SUCCESS STORY | UNIVERSITY OF ILLINOIS

ACCELERATING BIOMEDICAL IMAGING

OpenACC accelerates advanced MRI reconstruction model at University of Illinois.

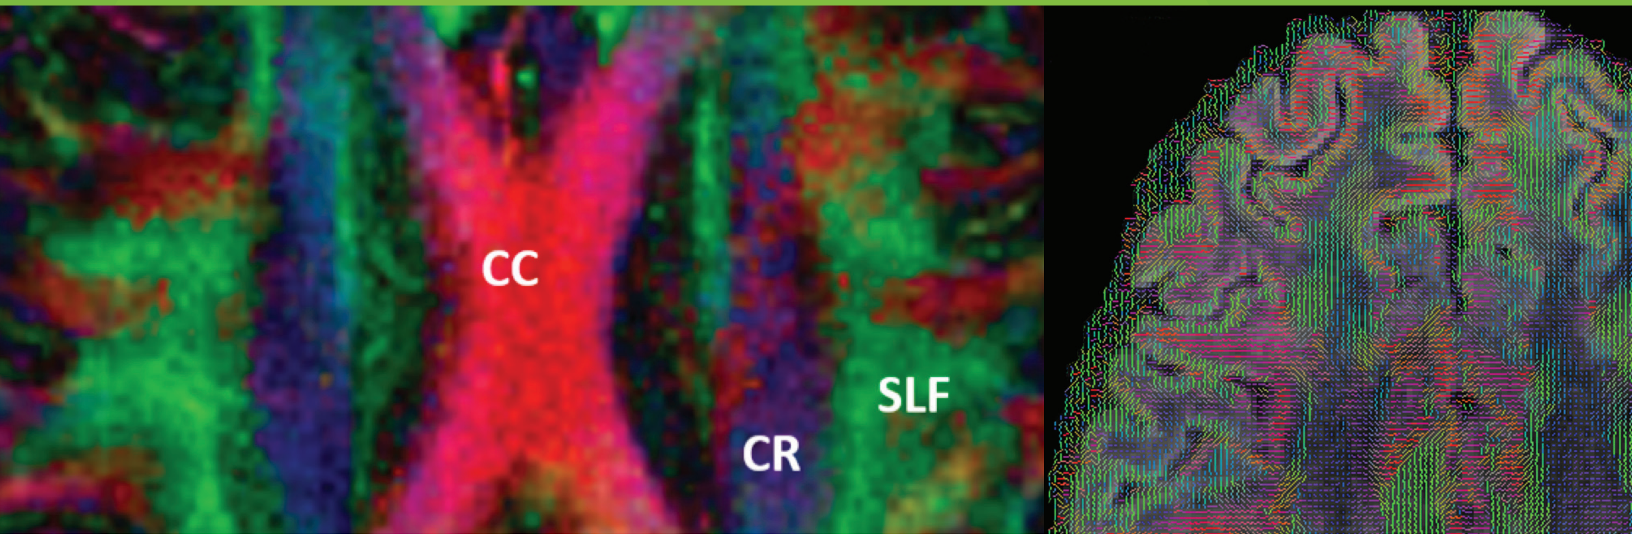

OpenACC acceleration helps reduce the time it takes to reconstruct a high-resolution scan from 40 days to a couple of hours.

SUMMARY

CHALLENGE

Accelerate an advanced imaging model that will combine information from the MRI scan with computationally intensive algorithms to produce a detailed, accurate image of the brain in minutes—, enabling faster diagnosis and treatment

SOLUTION

Used OpenACC in MRI reconstruction application using NVIDIA GPUs

RESULT

Reduced the computation time to reconstruct a high-resolution MRI scan from 40 days to a couple of hours

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. Primary applications include functional brain imaging, structural brain imaging, and neuromuscular.

Sutton's team used OpenACC to accelerate their advanced imaging model. OpenACC is a directive-based programming model designed for scientists and researchers looking to tap into the computational power of accelerators without significant programming effort.

CHALLENGE

"We get a small amount of data during a ten-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.

Typically, high-resolution detailed and accurate images may take up to 40 days to create. This is too long in many cases, especially if the image is being used to treat a patient.

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."

SOLUTION

Sutton's students created an advanced model for MRI reconstruction with OpenACC.

"OpenACC is an easy-to-learn API 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," said Sutton.

