

# Georgia Tech wins NSF award for next-gen supercomputing

**Public release date: 21-Oct-2009**

Contact: Stefany Wilson

[stefany@cc.gatech.edu](mailto:stefany@cc.gatech.edu)

404-894-7253

[Georgia Institute of Technology](#)

The Georgia Institute of Technology today announced its receipt of a five-year, \$12 million Track 2 award from the National Science Foundation's (NSF) Office of Cyberinfrastructure to lead a partnership of academic, industry and government experts in the development and deployment of an innovative and experimental high-performance computing (HPC) system. The award provides for the creation of two heterogeneous, HPC systems that will expand the range of research projects that scientists and engineers can tackle, including computational biology, combustion, materials science, and massive visual analytics. The project brings together leading expertise and technology resources from Georgia Tech's College of Computing, Oak Ridge National Laboratory (ORNL), University of Tennessee, National Institute for Computational Sciences, HP and NVIDIA.

NSF's Track 2 program is an activity designed to fund the deployment and operation of several leading-edge computing systems operating at or near the petascale. An underlying goal is to advance U.S. computing capability in order to support computational scientists and engineers in the pursuit of scientific discovery. The award announced today is the part of the fourth round of awards in the Track 2 program.

"Our goal is to develop and deploy a novel, next-generation system for the computational science community that demonstrates unprecedented performance on computational science and data-intensive applications, while also addressing the new challenges of energy-efficiency," said Jeffrey Vetter, joint professor of computational science and engineering at Georgia Tech and Oak Ridge National Laboratory.

"The user community is very excited about this strategy," Vetter continued. For example, James Phillips, senior research programmer at the University of Illinois who leads development of the widely-used NAMD application, says "Our experiences with graphics processors over the past two years have been very positive and we can't wait to explore the new Fermi architecture; this new NSF resource will provide an ideal platform for our large biomolecular simulations."

Georgia Tech's Vetter will lead the five-year project as principal investigator. The project team is comprised of luminaries in the HPC field, including a Gordon Bell Prize winner and previous recipients of the NSF Track 2B award. Co-principal investigators on the project are Prof. Jack Dongarra (University of Tennessee and ORNL), Prof. Karsten Schwan (Georgia Tech), Prof. Richard Fujimoto (Georgia Tech), and Prof. Thomas Schulthess (Swiss National Supercomputing Centre and ORNL).

The platforms will be developed and deployed in two phases, with initial system delivery planned for deployment in early 2010. This system's innovations in performance and power will be achieved through heterogeneous processing based on widely-available NVIDIA graphics processing units (GPUs). As industry partners, HP and NVIDIA will be providing the computational systems, platforms and processors needed to develop the system.

"Research institutions are looking for energy-efficient, high-performance computing architectures that can speed time to solution," said Ed Turkel, manager of business development in the Scalable Computing and Infrastructure business unit at HP. "The combination of HP's industry-standard HPC server technology with NVIDIA processors delivers increased performance and faster application development, accelerating higher education research

projects."

The initial system will pair hundreds of HP high-performance Intel processors with NVIDIA's new next-generation CUDA architecture, codenamed Fermi, designed specifically for high-performance computing. This project will be the first of the Track 2 awards to realize the vast potential of GPUs for HPC.

"Computational science is a key area driving the worldwide application of GPUs for high-performance computing," said Bill Dally, chief scientist at NVIDIA. "GPUs working in concert with CPUs is the architecture of choice for future demanding applications."

A critical component of the program is a focus on education, outreach and training to expand the knowledge and understanding of HPC among a broader audience. The Georgia Tech team will conduct workshops to attract and train new users for the system, engage historically underrepresented groups such as women and minorities, and educate future generations on the vast potential of high-performance computing as a career field.

###

More information on the project and its resources is available at <http://keeneland.gatech.edu>.

---