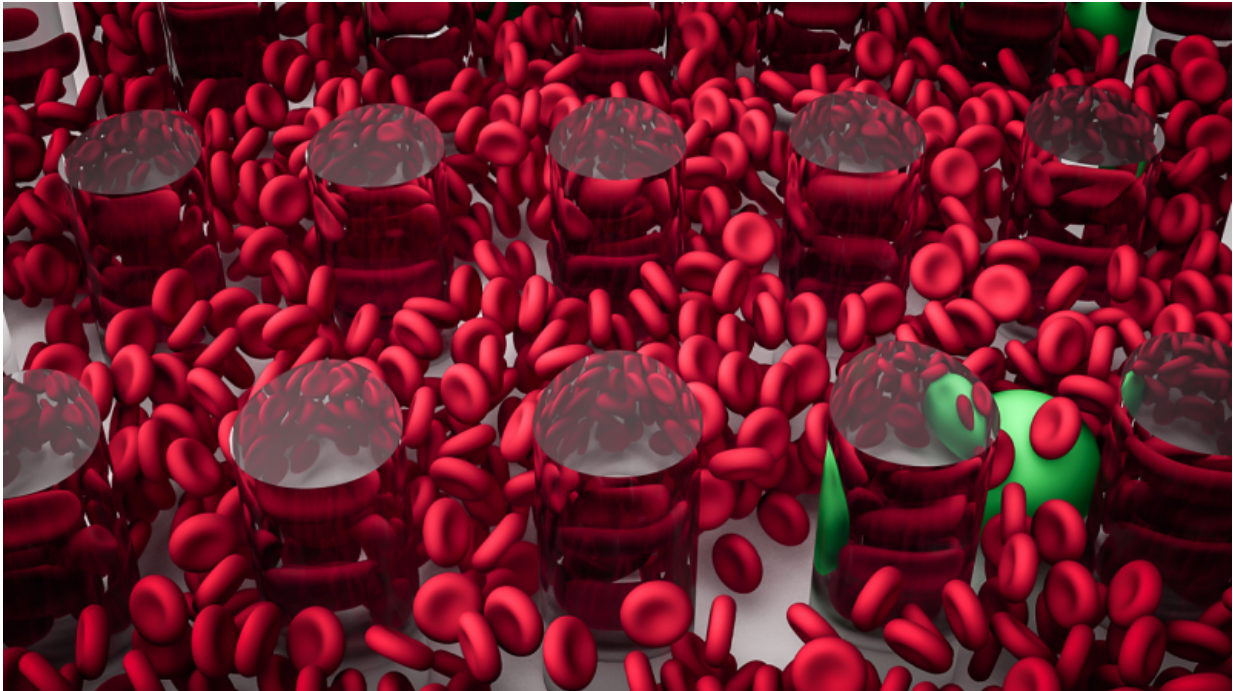# Flowing toward red blood cell breakthroughs

October 14 2015, by Eric Gedenk, Oak Ridge National Laboratory

Red blood cells (red) and circulating tumor cells (green) traveling through a microfluidic cell sorting device as simulated by uDeviceX. Credit: Yu-Hang Tang, Brown University

A team of researchers from Brown University and ETH Zurich the Universita da Svizzera Italiana (USI) and Consiglio Nazionale delle Ricerche (CNR) is using America's largest, most powerful supercomputer to help understand and fight diseases affecting some of the body's smallest building blocks.

The , led by Brown's George Karniadakis, is using the Cray XK7 Titan supercomputer at the Oak Ridge Leadership Computing Facility - a U.S. Department of Energy Office of Science User Facility located at Oak Ridge National Laboratory - to simulate hundreds of millions of red blood cells in an attempt to develop better drug delivery methods and predictors to fight against tumor formation and sickle cell anemia.

Karniadakis' group approaches these simulations from a unique place - despite the research focus on medical issues, the team actually is based in Brown's Applied Mathematics department. Karniadakis explained that mathematicians are in a good position to help guide computation for the multiple scales involved in simulating parts of the human body.

"All biological systems are multiscale systems," Karniadakis said. "The research goes from the protein, to the cell, to the tissue, all the way to the human. You have to cover scales from one nanometer to one meter."

During the first year of a three-year Innovative and Novel Computational Impact on Theory and Experiment, or INCITE, allocation at the OLCF and the Argonne Leadership Computing Facility, the team has been performing a suite of simulations related to different diseases and drug delivery methods to better predict, diagnose, and treat several mysterious hematological, or blood-based, diseases.

The team's research has made it a finalist for this year's Association for Computing Machinery Gordon Bell Prize - one of the most prestigious awards in high-performance computing - presented at the SC15 supercomputing conference, to be held this year in Austin, Texas.

