Applications of GPU package to soft matter simulations

Trung Dac Nguyen, Jan-Michael Y. Carrillo, W. Michael Brown

National Center for Computational Sciences Oak Ridge National Laboratory

LAMMPS Accelerator Library

- Available in LAMMPS as "GPU Package"
- Support for many commonly used pair styles
  - Pair-wise potentials,
  - Three-body potentials: EAM, Stellinger-Weber,
  - Particle-Particle Particle-Mesh
- Multiple MPI tasks share available Accelerators
  - Still get parallelization on CPU and compatibility with other LAMMPS features
- CPU/Accelerator concurrency for force computation
- CUDA and OpenCL compatibility via a generic API (*Geryon*)

Brown, Wang, Plimpton, Tharrington, *Comput. Phys. Comm.*, 182, 898-911, **2011**
Brown, Kohlmeyer, Plimpton, Tharrington, *Comput. Phys. Comm.*, 183, 449-459, **2012**
Brown, Yamada, *Comput. Phys. Comm.*, To Appear, **2013**

ORNL Computing resource

Jaguar Cray XK6 (2011)
18,688 nodes:
- 16-core AMD Opteron 6200
- 1 NVIDIA Tesla X2090 (Fermi)
- 32 GB RAM

Gemini interconnect (75 Gbits/s)

Titan Cray XK7 (2012)
18,688 nodes:
- 16-core AMD Opteron 6274
- 1 NVIDIA Tesla K20X (Kepler)
- 32 GB RAM

Soft matter systems

of which the energy scale is on the order of $k_B T$

Charged polyelectrolytes

Liquid metallic rings/stripes

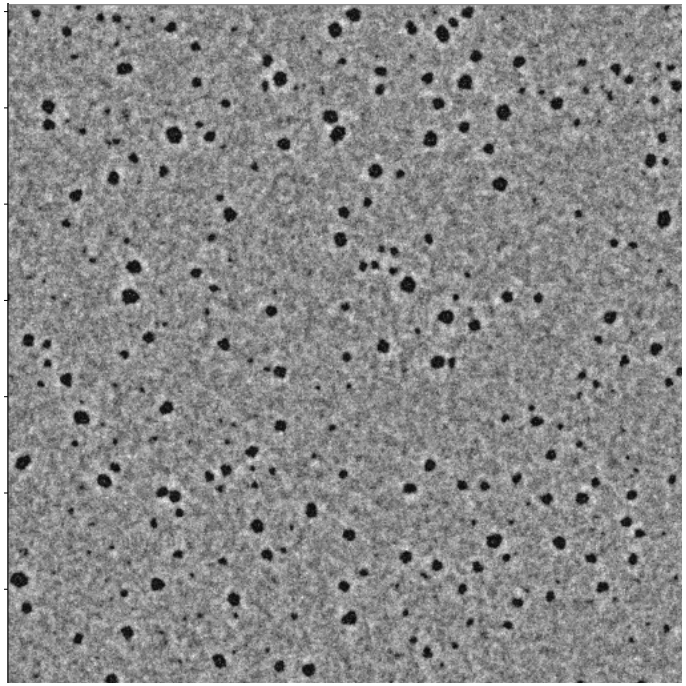

Liquid crystal films

Charged polyelectrolytes

- Objectives:
  - Structural and dynamic features as a function of the distance between two grafted surfaces
- Computational challenges:
  - Long-range solver with sparse charges
  - Slab geometry; periodicity in *z* undesirable
- Resolutions
  - Longer cutoff for real-space calculations (GPU) + coarser mesh for k-space calculations (CPU): **3x (GPU vs CPU-Only PPPM on XK6)**
  - Truncated shifted-force Coulombic potential (lj/cut/coul/dsf) for systems of uniformly distributed charges: **~12x (GPU vs CPU-Only PPPM on XK6)**

Fennel, Gazelter, *J. Chem. Phys.*, 124, 234104, **2006**
Carrillo, Russell, Dobrynin, *Langmuir*, 27, 14599, **2011**

