

The Effect of Memory Capacity on Time Warp Performance*

IAN F. AKYILDIZ,[†] LIANG CHEN,[‡] SAMIR RANJAN DAS,[§] RICHARD M. FUJIMOTO,[§] AND RICHARD F. SERFOZO[‡]

[†]School of Electrical Engineering, [‡]School of Industrial and Systems Engineering, and [§]College of Computing, Georgia Institute of Technology, Atlanta, Georgia 30332

The behavior of n interacting processes synchronized by the "Time Warp" rollback mechanism is analyzed under the constraint that the total amount of memory to execute the program is limited. In Time Warp, a protocol called "cancelback" has been proposed to reclaim storage when the system runs out of memory. A discrete state, continuous time Markov chain model for Time Warp augmented with the cancelback protocol is developed for a shared memory system with n homogeneous processors and homogeneous workload with constant message population. The model allows one to predict speedup as the amount of available memory is varied. The performance predicted by the model is validated through performance measurements on an operational Time Warp system executing on a shared-memory multiprocessor using a workload similar to that in the model. It is observed that if the sequential simulation requires m message buffers, Time Warp with a small fraction of message buffers beyond m performs almost as well as Time Warp with unlimited memory. © 1993 Academic Press, Inc.

1. INTRODUCTION

The Time Warp mechanism has been proposed as a general technique for synchronizing asynchronous parallel computations [9]. Unlike conventional, so-called *conservative* approaches to synchronization that utilize blocking to avoid the possibility of synchronization errors, Time Warp uses a mechanism that detects errors at runtime, and automatically recovers using rollback. Empirical studies have reported some success in speeding up the execution of discrete-event simulation applications using Time Warp [5].

A Time Warp program consists of a collection of *logical processes* (LPs) that communicate by exchanging

timestamped event messages.¹ We assume timestamps are unique, real values that are totally ordered by the relation " $<$." All computations are the result of processing messages, i.e., processes do not "spontaneously" begin new computations. Each process maintains a local *clock* variable that indicates the timestamp of the message now being processed, or the last message it processed if the LP is idle. Messages must be processed in non-decreasing timestamp order. If a message with timestamp T is received "in the past," the computations associated with messages with timestamps larger than T must be rolled back. This may involve sending *antimessages* to cancel previously sent messages, which in turn may induce additional rollbacks. Details of the mechanism are described in [9]. We assume that the reader is familiar with the definitions and basic terminology described in that work.

A substantial amount of effort has been devoted to developing analytic models to characterize the behavior of Time Warp. Several models have been developed to model execution on two processors [3, 12, 21, 23]. Recently some work has also attacked the n processor case [8, 11, 20]. Bounds have been derived to compare Time Warp's performance with that of other approaches [2, 14, 17, 22], and conditions have been identified under which an excessive number of rollbacks may occur [19].

All prior work in modeling Time Warp has assumed that there are no constraints on the amount of memory that is available to execute the program. In effect, these models ignore the *message sendback* aspect of Time Warp that was proposed by Jefferson to implement flow control [9]. Several extensions of message sendback have since been proposed in the literature. As discussed later, message sendback based protocols are often used when the simulation runs out of memory and the normal garbage collection procedure (called *fossil collection*) fails to recover additional storage. These protocols roll back some processes that are ahead in virtual time relative to others in order to reclaim memory and enable the simulation to progress. While Time Warp models that ignore

* A preliminary version of this paper appeared in *Proceedings, 1992 ACM Sigmetrics/Performance'92*, Newport, RI, June 1992. The work of Akyildiz, Das, and Fujimoto was supported by Innovative Science and Technology Contract DASG60-90-C-0147 provided by the Strategic Defense Initiative Office and managed through the Strategic Defense Command Advanced Technology Directorate Processing Division, and by NSF Grant CCR-8902362. The work of Chen and Serfozo was supported in part by AFOSR 89-0407 and NSF Grant DDM-9007532.

¹ We use the terms *event* and *message* synonymously.

message sendback can still yield accurate predictions when there is adequate memory, substantial deviations could arise when memory is limited. Moreover, little is known concerning the performance of Time Warp with limited memory. This is the central issue that we address here.

The performance of Time Warp with unlimited memory was discussed in a prior work [8]. Although many of the assumptions used here are similar to those in that work, substantial changes have been made to model the limited memory situation. Furthermore, the solution to the present model differs significantly from the earlier one.

This paper is organized as follows: in Section 2, we briefly overview different memory management schemes for Time Warp and discuss the scheme that we are analyzing. In Section 3, we describe the model. The analysis of this model is discussed in Section 4. In Section 5, we compare predictions made by the model with experimental measurements. In Section 6, we conclude and discuss directions for future research.

2. MEMORY MANAGEMENT IN TIME WARP

Time Warp consumes memory by storing three types of *objects*, viz., state vectors in the state queue, positive messages in the input queue, and negative (anti-) messages in the output queue. Several mechanisms can be used in Time Warp to control the amount of memory. Various approaches are enumerated below:

- *Fossil collection*: Storage used by objects that are older than *global virtual time* (GVT)² can be reclaimed and used for other purposes.
- *Message sendback and its extensions*: If fossil collection fails to recover any storage, but additional memory is needed, some other mechanism is required. Jefferson's message sendback protocol is one such approach [9]. Several variations and/or extensions of message sendback have also been proposed, as outlined below.
- *Infrequent state saving*: State vectors can be saved at a less frequent rate to reduce storage consumed by the state queue [15, 24].
- *Limiting optimism*: Variations on Time Warp that limit the degree to which processes can advance ahead of others implicitly reduce the amount of memory that is required [18, 20, 25, 26].

Several policies utilizing message sendback have been proposed. Jefferson's original proposal invokes message sendback when a process P receives a message, but finds

that there is no memory available to store it [9]. Then the message with the largest send timestamp in P 's input queue is returned to its sender. The sender rolls back on receipt of this message if its local virtual time is more than the send timestamp of the returned message, and possibly resends it in a later forward execution phase. Gafni's protocol [7] generalizes message sendback by removing any stored object (input message, state vector or output message) from the process P that runs out of memory. If the discarded object is an input message, it is returned to the sender, as in message sendback. If it is an output message, it is transmitted to its receiver where it will cancel the corresponding positive message, and P rolls back to the state before it sent the original positive message. If the stored object is a state, it is discarded. Typically, the object with the highest sendtime is selected for removal.

