HOOMD-blue: Scalable Molecular Dynamics on 1000's of GPUs

Jens Glaser, Joshua A. Anderson
Research Associate, Glotzer Group
Chemical Engineering, University of Michigan

XXVI IUPAP Conference on Computational Physics CCP2014, Boston, 08/11/14

Outline

1. Introduction

2. HOOMD-blue capabilities

3. Strong Scaling

HOOMD-blue overview

- **H**ighly-**O**ptimized, **O**bject-oriented **M**any-particle **D**ynamics, blue edition

- Open-source **GPU**-accelerated, written in CUDA/C++/Python

- Developed at UofM by Joshua Anderson (lead-developer) since March 2007 / CUDA 0.8 beta

- Initial 0.60 release February 2008

  J. A. Anderson, C. D. Lorenz, A. Travesset
  General purpose molecular dynamics simulations fully implemented on graphics processing units, Journal of Computational Physics 227 (2008), 5342–5359.

- April 2014: HOOMD-blue **1.0** - multi-GPU

Applications of HOOMD-blue

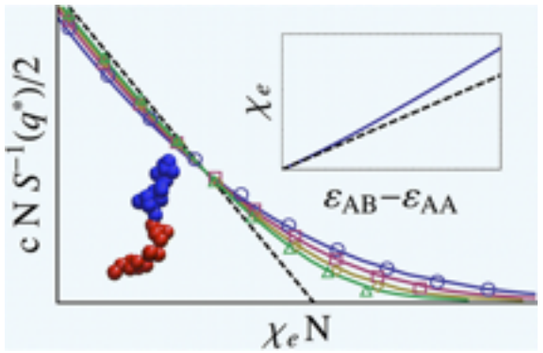

Glaser et al. Macromolecules 2014

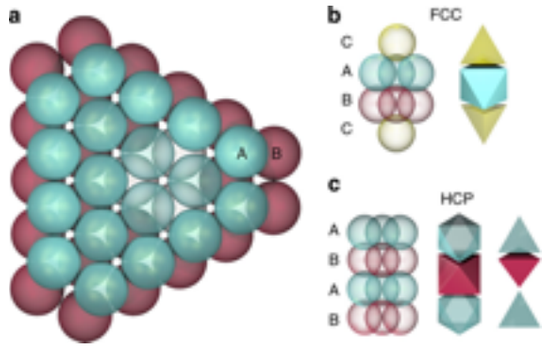

Mahynsk, A. Nat. Comm. 2014

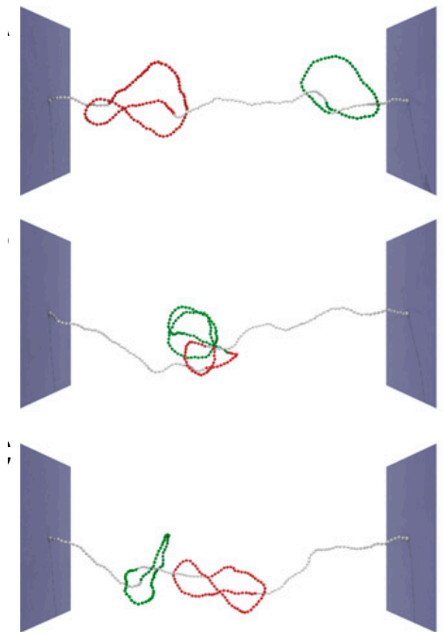

Trefz, B. et al. PNAS 2014

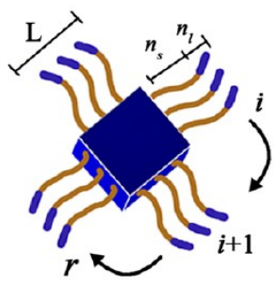

Knorowski, C. and Travesset, A. JACS 2014

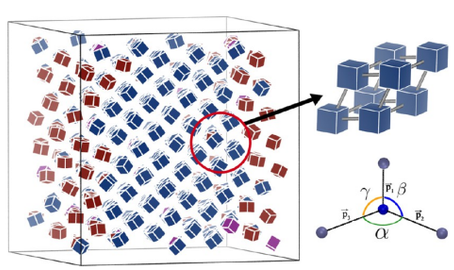

Beltran-Villegas et al. Soft Matter 2014

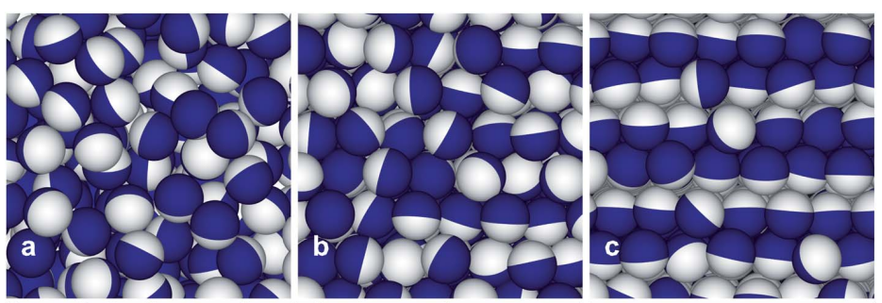

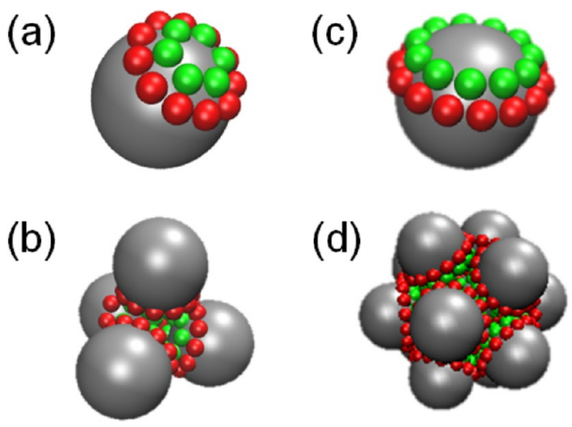

Long, A.W. and Ferguson, A.L. J. Phys. Chem. B 2014

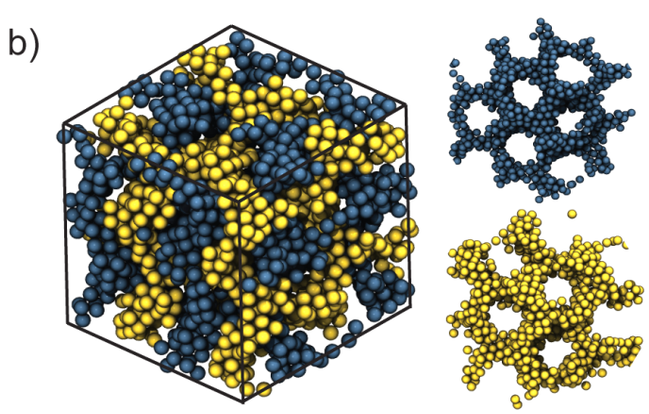

Marson, R. L. et al. Nano Lett. 2014

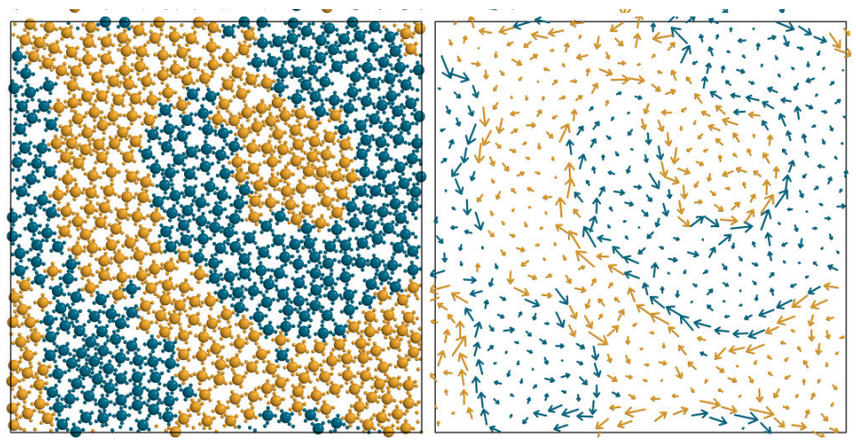

Nguyen et al. Phys Rev. Lett. 2014

75+ peer-reviewed publications using HOOMD-blue as of August 2014
>300 citations of 2008 paper (ISI)

.. in the Glotzer group

Molecular Dynamics

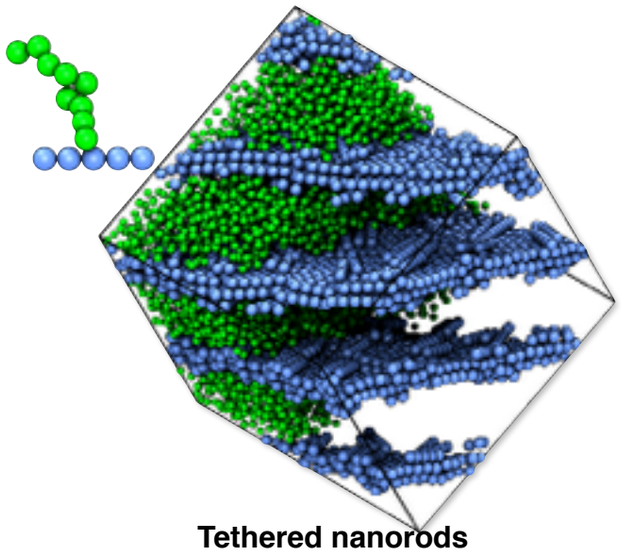

Tethered nanorods
Rigid-body dynamics
Nguyen, T. D. et al., *ACS Nano* **4**, 2585-94 (2010)

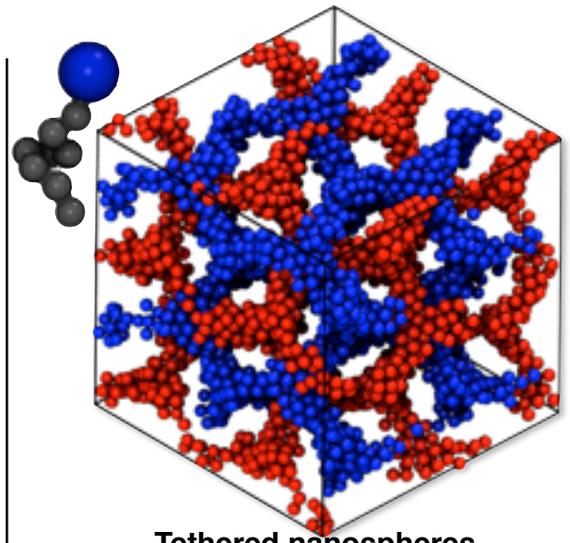

Tethered nanospheres
Langevin dynamics
Phillips, C. L. et al., *Soft Matter* **6**, 1693 (2010)

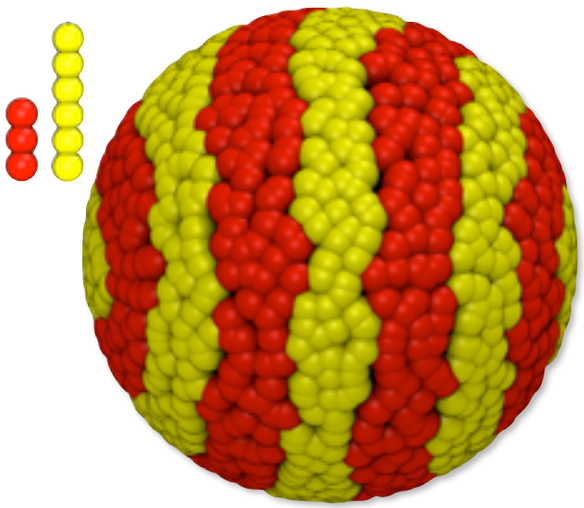

Surfactant coated surfaces
Dissipative particle dynamics
Pons-Siepermann, I. C., *Soft* matter **6** 3919 (2012)

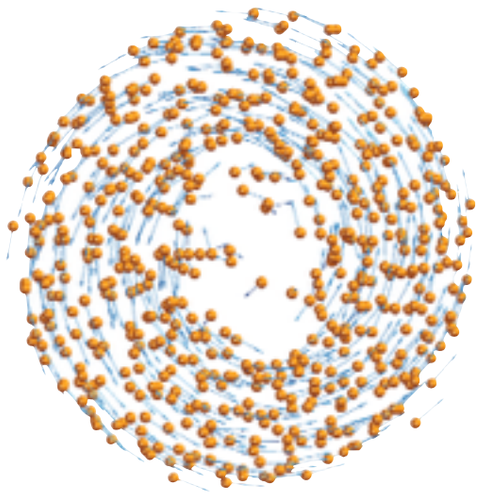

Self-propelled colloids
Non-equilibrium MD
Nguyen N., *Phys Rev E* **86** 1, 2012

Monte Carlo

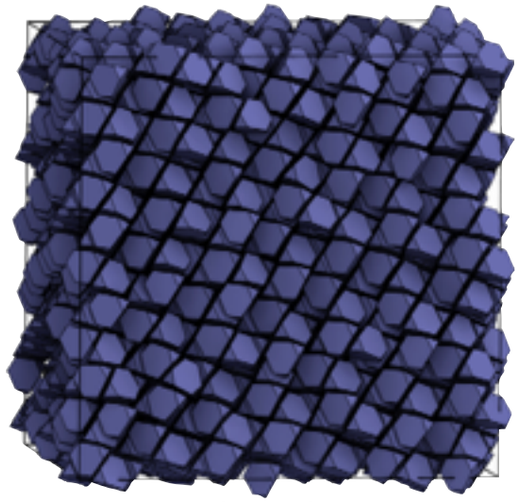

Truncated Tetrahedra
Hard particle MC
Damasceno, P. F. et al., *ACS Nano* **6**, 609 (2012)

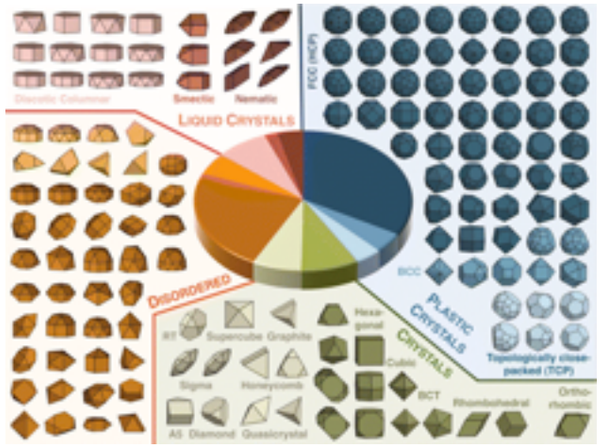

Arbitrary polyhedra
Hard particle MC
Damasceno, P. F. et al., *Science* **337**, 453 (2012)

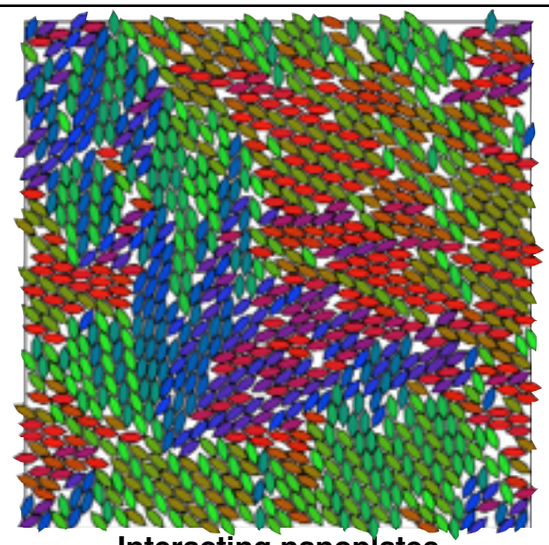

Interacting nanoplates
Hard particle MC with interactions
Ye X. et al., *Nature Chemistry* **cover article** (2013)

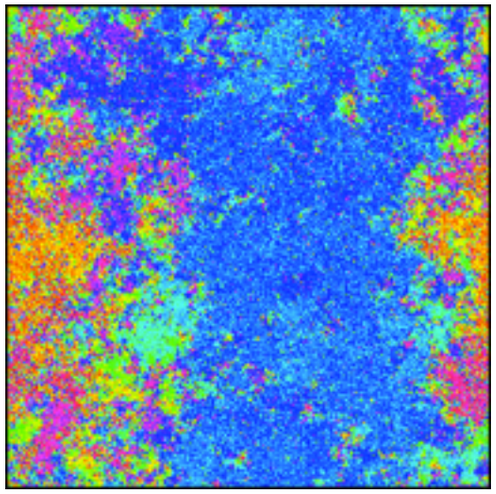

Hard disks - hexatic
Hard particle MC
Engel M. et al., PRE **87**, 042134 (2013)

