April 2012

Center for Adaptive Supercomputing Software

http://cass-mt.pnnl.gov

CASS Research Areas

Architecture Studies

- Next generation hardware designs
- Next generation hybrid systems
- Software Multithreading on next generation systems

System Software

- Compilers
- Runtime systems
- Programming tools
- Communication libraries

Algorithms

- Semantic databases
- Structural analysis
- Social networks
- Bayesian networks
- Natural language comprehension
- Cybersecurity
- Clustering
- Real-time methods
- Dynamic data structures

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

MODA

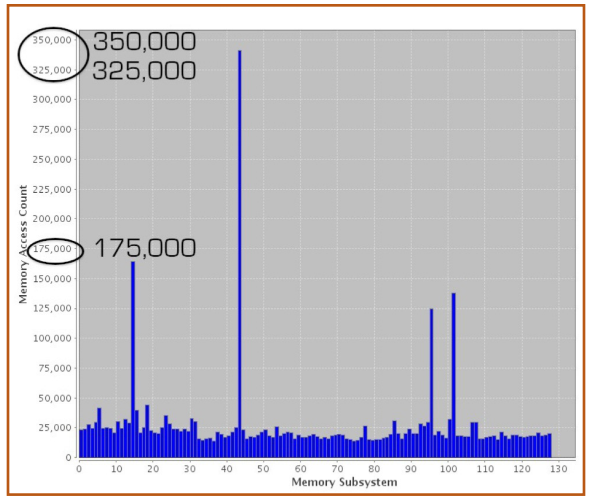

Graph 500 memory reference pattern by 4k threads on 128 XMT nodes showcasing anomalous behavior

The multi-core revolution aims to achieve performance gains from one generation to the next by exploiting ever more thread- and data-parallelism available in current applications. Contrary to the performance gains achieved by super-scalar design improvements in the past, multi-core systems require changes in programming models, system software and software tools, in order to harness the new found levels of parallelism. Tools, in particular, are required to automatically identify opportunities for parallelism, from the compilation phase up to the post-execution analysis phase. Current performance analysis tools tend to be control centered, a legacy that has its origin in the previous generation of architectures. These tools are good in pinpointing threads and code regions that are culprits of performance bottlenecks but provide little detail on the resources involved in the program execution. For example, current performance analysis tools can determine if a code section is memory bound by measuring CPU utilization and memory references over a specified time interval. These tools will attribute the performance bottlenecks to a set of threads but will not provide further details on the memory subsystem beyond caches.

This type of analysis was not necessarily a problem for previous architectures where resources like memory, network and I/O components were tightly coupled to a single or few cores in a bijective fashion. Identifying a resource problem, as in our previous example, automatically meant that the culprit could be identified as well. Yet with the advent of multi-core systems, to achieve some level of system balance, resources needed to be replicated as well. In the case of memory, that means multiple memory subsystems attached over multiple channels to a pool of cores and(or) processors. In addition, multicore systems created an execution paradigm shift from being compute bound to memory, bandwidth and I/O bound. Control centric analysis in such cases can identify a memory bound problem, but won't provide clues about the root causes.

Here we showcase a memory centric performance tool called the Memory Observant and Data Analysis Framework, or MODA for short. It is designed to reveal existing and potential algorithmic and architectural resource hot-spots by means of a sophisticated memory model. The tool helps to identify performance degradation factors at a small scale where debugging and performance analysis is more manageable. Salient features of this tool include (1) a memory trace collection with minimal perturbation of the application's behavior; (2) data management of multiple Giga and Tera byte size trace files; (3) efficient data analysis and presentation of traces; and (4) the introduction of the target architecture's memory model into the analysis module for a truly memory centric view.

UPCOMING EVENTS

Sinan al-Saffar

Oreste Villa

» Sinan al-Saffar, John Feo, and Oreste Villa will participate in a workshop "HPC/Big Graph Data" at the A*STAR Computational Resource Centre (A*CRC) in Singapore on April 25-27. The purpose of the workshop is to look at solutions to big data challenges through high performance computing. Villa will focus on supercomputer architectures and programming; Feo will focus on shared memory programming; and al-Saffar will focus on data intensive and semantic computing. A*CRC provides high performance computational (HPC) resources to the entire A*STAR research community. Currently A*CRC supports HPC needs of over seven hundred strong user community and manages several high-end computers. It is also responsible for very rapidly growing data storage resources.

John Feo

» John Feo and Oreste Villa will give the workshop "Big Data and Graph Analysis on the new Cray uRiKA" at the Swiss National Supercomputing Centre in Lugano, Switzerland on May 11-12. The aim of the workshop is to bring together potential users of the Cray uRiKa (XMT) system to become familiar with its capabilities, both for those who plan to develop code for the system as well as those who intend to use the machine to analyze big datasets.

RECENT EVENTS

Mahantesh Halappanavar

» Mahantesh Halappanavar, Ariful Azad, Umit Catalyurek, and Alex Pothen presented the paper "Parallel Algorithms for Matching and Coloring" at the Society for Industrial and Applied Mathematics (SIAM) Conference on Parallel Processing for Scientific Computing in Savannah in February. This paper compared the performance of the Aho-Corasick string matching algorithms across a variety of parallel computing systems including manycore and multicore systems, shared-memory servers, multithreaded systems, and GPUs.

Antonino Tumeo

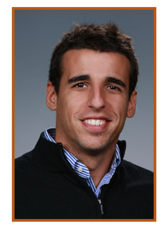

Simone Secchi

» Also at the SIAM Conference in February, Antonino Tumeo, Oreste Villa, and Simone Secchi presented the paper "Exploring Architectural Features for Supporting Parallel Graph Algorithms." The paper presented five multithreaded algorithms for computing maximum matchings in bipartite graphs and their evaluation on multithreaded and multicore platforms. It also presented two approximation algorithms for greedy initializations.

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

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*