Taking Fluid Models Beyond Navier Stokes

Amanda Peters Randles[1], Vivek Kale[2], William Gropp[2], & Efthimios Kaxiras[1]

[1]School of Engineering and Applied Sciences, Harvard University. [2]Department of Computer Science, University of Illinois Urbana Champaign

Summary

The lattice Boltzmann method is increasingly important in facilitating large-scale fluid dynamics simulations. To date, these simulations have been built on discretized velocity models of up to 27 neighbors. Recent work has shown that higher order approximations of the continuum Boltzmann equation enable not only recovery of the Navier-Stokes hydrodynamics, but also simulations for a wider range of Knudsen numbers, which is especially important in micro- and nano-scale flows. These higher-order models have significant impact on both the communication and computational complexity of the application. We present a performance study of the higher-order models as compared to the traditional ones, on both the IBM Blue Gene/P and Blue Gene/Q architectures. We study the tradeoffs of many optimizations methods such as the use of deep halo level ghost cells that, alongside hybrid programming models, reduce the impact of extended models and enable efficient modeling of extreme regimes of computational fluid dynamics.

Cardiovascular Disease

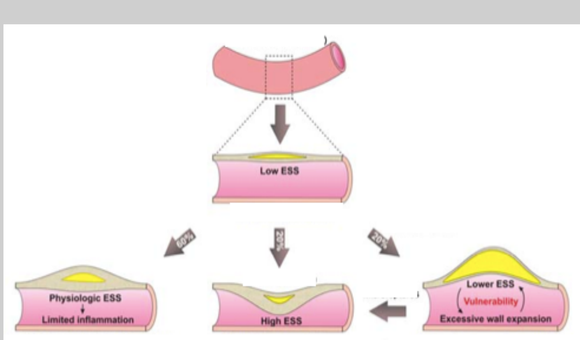

Plaque development in disease prone artery

- Leading cause of death in the Western countries
- 50% of these deaths occur without prior symptoms
- Low endothelial shear stress is a linked to heart disease progression

Currently, the only way to assess a patient's shear stress is through simulation

In collaboration with radiologists and cardiologists at the Brigham Women's Hospital in Boston, we retrieve patient specific coronary artery trees from MDCT data at a resolution of .1mm.

Lattice Boltzmann

- Simple, time-explicit, robust, and accurate
- Straight forward handling of extremely complex geometry
- Recovers Navier Stokes, and beyond
- Data locality
- Computational simplicity

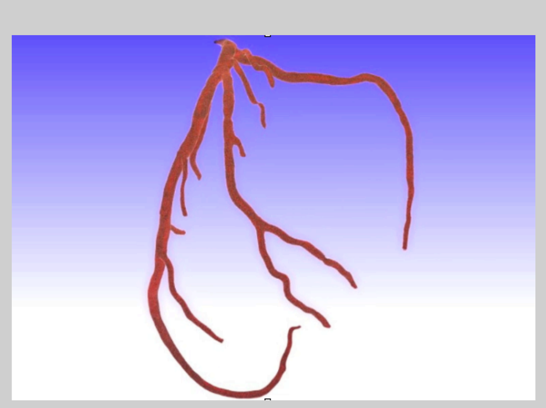

$$f_i(\vec{x}+\vec{c}_i\Delta t, t+\Delta t) = f_i(\vec{x},t) - \omega\Delta t(f_i - f_i^{eq})(\vec{x},t) + \vec{g}\cdot\partial v f$$

Streaming Collisiion External Force

HARVEY

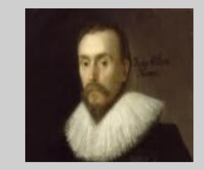

Building on our work with MUPHY, we have developed a new application, HARVEY, that is a highly extensible and efficient parallel lattice Boltzmann code.

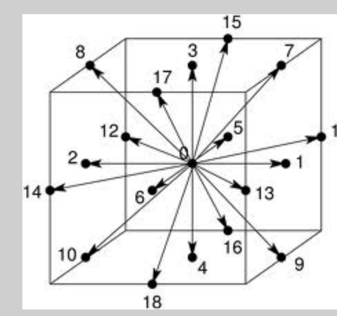

D3Q19 velocity model

- At small Knudsen numbers, high-order lattices are required.
- In HARVEY, we have implemented both the D3Q19 (second order accurate) and D3Q39 (third order accurate) velocity models.

Higher order lattice models impact both the communication and computational performance.

Is there a way to optimize the code for new architectures that can reduce this impact and enable efficient modeling of extreme regimes of CFD?

Deep Halo Ghost Cells

In order to improve communication performance through message aggregation, we introduced the idea of deep halo ghost cells. Often time applications such as those in molecular dynamics use the idea of ghost cells to limit the number of messages being sent between processors.

