# Implementing Dissipative Particle Dynamics for Hybrid Supercomputers

**T. D. Nguyen**[1] **and S. J. Plimpton**[2]

**Abstract:** Dissipative particle dynamics (DPD) is a coarse-grained particle-based simulation method that offers microscopic-scale insights into soft matter systems. We present here an efficient implementation of the standard DPD model for graphical processing units (GPUs). As implemented in the LAMMPS molecular dynamics package, it can run effectively on current-generation supercomputers, which often have hybrid compute nodes containing multi-core CPUs and (one or more) GPUs. In our implementation, 1) muliple CPUs can share a GPU device efficiently on a compute node, and 2) DPD interactions are computed on the GPU concurrently with other portions of a full simulation model (such as boundary conditions, constraints, bonded interactions, and diagnostic calculations) on the CPU. Our GPU-accelerated version shows a speedup of up to 9.5x versus many-core CPU simulations, and maintains a high parallel efficiency across thousands of compute nodes. We discuss how the new GPU implementation was validated against the CPU version for thermodynamics, diffusion, and hydrodynamic behavior. We also highlight large-scale models which the faster DPD implementation has enabled, for studies of block copolymer self-assembly and thin-film instability.

Keywords: Dissipative particle dynamics, LAMMPS, GPU acceleration, hybrid CPU/GPU, high-performance computing.

[1] National Center for Computational Sciences
Oak Ridge National Laboratory
Oak Ridge, Tennessee, 37831, USA
*nguyentd@ornl.gov, ndactrung@gmail.com*

[2] Sandia National Laboratories, Albuquerque, New Mexico, 87123, USA
*sjplimp@sandia.gov*