

NSF Awards \$20 Million to SDSC to Develop “Gordon”

Supercomputer Uses Flash to Solve “Data-Intensive” Problems 10 Times Faster

November 4, 2009

Jan Zverina

The San Diego Supercomputer Center (SDSC) at UC San Diego has been awarded a five-year, \$20 million grant from the National Science Foundation (NSF) to build and operate a powerful supercomputer dedicated to solving critical science and societal problems now overwhelmed by the avalanche of data generated by the digital devices of our era.

Among other features, this unique and innovative supercomputer will employ a vast amount of flash memory to help speed solutions now hamstrung by slower spinning disk technology. Also, new "supernodes" will exploit virtual shared-memory software to create large shared-memory systems that reduce solution times and yield results for applications that now tax even the most advanced supercomputers.

Called *Gordon*, SDSC's latest supercomputer is slated for installation by Appro International, Inc. in mid-2011, and will become a key part of a network of next-generation high-performance computers (HPC) being made available to the research community through an open-access national grid. Details of the new system were announced in advance of SC09, the leading international conference on high-performance computing, networking, storage and analysis, to be held in Portland, Oregon, November 14-20.

Gordon is the follow-on to SDSC's previously announced *Dash* system, the first supercomputer to use flash devices. *Dash* is a finalist in the Data Challenge at SC09.

"We are clearly excited about the potential for *Gordon*," said SDSC Interim Director Michael Norman, who is also the project's principal investigator. "This HPC system will allow researchers to tackle a growing list of critical 'data-intensive' problems. These include the analysis of individual genomes to tailor drugs to specific patients, the development of more accurate models to predict the impact of earthquakes on buildings and other structures, and simulations that offer greater insights into what's happening to the planet's climate."

"Data-driven scientific exploration is now complementing theory, experimentation and simulation as tools scientists and engineers use in order to make the scientific breakthroughs sought by the National Science Foundation," said José L. Muñoz, deputy director and senior science advisor for the National Science Foundation's Office of Cyberinfrastructure. "SDSC's *Gordon* will be the most recent tool that can be applied to data-driven scientific exploration. It was conceived and designed to enable scientists and engineers - indeed any area requiring demanding extensive data analysis - to conduct their research unburdened by the significant latencies that impede much of today's progress. *Gordon* will do for data-driven science what tera-/peta-scale systems have done for the simulation and modeling communities, and provides a new tool to conduct transformative research."

Gordon builds on technology now being deployed at SDSC, including the new *Triton Resource* and *Dash* systems. As part of the *Triton Resource*, *Dash* leverages lightning-fast flash memory technology already familiar to many from the micro storage world of digital cameras, thumb drives and laptop computers.

"For nearly a quarter century, SDSC has been a pioneer in the field of high-performance computing," said Art Ellis, UC San Diego's vice chancellor for research. "It is therefore fitting that this Center and its staff have been chosen to develop a one-of-a-kind HPC system that not only is powerful, but also will tackle data-intensive research applications that aren't easily handled by the current generation of supercomputers."

When fully configured and deployed, *Gordon* will feature 245 teraflops of total compute power (one teraflop or TF equals a trillion calculations per second), 64 terabytes (TB) of DRAM (digital random access memory), 256 TB of flash memory, and four petabytes of disk storage (one petabyte or PB equals one quadrillion bytes of data). For sheer power, when complete, *Gordon* should rate among the top 30 or so supercomputers in the world.

Though impressive, these statistics only explain part of the machine's special capabilities. *Gordon* is ideally suited to tackle a variety of problems involving large data sets that are less concerned with raw performance than productivity.

"This will be a state-of-the-art supercomputer that's unlike any HPC machine anywhere," said Anthony Kenisky, vice president of sales for Appro. "*Gordon*, to be built on Appro's Xtreme-X™ next-generation architecture, will provide an invaluable platform for academic and commercial scientists, engineers and others needing an HPC system that focuses on the rapid storage, manipulation and analysis of large volumes of data."

"Supernode" City

A key feature of *Gordon* will be 32 "supernodes" based on an Intel system utilizing the newest processors available in 2011, and combining several state-of-the-art technological innovations through novel virtual shared-memory software provided by Scale MP, Inc. Each supernode consists of 32 compute nodes, capable of 240 gigaflops/node (one gigaflop or GF equals a billion calculations per second) and 64 gigabytes (GB) of DRAM. A supernode also incorporates 2 I/O nodes, each with 4 TB of flash memory. When tied together by virtual shared memory, each of the system's 32 supernodes has the potential of 7.7 TF of compute power and 10 TB of memory (2 TB of DRAM and 8 TB of flash memory).

