The 2014 High Performance Computing & Simulation International Conference
(HPCS 2014)
July 21 - 25, 2014
Bologna, Italy

TUTORIAL VI

Parallel Discrete Event Simulation

Kalyan S. Perumalla
Oak Ridge National Laboratory, Tennessee, USA
perumallaks@ornl.gov

TUTORIAL DESCRIPTION

This tutorial introduces the fundamental principles and algorithms underlying parallel/distributed discrete event simulation (PDES), with a special focus on synchronization methods and their scalable implementation. The tutorial is designed to help the audience gain a detailed understanding of the core concepts and terminology in PDES. Further, classical conservative and optimistic algorithms such as null messages and Time Warp will be described. An overview of approaches customized for different computing platforms will be provided, including supercomputers, cloud computing, and GPUs. Implementations customized for different communication types will be described, including 1-sided and 2-sided messaging using Portals and Message Passing Interface (MPI) respectively. The critical influence of event granularity, timestamp distributions, and inter-logical process event exchange behaviors on runtime dynamics will be described. Important systems issues such as computation and communication overheads will be highlighted, along with illustrative performance results on benchmarks from multiple domains including transportation, epidemiology, and supercomputer simulation.

TUTORIAL OUTLINE

The following is the general outline of the topics covered in the tutorial.

- History of PDES, Motivation, Use cases, Applications
- Background
  - Pacing Types, Timelines, Orderings, Spatial & Temporal Parallelism, Granularity
  - Logical Processes, Memories, Simulation loops, Events, etc.
- Synchronization Problem and Solution Classes
  - Centralized Coordination, Shared Event Queues
  - Conservative Synchronization - Null Message scheme, Global Reduction schemes, Lookahead (LA), Dynamic/Static Changes to LA, LA Extraction
  - Optimistic Synchronization - Time Warp, Rollback, Global Virtual Time, State Saving Variants, Anti messages, Fossil collection
  - Unified Synchronization - Conservative+Optimistic, Micro-kernel approach, Lower Bound on incoming Time Stamps
  - Advanced Methods - Approximate Time, Cloning, Synchronization over Unreliable Network Transport
- Hardware-Optimized Methods
  - Network co-processors, GPUs, Hybrid cores
  - GPU-based SIMD Challenges, CPU-GPU Hybrid Algorithms, Agent-based Models on GPUs

- PDES on Cloud Computing
    - GVT-Vs-Fairness Scheduling Issues, PDES-Specific Hypervisor schedulers
    - Counter-intuitive Cost-time Tradeoffs
    - Benchmark Results on Xen-based Server, and on Amazon EC2 Cloud
- Scaling Results on Supercomputers
    - PDES execution on large Blue Gene and Cray XT machines
    - Interconnection Network Considerations, 1-Sided and 2-Sided Messaging
    - Synchronous and Asynchronous GVT Algorithms, MPI-based and Portals-based Implementation
    - Event Dynamics and Scaling Results with Discrete Event Models from Epidemiological Simulations, Transportation Simulations, Parallel Program Simulations

