HIGH-PERFORMANCE LINPACK BENCHMARK

HPL

HPL is a portable implementation of the High Performance LINPACK Benchmark for distributed memory computers.

- Algorithm: recursive panel factorization, multiple lookahead depths, bandwidth reducing swapping
- Easy to install, only needs MPI+BLAS or VSIPL
- Highly scalable and efficient from the smallest cluster to the largest supercomputers in the world

Compaq 64 Node AlphaServer SC (4 EV67 667 MHz CPUs per node); constant memory load/CPU = 335 MiB

CPU/Nodes	N ½	N max	R max (Gflop/s)
1/1	150	6625	1
4/1	800	12350	4
16/4	2300	26500	17
64/16	5700	53000	67.5
256/64	14000	106000	263.6

PARALLEL EFFICIENCY OF HPL

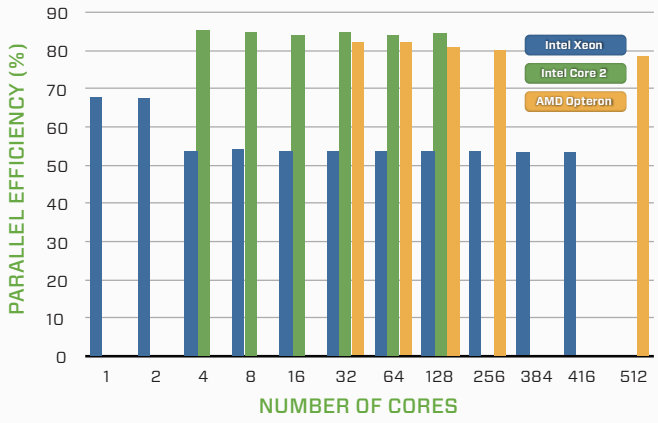

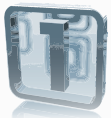

HPL was used to obtain a number of results in the current TOP500 list including the #1 entry.

HISTORY OF THE BENCHMARK

1974 **LINPACK software is released**
Solves systems of linear equations in FORTRAN 66

1977 **LINPACK 100 released**
Measures system performance in Mflop/s and solves 100x100 linear systems

1986 **LINPACK 1000 released**
Any language allowed and the linear system of size 1000 can be used

1989 **LINPACKDv2 released**
Extends random number generator from 16384 to 65536

1991 **LINPACK Table 3 (Highly Parallel Computing)**
Any size linear system is allowed

1993 **TOP500 first released**
With CM-5 running the LINPACK benchmark at nearly 60 Gflop/s

1996 **9th TOP500 is released**
With the 1st system breaking the 1 Tflop/s barrier: ASCI Red from Sandia National Laboratory

2000 **HPLv1 is released**
By Antoine Petitet, Jack Dongarra, Clint Whaley, and Andy Cleary

2008 **31st TOP500 is released**
With the 1st system breaking the 1 Pflop/s barrier: Roadrunner from Los Alamos National Laboratory

2008 **HPLv2 is released**
The new version features a 64-bit random number generator that prevents the benchmark failures from generating nearly singular matrices which was the problem with the old generator.

2009 **Peta flop/s are spreading**
The upgrades of the machines hosted at ORNL results in shattering the 2 Pflop/s theoretical peak barrier for Cray's XT5 Jaguar. Also, the first academic system reaches 1 Pflop/s theoretical peak: University of Tennessee's Kraken.

2010 **GPUs are coming**
The performance growth at Peta-scale is now fueled by a new breed of GPUs that now power systems with over 2 Pflop/s in LINPACK performance.

