MANY-GPU SIMULATIONS FOR SOFT MATTER DESIGN

Allocation: GLCPC/0.443 Mnh
PI: Sharon C. Glotzer[1]

[1]University of Michigan

EXECUTIVE SUMMARY:

We aim to leverage large-scale parallel simulations on GPUs to discover fundamental principles of how soft matter systems self-assemble and how to control the assembly process to engineer new materials. Our open-source molecular dynamics (MD) package HOOMD-blue [1], designed specifically for GPU architectures, uses completely device-resident data structures; all simulation work is performed on the GPU so the CPU merely acts as a driver for the GPU, and it offers a rich feature set for general-purpose particle-based simulations [2]. We implemented MPI spatial domain decomposition and additional code optimizations and demonstrated unprecedented performance for MD simulations running on thousands of GPUs in parallel. In addition to carrying out benchmarks, we used Blue Waters to run dissipative particle dynamics (DPD) on large systems for 50 million time-steps each to reproduce experimental findings for formation of nanofibrous microspheres and elucidate the mechanisms of self-assembly.

INTRODUCTION

Multi-GPU simulations are becoming a mainstay in scientific computing. Simulations of bulk polymeric liquids or large assemblies of anisotropic shapes involve large numbers of particles, which are best simulated using spatial domain decomposition. Such a work distribution allows the simulation of a single system on many processors—GPUs and CPUs—and results in high scaling efficiency. Often, speed-ups of several orders of magnitude can be achieved over execution on a single processor. The challenge in multi-GPU simulations is overcoming the inherent latencies of the PCIe bus, which limits communication between these extremely fast processors. Our multi-GPU enabled HOOMD-blue 1.0 contains many optimizations that were tested and benchmarked on Blue Waters. For the first time we were able to simulate self-assembly of hollow microspheres with a fibrous nanostructure and a hollow microstructure. The discovered mechanisms provide guidance on simultaneous control of nano- and micro-structure formation for synthesis of nanostructured particles, which may broadly impact biomedical and other emerging technologies.

METHODS & RESULTS

The first HOOMD releases up to 0.11.3 are highly tuned for single-GPU performance and do not support multi-GPU runs. Reducing latency is one of the biggest challenges in developing a code scalable to many GPUs. Data transferred between GPUs moves over the PCIexpress bus (PCIe), whose bandwidth (up to 16 GB/s) and latency (several μs) is much slower than on-board GPU memory (250 GB/s,~100 ns). Communicating over PCIe adds latency that is not present in single-GPU runs. In the strong scaling limit of increasing the number of GPUs, P, at constant number of particles, N, the work N/P performed by each GPU decreases to the point where it is too small to fully use the device. We addressed this challenge using the MPI protocol to support communication between different GPUs on the same node and between different nodes. Our communication routines are implemented on the GPU to reduce the data transferred over PCIe and to allow us to take advantage of CUDA-aware MPI libraries. We updated a majority of the classes in HOOMD-blue to support MPI, including the file I/O classes, integrators, pair and bond potentials, and analyzer classes. We also optimized for strong scaling on thousands of GPUs, which we achieved using a design for the neighbor list and force computation kernels based on cooperative thread arrays and an auto-tuning algorithm.

Our benchmarks on Blue Waters further establish GPUs as extremely fast engines for MD simulation compared to traditional CPU cores. GPUs not only realize an order of magnitude speed-up over current-generation CPUs, but they also scale very well using spatial domain decomposition.

This functionality is critically important in the study of real systems with interactions at the nanoscale that can dictate macroscale behavior. In our *Advanced Materials* paper, a collaborative effort with an experimental group at the University of Michigan, we studied the behavior of microdroplets with tunable porosity, which serve as implants to aid in tissue regeneration. By varying the star polymer architecture (arm number or length), we demonstrated that droplets change from solid to hollow or porous structures. These features are critically important for proper tissue adhesion. We systematically investigated the droplet formation for several lengths of polymer and arm number, each at numerous droplet sizes. These DPD simulations required over 10 million particles, divided between polymer and solvent, to approach the micron-sized length scale of smaller droplets. Our simulations would not have been possible without the scalable performance of HOOMD-blue and simultaneous access to up to 128 GPUs on Blue Waters to investigate the droplet phase diagram.

