FIND OUT MORE AT https://icl.utk.edu/hpl

HPL

HIGH PERFORMANCE LINPACK BENCHMARK

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

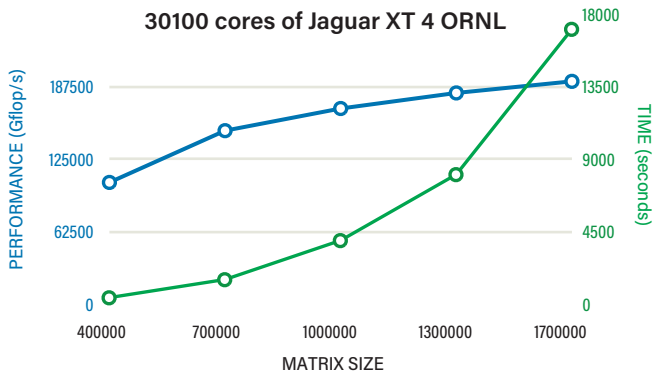

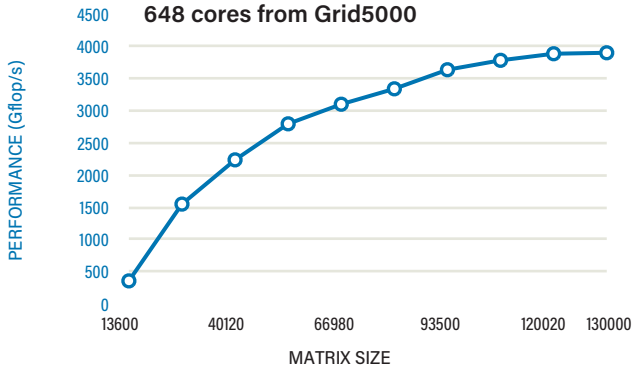

TOP 500 The List.

HPL was used to obtain a number of results in the current TOP500 list, including some of the entries at the top of the list.

SPONSORED BY

U.S. DEPARTMENT OF ENERGY

IN COLLABORATION WITH

OAK RIDGE National Laboratory

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

HPLv2 is released: The new version features a 64-bit random number generator that prevents the benchmark failures from generating 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 petascale is now fueled by a new breed of GPUs that now power systems with over 2 Pflop/s in LINPACK performance

2011 Multicore strikes back by exceeding 8 Pflop/s using over half a million of modified SPARC VIII cores

2012 Even more cores needed to be #1: a million and a half cores were required to break 15 Pflop/s barrier

2013 Intel back on top with over half of 100 Pflop/s of peak performance coming mostly from its Intel Xeon Phi accelerators

2014 Slow down in top ten of the TOP500 continues while the accelerators contribute increasing fraction of total Pflop/s

2016 New barriers broken by TaihuLight, a Chinese-built supercomputer, that featured over 10 million cores and gets close to 100 Pflop/s performance

2018 US supercomputer located at ORNL nearly reached 200 Pflop/s peak performance mark in double precision and over 1 Eflop/s in half precision

2019 Exascale machines were announced as part of the United States' CORAL2 program. These machines include Aurora (ANL), El Capitan (LLNL), and Frontier (ORNL). All three will be based on Cray's Shasta design and will feature hardware accelerators from AMD and Intel.

2020 Confirmed press releases announced exascale machines installed in China without an official TOP500 entry

2022 Exascale was finally reached: Frontier supercomputer at ORNL achieved over 1 Eflop/s while running HPL

2023 Frontier improved its performance by nearly 10% which is more than the Rmax performance of Summit housed in the same building at ORNL

2024 Two systems on TOP500 crossed the 1 Eflop/s mark, both located in the USA

2025 Three US-based systems now passed the 1 Eflop/s score on TOP500 and European HPC initiative prepares JUPITER to do the same in Germany