How parallelism helps computational physics discovery

Jens Glaser
Chemical Engineering, University of Michigan

3rd Annual MVAPICH User Group (MUG) Meeting, Columbus, OH
August 21, 2015

Moore’s law

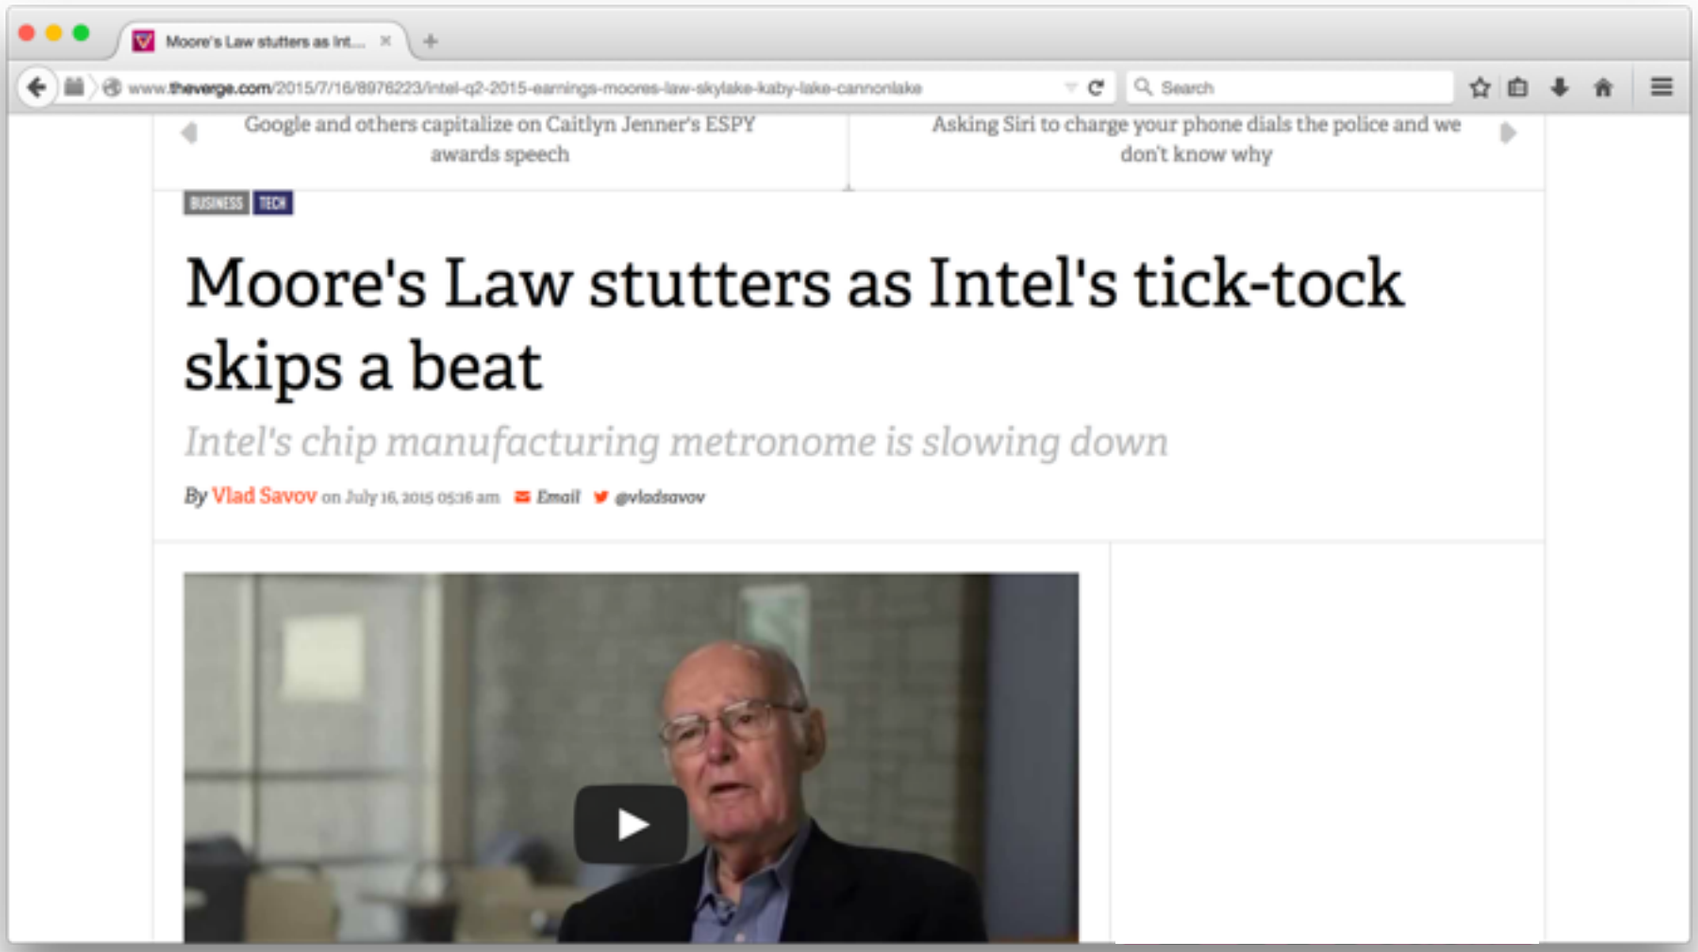

“Before founding Intel in the 1960s, Gordon Moore made a bold prediction about the exponential growth in the number of components on integrated circuits, which has been proven remarkably accurate by subsequent history and immortalized under the title of Moore's Law. Intel has been doggedly upholding Moore's Law by roughly doubling the number of transistors in its processors every couple of years, but now that schedule is starting to slip.”

Should this concern us computational physicists?

Computing paradigms

Serial processor

Parallel processor

Outline

- Parallelism in Molecular Dynamics (MD)
- Parallelism in Monte Carlo (MC)
- Strong Scaling of MD & GPUDirect RDMA
- Applications

Molecular dynamics

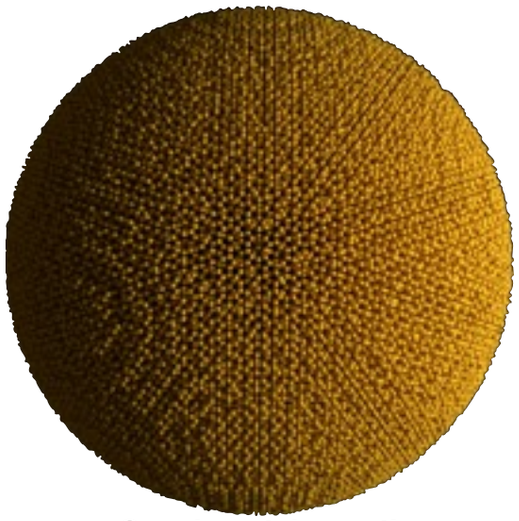

Quasicrystal growth
Molecular Dynamics
Engel M. et al., *Nature Materials* (in press)

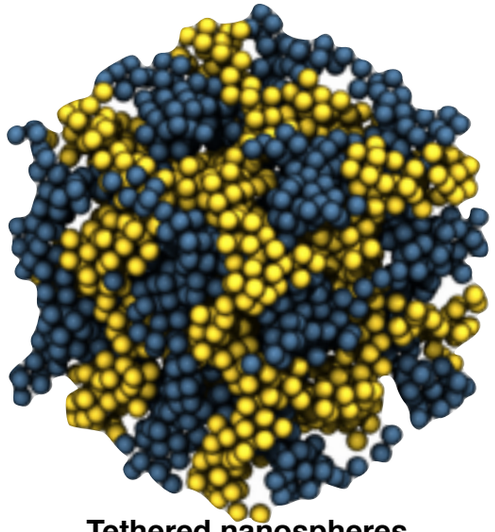

Tethered nanospheres
Langevin dynamics
Marson, R, *Nano Letters* **14**, 4, 2014

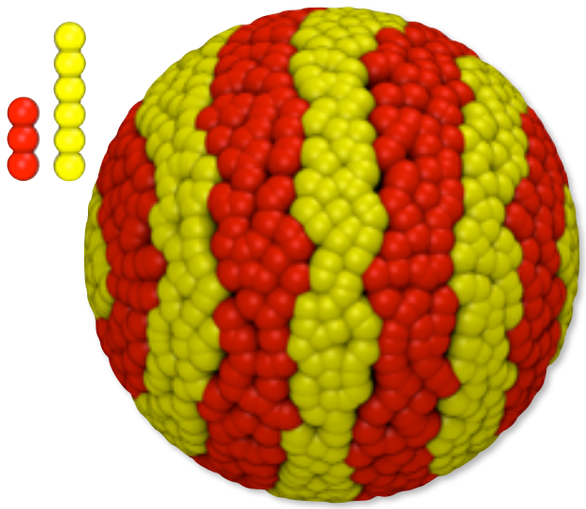

Surfactant coated surfaces
Dissipative particle dynamics
Pons-Siepermann, I. C., *Soft* matter **6** 3919 (2012)

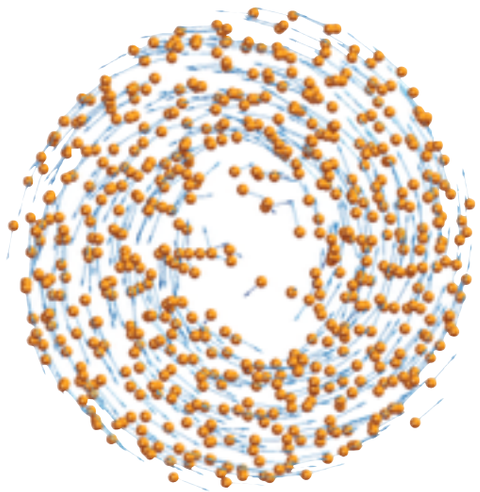

Self-propelled colloids
Non-equilibrium MD
Nguyen N., *Phys Rev E* **86** 1, 2012

Monte Carlo

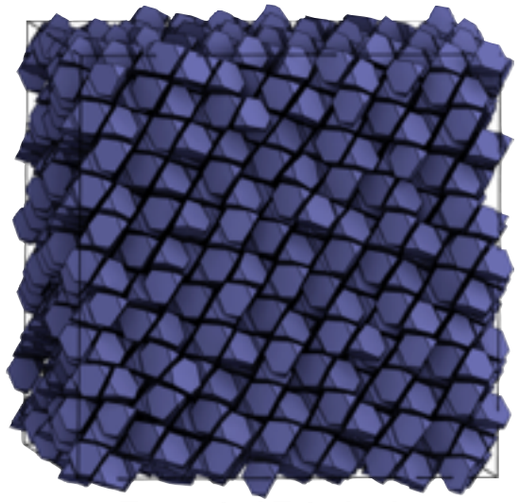

Truncated Tetrahedra
Hard particle MC
Damasceno, P. F. et al., *ACS Nano* **6**, 609 (2012)

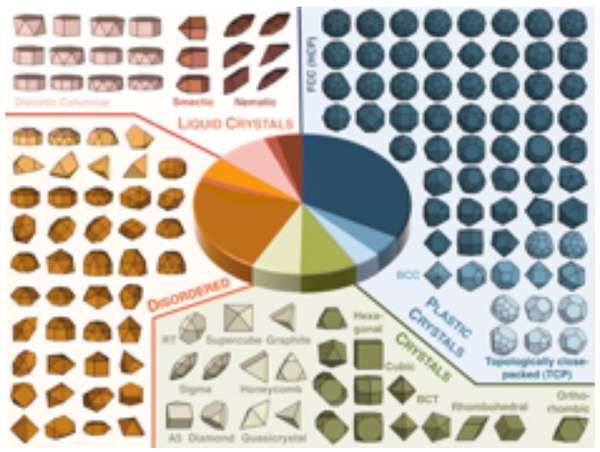

Arbitrary polyhedra
Hard particle MC
Damasceno, P. F. et al., *Science* **337**, 453 (2012)

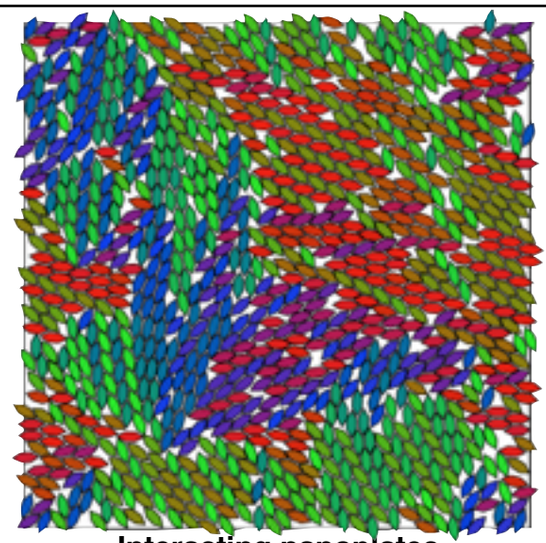

Interacting nanoplates
Hard particle MC with interactions
Ye X. et al., *Nature Chemistry* **cover article** (2013)

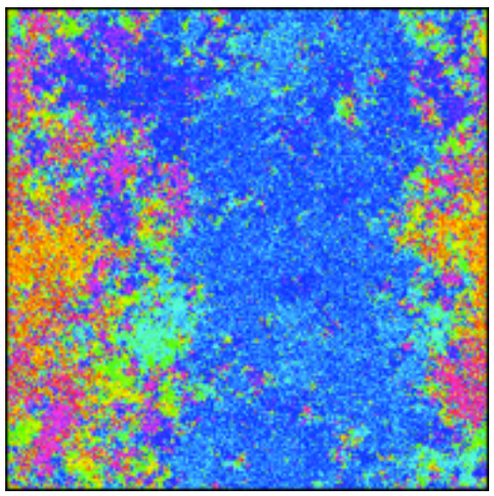

Hard disks - hexatic
Hard particle MC
Engel M. et al., PRE **87**, 042134 (2013)

Publications using HOOMD-blue

Is your publication not listed? *Contact us* and we will add it.

[109]: Fabrizio Benedetti, Aleksandre Japaridze, Julien Dorier, Dusan Racko, Robert Kwapich, Yannis Burnier, Giovanni Dietler, and Andrzej Stasiak. Effects of physiological self-crowding of DNA on shape and biological properties of DNA molecules with various levels of supercoiling. *Nucleic Acids Research*, pages 1–10, 2015. (doi:10.1093/nar/gkv055)

[108]: Hendrick W. de Haan and Tyler N. Shendruk. Force–Extension for DNA in a Nanoslit: Mapping between the 3D and 2D Limits. *ACS Macro Letters*, pages 632–635, 2015. (doi:10.1021/acsmacrolett.5b00138)

[107]: Filippo Federici Canova, Masashi Mizukami, Takako Imamura, Kazue Kurihara, and Alexander Shluger. Structural stability and polarisation in ionic liquids films on silica surfaces. *Phys. Chem. Chem. Phys.*, 2015. (doi:10.1039/C5CP02299A)

[106]: Elijah Flenner and Grzegorz Szamel. Long-range correlations in glasses and glassy fluids. *Journal of Physics: Condensed Matter*,

slide by Joshua Anderson

Molecular Dynamics

Iteratively solve **Newton's equations of motion**

$$F = ma$$

(1) $$\vec{r}(t+\Delta t) = \vec{r}(t) + \vec{v}(t)\Delta t + \frac{1}{2}\vec{a}(t)\Delta t^2$$

$$\vec{v}(t+\Delta t/2) = \vec{v}(t) + \frac{1}{2}\vec{a}(t)\Delta t$$

compute accelerations $a(t)$

(2) $$\vec{v}(t+\Delta t) = \vec{v}(t+\Delta t/2) + \frac{1}{2}\vec{a}(t+\Delta t)$$