Diblock Copolymers on Multi-GPU

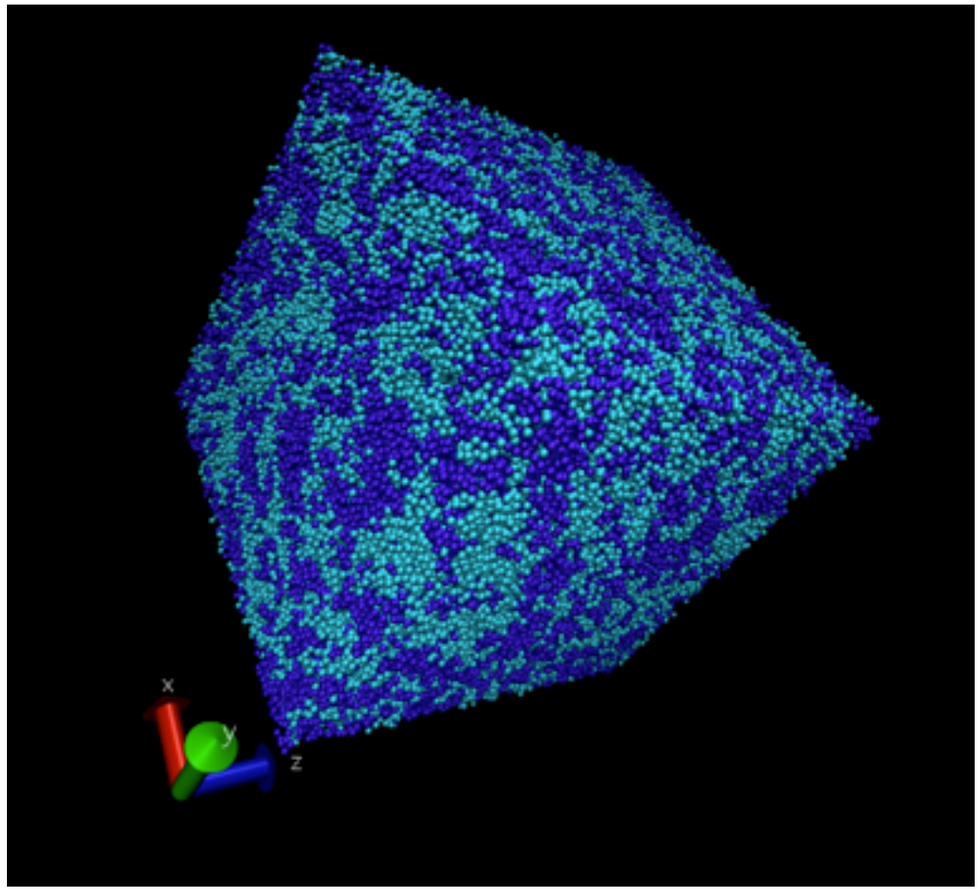

Diblock copolymer melts

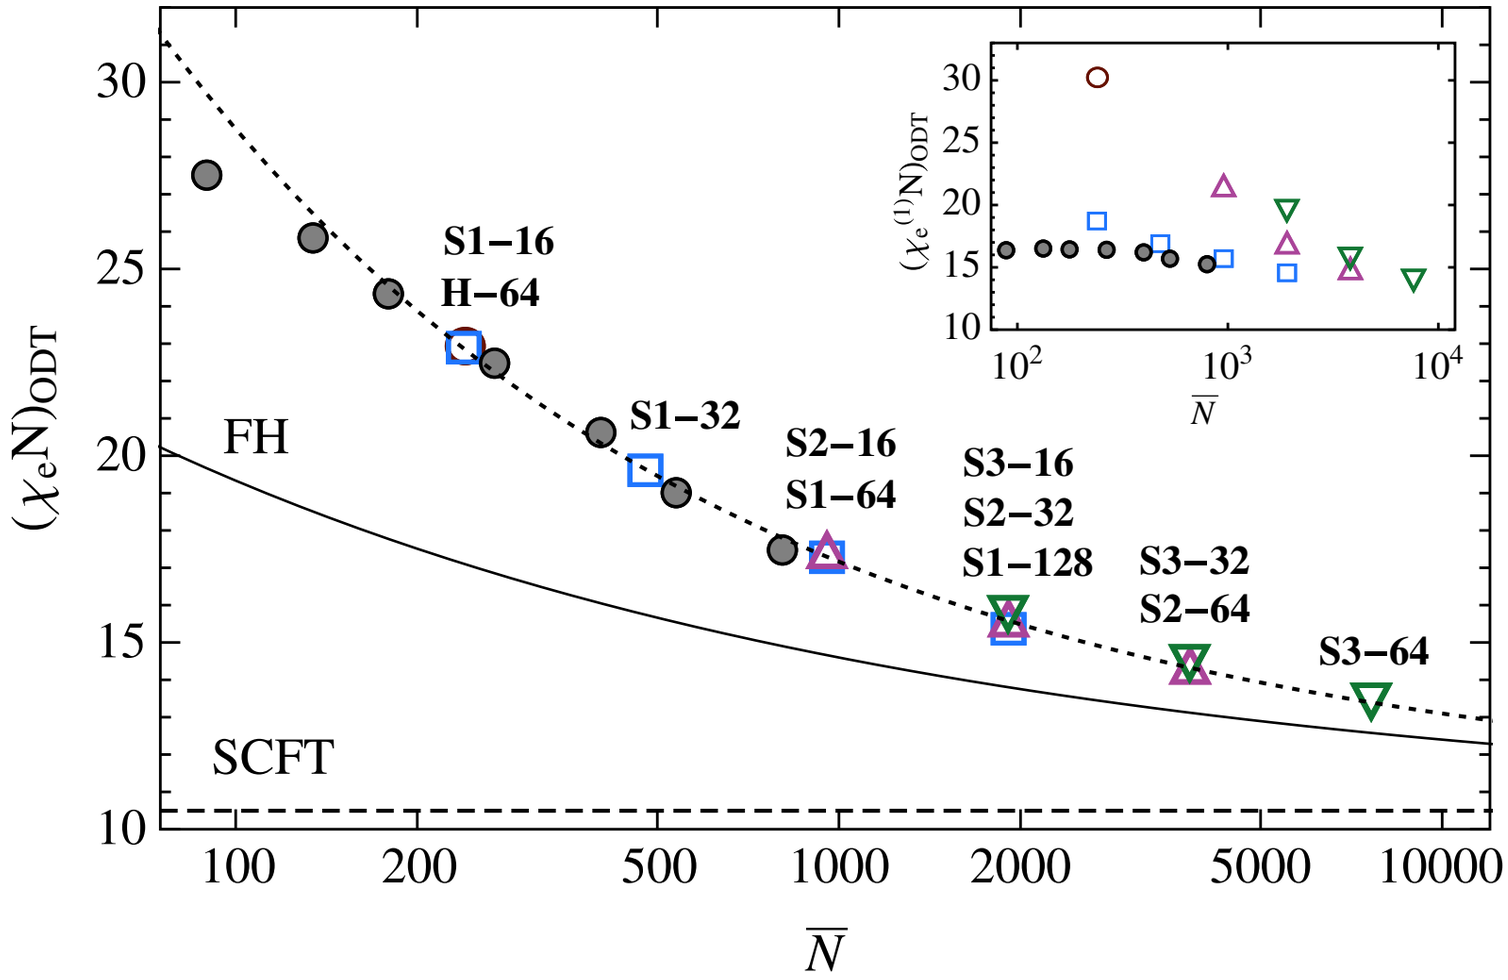

Glaser, J., Medapuram, P., Beardsley, T. M., Matsen, M. W., & Morse, D. C. Universality of Block Copolymer Melts. *PRL*, **113**, 068302 (2014)

Virus capsid particles

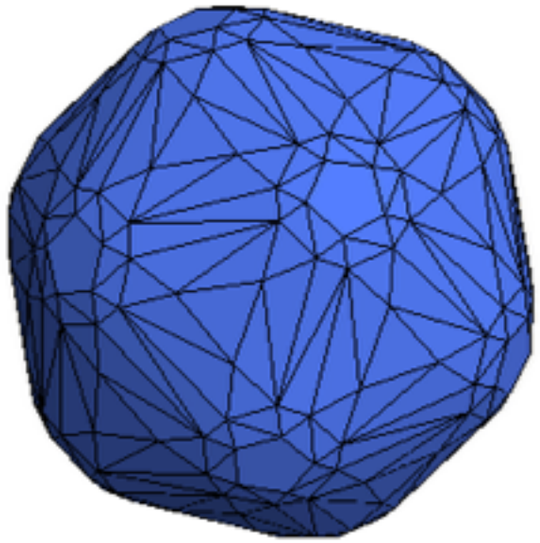

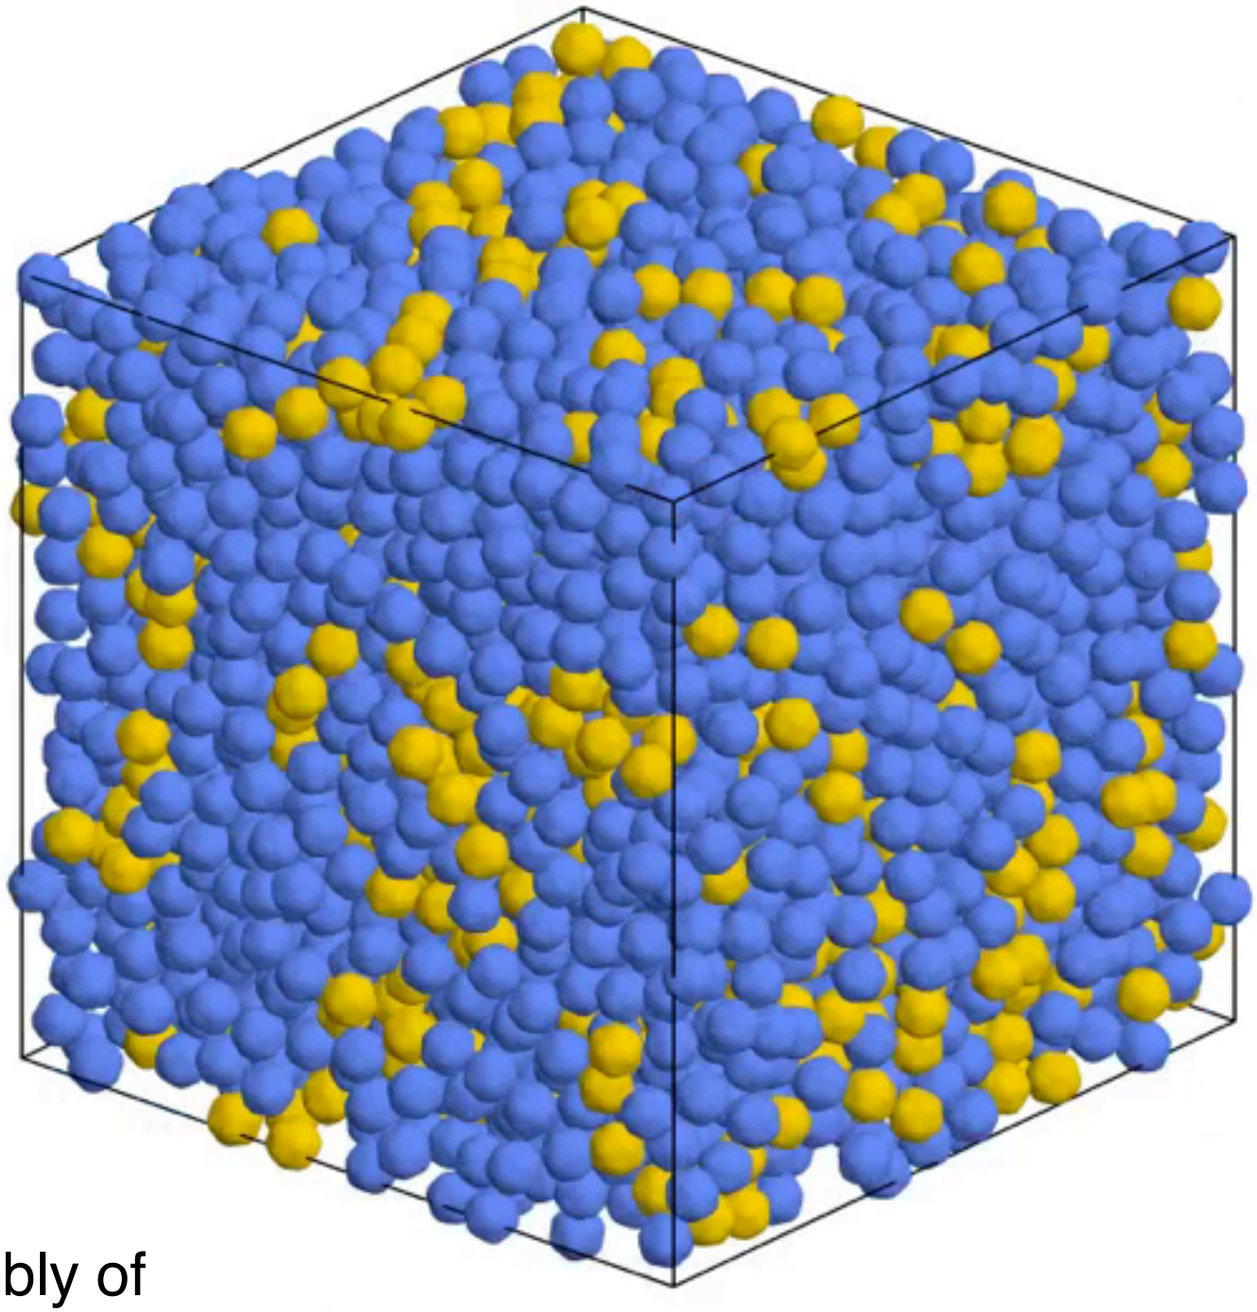

Crystalline self-assembly of
4096 hard convex polyhedra
(virus capsids)

GPU Molecular Dynamics in 1 slide

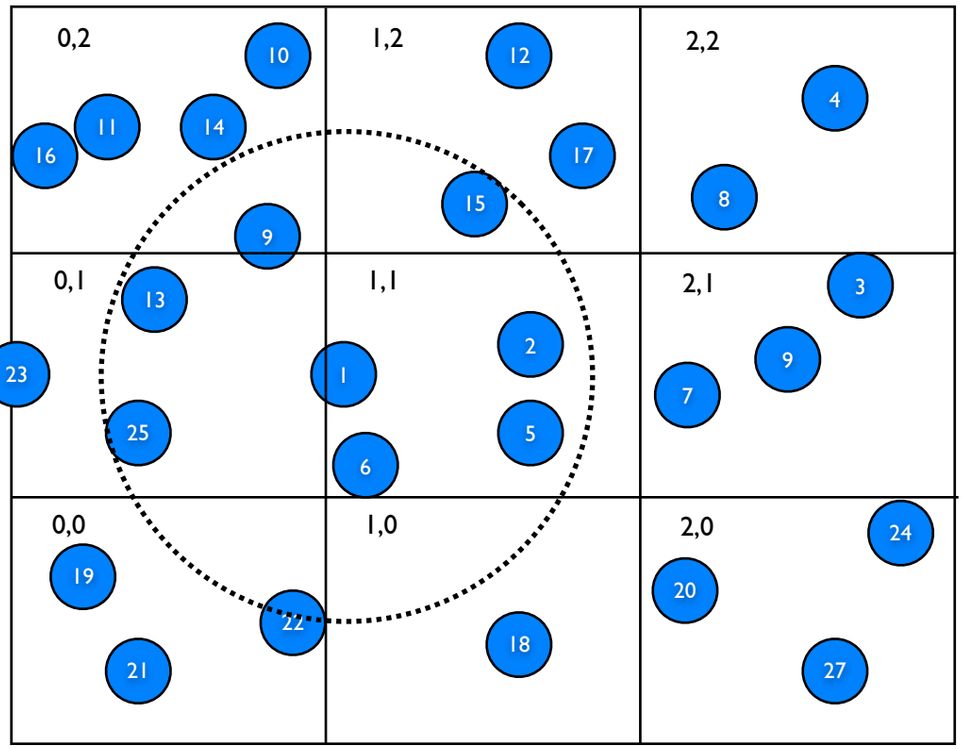

Cell list

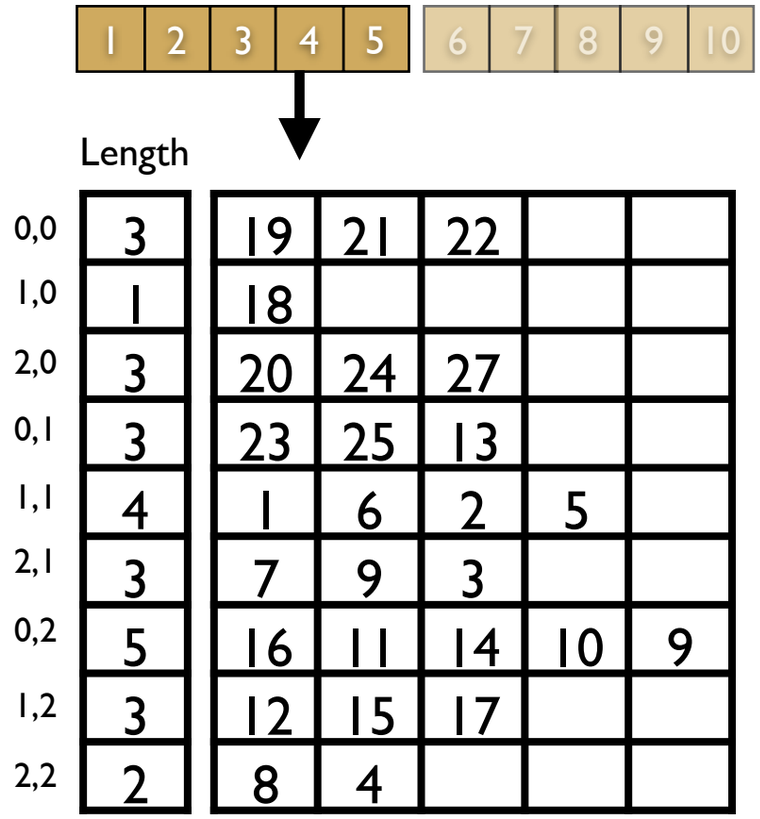

Neighbor list

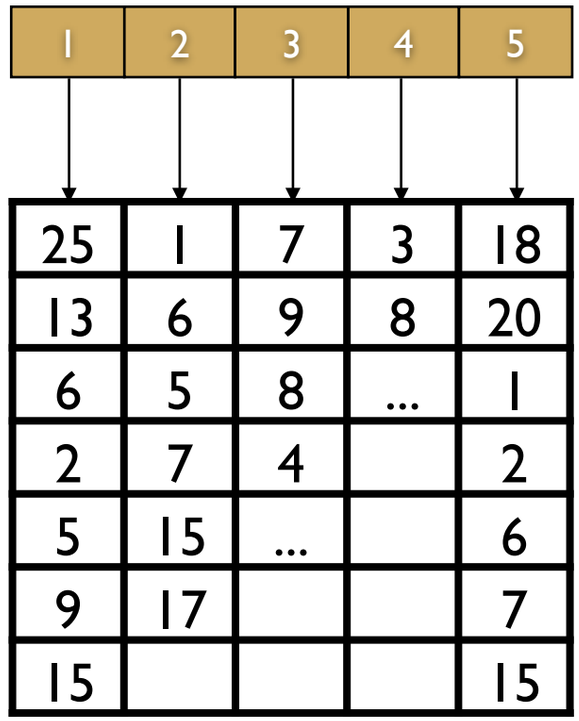

Neighbor list

Num Neighbors

Pair force

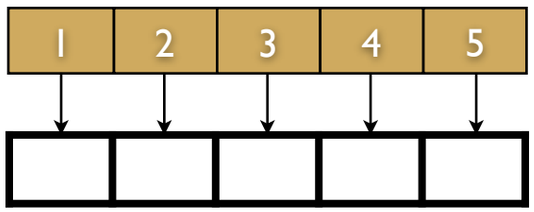

Integrate

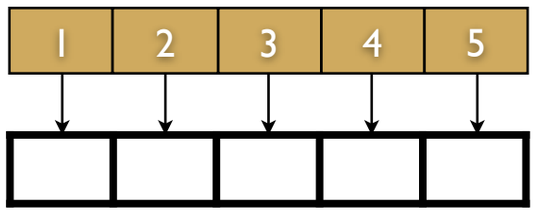

1 thread/particle

But this is all *you* need to write.

```
from hoomd_script import *

init.read_xml(filename='init.xml')

lj = pair.lj(r_cut=2.5)
lj.pair_coeff.set('A', 'A', epsilon=1.0,
                            sigma=1.0)

integrate.mode_standard(dt=0.005)
integrate.nvt(T=1.2, tau=0.5)

run(10e3)
```

1 4 X

one K40 GPU vs one 10c Xeon E5-2680v2

100x over a single CPU core

1/7 the power consumption

5x perf/$

Features in HOOMD-blue v1.0

Integration
- NVT (Nosé-Hoover)
- NPT
- NPH
- Brownian Dynamics
- Dissipative Particle Dynamics
- NVE
- FIRE energy minimization
- Rigid body dynamics

Snapshot formats
- MOL2
- DCD
- PDB
- XML

Simulation types
- 2D and 3D
- Triclinic box
- Replica exchange (via script)

Pair forces
- Lennard Jones
- Gaussian
- CGCMM
- Morse
- Table
- Yukawa
- PPPM electrostatics
- Orientation-Averaged Ewald (code not publicly available)
  Jha, P. K. et al., Journal of Chemical Theory and Computation 6, 3058-3065 (2010)

Many-body forces
- EAM

Bond forces
- Harmonic
- FENE
- Table

Angle forces
- Harmonic
- CGCMM
- Table

Dihedral/Improper forces
- Harmonic
- Table

Hardware support
- All recent NVIDIA GPUs
- **Multi-GPU with MPI**
- Multi-CPU with MPI

Next up, Hard particle Monte Carlo

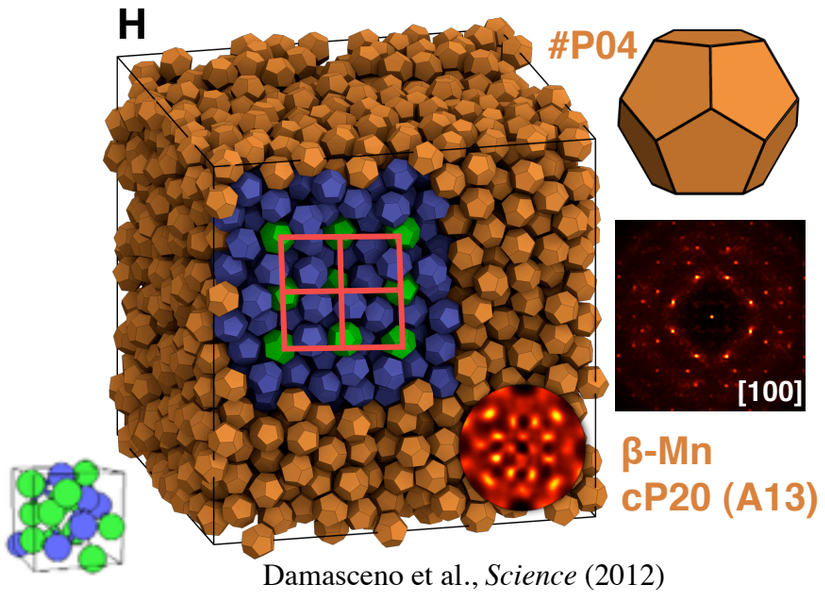

Damasceno et al., *Science* (2012)

