Title:
Optimizing Application Workflows for Power-Efficient Resilience

Authors:

Liang Wang, *University of Virginia*
Augusto Vega, *IBM T. J. Watson Research Center*
Hans Jacobson, *IBM T. J. Watson Research Center*
Chen-Yong Cher, *IBM T. J. Watson Research Center*
Alper Buyuktosunoglu, *IBM T. J. Watson Research Center*
Pradip Bose, *IBM T. J. Watson Research Center*
Kevin Skadron, *University of Virginia*

Abstract:
A workflow, consisting of segments of individual applications and represented as a dependence graph is targeted for *optimized* execution on an underlying compute node. The context could be a real-time constrained embedded system or a large-scale scientific or commercial system without an explicit real-time constraint. Depending on the context, the metric of optimization could be throughput performance, the net execution time or an efficiency feature like performance per watt. The optimization intent must necessarily also be guarded (or constrained) by specified system resilience metrics. Thus, in effect, the problem at hand can be formulated as an integrated power-performance-reliability modeling and optimization challenge for a given {workflow, processor} pair. In this talk we will present an integrated modeling framework, called PEARL, that is positioned to provide such a cross-metric optimization facility for application developers. Although this facility is being developed under the DARPA-sponsored PERFECT program, the features can be adapted or configured to cater to other computing contexts – e.g. a compute node of a next generation supercomputer system.