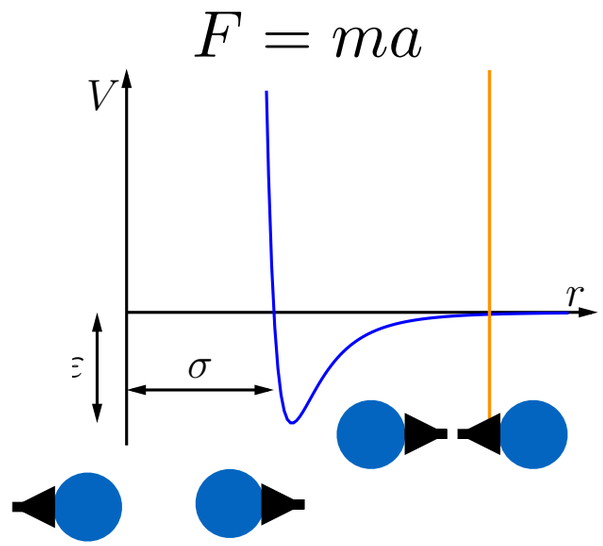

Lennard-Jones

$$V_{LJ}(r) = 4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right]$$

Parallel algorithm for N particles

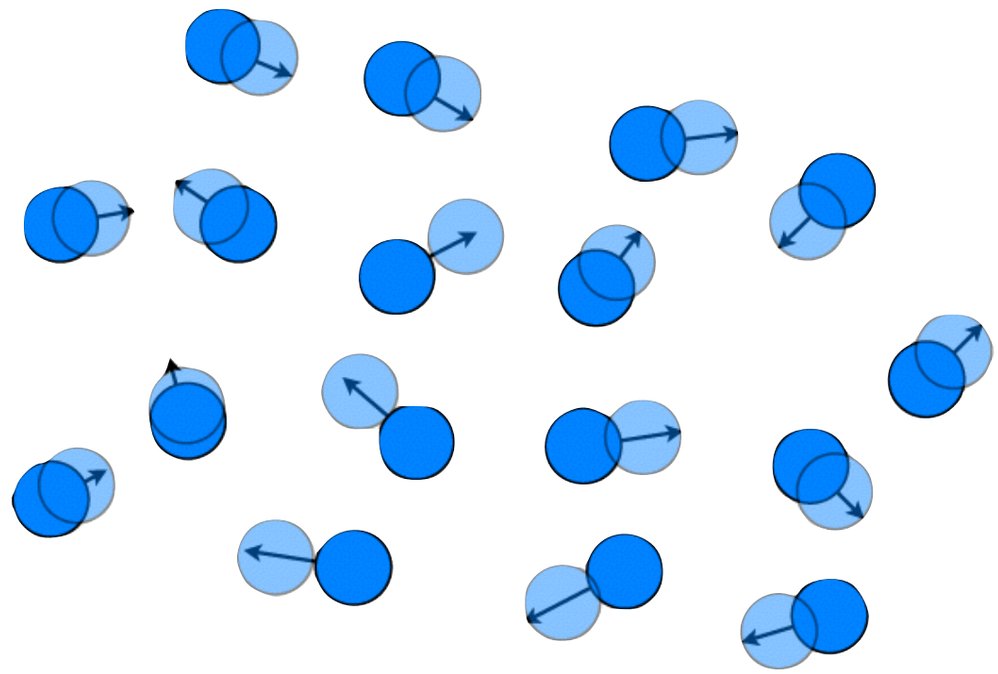

GPU Molecular Dynamics in 1 slide

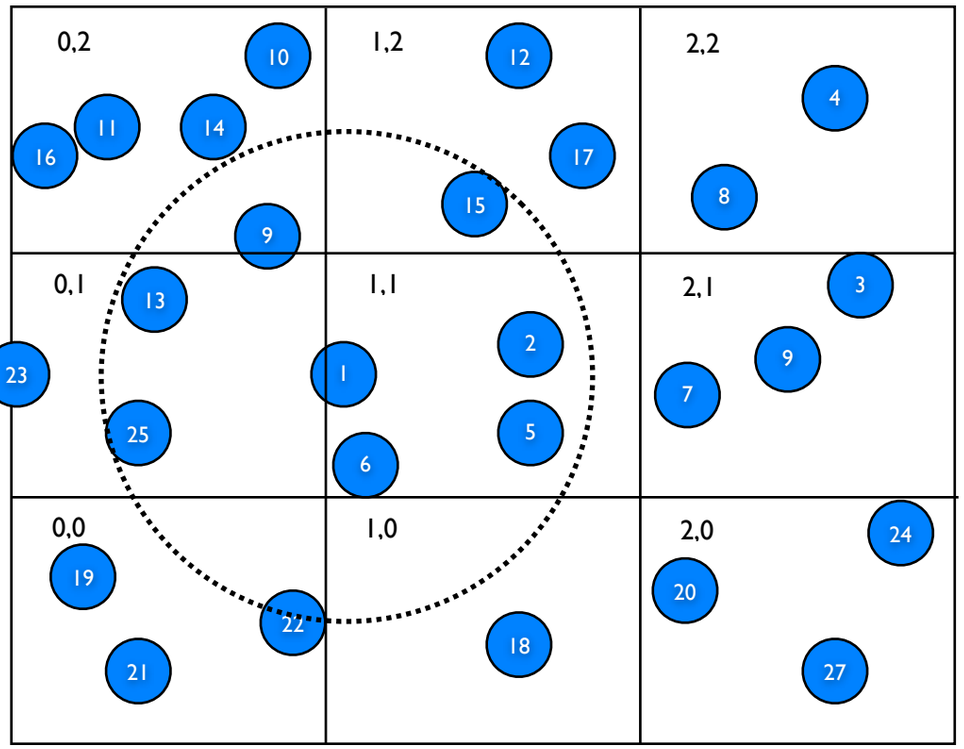

Cell list

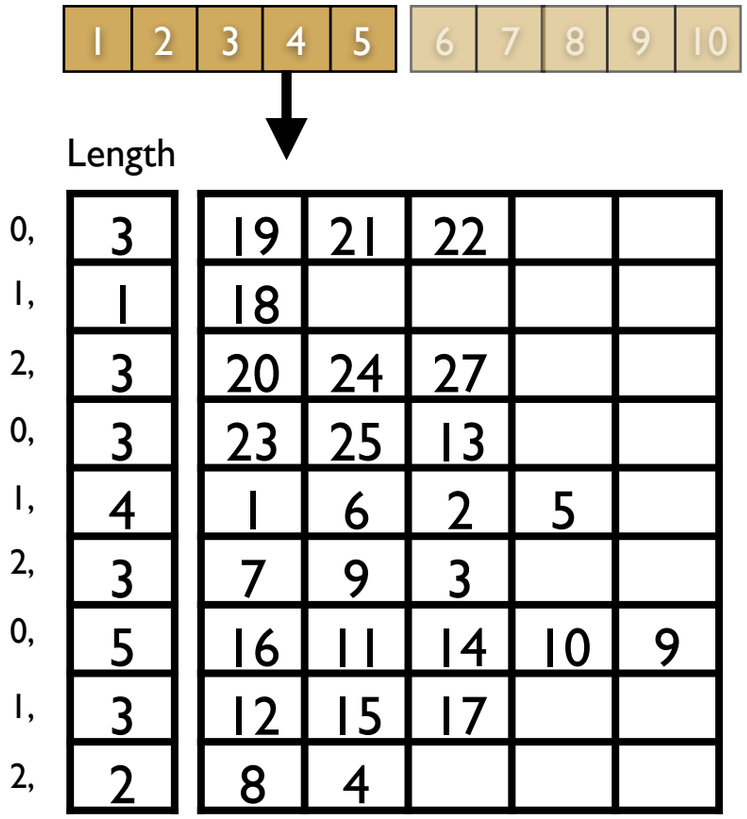

Neighbor list

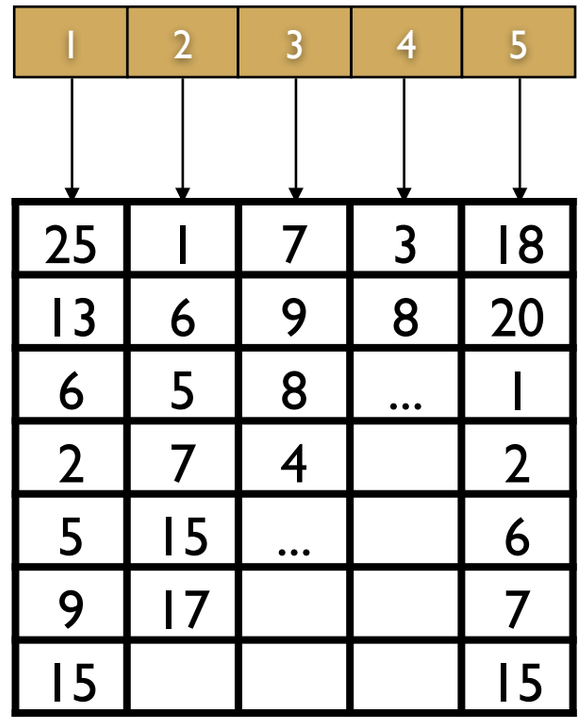

Neighbor list

Num Neighbors

Pair force

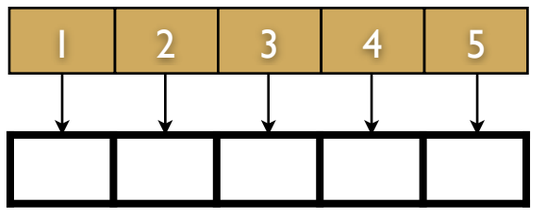

Integrate

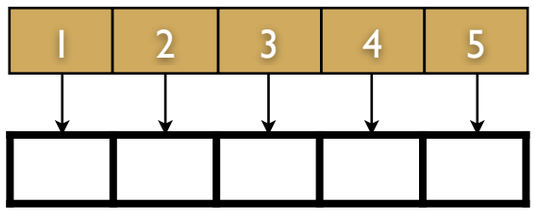

one thread/particle

slide by Joshua Anderson

Parallel domain decomposition

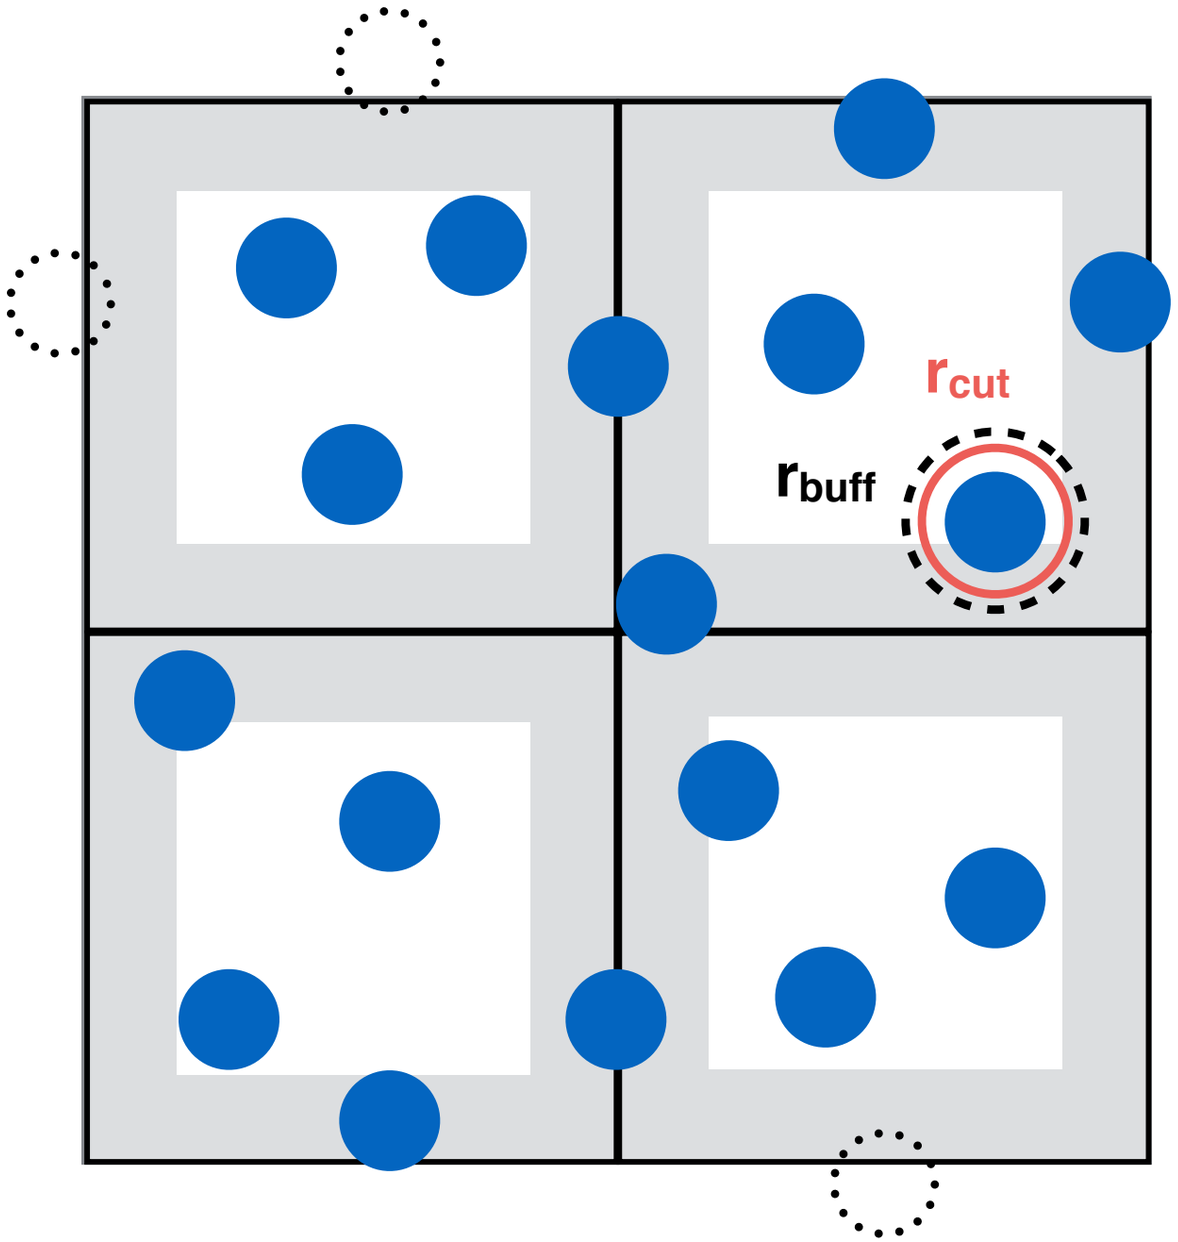

- Particles can **leave** and **enter** domains (periodic boundary conditions)
- **Ghost** particles required for force computation
- **Communicate** positions of ghost particles every time step

Easy and flexible to use

```
from hoomd_script import *

context.initialize()
init.read_xml('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(group=group.all(), T=1.2, tau=0.5)
run(1e5)
```

```
$ hoomd run.py --mode=cpu
$ hoomd run.py --mode=gpu
$ mpirun -n 256 hoomd run.py --mode=cpu
$ mpirun -n 64 hoomd run.py --mode=gpu
```

Hard particle Monte Carlo

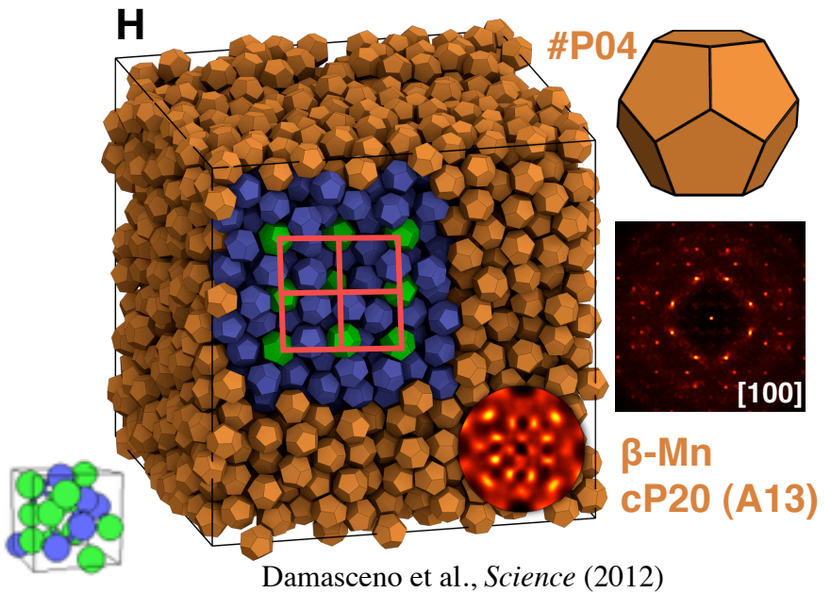

Damasceno et al., *Science* (2012)

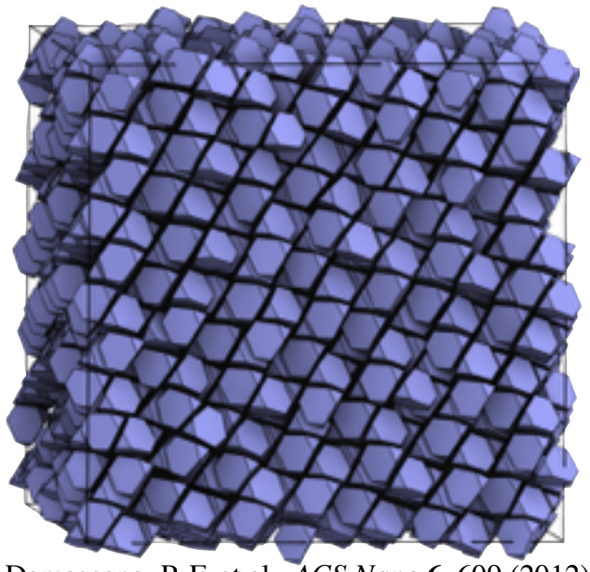

Damasceno, P. F. et al., *ACS Nano* **6**, 609 (2012)

- 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
- fast AABB trees on CPU
- Parallel execution on a single GPU
- Domain decomposition across multiple nodes (CPUs or GPUs)

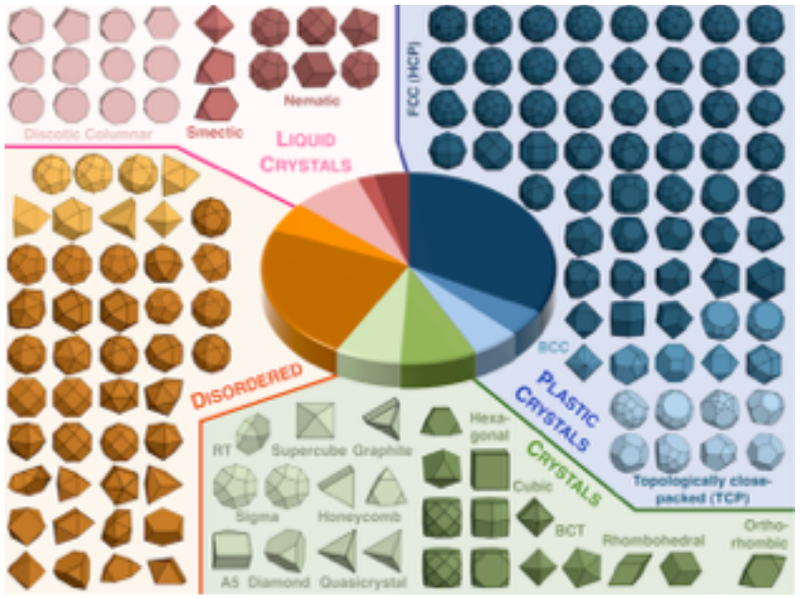

Damasceno et al., *Science* (2012)

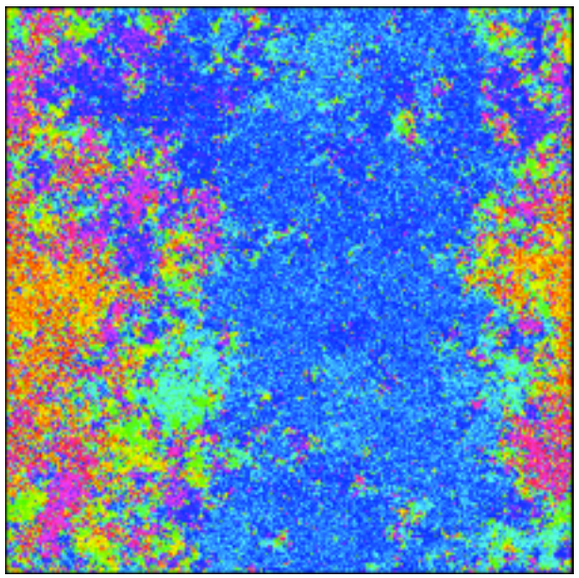

Engel M. et al., PRE **87**, 042134 (2013)

