Analyzing Large-Scale Heterogeneous Systems through the Ages

Darren Kerbyson
Pacific Northwest National Laboratory
Darren.Kerbyson@pnl.gov

Multi-core, heterogeneity, as well as memory and network hierarchies are already here. As a famous 20th thinker once said: "The future will be like the present only more so". In this talk I will examine a number of issues that have been explored at different system ages: from early design where tradeoffs in key architectural parameters are explored, through to early installations where configuration optimizations can be achieved before actual production. The system's healthiness and its ability to achieve application performance is of prime importance to production users. Performance modeling can (and has) been used in all of these stages of a system's development. Examples will be drawn from: design space exploration of heterogeneous accelerated systems which preceded the petascale Roadrunner system and algorithmic optimizations that efficiently use hierarchical communication structures. Several different programming models are available in heterogeneous systems including work-queue approaches portable Thread-building-blocks, and MPI on accelerator-cores. These will be examined in the context of the kernel application, Sweep3D. Achieving application performance is the key - it can be impeded or assisted by the capabilities of a socket, configuration of a node, or system connectivity. As the depth of system hierarchies and complexity increase, the challenges of achieving high application performance will increase many-fold also.

A brief biography:

Darren Kerbyson recently became a Laboratory Fellow at the Pacific Northwest National Laboratory. He was previously the lead of the Performance and Architecture Lab (PAL) at Los Alamos National Laboratory and prior to 2001 he was a senior Lecturer in Computer Science at the University of Warwick in the UK. He has been active in the areas of performance modeling, and performance evaluation of large-scale systems, for nearly 20 years. He has worked on many performance projects funded by the DOE, DARPA, NSF and European agencies. He has published over 140 papers, taught tutorials at many events as well as at undergraduate and postgraduate levels. He is currently involved in the modeling of large-scale applications on the largest systems available today and foreseen tomorrow. Recent work has included analysis of the

Roadrunner system at Los Alamos, the Blue Waters system at NCSA, and upgrades to Jaguar at ORNL. His work on the performance analysis of ASCI Q during its installation lead to the best paper award at Supercomputing in 2003. He is the recipient of the Los Alamos Achievement Award in 2003, 2004, 2005, 2006 and 2008, and several awards from the Department of Defense.