



Tougher rating system evaluates nine supercomputer capabilities(2)

Friday, November 19, 2010

Nine supercomputers have been tested, validated, and ranked by the new "Graph500" challenge, first introduced the week of Nov. 15th by an international team led by Sandia National Laboratories. The list of submitters and the order of their finish was released Nov. 17 at the supercomputing conference SC10 meeting.

The machines were tested for their ability to solve complex problems involving random-appearing graphs, rather than for their speed in solving a basic numerical problem, today's popular method for ranking top systems.

"Some, whose supercomputers placed very highly on simpler tests like the Linpack, also tested them on the Graph500, but decided not to submit results because their machines would shine much less brightly," said Sandia computer scientist Richard Murphy, a lead researcher in creating and maintaining the test.

Murphy developed the Graph500 Challenge with researchers at the Georgia Institute of Technology, University of Illinois at Urbana-Champaign, and Indiana University, among others.

Complex problems involving huge numbers of related data points are found in the medical world where large numbers of medical entries must be correlated, in the analysis of social networks with their huge numbers of electronically related participants, or in international security where huge numbers of containers on ships roaming the world and their ports of call must be tracked.

Such problems are solved by creating large, complex graphs with vertices that represent the data points—say, people on Facebook—and edges that represent relations between the data points—say, friends on Facebook. These problems stress the ability of computing systems to store and communicate large amounts of data in irregular, fast-changing communication patterns, rather than the ability to perform many arithmetic operations. The Graph500 benchmarks are indicative of the ability of supercomputers to handle such complex problems.

The Graph500 benchmarks present problems in different input sizes. These are described as huge, large, medium, small, mini, and toy. No machine proved capable of handling problems in the huge or large categories.

"I consider that a success," Murphy said. "We posed a really hard challenge and I think people are going to have to work to do 'large' or 'huge' problems in the available time." More memory, he said, might help.

The abbreviations "GE/s" and "ME/s" represented in the table below describe each machine's capabilities in giga-edges per second and mega-edges per second—a billion and million edges traversed in a second, respectively.

Competitors were ranked first by the size of the problem attempted and then by edges per second.

The rankings were:

Rank #1 – **Intrepid**, Argonne National Laboratory – 6.6 GE/s on scale 36 (Medium)

Rank #2 – **Franklin**, National Energy Research Scientific Computing Center – 5.22 GE/s on Scale 32 (Small)

Rank #3 – **cougarxmt**, Pacific Northwest National Laboratory – 1.22 GE/s on Scale 29 (Mini)

Rank #4 – **graphstorm**, Sandia National Laboratories' – 1.17 GE/s on Scale 29 (Mini)

Rank #5 – **Endeavor**, Intel Corporation, 533 ME/s on Scale 29 (Mini)

Rank #6 – **Erdos**, Oak Ridge National Laboratory – 50.5 ME/s on Scale 29 (Mini)

Rank #7 – **Red Sky**, Sandia National Laboratories – 477.5 ME/s on Scale 28 (Toy++)

Rank #8 – **Jaguar**, Oak Ridge National Laboratory – 800 ME/s on Scale 27 (Toy+)

Rank #9 – **Endeavor**, Intel Corporation – 615.8 ME/s on Scale 26 (Toy)

A more detailed description of the Graph500 benchmark and additional results are available at graph500.org.

[SOURCE](#)

Rate Article:  Average 0 out of 5

[Register](#) or [log in](#) to comment on this article!

0 COMMENTS