HCW Keynote Address
ParalleX: An Asynchronous Execution Model for Scalable Heterogeneous Computing

Thomas Sterling[1,2,3]

[1] Department of Computer Science
Center of Computation and Technology
Louisiana State University

[2] Center for Advance Computation Research
California Institute of Technology

[3] Computing and Computational Sciences Directorate
Oak Ridge National Laboratory

Abstract

Heterogeneous system architecture has long been appreciated as a potential strategy for achieving super-linear speedup with respect to some normalizing parameter like number of nodes, cost, or power. However the challenge of programming and managing the system resources has been a limiting factor for the application of such systems on a broad scale. The exception has been the use of special purpose processors such as graphics processing units that may yield dramatic increases for such functionality enabling capabilities largely impossible otherwise like realistic high resolution real time interactive games. However, with power emerging as the dominant constraint on high performance computing and the need to make better use of logic and storage resources such components as the ClearSpeed SIMD attached processor and the IBM cell architectures among others is forcing mainstream computing to adopt heterogeneous processing. This keynote presentation will describe a computational model, ParalleX that provides an asynchronous runtime framework for supporting effective execution in an environment comprising heterogeneous elements. ParalleX is based on a message-driven split-phase multithreaded transaction processing paradigm synthesizing a number of concepts represented in prior art that in ensemble will facilitate management of heterogeneous resources and provide the basis for a systematic programming methodology. Also discussed in this presentation is another example of a heterogeneous architecture, Gilgamesh II, that provides separate mechanisms for computations that exhibit disparate locality properties.

Speaker biography: Thomas Sterling holds joint appointments as Professor of Computer Science at Louisiana State University, Faculty Associate at the California Institute of Technology, and Distinguished Visiting Scientist at Oak Ridge National Laboratory. Since receiving his Ph.D from MIT as a Hertz Fellow in 1984, Dr. Sterling has engaged in applied research in the field of parallel computer architecture and systems at Harris Corporation, the Institute of Defense Analysis, NASA Goddard Space Flight Center and the Jet Propulsion Laboratory. A leader in the field, Professor Sterling created the Beowulf Project in 1993 that contributed to the current dominance of Linux-based commodity clusters and for which he was awarded the Gordon Bell prize with colleagues. An early contributor to the national Petaflops initiative, Sterling was the PI of the multi-agency HTMT project that was the first to explore the design space for Petascale systems employing advanced technologies and innovative architecture concepts. He has engaged in a series of projects to explore the potential of processor-in-memory architectures including the USC DIVA Project, the NASA Gilgamesh Project, and the Cray Cascade Project which together culminated in the MIND architecture concept. He is currently developing the "ParalleX" advanced execution model to guide the development of high performance system architecture and programming models. In collaboration with Sandia National Laboratory, University of New Mexico, University of Delaware, and Oak Ridge National Laboratory he is exploring potential practical implementations of ParalleX to achieve high efficiency and scalability towards Exaflops capability by the end of the next decade. Thomas Sterling is developing a new graduate level HPC Introductory course and new textbook, "High Performance Computing Concepts Methods and Means", in support of it. Dr. Sterling is a noted speaker and co-author of five books. He holds six patents.

1-4244-0910-1/07/$20.00 ©2007 IEEE