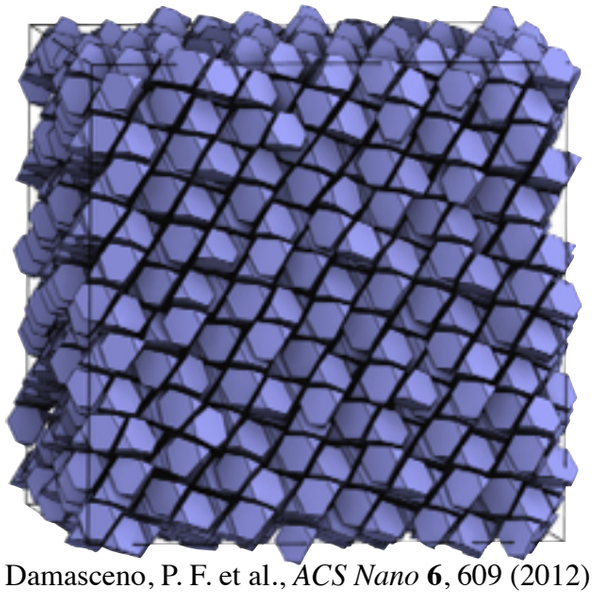

Damasceno, P. F. et al., *ACS Nano* **6**, 609 (2012)

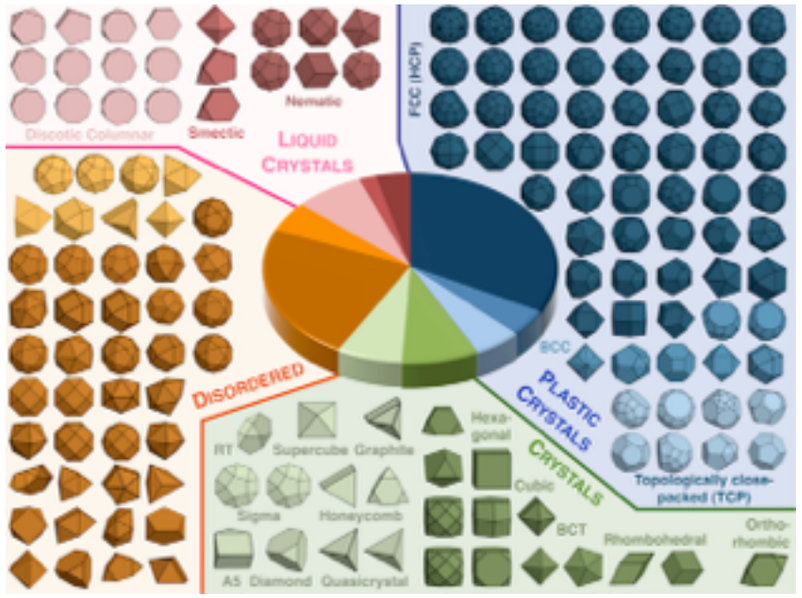

Damasceno et al., *Science* (2012)

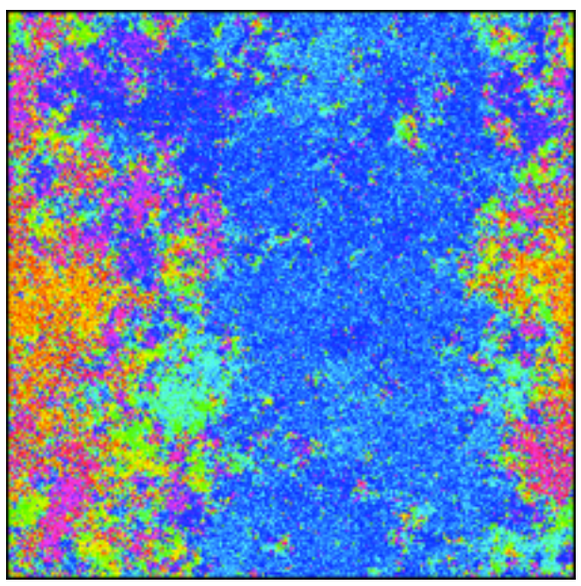

Engel M. et al., PRE **87**, 042134 (2013)

- Hard Particle Monte Carlo plugin for HOOMD-blue
- 2D Shapes
  - Disk
  - Convex (Sphero)polygon
  - Concave polygon
  - Ellipse
- 3D Shapes
  - Sphere
  - Ellipsoid
  - Convex (Sphero)polyhedon
- NVT and NPT ensembles
- Frenkel-Ladd free energy
- Parallel execution on a single GPU
- Domain decomposition across multiple nodes (CPUs or GPUs)

