Proceedings of the 1997 Winter Simulation Conference
ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson

EFFICIENCY OF TIME SEGMENTATION PARALLEL SIMULATION OF QUEUEING NETWORKS AS A FUNCTION OF THE SIZE OF THE NETWORK

Mehdi Hoseyni-Nasab

Department of Industrial Engineering
University of Wisconsin – Madison
Madison, Wisconsin 53706, U.S.A.

Sigrún Andradóttir

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332, U.S.A.

ABSTRACT

In this paper, we study the efficiency of a time parallel simulation method, namely the time segmentation approach, that can be applied to simulate long sample paths of a variety of discrete event systems. We show that the efficiency of the method is closely related to the amount of time required for sample paths of the system generated with a common sequence of events to couple (i.e., become identical). Then we provide bounds and estimates of the expected coupling time for Markovian queueing networks of loss and communication blocking stations as a function of the number of stations and buffer capacities in the network.

1 INTRODUCTION

Discrete event simulation has proven to be an effective tool for analysis and evaluation of stochastic systems. However, many simulation experiments require extensive computations to provide reliable estimates for the performance measures of interest. Therefore, developing efficient methods for distributing the computational load of such simulation experiments among multiple processing units is a crucial part of numerical evaluation of these systems. Numerous methods for implementation of simulation experiments on multiple processing units (i.e., parallel simulation) have been studied in recent years. The goal of most of these methods is to exploit the characteristics of the system under study to establish valid communication between multiple processors and simulate the system as efficiently as possible. A survey of recent advances in parallel simulation can be found in Fujimoto (1993).

In this paper, we investigate applications of the time segmentation method for efficient simulation of a class of Markovian queueing networks. This method is a time parallel simulation approach in which parallelism is achieved by partitioning the time domain of the simulation into smaller segments and distributing the processors among these segments. Time parallel simulation approaches using similar ideas have been studied by Heidelberger and Stone (1990), Greenberg, Lubachevsky, and Mitrani (1990), and Lin and Lazowska (1991), among others.

The time segmentation approach was originally proposed by Andradóttir and Ott (1995). They show that the time segmentation method is applicable to Markovian queueing networks that consist of either loss stations or communication blocking stations. Their results have been generalized to Markovian queueing systems containing both loss and communication blocking stations by Hoseyni-Nasab and Andradóttir (1996, 1997). The efficiency of the time segmentation approach is closely related to the amount of time required for the sample paths of the system to couple (see Section 2). Both Andradóttir and Ott (1995) and Hoseyni-Nasab and Andradóttir (1996, 1997) have studied the efficiency of the approach by investigating the dependence of the coupling times on the parameters of the system under study. In this paper we investigate the dependence of the expected coupling times of Markovian queueing networks of loss and communication blocking stations on the size of the network. In particular, we study the growth rate of the expected coupling times with respect to the number of stations and the buffer capacities in the network.

The remainder of this paper is organized as follows: In Section 2 we present the time segmentation method and briefly discuss the conditions under which this method is applicable. In Section 3 we explain how the time segmentation method is applicable to Markovian queueing networks of loss and communication blocking stations. In Section 4 we investigate the efficiency of the approach by providing bounds for the expected coupling times as a function of the size of the network under study. In Section 5 we further examine the results of Section 4 using numerical examples. Finally, Section 6 contains some concluding remarks.

2 THE TIME SEGMENTATION METHOD

In this section we present a parallel simulation approach, namely the time segmentation method, which can be applied to simulate long sample paths of a variety of discrete event systems using multiple processors. Generating long sample paths is of particular interest in steady-state simulation as well as in certain transient simulation experiments. We will also discuss how the efficiency of the time segmentation method is related to the magnitude of the coupling times of the sample paths of the system under study. For more detailed discussions of the time segmentation approach, the reader is referred to Andradóttir and Ott (1995) and Hoseyni-Nasab and Andradóttir (1997).

