UNIVERSITY OF TENNESSEE CHATTANOOGA

SimCenter: Center of Excellence in Applied Computational Science and Engineering

presents

Checkpointing the Un-checkpointable: A Tour of the Split-Process Approach and its Applications

given by **Dr. Gene Cooperman**

Professor, Khoury College of Computer Sciences
Northeastern University

Wednesday, October 30th, 2:00 p.m., UTC SimCenter Auditorium*

Public Invited

Checkpointing is the ability to save the state of a running process to stable storage, and later restart, perhaps on a different computer. Transparent checkpointing (or system-level checkpointing) is the ability to checkpoint a (possibly parallel or distributed) application, without modifying the application binaries. The speaker has led the DMTCP project in this area for 15 years. It is argued that a new split process approach can revolutionize checkpointing and its applications.

We begin with a simple motivating example of split processes from formal verification, based on a small SimGrid demonstration. A standard multithreaded application executes under SimGrid, in order to find bugs. Upon finding a bug, a schedule of thread execution is then produced to illustrate the bug. Next, we begin a tour of split processes and its applications, by describing how to use this to diagnose such bugs deep in the execution of the multithread application. Continuing the tour, ongoing work is described that will extend this SimGrid/split-process approach to diagnose deadlock and other MPI bugs deep in their execution. Some practitioners have joked that we don't debug high-end MPI applications, due to the expense of running on high-end clusters. This SimGrid/split-process approach may help in the future in bringing back debugability. From there, this tour of split processes goes on to demonstrate low-overhead migration of MPI applications among wildly different clusters (e.g., TCP versus InfiniBand, or varying the number of cores per host). Indeed, that was the original motivation or split processes. Finally, the tour concludes with a description of ongoing work intended to apply this same flexible framework for efficient checkpointing and software migration in the cases of CUDA applications, and of numerical toolkit libraries.

The split-process approach represents joint work with Rohan Garg. The SimGrid discussion is based on joint work with Martin Quinson.

Professor Cooperman works in high-performance computing and scalable applications for computational algebra. He received his B.S. from the University of Michigan in 1974, and his Ph.D. from Brown University in 1978. He then spent six years in basic research at GTE Laboratories. He came to Northeastern University in 1986, and has been a full professor there since 1992. His visiting research positions include a 5-year IDEX Chair of Attractivity at the University of Toulouse/CNRS in France, and sabbaticals at Concordia University, at CERN, and at Inria (France). He is one of the more than 100 co-authors on the foundational Geant4 paper, whose current citation count is at 25,000. The extension of the million-line code of Geant4 to use multi-threading (Geant4-MT) was accomplished in 2014 on the basis of joint work with his PhD student, Xin Dong. Prof. Cooperman currently leads the DMTCP project (Distributed Multi-Threaded CheckPointing) for transparent checkpointing. The project began in 2004, and has benefited from a series of PhD theses. Over 100 refereed publications cite DMTCP as having contributed to their research project.

**UTC MDRB, Auditorium, 701 E. M.L. King Blvd., Chattanooga TN, 37403*