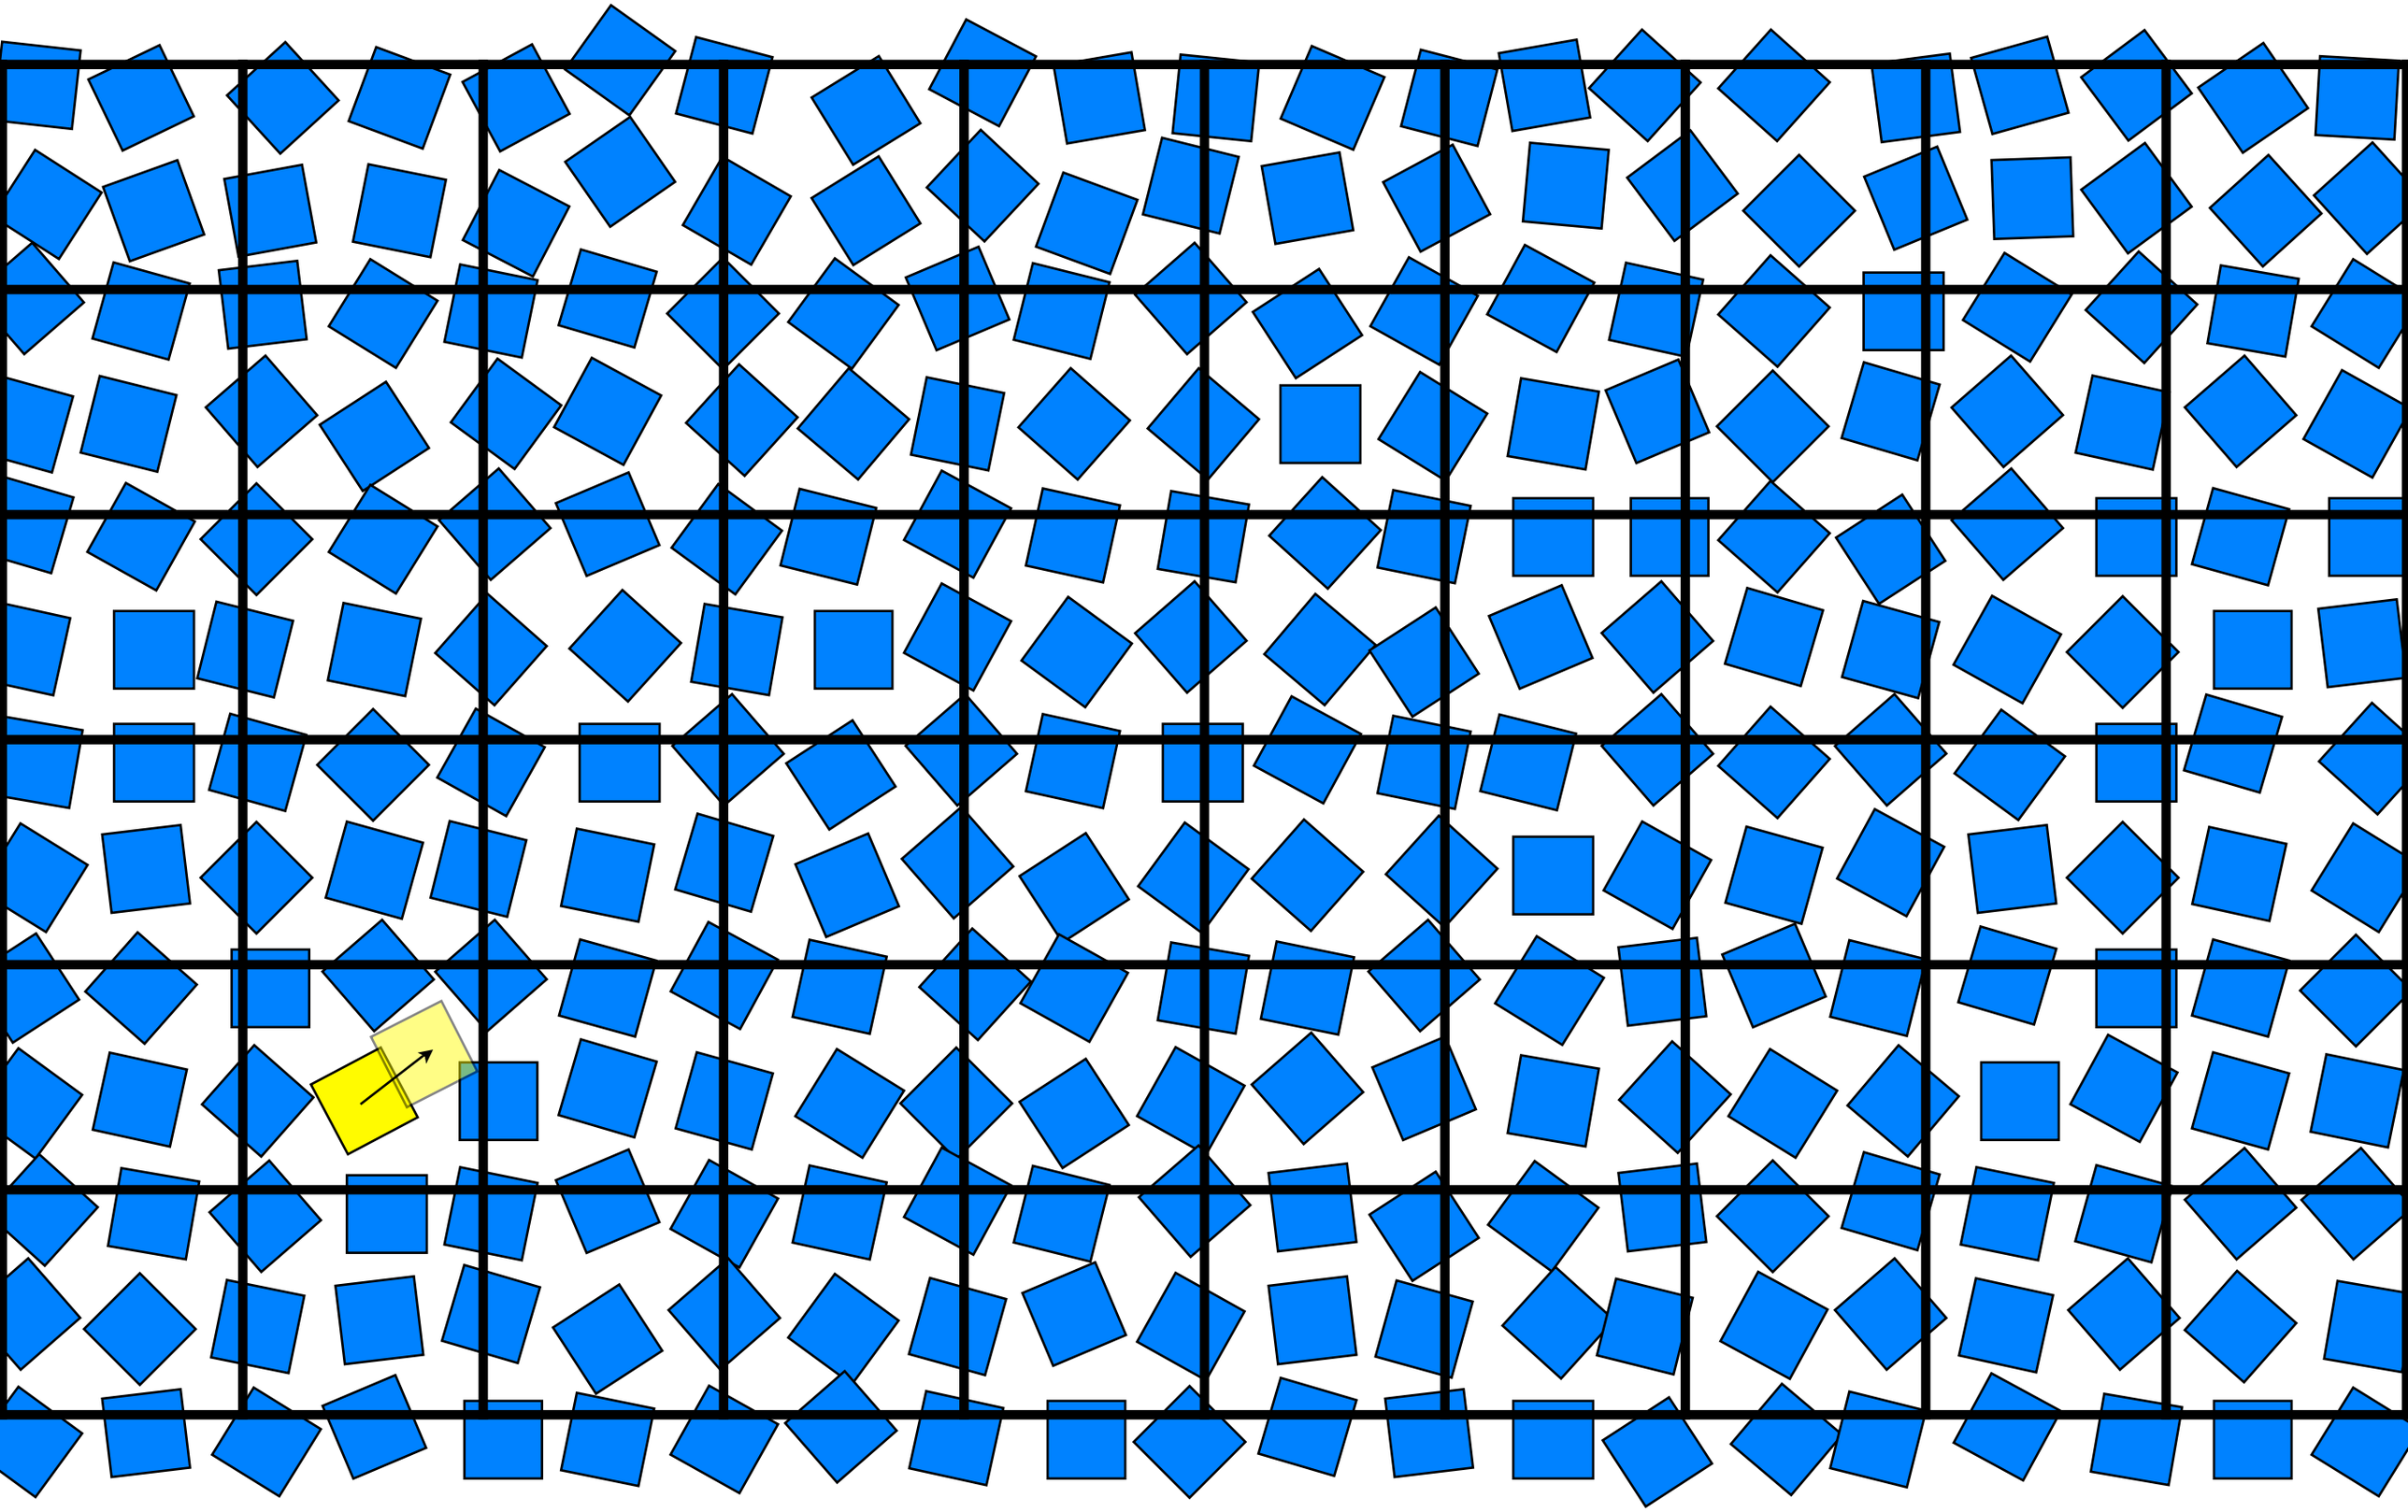

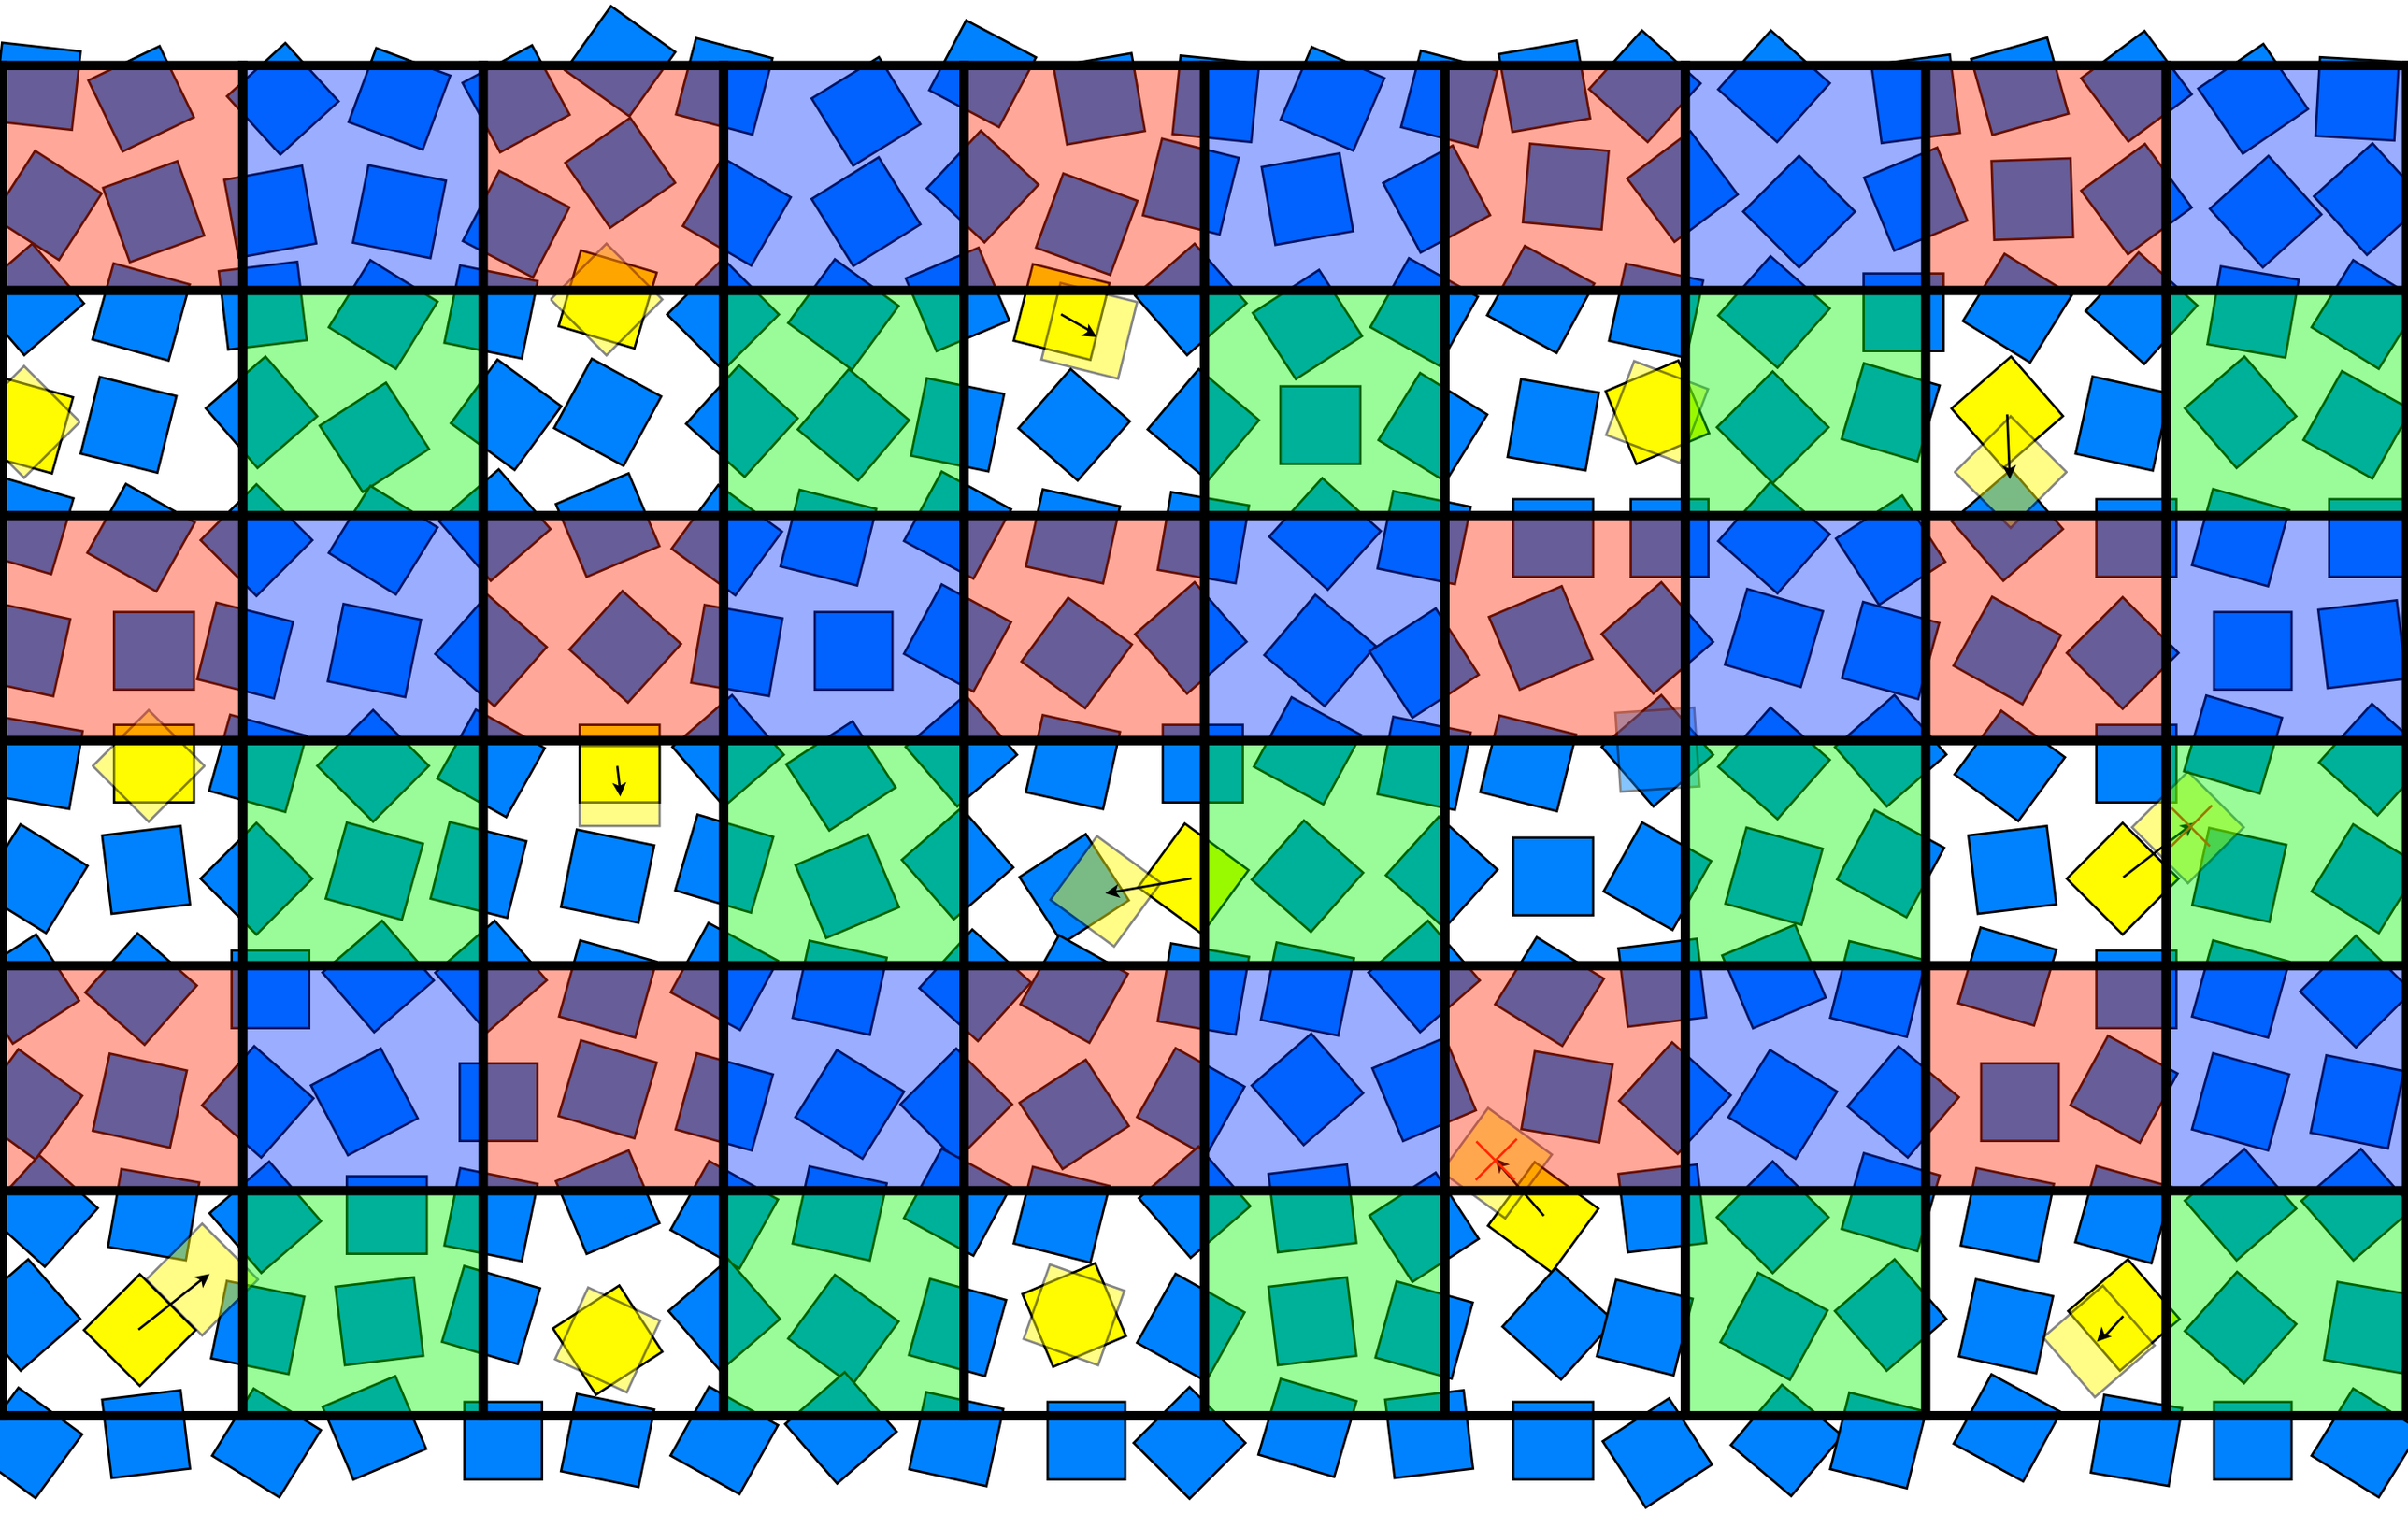

Overlap checks

Separating planes

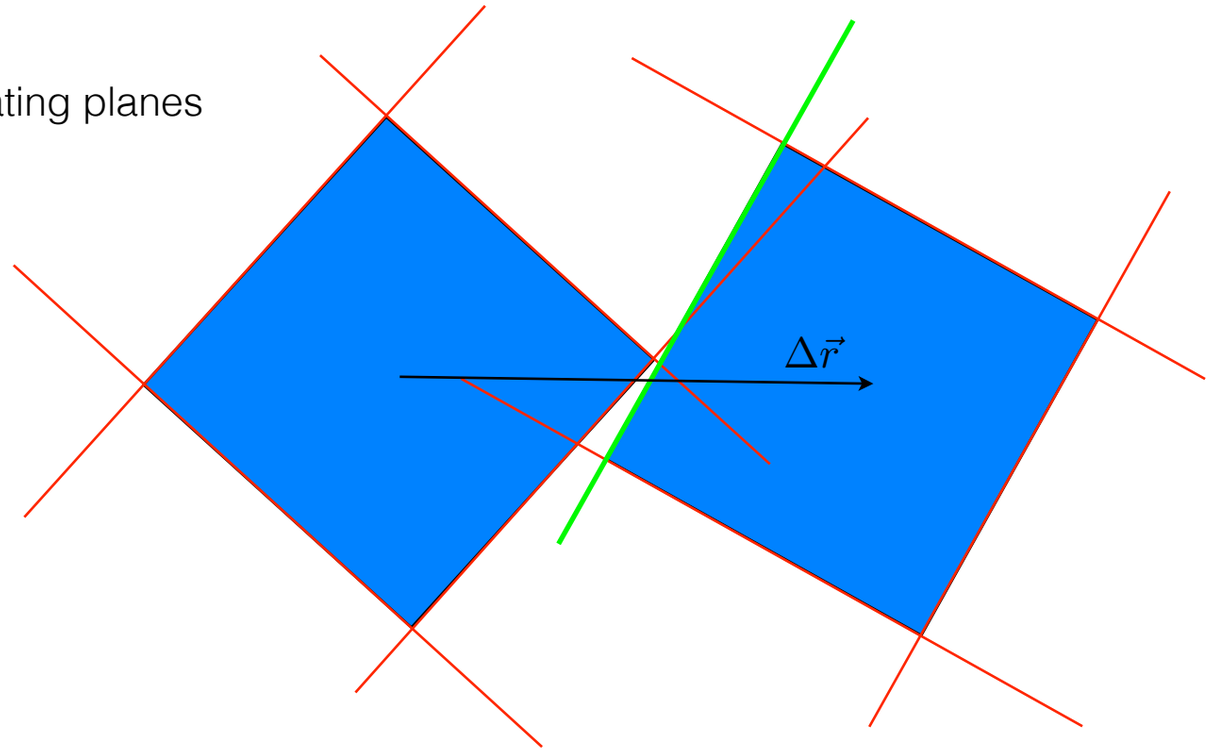

- Disk/sphere - trivial
- Convex polygons - separating axis
- Concave polygons - brute force
- Spheropolygons - XenoCollide/GJK
- Convex polyhedra - XenoCollide/GJK
- Ellipsoid / Ellipse: Matrix method
- Compute delta in double, convert to single for expensive overlap check

1001.842 - 1000.967 = 0.875

XenoCollide/GJK

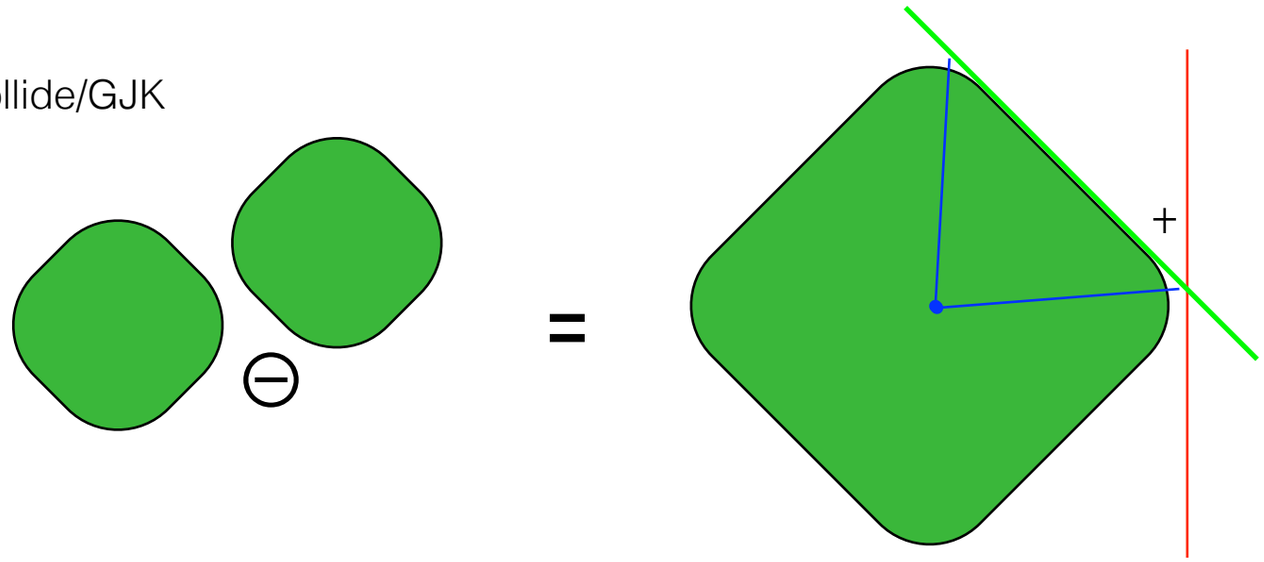

Easy and flexible to use

```
from hoomd_script import *
from hoomd_plugins import hpmc

init.read_xml(filename='init.xml')

mc = hpmc.integrate.convex_polygon(seed=10, d=0.25, a=0.3);
mc.shape_param.set('A', vertices=[(-0.5, -0.5), (0.5, -0.5),
                                  (0.5, 0.5), (-0.5, 0.5)]);

run(10e3)
```

4x

one K40 GPU vs one 10c Xeon E5-2680v2

40x over a single CPU core

1/2 the power consumption

1.4x perf/$

MC performance scaling

GPU: Tesla K20X, CPU: Xeon E5-2680 (XSEDE Stampede)