Jefferson proposed an alternative approach called *cancelback* [10] specifically targeted for shared memory systems, where there is a single shared pool of memory. In this protocol, if a process P needs storage for any object u , it is assumed that u is always allocated, but after allocation there may not be any free memory to continue the simulation. The protocol then discards a stored object from some process (not necessarily the same process that stores u) exactly as in Gafni's protocol to free memory. The discarded object must have its sendtime greater than GVT.

Lin observed that all these protocols rely on some process rolling back to free memory [13, 16]; so the ability to *artificially* roll back any process to an earlier virtual time will be a simple, but efficient, memory management scheme. He described the artificial rollback protocol in connection with a shared memory architecture utilizing a global memory pool. If any process runs out of memory, and fossil collection fails to reclaim enough storage, the process farthest ahead in virtual time is rolled back. How far to roll back is a parameter of the scheme. For efficiency reasons, Lin recommends rolling back the process with the latest local clock to the second latest local clock [13]. This continues until enough storage has been reclaimed.

Artificial rollback is semantically similar to cancelback. The syntactic difference, however, makes it somewhat easier to implement in most systems (there is no need to distinguish between messages in forward and reverse transit).

Here, we are concerned with the performance of Time Warp with cancelback. Cancelback is believed to be a complete solution for the Time Warp memory management problem, as it enables Time Warp to complete the simulation with the same amount of memory as the corresponding sequential execution [10].

² GVT is defined as a lower bound on the timestamps of all future rollbacks. For operational definitions of GVT, see [9, 13].

3. MODEL DESCRIPTION

We assume that the Time Warp program is partitioned into n processes, each of which executes on a separate processor. All processors and processes are identical. Since each process executes on a distinct processor we use the words "process" and "processor" interchangeably. We also assume that there are initially m unprocessed messages in the system, and that upon processing each message, exactly one new message is generated. The quantity m is referred to as the *message population*. This fixed message population assumption holds exactly for simulations such as closed queueing networks, and approximately for many other simulations where the size of the event list (in sequential execution) does not vary substantially throughout the simulation. The computation time associated with each message is assumed to be exponentially distributed with rate λ .

We assume rollback is nonpreemptive; if a message in the past is received while an event is being processed, the rollback does not take effect until the processing of the current event is finished. We assume that the time for rollback is negligible, and a message is received immediately after it is sent. We also assume that the process's state is saved prior to processing each event, and the time to save state is negligible. These latter assumptions are mild for medium to large grain simulations (e.g., many combat models), where the associated overheads are small relative to the computation time per event.

An event whose timestamp is less than GVT is called a *committed message*. Committed messages cannot be rolled back. Others are referred to as *uncommitted messages*. These may or may not have been processed previously and, with the exception of the message with timestamp equal to GVT, could later be rolled back.

We use a shared-memory model for memory allocation, i.e., we assume unallocated memory is stored in a global, shared, pool of buffers. Each message includes the associated state vector (state is saved before *each* message is processed) and is assumed to require one memory buffer. M denotes the total number of memory buffers available in the system. We assume no additional buffers to hold messages in transit. When a message is sent, the sender requests and obtains a memory buffer (if one is available), fills it in, and puts the buffer directly in the receiving process's input queue. Any buffer freed by a process (e.g., via fossil collection) is returned to the shared pool. We assume cancellation is done by a pointer traversal as in *direct cancellation* [4] rather than through antimessages. As demonstrated in [4], each pointer consumes a constant amount of space for each positive message in the system, so they are included as part of the

memory buffer. Thus, no additional buffers are required to hold antimessages.

The total number of messages in the system cannot exceed the memory capacity M . We assume that fossil collection is instantaneous and runs continuously in the background so that any fossil is immediately reclaimed. Thus, all messages in the system correspond to uncommitted events. When there are no free message buffers in the system and process P_s attempts to send a message to process P_r , the Time Warp system returns the message u with the largest send timestamp back to the process P_u that originally sent it. This causes P_u to roll back to the last state prior to sending u , and the memory buffer holding u is freed. The only exception occurs when the local time of P_s is greater than the sendtime of u . In this case, P_s aborts the send, and rolls back the current event.

The processing of a message with timestamp τ involves the following operations:

- (i) Read the contents of the message.
- (ii) Compute the update state variables.
- (iii) Send a new message to one of the n processors (including possibly itself) chosen from a uniform distribution. The timestamp of the new message is $\tau + \zeta$, where ζ is an exponentially distributed random variable with rate μ .

The assumptions regarding the choice of message recipient and the timestamp increment may not be true for many systems, but are necessary to make the analysis tractable. Empirical evidence suggests that the message routing function, computation time, the timestamp distributions have a secondary effect on performance for homogeneous applications [6].

A memory buffer needs to be allocated from the shared pool to hold the message to be sent. If none is free, the system already contains M uncommitted messages and the memory management protocol described above is invoked to free a buffer. We assume that the time to send a message, whether or not the protocol is invoked, is negligible. As noted earlier, this is consistent with the situation where computation grain is significantly larger than the associated overheads.

As a processor completes processing a message with timestamp τ , it looks for the lowest timestamped unprocessed message (or antimessage) in its input queue. Suppose the resulting timestamp is τ' . If $\tau' \leq \tau$ (i.e., the lowest timestamped unprocessed message is in the local past), the process rolls back to the most recent state earlier than τ' . The process then (whether or not it rolls back) (i) sets its local clock to τ' , (ii) copies the state vector from the most recent message with timestamp less than τ' , and (iii) processes the message with timestamp τ' as described before. If the message with timestamp τ' is

an antimessage, however, it annihilates the corresponding positive message (if it is already present in the queue), or is processed as a no-op. The time taken to do these operations in between processing of two messages is assumed to be negligible. It is assumed that there is always at least one unprocessed message in each processor's input queue, i.e., the message population is substantially larger than the number of processors n . This implies that all processors busy all the time. This is a reasonable assumption for simulation models that are much larger than the multiprocessor configuration.

Note that a change in the system occurs only when a processor completes processing a message and the above actions are taken. The number of processed uncommitted messages in the system may increase or decrease depending on whether or not there is a rollback, or whether the GVT advances (thus committing some messages). The number of unprocessed messages is always constant and is equal to the message population m . Thus the total number of processed uncommitted messages cannot exceed $(M - m)$. It follows that M must at least be m for the simulation to complete. Our objective is to determine the effect of memory capacity M on the speed of the system.