slide by Joshua Anderson

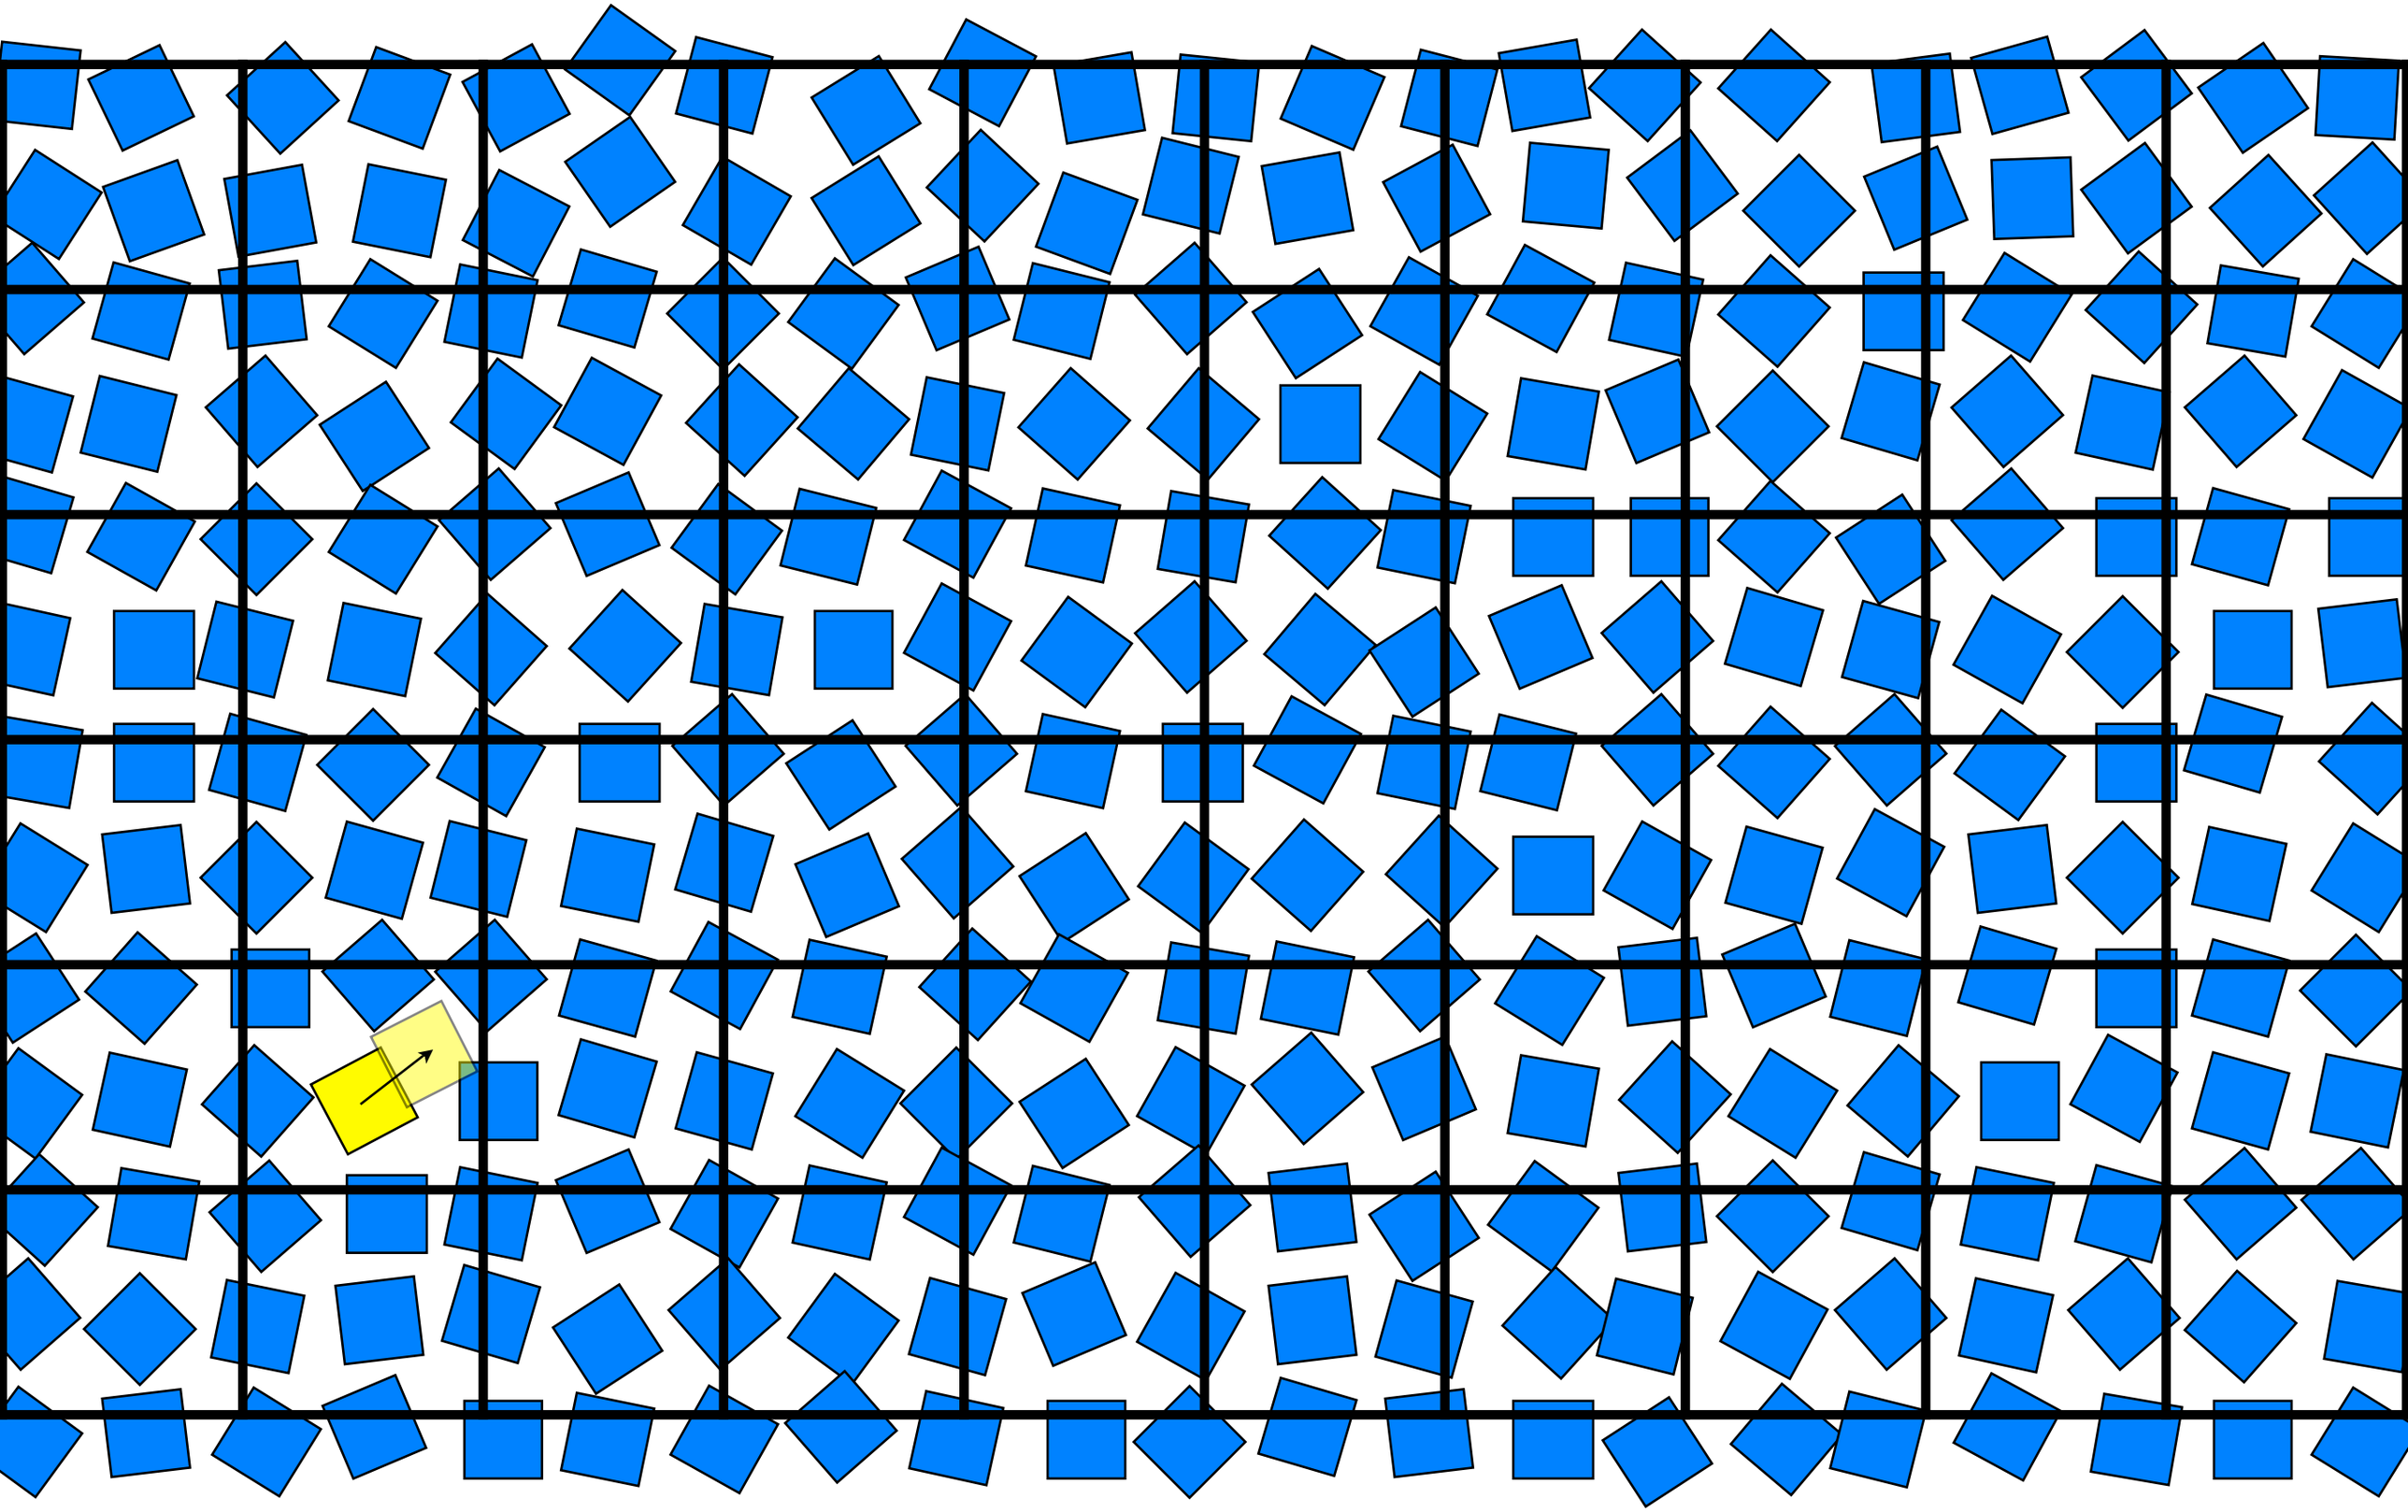

slide by Joshua Anderson

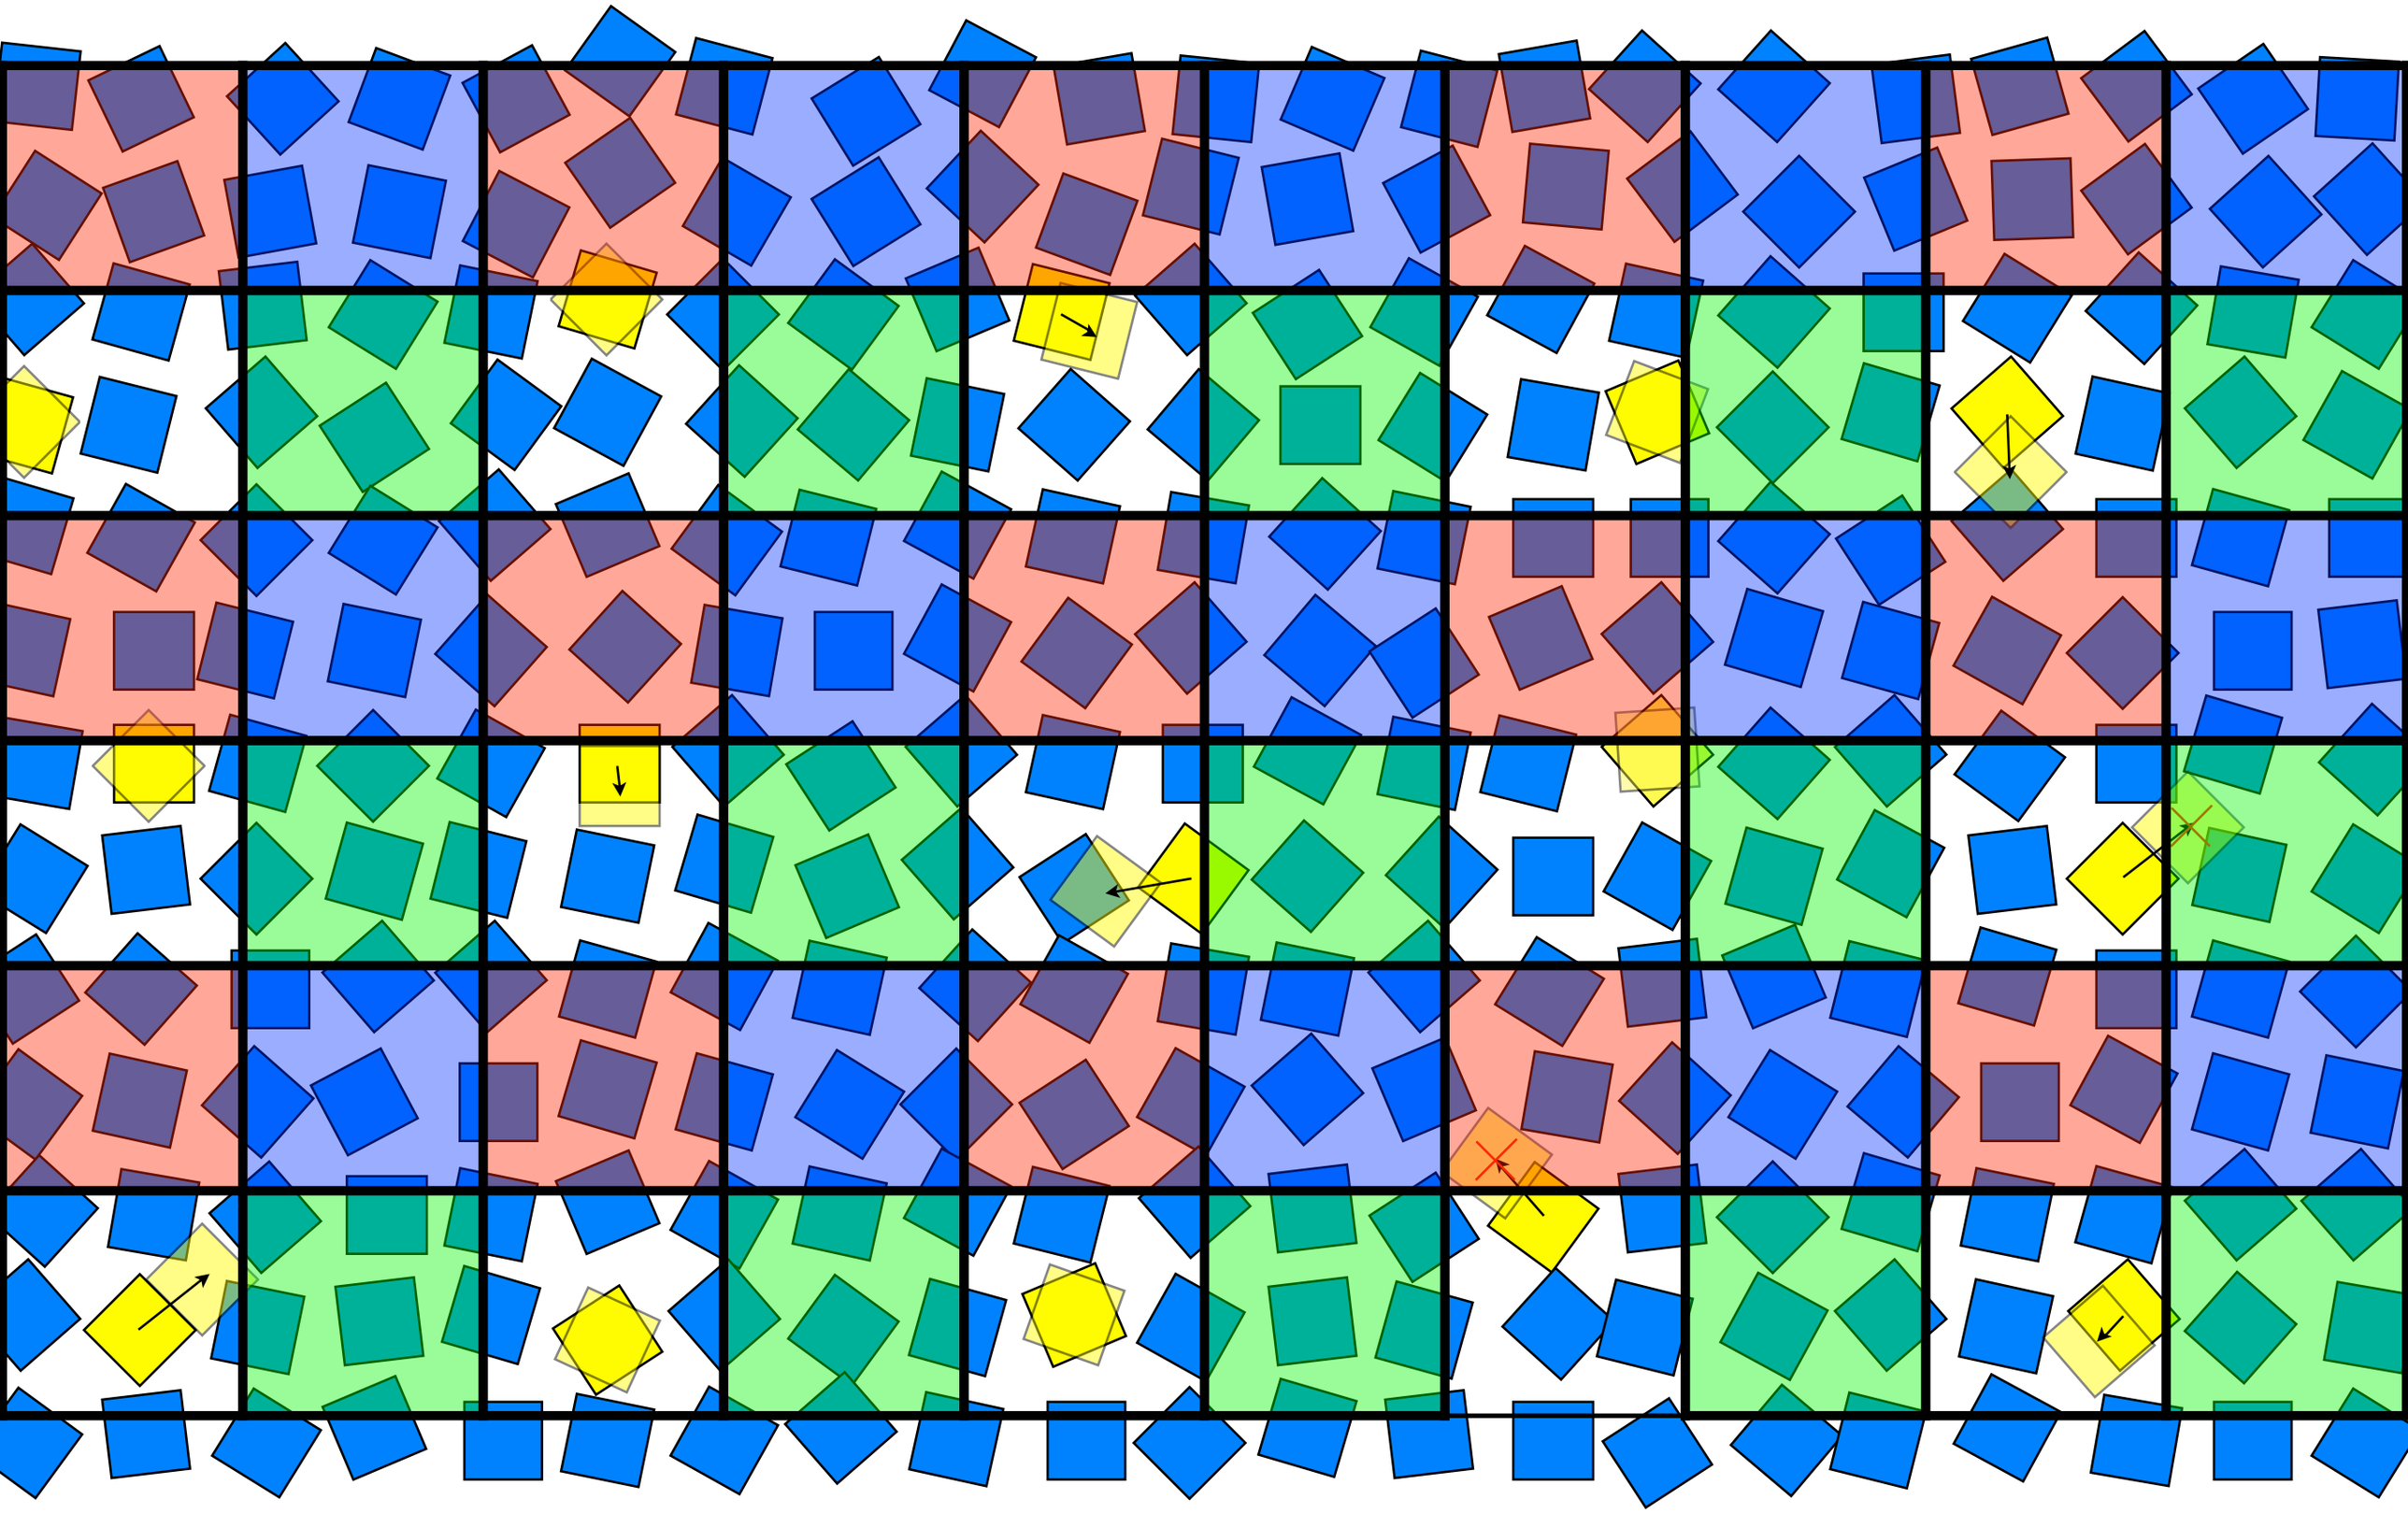

slide by Joshua Anderson

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)
```

Depletion Interaction

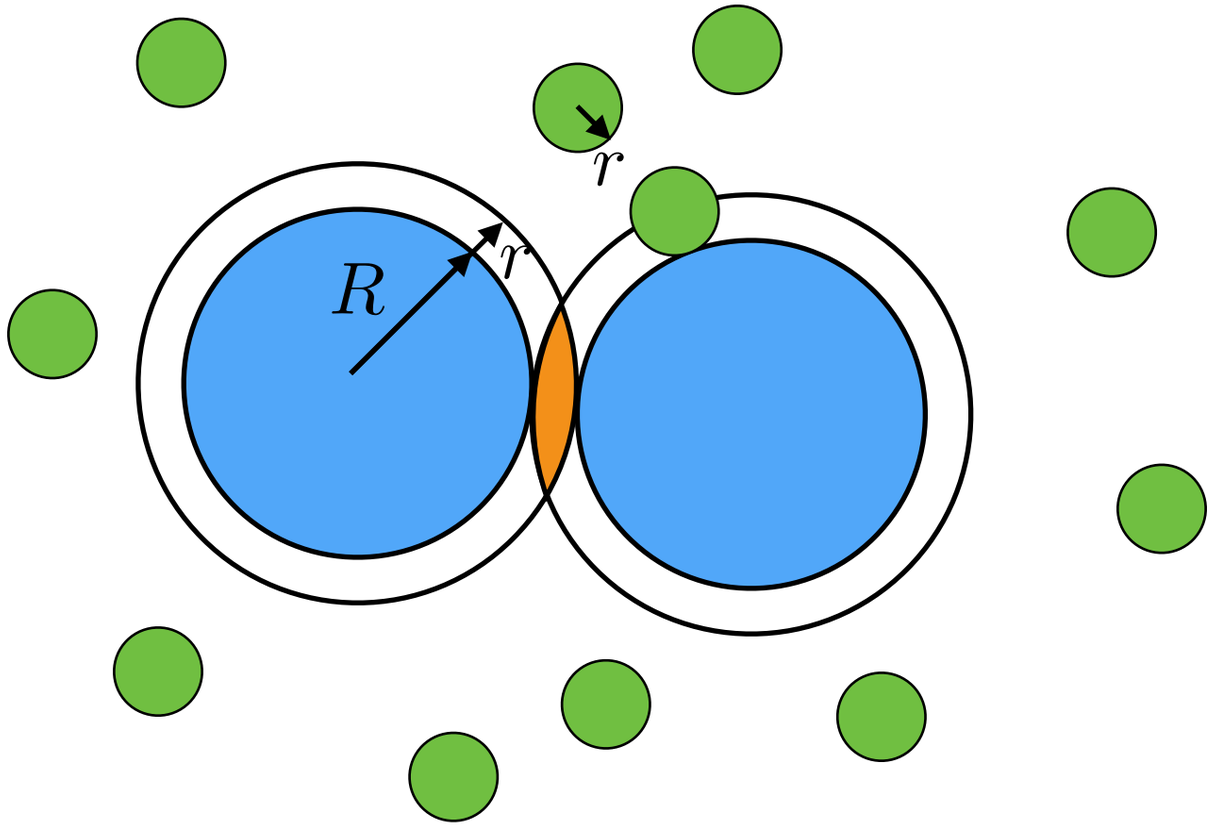

S. Asakura and F. Oosawa J. Chem. Phys. 1954.

Implicit penetrable hard-sphere depletants Monte Carlo

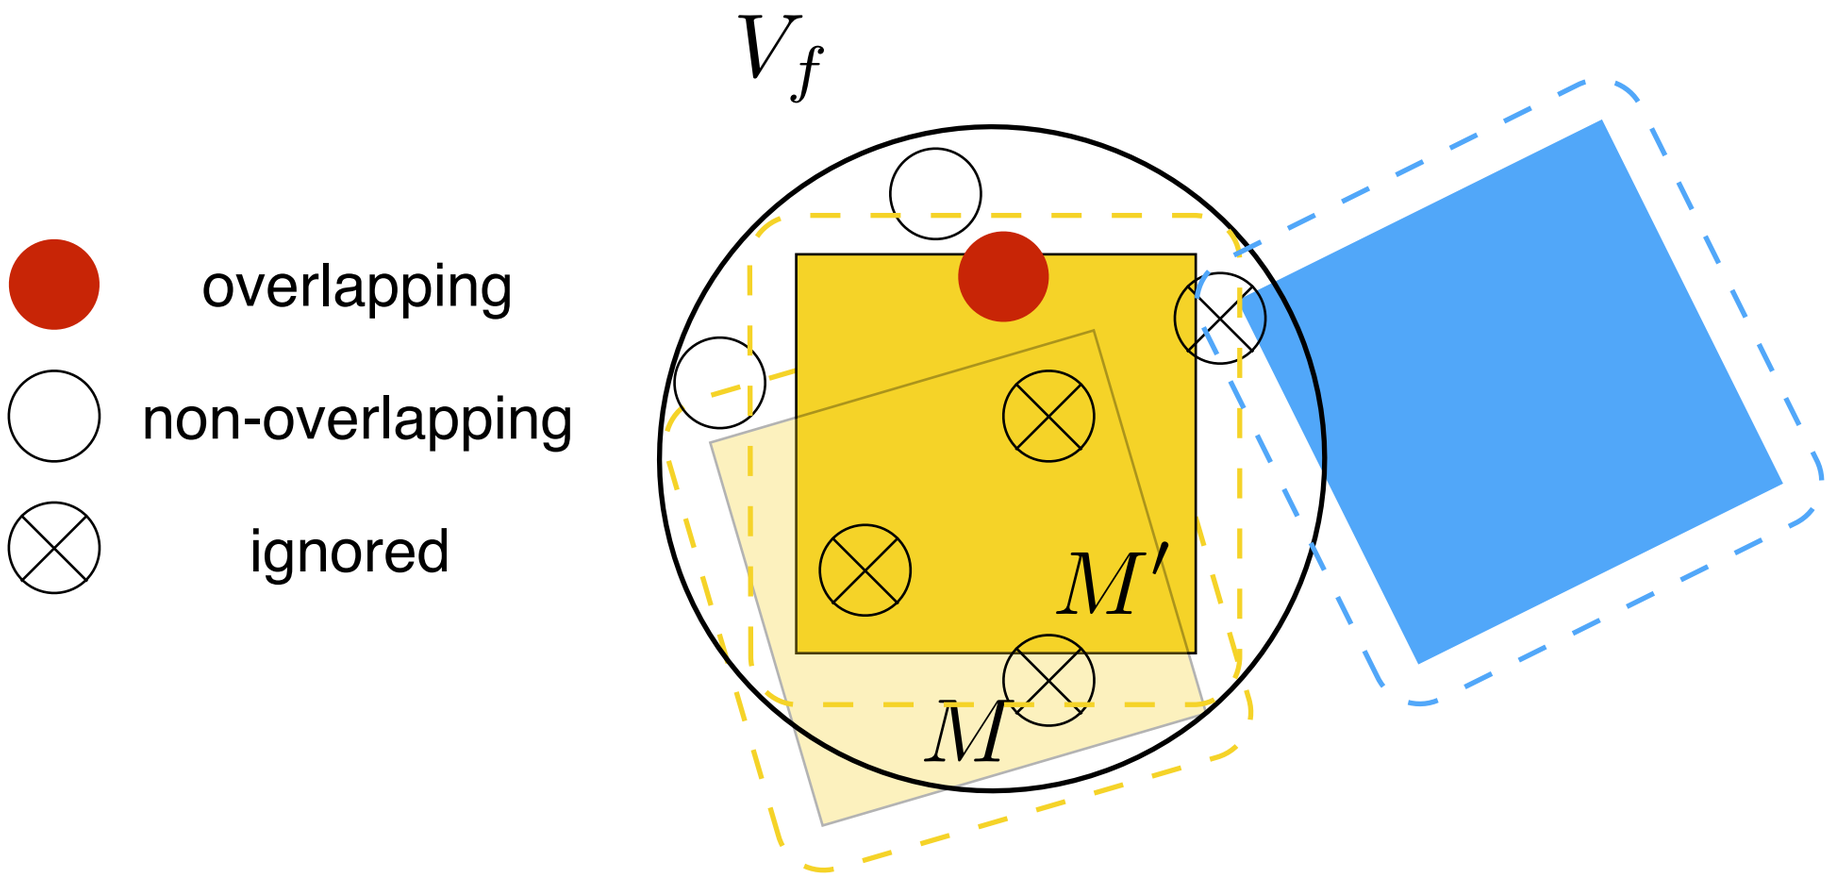

$$e^{-\beta \Xi\{\vec{r}_{c,i}\}} = e^{z_p V_f - \beta H_{cc}}$$

z_p depletant fugacity

parallel depletant insertion

1. Propose a trial move for the colloids $M \to M'$