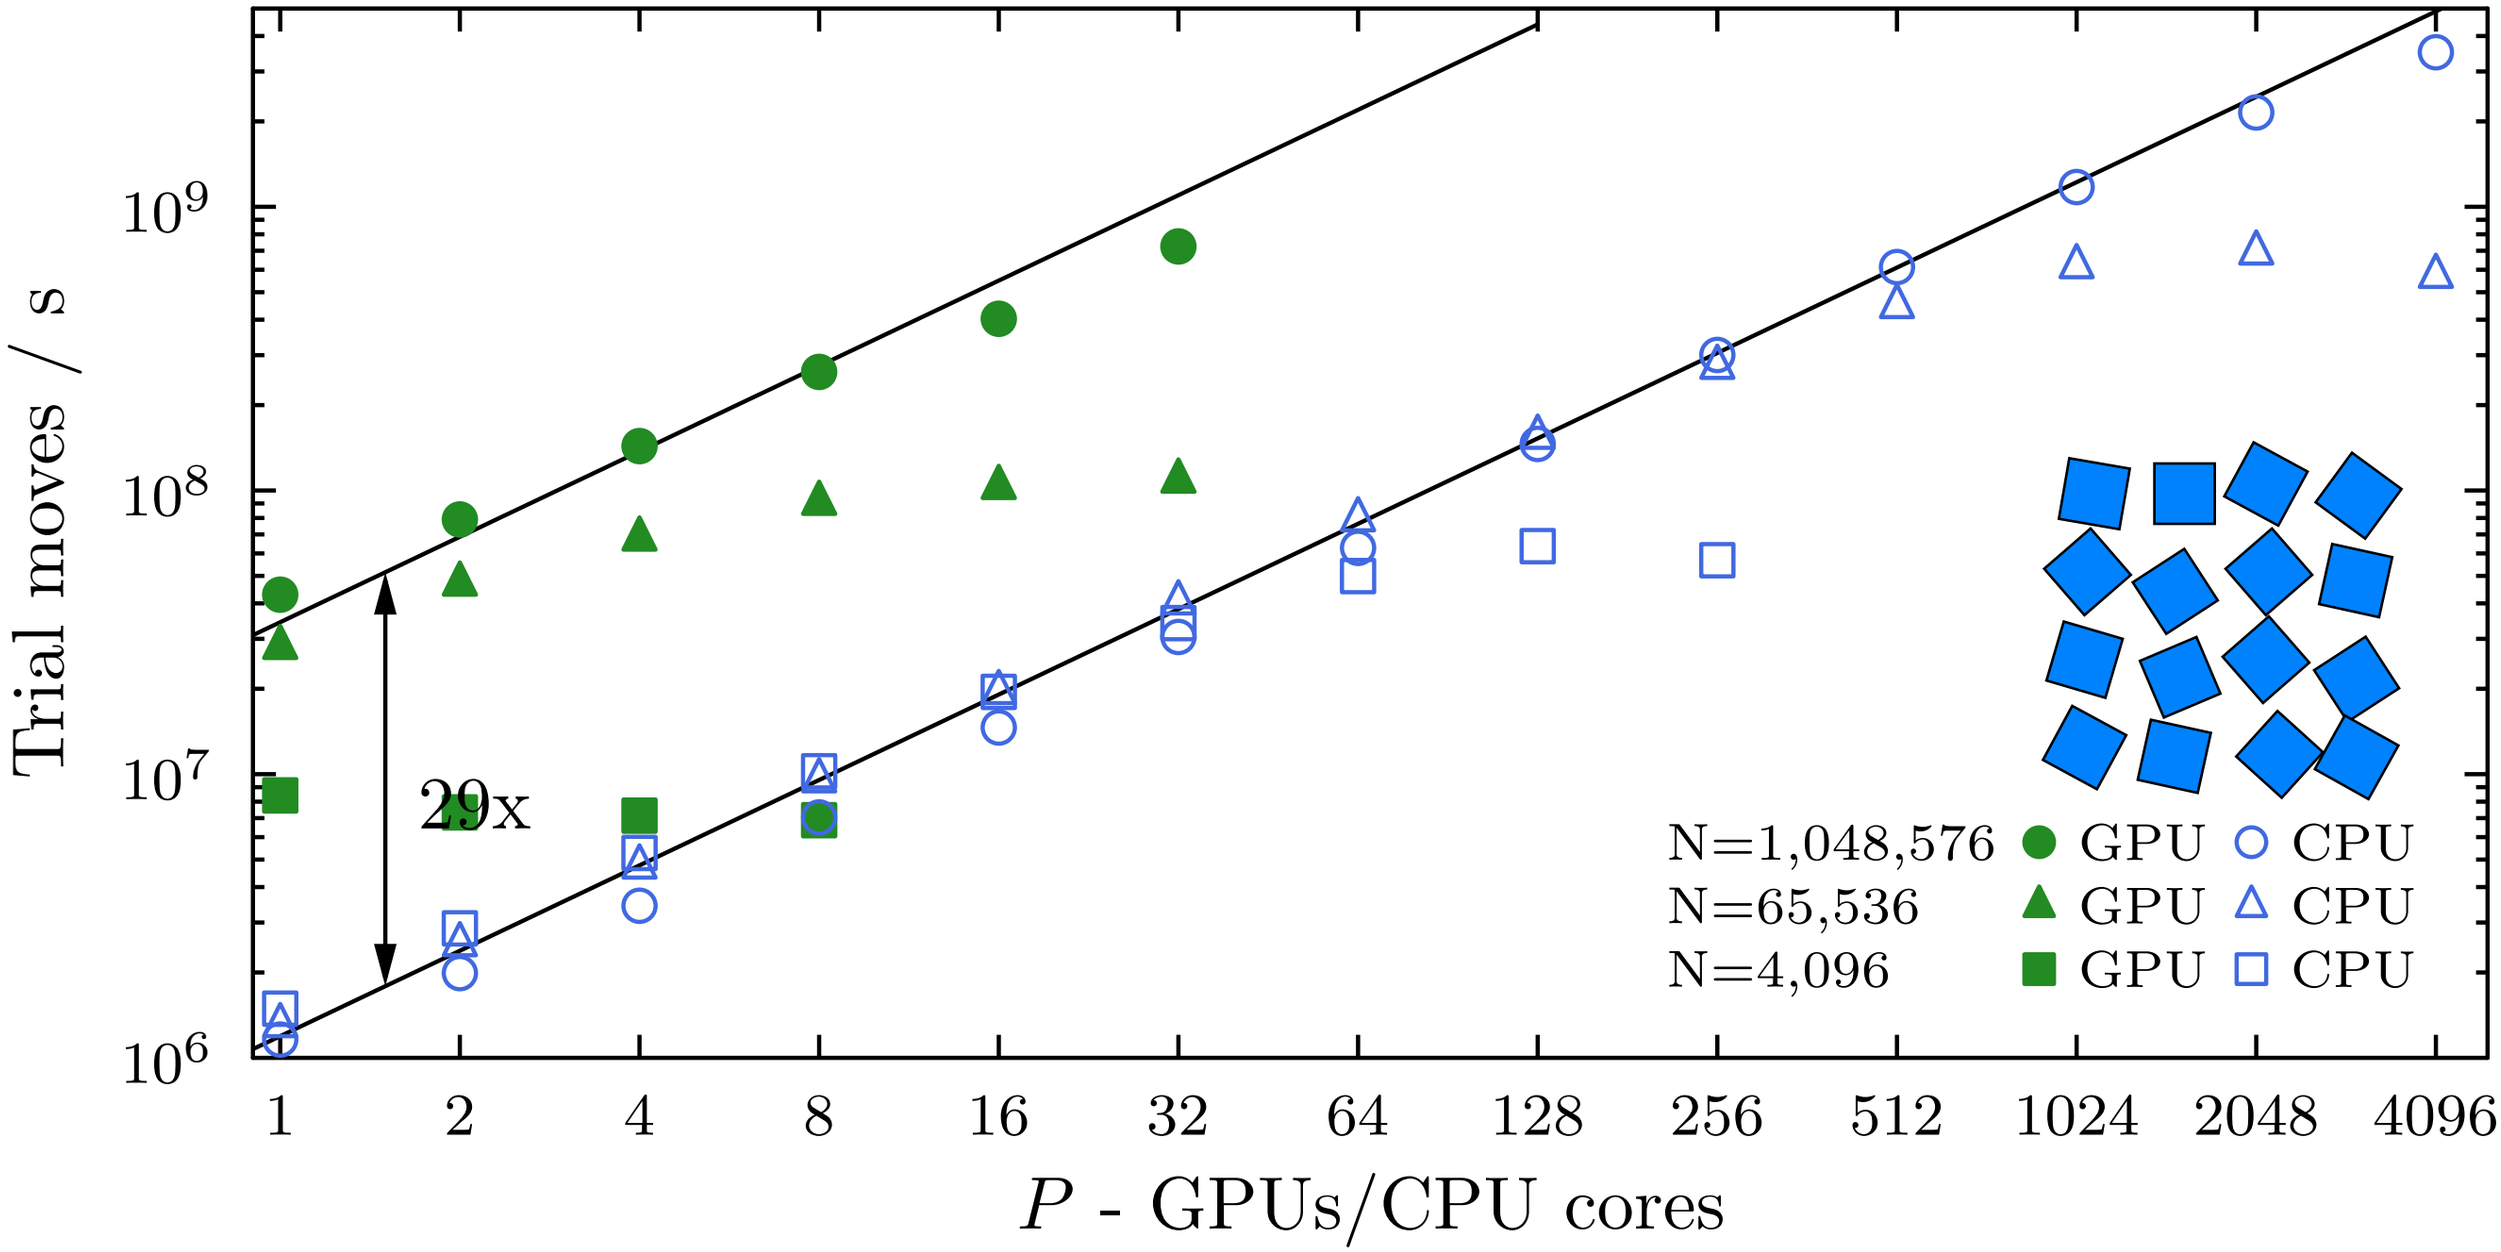

Summary capabilities

- Fast, versatile, easy-to-use Molecular Dynamics on GPUs and CPUs

- Open-source, download from **http://codeblue.umich.edu/hoomd-blue**

- HOOMD 2.0: Hard particle Monte Carlo

Advent of GPUs in HPC

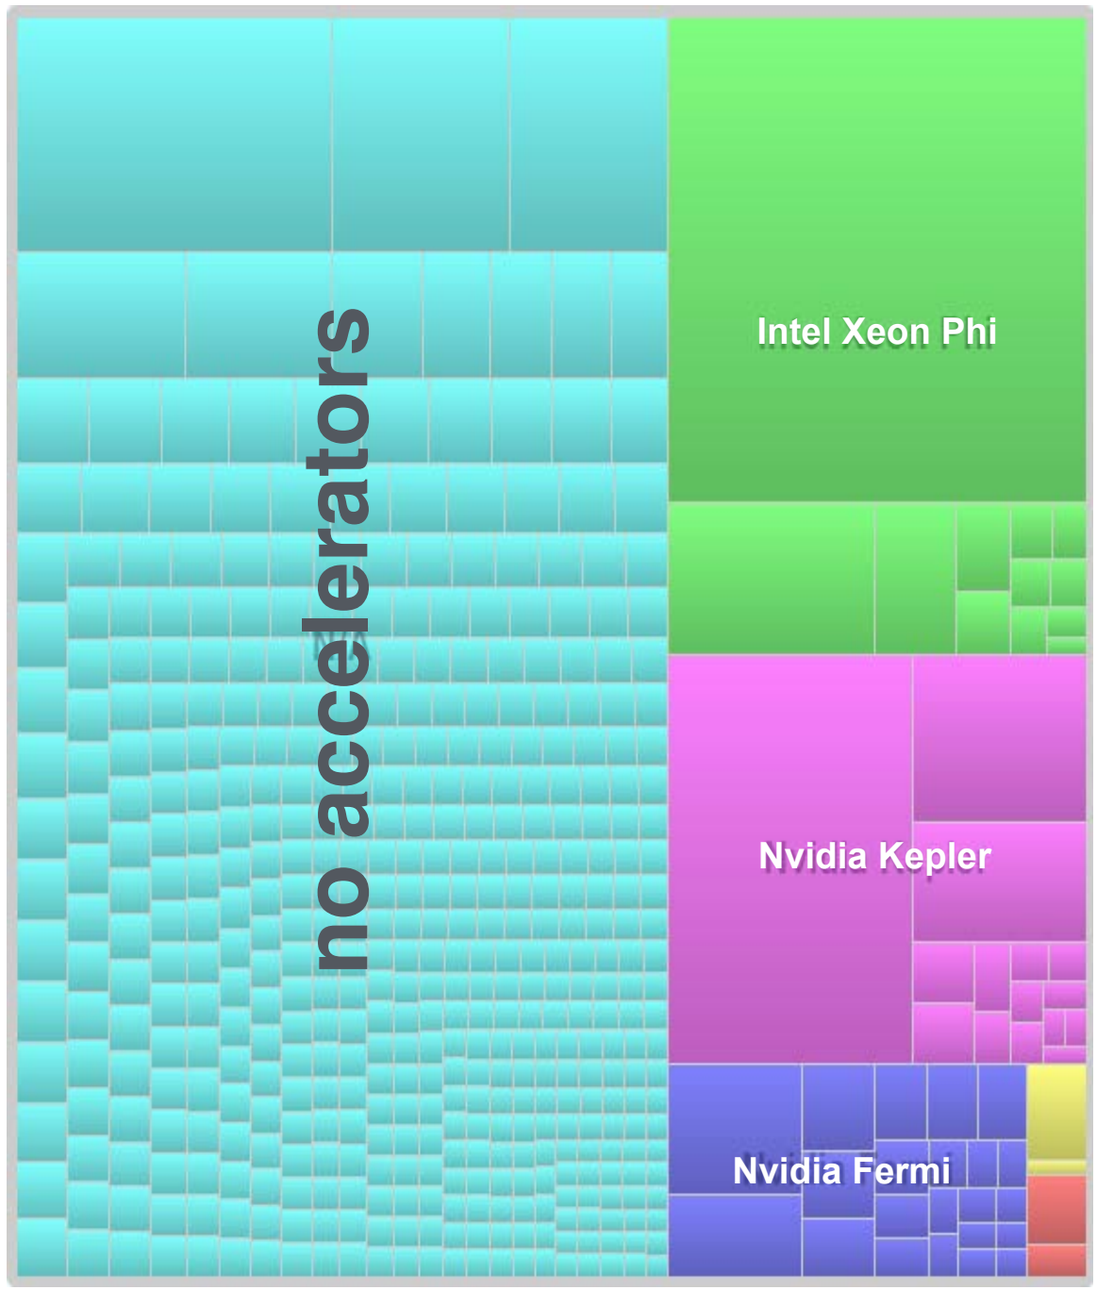

Performance share of accelerators in TOP500 supercomputers

18.1%

16.2%

TOP500 Nov ‘13
slide courtesy of D.K. Panda

Titan (#2 TOP500) at Oak Ridge National Laboratories

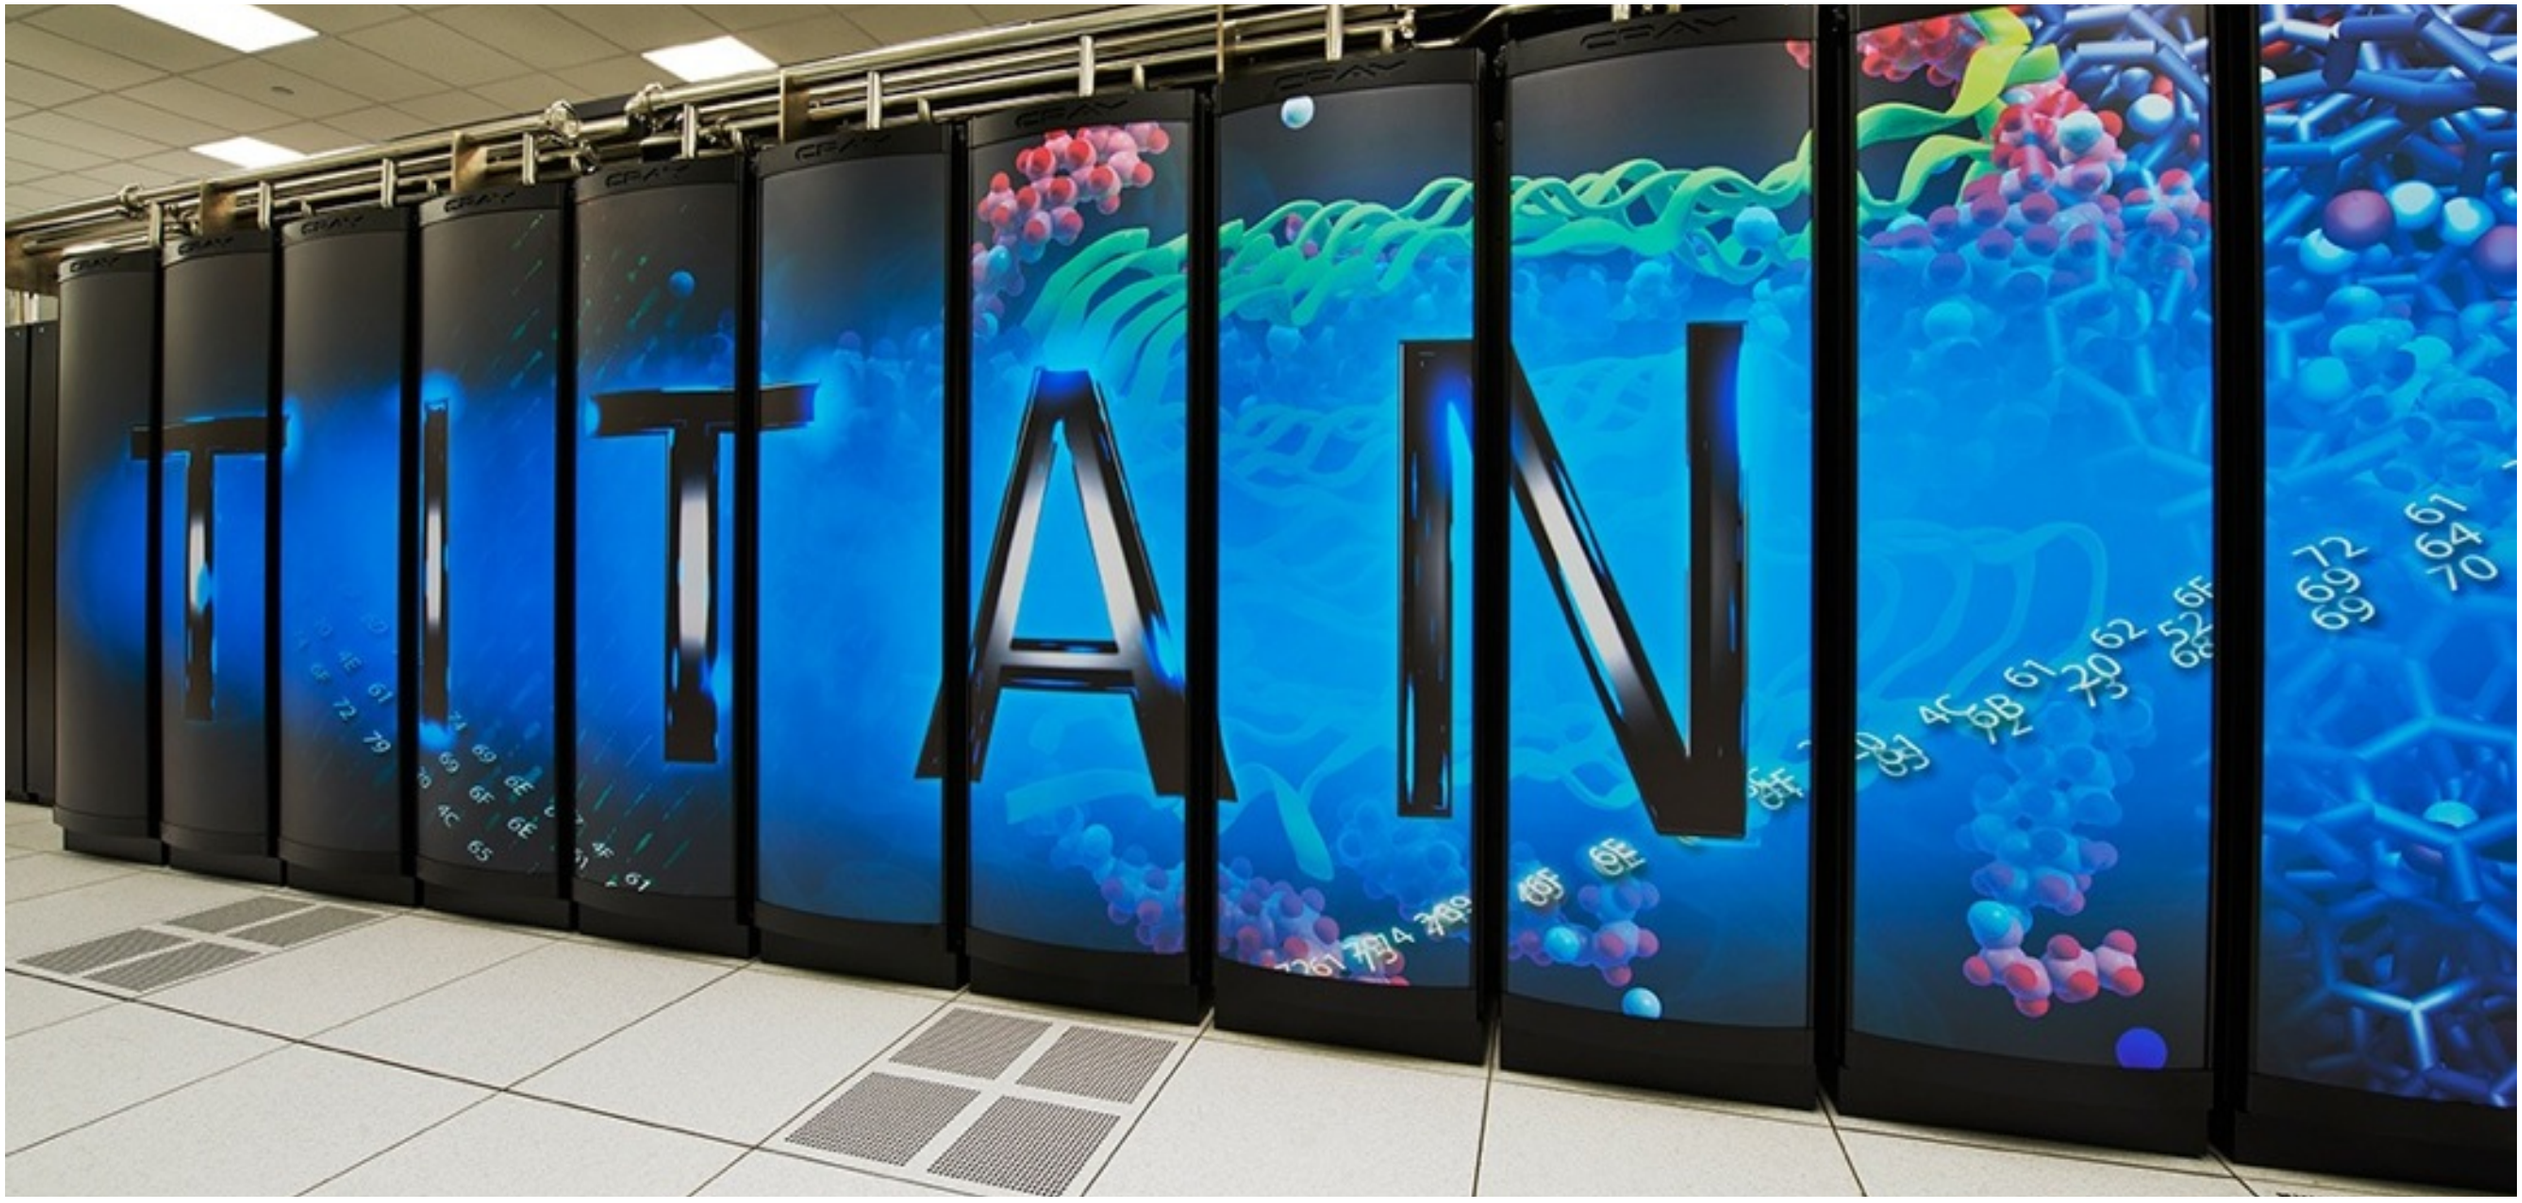

Cray XK7, 18,688 compute nodes 512 XIO service nodes
3D-torus Cray Gemini interconnect
Node configuration: 16-core 2.2 GHz AMD Opteron 6274, 32 GB ECC DDR3 SDRAM and NVIDIA Kepler K20X card connected via PCI-e2.0

