# 'Data motion metric' needed for supercomputer rankings, says SDSC's Snavely

August 10 2011

As we enter the era of data-intensive research and supercomputing, the world's top computer systems should not be ranked on calculation speed alone, according to Allan Snavely, associate director of the San Diego Supercomputer Center (SDSC) at the University of California, San Diego.

"I'd like to propose that we routinely compare machines using the metric of data motion capacity, or their ability to move data quickly," Snavely told attendees of the 'Get Ready for Gordon – Summer Institute' being held this week (August 8-11) at SDSC to familiarize potential users with the unique capabilities of SDSC's new Gordon data-intensive supercomputer.

Gordon, the result of a five-year, $20 million award from the National Science Foundation (NSF), is the first high-performance supercomputer to use large amounts of flash-based SSD (solid state drive) memory. With about 300 trillion bytes of flash memory and 64 I/O nodes, Gordon will be capable of handling massive data bases while providing up to 100 times faster speeds when compared to hard drive disk systems for some queries. Flash memory is more common in smaller devices such as mobile phones and laptop computers, but unique for supercomputers, which generally use slower spinning disk technology.

The system is set to formally enter production on January 1, 2012,

although pre-production allocations on some parts of the cluster will start as early as this month for U.S. academic researchers.

"This may be a somewhat heretical notion, but at SDSC we want a supercomputer to be data capable, not just FLOP/S capable," said Snavely, whom along with many other HPC experts now contend that supercomputers should also be measured by their overall ability to help researchers solve real-world science problems. Snavely's proposal includes a measurement that weights DRAM, flash memory, and disk capacity according to access time in a compute cycle.

A common term within the supercomputing community, peak speed means the fastest speed at which a supercomputer can calculate. It is typically measured in FLOP/S, which stands for FLoating point OPerations per Second. In lay terms, it basically means peak calculations per second. In June, a Japanese supercomputer capable of performing more than 8 quadrillion calculations per second (petaflop/s) was ranked the top system in the world, putting Japan back in the top spot for the first time since 2004, according the latest edition of the TOP500 List of the world's supercomputers. The system, called the K Computer, is at the RIKEN Advanced Institute for Computational Science (AICS) in Kobe, Japan, and replaced China's Tianhe-1A system as the fastest supercomputer in the rankings, which has been using this metric since 1993.

"Everyone says we are literally drowning in data, but here are some simple technical reasons," said Snavely. "The number of cycles for computers to access data is getting longer – in fact disks are getting slower all the time as their capacity goes up but access times stay the same. It now takes twice as long to examine a disk every year, or put another way, this doubling of capacity halves the accessibility to any random data on a given media.

"That's a pernicious outcome for Moore's Law," he said, noting that as the number of cycles for computers to access data gets longer, some large-scale systems are just "spending time twiddling their thumbs."

Provided by University of California - San Diego