

Developing a Tuned Version of ScaLAPACK's Linear Equation Solver

J. Dongarra

October 29, 2000

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

**University of Tennessee
1122 Volunteer Blvd
Knoxville TN, 37996-3450**

**Final Report for
Developing a Tuned Version of ScaLAPACK's Linear
Equation Solver**

Submitted by:
Jack Dongarra
University Distinguished Professor

FINAL REPORT
January 1, 2000

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

Under
B503913

October 29, 2000

Final Report on
LLNL Subcontract B503913
Developing a Tuned Version of ScaLAPACK's Linear Equation Solver

Jack Dongarra

1. Project Overview

The LINPACK Benchmark has been used as a yardstick in measuring the performance of the Top500 installed high-end computers. This benchmark was chosen because it is widely used and performance numbers are available for almost all relevant systems. The approach used in the LINPACK Benchmark is to solve a dense system of linear equations. For the Top500, the benchmark allows the user to scale the size of the problem and to optimize the software in order to achieve the best performance for a given machine. This performance does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations. Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good check of peak performance of a system.

By measuring the actual performance for different problem sizes n , a user can get not only the maximal achieved performance R_{\max} for the problem size N_{\max} but also the problem size $N_{1/2}$ where half of the performance R_{\max} is achieved. These numbers together with the theoretical peak performance R_{peak} are the numbers given in the Top500. In an attempt to obtain uniformity across all computers in performance reporting, the algorithm used in solving the system of equations must conform to the standard operation count for LU factorization with partial pivoting. In particular, the operation count for the algorithm must be $2/3 n^3 + O(n^2)$ floating point operations.

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

None

3. Project Highlight

As part of this proposal we developed a version of the Benchmark based on the hardware of the ASCI BLUE Pacific system to achieve high performance. It was our goal to produce the fastest implementation that will take advantage of the hardware and software infrastructure on the BLUE Pacific machine in achieving this mark. In order to accomplish this goal, our team drew on the expertise in this field developed with our work on the ScaLAPACK and ATLAS projects. We collaborated with researchers at LLNL to produce software for the Linpack benchmark that is fully optimized for the ASCI-Blue Pacific system.

The ATLAS project involves using timings coupled with code generation to automatically adapt linear algebra computations to run optimally on varying

architectures. The expertise developed in this project will be leveraged in order to make machine-specific performance tweaks for the computations involved in the benchmark. In particular, tunings for various levels of caches, floating point register and unit usage, and optimal use of threading. We presented a tutorial at LLNL on the ATLAS optimization implementation and described ways to optimize numerical software in general.

4. Concluding Remarks

Some of this work is on going with LLNL researchers.

5. References

See the Top 500 list for the results of this work. <http://www.top500.org/>