Licensable Technologies

Scalable and Efficient Computer Systems

Applications:

- Real-time handling of streaming data
- Signal processing
- Molecular dynamics
- Correlation of radio astronomy signals

Benefits:

- Improved programmability
- Higher efficiency
- Extremely scalable
- Optimized for performance per watt
- Failing nodes can be isolated and replaced while a large computation continues to run (high availability)

Contact:

Aaron G. Sauers , (505) 665-0132
asauers@lanl.gov

tmt-4@lanl.gov

Technology Transfer Division

Los Alamos NATIONAL LABORATORY ESTABLISHED 1943

Summary:

A scissors crossover provides flexible mobility for train cars between parallel tracks. LANL is pursuing a design to allow flexible and efficient movement of data in scalable computer systems, providing a novel capacity to manage high-volume data-streams.

Most supercomputer applications require some non-local communication. As a result, the relatively high-latency and low-bandwidth interconnection network becomes a limiting factor on the machine's efficiency. In addition, designers are extending the peak performance of supercomputers by adding multi-core accelerators such as Cell processors or Graphics Processing Units (GPUs). This introduces another high-latency and low-bandwidth bottleneck, at the point where data moves into and out of the accelerator, as well as another dimension of complexity in software.

These factors limit the kinds of applications that can run effectively on supercomputers, and increase the cost of developing or porting those applications. Algorithms that require intercommunication result in underutilized components, wasting energy and the potential of the machine. Furthermore, there appear to be some problems which perform poorly on these architectures, regardless of optimization.

Los Alamos National Laboratory (LANL) researchers have developed a new design to support construction of large machines, allowing the machines to perform closer to their theoretical peak. This approach emphasizes scalable throughput rather than attempting to tailor machines around the highest performing accelerators, and allows selection of individual components that maximize performance against energy draw or cost. The design makes use of commodity components that are modest in computing power and energy consumption.

The LANL hardware is being co-designed along with a powerful and expressive high-level programming language, adapted from a well-studied body of research languages. It is expected that applications written in this language will require no other system-level or low-level programming in order to run efficiently, but diagnostic feedback could allow selection of more efficient idioms.

LANL's design supports the data-intensive applications currently encountered in scientific computing, while opening the door to new levels of capability for communication-intensive and throughput-intensive applications such as molecular dynamics and signal correlation. In addition, researchers expect the LANL design can support transparent fail-over, allowing failed nodes to be replaced on-the-fly without stopping ongoing computations.

Development Stage: Early-stage development of a research prototype

Patent Status: Seeking patent protection

Licensing Status: Available for exclusive or non-exclusive licensing. Researchers are willing to partner with industry in order to develop the technology for specific applications.

www.lanl.gov/partnerships/license/technologies/

An Equal Opportunity Employer / Operated by Los Alamos National Security LLC for DOE/NNSA