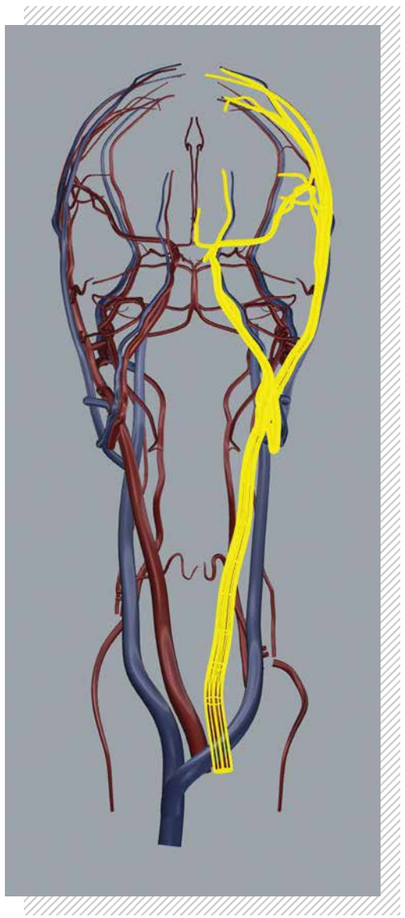

FIGURE 2: Frontal view of the patient specific geometric model of arterial system in the brain along with the superposed computational grid.

the high G forces on reduced perfusion of the brain and potential G-LOC. To summarize, following are the significant findings of our work:

1. We show that the Newtonian model for blood underestimates the wall shear-stress (WSS) by 16% in diastole, which is a non-conservative estimate from patient care perspective.

2. We show viscosity build-up in aneurisms in arteries where local ballooning effect causes a recirculation region, and viscosity build-up is significant at the end of diastole which can lead to clot formation that is caused by local stagnation.

3. We show that constant flow VADS with minimal pulsatility lead to standing waves that cause viscosity build-up which can increase the risk factor for stroke in the patients. This aspect of simulation can be used to introduce minimal pulsatility in the VAD device, thus optimizing the device for the patient, giving rise to patient-specific treatments.

4. We can calculate increased pressure fields in stiffened arteries, even when other geometric and flow conditions are held constant, thus providing insight into the effects of mechanical and rheological factors on hypertension.

5. High-performance aircraft pilots are routinely exposed to high levels of +Gz (head-to-foot) accelerations. G-LOC is a major threat to high performance aircraft pilots and astronauts. With recent advances in MRI technology and computing power, cerebral hemodynamic simulations can now be run using individual arterial geometries, and these can open the door to understanding G-LOC from an intracranial perspective.

WHY BLUE WATERS

This work relies on having many long simulations to achieve rigorous sampling of the variability in biological systems. Blue Waters was critical for both the development of **cutting-edge** software and the application of this software to perform large-scale biomechanics simulations. We found Blue Waters to be an extremely powerful and versatile computational resource that, in addition to powerful CPU and GPU hardware, provided fast interconnects that allowed us to do types of calculations that we could not have done on other platforms. Specifically, the large local memory of Blue Waters is ideally suited for our methods as we can exploit the resident memory on the processing nodes to make the macro elements "smart," reducing the size of the global problem and minimizing data communication. The Blue Waters project staff provided in-depth technical information and timely advice on the optimal deployment and performance tuning of our software.

NEXT GENERATION WORK

In the next Track-1 system we plan to extend and embed the method in a probabilistic framework for blood flow simulation in patient-specific arterial geometries, with the objective of optimization of VADs for patient-specific needs.

PUBLICATIONS AND DATA SETS

Weddell, J.C, J. Kwack, P.I. Imoukhuede, and A. Masud, Hemodynamic analysis in an idealized artery tree: Differences in wall shear stress between Newtonian and non-Newtonian blood models. *PLoS ONE*, 10(4): e0124575, 2015.

Kwack, J., S. Kang, G. Bhatt, and A. Masud, Outflow boundary conditions for non-Newtonian models for blood, (eds. K. Takizawa and Y. Bazilevs), Modeling and Simulation in Science, Engineering and Technology Book Series, Springer, 2015.