The Brown team is joined by co-principal investigator Petros Koumoutsakios and Diego Rossinelli - the researcher leading the Gordon Bell effort - from ETH Zurich

At the OLCF, the team primarily has focused its disease research on sickle cell anemia (SCA) and tumor cells, as well as on developing better drug delivery methods.

Accelerating state of the art

SCA is a disease that causes red blood cells to become rigid and "sickle-shaped," leading to chronic circulatory system problems and an increased risk for death.

Despite its prevalence - roughly 8 percent of the African-American population carries the trait, and more than 180,000 babies are born with the disorder every year - very little is known about how this red-blood-cell-related disorder interacts with human blood vessels.

The team uses dissipative particle dynamics in its simulations to study blood flow as a collection of individual particles rather than one fluid object. To model each individual particle's behavior accurately for any meaningful length of time, the team needed leadership-class supercomputing power.

"Our work is done with dissipative particle dynamics, meaning you basically model everything in the simulation to be either an individual particle or a collection of particles," project collaborator and Brown doctoral researcher Yu-Hang Tang said. "It's very easy for the number of particles in the system to grow wildly. For example, if we want to model just one red blood cell, we don't just put in particles for the red blood cells; we also need particles for the fluid surrounding it. That might get you 50,000 particles just to simulate a single red blood cell."

In addition to its research on SCA, the team is also leveraging Titan to understand how diseases could be treated. Thus far, the team simulated blood and cancer cell separation using microfluidic devices. Microfluidic

devices can manipulate extremely small amounts of fluids -typically microliters (one-millionth of a liter) or smaller.

Tang focused on how blood and cancerous tumor cells might be separated by microfluidic devices, and his simulations that are one to three times larger - in terms of the number of simulated cells and computational elements - than the current state of the art within the field.

"In our tumor cell study, we used specifically arranged obstacles in the microfluidic device," Tang said. "We use these obstacles to separate cells, because different types of cells have different shapes, so when two cells hit the same obstacle, their responses to this will be different, causing them to go different directions in the device."

Such microfluidic devices would allow doctors to take a very small sample of blood and quickly identify whether someone had a malignant tumor. This "lab on a chip" could help doctors test for illness in the least invasive way possible.

Tang and his collaborators exploited Titan's GPU accelerators and developed uDeviceX, a GPU-driven particle solver - an important part of the team's code that helps plot individual particles in the simulation. Tang's new solver showed a 45-fold decrease in time to solution compared with competing state-of-the-art methods.

Moreover, Tang's extensive work with GPUs has led to the team's newest computational tool -the multiscale universal interface, or MUI. The team's research interests require a variety of different codes, with certain computing architectures working better with certain codes and different architectures benefitting other codes. It also gives the team freedom to focus specific solvers for different scales simultaneously.

MUI allows the team to quickly connect its contrasting codes into one larger code, significantly cutting down the team's computational cost for running its simulations by focusing on different hardware configurations' strengths.

The team not only can target different solvers or other portions of a code toward certain parts of a supercomputer but also is able to offload different parts of a code on different supercomputers. These distinct pieces of code do not have to communicate during each time step but ultimately will share their results during the course of a simulation.

Tang credits computer graphics for inspiring MUI. "Basically, I made MUI by borrowing a concept from computer graphics, where when you want to render the color of a pixel, you're actually doing interpolation from the nearby pixels," Tang said. "We borrowed this into the MUI world when you want to do different kinds of simulations. We proposed a general framework where you can interpolate the data that you want from nearby points by easily inserting your own interpolation algorithms."

OLCF staff member Wayne Joubert helped the team scale MUI to make efficient use of Titan's large node count.

Microscopic scales, macroscopic implications

Karniadakis and Tang emphasized that multidisciplinary collaboration must remain central in keeping their research moving forward.

"We are targeting different pathologies, because having that as a canvas, we can develop interesting mathematics and computational algorithms that can be used in other contexts," Karniadakis said, adding that some of Tang's methods already were being adopted for materials science computations.

In addition to developing more efficient algorithms, the team has worked on developing methods that would be "hardware aware," allowing them to mount their codes quickly and seamlessly on a variety of supercomputing infrastructures. Karniadakis also led a team working on domain decomposition methods - which minimize communications between a supercomputer's nodes, increasing efficiency and time to solution. The work led to the team being named a finalist for the 2011 Gordon Bell Prize.

Despite significant advancements in the team's code development, Karniadakis still sees plenty of room for improvements. As high-performance computers continue to get more powerful, Karniadakis predicts his team will develop larger-scale simulations capable of gaining deeper insight into blood-based illnesses.

"We are pushing the envelope on using current computational resources, and one of the difficulties we still struggle with is whether we can actually compute from the onset of a disease to the effects of that disease," Karniadakis said. "It can take years between the onset and effect of a disease. In this context, you cannot just rely on computing or mathematics alone, so we rely on both."

Provided by Oak Ridge National Laboratory