



RESEARCH LANDSCAPE

.....

Ohio's strengths in basic and applied research are broad and deep, spanning a multitude of academic, business and industrial interests. The spectrum of clients served by the Ohio Supercomputer Center encompasses many fields of study, from establishing secure, firewall-free network architecture to improving injection molds for manufacturing, as highlighted on the following pages.



PARALLEL COMPARISONS

Bokhari conducts comparisons of system architectures

Surveying the wide range of parallel system architectures offered in the supercomputer market, a former Ohio State University graduate research student sought to establish some side-by-side performance comparisons.

“We explored the parallelization of the subset-sum problem on three contemporary but very different architectures: a 128-processor Cray massively multithreaded machine, a 16-processor IBM shared memory machine and a 240-core NVIDIA graphics processing unit,” said former computer science and engineering graduate student Saniyah Bokhari. “These experiments highlighted the strengths and weaknesses of these architectures in the context of a well-defined combinatorial problem.”

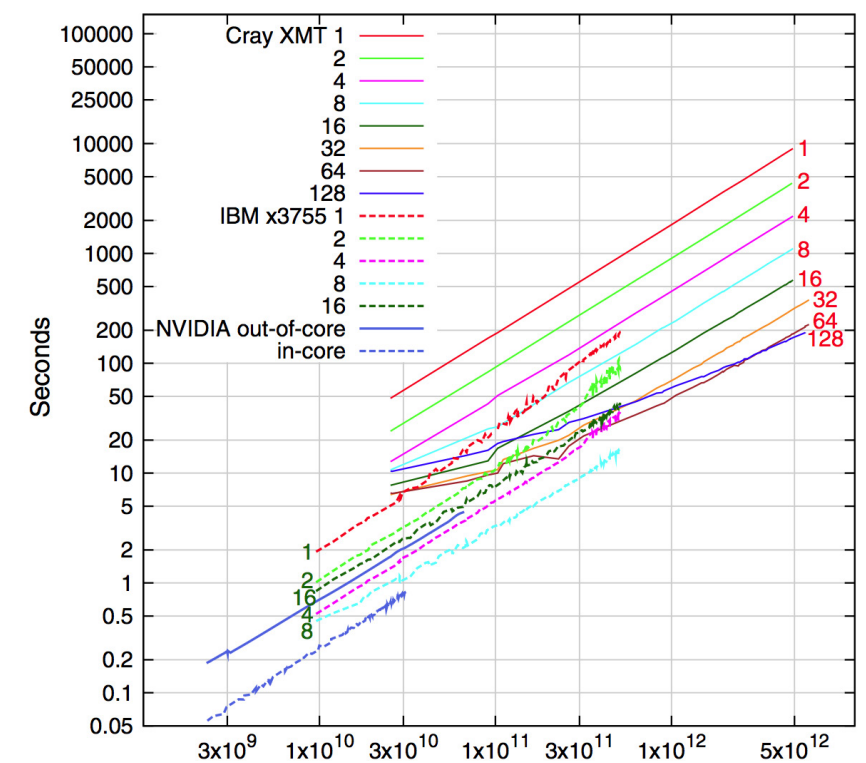
Bokhari evaluated the conventional central processing unit architecture of the IBM 1350 Glenn Cluster at the Ohio Supercomputer Center and the less-traditional general-purpose graphic processing unit architecture available on the same cluster. She also evaluated the multithreaded architecture of a Cray Extreme Multithreading (XMT) supercomputer at the Pacific Northwest National Laboratory’s Center for Adaptive Supercomputing Software.

Each of the architectures Bokhari tested fall in the area of computing where multiple processors are used to tackle pieces of complex problems “in parallel.” To solve the subset-sum problem she used an algorithm that takes a period of time that is proportional to the number of objects entered, multiplied

by the sum of their sizes. Also, she carefully timed the code runs for solving a comprehensive range of problem sizes.

Bokhari’s results illustrate that the subset-sum problem can be parallelized well on all three architectures, although for different ranges of problem sizes. Bokhari concluded that the graphical processing unit (GPU) architecture performs well for problems whose tables fit within the limitations of the device memory. Because GPUs typically have memory sizes in the range of 10 gigabytes (GB), such architectures are best for small problems that have table sizes of approximately 30 billion bits.

She found that the IBM x3755 performed very well on medium-sized problems that fit within its 64-GB memory, but had poor scalability as the number of processors increased and was unable to sustain its performance as the problem size increased beyond 300 billion bits. The Cray XMT showed very good scaling for large problems and demonstrated sustained performance as the problem size increased, she said. However, the Cray had poor scaling for small problem sizes, performing best with table sizes of a trillion bits or more.



Comparison of performance measures of parallel supercomputers that employ various architectures: (1) Cray XMT, 128 proc, 500 MHz, 1-TB; (2) IBM x3755, 2.4-GHz Opteron, 16-core, 64-GB; (3) NVIDIA FX 5800 GPU, 1.296 GHz, 240 cores, 4-GB device memory.

Project lead: Saniyah Bokhari, The Ohio State University
 Research title: Parallel solution of the subset-sum problem: an empirical study
 Funding source: The Ohio State University
 Web site: <http://bit.ly/OSC-RR-Bokhari>