GA
TN
MP
CI

# MULTISCALE SIMULATIONS OF COMPLEX FLUID RHEOLOGY

**Michael P. Howard, Princeton University**
2016-2017 Graduate Fellow

## EXECUTIVE SUMMARY

We developed the first massively parallel, open-source implementation of the multiparticle collision dynamics (MPCD) algorithm for graphics processing units (GPUs). MPCD is a mesoscale particle-based simulation method for hydrodynamics, and is particularly useful for modeling complex fluids and soft matter. Our implementation of the MPCD algorithm scales up to 1024 nodes on Blue Waters, and GPU acceleration gives a 3x speedup for an XK node compared to an XE node. The developed software will enable studies of complex fluids at length and time scales that would be otherwise inaccessible.

## RESEARCH CHALLENGE

Complex fluids, readily encountered in biology, consumer products, and industrial processing, are multicomponent mixtures that exhibit a rich variety of flow behaviors. A classic example is the cornstarch-water "oobleck" mixture, which acts like a liquid when pressed slowly but can thicken to support the weight of a person when struck quickly. Such peculiar macroscopic flow properties of complex fluids are fundamentally controlled by microscopic molecular structures and interactions. Computer simulations are ideal tools for studying this nontrivial and difficult-to-predict relationship; however, performing simulations of complex fluids at physically relevant scales presents a considerable challenge.

Many complex fluids consist of a mixture of solute macromolecules (polymers) and hard spherical particles (colloids) suspended in a liquid solvent such as water. To obtain the correct dynamics, it is essential to resolve both the direct solute-solute interactions as well as the solvent-mediated interactions. The latter dominate the computational cost for a molecular model of such a mixture; however, a molecular-level description of the solvent itself is often not of interest. A multiscale approach that simplifies the solvent model while preserving its most important interactions is required to study complex fluids at relevant length and time scales.

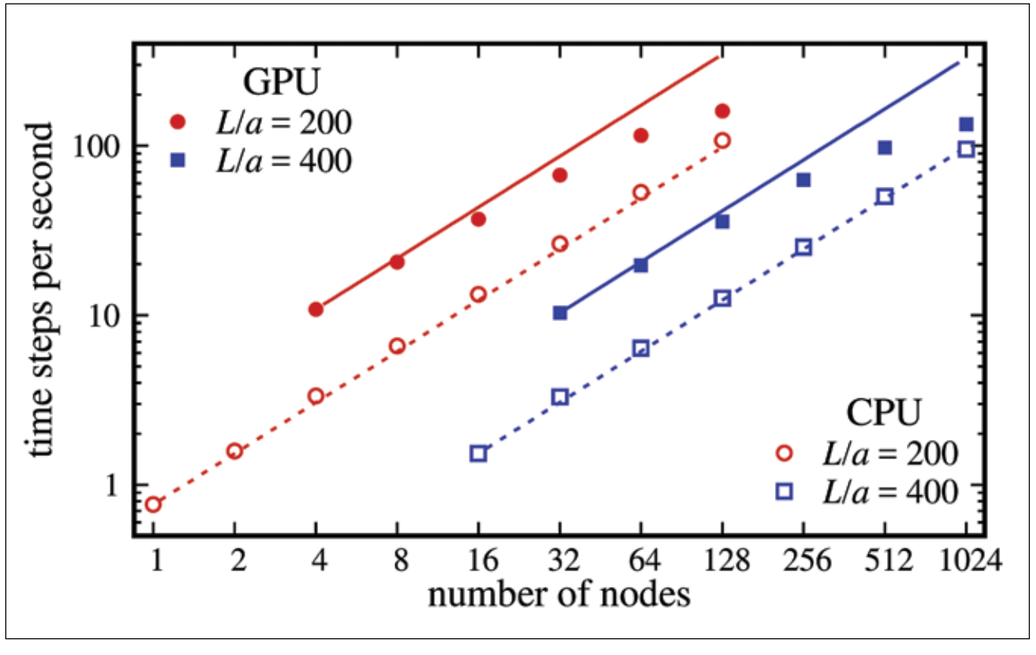


Figure 1: Strong-scaling benchmarks of MPCD software on Blue Waters for CPU-only (XE nodes) and GPU (XK nodes) implementations.

## METHODS & CODES

Multiparticle collision dynamics (MPCD) [1] is a mesoscale simulation method that combines classical molecular dynamics (MD) simulations of microscopic solutes with a coarse-grained solvent representation that faithfully reproduces long-range hydrodynamics and thermal fluctuations. In MPCD, the solvent is modeled by a set of point particles that alternate between ballistic streaming steps and cell-based, momentum-conserving, stochastic multiparticle collisions. The frequency and nature of the collisions control the transport coefficients of the solvent. Solute molecules propagate using the MD equations of motion, and are coupled to the MPCD solvent during either the streaming or collision step [2, 3]. Significantly larger length and time scales can be accessed in MPCD than with an explicit-solvent MD model because MPCD solvent particles have simple interactions with each other. However, in practice, tens of millions of MPCD particles may be required for a simulation, necessitating a parallel computational approach. We developed the first massively-parallel, open-source implementation of MPCD for graphics processing units (GPUs).

We implemented the MPCD algorithm as part of the open-source simulation package HOOMD-blue [4, 5], which has optimized MD methods designed for NVIDIA GPU architectures. All MPCD data are stored independently from HOOMD-blue's MD data to ensure high performance. Nearly all computations are performed exclusively on the GPU, which minimizes latency from host-device data transfers. We employ a spatial domain decomposition strategy to extend support to multiple GPUs [5]. One MPI process is assigned per GPU, and we perform runtime autotuning to ensure optimal CUDA kernel launch parameters. Flexible initialization and simulation setup is supported through a scriptable Python user interface.

## RESULTS & IMPACT

We performed strong-scaling benchmarks of our MPCD implementation on Blue Waters. We benchmarked two cubic simulation boxes with edge lengths $L = 200a$ and $L = 400a$, where $a$ is the size of an MPCD cell. There were 10 MPCD solvent particles per cell, giving 80 million and 640 million total particles in each simulation box, respectively. We benchmarked a CPU-only implementation of our code using 16 processes per XE node and our GPU implementation using 1 process per XK node. The CPU-only code showed excellent strong scaling up to 1024 nodes. The GPU code also showed good scaling to 1024 nodes, with some performance lost at the highest node counts owing to communication latency for the MPCD cells. GPU acceleration on the XK nodes gave a roughly 3x speedup compared to the XE nodes. The complex fluid and soft matter research communities will significantly benefit from the developed MPCD software, especially the GPU implementation, which will permit studying processes at scales that would otherwise be inaccessible.

## WHY BLUE WATERS

Blue Waters is the only system available to us that delivers both the CPU and GPU resources necessary to develop and optimize our software at scale. The large number of GPUs available in the XK nodes also significantly increases our overall scientific productivity and allows us to study process dynamics that would be challenging or impossible to obtain with fewer resources.

Michael P. Howard is currently completing the fourth year of a Ph.D. program in chemical engineering at Princeton University, under the direction of Athanassios Z. Panagiotopoulos. He expects to graduate in 2018.