4. MODEL ANALYSIS

A complete Markov modeling of the system would entail keeping track of all the messages in the system, i.e., recording their timestamps, locations and whether they are processed or not. Since this is impractical we take another approach. The key idea in our approach centers on the assumptions that the processors are identical and the workload is homogeneous. We use these assumptions to extrapolate or generate information about the messages knowing only their quantities. For instance, each processed message in the system has probability $1/n$ of being located at a particular processor. This is valid in equilibrium and hence is a reasonable assumption for nonequilibrium states when one has little information about the past.

Recall that in our model the timestamp increment (difference of the receive and send timestamps of a message and its source message) is exponentially distributed with rate m . This assumption implies that the interdistance in virtual time of the processed uncommitted messages in any process is also exponentially distributed with rate α (say), where α may differ from μ (See [8] for an explanation). Similarly, we assume that the distribution of timestamp differences is also exponential for the unprocessed messages with rate γ , and for the antimessages with rate β . We use these assumptions to model the distribution in virtual time of the processed and unprocessed messages at a processor in order to determine the number of pro-

cessed messages that must be undone in a rollback. Note that in Time Warp actual values of the timestamps are not important aside from their relative ordering.

4.1. Equilibrium State Probabilities

We will represent the system by the continuous-time stochastic process $\{X(t): t \geq 0\}$, where $X(t)$ denotes the total number of processed but uncommitted events in the system at real time t . Note that $X(t)$ does not include messages currently being processed (which are partially processed messages). Recall that $X(t)$ can be at most $(M - m)$, where M is the memory capacity in number of message buffers, and m is the message population. Information concerning the locations and types (processed or unprocessed) of messages will be generated from $X(t)$.

Under our assumptions, the process X is an irreducible Markov chain with state space $S = \{0, 1, \dots, (M - m)\}$. The evolution of X is characterized by its transition rates

$$q_{jk} = \lim_{t \downarrow 0} t^{-1} P\{X(t) = k | X(0) = j\}, \quad j \neq k \in S, \quad (1)$$

The process X is irreducible and its equilibrium distribution π is the solution to the balance equations

$$\pi_j \sum_{k \neq j} q_{jk} = \sum_{k \neq j} \pi_k q_{kj}, \quad j \in S. \quad (2)$$

Under our assumption that all processors are always busy, the exponentially distributed time that the process X remains in any state is $n\lambda$, where λ is the rate of the exponential processing time of each processor. Then, $q_{jk} = n\lambda P_{jk}$, where P_{jk} is the probability that X moves from state j to state k at a transition. Therefore the balance equations (2) simplify to

$$\pi_j = \sum_k \pi_k P_{kj}, \quad j \in S. \quad (3)$$

The following subsections describe the transition probabilities P_{jk} and a computational procedure for obtaining the equilibrium probabilities π_j from (3).

4.2. Expressions for Transition Probabilities

The transitions of the process X involve the following events:

Normal Rollback: $\mathcal{R}_k = \{X \text{ moves to state } k \text{ due to a rollback}\}$

Cancelback Induced Rollback: $\tilde{\mathcal{R}}_k = \{X \text{ moves to state } k \text{ from state } M - m \text{ after cancelback}\}$

GVT Advance: $\mathcal{D}_k = \{X \text{ decreases to state } k \text{ due to a GVT advance}\}$

Unit Increase: $\mathcal{I} = \{X \text{ increases by one unit}\}$

The nonzero transition probabilities of the process X where are therefore

$$P_{j, j+1} = P_j\{\mathcal{F}\}, \quad j, j+1 \in S. \quad (4)$$

$$P_{jk} = P_j\{\mathcal{R}_k\} + P_j\{\mathcal{D}_k\}, \quad 0 \leq k \leq j, j \in S, j \neq M-m. \quad (5)$$

$$P_{jk} = P_j\{\mathcal{R}_k\} + P_j\{\mathcal{D}_k\} + P_j\{\tilde{\mathcal{R}}_k\}, \quad 0 \leq k \leq j = M-m. \quad (6)$$

Here $P_j\{\cdot\}$ is the conditional probability given $X(0) = j$. We derive expressions for these probabilities in the following subsections. We make frequent reference to the following random variables associated with processor i at a transition:

X_i is the number of processed uncommitted messages (in processor i).

Y_i is the number of unprocessed messages in the local past.

Z_i is the number of antimessages in the local past.

$U_i = Y_i + Z_i$ is the number of messages in the local past. These are the messages that cause rollback in processor i . We assume Y_i and Z_i are independent.

4.2.1. Probability Distribution of Messages in the Local Past. Our expressions for the transition probabilities will involve the following conditional distribution for the number of past messages U_i at processor i . Conditioning on X_i ,

$$P\{U_i = u | X_i = x\} = \frac{g(u)}{\sum_{k=0}^{x+m} g(k)}, \quad u = 0, \dots, x+m, \quad (7)$$

where

$$\begin{aligned} g(u) &= P\{U_i = u, X_i = x\} \\ &= \int_0^x e^{-(\beta+\gamma)s} \frac{[(\beta+\gamma)s]^u}{u!} \frac{\alpha(\alpha s)^x}{x!} e^{-\alpha s} ds \\ &= \binom{x+u}{u} \frac{(r_\beta + r_\gamma)^u}{(1+r_\beta+r_\gamma)^{x+u+1}}, \\ & \quad u = 0, \dots, x+m. \quad (8) \end{aligned}$$

Here $r_\beta = \beta/\alpha$ and $r_\gamma = \gamma/\alpha$. The integral follows since the differences between the timestamps of the $X_i = x$ messages are exponential with rate α and the differences between the timestamps of the $U_i = u$ past messages are exponential with rate $\beta + \gamma$ (because of assumption that Y_i and Z_i are independent).