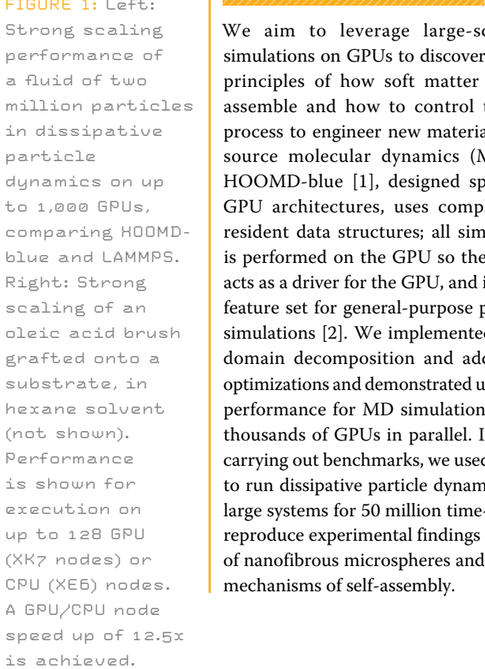

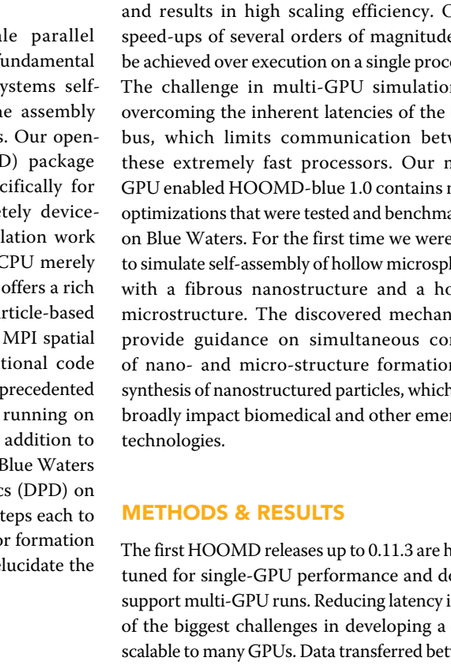

FIGURE 1: Left: Strong scaling performance of a fluid of two million particles in dissipative particle dynamics on up to 1,000 GPUs, comparing HOOMD-blue and LAMMPS. Right: Strong scaling of an oleic acid brush grafted onto a substrate, in hexane solvent (not shown). Performance is shown for execution on up to 128 GPU (XK7 nodes) or CPU (XE6) nodes. A GPU/CPU node speed up of 12.5x is achieved.

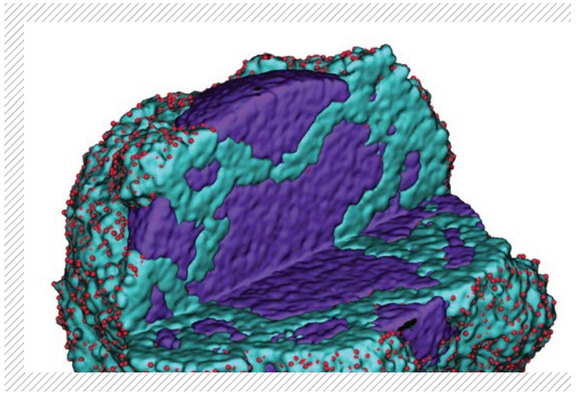

FIGURE 2: Isosurface of large-scale DPD simulation of 15-arm star-shaped polymers of varying arm length, showing spongy structure formation.

WHY BLUE WATERS

Blue Waters is the only NSF-funded system that offers large-scale access to GPUs. GPUs can accelerate soft matter simulations by more than an order of magnitude and thereby drastically improve the turnaround for research. We used Blue Waters to benchmark the performance of the HOOMD-blue code on up to 1,024 nodes, and the results have been published in Glaser et al [2]. Additionally, as recently reported in *Advanced Materials*, our Blue Waters allocation allowed us to run large systems on up to 128 nodes to reproduce experimental findings for structure and formation of nanofibrous microspheres and elucidate the mechanisms of self-assembly for the system.

PUBLICATIONS

Glaser J., Nguyen T.D., Anderson J.A. et al., Strong scaling of general-purpose molecular dynamics simulations on GPUs. *Comput. Phys. Commun.* 192, (2015), pp. 97-107. DOI:10.1016/j.cpc.2015.02.028.

Zhang Z., R. L. Marson, Z. Ge, S. C. Glotzer, and P. X. Ma, Simultaneous Nano- and Microscale Control of Nanofibrous Microspheres Self-Assembled from Star-Shaped Polymers, *Adv. Mater.* 27:26, (2015), pp. 3947–3952. DOI: 10.1002/adma.201501329