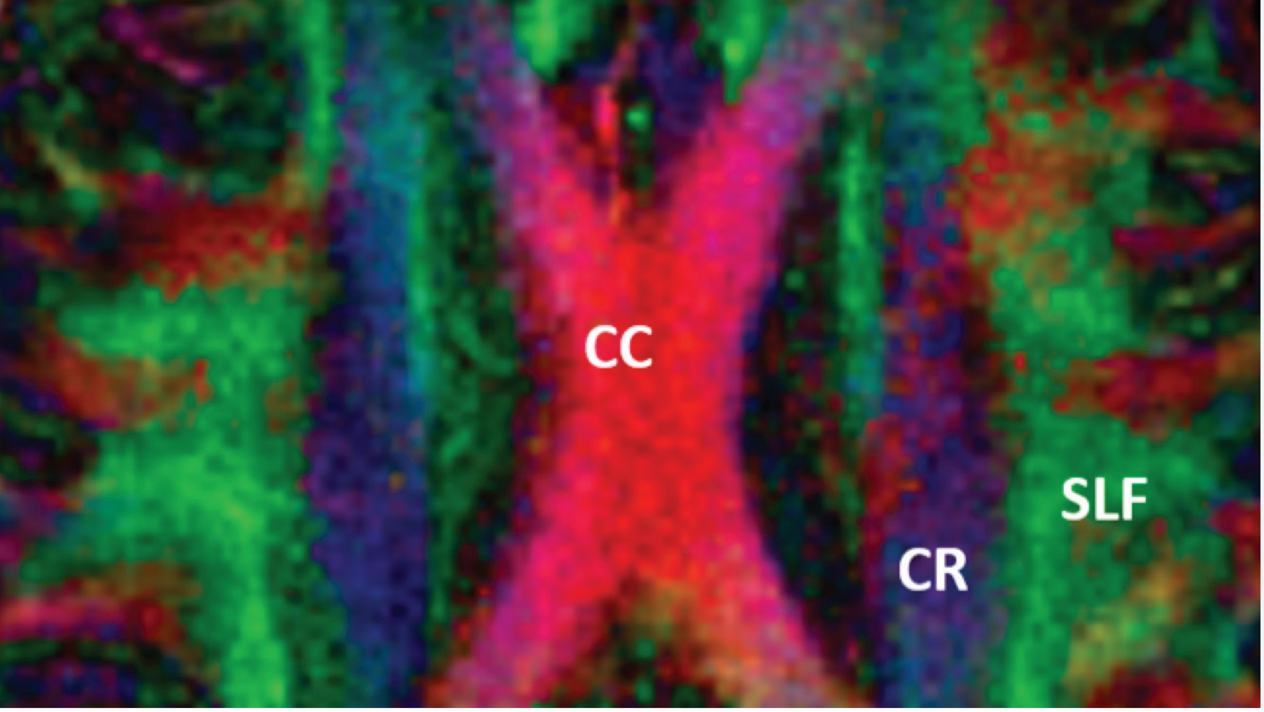

Left: A high-resolution (0.8 mm isotropic) diffusion tensor image of a human brain. High-resolution imaging of diffusion in the brain requires advanced acquisition and image reconstruction tools to keep the acquisition time short for a patient.

"We chose to use OpenACC because it is very flexible and easy to use. My students could all understand it, and therefore work together and contribute effectively."

Having taken only a two-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 five-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."

RESULT

Sutton's team realized significant speed-ups using their new PGI compiled software on the NVIDIA GPU. "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.

OpenACC also allowed the team to run on one of the fastest supercomputers in the world, Blue Waters at NCSA. The team reconstructed 3,000 brain images in under 24 hours during one of the runs 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," he said.

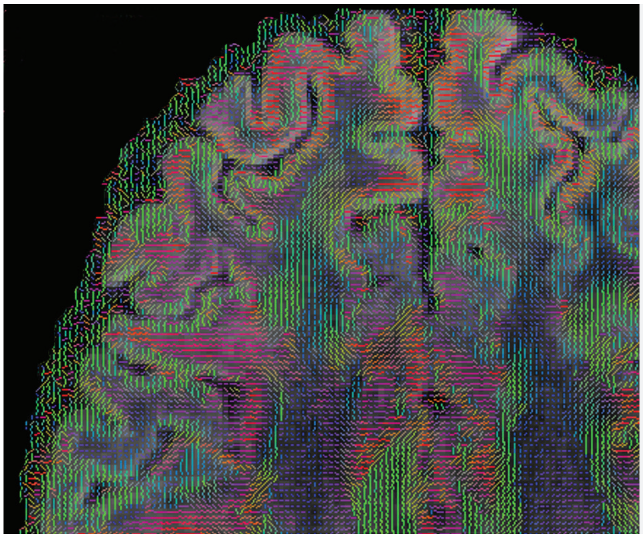

Left: High-resolution reconstructed diffusion tensor image across the brain. Source: University of Illinois.

"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," said Sutton.

"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."

OpenACC
Directives for Accelerators

OpenACC is a directives-based parallel programming model designed for simplifying programming of massively parallel processors. The model provides performance portability across a wide range of platforms, including host-GPU, multi-core, and many-core processors. OpenACC is complementary to and interoperates with existing HPC programming models, including OpenMP, MPI, and CUDA.

To learn more and to download the free OpenACC Toolkit visit **developer.nvidia.com/openacc**

Join the discussion on Stack Overflow **bit.ly/openaccforum**

NVIDIA