Guest Editor's Introduction

Joint Special Issue on Parallel and Distributed Simulation

by Christopher Carothers, Special Issue Guest Editor

Webster's Dictionary defines simulation as "the act of pretending or creating a falsehood." In order to turn these "lies" into "truths," modelers must ensure that the simulation model correctly reproduces the behavior of the system being modeled. As systems become arbitrarily complex, so does the simulation model. The caveat here is that despite the continuing advances being made in processor speed, complex simulation models consistently require more CPU cycles than the state of the art can provide. Today, many simulation applications require days or weeks of CPU time on the fastest uniprocessor machines. The primary research objective of the parallel and distributed simulation community is to reduce the execution time of these large and complex simulations by employing many computers or processors simultaneously. To that end, this joint special issue is dedicated to presenting some of the latest advances in the area of parallel and distributed simulation.

In constructing the first SCS joint *SIMULATION* / *TRANSACTIONS* Special Issue, we selected eight papers from a pool of 20 submissions. Of the eight, we placed four of the finalists in *SIMULATION* and the other four in *TRANSACTIONS* of the SCS. The decision of where to place each article was determined by the reviewers, myself, Bernard Zeigler, the Editor-in-Chief for *TRANSACTIONS*, and Lou Birta, the Editor for Special Issues for *SIMULATION*. These decisions were based on matching the type of paper to the audience of the respective publication. As a rule of thumb, articles that are more theoretical and methodological in nature are placed in *TRANSACTIONS*, and experimental studies and survey articles appear in *SIMULATION*.

For SCS members who are currently not subscribing to *TRANSACTIONS*, we encourage you to do so. To allow anyone (not just current SCS subscribers) a "no-cost" examination of *TRANSACTIONS*, the articles appearing in *TRANSACTIONS* for this special issue will be available for viewing and downloading via the SCS Website at *www.scs.org*.

The first paper appearing in *SIMULATION*, "Parallel Simulation of a Large Number of Processes," by Edwin Naroska and Uwe Schwiegelshohn, describes a novel conservative algorithm for exploiting model parallelism in the presence of a large number of processes. As a demonstration of this approach, they conduct a series of experiments on an IBM SP2 distributed memory machine using a pipelined multiplier model as the example application.

Next, the paper "On the CRAY-System Random Number Generator," by Karl Entacher, analyzes the RANF random number generator found on a CRAY supercomputer from both a theoretical as well as an empirical point of view. This article serves as an example of how a parallel random number generator should be analyzed in practice.

The paper, "Survey of Languages and Runtime Libraries for Parallel Discrete Event Simulation," by Yoke-Hean Low, Chu-Cheow Lim, Wentong Cai, Shell-Ying Huang, Wen-Jing Hsu, Sanjay Jain and Stephen Turner, provides an in-depth survey of current research and commercial PDES systems. These systems are compared based on user model, programming framework, language features, and protocol, to name a few.

In addition to making simulations execute faster, simulator interoperability is another major thrust area in the field of parallel and distributed simulation, such as efforts in the DoD High-Level Architecture (HLA). With that in mind, the final paper appearing in

SIMULATION is "Synchronization Mechanisms for Integration of Distributed Manufacturing Simulation Systems," by Susumu Fujii, Yasushi Kidani, Atsushi Ogita and Toshiya Kaihara. It presents a new approach integrating different subcomponents of manufacturing models, enabling a complete distributed manufacturing simulation system.

The first paper appearing in *TRANSACTIONS* of the SCS is "High-Performance Hardware Description Language Simulation: Modeling Issues and Recommended Practices," by Gregory Peterson and John Willis. This article presents some "rules of thumb" for parallelizing hardware description language (HDL) simulations. These modeling practices are part of a larger effort by the IEEE Design Automation Standards Committee (DASC) on High Performance Modeling and Simulation to develop a set of guidelines for HDL modeling.

Next, mean value analysis (MVA) is a well-known technique for analyzing closed queueing networks. In the paper "Parallelising the Mean Value Analysis Algorithm," by Claudio Gennaro and Peter King, a novel parallel MVA algorithm is presented. Their approach obtains speedups in excess of 200 on a Cray T3D multiprocessor system.

One of the critical issues to making not only parallel simulation but any simulation accurate is processing events in the correct timestamp order. However, in the presence of events with identical timestamps, the question becomes what is that correct order. In the paper "The Threshold of Event Simultaneity," by Frederick Wieland, a new view of simultaneous event processing is presented. Previous work in this area has largely concentrated on tie-breaking strategies. This paper shows how in many cases these strategies lead to inaccurate or even invalid simulations.

Crucial to achieving good performance from any Time Warp parallel simulation is the minimization of state-saving overheads. In Time Warp systems there are largely two state-saving methodologies: (1) periodic state-saving, which copies the state of a logical process every "k" events, and (2) incremental state-saving, which records only the changes made to logical processes during event processing. There exist cases where one method is more appropriate than the other. The final paper, "On the Selection of the State Saving Strategy in Time Warp Parallel Simulation," by Hussam Soliman, derives the new selection criteria for determining which mechanism should be employed.

In closing I want to thank Lou Birta, Bernie Zeigler and all the reviewers for their assistance in the development of this first-ever joint special issue. I would especially like to thank the reviewers listed below for their time and effort in helping to make this special issue. Without them, this issue would not be possible.

Christopher D. Carothers
Assistant Professor, Dept. of Computer Science
Rensselaer Polytechnic Institute
110 Eighth Street, Troy, New York, 12180-3590
chrisc@cs.rpi.edu