Scaling bottlenecks in spatial domain decomposition

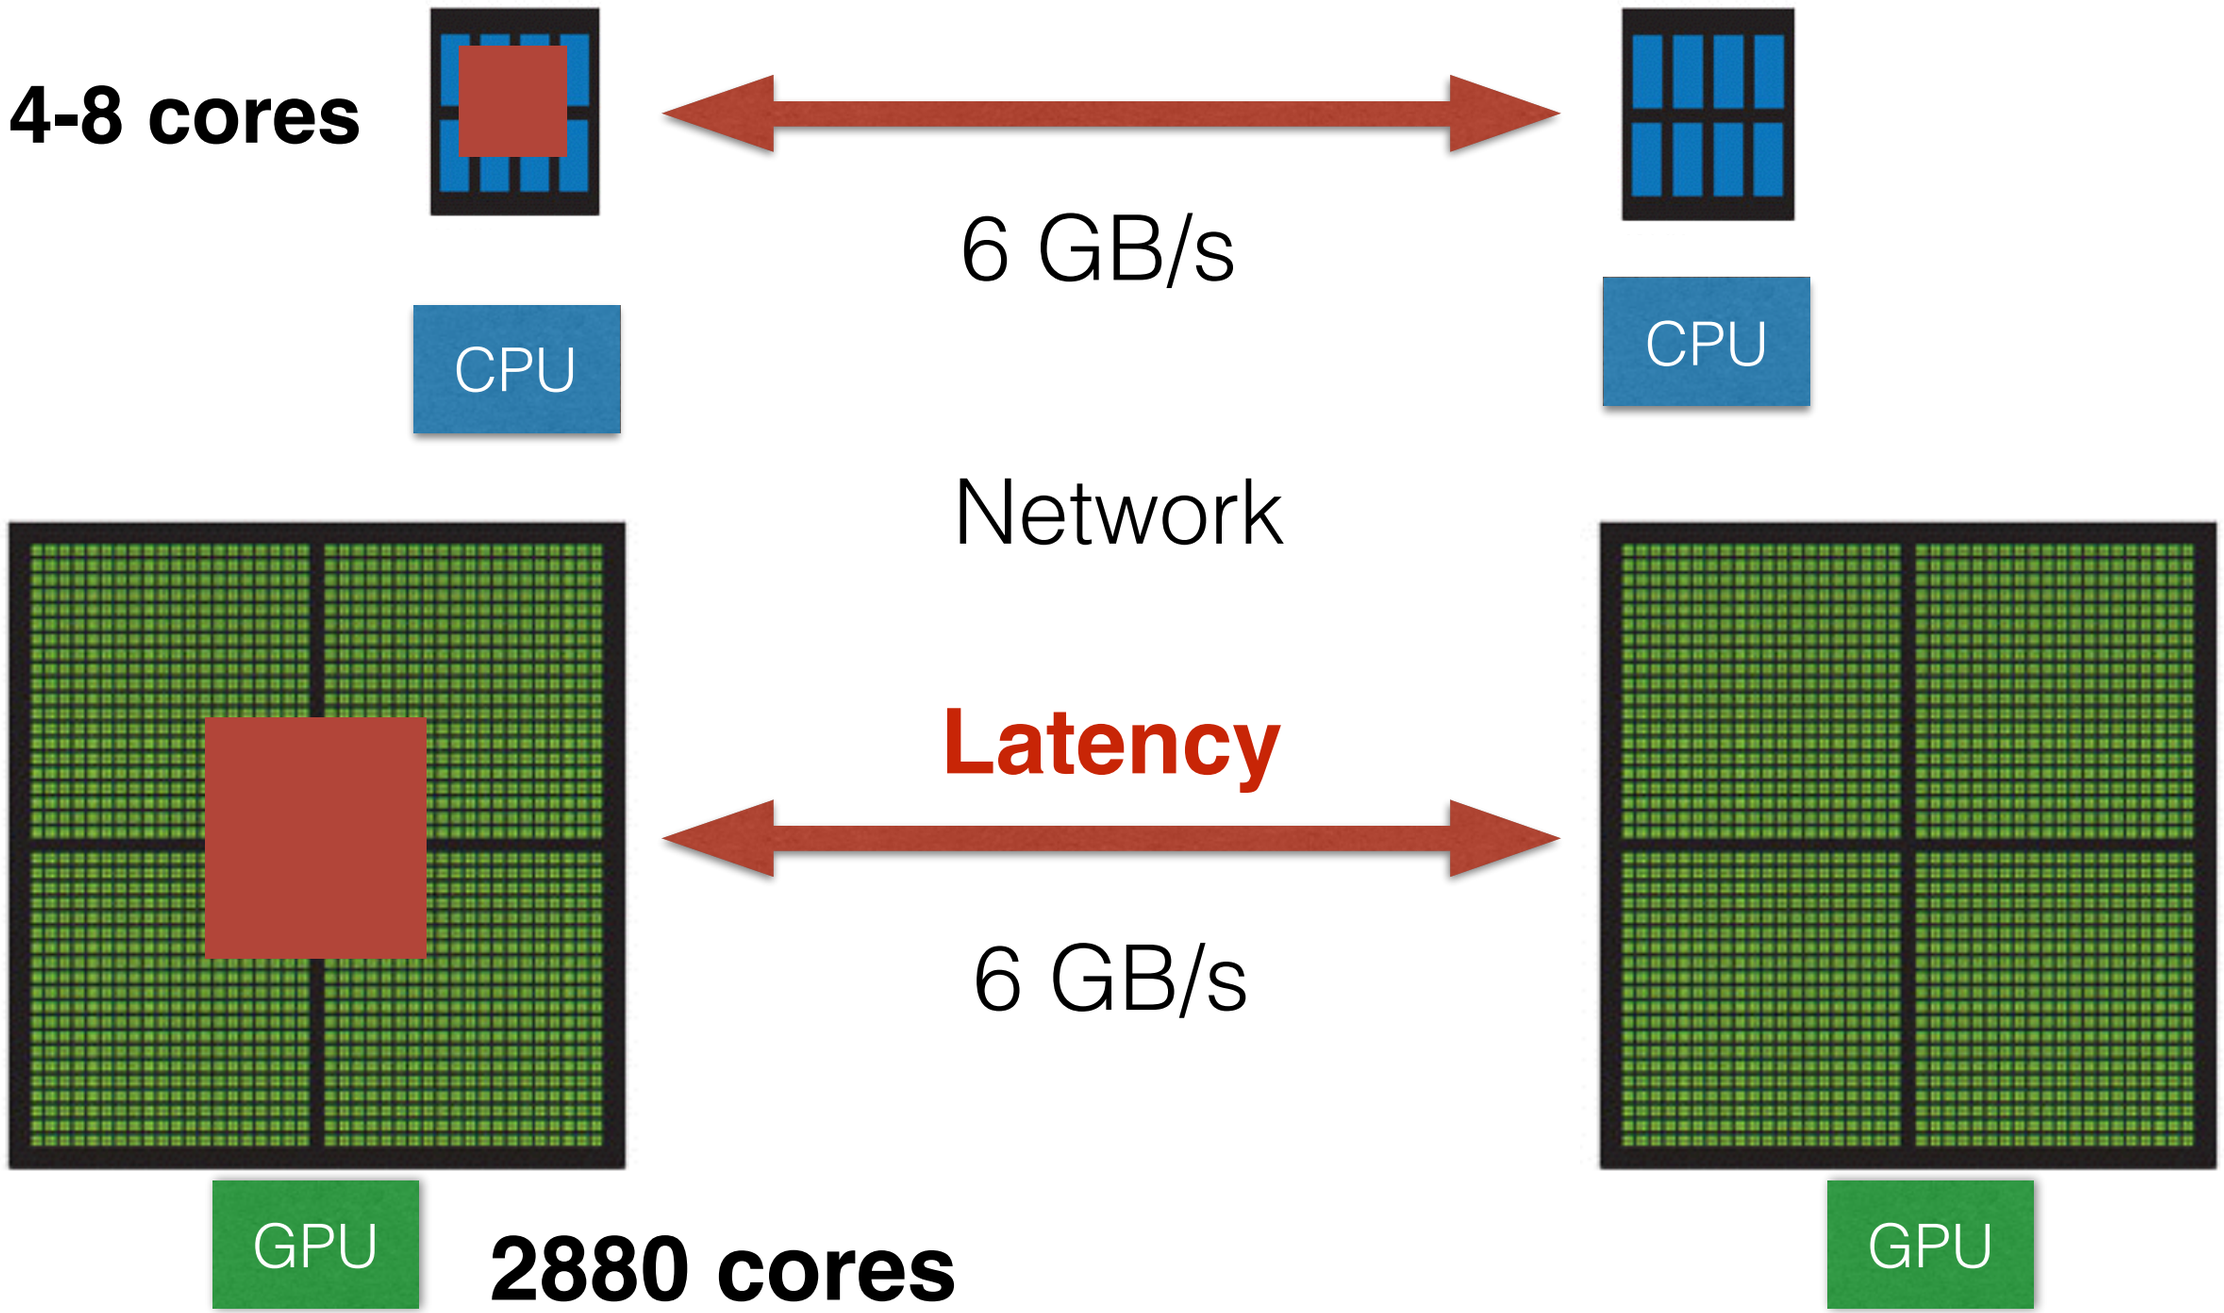

Latency optimizations

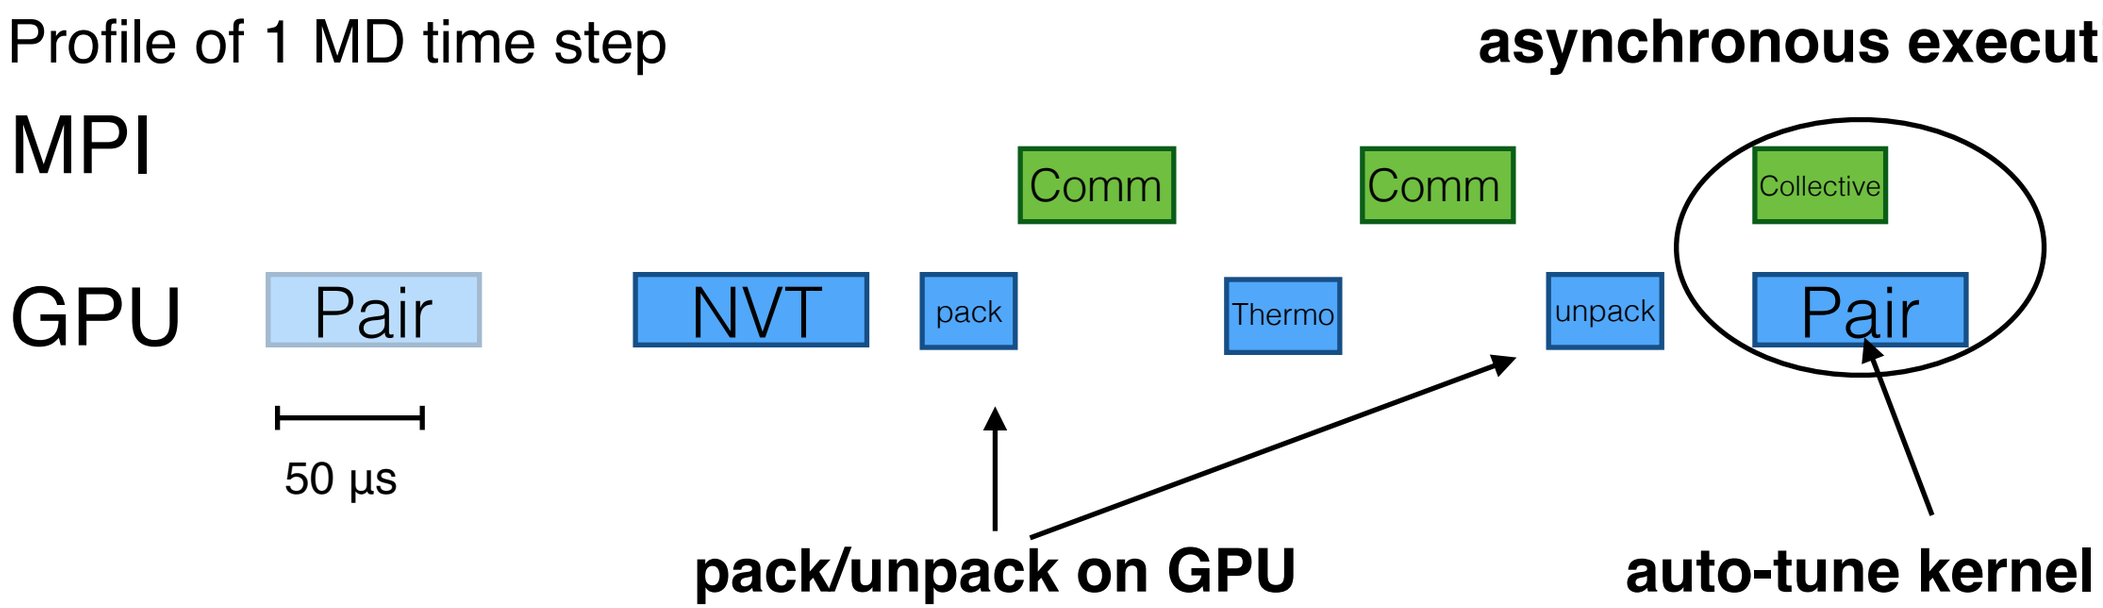

- **Launch** latency (many kernels)
- **Communication** latency (PCIe, Infiniband)
- **Synchronization** latency (MPI collectives, many ranks)

Strong Scaling of a LJ Liquid (N=10,976,000)