Disjoining pressure vs. gap

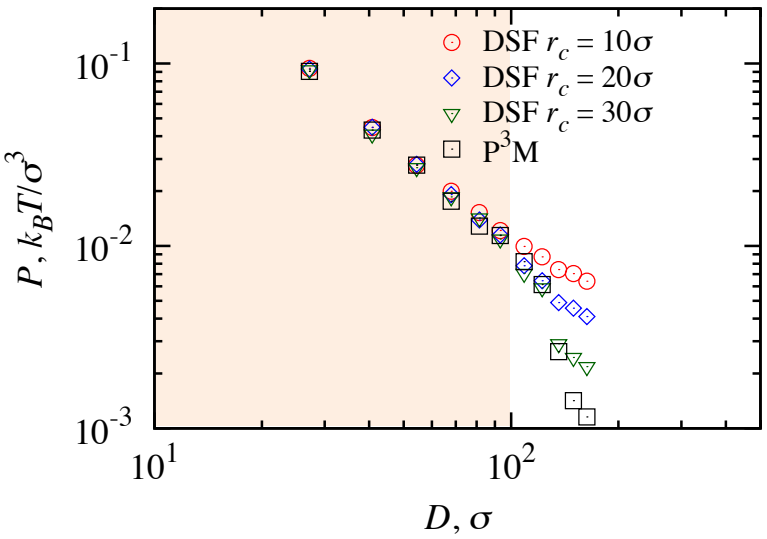

Sliding velocity vs. shear stress

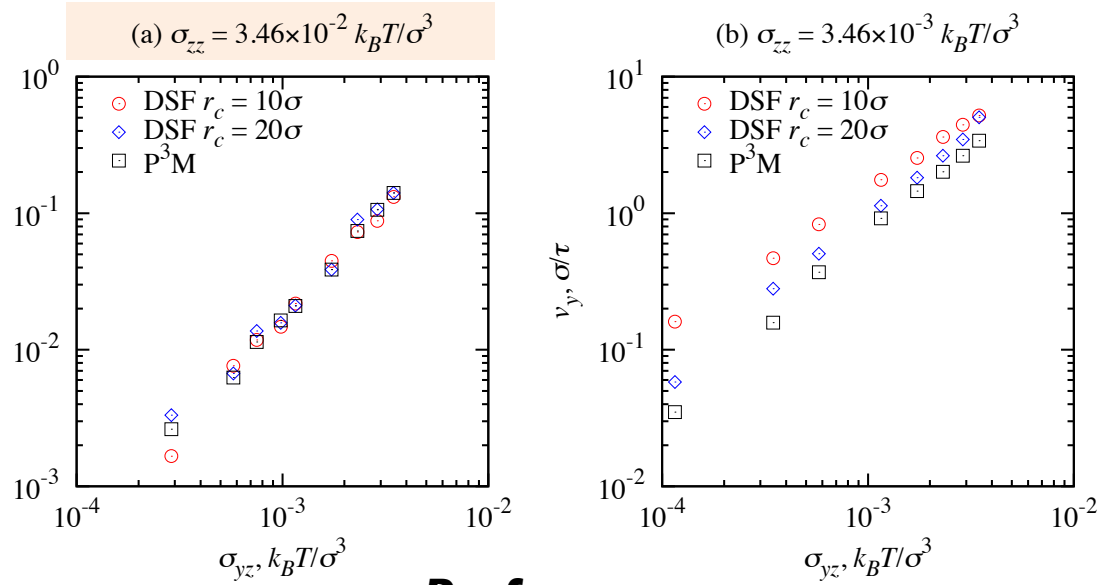

Performance

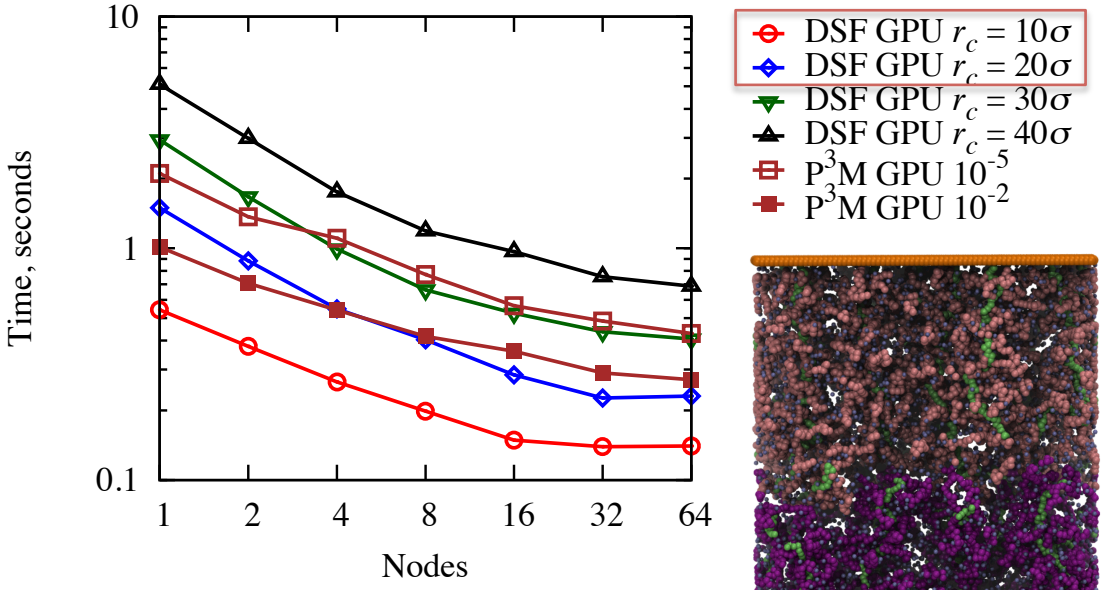

Model:

- Polymer chains randomly grafted to opposite surfaces
- Charges on brush monomers; neutralizing counterions freely floating

Simulation:

- NVT (Langevin thermostat)

Nguyen, Carrillo, Dobrynin, Brown, *J. Chem. Theory Comput.*, 9, 73-83, **2013**

Lessons learned

- Shift the computational work to real-space calculations by increasing cutoff
- Use verlet/split to perform k-space and real-space calculations concurrently
- Use GPU acceleration to decrease the number of MPI tasks performing long-range calculations
- Increase the number of threads per atom with the real-space cutoff
- Enhanced truncation methods are accurate for certain systems but need further work to characterize charge screening in condensed phase

Liquid metallic rings and stripes

- Objectives:
  - Competition between instability and surface tension minimization that leads to breakup or collapse of thin rings and stripes on a solid substrate.
- Computational challenges:
  - Many-body interaction between liquid atoms (modeled by the EAM potential)
- Resolution
  - EAM (liquid-liquid) and LJ (solid-liquid) are performed on GPUs
  - Performance gain: **~3-4x (GPU vs CPU-only on XK7)**

Brown, Nguyen, Fuentes-Cabrera, Fowlkes, Rack, Berger, Bland, *Procedia Computer Science*, 9, 186, **2012**

Breakup vs. collapse

Model:
- Liquid atoms interact via Cu EAM
- Liquid-substrate atoms interact via Lennard-Jones 12-6
- Substrate atoms are immobile

