

Improved ranking test for supercomputers to be released by Sandia

November 18 2013

Sandia National Laboratories researcher Mike Heroux has helped craft a new benchmark that more accurately measures the power of modern supercomputers for scientific and engineering applications. Heroux collaborated with the creator of the widely used LINPACK benchmark, Jack Dongarra, and his colleagues at the University of Tennessee and Oak Ridge National Laboratory.

The LINPACK TOP 500 test, devised by Dongarra's team, for decades has certified which new machines were among the 500 fastest in the world.

The new <u>benchmark</u>—a relatively small program called a High Performance Conjugate Gradient (HPCG)—is undergoing field tests on several National Nuclear Security Administration (NNSA) supercomputers. It will be formally released this month in Denver at SC13, an annual supercomputing conference, at the TOP500 Birds-of-a-Feather meeting, from 7:30-9:30 p.m. Tuesday at the Convention Center's Mile-High Ballroom.

"Supercomputer designers have known for quite a few years that LINPACK was not a good performance proxy for many complex modern applications," said Heroux.

NNSA's Office of Advanced Simulation and Computing (ASC) is funding the HPCG work because it wants a more meaningful measure that reflects how increasingly complex scientific and engineering codes



would perform on upcoming supercomputing architectures, said Heroux.

LINPACK's influential semi-annual TOP500 listing of the 500 fastest machines has been the global benchmark for more than 25 years, initially because even nonexperts considered it a simple and accurate metric.

"The TOP500 was and continues to be the best advertising supercomputing gets," Heroux said. "Twice a year, when the new rankings come out, we get articles in media around the world. My 6-yearold can appreciate what it means."

In the early years of supercomputing, applications and problems were simpler and better matched the algorithms and data structures used in the LINPACK benchmark. Since then, applications and problems have become much more complex, demanding a broader collection of capabilities from computer systems. Thus, the gap between LINPACK performance and performance in real applications has grown dramatically in recent years.

"LINPACK specifications are like telling race car designers to build the fastest car for a completely flat, open terrain," Heroux said. "In that setting, the car has to satisfy only a single design goal. It does not need brakes, a steering wheel or other control features, making it impractical for real driving situations. It still gets there when needed, but is it the best for today?"

Additionally, computer designers have built systems with many arithmetic units but very weak data networks and primitive execution models.

"The extra arithmetic units are useless," Heroux said, "because modern applications cannot use them without better access to data and more



flexible execution models."

He developed the new benchmark by starting with a teaching code he wrote to instruct students and junior staff members on developing parallel applications. This code later became the first "miniapp" in Mantevo, a project that recently won a 2013 R&D 100 Award.

The technical challenge of HPCG is to develop a very small program that captures as much of the essential performance of a large application as possible without making it too complicated. "We created a program with only 4,000 lines that behaves a lot like a real code of 1 million lines but is much simpler," Heroux said. "If we run HPCG on a simulator or new system and modify the code or computer design so that the code runs faster, we can make the same changes to make the real code run faster. The beauty of the approach is that it really works."

HPCG generates a large collection of algebraic equations that must be satisfied simultaneously. The conjugate gradient algorithm used in HPCG to solve these equations is an iterative method, which hones closer to the solution by repeated trials. It's the simplest practical method of its kind, so it is both a real algorithm that people need and not too complicated to implement.

It also uses data structures that more closely match real applications. LINPACK's data structures are no longer used for large problems in real applications because they require storing many zero values, which worked when application problems and computer memory sizes were much smaller. Today's problems are so large that data structures must pay attention to what is zero and what is not zero, which HPCG does.

For example, in a simulation of a car's structural integrity, the terms represent how points on the frame of a vehicle directly interact with each other. Most points have no direct interaction with each other, so



most terms are zero. LINPACK data structures would require storing all terms, while HPCG only stores the nonzero terms.

"By providing a new benchmark, we hope system designers will build hardware and software that will run faster for our very large and complicated problems," Heroux said.

The HPCG code tests science and engineering problems involving complex equations, and is not related to another Sandia-led benchmark code known as Graph 500, which assesses and ranks the capabilities of supercomputers involved in so-called "big data" problems that search for relationships through graphs.

Provided by Sandia National Laboratories

Citation: Improved ranking test for supercomputers to be released by Sandia (2013, November 18) retrieved 23 May 2024 from <u>https://phys.org/news/2013-11-supercomputers-sandia.html</u>

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.