Similarly, we can write

$$P\{Z_i = z | X_i = x\} = \frac{h(z)}{\sum_{l=0}^x h(l)}, \quad z = 0, \dots, x, \quad (9)$$

$$\begin{aligned} h(z) &= P\{Z_i = z, X_i = x\} \\ &= \int_0^x e^{-\beta s} \frac{(\beta s)^z}{z!} \frac{\alpha(\alpha s)^x}{x!} e^{-\alpha s} ds \\ &= \binom{x+z}{z} \frac{(r_\beta)^z}{(1+r_\beta)^{z+x+1}}, \quad z = 0, \dots, x \end{aligned} \quad (10)$$

We also have

$$\begin{aligned} P\{Y_i = y | X_i = x\} &= \sum_{w=y}^m P\{Y_i = y | X_i = x, W_i \\ &= w\} P\{W_i = w | X_i = x\}, \quad y = 0, \dots, m, \end{aligned} \quad (11)$$

where W_i is the number of unprocessed messages in processor i . The m unprocessed messages are equally likely to be located at any processor, and so

$$\begin{aligned} P\{W_i = w | X_i = x\} &= \binom{m}{w} \left(\frac{1}{n}\right)^w \left(1 - \frac{1}{n}\right)^{m-w}, \\ & \quad w = 0, \dots, m. \end{aligned}$$

Consequently, W_i is independent of X_i . Then, similarly to (7) and (9),

$$P\{Y_i = y | X_i = x, W_i = w\} = \frac{f(y)}{\sum_{l=0}^w f(l)}, \quad y = 0, \dots, w.$$

where

$$f(y) = P\{Y_i = y, X_i = x\} = \binom{x+y}{y} (r_\gamma)^y / (1+r_\gamma)^{y+x+1}. \quad (12)$$

4.2.2. Unit Increase. Consider the event \mathcal{F} that there is a one-step increase in the state of the process X . Let I denote the processor that initiates the transition. Note that I is not the GVT regulator [8]. (GVT regulator is the processor that has the GVT event in its input queue. The GVT event is the uncommitted event with the minimum timestamp among all such events.) This is because whenever the regulator advances it immediately commits the event it has just processed. Since all processors are identical, we may assume the n th processor is the regulator at a transition, and let $P_j\{\cdot\}$ denote the conditional probability under this additional condition. Then by standard conditioning,

$$P_j\{\mathcal{F}\} = \sum_{i=1}^{n-1} \sum_{x=0}^j P_j\{\mathcal{F}, I = i | X_i = x\} P_j\{X_i = x\}, \quad j \in S. \quad (13)$$

Whenever $X = j$, we assume that these j processed uncommitted events are independently located at processors $1, \dots, n - 1$ (no processed uncommitted event can be in processor n , which is the regulator), and the probability of one being at processor $i \neq n$ is $1/(n - 1)$. Then

$$P_j\{X_i = x\} = \binom{j}{x} \left(\frac{1}{n-1}\right)^x \left(1 - \frac{1}{n-1}\right)^{j-x}, \quad x = 0, \dots, j. \quad (14)$$

Also, the process X can increase by one unit at processor $I = i$ if and only if $U_i = 0$ (there are no past messages or antimessages at processor i). Then

$$P_j\{\mathcal{F}, I = i | X_i = x\} = n^{-1} P\{U_i = 0 | X_i = x\}, \quad (15)$$

where n^{-1} is the probability that processor i is the first of the n processors to finish its processing. The last probability in (15) is given by expression (7). Thus, $P_j\{\mathcal{F}\}$ can be computed from (13)–(15) and (7).

4.2.3. Rollback. Consider the event \mathcal{R}_k that process X moves to state k as a result of a rollback. As above, assume that processor n is the regulator at the transition and that $P_j\{\cdot\}$ is conditioned on this event. Let I denote the processor that finishes processing first and triggers the transition. If $I = n$ and a rollback occurs, then the process X can move only from j back to j . Then conditioning on X_i and I , we have

$$P_j\{\mathcal{R}_k\} = \sum_{i=1}^{n-1} \sum_{x=j-k}^j P_j\{X_i = x\} P_j\{I = i | X_i = x\} P_j\{\mathcal{R}_k | X_i = x, I = i\} + 1_{(k=j)} P_j\{\mathcal{R}_j, I = n, X_n = 0\}, \quad k = 0, \dots, j. \quad (16)$$

Here, $1_{(k=j)} = 1$ if $k = j$, and 0 otherwise, and

$$P_j\{\mathcal{R}_j, I = n, X_n = 0\} = \frac{1}{n} (1 - P\{U_n = 0 | X_n = 0\}).$$

In the summation of (16), the first probability is given by (14) and the second probability is

$$n^{-1} [1 - P\{U_i = 0 | X_i = x\}],$$

which is evaluated by (7). Here n^{-1} is the probability that processor i finishes processing the current event before the other $n - 1$ processors and this processor can rollback if and only if $U_i \neq 0$.

The third probability in the sum (16) (conditioned on U_i) is

$$P_j\{\mathcal{R}_k | X_i = x, I = i\} = \sum_{u=1}^{x+m} P_j\{\mathcal{R}_k | C_i\} P\{U_i = u | X_i = x, I = i\}, \quad (17)$$

where $C_i = \{U_i = u, X_i = x, I = i\}$. The last probability in (17) is

$$P\{U_i = u | X_i = x, 1 \leq u \leq x + m\} = \frac{P\{U_i = u | X_i = x\}}{1 - P\{U_i = 0 | X_i = x\}},$$

where the probabilities on the right hand side are given by (7).

To complete the evaluation of (16), it remains to obtain an expression for the probability in (17). We can express it as

$$P_j\{\mathcal{R}_k | C_i\} = \int_0^\infty P_j\{\mathcal{R}_k | C_i, T_i = t\} P\{T_i \in dt | C_i\}, \quad (18)$$

where T_i is defined as the local virtual time of processor i minus the GVT. To evaluate this integral, note that for processor i to rollback to move process X to state k when $(X = j, X_i = x, U_i = u, T_i = t)$, exactly $[x - (j - k)]$ processed uncommitted messages must have timestamps below the minimum of the timestamps of the u past messages and the rest of the processed uncommitted messages ($(j - k)$ of them) must have timestamps above this minimum. The minimum timestamp of the u past messages has the distribution

$$1 - (1 - s/t)^u, \quad 0 \leq s \leq t. \quad (19)$$

This is the probability that the minimum of u samples from a uniform distribution on $[0, t]$ is s . The probability of this rollback event is therefore

$$P\{\mathcal{R}_k | C_i, T_i = t\} = \binom{x}{x-j+k} \int_0^t (s/t)^{x-j+k} (1 - s/t)^{j-k} [u(1 - s/t)^{u-1} t^{-1}] ds. \quad (20)$$

The last term in the integral is the density of the distribution (19) that the minimum of the u timestamps for past messages equals s . (We are simply conditioning on the distribution (19)). The rest of the right side of (20) is the binomial probability that exactly $(x - j + k)$ out of the x processed uncommitted messages have timestamps below s . Substituting $y = s/t$, the integral reduces to

$$u \int_0^1 y^{x-j+k} (1 - y)^{u+j-k-1} dy = uB(x - j + k + 1, u + j - k), \quad (21)$$

where

$$B(a, b) = \int_0^1 y^{a-1}(1-y)^{b-1} dy$$

is the beta function. Note that this integral is independent of t . Then (20) is also independent of t and hence it factors out of the integral (18). In other words, substituting (20), (21) in (18), we obtain

$$P_j\{\mathcal{R}_k|C_i\} = u \binom{x}{x-j+k} B(x-j+k+1, u+j-k) \quad (22)$$

$$\begin{aligned} \text{for } 0 \leq u \leq (x+m), \quad 0 \leq x \leq j, \\ 0 \leq j \leq (M-m), \quad 0 \leq k \leq j, \quad 1 \leq i \leq n. \end{aligned}$$

This completes our evaluation of the rollback probability $P_j\{\mathcal{R}_k\}$.

