

“Reimplementing the Cray XMT programming model in software to execute on commodity components promises both to broaden applicability and to perpetuate scalable performance on large irregular problems.”

- Pacific Northwest National
Laboratory Task Lead
Simon Kahan


SoftXMT

Evaluating commodity-based approaches to scaling irregular computation

At a glance

Today, the Cray XMT solves certain large “irregular problems” arising in today’s national security, engineering and science applications with performance unmatched by any other computer system. Going forward, these applications will demand ever increasing performance. To meet this demand requires further advances in computing technology. Unfortunately, redesigning the Cray XMT’s

Threadstorm processor from the ground up would be prohibitively expensive. This evaluation task explores cost-effective alternatives.



What we do

Though the cost of developing computer systems is distributed across both hardware and software, most of the innovation and execution risk is in the hardware path. Therefore, custom software solutions leveraging commodity hardware tend to minimize development risk. Over time, they also reap price-performance improvements driven by competitive pressures within the broad commodity marketplace.

Therefore, we focus our evaluation of cost-effective alternatives to the Cray XMT on systems built largely or wholly from commodity hardware components. What we call “SoftXMT” are software-based system solutions leveraging Multicore Processors and Field Programmable Gate Arrays (FPGAs).

How we do it

Multicore Processors are prevalent in today’s computers: from handhelds to warehouse-sized supercomputers, they represent the dominant paradigm in computing. Though the Cray XMT’s Threadstorm processors are slower on almost all computations than Multicore Processors, they do possess features that enable them to work together efficiently on large irregular problems. One of the key features is latency tolerance: Threadstorm processors keep themselves busy even while waiting for remote fetches of faraway data by juggling up to 128 subcomputations simultaneously. Doing so requires both an efficient mechanism for switching amongst these computational contexts -- the hands of the juggler -- and the capacity to keep track of the thousand or more references in flight -- the juggler’s eyes. Multicore Processors lack both hands and eyes. In performing our evaluation, we must determine whether or not software can compensate adequately for these missing features.

FPGAs are configurable components that can be programmed to deliver specific functionality at greater efficiency than multicore processors, though as general purpose computing devices they are inherently slower. By programming them to provide the hands and eyes -- as well as other Threadstorm features -- that Multicore Processors lack, we may be able to provide hardware compensation that puts commodity systems on par or ahead of the Cray XMT, even within its sweetspot for performance. Our evaluation will explore architectures incorporating FPGAs to see what it would take for them to provide a viable alternative.

Applications

Should we establish a viable alternative to reimplement of Threadstorm, we will champion its development. This would ensure a strong and cost-effective future for Cray XMT applications.

CASS-MT is dedicated to research on systems software, programming environments, and applications in a High-Performance Computing (HPC) multithreaded architecture environment.


We offer the only Open Science Cray XMT system, a one-of-a-kind supercomputer consisting of 128 multithreaded processors, 1 TB RAM, and a 7.7 TB Lustre parallel filesystem.

The Cray XMT supercomputer has the potential to substantially accelerate data analysis and predictive analytics beyond the limitations of traditional computing. Multithreaded processors allow multiple, simultaneous processing, helping researchers find solutions to the world's most complex challenges faster. The XMT can process irregular, data-intensive applications that have random memory access patterns. Unlike many applications where data delivery is dependent on memory speed, the Cray XMT's multi-threaded architecture tolerates memory access latencies by switching context between multiple threads that work continuously, overlapping the memory latency and preventing the processor from being held up while it waits for data to arrive.

The multithreaded technology powering our Cray XMT is ideally suited to perform pattern matching, scenario development, behavioral prediction, anomaly identification, and graph analysis.

Try it for yourself. We seek to create collaborations and provide expertise for porting and optimizing applications. The opportunity to use our Cray XMT system is available to internal and external research partners.

John Feo,
CASS-MT Director
(509) 375-3768
John.feo@pnl.gov
cass-mt.pnl.gov/



Simon Kahan
Task Lead
Pacific Northwest National Laboratory
206-685-3853
simon.kahan@pnl.gov


Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by **Battelle** Since 1965