Sandia National Laboratories

Measurement of Time Instability (Heterogeneity) of Parallel Computer and Its Impact on Parallel Algorithm Scalability

Principal Investigators: Robert Benner (SNL), A.M. Vargin, and Yuri Bartenev (VNIIEF)

Project Description

Timing instabilities diminish the efficiency of supercomputers. Timing instabilities of interest include unequal central processing unit speeds, unequal memory speeds, operating system noise, and message contention on interprocessor interconnects. This project seeks to characterize these timing instabilities though a measurement program. The project will develop a parallel software program to measure and report timing instabilities on a parallel computer, including validation that the program's results are accurate, release of the program as general tool with "open source" license, and publication of a scientific paper.

Six tasks will be performed at the discretion of the Sandia principal investigator. The first three tasks have been completed as of December 2006. The remaining tasks are scheduled for completion by September 2007, with successive tasks being completed at 3 or 4 month intervals:

- Development of proposals, project specifications, and prototyping of activities.
- Develop measurers of noise on a single cluster node.
- Develop the measurers of noise on all cluster noted and regular scalable communication operation instability.
- Develop the measurers of cluster regular communication instabilities on algorithms possessing weak scalability degradation, including testing and reduction of tests to a complex with limited shell possibilities and verification on parallel clusters.
- Develop measurers of regular communication instabilities in a part of a cluster in a random communication system load with communication and I/O operations on the side of the other part of the cluster, including testing and reduction of tests to a complex.
- Develop full-value shell and perform complex testing of project. Verify measurer on actual parallel clusters and develop follow-on proposals.

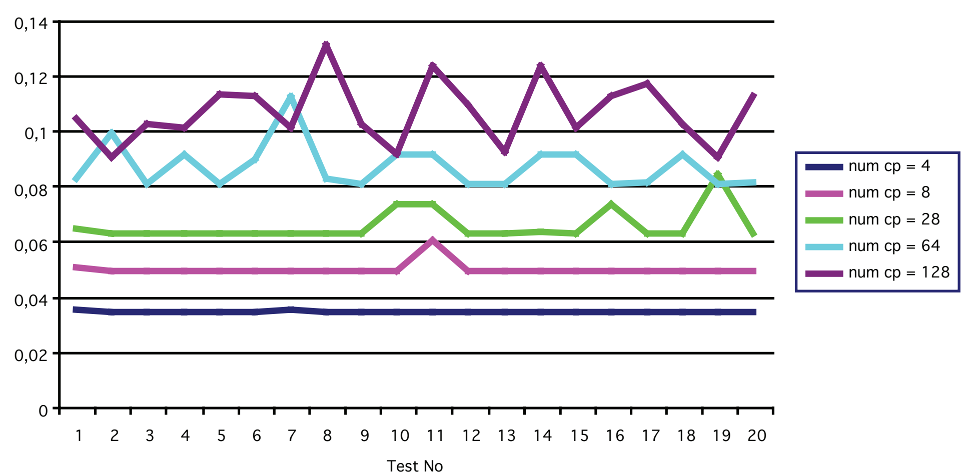

Runtime of solver Bicgstab with preconditioner block Gauss-Seidel.
Approximation template = 7. Number iterations = 1.
Local size of matrix = 1000 rows. Number of calls = 40.

Technical Purpose and Benefits

As a result of this work, the measurement program will run loops on all processors of a supercomputer to identify small timing differences. Then the measurement program will process the data into parameters for a model of timing instabilities. The program can be used in purchasing computers to select those with better timing stability. It can be used as a diagnostic tool for improving hardware and software, and it can help to assess the feasibility of increasing the size of a supercomputer or parallel applications. In particular, the measurement and characterization of operating system noise would lead to improvements in the runtime performance and scalability of these systems as well as future systems and the applications that run on them.

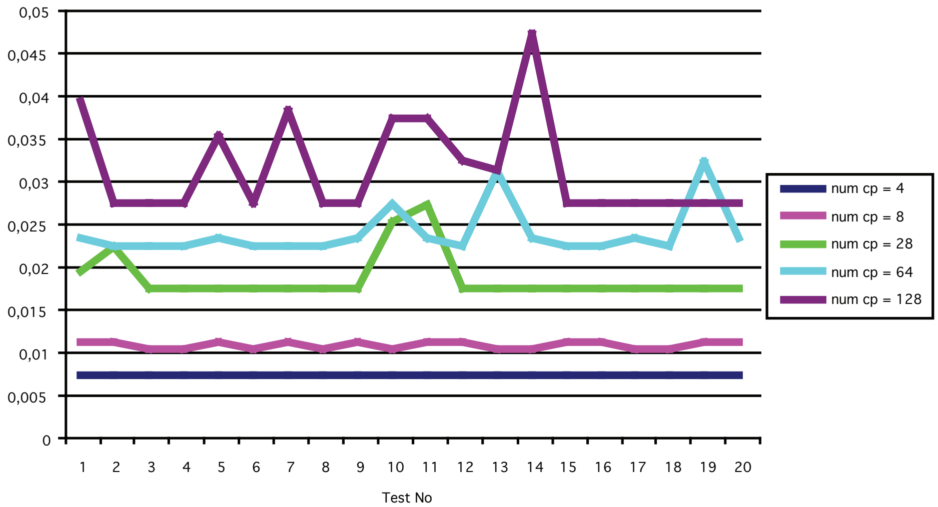

Runtime of function Matvec. Approximation template = 27. Matrix size = 1000 rows. Number of calls = 40.

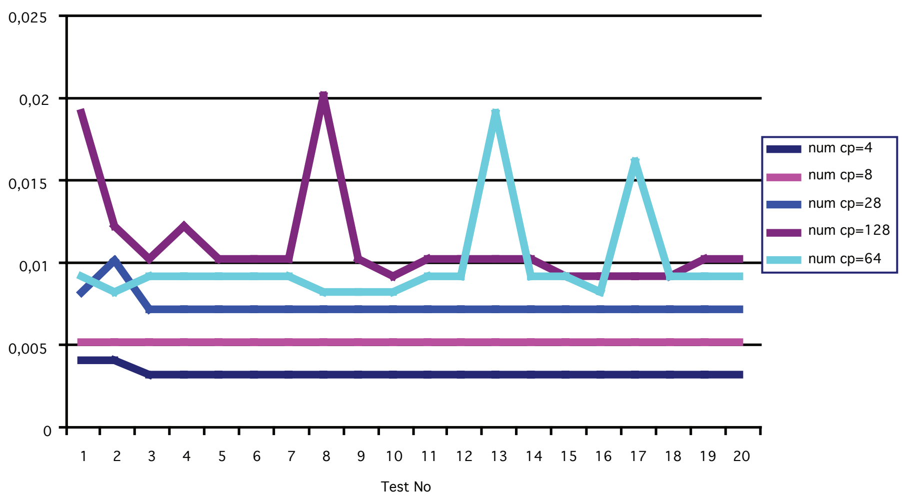

Runtime of function Matvec. Approximation template = 7. Matrix size = 1000 rows. Number of calls = 40.

Collaboration between Sandia National Laboratories (SNL), Livermore, CA, USA, and the Russian Federal Nuclear Center – All Russian Research Institute of Experimental Physics (RFNC-VNIIEF), Sarov, Russia

Sandia National Laboratories

A Publication of the Office of Advanced Simulation & Computing, NNSA Defense Programs. SAND 2007-4400P
Issued by Sandia National Laboratories for NNSA's Office of Advanced Simulation & Computing, NA-114. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.