HPCwire   PRINT THIS

November 10, 2010

Georgia Tech Keeps Sights Set on Exascale at SC10

Nov. 10 -- The road to exascale computing is a long one, but the Georgia Institute of Technology, a new leader in high-performance computing research and education, continues to win new awards and attract new talent to drive technology innovation. From algorithms to architectures and applications, Georgia Tech's researchers are collaborating with top companies, national labs and defense organizations to solve the complex challenges of tomorrow's supercomputing systems. Ongoing projects and new research initiatives spanning several Georgia Tech disciplines directly addressing core HPC issues such as sustainability, reliability and massive data computation will be on display November 13-19, 2010, at SC10 in New Orleans, La.

Led by Jeffrey Vetter, joint professor of computational science and engineering at Georgia Tech and Oak Ridge National Laboratory, Keeneland is an NSF-funded project to deploy a high-performance heterogeneous computing system consisting of HP servers integrated with NVIDIA Tesla GPUs. Entering its second-year, the project will deploy its initial delivery system -- the first of two experimental systems -- this month. During the initial performance runs, the Keeneland system was clocked at running 64 teraflops per second, placing it well within the top 100 systems in the world on the most recent TOP500 list of supercomputers (June 2010). Given the system's excellent energy efficiency of approximately 650 megaflops per second per watt on the TOP500 Linpack, the team is hoping to secure a strong position on the Green500 list of the most energy efficient supercomputers in the world. Keeneland is supported by a $12 million grant from NSF's Track 2D program, a five-year activity designed to fund the deployment and operation of two innovative computing systems, with an overarching goal of preparing the open computational science community for emerging architectures that have high performance and are energy efficiency.

"Heterogeneous computing will play an important role in the future of high performance computing due to the new challenges of extreme parallelism and energy efficiency," said Vetter. "The Keeneland partnership is providing hardware and software resources, training, and expertise to the computational science community at a critical time in this transition to new computing architectures."

A Georgia Tech team led by George Biros is a Gordon Bell Prize finalist at SC10 for their work demonstrating the simulation of blood flow using heterogeneous architectures and programming models at the petascale using CPU and hybrid CPU-GPU platforms, including the new NVIDIA Fermi architecture and 200,000 cores of ORNL's Jaguar system.

Reliable and sustainable computing are core aspects of DARPA's recently announced

Ubiquitous High Performance Computing (UHPC) program, a $100 million initiative to build future systems that dramatically reduce power consumption while delivering a thousand-fold increase in processing capabilities. Georgia Tech researchers are supporting several components of the NVIDIA-led UHPC team, ECHELON, while the Georgia Tech Research Institute (GTRI) will lead a fifth group, CHASM, that will develop applications, benchmarking and metrics to drive UHPC system design considerations and support performance analysis of the developing system designs.

"The key to solving the energy requirement roadblock to future systems is massive parallelism, which requires an entirely new way of thinking about today's algorithms and architectures," said Dan Campbell, senior researcher at GTRI and a co-PI of CHASM.

"UHPC provides an opportunity for anticipated application challenges to influence the high-end system designs, in ways that are outside the traditional planning of industrial roadmaps in high performance computing," said David Bader, professor of Computational Science & Engineering at Georgia Tech and Applications Lead for ECHELON.

Georgia Tech was also named an NVIDIA CUDA Center of Excellence in August 2010, further empowering the Institute to conduct game changing research and increase the computing power available to scientists and engineers through massively parallel computing.

While computing systems one thousand times faster than current petascale levels is still 10 years away, massive amounts of data are currently being generated every day in health care, computational biology, homeland security, commerce, social media and many other industries. Georgia Tech is attacking the massive data analytics challenge. The Georgia Tech-led Foundations on Data Analysis and Visual Analytics (FODAVA) research initiative is in its third year, developing state-of-the-art approaches for analyzing massive and complex data sets. In September 2010, Edmond Chow joined the Georgia Tech School of Computational Science and Engineering as an associate professor to continue his work applying numerical and discrete algorithms to simulated physical and scientific systems such as microbiology and quantum chemistry as part of Georgia Tech's new Institute for Data and High Performance Computing (GTIDH).

Georgia Tech is making the investments in personnel and infrastructure required to be positioned competitively alongside the nation's top HPC institutions. The Institute will continue to support research and educational initiatives that push the boundaries of technological capabilities and broaden the reach of computing innovation.

Visit booth 1561 at the SC10 show in New Orleans, La., November 13-19, 2010.

Source: Georgia Tech

Find this article at:
http://www.hpcwire.com/offthewire/Georgia-Tech-Keeps-Sights-Set-on-Exascale-at-SC10-107038868.html

☐ Check the box to include the list of links referenced in the article.