We now consider the event $\tilde{\mathcal{R}}_k$ that process X moves to state k from state $M-m$ after invocation of cancelback. Let \tilde{I} denote the processor that finishes processing first and triggers cancelback that causes a rollback at processor l . Suppose, $\tilde{I} = i$. Then due to the homogeneous assumption, X_i and X_j have the same distribution. As an approximation, we assume that $X_i = X_j$ as cancelback is invoked. Thus, we have

$$P_{M-m}\{\tilde{\mathcal{R}}_k\} = \sum_{i=1}^{n-1} \sum_{x=j-k}^j P_{M-m}\{X_j = x\} P_{M-m}\{\tilde{I} = i | X_i = x\} P_{M-m}\{\tilde{\mathcal{R}}_k | X_i = x, \tilde{I} = i\}, \quad (23)$$

where

$$P_{M-m}\{\tilde{I} = i | X_i = x\} = \frac{1}{n} P\{U_i = 0 | X_i = x\}$$

and $P_{M-m}\{\tilde{\mathcal{R}}_k | X_i = x, \tilde{I} = i\} = P_{M-m}\{\mathcal{R}_k | X_i = x, I = i\}$ is given by (17) through (22).

4.2.4. GVT Advance. As above, we assume that, at a transition of X , the processor n is the regulator and $P_j\{\cdot\}$ is conditioned on this assumption. Let i be the new regulator. As before T_i denotes the virtual time of processor i minus the GVT. Set $T = \min\{T_1, \dots, T_{n-1}\}$. Then the probability of a decrease of the process from j to k due to a GVT advance is

$$P_j\{\mathcal{D}_k\} = \sum_{x=0}^j P_j\{X_i = x\} P_j\{T = T_i | X_i = x\} P_j\{\mathcal{D}_k | X_i = x, T = T_i\} \quad (24)$$

The first probability in the sum is given by (14). The second probability in the sum has a closed form expression, but it is not practical for computations. We therefore use the approximation: The quantities X_l are the same for all the processors l that are not the regulator before or after the transition. Then, for such a processor, $X_l = [(j-x)/(n-2)]$ (the integer part of the number). Note also that

$$\begin{aligned} p(x, y) &= P_j\{T_i \leq T_l | X_i = x, X_l = y\} \\ &= \sum_{k=x+1}^{x+1+y} \binom{x+1+y}{k} \left(\frac{1}{2}\right)^{x+1+y}. \end{aligned}$$

Then our approximation is

$$P_j\{T = T_i | X_i = x\} = \frac{q(j, x)}{\sum_{k=0}^j q(j, k)}, \quad (25)$$

where

$$q(j, k) = \left[p \left(k, \left[\frac{j-k}{n-2} \right] \right) \right]^{n-2}, \quad k = 0, \dots, j.$$

To evaluate the last probability in (24), we will use the following events:

\mathcal{A} = {The GVT advances}

\mathcal{P} = {The message at the new GVT was partially processed before the transition}.

Since we consider the last probability in (23) for fixed i, j, x , we write it as $\bar{P}\{\mathcal{D}_k\}$, where \bar{P} is the probability defined for any event B by

$$\bar{P}\{B\} = P_j\{B | X_i = x, T = T_i\}, \quad (i, j, x \text{ are fixed})$$

Conditioning on \mathcal{A} and \mathcal{P} , we have

$$\bar{P}\{\mathcal{D}_k\} = \bar{P}\{\mathcal{A}\} (\bar{P}\{\mathcal{P}\} \bar{P}\{\mathcal{D}_k | \mathcal{A}, \mathcal{P}\} + \bar{P}\{\mathcal{P}^c\} \bar{P}\{\mathcal{D}_k | \mathcal{A}, \mathcal{P}^c\}). \quad (26)$$

Here \mathcal{P}^c is the complement of \mathcal{P} . If T^* is the minimum timestamp of all unprocessed events, then $T^* - \text{GVT}$ is exponentially distributed with rate $n\gamma$ (it is the smallest timestamp from the n merged streams whose rates are γ). It follows that

$$\bar{P}\{\mathcal{P}^c\} = \left(\frac{\alpha}{\alpha + n\gamma} \right)^{x+1} = (1 + nr_\gamma)^{-x-1}, \quad x = 0, \dots, j, \quad (27)$$

$r_\gamma = \gamma/\alpha$. This is the probability that $x+1$ events from a Poisson stream of rate α occurs before one event occurs

from a Poisson stream of rate $n\gamma$. $\bar{P}\{\mathcal{A}\}$ is given by the following:

$$\bar{P}\{\mathcal{A}\} = \frac{1}{n} P\{U_n = 0 | X_n = 0\}. \quad (28)$$

We now consider $\bar{P}\{\mathcal{D}_k | \mathcal{A}, \mathcal{P}\}$. As the process X moves from j to k due to a GVT advance and the new GVT event was partially processed in processor i with $X_i = x$ before the advance, then $(j - k - x)$ processed uncommitted events will be committed from the other $n - 2$ processors (excluding i and n). Thus

$$\begin{aligned} \bar{P}\{\mathcal{D}_k | \mathcal{A}, \mathcal{P}\} &= \int_0^\infty e^{-(n-2)\alpha s} \frac{[(n-2)\alpha s]^{j-k-x}}{x!} \alpha(\alpha s)^x \frac{e^{-\alpha s}}{x!} ds \quad (29) \\ &= c \binom{j-k}{x} \frac{(n-2)^{j-k-x}}{(n-1)^{j-k+1}}, \quad k = 0, \dots, j-k. \end{aligned}$$

where

$$c = \left[\sum_{k=0}^{j-k} \binom{j-k}{x} \frac{(n-2)^{j-k-x}}{(n-1)^{j-k+1}} \right]^{-1}.$$