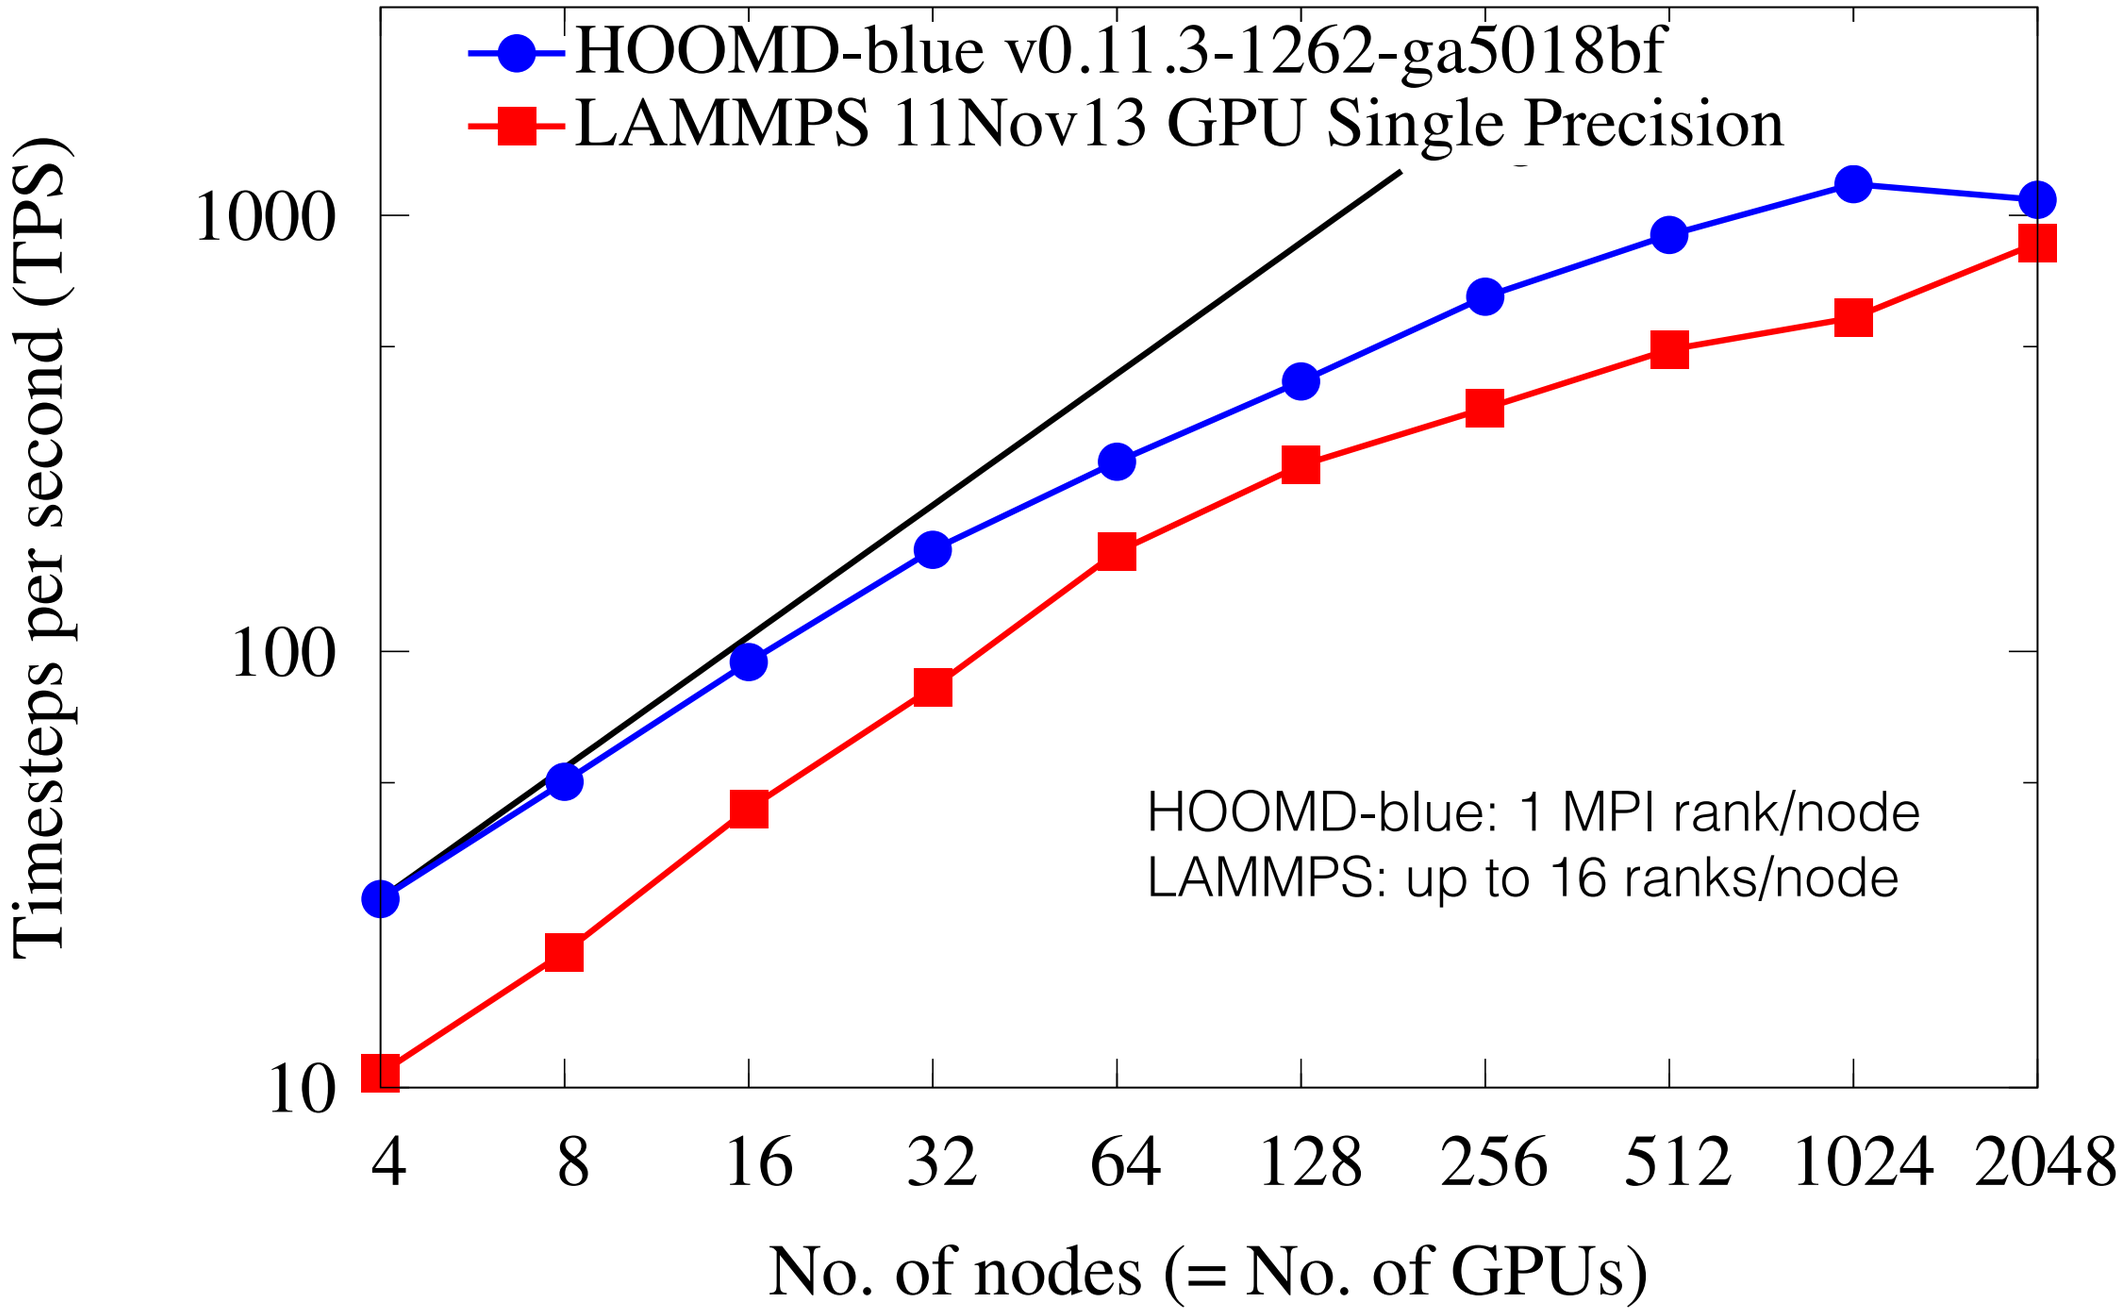

Benchmarks courtesy of Trung Nguyen, ORNL

Weak scaling up to 108,000,000 particles

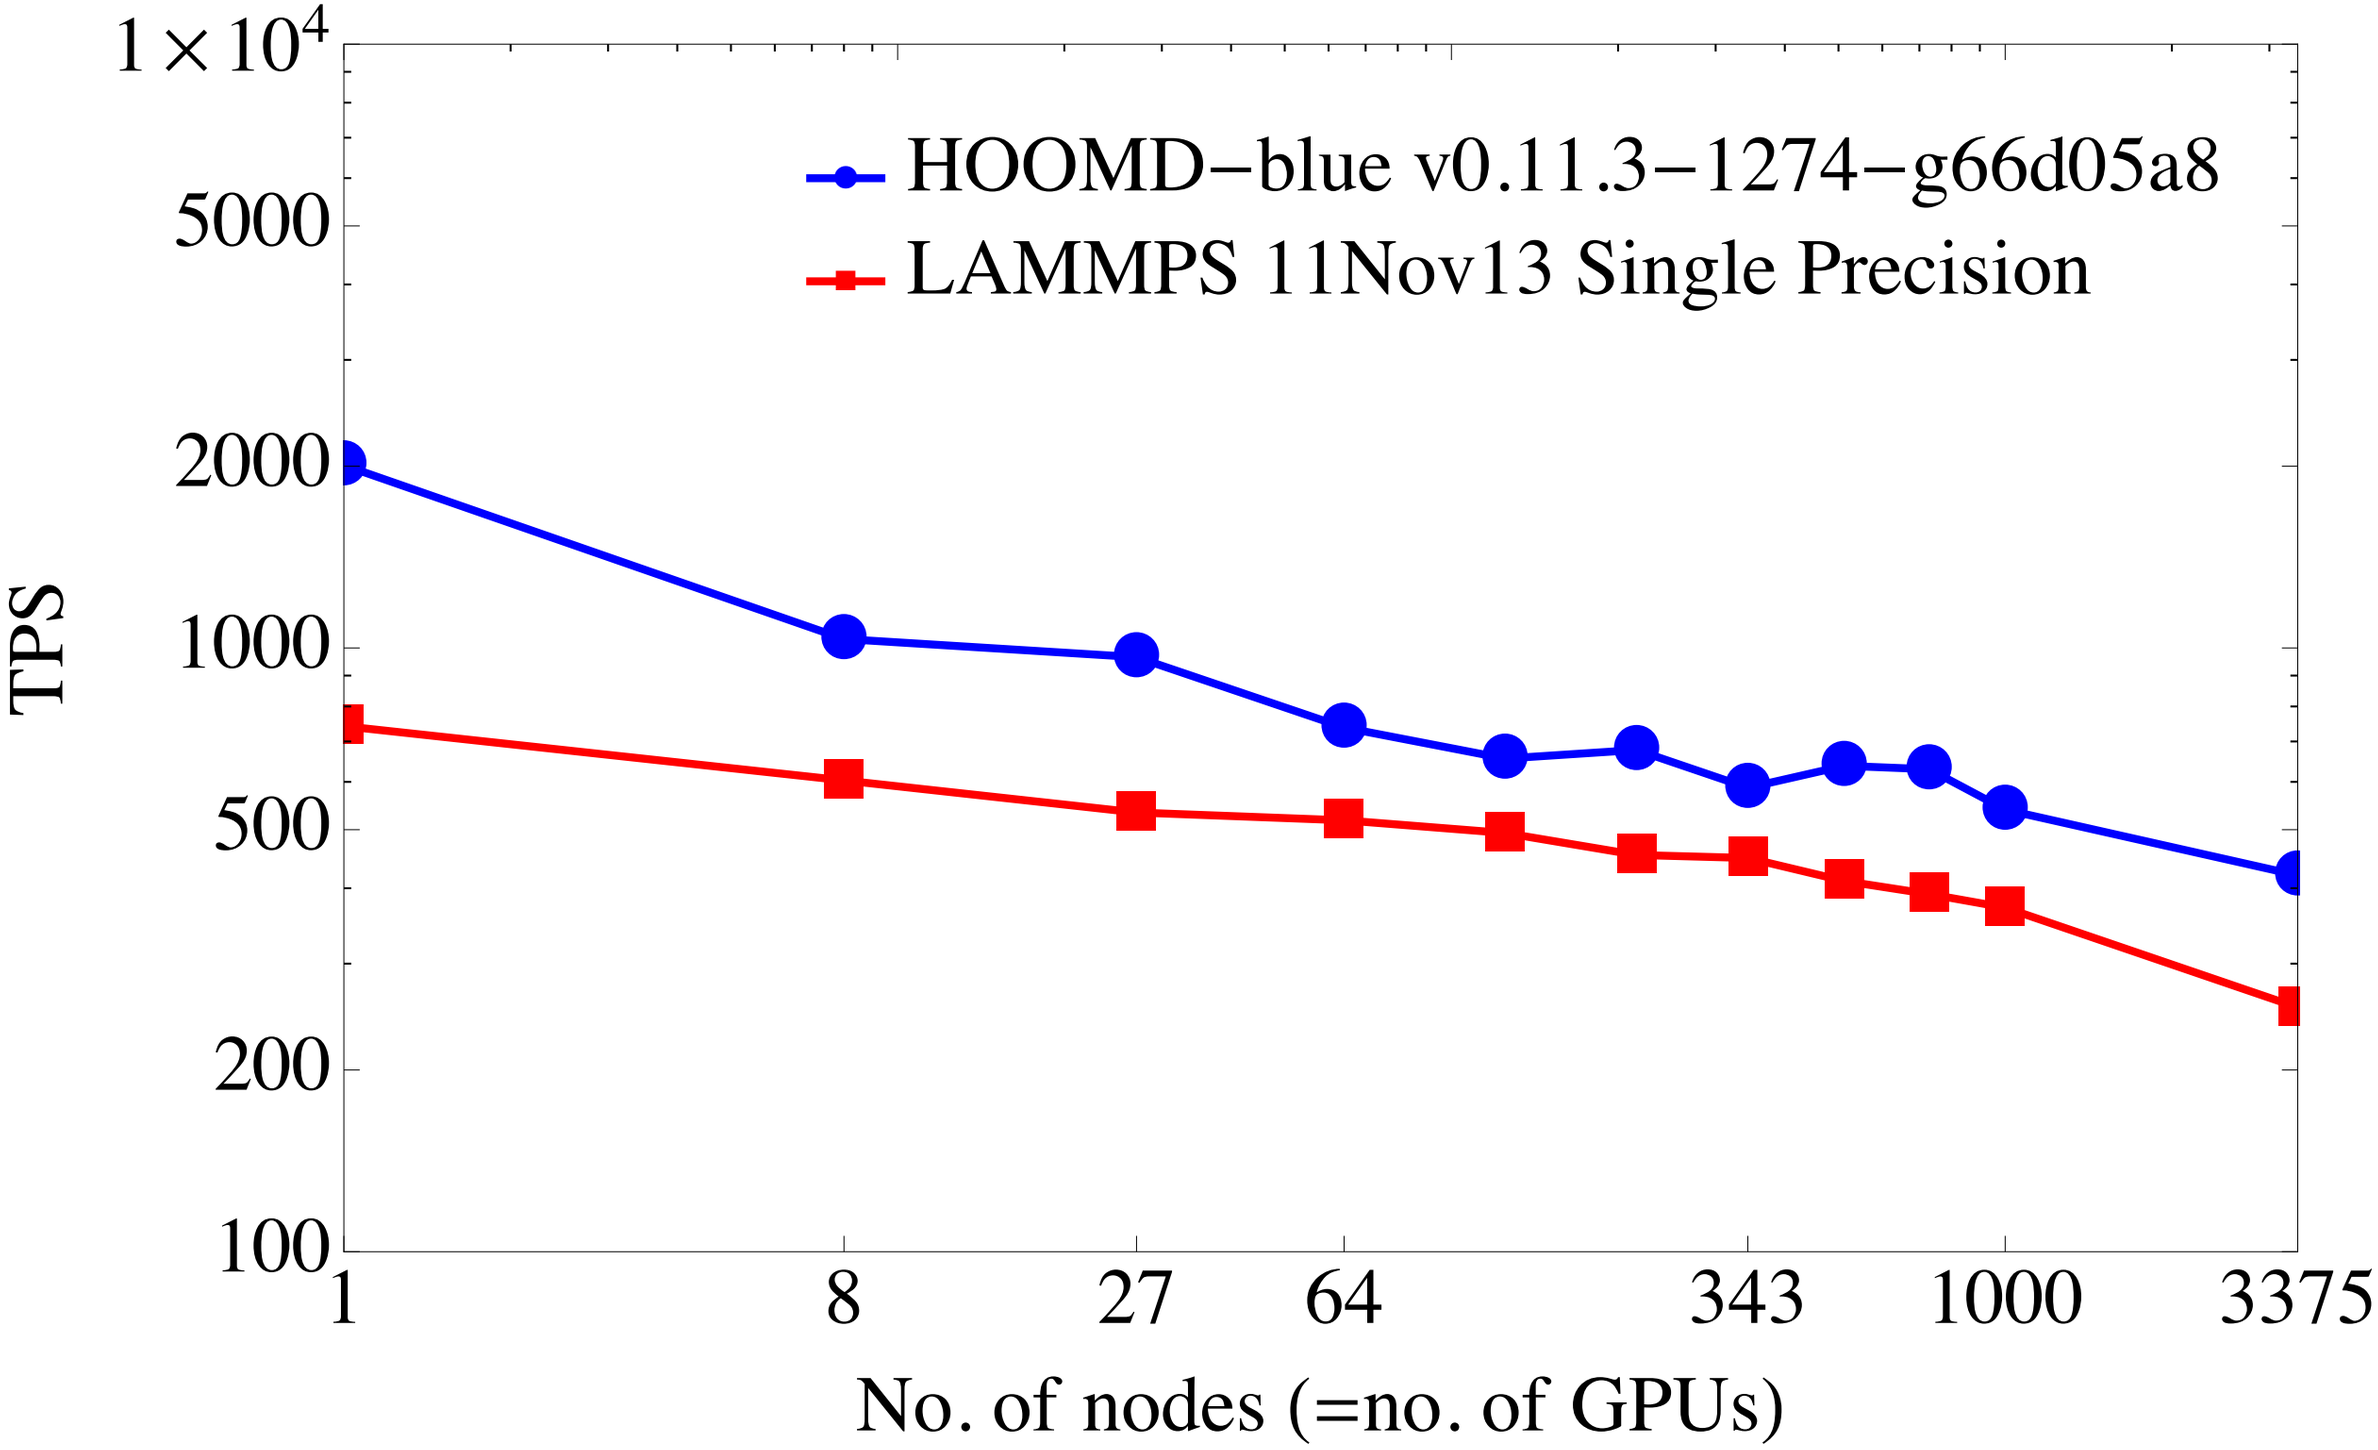

32,000 particles/GPU

Scaling efficiency

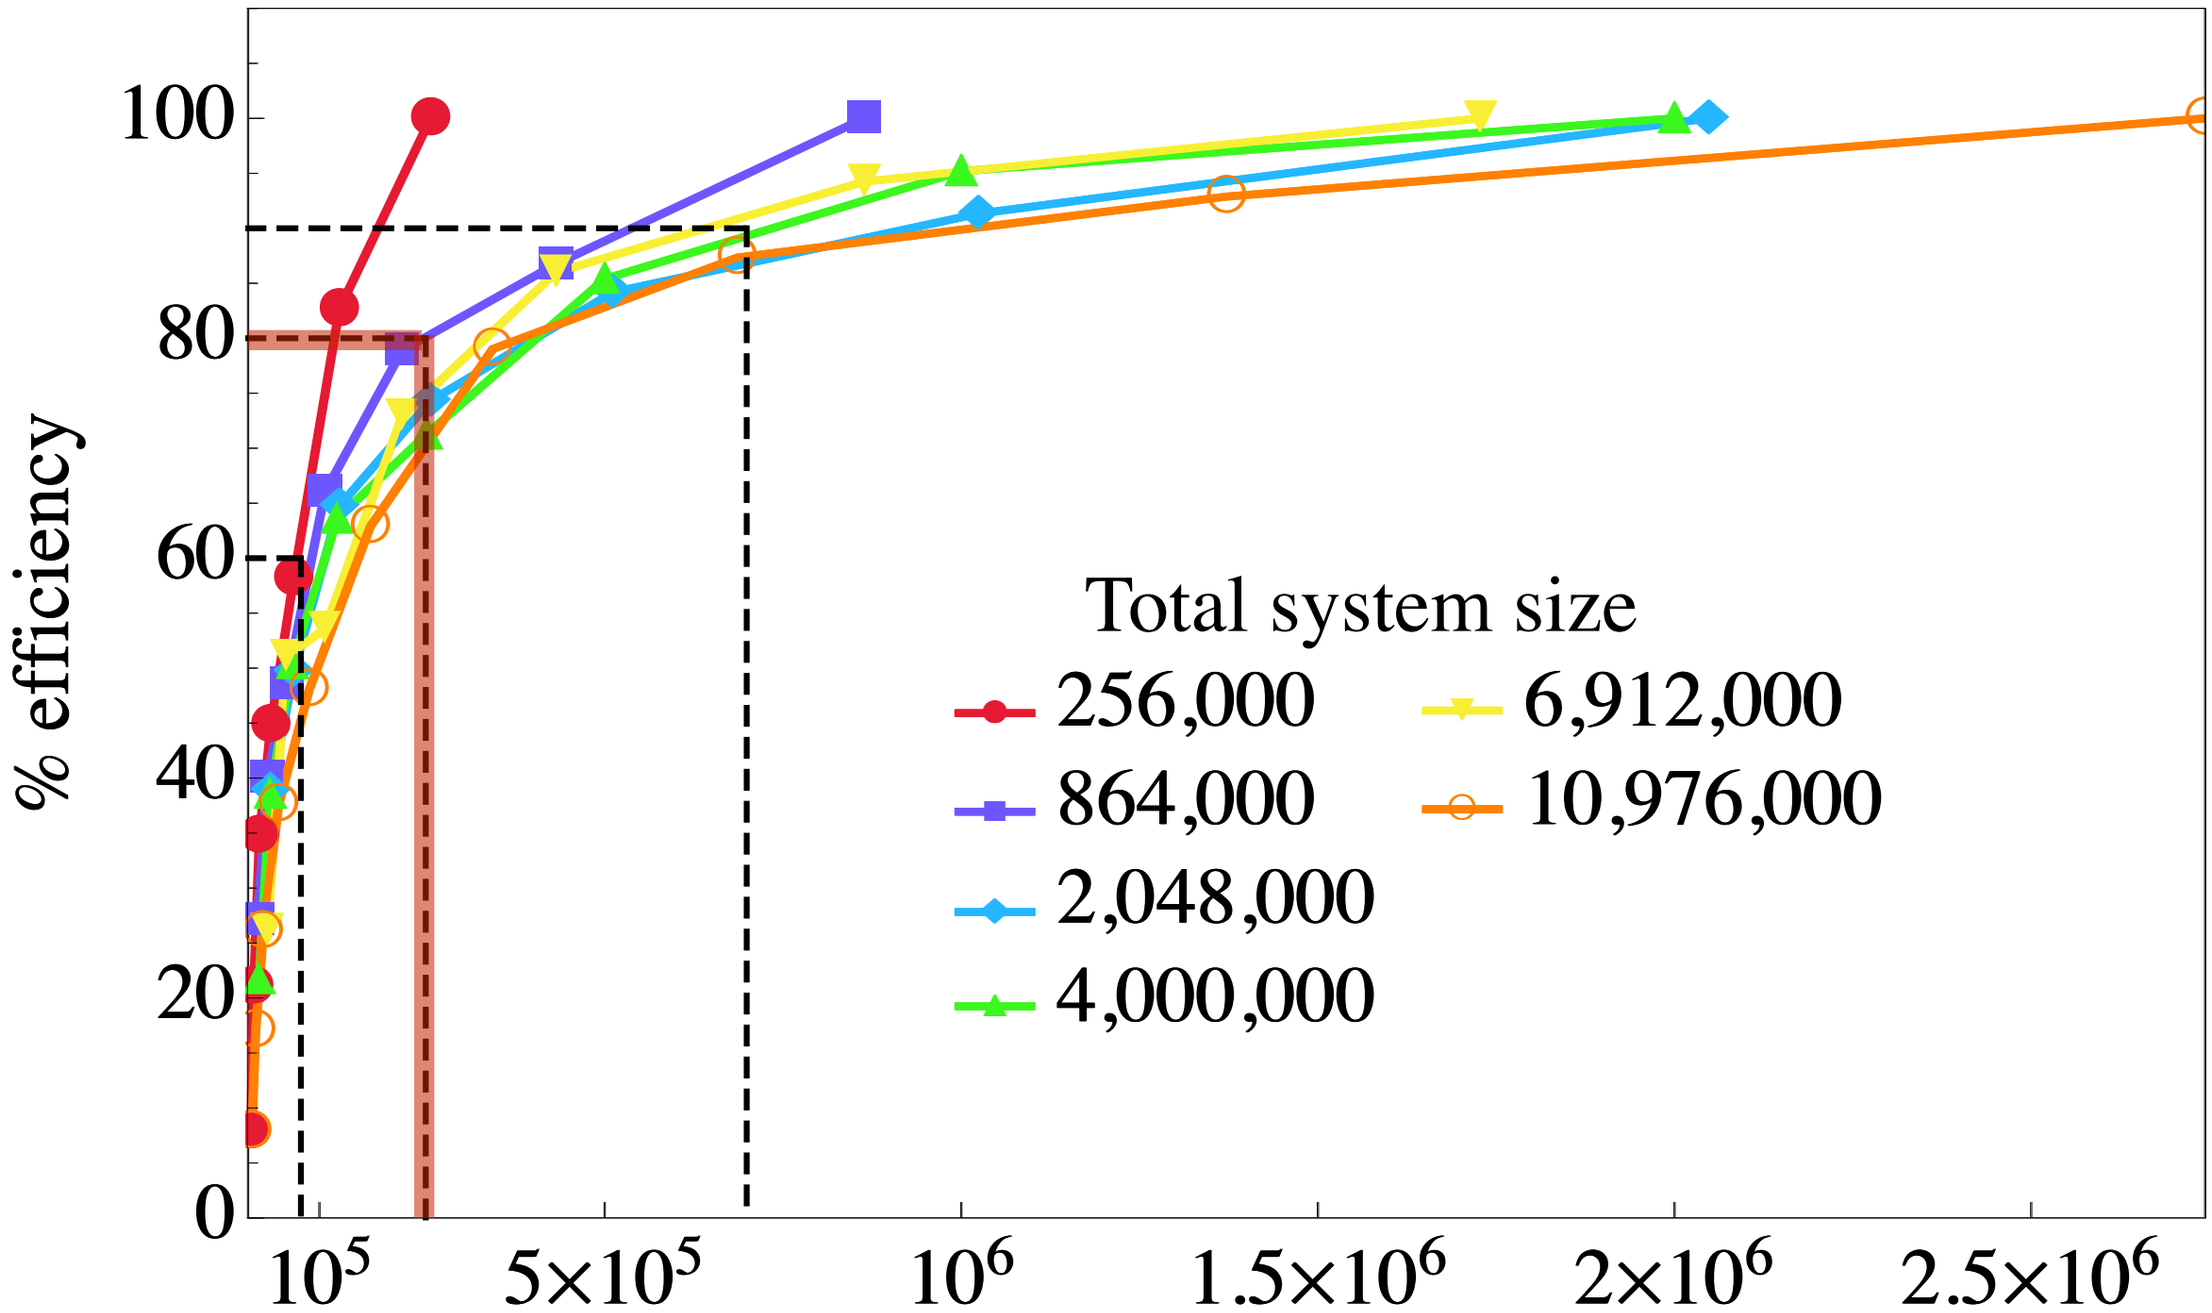

80% efficiency for 250,000 ptls/GPU

Remote Direct Memory Access

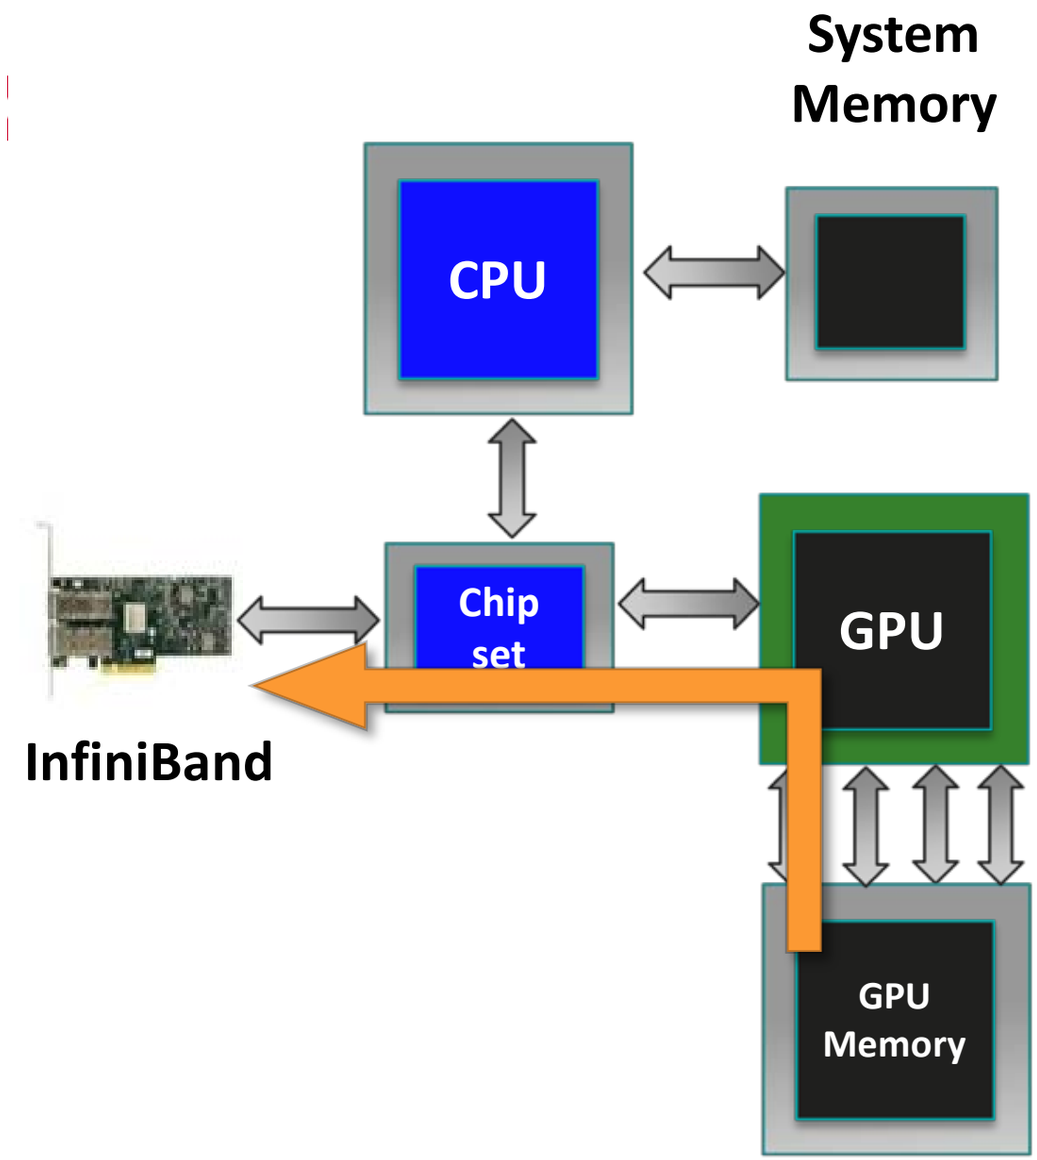

GPUDirect RDMA

Wilkes cluster (#2 Green 500) at University of Cambridge

Dell T620 Cluster, 128 nodes, 3631 MFLOP/W
Intel Xeon E5-2630v2 6C 2.600GHz
Node configuration: 2x Infiniband FDR, 2x NVIDIA K20

HOOMD with GPUDirect RDMA