Suppose that we would like to generate a sample path of a discrete event system S on the interval $[0, T]$. Let A and B be arbitrary sample paths of the system and let $N_A(t)$ and $N_B(t)$ denote the state of the system at time t in sample paths A and B, respectively. Suppose that it is possible to generate valid sample paths of the system in parallel using a common sequence of potential events in such a way that if $N_A(0) \leq N_B(0)$, then $N_A(t) \leq N_B(t)$, for all $t > 0$. Moreover, suppose that there exist two sample paths of the system, sample paths l and u, such that for any sample path A, $N_l(0) \leq N_A(0) \leq N_u(0)$. Then $N_l(t) \leq N_A(t) \leq N_u(t)$, for all $t > 0$ and all sample paths A. This means that the states of the system in sample paths l and u bound the state of the system in all sample paths, at all times, from below and above, respectively. This property implies that by the time the bounding sample paths l and u *couple* (i.e., when $N_l(t) = N_u(t)$), all other sample paths that are being generated with the same sequence of events will have also coupled with sample paths l and u.

To generate a sample path of the system on the interval $[0, T]$, we proceed as follows: We partition the time horizon of the simulation experiment into P equal segments, $[0, T/P], \ldots, [(P-1)T/P, T]$, where P is the number of available processors. Suppose interval i refers to the interval $[(i-1)T/P, iT/P]$, $1 \leq i \leq P$. Each processor is assigned to one segment of the sample path and is responsible for generating the sample path over a time period of length T/P corresponding to that segment. In order for a processor to initiate the simulation of the sample path on its corresponding segment, the processor needs to know the state of the system at the end of the sample path on the previous segment (i.e., to generate a valid sample path we need to ensure that the state of the system at the end of each segment matches the state of the system at the beginning of the next segment). We start the simulation of the sample path on the interval $[0, T/P]$ from the true initial state of the sample path. On the other subintervals, the initial states of the sample path are not known at first (note that all processors start processing at the same time). Therefore, the corresponding processors start simulating sample paths l and u using a common sequence of potential events for the sample paths of each interval. For $i = 2, \ldots, P$, let

$$T_c^i = \inf\left\{t \in \left[\frac{(i-1)T}{P}, \frac{iT}{P}\right) : N_l(t) = N_u(t)\right\}$$

be the coupling time of sample paths l and u generated by processor i. If $T_c^i < iT/P$, then $N_l(t) = N_u(t)$, for $t \in [T_c^i, iT/P)$. This means that after the bounding paths of the system on subinterval i couple, then the processor corresponding to that interval can start collecting data as the state of the system no longer depends on the initial state of the sample path. Therefore the information collected on the interval $[T_c^i, iT/P)$ is valid data for the true sample path. Also, assuming that sample paths l and u on interval $i-1$ couple prior to the end of that interval, then the real initial state of the sample path on interval i is the same as the final state of the coupled sample paths on interval $i-1$. To complete the sample path on interval i, we start the simulation from the true initial state (given by the final state of the coupled bounding paths on interval $i-1$) and simulate the sample path on the interval $[(i-1)T/P, \min\{T_c^i, iT/P\})$ using the same sequence of events as the bounding sample paths of the interval. On the other hand, if sample paths l and u on interval $i-1$ are not coupled prior to the end of that interval, then we again simulate two sample paths l' and u' by initiating two sample paths on interval i starting at the final states $N_l((i-1)T/P)$ and $N_u((i-1)T/P)$ of sample paths l and u on interval $i-1$ and using the same sequence of potential events as the one used to generate sample paths l and u on interval i. By repeating this process, possibly several times, and combining the data collected on all the subintervals, we can generate a complete sample path of the system on the interval $[0, T]$.

In the above procedure, if all the bounding paths l and u couple prior to the end of their corresponding intervals, then the true initial state of the sample path on each subinterval will be known upon completing the simulation of the coupled bounding sample paths on the previous interval. Also, it is clear that if the coupling times are small, then a larger portion of the sample path will be generated during the simulation of the bounding sample paths l and u of the subintervals. On the other hand, if the bounding sample paths on the subintervals do not couple, then we will collect no data during the simulation of the bounding paths l and u and the true initial points of the sample paths on the subintervals are yet to be determined by repeating the procedure with new bounding paths. This discussion suggests that the amount of computer time spent on generating a sample path of S on the interval $[0, T]$ using the time segmentation approach is an increasing function of the coupling times of the bounding paths of the system. Therefore, the magnitude of these coupling times can be used to measure the efficiency of the time segmentation approach. In the following sections we study the behavior of the coupling times of a class of Markovian queueing networks and show that the time segmentation method can be applied to efficiently simulate sample paths of such systems under weak conditions.

