San Diego Supercomputer Center First Academic Institution To Install IBM Blue Gene Computer

November 8, 2004

Ashley Wood

The San Diego Supercomputer Center (SDSC) announced today that it will be the first academic institution in the world to install the new IBM eServer Blue Gene Solution computing system. Designed to attack a broad range of scientific problems, the Blue Gene system will be installed in December and used for high-end, data-intensive applications in the academic environment. Funding is being provided by the National Science Foundation and the Department of Energy.

"For the past four years SDSC has been collaborating with IBM researchers on the design and testing of Blue Gene technology to help broaden its applicability," said Phil Andrews, director of high-performance computing technologies at SDSC. "SDSC believes that with the right configuration, Blue Gene can be an extremely innovative and powerful data-processing system, with potentially enormous impact in the data-intensive world."

"IBM's commitment to developing the most powerful and flexible supercomputers in the world has always included a focus on breaking new ground for scientific research," said Tilak Agerwala, vice president of systems, IBM Research. "By applying Blue Gene to tackle a range of applications including astrophysics and molecular dynamics, SDSC scientists have a flexible, affordable and immensely powerful resource to advance their work."

Among the first users of the Blue Gene system at SDSC will be astrophysicists Michael Norman of the University of California, San Diego and Robert Harkness of SDSC. In early 2004, this team ran their Enzo code on DataStar, SDSC's current IBM system, to complete the largest hydrodynamical cosmology simulation performed to date. This simulation produced over 30 TB of primary output and is typical of fluid dynamic computations that can generate huge amounts of data. Norman and co-workers will initially use the Blue Gene system to carry out cosmological parameter surveys with Enzo and to tune Enzo for the next-generation of ultra-large-scale cosmology simulations, perhaps two orders of magnitude larger than the present state-of-the-art.

The IBM eServer Blue Gene system being installed at SDSC will be housed in a single rack with 1,024 compute nodes and 128 I/O nodes, which is the maximum ratio of I/O to compute nodes to support data-intensive computing. Each node consists of two PowerPC processors that run at 700 MHz and share 512 MB of memory, giving an aggregate peak speed of 5.7 teraflops and a total memory of 512 GB. To ensure effective parallel processing, all compute nodes are connected by two high-speed networks: a 3-D torus for point-to-point message passing and a global tree for collective message passing. All I/O nodes are connected internally to the global tree and externally via gigabit Ethernet. This gives an aggregate I/O rate of 16 GBps in SDSC's data-optimized configuration.

With its large number of processors in a compact footprint, Blue Gene enables reductions in power consumption, cooling and space requirements for institutions requiring immense computing power. The new architecture's ability to produce cost-effective compute power in such a small package provides a glimpse into the future of supercomputing.

Media contacts: Ashley Wood, SDSC Communications, (858) 534-8363 Adam Emery, IBM Research, (914) 945-2121