A similar argument yields

$$\begin{aligned} \bar{P}\{\mathcal{D}_k | \mathcal{A}, \mathcal{P}^c\} &= \int_0^\infty e^{-(n-1)\alpha s} \frac{[(n-1)\alpha s]^{j-k}}{(j-k)!} n\gamma e^{-n\gamma s} ds \\ &= \bar{c} \left[\frac{n-1}{nr_\gamma + (n-1)} \right]^{j-k}, \quad k = 0, \dots, j, \end{aligned} \quad (30)$$

where

$$\bar{c}^{-1} = \sum_{k=0}^j \left(\frac{n-1}{nr_\gamma + n-1} \right)^k.$$

Here if processor i is the new regulator and the new GVT event is unprocessed, then all $j - k$ processed uncommitted events will be committed from the $n - 1$ processors (excluding processor n).

4.3. Estimation of Parameters

The transition probabilities we have derived above are functions of the parameters $r_\beta = \beta/\alpha$, $r_\gamma = \gamma/\alpha$, where α , β are unknown. We now present an iterative procedure for computing r_β , r_γ and the equilibrium distribution π . We first express r_β , r_γ as a function of π . Whenever $X_i = x$ in processor i , its virtual time (with respect to GVT) T_i is the location of the $(x + 1)$ -th message in a Poisson process with rate α . Therefore,

$$E(T_i | X_i = x) = (x + 1)/\alpha. \quad (31)$$

To estimate r_β , we note that

$$\frac{E(Z_i | X_i = x)}{\beta} \leq E(T_i | X_i = x) \leq \frac{E(Z_i | X_i = x) + 1}{\beta}.$$

We approximate $E(T_i | X_i = x)$ by taking the weighted average of the left and right sides of the above inequalities, using the probability of rollback (and not rollback) of process i as the respective weights. Thus, we have

$$\begin{aligned} E(T_i | X_i = x) &= \frac{1}{\beta} (E(Z_i | X_i = x)(1 - P\{U_i = 0 | X_i = x\}) \\ &\quad + (E(Z_i | X_i = x) + 1)P\{U_i = 0 | X_i = x\}) \\ &= \frac{1}{\beta} (E(Z_i | X_i = x) + P\{U_i = 0 | X_i = x\}). \end{aligned} \quad (32)$$

Let $d(x) = P\{U_i = 0 | X_i = x\} = (1 + r_\beta + r_\gamma)^{-(x+1)}$. Then from (31) and (32), it follows that

$$r_\beta = \frac{E(Z_i | X_i = x) + d(x)}{x + 1}.$$

Now viewing X_i as a random variable whose distribution is determined from the equilibrium distribution π , a reasonable estimate of r_β is

$$r_\beta = \sum_{x=0}^{M-m} \frac{E(Z_i | X_i = x) + d(x)}{x + 1} P\{X_i = x\}, \quad (33)$$

where

$$P\{X_i = x\} = \sum_{j=x}^{M-m} P_j\{X_i = x\} \pi_j.$$

Here, $E(Z_i | X_i = x)$ can be computed using (9).

Similarly, we observe that

$$\frac{E(Y_i | X_i = x)}{\gamma} \leq E(T_i | X_i = x) \leq \frac{E(Y_i | X_i = x) + 1}{\gamma}.$$

As Y_i is expected to be much larger than Z_i , as an approximation $E(T_i | X_i = x) = E(Y_i | X_i = x)/\gamma$ should be fine. Thus, using (31) we have

$$r_\gamma = \gamma/\alpha = E(Y_i | X_i = x)/(x + 1),$$

and a reasonable estimator of r_γ is

$$r_\gamma = \sum_{x=0}^{M-m} \frac{E(Y_i | X_i = x)}{(x + 1)} P\{X_i = x\}. \quad (34)$$

Here, $E(Y_i | X_i = x)$ can be computed using (11).

Our iterative procedure for computing r_β , r_γ , π is as follows. For an initial setting of r_β , r_γ , we compute π from the balance equations (3). Using this π , we compute new values of r_β , r_γ from (33) and (34), and use these in the balance equations (3) to compute a new π . We repeat this procedure until the new values of r_β , r_γ are within a distance of 0.001 from their previous values. Performance measures can then be computed as described in the following subsection.

Extensive computations with this procedure showed it to be very efficient. Although we cannot prove that there exist unique fixed points for r_β , r_γ , the data suggests that this might be true. In runs of the procedure for 100 initial values of r_β , r_γ ranging from 0.001 to 1,000, these parameters always converged to the same values. The number of iterations for the convergence depends strongly on their initial values. This number is very low, approximately 2 to 4, when $(M - m)$ is below 10. For larger values of $(M - m)$, we found that the number of iterations can be decreased significantly by using the previously computed values of r_β and r_γ as initial values for the new value of $(M - m)$.

4.4. Performance Measures

The main performance measure of the system is the *number of messages committed per unit time*. This is the same as the rate of GVT advancement per unit time. This quantity is given by

$$\rho = n\gamma \sum_{j=0}^{M-m} \sum_{k=0}^j (j+1-k) P_j\{\mathcal{D}_k\} \pi_j \quad (35)$$

This follows by a standard law of large numbers for a Markov process. The sum represents the expected decrease in the state of X which is the expected number of committed events per transition, and $n\lambda$ is the total rate of the transitions per unit time.

One can compare this n -processor model to a single one operating in series as follows. Consider a single processor working in series that processes messages according to a Poisson process with rate λ . Then the *speedup* of the n -processor system is expressed as

$$S = \frac{\rho}{\lambda} \quad (36)$$

5. EXPERIMENTAL RESULTS AND VALIDATION OF THE ANALYTIC MODEL

The analytic model makes a number of simplifying assumptions (for instance, we assume that the timestamps of the processed uncommitted events in each processor follow a Poisson distribution) in order to make the analy-

sis tractable. Measurements on a Time Warp kernel running on a shared memory multiprocessor (specifically, a 32-processor KSR1³ manufactured by the Kendall Square Research Corporation) were made and compared with the performance predicted by the analytic model in order to test the validity of the approximations underlying the analysis.

