

Research at University of Illinois Leads to an Advanced Model for MRI Reconstruction Using OpenACC



CHALLENGE

Create an advanced imaging model to enable the rapid reconstruction of data from MRI scans, to produce a detailed, accurate image of the brain.

SOLUTION

Used OpenACC to develop software at a hackathon held at the National Center for Supercomputing Applications (NCSA) to enable MRI reconstruction using NVIDIA GPUs.

RESULTS

- Reduced the time it takes to reconstruct an MRI scan from 40 days down to a couple of hours
- Using parallel node implementation, reconstructed 3,000 brain images in less than 24 hours
- Enabled 70 times faster processing times in a supercomputing environment, compared to using a workstation

Project Overview

Brad Sutton is the Associate Professor of Bioengineering and Technical Director of the Biomedical Imaging Center at the University of Illinois at Urbana-Champaign. He works in the Magnetic Resonance Functional Imaging Laboratory (MRFIL), with a focus on the development of acquisition and image reconstruction strategies, using MRI, to accurately and quantitatively image physiology with primary applications in functional brain imaging, structural brain imaging and neuromuscular dynamics during normal speech and swallowing. Sutton's research group is also interested in biomarkers associated with declines in cognitive and motor performance that come with healthy aging.

In April 2015, Sutton's team participated in a hackathon organized by Oak Ridge National Laboratory. During this event, programmers gathered at the National Center for Supercomputing Applications (NCSA) for an OpenACC GPU accelerator workshop. OpenACC is a directive-based accelerator programming model targeted at scientists, engineers and other domain experts who are not full-time software developers.

At the hackathon, four teams worked on important applications—SIMU-LOCEAN, Nek5000, VWM and PowerGrid. The teams were given accounts on NCSA's JYC test system and the Blue Waters supercomputer, as well as NVIDIA's internal cluster, and Titan and Chester (a Cray XK7 supercomputer and a single cabinet Cray XK7, respectively) at the U.S. Department of Energy's (DOE) Office of Science User Facility at Oak Ridge National Laboratory. Teams spent the first four days learning various methods to port their codes to GPUs.

"Now that we've seen how easy it is to program the GPU using OpenACC and the PGI compiler, we're looking forward to translating more of our projects. And, because of the nature of OpenACC, I believe we'll begin to see more collaboration between our grad students, and fewer 'lone' programmers—enabling a richer research and learning environment."

Brad Sutton, Associate Professor, Bioengineering and Technical Director, Biomedical Imaging Center, University Of Illinois at Urbana-Champaign

Challenge

The challenge for Sutton's team was to create an advanced imaging model that would combine information from the MRI scan with existing data, to produce a detailed, accurate image of the brain. Typically, such high-resolution images may take up to 40 days to create, which is too long in many cases, especially if the image is being used to treat a patient.

"We get a small amount of data during a 10-minute scan, and our algorithms help to fill in the blanks for a much more detailed image, but this takes time, because there is so much data to process" he said. "We wanted to leverage OpenACC to achieve significant speed-ups in processing times and incorporate the new physics of the MRI scanner."

However, Sutton and his team are bioengineers, not computer scientists. "We have some computing skills, but they're fairly limited," he said. "We're not software programmers. Still, we need to leverage high-performance computing to complete our research." The hackathon would provide Sutton's group with the opportunity to create a simulation for the new imaging model, with support from mentors from NVIDIA, PGI and the NCSA as well as a group led by Wen-Mei Hwu from the [CUDA Center of Excellence](#) at the University of Illinois.


"The purpose of the hackathon was to try out OpenACC for coding GPUs, which are difficult to program," said Sutton. "OpenACC is an easy-to-learn language that enables significant speed ups in a short amount of time, because you can make 'suggestions' in the code regarding how to best program the GPU to accelerate processing."

Solution

Working with the PowerGrid team, Sutton's students created an advanced model for MRI reconstruction with OpenACC. A key contributing factor to the teams' success was the PGI OpenACC compiler which provided significant insight and feedback to the group throughout the process. PGI is at the forefront in developing programming models for GPU accelerators.

"We chose to use OpenACC because it is very flexible and easy to use," said Sutton. "My students could all understand it, and therefore work together and contribute effectively. Additionally, PGI's compiler enabled the team to make suggestions for how to parallelize the code, which had a significant impact on decreasing processing times."

Having only taken a quick, 2-hour online course on OpenACC, Sutton's students were able to develop working code in just a couple of weeks. The team performed image reconstructions on three environments—the Blue Waters supercomputer, a single-cabinet Cray XK7, and an engineering workstation—each of which might be appropriate for different situations. "Which environment we use would depend on demand," said Sutton. "If we have a 5-minute scan and the patient needs information back immediately, we could reconstruct the image in the lab. On the other hand, if we have several scans to process, we may use a supercomputing environment."



A high resolution (0.8 mm isotropic) diffusion tensor image of a human brain. High-resolution imaging of diffusion in the brain requires advanced acquisition and image reconstruction tools in order to keep the acquisition time short for a patient.

Results

Using their new PGI compiled software on the NVIDIA GPU, Sutton's team realized significant speed-ups. "We were able to develop some very nice software that reduced the time it would normally take to reconstruct the MRI scan from 40 days down to a couple hours," he said.

According to Sutton, during one of the runs, the team reconstructed 3,000 brain images in under 24 hours, by using many of the Blue Waters GPU nodes simultaneously—a task that would have taken months without OpenACC. Moreover, image quality has improved.

"In the past we had to approximate our calculations because they were so computationally intensive," said Sutton. "In certain situations these approximations have a negative impact on image quality; now we can achieve accurate solutions using OpenACC to maximize image quality and performance."

Brad Sutton, Associate Professor, Bioengineering and Technical Director, Biomedical Imaging Center, University Of Illinois at Urbana-Champaign

Sutton added that if the goal is to run it on hundreds of patient scans together, parallel node implementation can complete the task quickly. "Running on a single node with a GPU, the solution was 70 times faster than it would have been using just a workstation. With multiple patient scans, we can easily leverage multiple nodes on a cluster," he said.

PGI now offers a free educational license, so Sutton and his team can leverage the compilers on the Blue Waters supercomputer as well as in the lab. This will enable his team to continue to scale the solution and later integrate OpenACC into other projects, such as blood flow analysis.

"Now that we've seen how easy it is to program the GPU using OpenACC and the PGI compiler, we're looking forward to translating more of our projects," he said. "And, because of the nature of OpenACC, I believe we'll begin to see more collaboration between our grad students, and fewer 'lone' programmers—enabling a richer research and learning environment."