2. Generate N_p of depletant positions $\vec{r}_i^{(p)}$ randomly in the free volume of the old configuration M according to $P_{z_p V_f}(N_p) \sim \text{Poisson}(V_f z_p)$. One possibility is to use rejection sampling in a larger volume $V_0 \supset V_f$.
3. Reject the trial move if any depletant overlaps with new colloid configuration M', otherwise accept.

Validation: equation of state of hard spheres

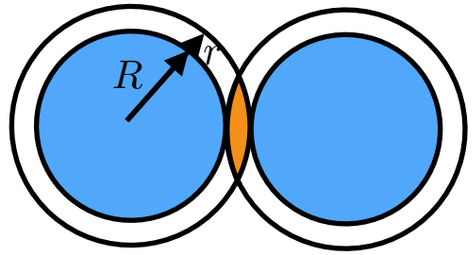

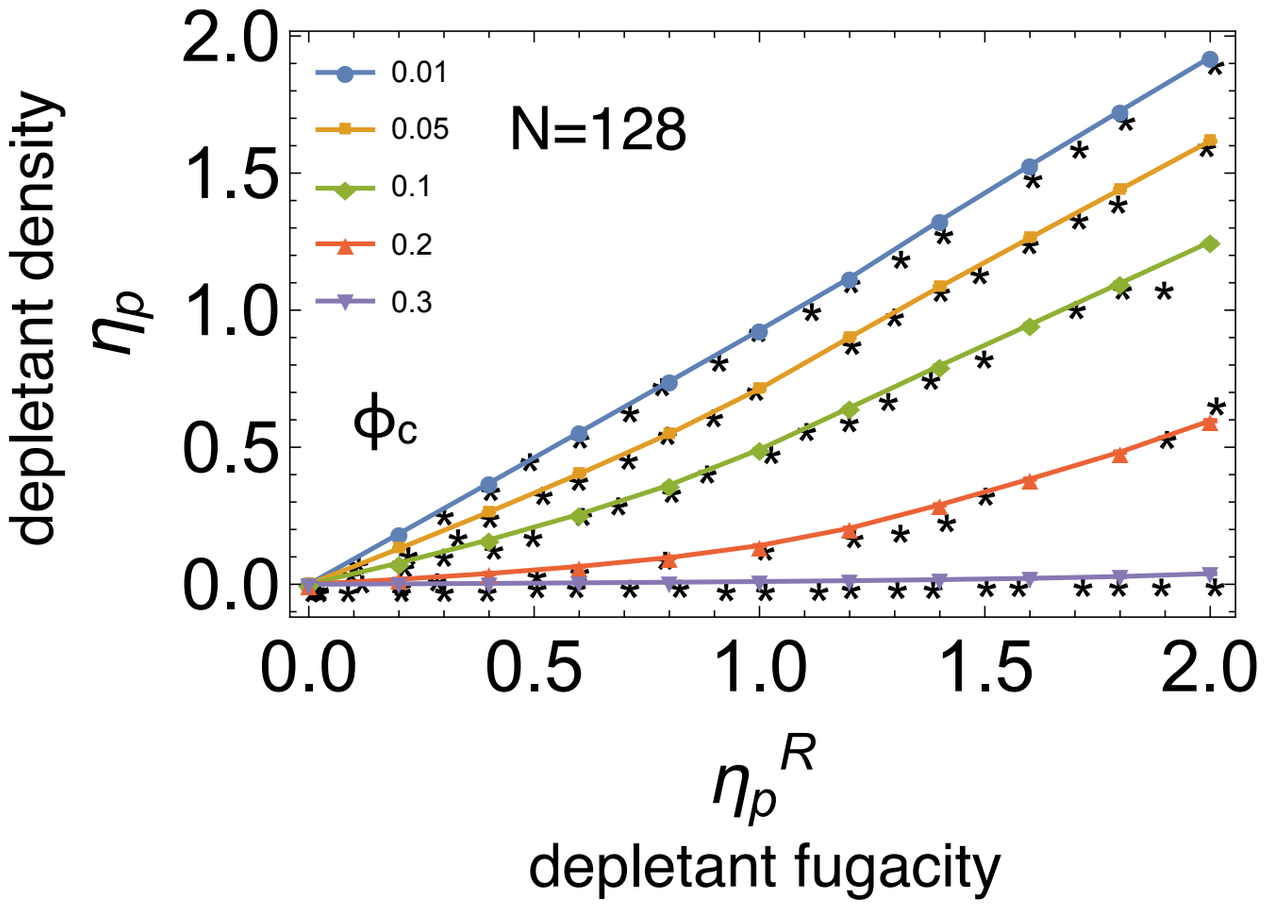

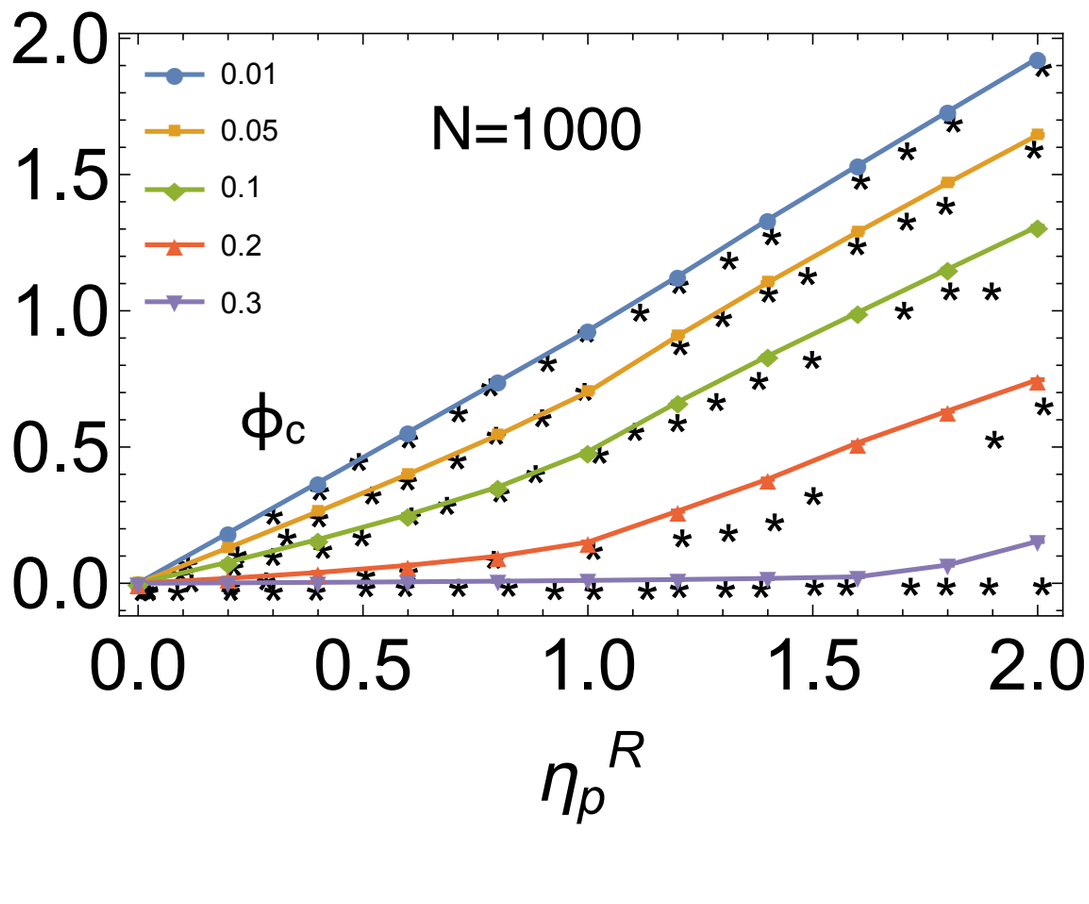

* Dijkstra M, van Roij R, Roth R, Fortini A.

Phys. Rev. E 2006 **73** 041404
doi: 10.1103/PhysRevE.73.041404

Summary - algorithms

- MD is the perfect parallel algorithm