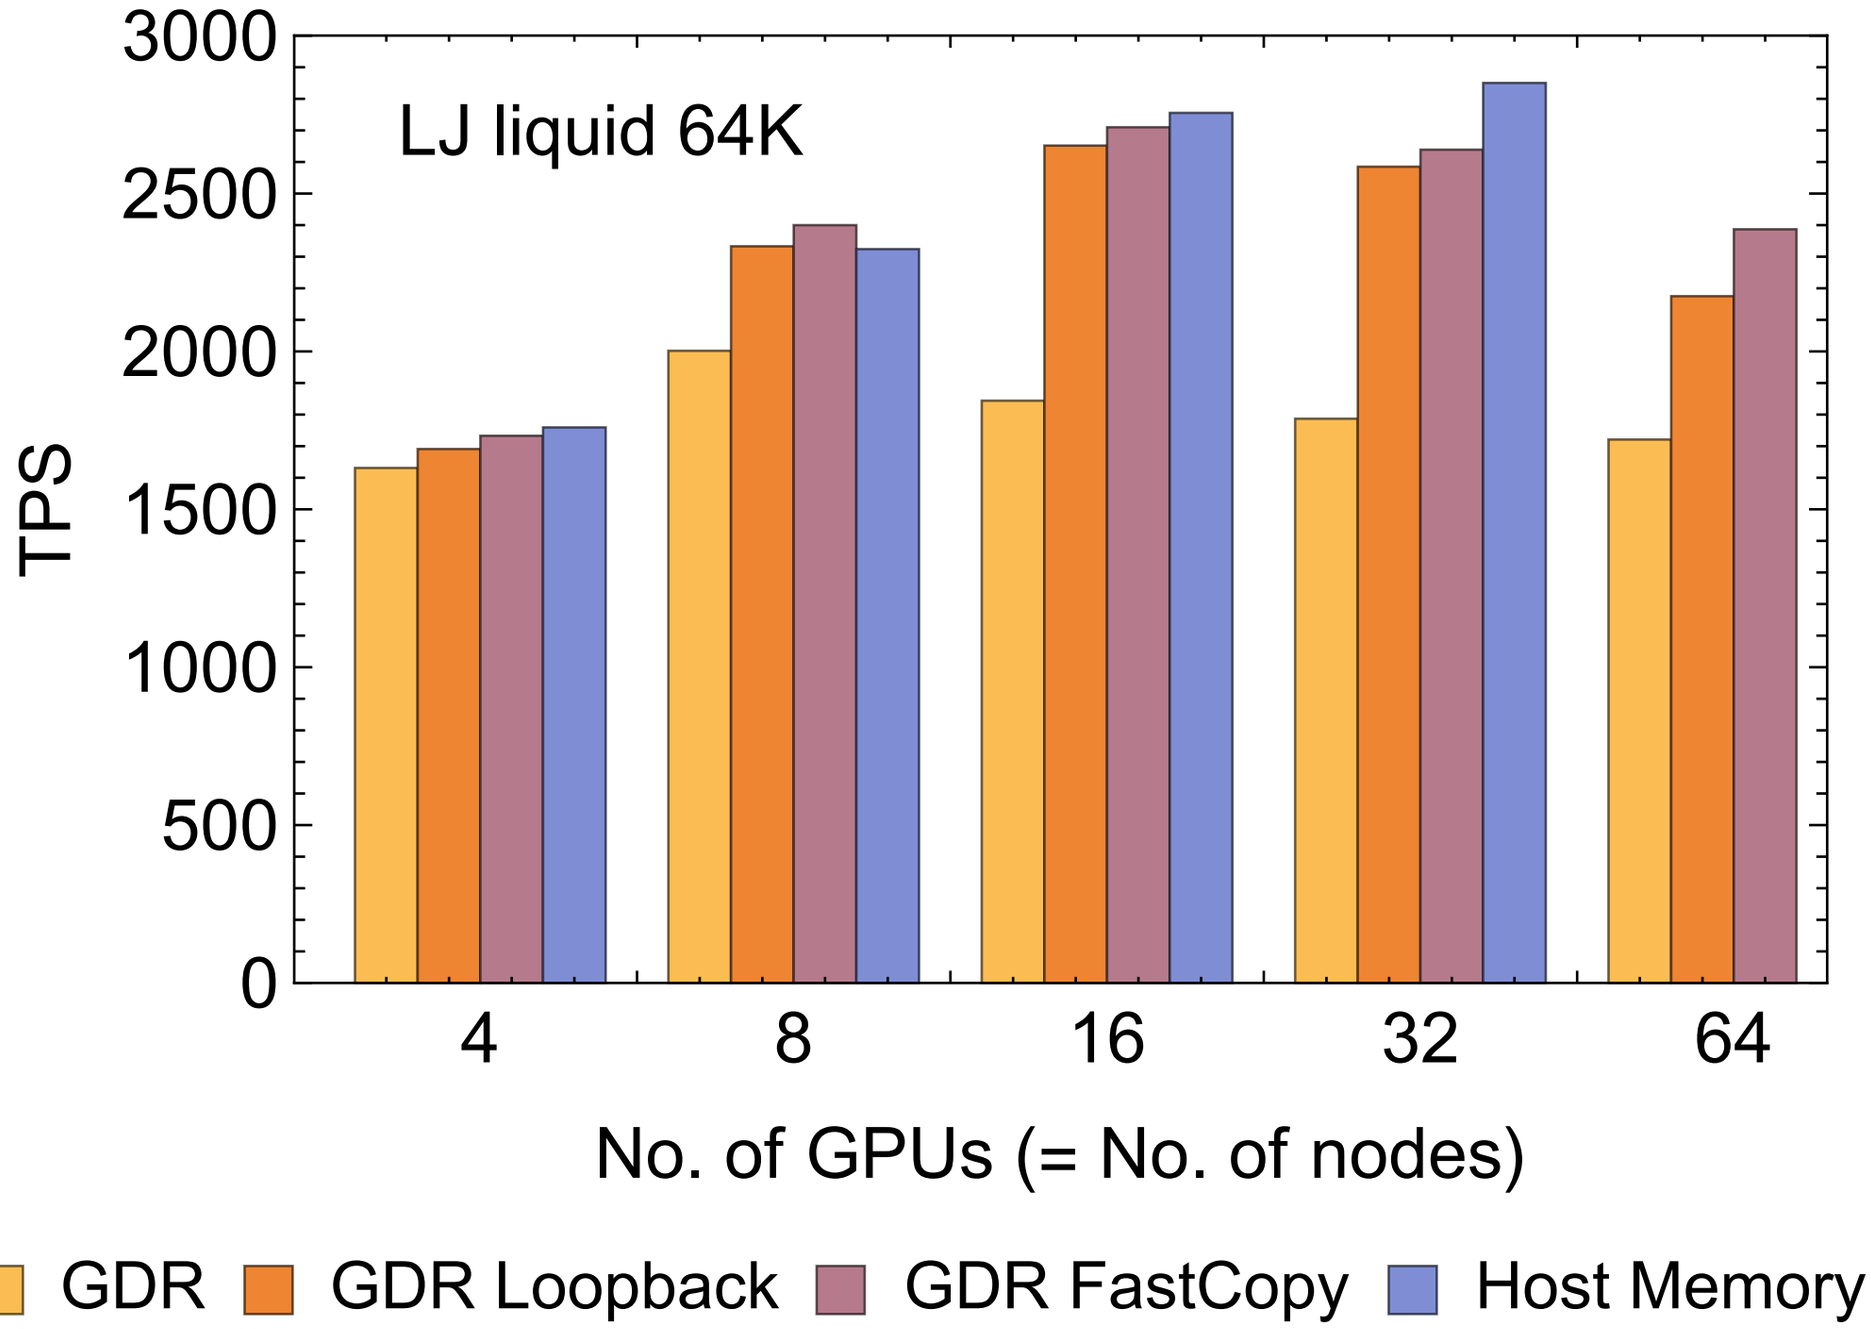

XSEDE Stampede GPU partition (Host Memory)

MVAPICH2 2.0 GDR data
courtesy of D.K. Panda and Rong Shi, OSU

Summary

- **Multi-GPU** support in **HOOMD 1.0** enables large-scale simulations

- **Strong Scaling** is latency-limited, works on 1000's of GPUs

- **GPUDirect RDMA** is a promising technology, and MPI3 RMA may be the way to fully take advantage of it

URL & Outlook

http://codeblue.umich.edu/hoomd-blue

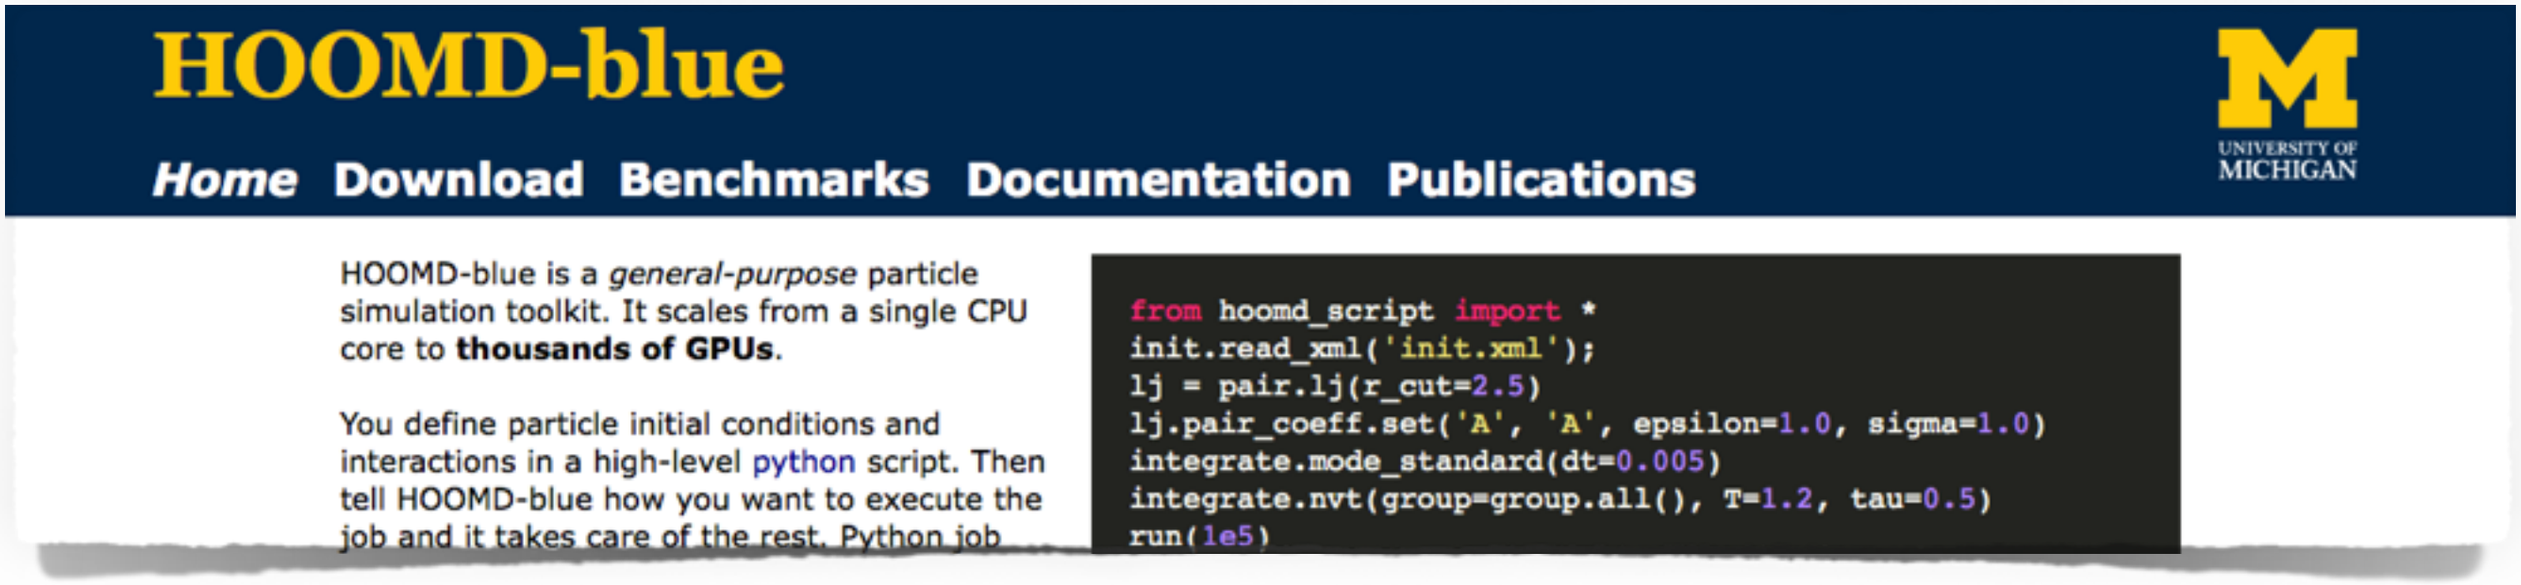

HOOMD 1.1+:

- Anisotropic pair potentials
- Multi-GPU rigid bodies
- … your feature!

Questions?