3 APPLICATION TO QUEUEING NETWORKS

In this section we show that the time segmentation approach is applicable to a class of queueing networks with loss and communication blocking stations. Before stating the results, we need to introduce our model. Let S be a network of n queueing stations, and for $i = 1, \ldots, n$, let $s_i < \infty$ and $B_i < \infty$ denote the number of servers and the buffer capacity at station i, respectively, and let $C_i = B_i + s_i$. Suppose the service times at all servers of station i are independent and exponentially distributed with rate μ_i, $i = 1, \ldots, n$, and let the arrivals to stations $1, \ldots, n$ be distributed according to independent Poisson processes with rates $\lambda_1, \ldots, \lambda_n$, respectively. For $i = 1, \ldots, n$, suppose p_{ij} is the probability that a job will attempt to join the queue at station j, immediately after a service completion at station i, for $j = 1, \ldots, n$, and let $p_{i,n+1}$ be the probability that a job will leave the system after being served at station i. Moreover, suppose each station in the system is either a *loss* station (i.e., an arriving job to the station that finds a full buffer leaves the system immediately), or a *communication blocking* station (i.e., an arriving job to the station that finds a full buffer leaves the system immediately, unless it is arriving from another station in the network, in which case it undergoes another service time at the station of its most recent service completion and then gets rerouted using the probabilities $\{p_{ij}\}$). This model will be used throughout the rest of this paper.

Let $N(t) = (N_1(t), \ldots, N_n(t))$ denote the state of the system at time t, where $N_i(t)$ is the number of jobs in station i at time t, for $i = 1, \ldots, n$. It is clear that $\{N(t)\}$ is a continuous time Markov chain. It has been shown in Hoseyni-Nasab and Andradóttir (1996, 1997) that we can generate multiple sample paths of the Markov chain $\{N(t)\}$ in parallel using a common sequence of events as follows: We generate exponentially distributed arrival times and service times with rates λ_i and $s_i\mu_i$, respectively, $i = 1, \ldots, n$. Upon determining the next scheduled event (i.e., determining the station at which the event is to be executed and whether the event is an arrival or a service completion), we check each sample path to see if the event is feasible for that sample path. An arrival is a feasible event at station i in a sample path if the buffer of station i is not full. Also, a service completion time generated with rate $s_i\mu_i$ will be accepted (i.e., considered to be feasible) or rejected using thinning by rejection according to the number of busy servers at station i in each sample path (i.e., a scheduled service completion at time t in station i is accepted in sample path A with probability $s_i^A(t)/s_i$, where $s_i^A(t)$ is the number of busy servers in station i in sample path A just prior to time t). If an event is feasible for a sample path, then we execute the event and update the state of the system accordingly. If an event is not feasible for a sample path, we will simply ignore that event and the state of the system in that sample path will remain unchanged. This procedure for parallel simulation of multiple sample paths of discrete event systems using a common sequence of potential events is essentially based on *uniformization* of the Markov chain $\{N(t)\}$ and has been studied also by Vakili (1991).

Next, we need to show that sample paths of the system that are simulated in parallel using a common sequence of potential events satisfy the coupling properties that are required for the time segmentation approach to work. In particular, we need to identify two sample paths whose coupling guarantees the coupling of all the other sample paths. Let sample path l be the sample path that starts at the state where all servers of all stations are idle and sample path u be the sample path that starts at the state where the buffers of all stations are full. Let $T_c^{l,u}$ denote the coupling time of sample paths l and u. Hoseyni-Nasab and Andradóttir (1996, 1997) show that the sample paths l and u serve as the bounding paths required for the time segmentation method to work, and that $E\left\{T_c^{l,u}\right\} < \infty$ (provided that the Markov chain $\{N(t)\}$ is irreducible). Therefore, we can use the procedure described in Section 2 with sample paths l and u as the bounding paths on each interval and generate a long sample path of a system that satisfies the conditions of our model using multiple processors.

4 THE BEHAVIOR OF THE COUPLING TIMES

The efficiency of the time segmentation method is closely related to the magnitude of the coupling times of the system under study. Therefore, in order for the time segmentation approach to be efficiently applicable to simulate large systems, we need to ensure that the coupling time of the system grows reasonably slowly with respect to the size of the system. In this section we discuss the dependence of the expected coupling times on the number of stations and the buffer capacities in the network. All the networks under study satisfy the conditions of our model. The proofs of the results in this section can be found in Hoseyni-Nasab and Andradóttir (1997).