- Hard Particle Monte Carlo is serial, but can be parallelized using checkerboard decomposition
- Simulations of large and small particles can be further parallelized

Scaling bottlenecks in spatial domain decomposition

4-12 cores

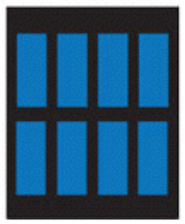

6 GB/s

CPU

CPU

Network

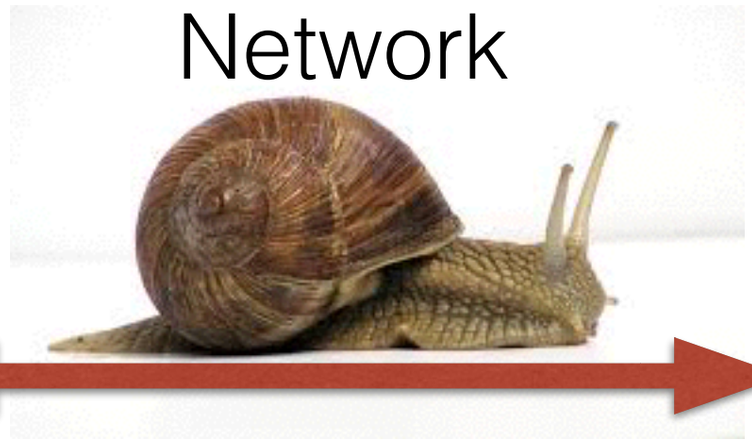

6 GB/s

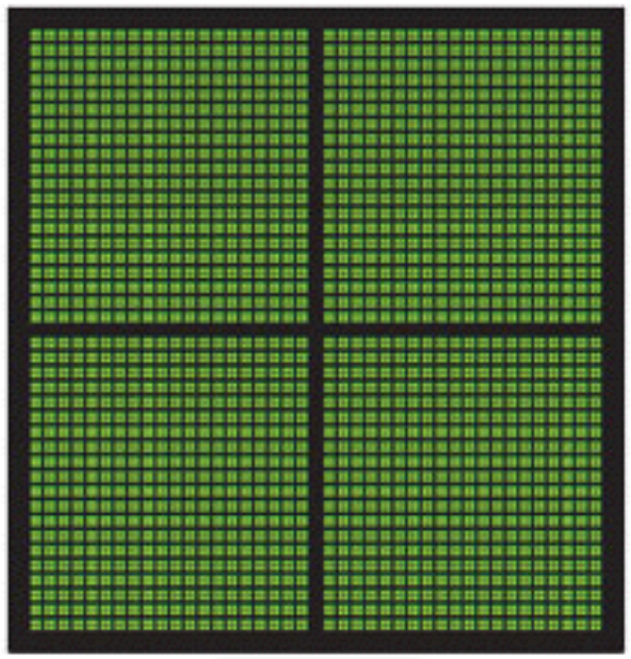

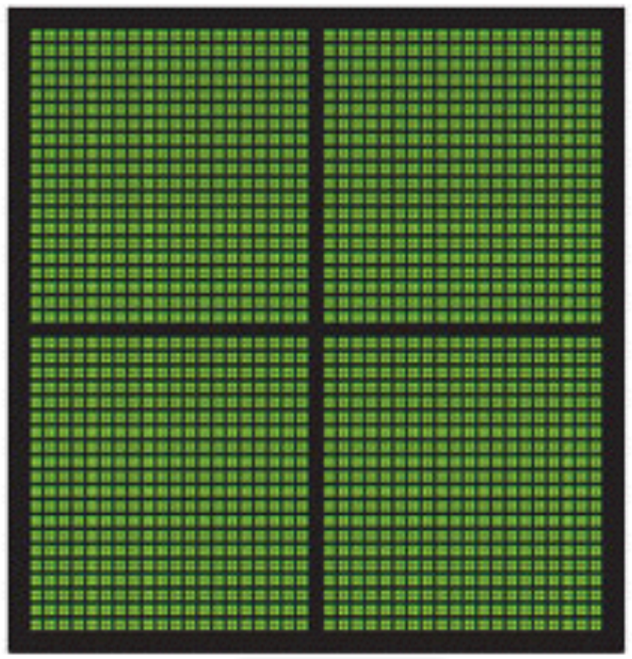

GPU

GPU

1000’s of cores

Weak scaling up to 108,000,000 particles

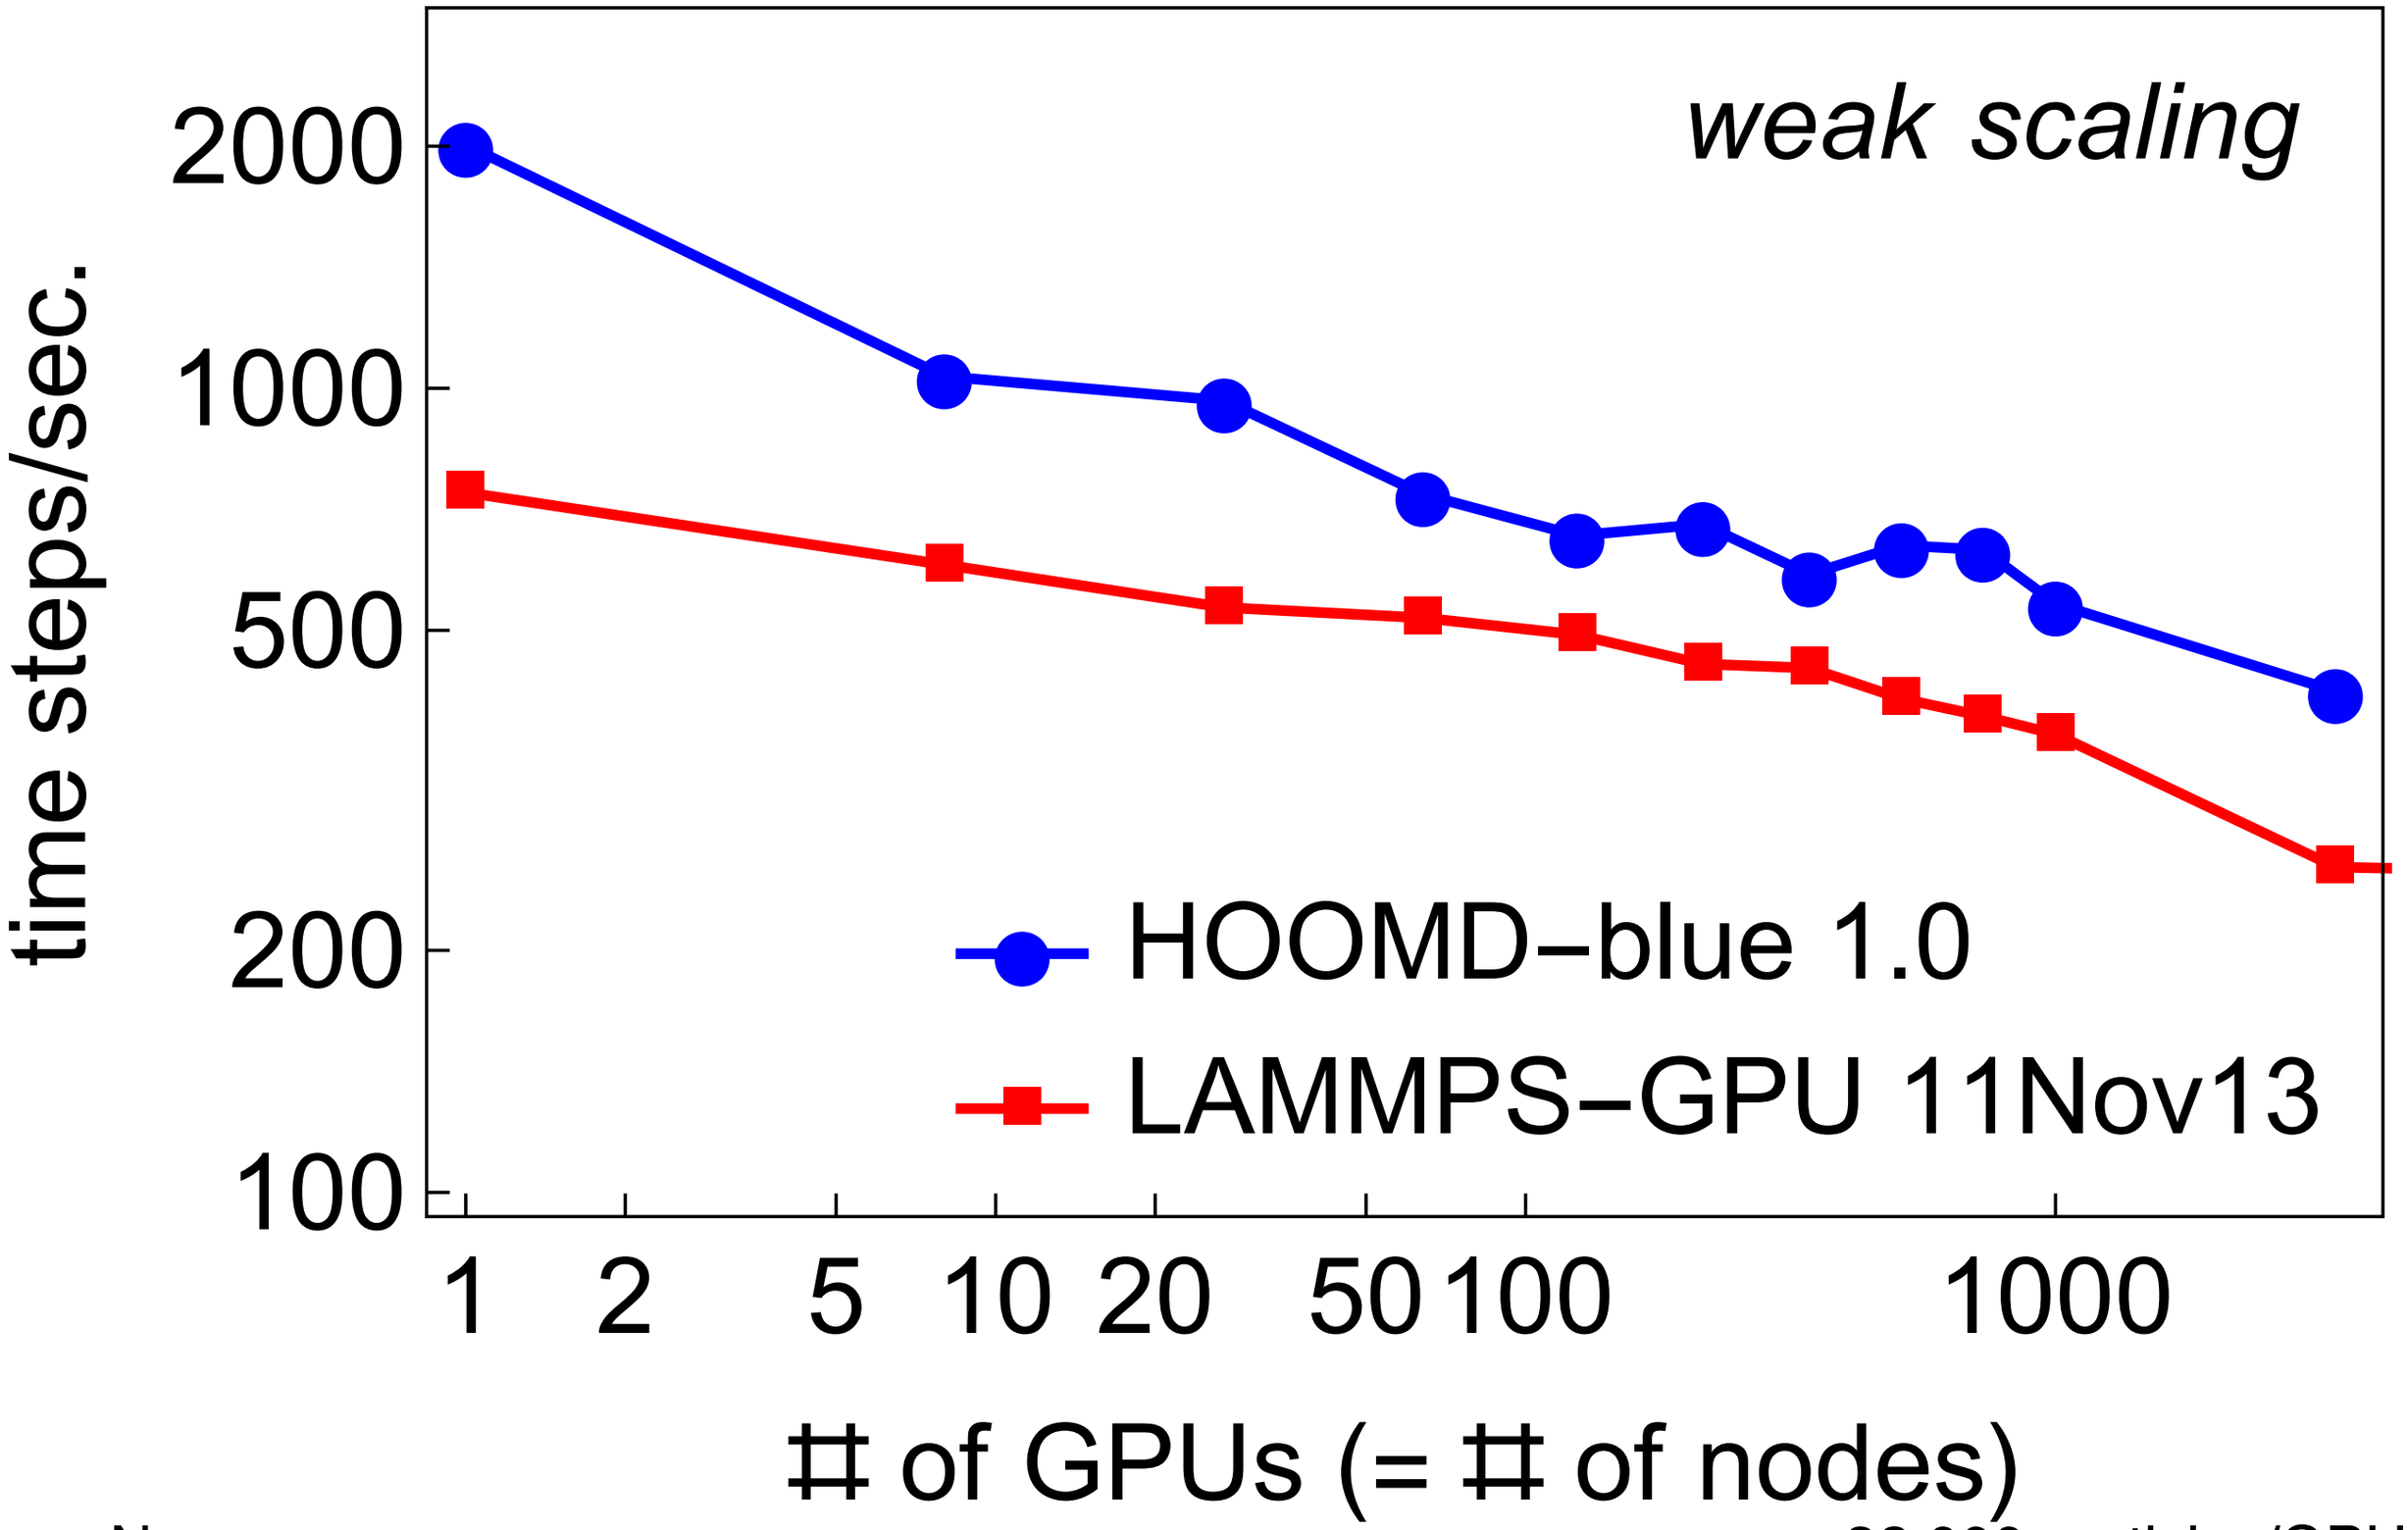

Trung Nguyen

32,000 particles/GPU

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

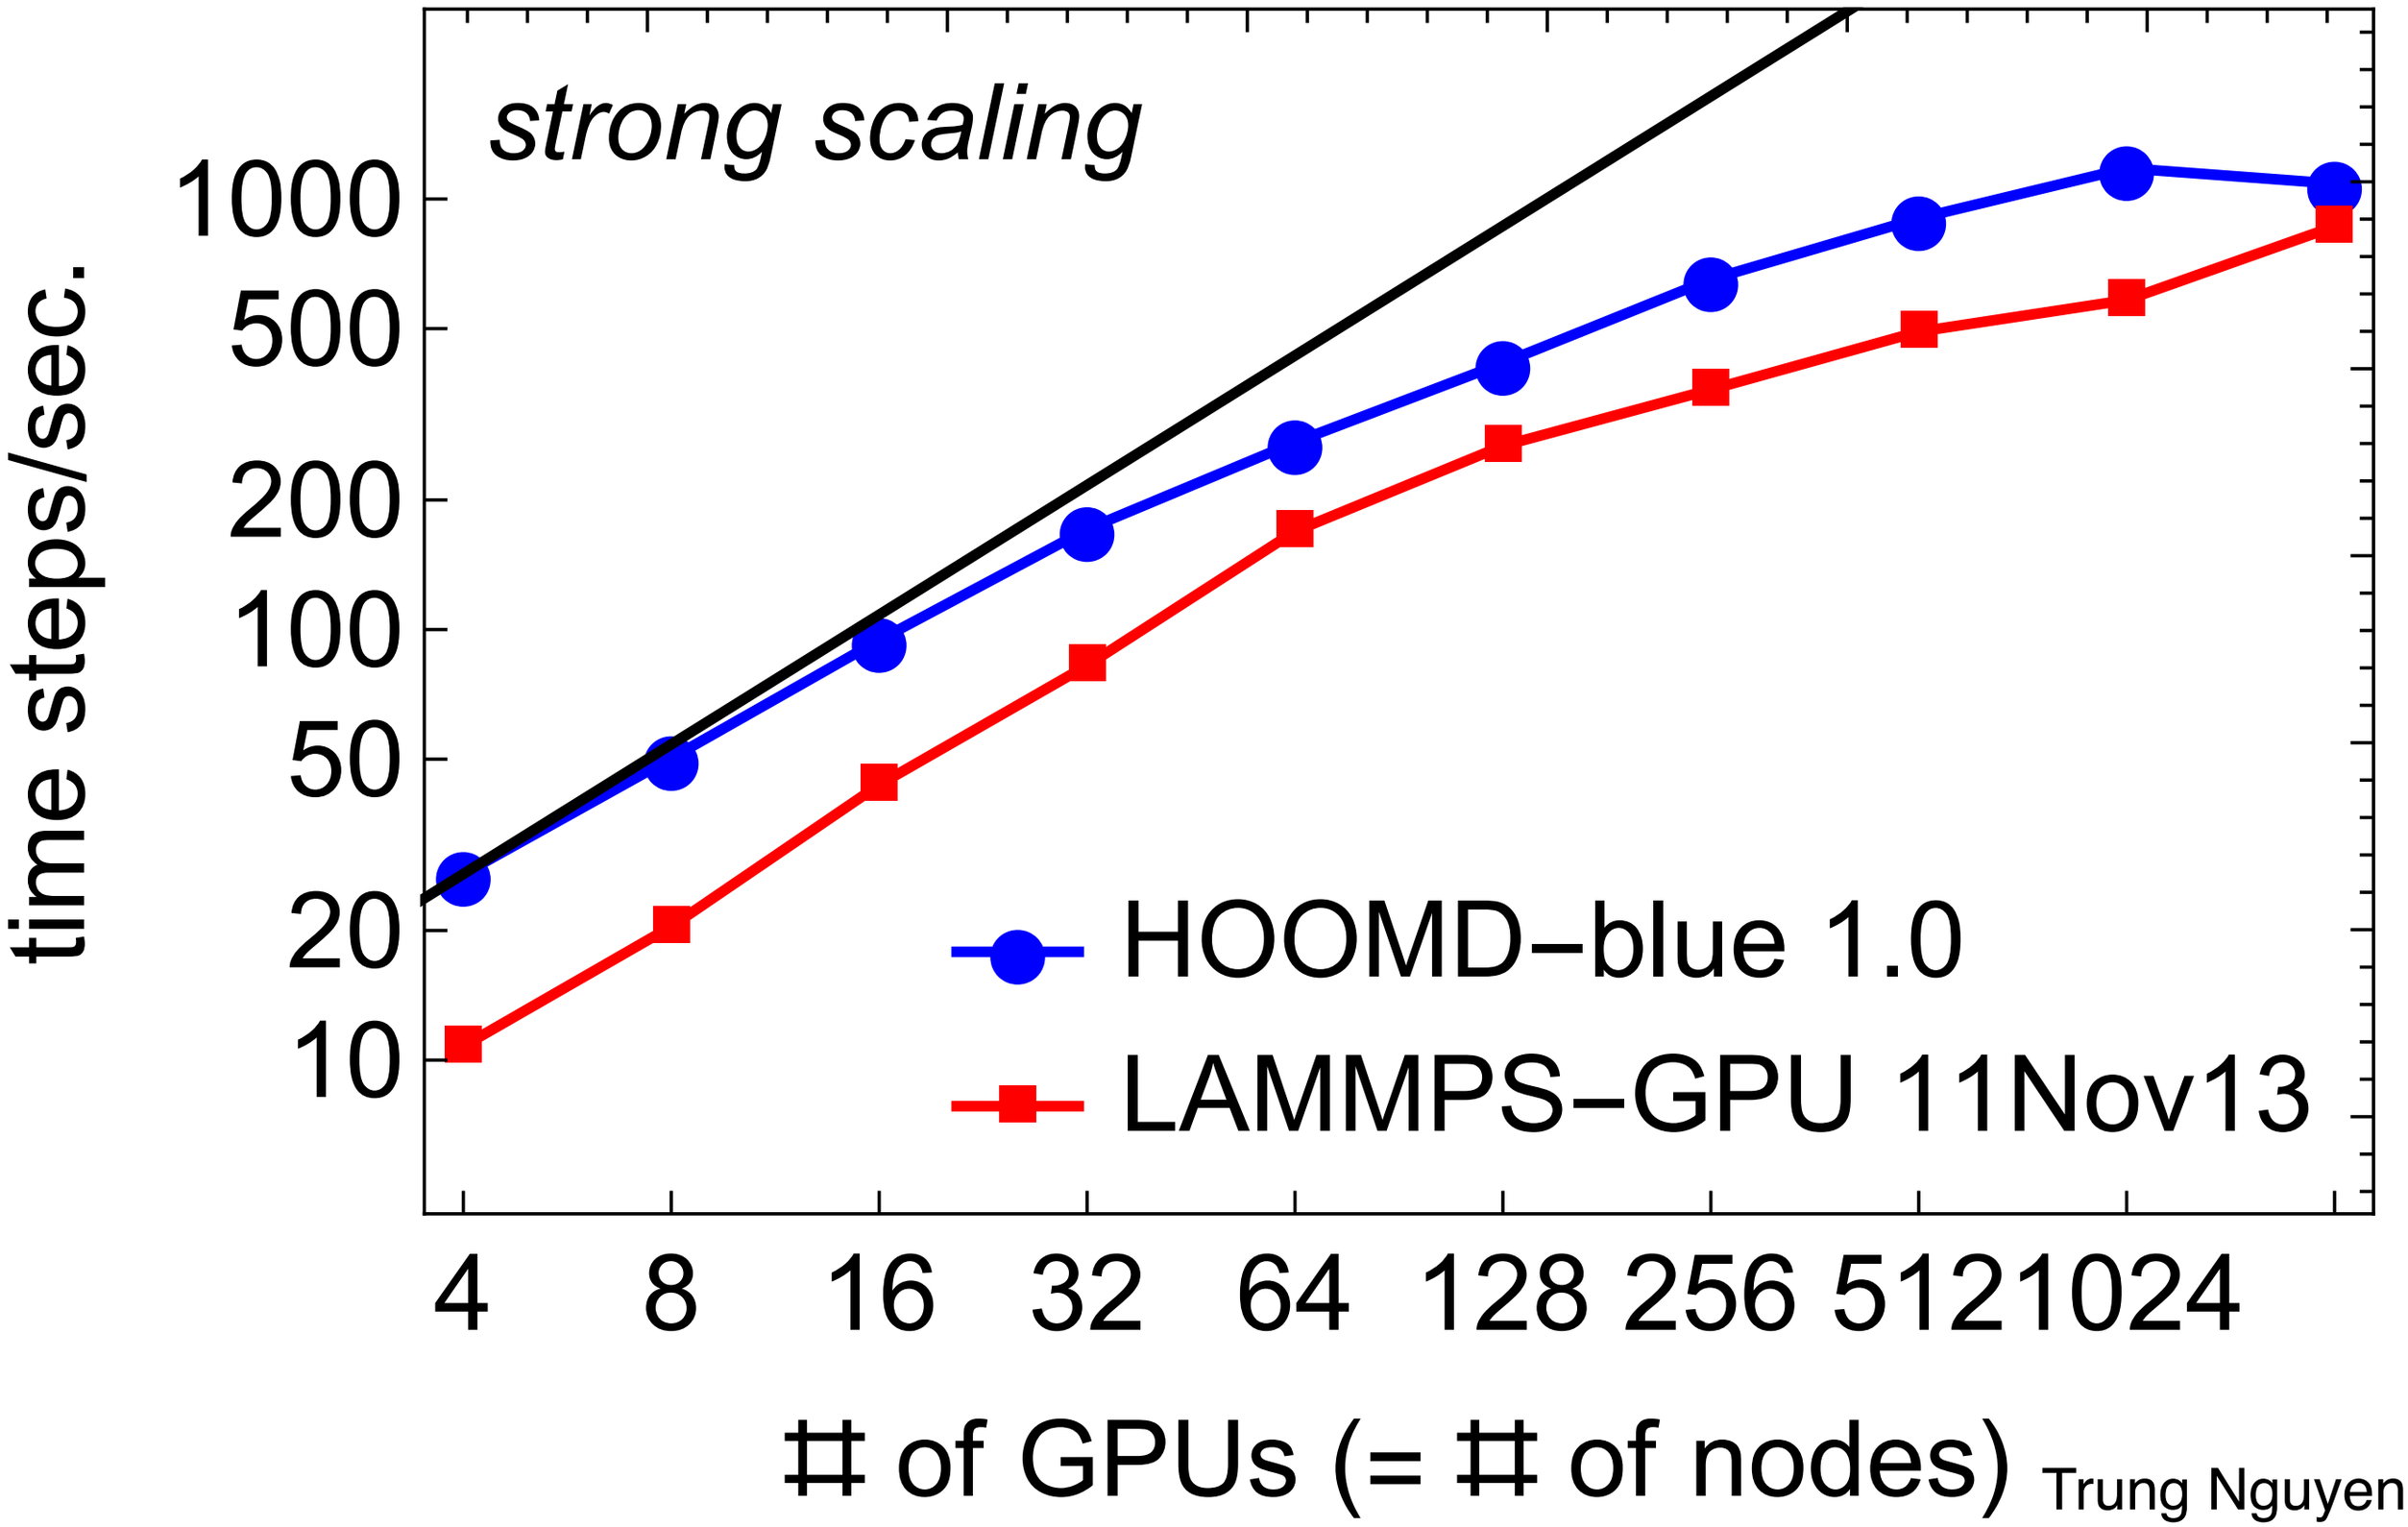

Trung Nguyen

Polymer Brush Scaling

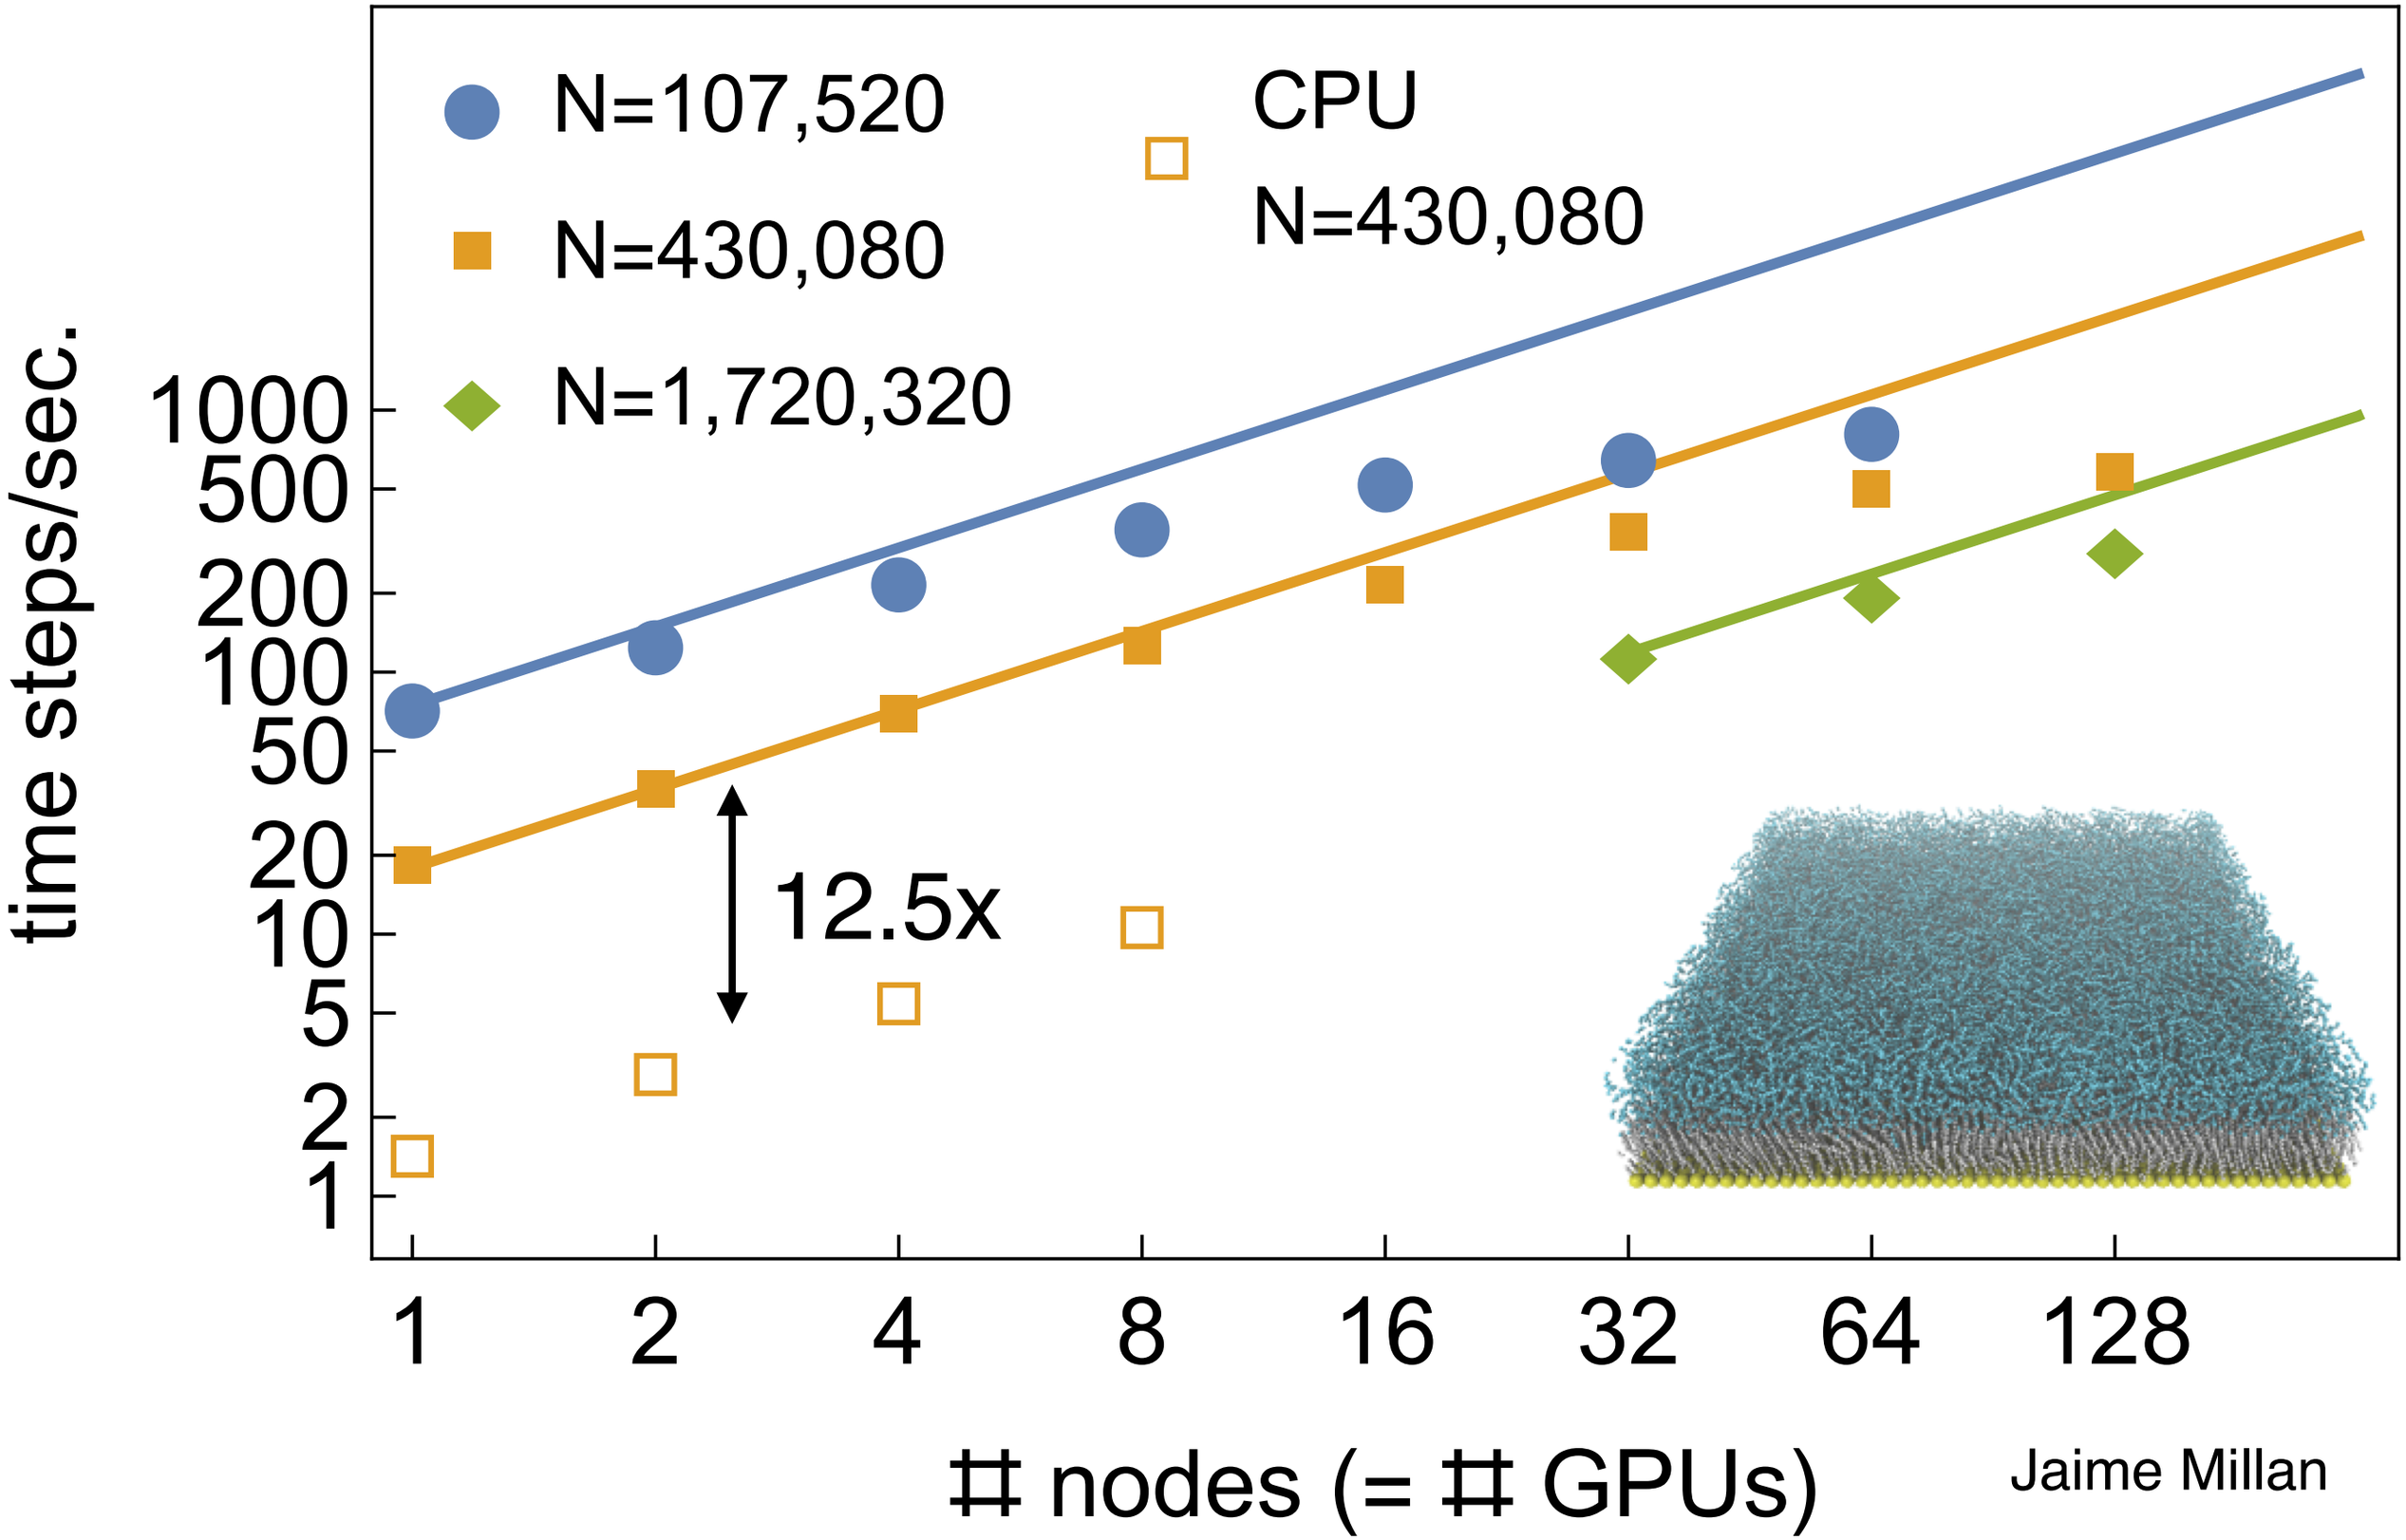

Jaime Millan

Dissipative Particle Dynamics on Blue Waters and Titan

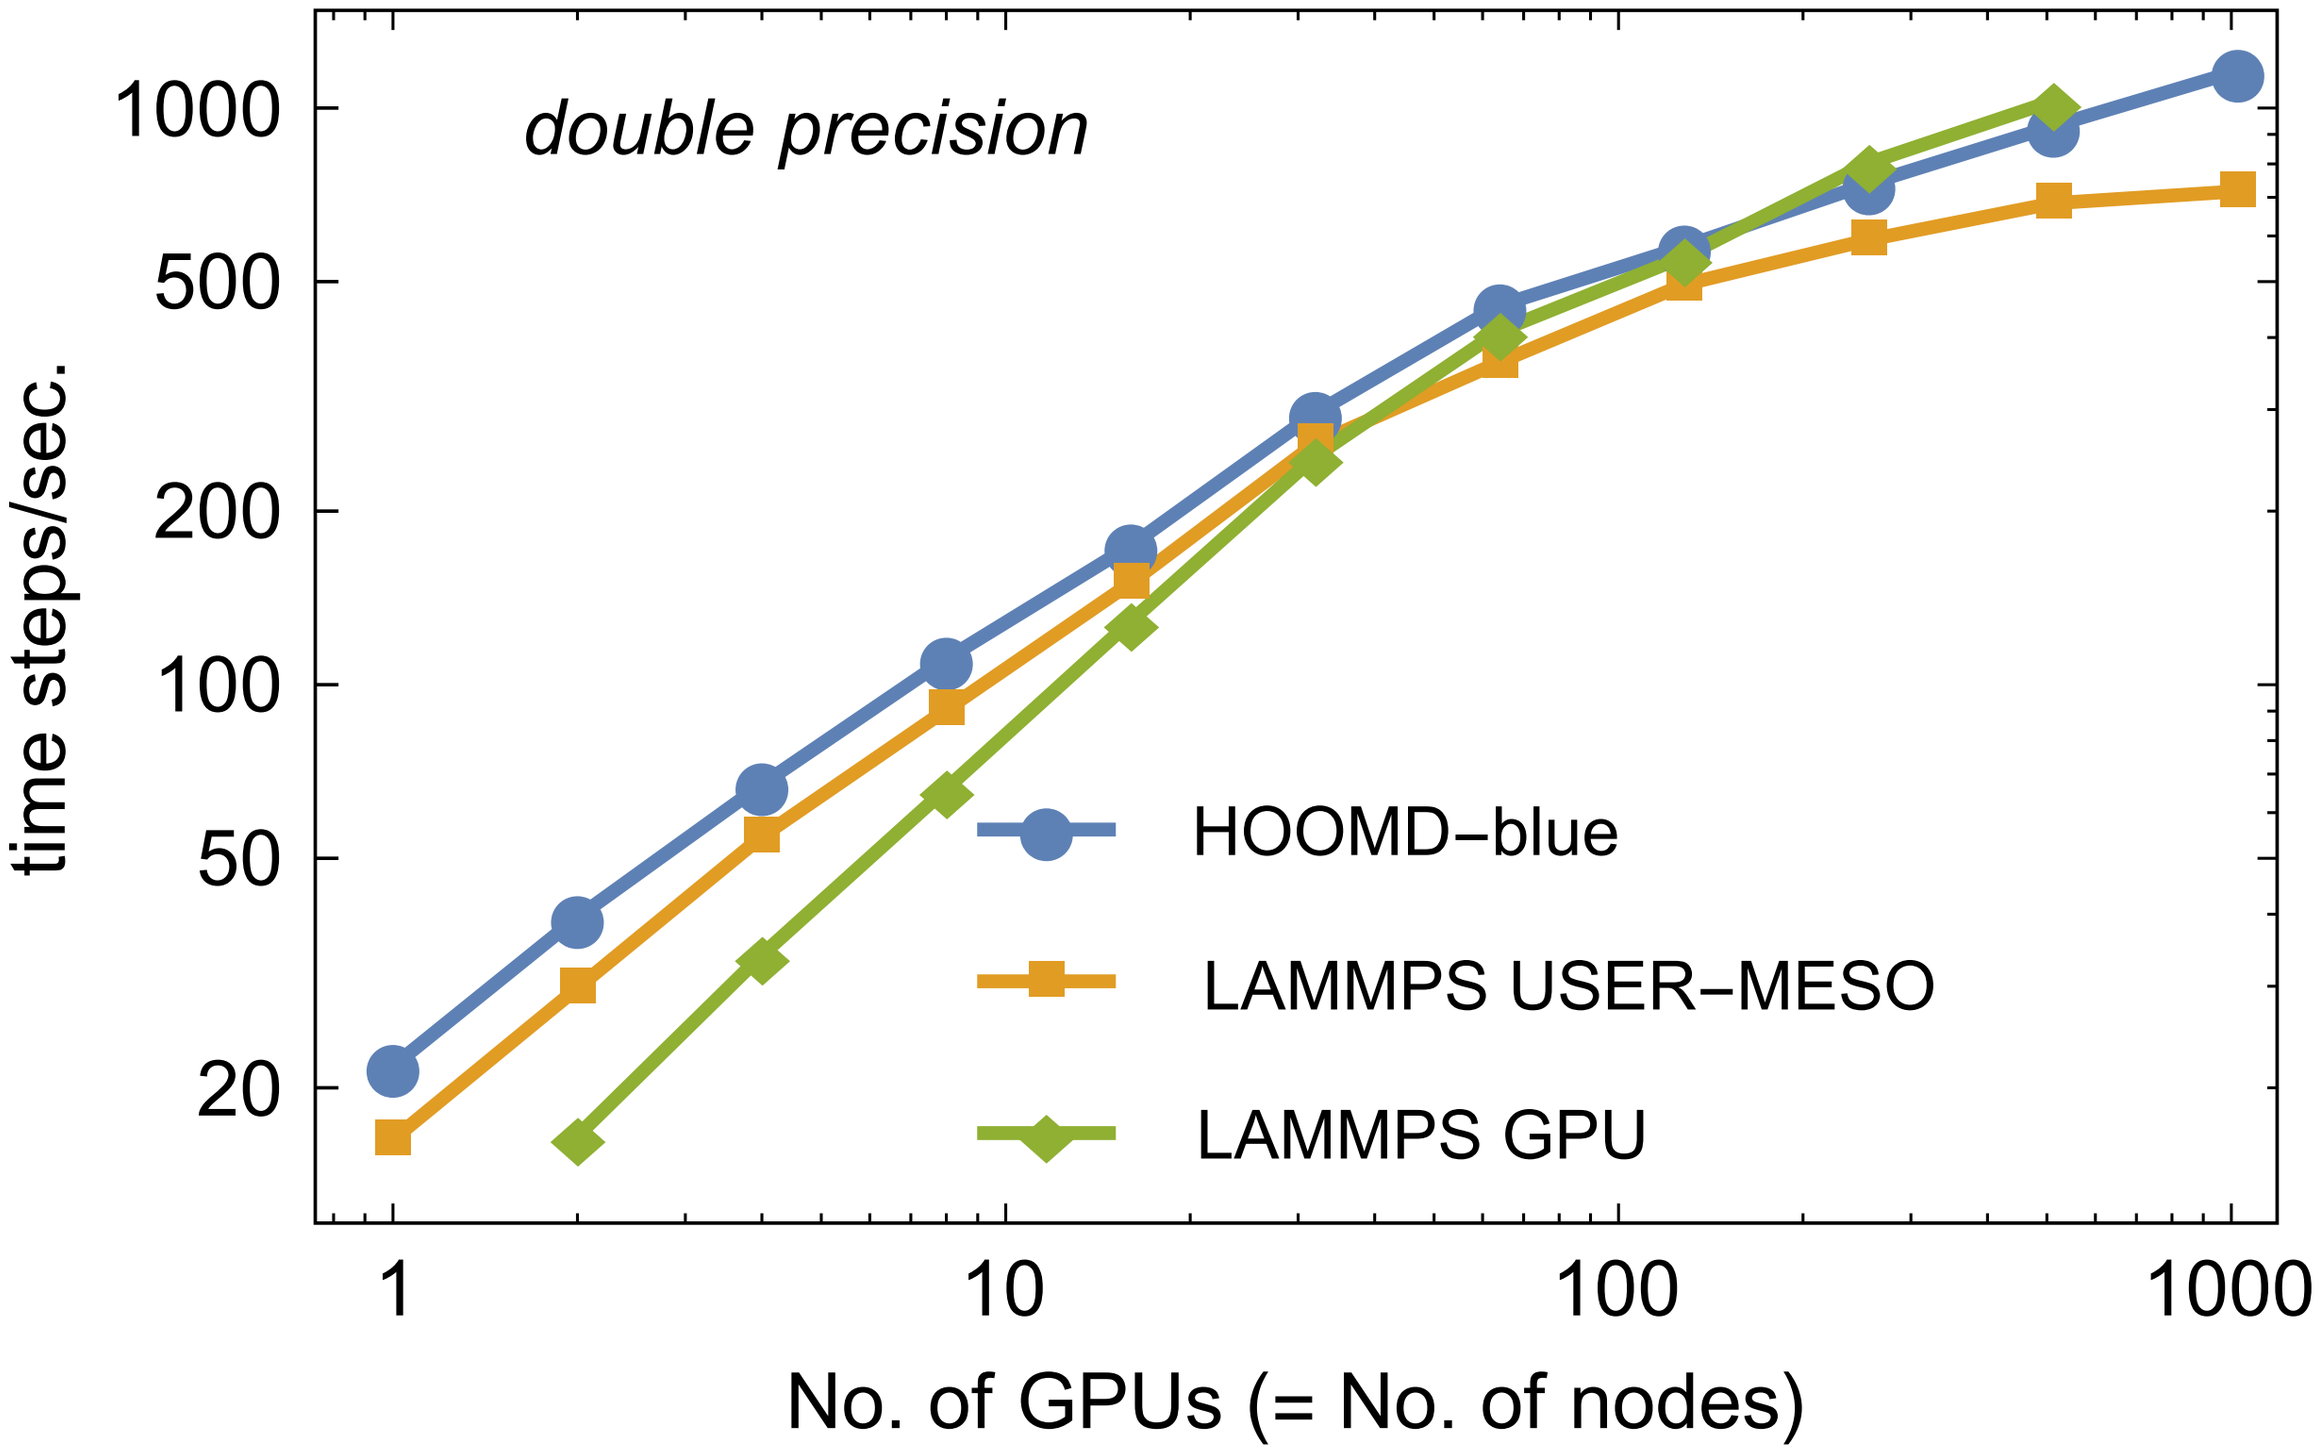

GPUDirect RDMA on Wilkes

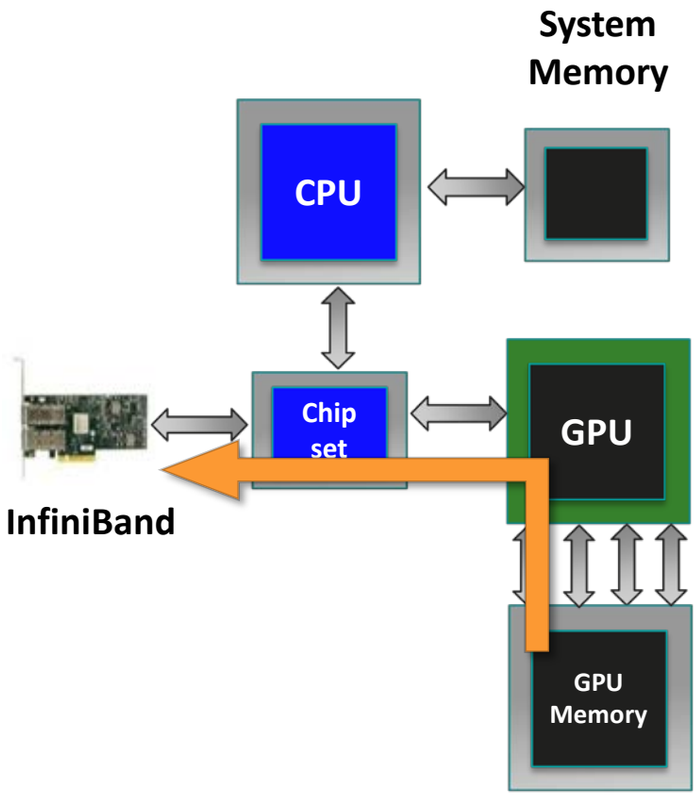

MVAPICH 2.1rc2?

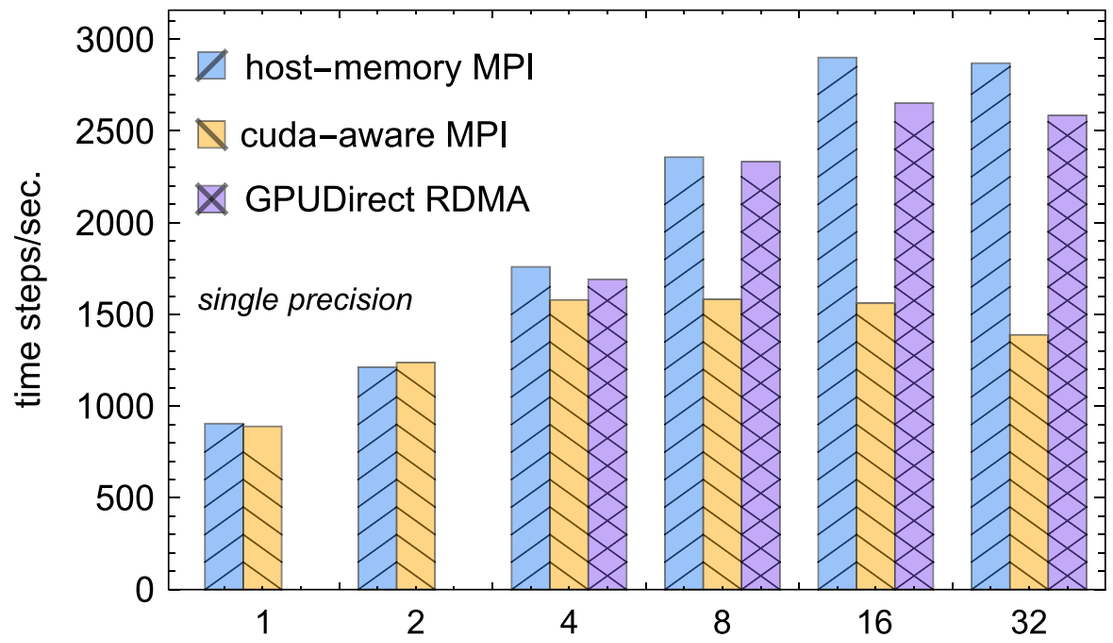

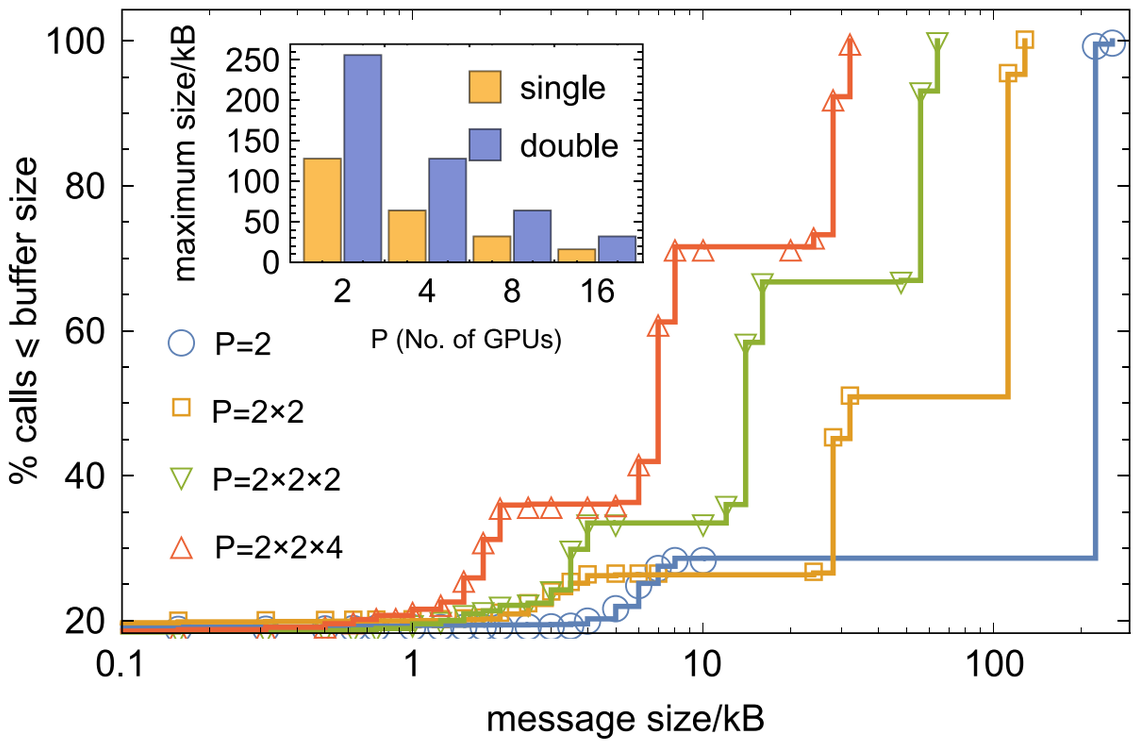

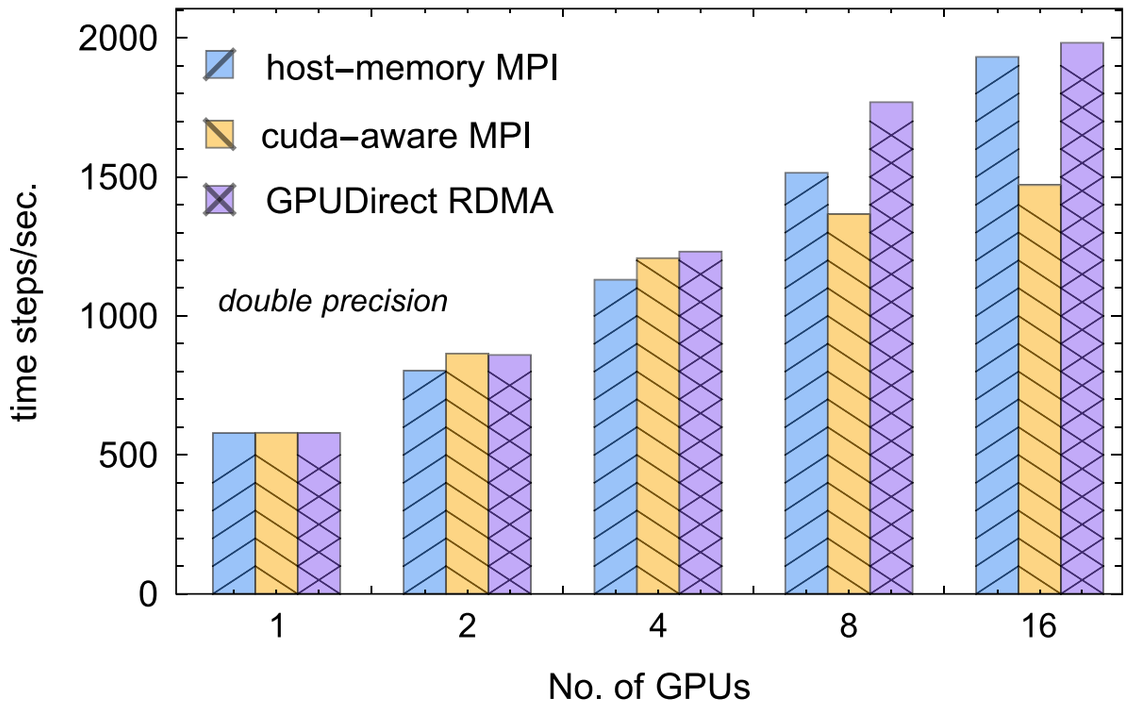

Pak Lui, Filippo Spiga, Rong Shi

Strong scaling - squares

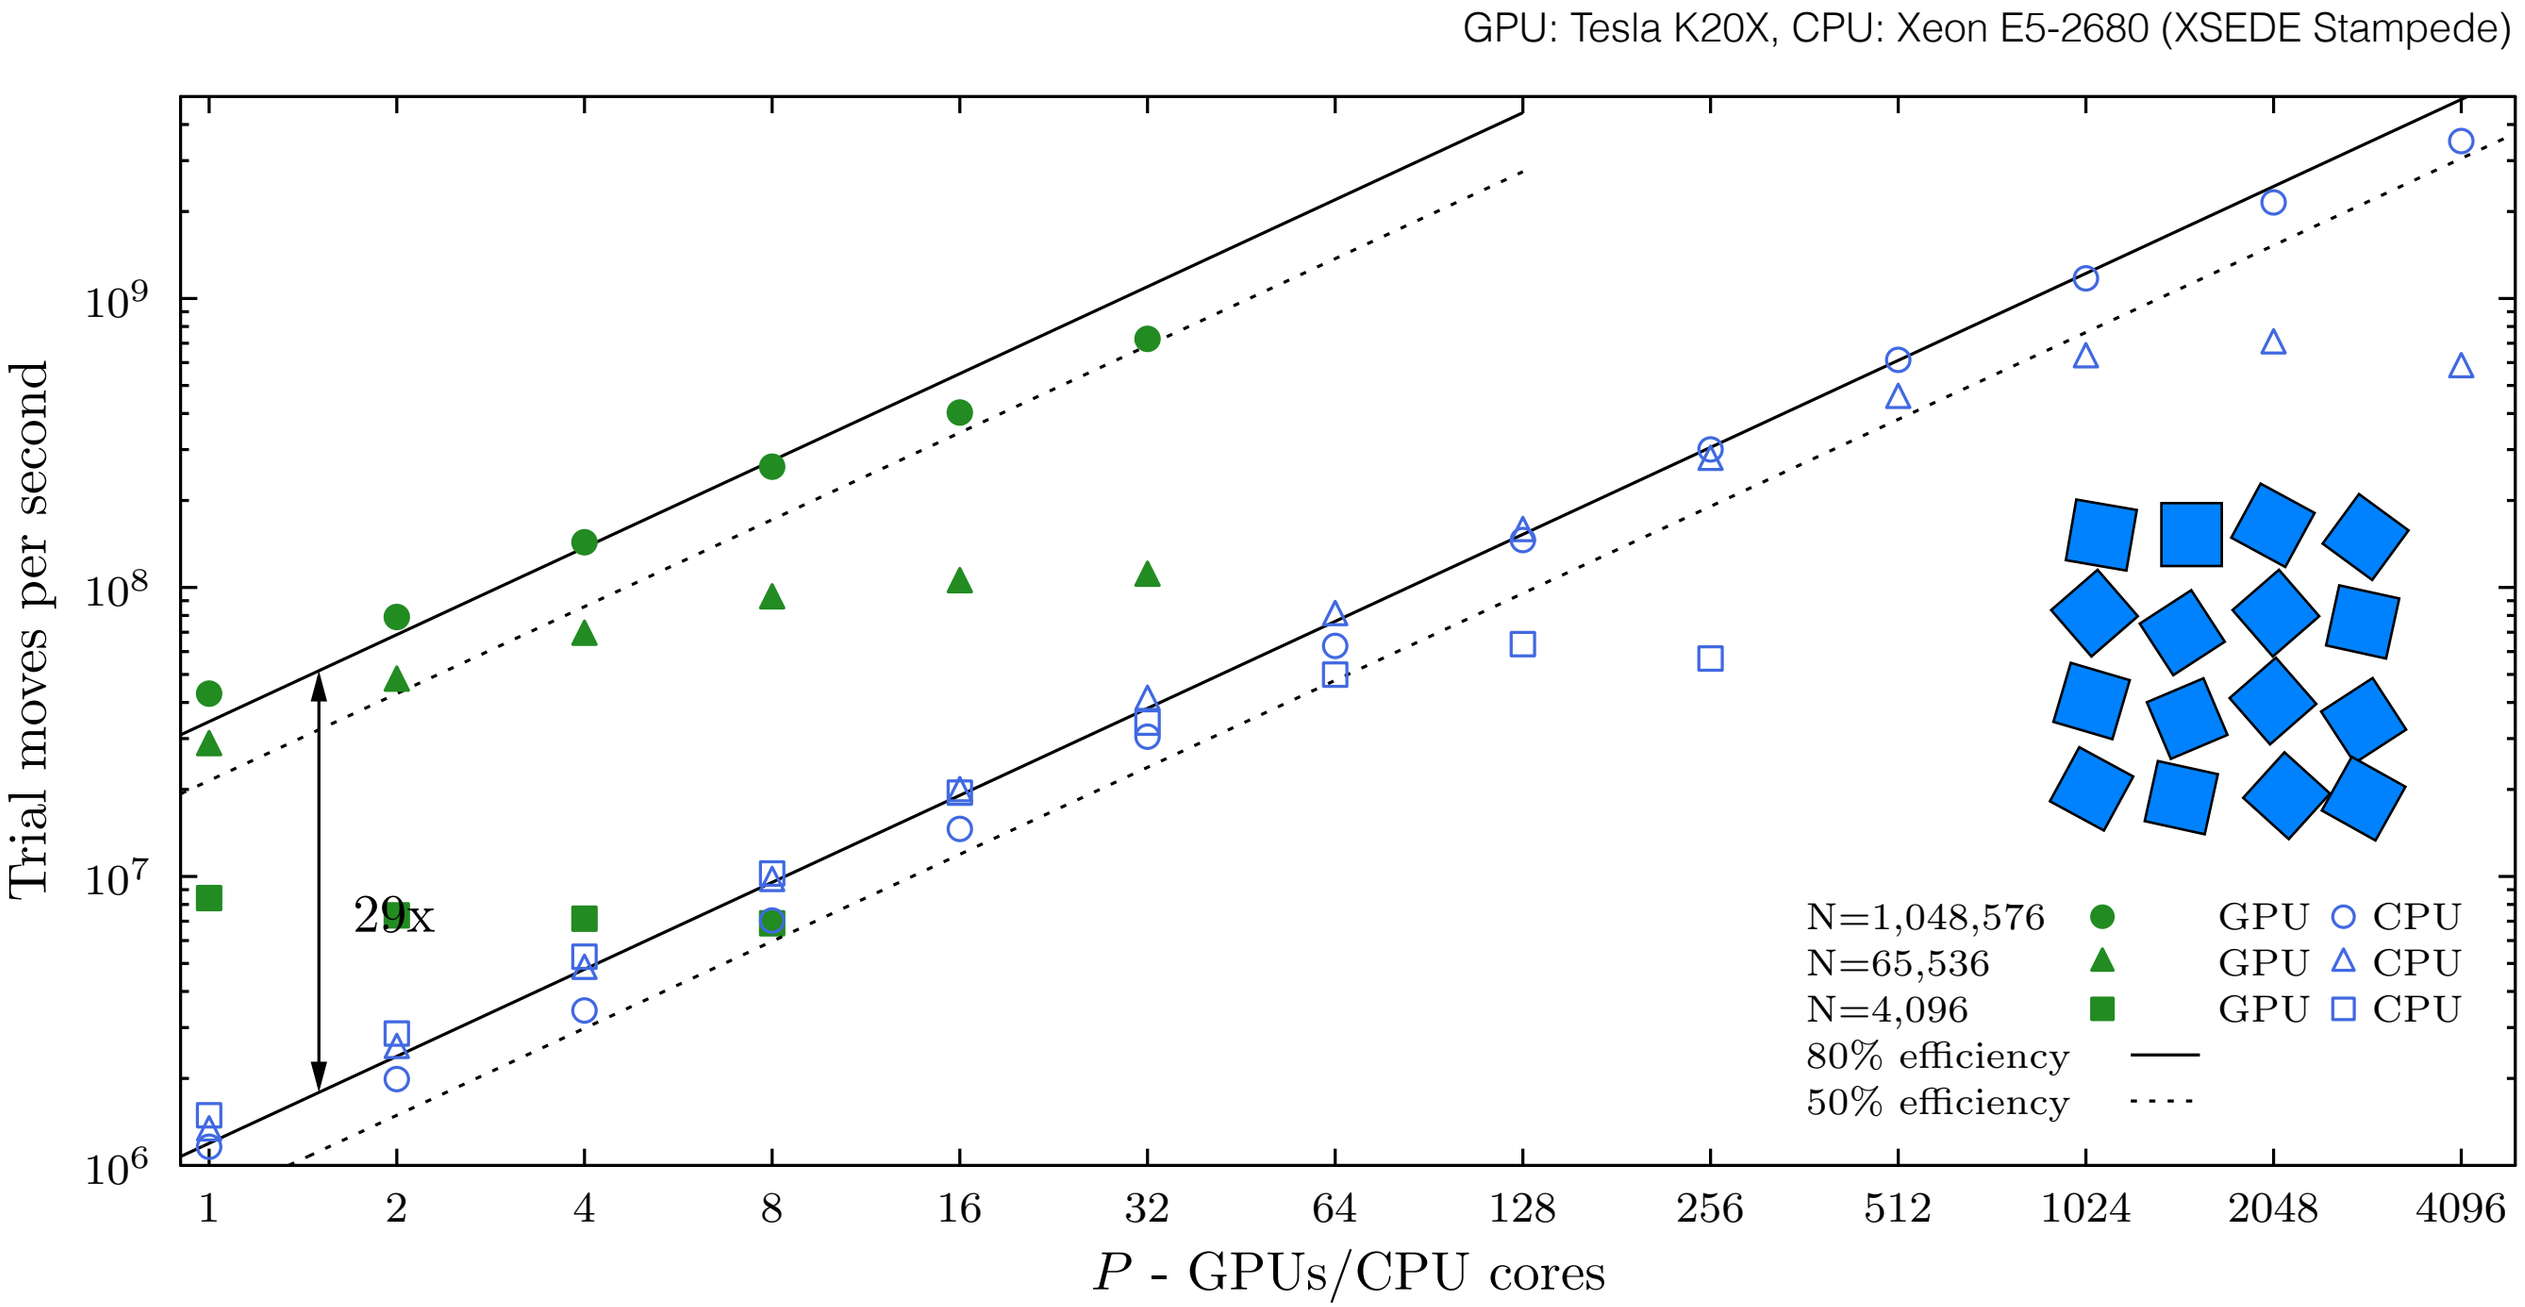

slide by Joshua Anderson

Summary - strong scaling

- For MD on GPUs, communication latency is the scaling bottleneck

- Strong scaling extends to 1000's of GPUs

- With Hard Particle Monte Carlo, CPU cores have higher scalability

Self-assembly of diblock copolymers

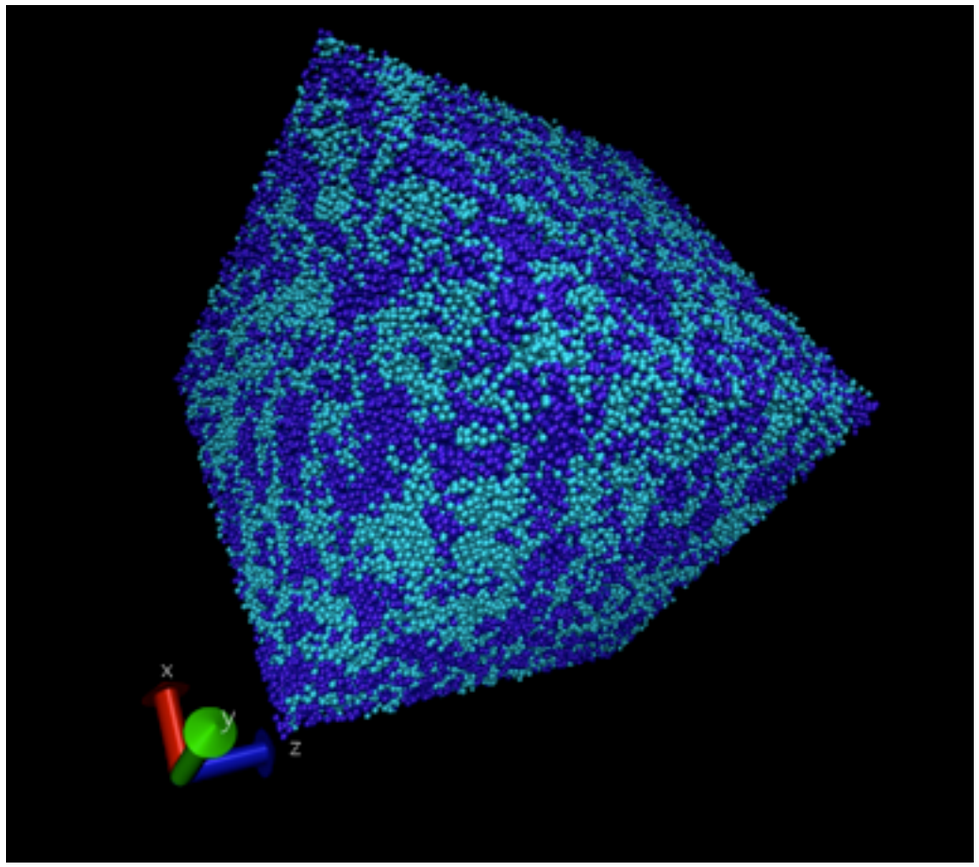

AB Diblock copolymer melt

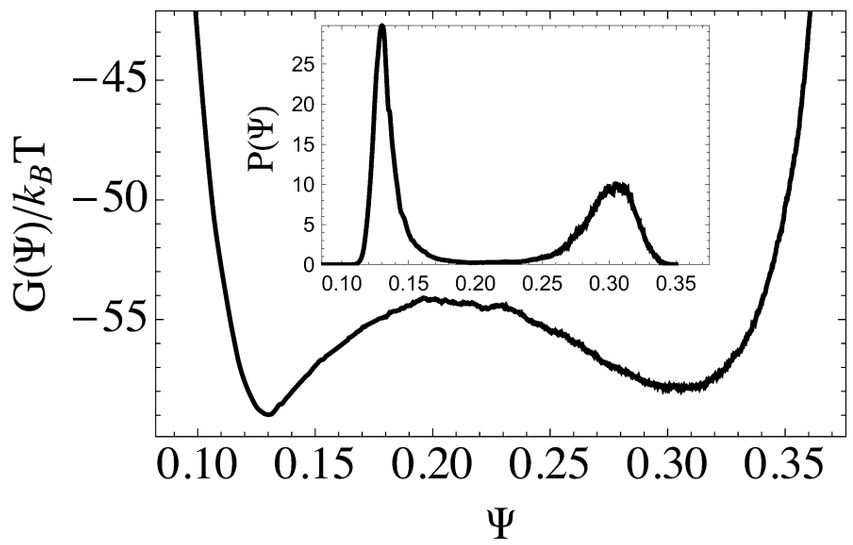

MD on 384 GPUs

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

Medapuram P., Glaser J., Morse D. C. Macromolecules 2015, **48**, 819-839.

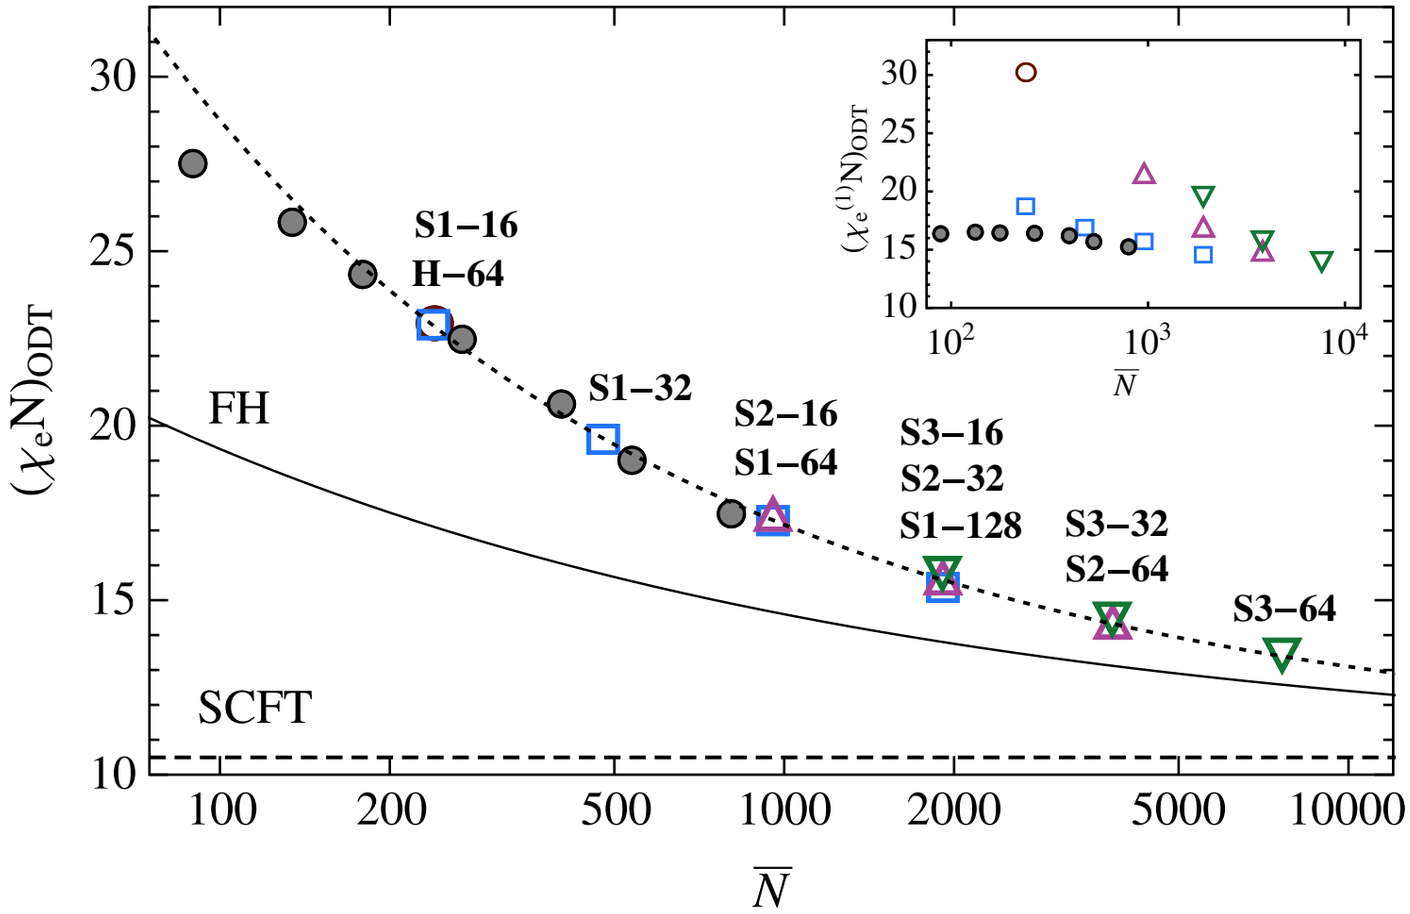

Discoids in depletants

200,000 explicit depletants - no way!
(even with MPI)

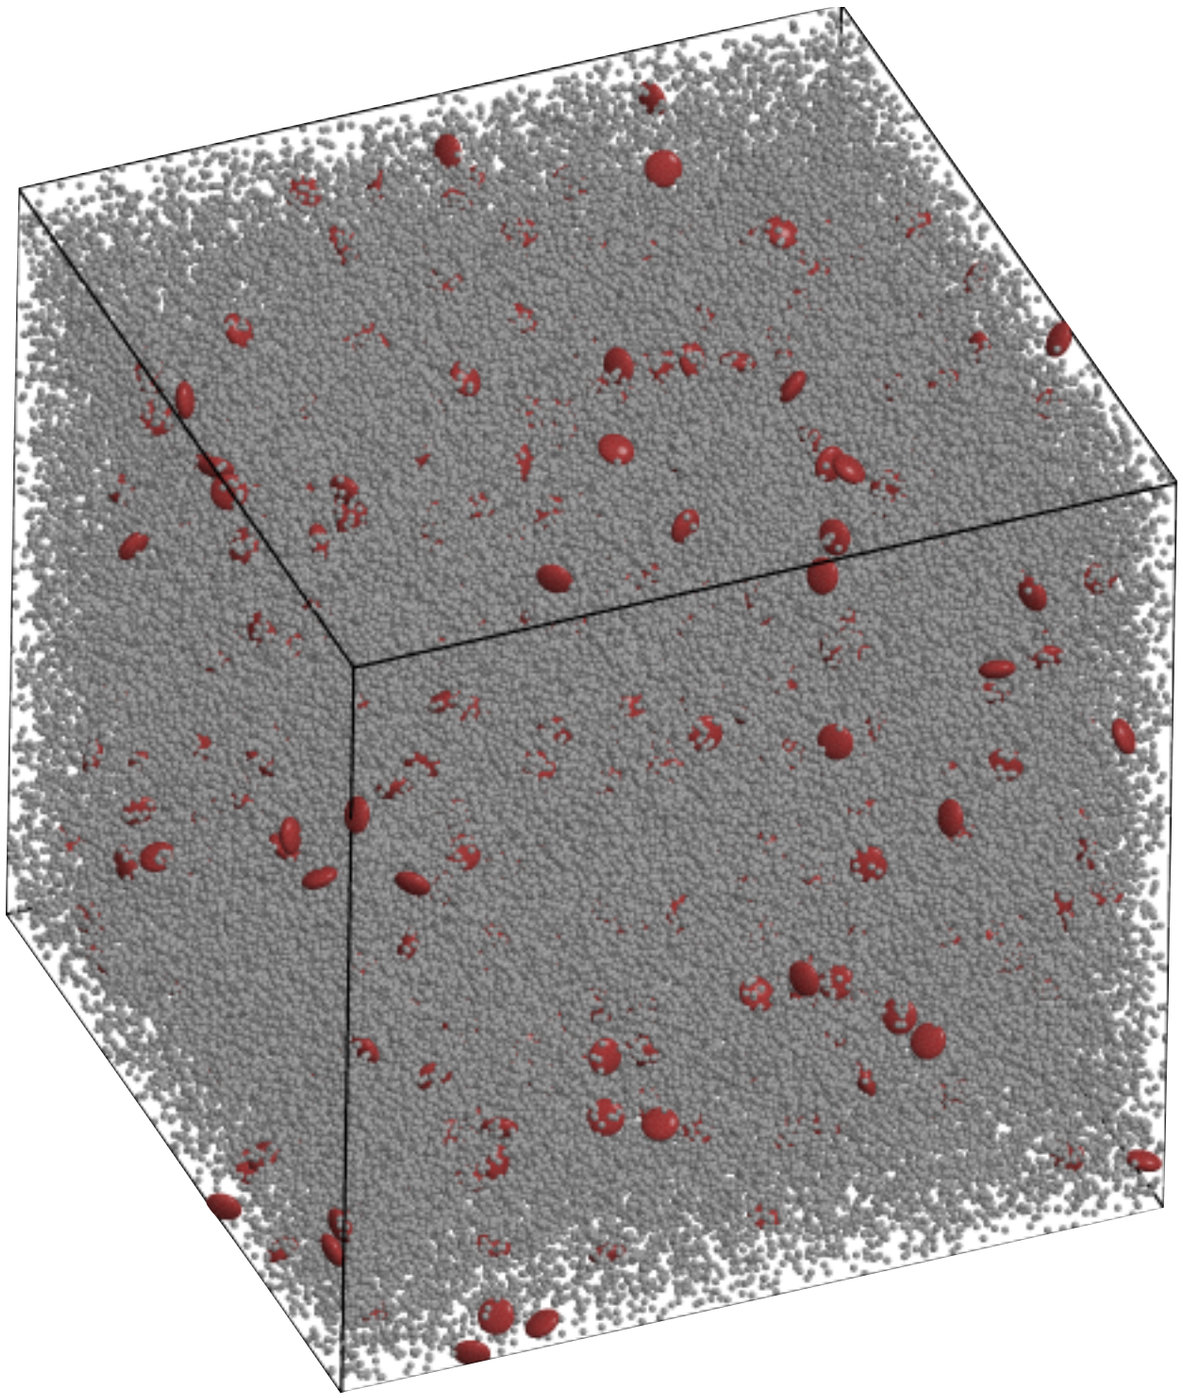

500 colloids in implicit depletants - clustering

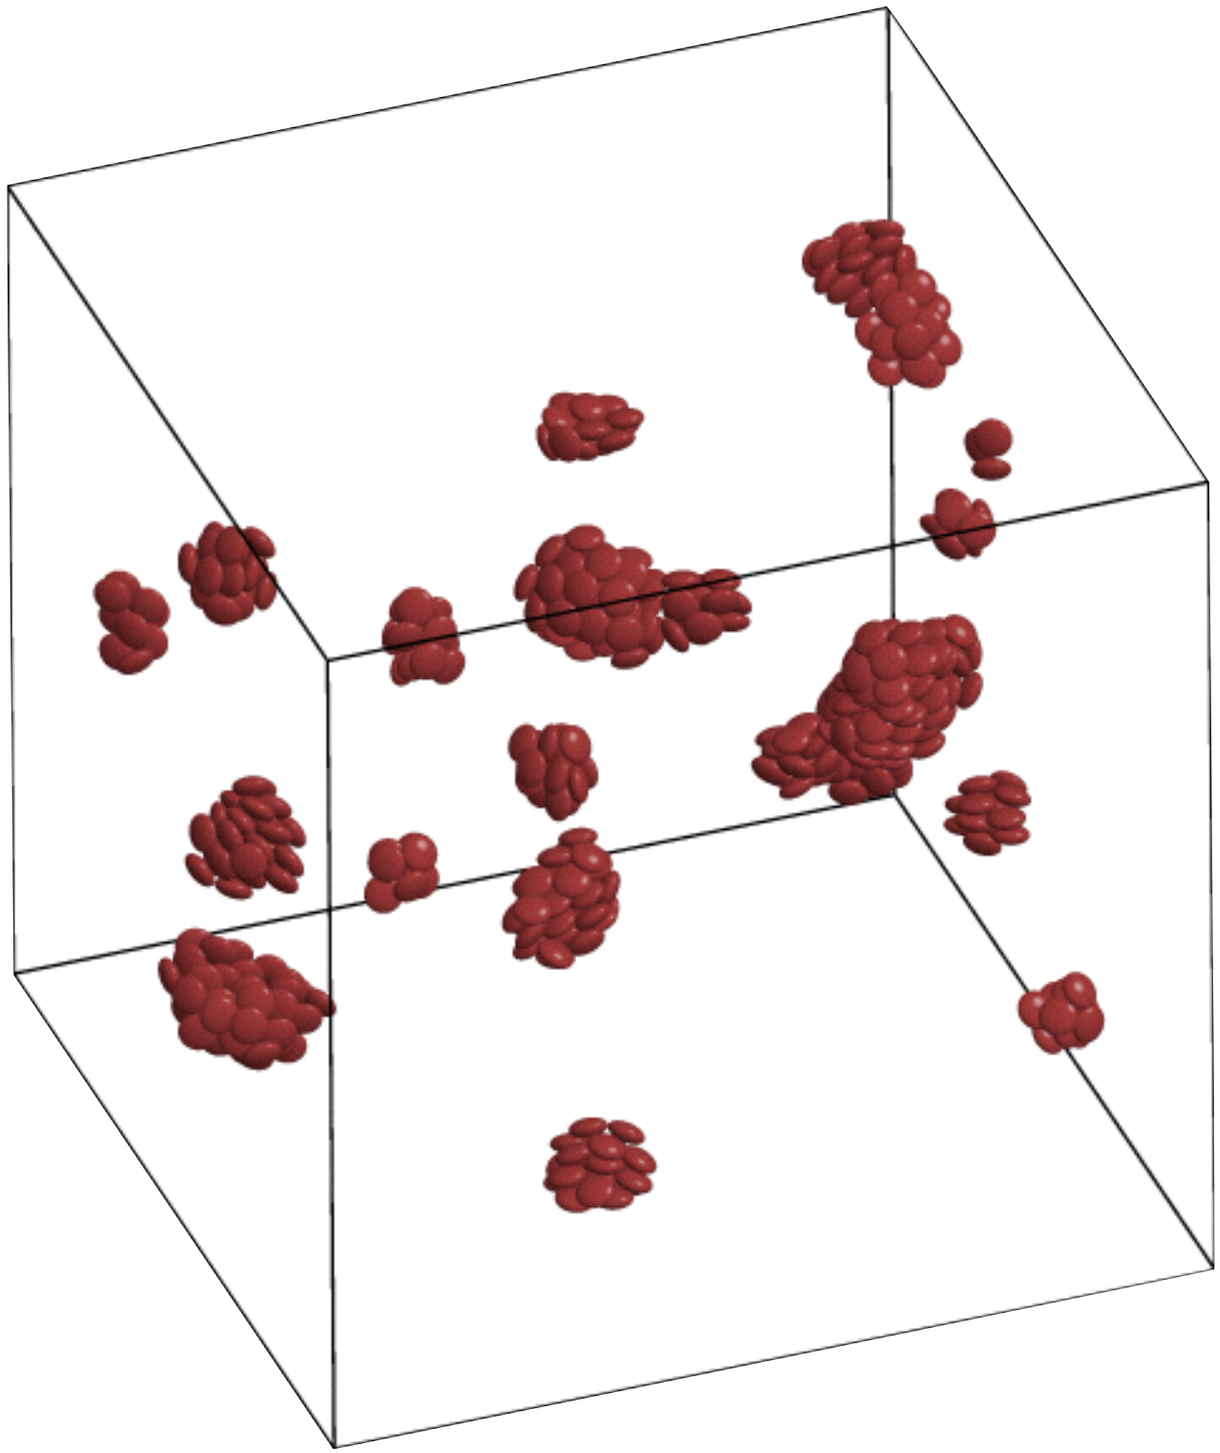

