The Numerical Reproducibility Fair Trade: Facing the Concurrency Challenges at the Extreme Scale

Michela Taufer, University of Delaware

Trends in execution concurrency make a compelling case for the development of methods able to automatically and efficiently model and mitigate irreproducibility beyond petascale architectures and into the exascale. It is expected that high performance computers at the exascale will exhibit a massively large level of concurrency - a factor of 10,000 greater than on current platforms - which will move computer simulations from bulk-synchronous executions to multithreading approaches and asynchronous I/O. Simulation calculations and analysis routines will also be tightly coupled on exascale platforms, requiring these two workflow components to work at extremely high levels of concurrency. As concurrency levels increase, the impact of rounding errors on numerical reproducibility also increases, ultimately affecting the ability of scientific simulations to reproduce program executions and numerical results. Under these circumstances, irreproducible results may not be trusted by a scientific community expecting reproducible behaviors and any attempt to pursue reproducibility may come at a cost in performance that is too high.

In this talk we present studies on the impact of rounding errors on result reproducibility when concurrent executions burst and workflow determinism vanishes in cutting-edge multicore architectures. Specifically, we discuss a mathematical method called "composite precision floating-point arithmetic" to model rounding-errors in scientific applications and show how this method can mitigate error drifting. We assess the impact of the composite precision on new generations of multicore architectures by measuring the cost and mitigation factors of the proposed method on error propagations for a diverse set of benchmarks and platforms.