The following proposition shows that under certain conditions, the expected coupling time of the queueing system S grows no faster than linearly with respect to the number of stations in the network.

Proposition 4.1 *For the Markovian queueing network S, suppose $C_i \leq C < \infty$, for $i = 1, \ldots, n$. Moreover, suppose that the Markov chain $\{N(t)\}$ is irreducible and that there exists a real number p such that $0 < p < 1$, and that for every station i at least one of the following two conditions is satisfied (for all $n \geq 1$):*

1. $\lambda_i \geq p$ *and* $\frac{\lambda_i}{s_i\mu_i(1-p_{ii})+\sum_{j\neq i} p_{ji}s_j\mu_j} \geq p$; *or*
2. $\mu_i(1-p_{ii}) \geq p$ *and* $\frac{\mu_i p_{i,n+1}}{\lambda_i+\sum_{j\neq i} p_{ji}s_j\mu_j} \geq p$.

Then the expected coupling time grows at most linearly with respect to n.

The conditions of Proposition 4.1 seem to be rather restrictive as they require the network to have either significant arrivals or significant departures at every station. The following result is a generalization of Proposition 4.1 which shows that the expected coupling time of the system grows no faster than linearly with respect to the number of stations even if some stations in the network do not satisfy the conditions of Proposition 4.1.

Proposition 4.2 *Suppose S is a Markovian network of queues as defined in the statement of Proposition 4.1. Suppose there exists a real number p such that $0 < p < 1$, and that for every station i at least one of the following three conditions is satisfied (for all $n \geq 1$):*

1. $\lambda_i \geq p$ *and* $\frac{\lambda_i}{s_i\mu_i(1-p_{ii})+\sum_{j\neq i} p_{ji}s_j\mu_j} \geq p$; *or*
2. $\mu_i(1-p_{ii}) \geq p$ *and* $\frac{\mu_i p_{i,n+1}}{\lambda_i+\sum_{j\neq i} p_{ji}s_j\mu_j} \geq p$; *or*
3. $\mu_i(1-p_{ii}) \geq p$ *and* $\frac{\mu_i(1-p_{ii})}{\lambda_i+\sum_{j\neq i} p_{ji}s_j\mu_j} \geq p$.

Furthermore, let L be the maximum number of stations that are visited by a customer and do not satisfy Condition 1 or 2 between any two successive service completions at stations that do satisfy either Condition 1 or Condition 2, and suppose that $L < \infty$. Then the expected coupling time grows at most linearly with respect to n.

The size of the queueing network S depends on both the number of stations and the buffer capacities of the network. Propositions 4.1 and 4.2 show that, under certain conditions, the expected coupling time does not grow faster than linearly with respect to the number of stations. The following proposition indicates that for certain networks of loss stations, a similar result holds for the behavior of the expected coupling times with respect to the buffer capacities of the stations.

Proposition 4.3 *Suppose S is a feed-forward network of n loss queueing stations (i.e., $p_{ij} = 0$ for $i > j$). For $i = 1, \ldots, n$, let $s_i \leq s$ and $C_i = B_i + s_i \leq C$, where $s > 0$ is a fixed, positive integer. Then the following hold:*

1. *If $n = 1$ (i.e., there is only one station in the system) and if either $\lambda_1 > 0$ or $\mu_1(1-p_{11}) > 0$, then the expected coupling time grows at most linearly or quadratically with respect to C_1, if $\lambda_1 \neq s_1\mu_1(1-p_{11})$ or $\lambda_1 = s_1\mu_1(1-p_{11})$, respectively.*
2. *For any $n \geq 1$, the expected coupling time of S grows at most quadratically with respect to C, provided that either $\lambda_i > 0$ or $\mu_i(1-p_{ii}) > 0$, for $i = 1, \ldots, n$.*
3. *If for every station i either $\lambda_i > s_i\mu_i(1-p_{ii})$ or $\mu_i(1-p_{ii}) > \lambda_i + \sum_{j\neq i} p_{ji}s_j\mu_j$, then the expected coupling time of S grows at most linearly with respect to C.*

Propositions 4.1-4.3 indicate that for certain queueing networks, the expected coupling time grows reasonably slowly with respect to the size of the network. In the next section, we present some numerical results that support the conclusions of Propositions 4.1-4.3.

5 NUMERICAL RESULTS

In this section we study the behavior of the expected coupling times of networks of queues satisfying our model (see Section 3) as a function of the size of the network through a number of numerical experiments. Our goal is to examine the necessity of the conditions of Propositions 4.1-4.3.