A COLLABORATION OF THE UNIVERSITY of TENNESSEE · OAK RIDGE National Laboratory

SPONSORED BY

CITR CENTER FOR INFORMATION TECHNOLOGY RESEARCH

FIND OUT MORE AT http://icl.eecs.utk.edu/hpl/

INNOVATIVE COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE

HIGH-PERFORMANCE LINPACK BENCHMARK

HPL

The original LINPACK Benchmark is, in some sense, an accident. It was originally designed to assist users of the LINPACK package by providing information on execution times required to solve a system of linear equations. The first "LINPACK Benchmark" report appeared as an appendix in the LINPACK Users' Guide in 1979. Over the years additional performance data was added, more as a hobby than anything else, and today the collection includes over 1300 different computer systems. In addition to the number of computers increasing, the scope of the benchmark has also expanded. The benchmark report describes the performance for solving a general dense matrix problem Ax=b at three levels of problem size and optimization opportunity: 100 by 100 problem (inner loop optimization), 1000 by 1000 problem (three loop optimization – the whole program), and a scalable parallel problem – the HPL. New statistics were required to reflect the diversification of supercomputers, the enormous performance difference between low-end and high-end models, the increasing availability of massively parallel processing (MPP) systems, and the strong increase in computing power of the high-end models of workstation suppliers (SMP). To provide this new statistical foundation, the TOP500 list was created in 1993 to assemble and maintain a list of the 500 most powerful computer systems. The list is compiled twice a year with the help of high-performance computer experts, computational scientists, manufacturers, and the Internet community in general. In the list, computers are ranked by their performance on the HPL benchmark. With the unprecedented increase in performance comes the price of increased time the benchmark takes to complete. As the chart below shows, the time to completion is increasing with the number of cores. And in June 2009, the #2 entry reported a run that took over 18 hours.

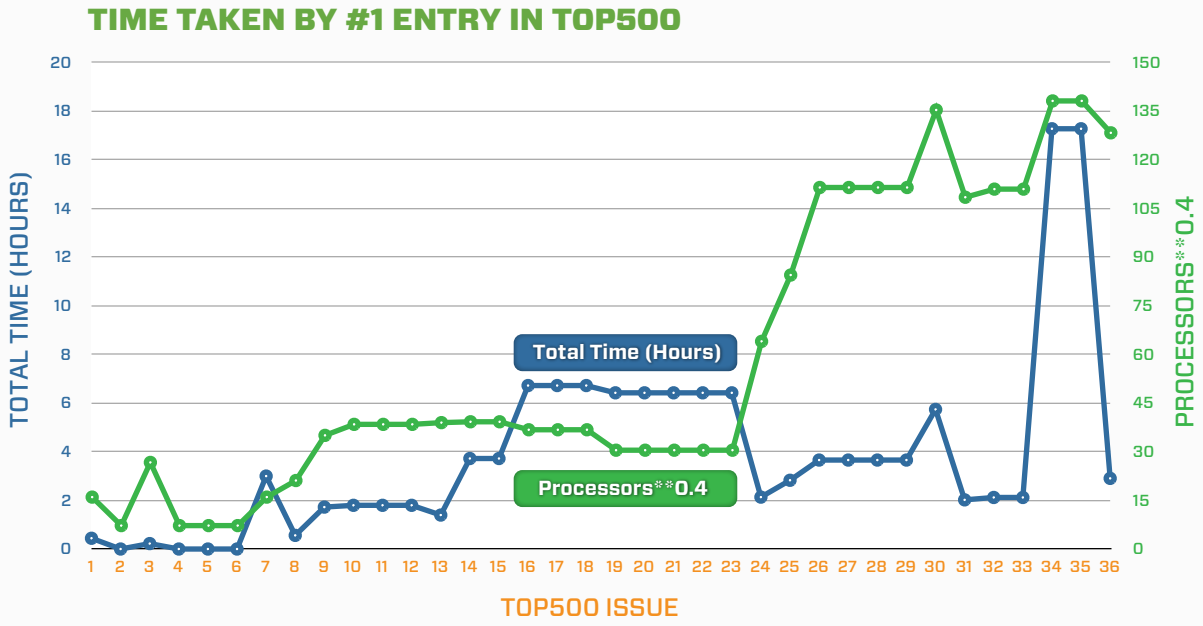

To deal with this problem a new version of HPL will offer a partial run option. This will allow to run the benchmark only a fraction of the time it takes to run the full version of the benchmark without losing much of the reported performance. Initial runs on a system with over 30000 cores indicate only a 6% drop in performance with a 50% reduction in time. This becomes a viable alternative to, say, checkpointing to solve the problem of the the rapidly increasing number of cores and the decrease of Mean Time Between Failures (MTBF) for large supercomputer installations.

A COLLABORATION OF THE UNIVERSITY of TENNESSEE · OAK RIDGE National Laboratory

SPONSORED BY

CITR CENTER FOR INFORMATION TECHNOLOGY RESEARCH

FIND OUT MORE AT http://icl.eecs.utk.edu/hpl/

INNOVATIVE COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE