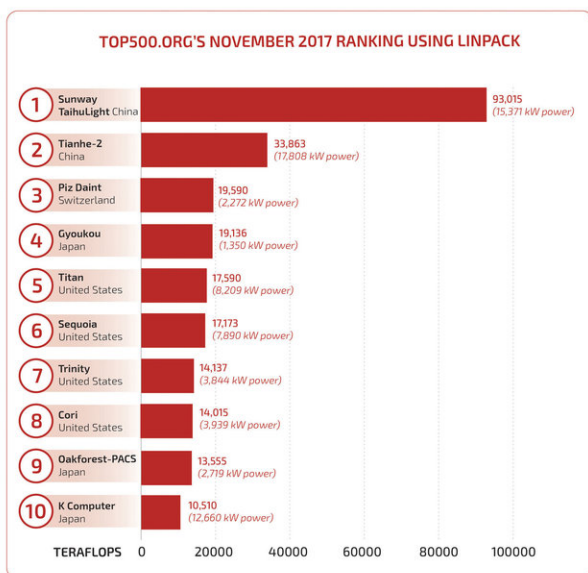


# Supercomputing under a new lens: A Sandia-developed benchmark re-ranks top computers

February 27 2018



TOP500 LINPACK and HPCG charts of the fastest supercomputers of 2017. The rearranged order and drastic reduction in estimated speed for the HPCG benchmarks are the result of a different method of testing modern supercomputer programs. Credit: Sandia National Laboratories

A Sandia National Laboratories software program now installed as an additional test for the widely observed TOP500 supercomputer challenge

has become increasingly prominent. The program's full name—High Performance Conjugate Gradients, or HPCG—doesn't come trippingly to the tongue, but word is seeping out that this relatively new benchmarking program is becoming as valuable as its venerable partner—the High Performance LINPACK program—which some say has become less than satisfactory in measuring many of today's computational challenges.

"The LINPACK [program](#) used to represent a broad spectrum of the core computations that needed to be performed, but things have changed," said Sandia researcher Mike Heroux, who created and developed the HPCG program. "The LINPACK program performs compute-rich algorithms on dense data structures to identify the theoretical maximum speed of a supercomputer. Today's applications often use sparse data structures, and computations are leaner."

The term "sparse" means that a matrix under consideration has mostly zero values. "The world is really sparse at large sizes," said Heroux. "Think about your social media connections: there may be millions of people represented in a matrix, but your row—the people who influence you—are few. So, the effective matrix is sparse. Do other people on the planet still influence you? Yes, but through people close to you."

Similarly, for a scientific problem whose solution requires billions of equations, most of the matrix coefficients are zero. For example, when measuring pressure differentials in a 3-D mesh, the pressure on each node is directly dependent on its neighbors' pressures. The pressure in faraway places is represented through the node's near neighbors. "The cost of storing all matrix terms, as the LINPACK program does, becomes prohibitive, and the computational cost even more so," said Heroux. A computer may be very fast in computing with dense matrices, and thus score highly on the LINPACK test, but in practical terms the HPCG test is more realistic.

To better reflect the practical elements of current supercomputing application programs, Heroux developed HPCG's preconditioned iterative method for solving systems containing billions of linear equations and billions of unknowns. "Iterative" means the program starts with an initial guess to the solution, and then computes a sequence of improved answers. Preconditioning uses other properties of the problem to quickly converge to an acceptably close answer.

"To solve the problems we need to for our mission, which might range from a full weapons simulation to a wind farm, we need to describe physical phenomena to high fidelity, such as the pressure differential of a fluid flow simulation," said Heroux. "For a mesh in a 3-D domain, you need to know at each node on the grid the relations to values at all the other nodes. A preconditioner makes the iterative method converge more quickly, so a multigrid preconditioner is applied to the method at each iteration."

Supercomputer vendors like NVIDIA Corp., Fujitsu Ltd., IBM, Intel Corp. and Chinese companies write versions of HPCG's program that are optimal for their platform. While it might seem odd for students to modify a test to suit themselves, it's clearly desirable for supercomputers of various designs to personalize the test, as long as each competitor touches all the agreed-upon calculation bases.

"We have checks in the code to detect optimizations that are not permitted under published benchmark policy," said Heroux.

On the HPCG TOP500 list, the Sandia and Los Alamos National Laboratory supercomputer Trinity has risen to No. 3, and is the top Department of Energy system. Trinity is No. 7 overall in the LINPACK ranking. HPCG better reflects the Trinity design choices.

Heroux says he wrote the base HPCG code 15 years ago, originally as a

teaching code for students and colleagues who wanted to learn the anatomy of an application that uses scalable sparse solvers. Jack Dongarra and Piotr Luszczek of the University of Tennessee have been essential collaborators on the HPCG project. In particular, Dongarra, whose visibility in the high-performance computing community is unrivaled, has been a strong promoter of HPCG.

"His promotional contributions are essential," said Heroux. "People respect Jack's knowledge and it helped immensely in spreading the word. But if the program wasn't solid, promotion alone wouldn't be enough."

Heroux invested his time in developing HPCG because he had a strong desire to better assure the U.S. stockpile's safety and effectiveness. The supercomputing community needed a new benchmark that better reflected the needs of the national security scientific computing community.

"I had worked at Cray Inc. for 10 years before joining Sandia in '98," he says, "when I saw the algorithmic work I cared about moving to the labs for the Accelerated Strategic Computing Initiative (ASCI). When the US decided to observe the Comprehensive Nuclear Test Ban Treaty, we needed high-end computing to better ensure the nuclear stockpile's safety and effectiveness. I thought it was a noble thing, that I would be happy to be part of it, and that my expertise could be applied to develop next-generation simulation capabilities. ASCI was the big new project in the late 1990s if I wanted to do something meaningful in my area of research and development."

Heroux is now director of software technology for the Department of Energy's Exascale Computing Project. There, he works to harmonize the computing work of the DOE national labs—Oak Ridge, Argonne, Lawrence Berkeley, Pacific Northwest, Brookhaven and Fermi, along with the three National Nuclear Security Administration labs.

"Today, we have an opportunity to create an integrated effort among the national labs," said Heroux. "We now have daily forums at the project level, and the people I work with most closely are people from the other labs. Because the Exascale Computing Project is integrated, we have to deliver software to the applications and the hardware at all labs. The Department of Energy's attempt at a multi-lab, multi-university project gives an organizational structure for us to work together as a cohesive unit so that software is delivered to fit the key applications."

Provided by Sandia National Laboratories

Citation: Supercomputing under a new lens: A Sandia-developed benchmark re-ranks top computers (2018, February 27) retrieved 26 April 2024 from <https://phys.org/news/2018-02-supercomputing-lens-sandia-developed-benchmark-re-ranks.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.