In Hoseyni-Nasab and Andradóttir (1996) we have presented numerical results of simulation experiments with tandem networks of communication blocking queueing stations with $s_i = 1$, $B_i = 10$, $\mu_i = 1$, $p_{i,i+1} = 1$, for $i = 1, \ldots, n$,

$\lambda_1 = 1$, and $\lambda_i = 0$, for $i \neq 1$. Note that for large n, this system does not satisfy the conditions of Propositions 4.1 and 4.2. The numerical results given in Hoseyni-Nasab and Andradóttir (1996) suggest that the coupling times of this network grow faster than linearly with respect to the number of stations in the network and hence the conditions of Proposition 4.1 appear to be at least to some extent necessary for linear growth of the expected coupling times with respect to the number of stations in the network. The results given in Hoseyni-Nasab and Andradóttir (1996) also confirm that the growth rate of the coupling times with respect to the number of stations appears to be linear if either $\lambda_i > 0$ or $\mu_i(1 - p_{ii}) > 0$, for all $i = 1, \ldots, n$ (see Proposition 4.1). To continue the numerical studies presented in Hoseyni-Nasab and Andradóttir (1996), we have simulated queueing networks with the same parameters as the network described above except that $\lambda_i = 0.1$, for $i \in \{20, 40, 60, 80\}$, provided that $n > i$. In each experiment we simulate the bounding sample paths using a common sequence of events and obtain confidence intervals for the expected coupling times by generating 100 independent replications of the two bounding sample paths. The results of our experiments are presented in Table 1.

Table 1: Dependence of the Expected Coupling Times of Tandem Networks of Communication Blocking Stations on the Number of Stations in the System

Number of Stations, n	Coupling Times (95% Confidence Interval)
1	34.63 ($\pm$ 3.33)
2	67.39 ($\pm$ 5.98)
5	169.30 ($\pm$ 13.96)
10	487.51 ($\pm$ 36.34)
20	1,249.09 ($\pm$ 89.68)
40	1,366.08 ($\pm$ 93.59)
60	1,277.05 ($\pm$ 80.69)
80	1,333.19 ($\pm$ 79.97)
100	1,320.97 ($\pm$ 90.31)

The results presented in Table 1 suggest that the expected coupling times do not appear to grow linearly with respect to the number of stations in the system for $n \leq 20$. However, for $n > 20$ the growth rate appears to be sublinear. The results clearly agree with the conclusion of Proposition 4.2.

Our next set of numerical results investigates the behavior of the expected coupling times with respect to the buffer capacities in the system. We have simulated tandem queueing networks of $n = 5$ stations, with $s_i = 1$, $\mu_i = 1$, $p_{i,i+1} = 1$, for $i = 1, \ldots, 5$, $\lambda_1 = 1$, $\lambda_i = 0$, for $i = 2, \ldots, 5$, and $B_i = B \in \{1, 2, 5, 10, 20, 40, 60, 80, 100\}$, for $i = 1, \ldots, 5$. The experiments have been conducted for both systems of communication blocking stations and systems of loss stations. Again, the confidence intervals are obtained by simulating 100 independent replications of the bounding sample paths. Note that these systems do not satisfy the conditions of part 3 of Proposition 4.3, because for all $i \in \{1, \ldots, 5\}$, neither $\lambda_i > s_i\mu_i(1-p_{ii})$ nor $\mu_i(1-p_{ii}) > \lambda_i + \sum_{j \neq i} p_{ji} s_j \mu_j$. The results are presented in Tables 2 and 3.

Table 2: Dependence of the Expected Coupling Times of Tandem Networks of Communication Blocking Stations on the Buffer Capacities in the System

Buffer Capacity, B	Coupling Times (95% Confidence Interval)
1	14.77 ($\pm$ 1.21)
2	25.55 ($\pm$ 2.05)
5	70.66 ($\pm$ 5.02)
10	174.23 ($\pm$ 12.97)
20	575.62 ($\pm$ 39.74)
40	2,027.81 ($\pm$ 159.98)
60	4,347.95 ($\pm$ 344.66)
80	7,905.56 ($\pm$ 611.16)
100	11,954.75 ($\pm$ 984.09)

The numerical results presented in Tables 2 and 3 indicate that the growth rates of the expected coupling times of the networks under study appear to be superlinear with respect to the buffer capacities. This suggests that the conditions of part 3 of Proposition 4.3 are to some extent necessary for the expected coupling times to grow linearly with respect to the buffer capacities. More numerical experiments, in addition to analytical studies, are required to determine the growth rates of the expected coupling times with respect to the size of the system for systems that do not satisfy the conditions of Propositions 4.1-4.3.

