Benchmarking Computing Clusters Using High Performance Linpack

Cindy Lopez
City University of New York-York College

Research Alliance in Math and Science (RAMS)
Computer Science and Mathematics Division
Mentors: John Mugler, Thomas Naughton, Stephen Scott
http://www.csm.ornl.gov/Internships/abstracts/CindyLopez.pdf

Abstract

High performance computing clusters are a significant part of modern scientific computing. At the Oak Ridge National Laboratory (ORNL), two such systems exist for the advancement of science. The first is a 64-node Pentium IV cluster named XTORC. The second is Crane, which is a 64-node Opteron cluster. The goal for the summer project is to learn about benchmarking and cluster computing systems, and to benchmark XTORC and Crane using the TOP500 benchmark suite. The TOP500 list is a project started to provide a reliable basis for tracking and detecting trends in high-performance computing. Since these two systems have different interconnects and different processors, it is important to understand the types of applications that will perform well on each. To benchmark these systems we chose to use the High Performance Linpack (HPL). HPL is one of the standard tools used for benchmarking clusters that solves a random dense linear system using double precision. HPL uses Basic Linear Algebra Subprograms (BLAS) that are high quality "building block" routines for performing basic vector and matrix operations. Automatically Tuned Linear Algebra Software (ATLAS) offers the ability to custom tune the BLAS subroutines to a particular system. Using ATLAS enables the HPL benchmark to achieve the highest performance on a given platform. HPL relies on the Message Passing Interface (MPI) to perform any communication across the cluster. Results obtained may be used as a standard for those who are competing to get on to or to remain on the TOP500 list.

Objective

To compare benchmark results between XTORC and Crane

Formulas

(MIPS) (CPU's) = Theoretical Peak

Highest GFLOPS/Theoretical Peak = Efficiency Rate

Efficiency Rates

37/169 = 21.8 % (XTORC)

208/319 = 65.2 % (Crane)

Benchmark Results

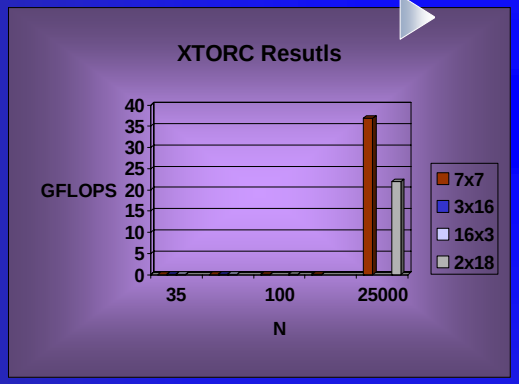

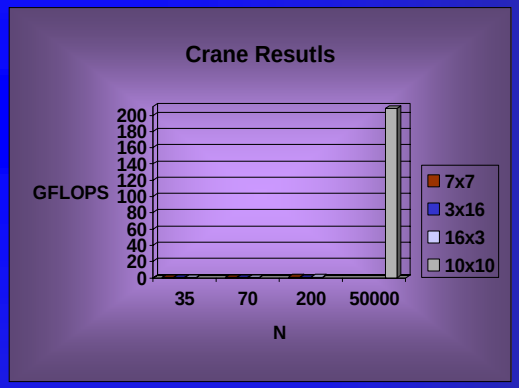

Results based on highest GFLOPS per given test

XTORC (64-node)

Pentium IV (single processor)
100/1000 Mbps Ethernet
Memory: 768KB
L2/L3 Cache: 265 KB
Speed: 1695.87 MHz

Crane (64-node)

AMD Opteron (dual processor 240)
1000 Mbps Ethernet
Elon-3 Quadrics
Memory: 2 GB
L2/L3 Cache: 1024 KB
Speed: 1395.67 MHz

Software

High Performance Linpack (HPL)

Basic Linear Algebra Subprograms (BLAS)

Automatically Tuned Linear Algebra Software (ATLAS)

Message Passing Interface (MPI)

Conclusion

Parameters used for tuning HPL make a huge difference in results.

Crane results are significantly higher than XTORC.

The Research Alliance in Math and Science program is sponsored by the Mathematical, Information, and Computational Sciences Division, Office of Advanced Scientific Computing Research, U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725. This work has been authored by a contractor of the U.S. Government, accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

This research used resources of the Center for Computational Sciences at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

OAK RIDGE NATIONAL LABORATORY
U.S. DEPARTMENT OF ENERGY

Acknowledgements: I would like to thank, first and foremost, Dr. Michelangelo Salcedo for the encouragement and drive to apply for the internship. Of course, I'd like to thank MICS along with Debbie McCoy for giving me the opportunity to intern this summer here at Oak Ridge National Laboratory. I would also like to thank ORNL for giving me the opportunity to conduct research using their facility. Next, I would like to thank my mentors John Mugler, Thomas Naughton and Stephen Scott for all their support, help and encouragement in finalizing this summer research project. Finally, I would like to thank Cheryl Hamby, Cynthia Latham and all the supportive and helpful staff here at ORNL.