UCRL-CR-142047
B508995

Final Report for "Automated Diagnosis of Large Scale Parallel Applications"

K. L. Karavanic

November 17, 2000

U.S. Department of Energy
Lawrence Livermore National Laboratory

DISCLAIMER

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

Portland State University
P.O. Box 751
Portland, OR 97207-0751

Final Report for "Automated Diagnosis of Large Scale Parallel Applications"

Submitted by:
Karen L. Karavanic
Assistant Professor, Department of Computer Science

FINAL REPORT
For the period ending 9/30/00

Prepared for:
University of California
Lawrence Livermore National Laboratory
Attn: Jill Dunaway
P.O. Box 808, L-561
Livermore, CA 94551

Under
B508995

Date prepared
11/17/00

Final Report on

LLNL Subcontract B508995

Automated Diagnosis of Large-Scale Parallel Applications

Karen L. Karavanic

1. Project Overview

The work performed is part of a continuing research project, PPerfDB, headed by Dr. Karavanic [1, 2, 3]. We are studying the application of experiment management techniques to the problems associated with gathering, storing, and using performance data with the goal of achieving completely automated diagnosis of application and system bottlenecks. This summer we focused on incorporating heterogeneous data from a variety of tools, applications, and platforms, and on designing novel techniques for automated performance diagnosis.

Three students participated in the project over the summer: Christian Hansen and Reena John are in the PSU CS Department masters program, and Andrew Stoneman was an area high school student who is now a freshman at the University of Washington.

2. Papers and Book Chapters Supported in Part by the Subcontract

Christian Hansen and Karen L. Karavanic, "Integrating Performance Data from Heterogeneous Sources" in progress, target submission date, January 2001.

3. Project Highlights

The first task includes representation and API for use and storage of performance data from a variety of tools, applications, and collection methods (tracing or profiling, online or offline, etc.). Eventually, we will accommodate data gathered and used in both correctness and performance debugging. We investigated currently available parallel performance tools: Vampir, Jumpshot, Paradyn, Pablo, AIMS, VT, Scala, and Guideview. We implemented extensions for the PPerfDB prototype to incorporate resource descriptions and performance data gathered with VT, Vampir, and Jumpshot on the LLNL ASCI Blue and Compass platforms and a locally available cluster of SPARC Solaris workstations. This was the first inclusion of data from trace-based tools into the experiment management prototype.

The second task is developing new techniques for the diagnosis itself. Project member Reena John investigated the Paradyn performance consultant in the context of historical data and heterogeneous resource hierarchies. We began designing a diagnostic framework that will incorporate automatically determined thresholds. This expands upon earlier work in which we developed a new automated diagnosis technique to incorporate

prior knowledge. We look forward to implementation and analysis of these ideas, which wasn't completed in the limited time period available over the summer.

Our work included a LLNL site visit in September, during which project member Christian Hansen presented a description of our work in progress. A copy of the slides from his talk is also included.

4. Concluding Remarks

The Experiment Management paradigm is a useful approach for designing a tool that will automatically diagnose performance problems in large-scale parallel applications. The ability to gather, store, and use performance data gathered over time from different executions and using different collection tools enables more sophisticated approaches to performance diagnosis and to performance evaluation more generally. We look forward to continuing our efforts by further development and analysis of online diagnosis using historical data, and by investigating performance data and diagnosis gathered from mixed MPI/OpenMP applications.

5. References

[1] Karen L. Karavanic, "Experiment Management Support for Parallel Performance Tuning" Ph.D. Dissertation, University of Wisconsin, 1999.

[2] Karen L. Karavanic and Barton P. Miller "Improving Online Performance Diagnosis by the use of Historical Performance Data." SC'99, Portland, OR.

[3] Karen L. Karavanic and Barton P. Miller, "Experiment Management Support for Performance Tuning", SC '97, San Jose, CA.