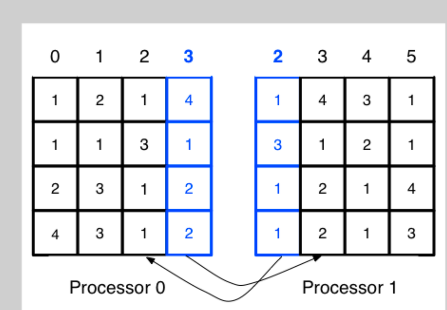

- Each processor adds extra data and computes the steps for the some rows belonging to a neighboring processor.
- The blue information shows the regions of redundant computation. This enables messages to be sent more infrequently.

For deep halo ghost cells, more data beyond simply one row was retained. This resulted in a tradeoff between memory usage, redundant computation, and communication. For systems with fairly dense grids, multiple ghost cell layers proved advantageous.

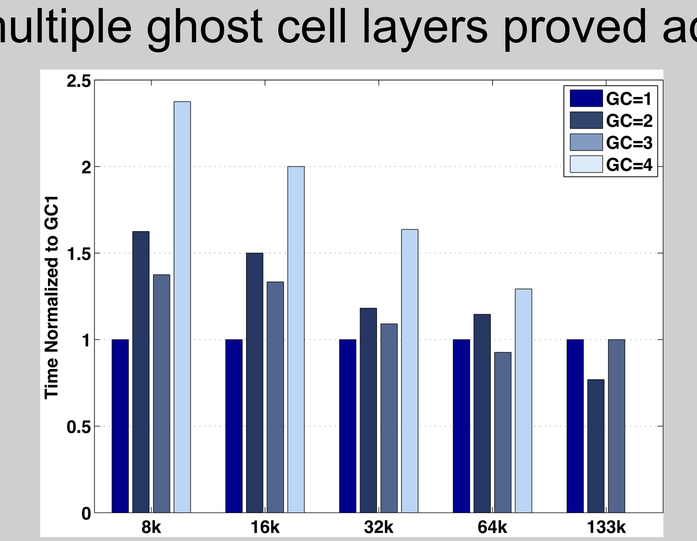

D3Q19 velocity model

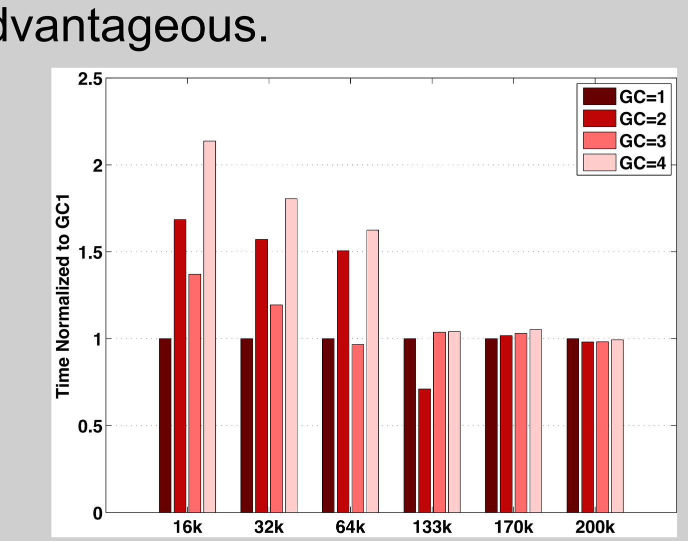

D3Q39 velocity model

Lattice Points/Proc	Ghost Cell Depth
R <= 16	1
16<R<=32	3
32<R<=66	2

Lattice Points/Proc	Ghost Cell Depth
R <= 256	1
532<R<=256	3
680<R<=532	2
800<R<=680	2 or 3

Optimizations

In this work we implemented several levels of optimizations to athe application including: ghost cells (GC), data handling (DH), compiler optimizations (CF), loop restructuring and branching reduction (LoBr), non-blocking communication (NB-C), use of a separate collide function for ghost cells (GC-C), use of intrinsics instructions (SIMD), and OpenMP. In the plots below, the solid lines represent the predicted upper bound on the performance.

Blue Gene/P

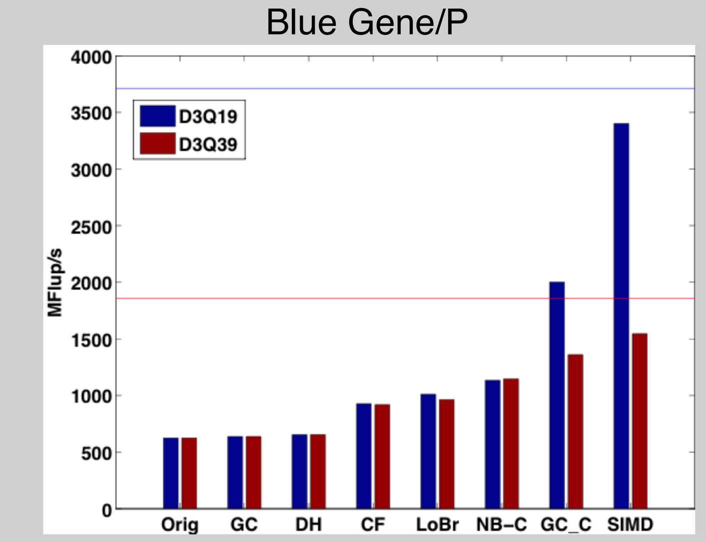

31% hardware efficiency full simulation
43% hardware efficiency for collide function

