Stagnation at Moore's Nano-scale Barrier and an HPC Architecture Reboot

Thomas Sterling
Indiana University

Abstract

After some four decades of exponential growth in processor density and performance, the asymptotic approach of von Neumann architecture to nano-scale feature size is imposing fundamental barriers to further improvements through conventional practices. Almost all conventional commercial high performance computing (HPC) systems over this time period are von Neumann derivative with the optimal designs based on obsolete premises from earlier enabling technologies such as vacuum tubes and core (magnetic) memories. Such assumptions include: emphasis on floating point logic utilization, sequential instruction issue, separation of processor structure from main memory, sequential memory consistency, program counter control state, and inter-processor communication as software message passing through I/O channels. Current priorities driven by technology advances demand a significantly different set of optimizations where memory bandwidth and latency as well as reduced overhead for parallel control flow mechanisms must be highlighted to yield substantial increases in parallelism for reduction to time to solution. This presentation will lay out the fundamental motivations for a dramatic shift from HPC system von Neumann derivative architectures to a post Moore's Law era non-von Neumann architectures. A specific example, "Continuum Computer Architecture" (CCA), will be considered to demonstrate one possible way to address the new priorities of future nano-scale based semiconductor for scalable architectures in the exascale performance regime. Questions from the audience will be welcome throughout the presentation and afterwards.