	10	100	-0.0007	0.017	-0.005	0.004
direct	(5,5,5)	0.1	4	0	1	3.62	0.663	3.449	3.791
direct	(5,5,5)	0.1	4	1	2	0.75	1.086	0.470	1.030
direct	(5,5,5)	0.1	4	10	11	-0.01	0.266	-0.078	0.058
direct	(5,5,5)	0.1	20	0	1	4.97	2.115	4.425	5.515
direct	(5,5,5)	0.1	20	3	4	3.13	2.273	2.545	3.715
direct	(5,5,5)	0.1	20	4	5	2.2	2.270	1.615	2.785
direct	(5,5,5)	0.1	20	10	11	-0.25	2.129	-0.798	0.298
LLR	(5,5,5)	1	4	1	2	0.43	1.552	0.030	0.830
LLR	(5,5,5)	0.1	20	1	2	5.28	2.151	4.726	5.834
LLR	(5,5,5)	0.1	20	3	4	3.32	2.274	2.734	3.906
LLR	(5,5,5)	0.1	20	4	5	2.48	2.830	1.751	3.209

(Shaded rows imply that the process is <u>not</u> in steady state)

5 CONCLUSIONS

We have proposed a simple technique for determining when a simulation of a standard telephone network routing process has reached a stochastic steady state. This approach may also be useful in other types of queueing situations, when there is uncertainty involved with the analytic calculation of arrival and/or departure rates, but for which a good representative simulation is available. The approach is particularly appropriate for situations in which there is a lack of problem-specific information or which involve complex interactions between arriving entities.

REFERENCES

Burke, P. J. 1956. The output of a queueing system. *Operations Research* 4:699-704.

Cash, C. R., B. L. Nelson, D. G. Dippold, J. M. Long, and W. P. Pollard. 1992. Evaluation of tests for initial-condition bias. In *Proceedings of the 1992 Winter Simulation Conference*, 577-585.

Goldsman, D., L. W. Schruben, and J. J. Swain. 1994. Tests for transient means in simulated time series. *Naval Research Logistics* 41: 171-187.

Kawashima, K.. 1995. State- and time-dependent routing in the NTT network. *IEEE Communications Magazine* 33: 40-47.

Kelly, F. P. 1995. The Clifford Patterson Lecture, 1995: Modeling communications networks, present and future. *Proceedings of the Royal Society of London A* 444:1-20.

Law, A. L., and W. D. Kelton. 1991. *Simulation Modeling and Analysis*. 2d ed. New York: McGraw-Hill.

Pawlikowski, K. 1990. Steady-state simulation of queueing processes: a survey of problems and solutions. *ACM Computing Surveys* 22:123-170.

Stamoulis, G. D., and J. N. Tsitsiklis. 1990. On the settling time of the congested GI/G/1 queue. *Advances in Applied Probability* 22:929-956.

White, K. P., Jr. 1997. An effective truncation heuristic for bias reduction in simulation output. *Simulation* 69:323-334.

Zobel, C. 1998. An Investigation into the Use of Simulation for Solving Large-scale Markov Decision Processes. Ph.D. dissertation, Department of Systems Engineering, University of Virginia, Charlottesville, Virginia.

AUTHOR BIOGRAPHIES

CHRISTOPHER W. ZOBEL is an Assistant Professor of Management Science and Information Technology at the Virginia Polytechnic Institute and State University. He holds a Ph.D. in Systems Engineering from the University of Virginia, an M.S. in Mathematics from the University of North Carolina at Chapel Hill, and a B.A. in Mathematics from Colgate University. Professor Zobel's primary research interests are in the areas of large-scale stochastic decision problems, intelligent decision systems, heuristic problem-solving, and simulation. He is a member of INFORMS and DSI.

K. PRESTON WHITE, JR. is Associate Professor of Systems Engineering at the University of Virginia. He received the B.S.E., M.S., and Ph.D. degrees from Duke University. He was formerly Assistant Professor at Polytechnic University and at Carnegie-Mellon University, Faculty-in-Residence at Newport News Shipbuilding, and Distinguished Visiting Professor at SEMATECH. His research interests include modeling, simulation, and control of discrete-event and continuous dynamic systems, with emphasis on the integration of decision and information technologies and applications in semiconductor manufacturing. He has published over 100 scholarly articles in these areas. He is U.S. Editor for *International Abstracts in Operations Research* and Associate Editor for *International Journal of Intelligent Automation* and *IEEE Transactions on Components, Packaging, and Manufacturing Technology.* Dr. White is a member of INFORMS and SCS and a senior member of IEEE and IIE. He has served on the AdCom of the IEEE Systems, Man, and Cybernetics Society and as chairman of the IEEE/SMC Technical Committee on System Simulation. He sits on the Boards of the Virginia Modeling, Simulation, and Analysis Center; the University of Virginia Institute for Microelectronics; and the University of Virginia Financial Engineering Group. Dr. White currently represents IEEE/SMC on the WSC Board.