Blue Gene/Q

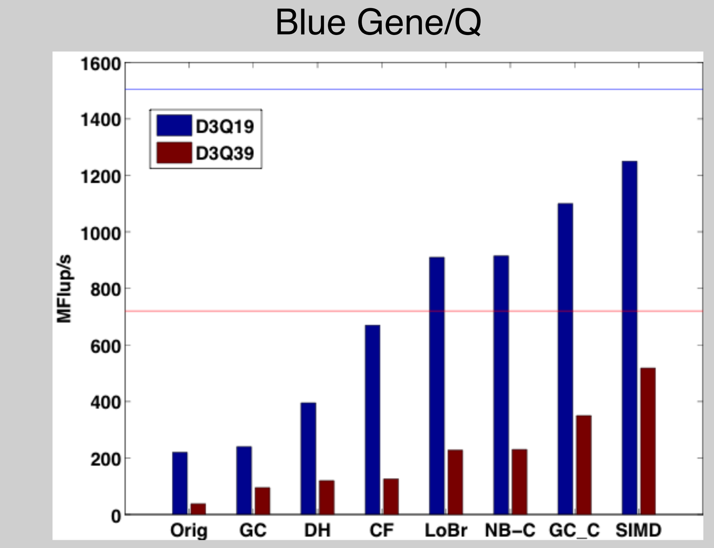

2.3% hardware efficiency full simulation
3.1% hardware efficiency for collide function

Conclusions

- Developed an efficient hybrid implementation of LBM for extreme condition fluid dynamics
- Demonstrated benefit of deep halo ghost cells
- Highlighted the potential limitations as the imbalance between bandwidth and floating point capabilities increases
- Enabled both larger systems and longer runtimes to be modeled

Future Work

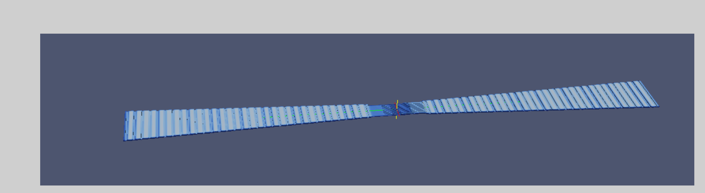

Microfluidics device simulation, Weitz Lab

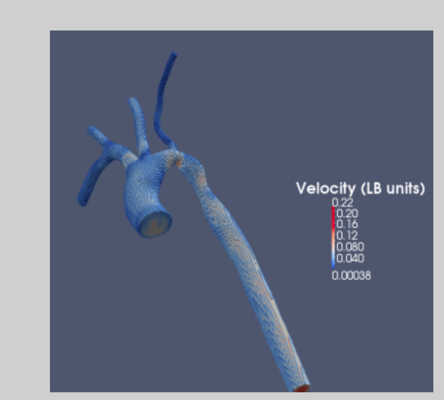

Aorta coarctation studies
accepted: STACOM'12 CFD Challenge

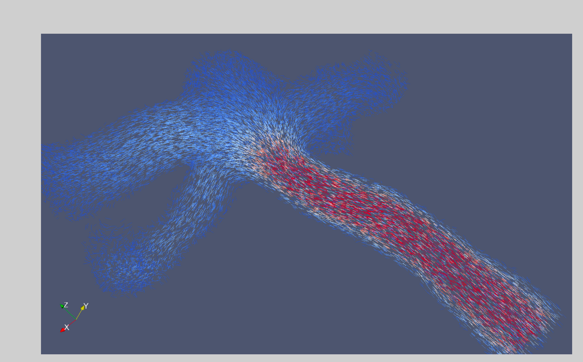

Aneurysm simulation, Auguste Lab

- Extended performance analysis.
- Endothelial shear stress studies with cardiologists from Brigham Women's Hospital.
- Parallel visualization component to be implemented in collaboration with Argonne National Lab.

Acknowledgements: This work was supported by the DOE CSGF, grant number DOE CSGF DE-FG02-97ER25308. We would like to thank F. Mintzer and D. Singer for help and support with use of the IBM Watson Blue Gene/P system and the LLNL Support Team for access to Blue Gene/Q.