Accelerating Astrophysics Research with GPU Supercomputing

Swinburne University Doubles Performance with SGI Rackable and NVIDIA Tesla K10 GPUs

Key Facts

Organization:
Swinburne University of Technology

Location:
Melbourne, Australia

Application:
Educational Institution

Swinburne University of Technology in Melbourne, Australia is home to the Centre for Astrophysics and Supercomputing (CAS). Their research includes observational, theoretical and computational astronomy with priority for galaxy evolution and cosmology, globular clusters, pulsars, star and planet formation, radio astronomy instrumentation simulations, and astronomy visualization. CAS manages the supercomputing needs of Swinburne and is committed to providing an international-standard computing facility based on the latest technology, enabling Swinburne staff and students to meet their research goals while maintaining Swinburne as a premier supercomputing site.

The technology fueling this research is a next-generation supercomputer called swinSTAR, based on SGI® Rackable® standard-depth servers. Each of the 86 swinSTAR nodes is powered by two Intel® Xeon® processor E5-2600 series, and 64 of the nodes contain NVIDIA® Tesla® K10 GPU Accelerators. The swinSTAR system is the latest in an impressive line of high-performance computing (HPC) solutions employed by Swinburne, and already their second such iteration to leverage GPUs.

Pushing the Technology Envelope for Performance and Efficiency Gains

Swinburne has been at the forefront of supercomputing since the inception of CAS, constantly seeking the best technology to meet their computational requirements. They started in 1998 with a Linux Beowulf cluster. "For the first several years, the approach was to cobble together PCs, as many as possible," said Dr. Jarrod Hurley, associate professor and manager of Swinburne's supercomputer. "Some great results were obtained at the time but it was not the easiest to manage and maintain."

In 2007, they achieved significant performance gains by installing a fully integrated rack-mounted system with a theoretical peak speed in excess of 10 Tflop/s (10 trillion floating point operations per second). In addition to its speed, the new system, named The Green Machine, also increased performance-per-watt. "The Green Machine provided a major boost to our ability to process telescope data," said Hurley, "also to tackle ambitious large-scale simulations of the assembly of dark matter in the universe and galaxy formation, for example."

GPU Supercomputer Changes the Game

In 2011, Swinburne installed the first GPU Supercomputer for Theoretical Astrophysics Research (gSTAR). The system consisted of 53 SGI Rackable C3108 computer servers and four SGI® UV™ 10 systems as large-memory nodes with over 764 total processor cores built with both Intel Xeon processor 5600 series and NVIDIA Tesla GPU Accelerators. gSTAR delivered more than 130 Tflop/s of computing power, a 10X improvement over The Green Machine. "When GPUs came along, our research team was well placed to charge in early and take advantage of their benefits," said Hurley.

Star cluster image courtesy of Dr. Jarrod R. Hurley, Centre for Astrophysics & Supercomputing, Swimburne University of Technology

As an early adopter of GPU Accelerators, the CAS team ported their code to CUDA®, NVIDIA's parallel computing platform, which enables dramatic increases in computing performance by harnessing the power of the GPU. "One of our areas of research is star clusters and their evolution," said Hurley. "These clusters are very dense environments, containing millions of stars in some instances, and they are key in understanding how our Galaxy formed. The challenge from a computational perspective is to simulate 12 billion years of evolution with the stars interacting and evolving over that timestamp. It is clearly an enormous task. With a CPU-only approach we are limited to simulations of tens of thousands of stars at most." The availability of GPUs offered a readily available engine to speed-up the approach. "The translation of libraries to CUDA was relatively seamless for the N-body star cluster community," continued Hurley. "We already had the code constructed to accelerate the gravity portion of the calculations. Other members of the CAS astronomy team have been more focused on a parallel CPU-based coding approach in the past, so transition of their codes has involved more work, but will prove extremely beneficial."

The gSTAR system has further changed the game for the Swinburne CAS team. "With our previous stand-alone GPU machines, we were able to run one simulation at a time," said Hurley. "With gSTAR, we could run 100 simulations in parallel, and they're each running faster. There's a lot of uncertainty in trying to encapsulate the evolution of these star clusters. It's not an exact science. If you're running one or two simulations at a time, they may take two or three months for each simulation to complete. If you want to analyze the effects of what you're seeing, you're a bit hamstrung. Now we have a greater ability to expand our view to address these issues in a timely manner."

GPUs Enable Real-Time Data Analysis for Radio Telescopes

Hurley adds, "For a group that's processing telescope data, such as our pulsar research team, working with GPUs has had an amazing change in the way they work. In the old process, you would go to your telescope, collect your data, come back to your institution, and analyze it. If you can speed up your analysis process to be able to do real-time processing of the data as it's coming in, you can be informed by your data as you're collecting it, which can change the course of what you do with your data collection. In the past, you might have seen something during your analysis, perhaps a strange object or other anomaly, and wish you could have followed up and done something about it. Now the team is able to do some real-time analysis as the data comes in, providing a rough-cut view up front that tells them whether they should go back and look further. This capability is opening up new discoveries and informing what we're doing in real time."

Newest SGI System Delivers 2X Performance Gain

In 2012, Swinburne completed another upgrade to their SGI system, swinSTAR, which incorporates the latest generation NVIDIA Tesla K10 GPU Accelerator. The Tesla K10 is particularly well suited to high single precision workloads, which makes it a strong fit for CAS research. It is designed for throughput and performance per watt and features two ultra-efficient GPUs that provide up to 2X performance per watt for single precision applications than its predecessor, the Tesla M2090 GPU. The NVIDIA Tesla K10 features two GK104 GPUs with 1536 CUDA cores each and 4GB of memory operating at 160 GB/second per GPU.

This technology has opened many new avenues for cutting-edge simulations for studying star clusters as well as rapid processing of telescope data. "With these large-scale Rackable systems, we're light-years from where we started in 1998," said Hurley. "For a long time, the Holy Grail in the field of star cluster research has been to do a simulation of a million stars. We've been hammering away to get simulations of a few thousand stars and gather what we could from those. It seemed like the goal was going to stay out of reach. With Tesla K10s on the SGI system, I'm now really thinking that within one or two years we'll crack this and be doing million star, million body simulations."

About Swinburne University of Technology

Swinburne is a progressive university that aims to increase Australia's capacity in science, technology and innovation as the drivers of modern, internationalized economies and workplaces. The university is focused on high-impact global research, high-quality teaching and active engagement with both industry and the community. Swinburne's Centre for Astrophysics and Supercomputing (CAS) is dedicated to inspiring a fascination in the universe through research and education. Established in 1998, CAS is the largest astronomical research group in Victoria, and one of the largest in Australia.

Global Sales and Support: sgi.com/global