Simulation:
- NVT (Nose-Hoover thermostat)
- N = ~100K atoms; up to 64 nodes

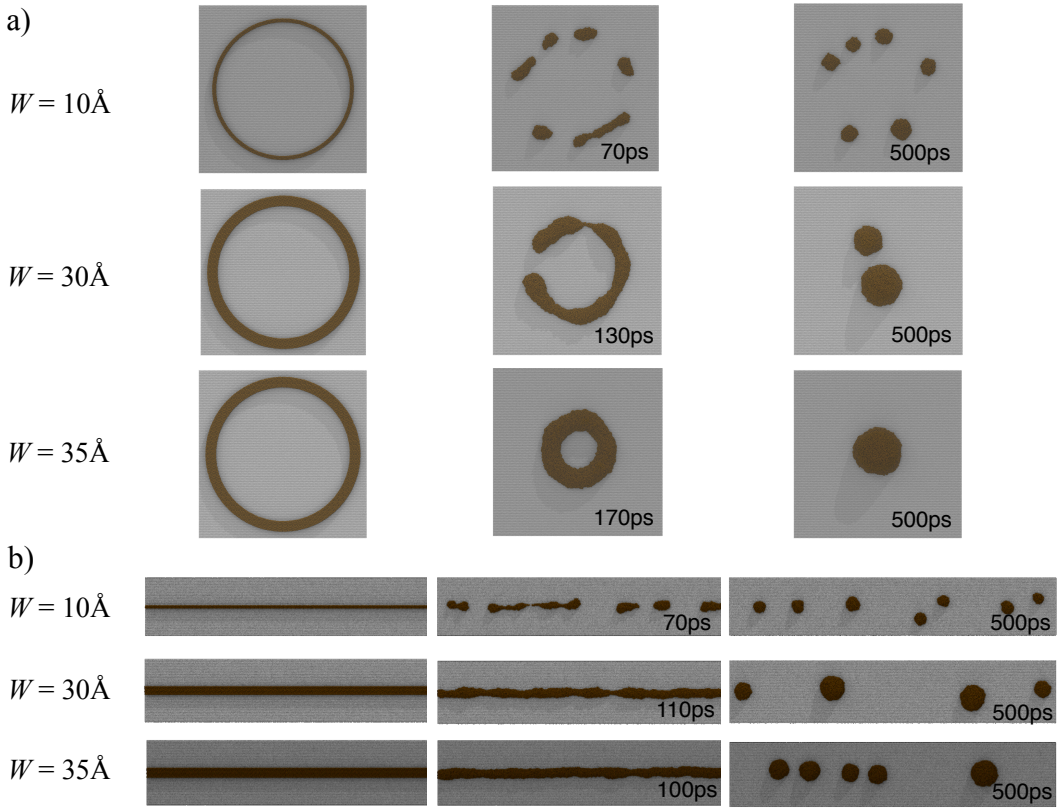

MD vs. continuum hydrodynamic simulations

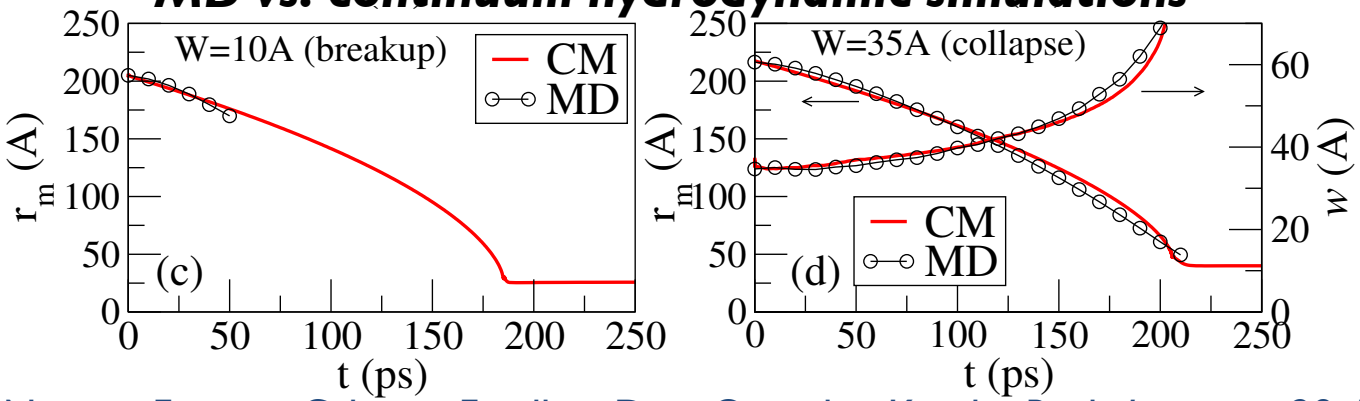

Nguyen, Fuentes-Cabrera, Fowlkes, Diez, Gonzalez, Kondic, Rack, *Langmuir*, 28, 13960, **2012**

Lessons learned

- Performance gain with GPU acceleration for EAM is modest because of the short-ranged EAM (0.5 nm for Cu) and an additional communication for electron density
- **processors** and **fix balance** are helpful for load balancing.

Liquid crystal dewetting

- Objectives:
  - Dewetting of liquid crystal films
- Computational challenges:
  - Anisotropic interaction between liquid crystal molecules
- Resolution
  - Gay-Berne interactions are performed on the GPUs
  - Performance speedup: **>7x (GPU vs CPU-only on XK7)**

Brown, Nguyen, Fuentes-Cabrera, Fowlkes, Rack, Berger, Bland, *Procedia Computer Science*, 9, 186, **2012**

Model:

- 3:1 LCs interact via Gay-Berne
- Spherical substrate atoms are immobile

Simulation:

- NVT (Nose-Hoover thermostat)
- fix ave/spatial

N = 10-30M LC molecules + ~10M substrate atoms; 400-4900 nodes

Film thickness h(x, y)

