

HPC in Phase Change: Towards a New Execution Model

Thomas Sterling

Louisiana State University, USA

Abstract. HPC is entering a new phase in system structure and operation driven by a combination of technology and architecture trends. Perhaps foremost are the constraints of power and complexity that as a result of the flat-lining of clock rates relies on multicore as the primary means by which performance gain is being achieved with Moore's Law. Indeed, for all intense and purposes, "multicore" is the new "Moore's Law" with steady increases in the number of cores per socket. Added to this are the highly multithreaded GPU components moving HPC into the heterogeneous modality for additional performance gain. These dramatic changes in system architecture are forcing new methods of use including programming and system management. Historically HPC has experienced five previous phase changes involving technology, architecture, and programming models. The current phase of two decades is exemplified by the communicating sequential model of computation replacing previous vector and SIMD models. HPC is now faced with the need for new effective means of sustaining performance growth with technology through rapid expansion of multicore with anticipated structures of hundreds of millions of cores by the end of this decade delivering Exaflops performance. This presentation will discuss the driving trends and issues of the new phase change in HPC and will discuss the ParalleX execution model that is serving as a pathfinding framework for exploring an innovative synthesis of semantic constructs and mechanisms that may serve as a foundation for computational systems and techniques in the Exascale era. This talk is being given just as DARPA is initiating its UHPC program and DOE is launching additional programs such as their X-stack all aimed at catalyzing research in to the challenging area.