

Abstract

The MIND Architecture for Practical Trans-Petaflops Computing

Even as the Earth Simulator continues to dominate high end computing after more than two years of successful operation, plans are underway by several teams to design and implement Petaflops-scale computing systems by the end of this decade; a performance gain of more than 25X. Extrapolation of the Top-500 list over the last decade implies that the first general purpose 1 Petaflops system as measured by the Linpack benchmark will be deployed sometime in 2009. One or more special purpose machines may deliver 1 Petaflops sustained performance on their respective target applications as early as 2007. However, in spite of these indicators, it is unclear when practical Petaflops scale systems will become available; systems that exhibit the necessary properties to enable them to be employed in existing environments, including: 1) total power consumption on the order of 1 Megawatt, 2) floor space of less than 4000 square feet, 3) sustained efficiency of > 50%, 4) cost < \$10M, 5) general purpose operation across a wide range of user applications, and 6) high programmability. Conventional computer architectures that have dominated the last decade are inadequate to meet these difficult but necessary goals, because in most cases they do not incorporate components designed to address the major challenges to effective parallel execution: latency, contention, starvation, and overhead. A proposed class of advanced processor in memory architecture is being developed at the California Institute of Technology to provide practical Petaflops scale computing. The MIND architecture integrates multiple combined processor/memory nodes with wide-word access on a single chip for memory bandwidth 100 – 1000X that of conventional systems of the same memory capacity. Combining multithreaded execution control, parcel message-driven computation, and in-memory virtual to physical address translation, the MIND architecture will enable very large arrays of such devices in a system exporting a single system image and achieving very high efficiency due to the intrinsic latency hiding methods and low overhead mechanisms. Using conventional packaging techniques, a Petaflops computer built with MIND chips would require a floor space of less than 1000 square feet. With advanced 3-D packaging, a peak Petaflops computer based on the MIND architecture would require approximately 10 square feet. This presentation will describe the advanced concepts of the MIND architecture and demonstrate through parametric simulation studies and other analysis the level of capability that may be achieved with it.