"Moving a physical disk-head to accomplish random I/O is so last-century," said Allan Snavely, associate director of SDSC and co-principal investigator for this innovative system. "Indeed, Charles Babbage designed a computer based on moving mechanical parts almost *two* centuries ago. With respect to I/O, it's time to stop trying to move protons and just move electrons. With the aid of flash solid-state drives (SSDs), this system should do latency-bound file reads 10 times faster and more efficiently than anything done today."

Flash memory is designed to reduce the "latency time" of passing data to and from processor and spinning disk, and thus provides the missing link between DRAM on each processor and disk.

"Intel® High-Performance Solid-State Drives utilize Intel's NAND flash memory that is optimized for the computing platform to deliver a robust, reliable, high-performance storage solution providing this missing link," said Pete Hazen, director of marketing, Intel NAND Solutions Group. "Using Intel SSD solutions, system designers can architect high-performance storage subsystems for these types of data-intensive applications."

Gordon's 32 supernodes will be interconnected via an InfiniBand network, capable of 16 gigabits per second of bi-directional bandwidth - that's eight times faster than some of the most powerful national supercomputers to come on-line in recent months. The combination of raw power, flash technology, and large-shared memory on a single supernode, coupled with high-bandwidth across the system, is expected to reduce the time and complexity often experienced when researchers tackle data-intensive problems that don't scale well on today's massively parallel supercomputers.

"Many of these problems cannot use the 'big flop' machines effectively," said SDSC's Norman. "The reason is simple: data volume is exploding while methods to mine these data are not becoming massively parallel at the same rate; in other words, generating data is relatively easy, while inventing new parallel algorithms is hard."

Moreover, *Gordon* will be configured to achieve a ratio of addressable memory in terabytes to peak teraflops on each supernode that is greater than 1:1. By contrast, the same metric for many HPC systems is less than 1:10.

"This provides a radically different system balance point for meeting the memory demands of data-intensive applications that may not need a lot of 'flops' and/or may not scale well, but do need a large addressable space," noted Snavelly.

Potential Scientific Applications

The new SDSC system will provide benefits to both academic and industrial researchers in need of special "data-mining" capabilities. For example, scientific databases in astronomy and earth science already contain terabytes of data and continue to grow. These databases currently are stored largely in disk arrays, with access limited by disk read rates.

Gordon should also be invaluable for what's known as "predictive science," whose goal is to create models of real-life phenomena of research interest. Geophysicists within the Southern California Earthquake Center (SCEC), for instance, are using full three-dimensional seismic tomographic images to obtain realistic models of the earth's crust under Southern California. Such models are critical to planners seeking to understand and predict the impact of large-scale earthquakes on buildings and other structures along major fault lines. However, this research is now limited by the huge quantities of raw data needed to simulate this activity. *Gordon*, with its large-scale memory on a single node, should speed these computations, creating models that more closely mimic real-life temblors.

Gordon will also enable the manipulation of massive graphs that arise in many data-intensive fields, including bioinformatics, social networks and neuroscience. In these applications, large databases could be loaded into flash memory and queried with much lower latency than if they were resident on disk.

"Many scientific applications need fast, interactive methods to manipulate large volumes of structured data," said Amarnath Gupta, director of SDSC's Advanced Query Processing Laboratory. "With *Gordon*, it will be possible to place database management systems such as PostgreSQL on the flash drive and get a three-to-four fold improvement in the execution time for many long-running queries that need a high volume of I/O access."

Gordon will be housed in SDSC's 18,000 square-foot, energy-efficient data center on the UC San Diego campus, and will build on SDSC's nearly 25-year experience in deploying, operating and supporting production-quality resources that serve the national community. Faculty from UC San Diego's Computer Science and Engineering department, and Computational Science, Mathematics and Engineering program will integrate *Gordon* into undergraduate and graduate education courses and curriculum.

About SDSC:

As an Organized Research Unit of UC San Diego, SDSC is a national leader in creating and providing cyberinfrastructure for data-intensive research. Cyberinfrastructure refers to an accessible and integrated network of computer-based resources and expertise, focused on accelerating scientific inquiry and discovery. In late 2008, SDSC doubled its size to 160,000 square feet with a new building and data center extension, and is a founding member of TeraGrid, the nation's largest open-access scientific discovery infrastructure.

Media Contacts: Warren R. Froelich, SDSC Communications, 858-822-3622 or froelich@sdsc.edu Jan Zverina, SDSC Communications, 858-534-5111 or jzverina@sdsc.edu