$h_0 = 4\sigma_0$

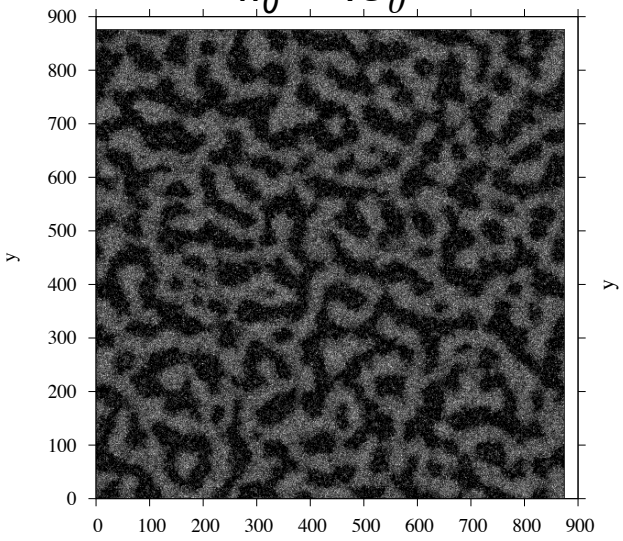

$h_0 = 6\sigma_0$

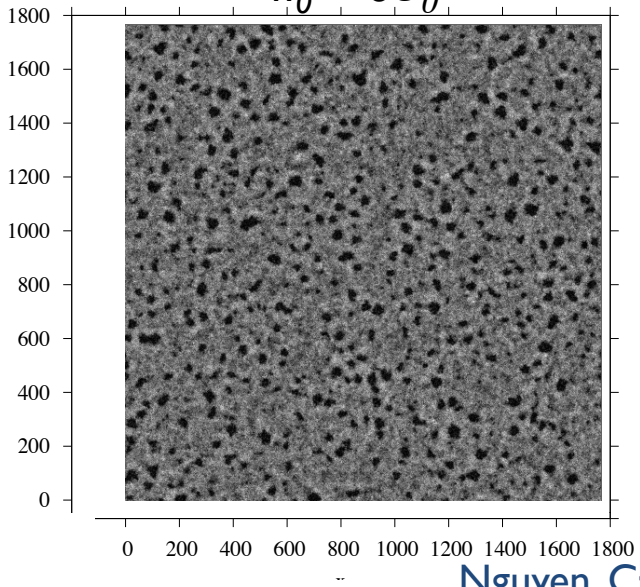

$h_0 = 8\sigma_0$

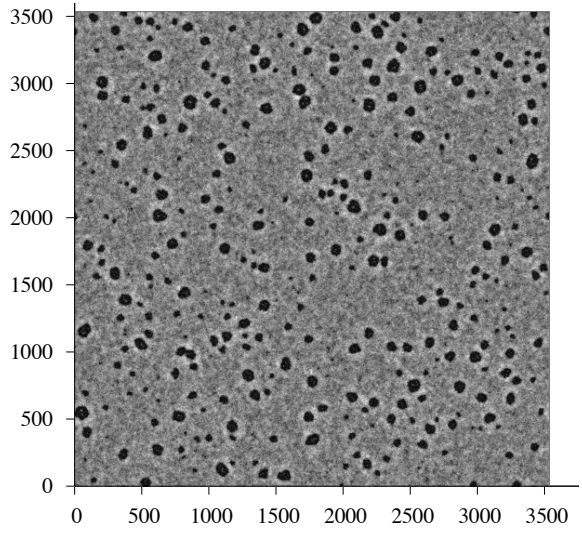

Nguyen, Carrillo, Brown, *in progress*

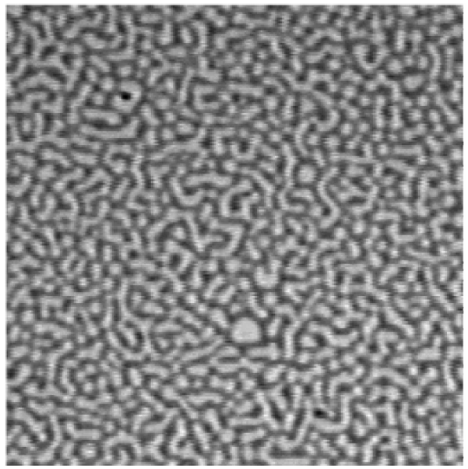

Experiment: Isotropic films of 5CB LCs dewetting

Vandenbrouck, Valignat, Cazabat, *Phys. Rev. Lett.*, 82, **1999**

FIG. 1. View of spinodal dewetting of a 42.8-nm-thick film ($T = 33.5 \pm 0.5$ °C). This pattern is reminiscent of those observed in spinodal decomposition studies. The image size is 460 μm × 460 μm.