Kwack, J., A. Masud and K.R. Rajagopal, Stabilized Mixed Three-field Formulation for a Generalized Incompressible Oldroyd-B Models, *Int. J. Num. Meth. Fl.*, (2016), doi: 10.1002/fld.4287

IMPROVING THE RESOLUTION OF BRAIN BLOOD FLOW IMAGING WITH ADVANCED MRI ACQUISITIONS AND COMPUTATION

Allocation: Illinois/26.0 Knh
PI: Brad Sutton[1]
Collaborators: Alex Cerjanic[1], Joseph Holtrop[1], and Giang Chau Ngo[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Adequate blood flow in the brain is critical for supporting healthy brain function into old age. Imaging of the brain's blood flow and obtaining information about the age-related changes in the structure of blood vessels, noninvasively using magnetic resonance (MR) imaging, requires advanced approaches for measuring this small signal. The small volume fraction of blood in the brain allows for physics-based image reconstruction models to be used to improve quality and usefulness of the images. These reconstruction models can be computationally demanding, and hundreds of images may need to be acquired to estimate information about the blood flow and vessels. As part of this project, PowerGrid, a toolkit for accelerating MR image reconstructions using graphics processing units (GPUs) and distributed computing, was created in C++ and OpenACC, while also leveraging message passing interface (MPI) to distribute across multiple GPUs. Using PowerGrid, we can reconstruct full datasets of images from a patient in a time frame similar to the acquisition of the images.

INTRODUCTION

Sufficient and reactive blood flow in the brain is a critical component for the health of neurons and their supporting cells. However, advanced aging is accompanied by critical changes to the vasculature [1], including the microvasculature that is involved in exchanging nutrients and waste with tissues. Measuring changes and degradations in the microvascular architecture of the human brain is

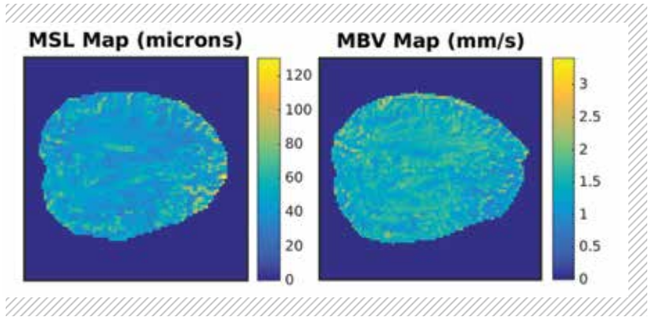

FIGURE 1: (left) Microvascular segment length (MSL) for a young subject in a slice high in the brain. (right) Microvascular blood velocity (MBV) for the same subject on the same slice. Resolution: 2mm isotropic, Scan time is 24:00 minutes total for both required scans.

limited to postmortem samples by the destructive nature of microscopy and histology used to image human brain tissue [2, 3]. Relying on destructive measurements of microvascular parameters prevents studies from assessing cognitive function of subjects, microvascular function, and observing changes longitudinally. New noninvasive, neuroimaging-based biomarkers for the state and function of the microvasculature in the brain are needed.

MR imaging provides an excellent neuroimaging platform to develop novel biomarkers for microvascular changes from age and pathology due to the flexible contrast mechanism, minimal risk profile, and broad clinical availability and applicability. Diffusion-weighted MR (DW-MR) imaging, MR imaging deriving sensitivity from the microscopic motion of water molecules, is an ideal platform for the development of novel biomarkers of blood flow in the microvasculature. The use of DW-MR imaging to characterize blood flow has been proposed with the Intravoxel Incoherent Motion (IVIM) method as introduced by Le Bihan[4]. While DW-MR imaging is exquisitely sensitive to blood motion in the microvasculature of the brain, it is difficult in practice due to low signal to noise ratio and challenges in sampling blood motion.

METHODS & RESULTS

Our team developed PowerGrid, a toolkit for accelerating iterative, model-based MR image reconstructions using GPUs and distributed memory computing. Implemented in C++, PowerGrid allows researchers and clinicians to retain the familiar structure arising from years of work developing advanced MR image reconstruction algorithms in MATLAB while leveraging high-performance computing (HPC) resources, such as Blue Waters and OpenACC.

The object-oriented structure of PowerGrid was designed to combine **state-of-the-art** image reconstruction in MR with MPI and OpenACC. This approach was necessary, as MR physicists traditionally do the majority of software development for image reconstructions inside a high-level interactive language like MATLAB.