6 SUMMARY AND CONCLUDING REMARKS

This paper is concerned with the time segmentation method for parallel simulation of discrete event systems. We discuss how the expected coupling time of the system under study can be used to evaluate the efficiency of the method and present a number of results aiming at understanding the behavior of the expected

Table 3: Dependence of the Expected Coupling Times of Tandem Networks of Loss Stations on the Buffer Capacities in the System

Buffer Capacity, B	Coupling Times (95% Confidence Interval)
1	14.77 (± 1.21)
2	21.85 (± 1.58)
5	30.72 (± 1.73)
10	85.80 (± 5.38)
20	316.65 (± 18.02)
40	1,029.29 (± 54.14)
60	2,508.57 (± 134.00)
80	3,928.49 (± 239.68)
100	6,773.52 (± 401.04)

coupling times for a class of Markovian queueing networks. We show that, under certain conditions, the expected coupling times grow linearly with respect to the size of the network (i.e., the number of stations and buffer capacities of the network) and examine the conditions of our results through a number of numerical examples. Further studies of the extent of applicability and efficiency of the time segmentation method are subjects of our current research.

ACKNOWLEDGMENT

This research was supported in part by the National Science Foundation under grants DDM–9210679 and DMI–9523111.

REFERENCES

Andradóttir, S., and T. Ott. 1995. Time segmentation parallel simulation of networks of queues with loss or communication blocking. *ACM Transactions on Modeling and Computer Simulation* 5(4):269–305.

Fujimoto, R. M. 1993. Parallel and distributed discrete event simulation: Algorithms and applications. In *Proceedings of the 1993 Winter Simulation Conference*, ed. G. W. Evans, M. Mollaghasemi, E. C. Russell, and W. E. Biles, 106–114. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Greenberg, A., B. Lubachevsky, and I. Mitrani. 1990. Unboundedly parallel simulation via recurrence relations. In *Proceedings of the Conference on Measurement and Modeling of Computer Systems*, 1–12. The Association for Computing Machinery, New York, New York.

Heidelberger, P., and H. Stone. 1990. Parallel trace driven cache simulation by time partitioning. In *Proceedings of the 1990 Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, and R. E. Nance, 734–737. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Hoseyni-Nasab, M., and S. Andradóttir. 1996. Parallel simulation by time segmentation: Methodology and applications. In *Proceedings of the 1996 Winter Simulation Conference*, ed. J. M. Charnes, D. M. Morrice, D. B. Brunner, and J. J. Swain, 376–381. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Hoseyni-Nasab, M., and S. Andradóttir. 1997. Efficiency of time segmentation parallel simulation of finite Markovian queueing networks. Working paper.

Lin, Y., and D. Lazowska. 1991. A time division algorithm for parallel simulation. *ACM Transactions on Modeling and Computer Simulation* 1(1):73–83.

Vakili, P. 1991. Using a standard clock technique for efficient simulation. *Operations Research Letters* 10(8):445–452.

AUTHOR BIOGRAPHIES

MEHDI HOSEYNI-NASAB is a Ph.D. student in the Department of Industrial Engineering of the University of Wisconsin – Madison, and a guest student in the School of Industrial and Systems Engineering of the Georgia Institute of Technology. His research interests include parallel simulation and applied probability. He is a member of INFORMS and IIE.

SIGRÚN ANDRADÓTTIR is an Associate Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. Previously, she was an Assistant Professor and later an Associate Professor in the Departments of Industrial Engineering, Mathematics, and Computer Sciences at the University of Wisconsin – Madison. She received a B.S. in Mathematics from the University of Iceland in 1986, an M.S. in Statistics from Stanford University in 1989, and a Ph.D. in Operations Research from Stanford University in 1990. Her research interests include stochastic optimization, simulation, and stochastic processes. She presently serves as Associate Editor for *IIE Transactions, Stochastic Models,* and the *ACM Transactions on Modeling and Computer Simulation.* She is a member of INFORMS and Editor of the *Proceedings of the 1997 Winter Simulation Conference.*