Lessons learned

- Optimal number of MPI tasks sharing a GPU is 4
- Optimal number of atoms per GPU is ~15-30K
- Use LAMMPS enhanced MPI task placement for Cartesian grid (via processors grid numa)
- Use the CUDA Proxy server on Titan so MPI processes can efficiently share the GPU
  - Compile lib/gpu with –DCUDA_PROXY and 'export CRAY_CUDA_PROXY=1' on Titan

Summary

- **GPU package can offer 3-7x speedup**
  - Coulombic interactions
  - Many-body interaction (EAM)
  - Anisotropic interaction (Gay-Berne)
- **Efficiency notes**
  - Shift work load to GPU-accelerated real-space computation
    - Especially for K-space "Slab" Calculations
  - GPU acceleration benefit the most with heavy force computation (e.g. Gay-Berne, three-body)
  - Options to probe performance before production runs
    - Number of atoms per GPU (estimate)
    - Number of MPI tasks sharing a GPU
    - Number of threads per atom (based on the number of neighbors per atom)
    - Neighbor list build related parameters: skin, every/delay (zero dangerous builds)

New/Coming Features

- Support for Multi-Process Service (MPS) on NVIDIA Kepler GPUs to use Hyper-Q with MPI
- Acceleration for 3-body potentials
  - Currently available for Stillinger-Weber
  - mW Water Model (*Nature* **479**, 506–508) simulation is **6.6** faster for single-node and **>5X** for production simulations being performed on Titan
  - **Deterministic** Algorithms on Accelerator (Brown, Yamada, *CPC*, to appear.
- Improved support for OpenCL build
  - Vectorization for LAMMPS routines on Intel/AMD CPUs, Integrated Graphics, Discrete Graphics, Intel MIC

New/Coming Features

- Just-In-Time Potentials
  - Specify any equation in the input script
  - Compiled at run-time for the CPU or accelerator
    - No run-time parsing of equations is required
- Automated potential generation for coarse-grain models
- Monte Carlo
- Performance Improvements

LAMMPS at the OLCF

- **Molecular Dynamics of Thin Films** (Early Science - Nguyen, et al.)
  - First simulations of liquid crystal layers at length and time-scales sufficient for investigation of the evolution of spinodal instability
  - Simulations are using up to half of Titan
- **Rational Design of Organic Photovoltaics** (Early Science - Carrillo, et al.)
  - First simulations of OPV solar cells at experimental length/time scales
  - Collaboration between computational scientists, theorists, and experimentalists at ORNL and UT to improve the efficiency of solar devices
- **Accelerated Modeling of Non-icing Surfaces for Cold Climate Wind Turbines** (ALCC – Yamada, et al.)
  - GE-Global Research - passive coatings for ice mitigation
  - Large-scale simulation of water with 3-body potentials
- **You?**

Getting Access to Titan

- INCITE - Novel Computational Impact on Theory and Experiment Program
  - Typically awards millions of processor-hours per project
  - Addresses grand challenges in science and engineering
  - There is an annual call for INCITE proposals and awards are made on an annual basis
  - http://www.doeleadershipcomputing.org
- ALCC – The ASCR Leadership Computing Challenge
  - Emphasis on high-risk, high-payoff simulations
  - DOE energy mission areas
    - advancing the clean energy agenda
    - understanding the Earth's climate, …
  - open to national labs, academia, industry
  - http://science.energy.gov/ascr/facilities/alcc/
- DD – Director's Discretion projects
  - Dedicated to leadership computing preparation, INCITE and ALCC scaling, and application performance to maximize scientific application efficiency and productivity
  - https://www.olcf.ornl.gov/support/getting-started/olcf-director-discretion-project-application/

Titan Specs (2012)	
Compute Nodes	18,688
Login & I/O Nodes	512
Memory per node	32 GB + 6 GB
# of Opteron cores	299,008
# of NVIDIA K20 "Kepler" accelerators (2013)	18,688
Total System Memory	710 TB
Total System Peak Performance	20+ Petaflops

Acknowledgements

Current LAMMPS Accelerator Library Developers:

Mike Brown ORNL brownw domain: ornl.gov	Trung Nguyen ORNL nguyentd domain: ornl.gov	Steve Plimpton SNL sjplimp domain: sandia.gov	Axel Kohlmeyer ICTP akohlmey domain: gmail.com

Department of Energy, Office of Science under Contract No. DE-AC05-00OR22725
OLCF Center for Accelerated Application Readiness (CAAR), OLCF-3 project

Thanks for attention-
Questions and comments?