SELECTED REFERENCES

1. Chandy, K. M. and J. Misra, "Distributed Simulation: A Case Study in Design and Verification of Distributed Programs," IEEE Transactions on Software Engineering, Vol. SE-5(5), pp. 440-452, 1978.
2. Chandy, K. M. and J. Misra, "Asynchronous Distributed Simulation via a Sequence of Parallel Computations," Communications of the ACM, vol. 24(4), pp. 198-205, 1981.
3. Jefferson, D., "Virtual Time," ACM Transactions on Programming Languages and Systems, vol. 7(3), pp. 404-425, 1985.
4. Fujimoto, R. M., "Parallel Discrete Event Simulation," Communications of the ACM, Vol. 33(10), pp. 30-53, 1990.
5. Mattern, F., "Efficient Algorithms for Distributed Snapshots and Global Virtual Time," Journal of Parallel and Distributed Computing, Vol. 18(4), pp. 423-434, 1993.
6. Kuhl, F., R. Weatherly and J. Dahmann, "Creating Computer Simulation Systems: An Introduction to the High Level Architecture for Simulation," Prentice Hall, 1999.
7. Fujimoto, R. M., "Parallel and Distributed Simulation Systems," Wiley-Interscience, 1st Edition, ISBN 978-0471183839, 2000.
8. Perumalla, K. S. and R. M. Fujimoto, "Virtual Time Synchronization over Unreliable Network Transport," IEEE Workshop on Parallel and Distributed Simulation, 2001.
9. Perumalla, K. S., "µsik - A Micro-Kernel for Parallel/Distributed Simulation Systems," Workshop on Principles of Advanced and Distributed Simulation, 2005.
10. Perumalla, K. S., "Discrete Event Execution Alternatives on General Purpose Graphical Processing Units (GPGPUs)," International Workshop on Principles of Advanced and Distributed Simulation, 2006.
11. Holder, A., and C. Carothers, "Analysis of Time Warp on a 32,768 Processor IBM Blue Gene/L Supercomputer," European Simulation Symposium, 2008.
12. Bauer, D., C. Carothers, and A. Holder, "Scalable Time Warp on Blue Gene Supercomputers," ACM/IEEE Workshop on Parallel and Distributed Simulation, 2009.
13. C. Carothers and K. Perumalla, "On Deciding between Conservative and Optimistic Approaches on Massively Parallel Platforms," Winter Simulation Conference, 2010.
14. Wenjie T., Y. Yao, "A GPU-Based Discrete Event Simulation Kernel," SIMULATION, November 2013; Vol. 89 (11), pp. 1335-1354.
15. Perumalla, K. S., A. J. Park, and V. Tipparaju, "Discrete Event Execution with One-sided and Two-sided GVT Algorithms on 216,000 Processor Cores," to appear in ACM Transactions on Modeling and Computer Simulation (TOMACS), 2014.

REQUIREMENTS AND TARGET AUDIENCE

- Familiarity with parallel or distributed computing (at senior undergraduate or higher level) will be assumed.
- The tutorial is suitable for anyone interested in gaining an understanding about the state-of-the-art in PDES.
- Those intending to develop PDES applications will benefit from learning key considerations for runtime efficiency, scalability, and model organization.
- Researchers interested in building new PDES engines or customizing existing ones will benefit from understanding the common data structures and runtime operation of the underlying PDES algorithms.
- Parallel computing system designers will benefit from learning the unique runtime dynamics and performance requirements of PDES engines and applications.

TUTORIAL DURATION

The tutorial material will be presented in a 2.5-hour session.

A/V AND EQUIPEMNT

A standard projector and screen is sufficient for this tutorial.

INSTRUCTOR BIOGRAPHY

Kalyan Perumalla, Ph.D., is a senior R&D staff member and manager at the Oak Ridge National Laboratory. He is the founding Group Leader of the Discrete Computing Systems Group in the Computational Sciences and Engineering Division. He also serves as an adjunct professor in the School of Computational Sciences and Engineering at the Georgia Institute of Technology.

Dr. Perumalla is a winner of the US Department of Energy Early Career Award in Advanced Scientific Computing Research, 2010-2015. His primary research contributions are in the application of reversible computation to high performance computing and in advancing the vision of a new class of supercomputing applications using parallel discrete event simulations. He has published articles and delivered distinguished lectures and invited tutorials on these topics. His recent book *Introduction to Reversible Computing* is among the first few in its area. He co-authored another book, three book chapters, and over 100 articles in peer-reviewed conferences and journals. Four of his co-authored papers received the best paper awards, in 1999, 2002, 2005 and 2008, and two were finalists in 2010. His research prototype tools in parallel and distributed computing have been disseminated to research institutions worldwide. He earned his Ph.D. in computer science from the Georgia Institute of Technology in 1999. Dr. Perumalla serves as program committee member, editorial board member, and reviewer for multiple international conferences and journals in computing.