The assumptions used in the analytic model pertaining to the workload (exponential execution time per event, exponential timestamp increment, fixed message population, etc.) correspond to a specific instance of the parallel HOLD (PHOLD) workload model [6]. This model was used in the experiments performed here, and a synthetic application program was designed that corresponds to the workload assumed in the model. Also the mean computation time per event was assigned a value of 25 ms, and the experiments here utilize a high message population m (32 times the number of processors) to satisfy the nonidle processor assumption.

The original Time Warp kernel we have used is described in detail in [4]. This kernel has been modified to make the memory to store messages accessible as a globally shared pool and has been ported to the KSR architecture. Several data structure and algorithmic improvements have also been done to reduce different Time Warp overheads. As described in [4] the antimessages in the kernel are implemented as pointers to the corresponding positive messages organized in a *causality record*, and these pointers require a constant amount of space per event. The state vector and the pointers implementing the causality record are included in each message (recall that the state vector is saved before processing each message), and are not separately allocated/deallocated. Memory buffers to store messages are dynamically allocated from the global pool when a process wants to send a message to another process. The only way storage can be freed is when a message is cancelled or sent back, or when it is fossil collected. GVT computation and fossil collection are performed atomically and on-demand. All processors are involved in fossil collection using a global synchronization. When memory is needed and no free buffer is available in the shared pool, fossil collection is invoked to reclaim memory used by possible fossils. If no fossil exists, cancelback (as described in Section 3) is invoked to free storage.

We have run the experiments with 4, 8, and 12 processors on the KSR1. The effect of memory capacity on speedup is shown in Fig. 1. Both the experimental data and the corresponding analytical estimate are shown. If the total memory capacity M in the shared pool is larger than some threshold value, fossil collection alone is suffi-

³ KSR architecture and KSR1 are trademarks of the Kendall Square Research Corporation.




FIG. 1. Effect of limited memory capacity on speedup.

cient to produce enough free memory for the simulation to progress. However for smaller amounts of memory, cancelback needs to be invoked occasionally to enable progress. Frequency of cancelback invocations increases as memory is reduced further, and speed of execution is reduced as the progress of many processes in virtual time is throttled by frequent cancelbacks. Memory capacity M , however, cannot be reduced indefinitely. Recall that M must at least be m as this is the number of message buffers required in the corresponding sequential execution.

It is noted that with increases in memory capacity from the value required by the sequential simulation, performance increases very sharply, and then becomes almost flat. Thus the speedup curve has a well-defined "knee." The location of this knee indicates the minimum number of message buffers needed to complete the Time Warp simulation almost as rapidly as in the unbounded memory situation. It is interesting to note that the knee occurs at a very low value of memory capacity. The additional memory (over that required for sequential simulation) needed to achieve nearly the unbounded memory performance is a small fraction of the message population.

The small discrepancy between the analytical prediction and experimental results (the difference is mostly within 5% and the maximum difference is about 10%) shown in Fig. 1 can be attributed to the simplifying assumptions used in developing the analytic model, which are only approximately true. In particular, in our analysis we have assumed that the distribution of timestamps for unprocessed positive messages and antimessages in the

local past are independent, which may not be true in reality. Hence we may have overestimated the number of messages in the local past thus counting more rollbacks than there actually are. This explains why the model is underestimating performance by a small amount. However, the knee for both the analytical and experimental curves occur almost at the same value of memory.

It can be noted here that these results differ from those which we previously reported in [1]. There we subtracted certain Time Warp overheads from the total execution time observed in the Time Warp system to compute speedup. Here we report actual, observed speedup, including all overheads. The present implementation was optimized to reduce these overheads, particularly those associated with GVT computation and fossil collection. In addition, several improvements were made in the model, particularly the way (i) rollbacks due to cancelback are modeled and (ii) the parameters r_β and r_γ are estimated.

6. CONCLUSIONS AND FUTURE RESEARCH

The principal contribution of this work is the development of an analytic model of Time Warp augmented with cancelback operating with limited memory. To our knowledge this is the first attempt to model the limited memory behavior of Time Warp. We developed a Markov chain model and derived expressions to predict the performance of the Time Warp system in terms of speedup.

In addition, we validated our analytical model through measurements on an operational Time Warp system and observed a close match between the model predictions and actual measurements. It was demonstrated that for homogeneous workloads with constant message population Time Warp performs reasonably well with a very limited amount of memory.

Our future work includes the following topics:

- performance comparisons of different memory management schemes (for example, infrequent state saving schemes versus flow control based schemes such as cancelback);
- investigating the cases where Time Warp overheads and heterogeneous processors and/or workloads are taken into account in order to model computations with small granularity; and
- estimating the minimum amount of memory required by Time Warp to achieve nearly the unlimited memory performance for nonhomogeneous workload models.

ACKNOWLEDGMENT

We thank the referees whose insightful comments significantly improved this paper.

