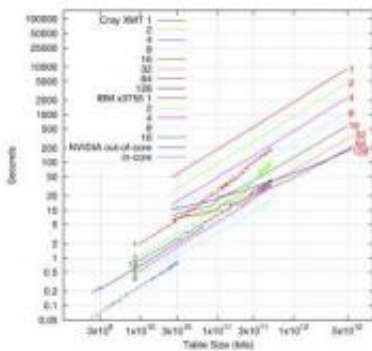


Researcher tests performance of diverse HPC architectures

March 29 2012



Saniyah Bokhari, an Ohio State University graduate student, compared performance measures of parallel supercomputers that employ various architectures, testing systems located at the Ohio Supercomputer Center and Pacific Northwest Laboratory: (1) a Cray XMT, 128 proc., 500 MHz, 1-TB; (2) an IBM x3755, 2.4-GHz Opteron, 16-core, 64-GB; and (3) an NVIDIA FX 5800 GPU, 1.296 GHz, 240 cores, 4-GB device memory. Credit: Bokhari

Surveying the wide range of parallel system architectures offered in the supercomputer market, an Ohio State University researcher recently sought to establish some side-by-side performance comparisons.

The journal, *Concurrency and Computation: Practice and Experience*, in February published, "Parallel solution of the subset-sum problem: an empirical study." The paper is based upon a master's thesis written last

year by computer science and engineering graduate student Saniyah Bokhari.

"We explore 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 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](#) (OSC) and the less-traditional general-purpose graphic processing unit (GPGPU) 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 (PNNL) Center for Adaptive Supercomputing Software.

"Ms. Bokhari's work provides valuable insights into matching the best high performance computing architecture with the computational needs of a given research community," noted Ashok Krishnamurthy, interim co-executive director of OSC. "These systems are continually evolving to incorporate new technologies, such as GPUs, in order to achieve new, higher performance measures, and we must understand exactly what each new innovation offers."

Each of the architectures Bokhari tested fall in the area of parallel computing, where multiple processors are used to tackle pieces of complex problems "in parallel." The subset-sum problem she used for her study is an algorithm with known solutions that is solvable in 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.

The results from Bokhari's study illustrate that the subset-sum problem can be parallelized well on all three architectures, although for different ranges of problem sizes. The performances of these three machines under varying problem sizes showed the strengths and weaknesses of the three architectures.

Bokhari concluded that the GPU 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 thirty 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. The machine tended to saturate for problem with table sizes of 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.

"In conclusion, we can state that the NVIDIA GPGPU is best suited to small problem sizes; the IBM x3755 performs well for medium sizes, and the Cray XMT is the clear choice for large problems," Bokhari said. "For the XMT, we expect to see better performance for large problem sizes, should memory larger than 1 TB become available."

Provided by Ohio Supercomputer Center

Citation: Researcher tests performance of diverse HPC architectures (2012, March 29) retrieved 19 April 2024 from <https://phys.org/news/2012-03-diverse-hpc-architectures.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.