Our use of MPI is enabled both by the code written in PowerGrid and the exploitation of natural parallelism that exists in the parallel, multi-coil imaging of modern MR scanners. These separate streams of data from each receiver coil provide a natural work unit that can be assigned to a GPU. The resulting parallelism is not complete, requiring the use of MPI communication routines inside of each global iteration of image reconstruction. Preliminary results enabled by Blue Waters and our PowerGrid project are demonstrated in the form of microvascular blood velocity (MBV), and microvessel length (MSL) maps for a single slice high in the brain of a young male volunteer are shown in Figure 1.

WHY BLUE WATERS

Blue Waters and NCSA provided excellent project support for moving our MR image reconstructions from desktop class machines to an HPC environment. Through a Blue Waters sponsored OpenACC Hackathon, our team, with hands-on assistance from two members of Blue Waters project staff, used OpenACC to accelerate the core routines of our image reconstruction utility as part of our PowerGrid project. Using PowerGrid and the Blue Waters multiple GPU support, we have shown speed up factors of up to ~11 times above the single GPU case through the use of MPI for distributed computing as shown in figure 2 for a small benchmark data set. We anticipate increased speed up factors as the dataset size and complexity increases.

Furthermore, the scale of Blue Waters supports exploring the inherent parallelism that exists in MR imaging. For each blood flow dataset, we need to reconstruct multiple 3D images from the 10-20 measurements that vary the diffusion weighting and are received from 32 parallel receiver coils [5]. This creates a large amount of data that is well-suited to parallel implementation across GPU-equipped nodes equipped with high speed interconnect. Also, the OpenACC support present in the Cray Programming Environment makes Blue Waters ideal for this work.

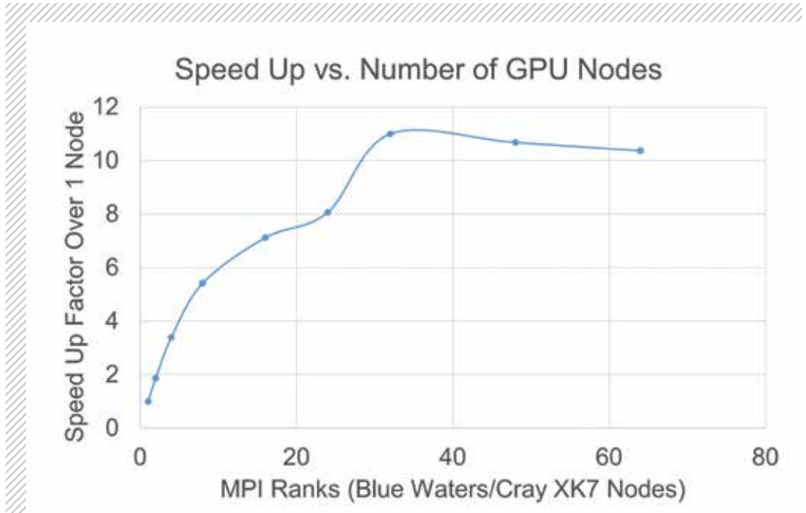

FIGURE 2: Results showing speed up versus number of MPI ranks (K20x GPUs) on Blue Waters showing peak speed up of ~11x with 32 nodes and saturation for additional ranks for a small benchmark case distributed with PowerGrid. This benchmark case represents 1/300th of a complete full brain dataset for blood flow imaging.

NEXT GENERATION WORK

Currently, it is common for MR acquisitions to use techniques to separate the encoding of a full 3D volume into smaller pieces to reduce the size of the smallest interdependent work unit or imaging volume[6]. However, recent advances in MR hardware have introduced 64 channel receivers, enabling imaging of larger volumes in less time with new data sampling patterns[7]. However, these new sampling patterns, such as a full brain 3D encoding strategy with 64-channel head coils would increase our smallest work unit by a factor of 30, presenting new demands on hardware accelerators and communication hardware.

Moving beyond larger imaging volumes, we have worked with strategies to exploit information across time or contrast weightings using 4D or higher model-based reconstructions [8, 9]. Applying these models to blood flow would easily increase the smallest work unit by a further factor of 15 – 20 for a predicted increase of 600 times the existing work unit, requiring 9,600 K20x class GPUs to reconstruct images in the same amount of time as it takes to acquire the data with a clinical MRI scanner. A next-generation Track-1 system will provide a testbed for future advanced reconstruction strategies for blood flow imaging in the human brain at high resolution at a scale not currently possible with Blue Waters.

PUBLICATIONS AND DATA SETS

Cerjanic, A., et al., PowerGrid: A open source library for accelerated iterative magnetic resonance image reconstruction. *ISMRM Annual Meeting*, Singapore, May 7th-13th 2016.