REFERENCES

1. Akyildiz, I. F., Chen, L., Das, S. R., Fujimoto, R. M., and Serfozo, R. F. Performance analysis of Time Warp with limited memory. In *Proceedings of the 1992 ACM SIGMETRICS Conference on Measurement and Modeling Computer Systems*, May 1992, pp. 213–224.
2. Felderman, R. E., and Kleinrock, L. An upper bound on the improvement of asynchronous versus synchronous distributed processing. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. 22, No. 1, January 1990, pp. 131–136.
3. Felderman, R. E., and Kleinrock, L. Two processor Time Warp analysis: Some results on a unifying approach. In *Proceedings of the Multiconference on Advances in Parallel and Distributed Simulation*, Vol. 23, No. 1 January 1991, pp. 3–10.
4. Fujimoto, R. M. Time Warp on a shared memory multiprocessor. *Trans. Soc. Comput. Simulation* **6**, 3 (July 1989), 211–239.
5. Fujimoto, R. M. Parallel discrete event simulation. *Comm. ACM* **33**, 10 (Oct. 1990), 30–53.
6. Fujimoto, R. M. Performance of Time Warp under synthetic workloads. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. 22, No. 1 January 1990, pp. 23–28.
7. Gafni, A. Rollback mechanisms for optimistic distributed simulation systems. *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. 19, No. 3 July 1988, pp. 61–67.
8. Gupta, A., Akyildiz, I. F., and Fujimoto, R. M. Performance analysis of Time Warp with multiple homogenous processors. *IEEE Trans. Software Eng.* **17**, 10 (Oct. 1991), 1013–1027.
9. Jefferson, D. R. Virtual time. *ACM Trans. Programming Languages Systems* **7**, 3 (July 1985), 404–425.
10. Jefferson, D. R. Virtual time II: The Cancelback protocol for storage management in distributed simulation. In *Proc. 9th Annual ACM Symposium on Principles of Distributed Computing*, August 1990, pp. 75–90.
11. Jefferson, D. R., and Witkowski, A. An approach to performance analysis of timestamp driven synchronization mechanisms. In *Proceedings of the 3rd Annual ACM Symposium on Principles of Distributed Computing*, 1984, pp. 243–253.
12. Lavenberg, S., Muntz, R., and Samadi, B. Performance analysis of a rollback method for distributed simulation. In *Performance '83* Elsevier, 1983, pp. 117–132.
13. Lin, Y.-B. Memory management algorithms for optimistic parallel simulation. In *Proceedings of the SCS Multiconference on Parallel and Distributed Simulation*, Vol. 24, No. 3, January 1992, pp. 43–52.
14. Lin, Y.-B., and Lazowska, E. D. Optimality considerations of "Time Warp" parallel simulation. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. 22, No. 1, January 1990, pp. 29–34.
15. Lin, Y.-B., and Lazowska, E. D. Reducing the state saving overhead for Time Warp parallel simulation. Technical Report 90-02-03, Dept. of Computer Science, University of Washington, Seattle, Washington, February 1990.
16. Lin, Y.-B., and Preiss, B. R. Optimal memory management for Time Warp parallel simulation. *ACM Trans. Modeling Comput. Simulation* **1**, 4 (Oct. 1991), 283–307.
17. Lipton, R. J., and Mizell, D. W. Time Warp vs. Chandy–Misra: A worst-case comparison. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. 22, No. 1, January 1990, pp. 137–143.
18. Lubachevsky, B. D., Shwartz, A., and Weiss, A. Rollback sometimes works ... if filtered. *1989 Winter Simulation Conference Proceedings*, December 1989, pp. 630–639.
19. Lubachevsky, B. D., Shwartz, A., and Weiss, A. An analysis of rollback-based simulation. *ACM Trans. Modeling Comput. Simulation* **1**, 2 (Apr. 1991), 154–193.
20. Madiseti, V., Walrand, J., and Messerschmitt, D. Synchronization in message-passing computers-models, algorithms, and analysis. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. 22, No. 1, January 1990, pp. 35–48.
21. Mitra, D., and Mitrani, I. Analysis and optimum performance of two message passing parallel processors synchronized by rollback. *Performance Evaluation J.* **7** (1987), 111–124.
22. Nicol, D. M. Performance bounds on parallel self-initiating discrete-event simulations. *ACM Trans. Modeling Comput. Simulation* **1**, 1 (Jan. 1991), 24–50.
23. Plateau, B. D., and Tripathi, S. K. Performance analysis of synchronization for two communicating processes. *Performance Evaluation J.* **8** (1988), 305–320.
24. Preiss, B. R., MacIntyre, I. D., and Loucks, W. M. On the trade-off between time and space in optimistic parallel discrete-event simulation. In *Proceedings of the SCS Multiconference on Parallel and Distributed Simulation*, Vol. 24, No. 3, January 1992, pp. 33–42.
25. Sokol, L. M., Briscoe, D. P., and Wieland, A. P. MTW: a strategy for scheduling discrete simulation events for concurrent execution. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Vol. No. 19, 3, July 1988, pp. 34–42.
26. Turner, S. J., and Xu, M. Q. Performance evaluation of the bounded Time Warp algorithm. In *Proceedings of the SCS Multiconference on Parallel and Distributed Simulation*, Vol. 24, No. 3, January 1992, pp. 117–126.

IAN F. AKYILDIZ received his B. S., M. S., and Doctor of Engineering degrees in computer engineering from the University of Erlangen-Nuernberg, Germany, in 1978, 1981, and 1984, respectively. Currently, he is an associate professor with the School of Electrical Engineering, Georgia Institute of Technology. He is an editor for *Computer Networks and ISDN Systems Journal* and *IEEE Transaction on Computers*. His research interests are computer networks, data communications, performance evaluation, parallel simulation, distributed systems, and computer security. Dr. Akyildiz is an ACM National Lecturer, a member of ACM (Sigcomm, Sigmetrics, and Sigops), and a senior member of IEEE.

LIANG CHEN is a Ph.D. student of the School of Industrial and Systems Engineering, Georgia Institute of Technology. He received B.S. (1981) degree in applied mathematics from Beijing Institute of Posts and Telecommunications, M.S. (1984) degree in systems analysis and control from Beijing University, and Master's degree (1990) in manufacturing systems engineering from the University of Wisconsin-Madison. He was an assistant researcher in the institute of Automation, Chinese Academy of Sciences, from 1984 to 1987. His current research interests focus on performance analysis of telecommunication network and parallel simulation of discrete-event system.

SAMIR RANJAN DAS is a Ph.D. student at the College of Computing at the Georgia Institute of Technology. He received his B.E. degree in electronics and telecommunication engineering from Jadavpur Uni-

versity, Calcutta, in 1986, and M.E. degree in computer science and engineering from the Indian Institute of Science, Bangalore, in 1988. During 1988-1989 he worked in the Indian Statistical Institute, Calcutta, as computer engineer in a research project on knowledge based computer systems. His research interests are in parallel and distributed simulation and computer architecture. He is a student member of ACM and IEEE.

RICHARD FUJIMOTO is an associate professor at the College of Computing at the Georgia Institute of Technology. He received B.S. degrees in computer science and computer engineering from the University of Illinois in Urbana in 1977 and 1978, and M.S. and Ph.D. degrees from the University of California at Berkeley in 1980 and 1983, respectively. His current research interests include parallel and distributed simulation, computer architecture, and parallel processing. He has served as the program and general chairs for the annual *Workshop on Parallel and Distributed Simulation* and currently chairs the steering committee for that conference. He is also an area editor for *ACM Transactions on Modeling and Computer Simulation*.

RICHARD SERFOZO has been professor of industrial and systems engineering at Georgia Tech for 10 years. Prior to that he held positions at Bell Laboratories, Syracuse University, and the Boeing Company. His research area is stochastic processes, namely, modeling of queueing systems, point processes, regenerative and semi-stationary processes, Markov decision processes, and extreme values of processes. His recent work is on queueing networks and parallel simulation.

Received June 22, 1992; revised March 1, 1993; accepted April 22, 1993