Message-driven Split-transaction Execution to Enable Scalable Parallel Architecture for trans-Petascale Domain Computing

Thomas Sterling

Louisiana State University
And
California Institute of Technology

Abstract

Passed methods of organizing parallel hardware structures and applying them to parallel application algorithms have allowed the exploitation of course grained ensembles of sequential commodity components but will not scale for efficient and cost effective computing across the trans-Petaflops performance regime towards an Exaflops and beyond. This limitation is due to the growing need to exploit an increasing amount of algorithm parallelism requiring effective use of finer grain parallelism while at the same time overcoming the increase in global latency (measured in processor cycles). An aggravating factor is a trend to more memory intensive applications for manipulating meta-data such as sparse matrices, adaptive mesh refinements, particle codes (e.g., n-body tree codes or particle in cell codes), or knowledge management through dynamic directed graph processing (e.g., semantic nets) for data mining. Mechanisms derived and implemented specifically to realize efficient management of fine-grain parallelism and synchronization while intrinsically responding to communication latency (and contention due to inadequate bandwidth). A second aggravating factor is the rapid trend to multi-core components flooding the market putting more memory demand on the wrong side of the I/O pins. By the end of the Petaflops era, concurrency will exceed billion-way parallelism which in many cases will come from enormous data structures as defined by relational meta-data. Future parallel architectures will have to respond to these challenges to continue performance growth while improving efficiencies that are manifested as dramatic reductions in power, size, cost, and failure rates per unit of operating speed. The old model of computation widely employed does not provide a clear path to meeting these requirements or address these trends and may be eclipsed by a new strategy including new execution models and new architectures that supercede simple aggregations of sequential microprocessor cores. This talk will present one possible strategy based on a message-driven split-transaction execution model, ParalleX, and a corresponding new class of processor-in-memory architectures (PIM), MIND, that provides efficient realization of that parallel model. A transaction model which performs only local operations per transaction event driven by action specifying messages supported by hardware mechanisms and tightly coupled (on-chip) to main memory stacks address all of the problems cited above and may provide the basis of future Petascale computing.