Colloidal discoids

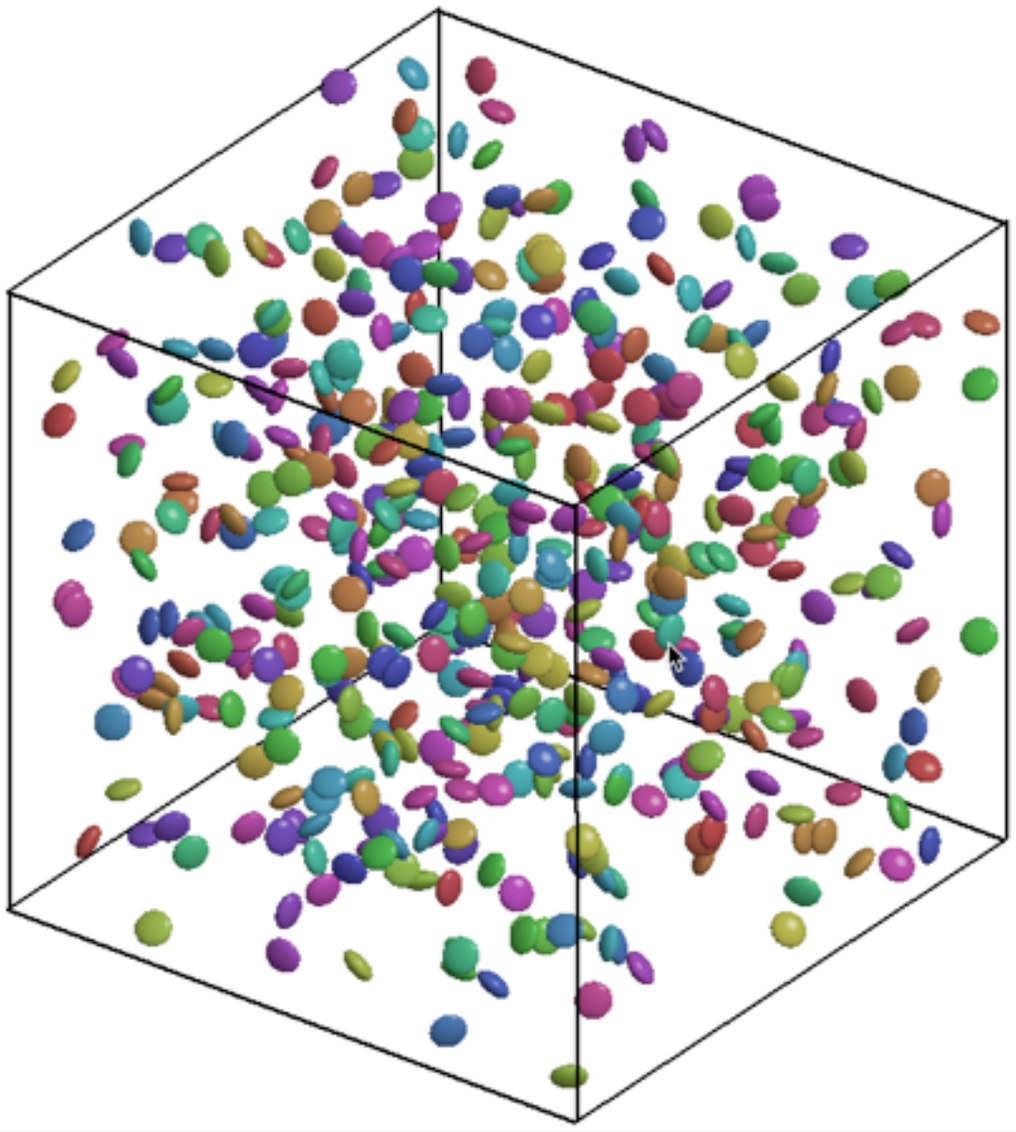

We employ Hard Particle Monte Carlo with implicit simulation of depletion interactions running on two K80 nodes using **eight GPU**s or on **64 Haswell cores**, using GPU-level and MPI parallelism. We perform long running simulations to equilibrate the system over 48h, for 35 different state points.

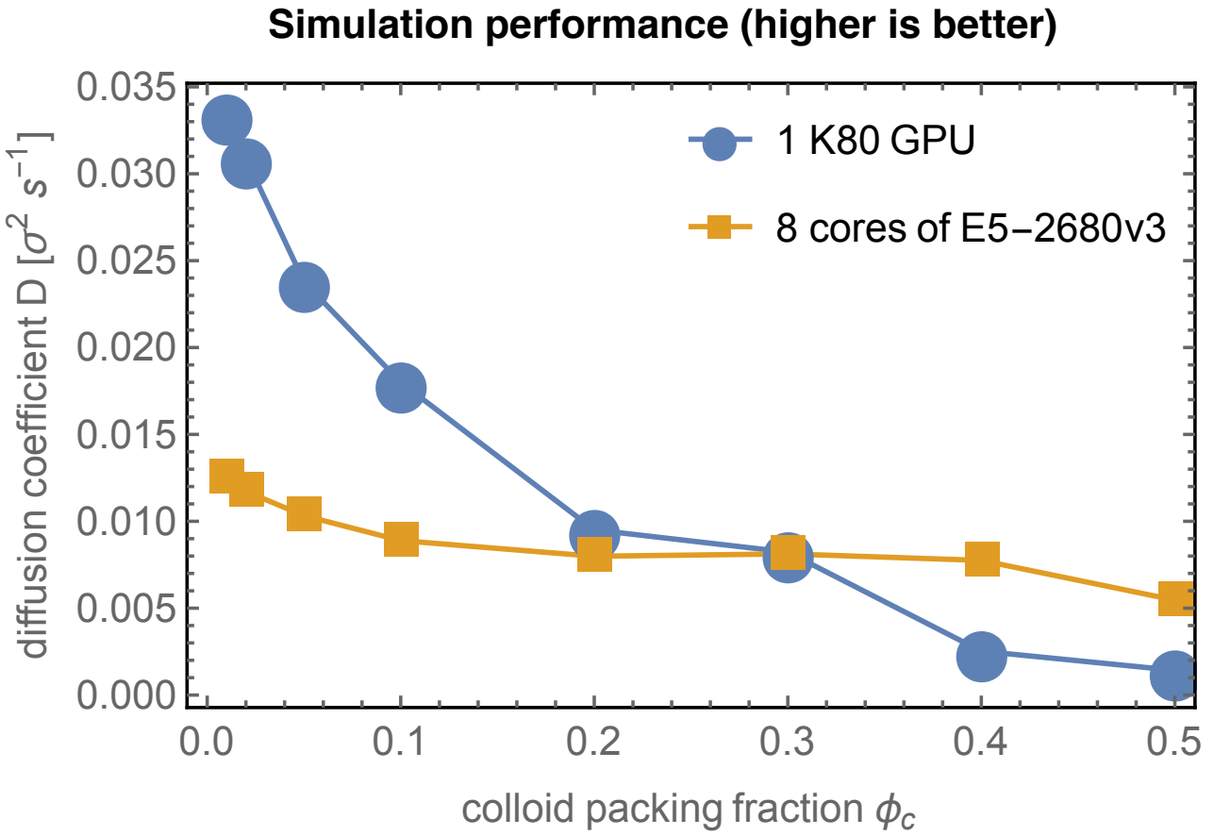

64 Haswell cores
MVAPICH2 2.1

Hemispheres with implicit depletants

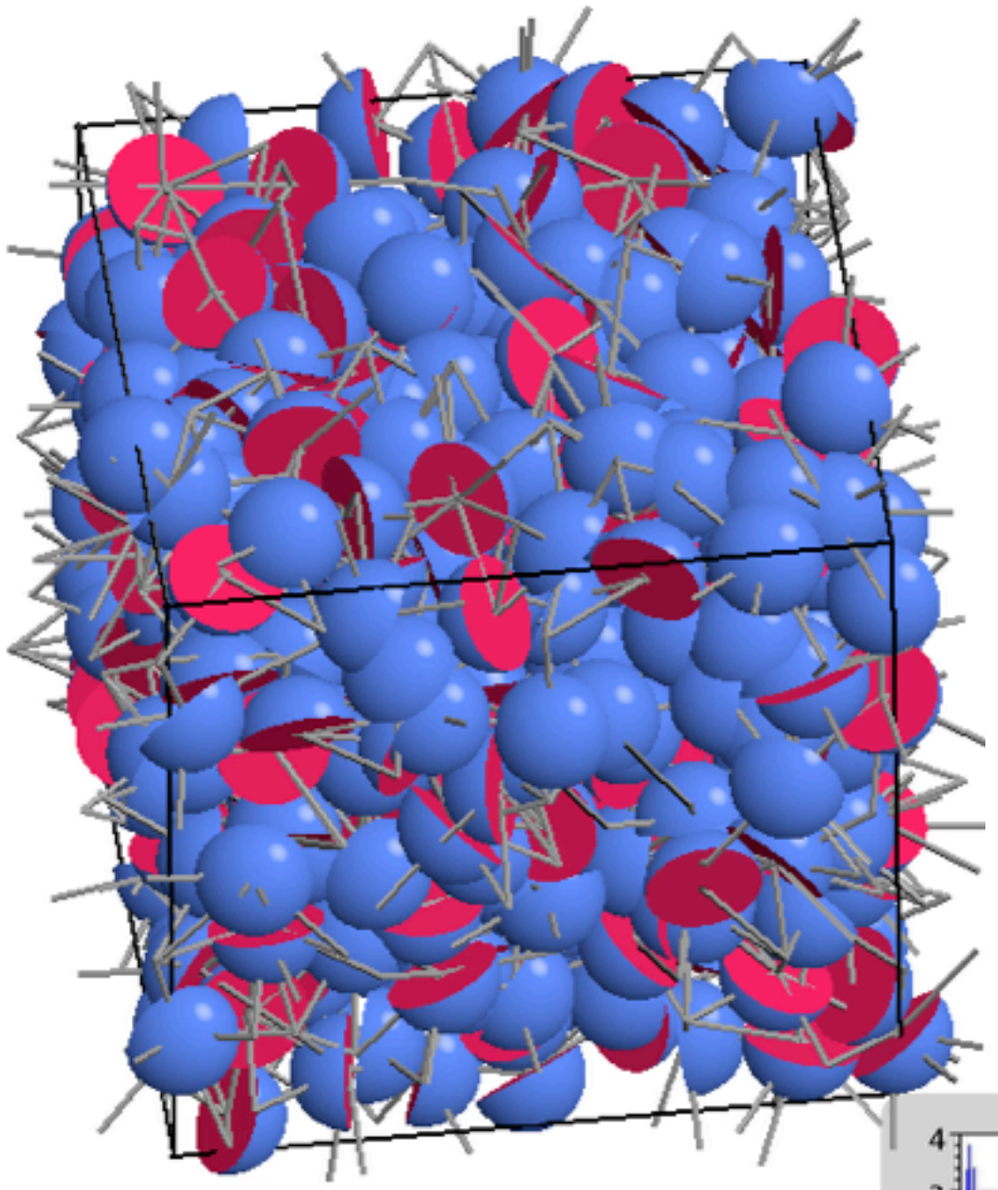

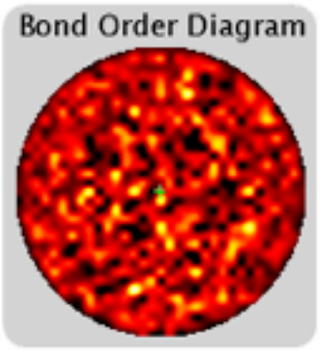

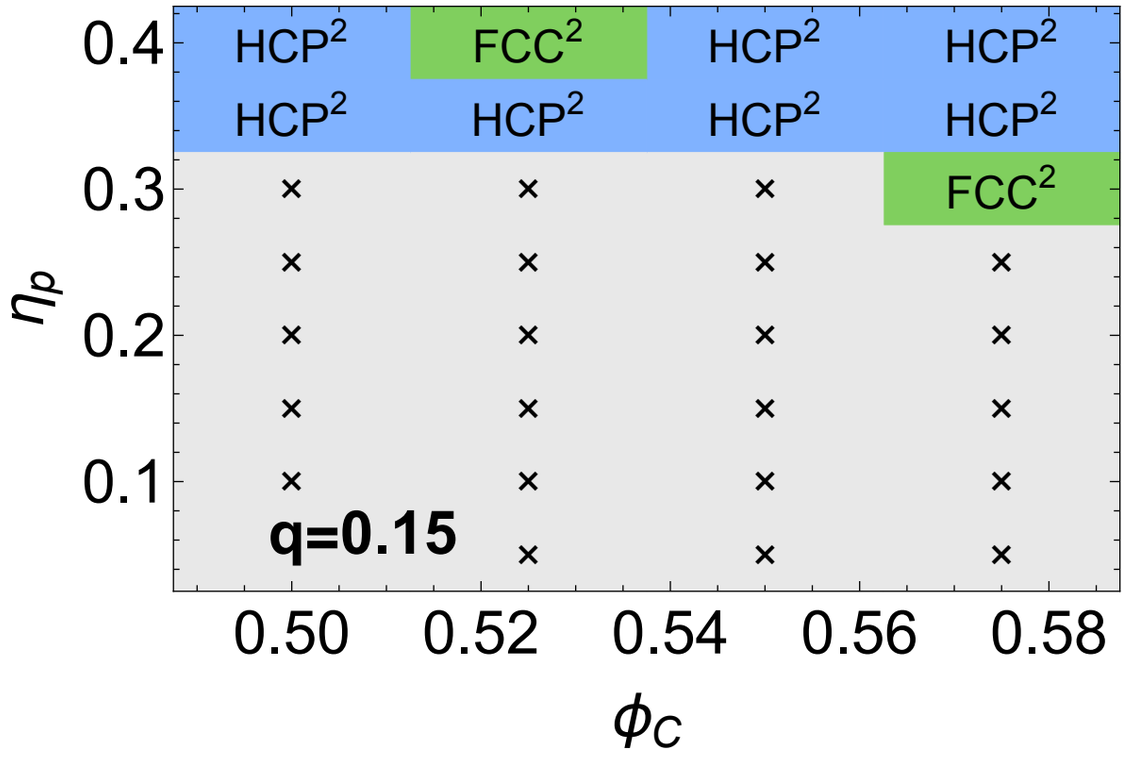

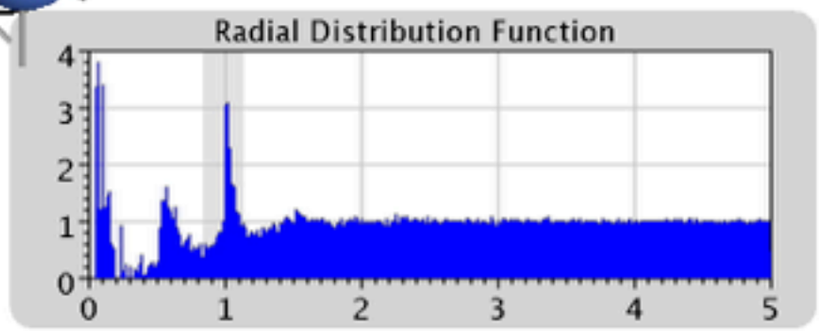

Clustering kinetics

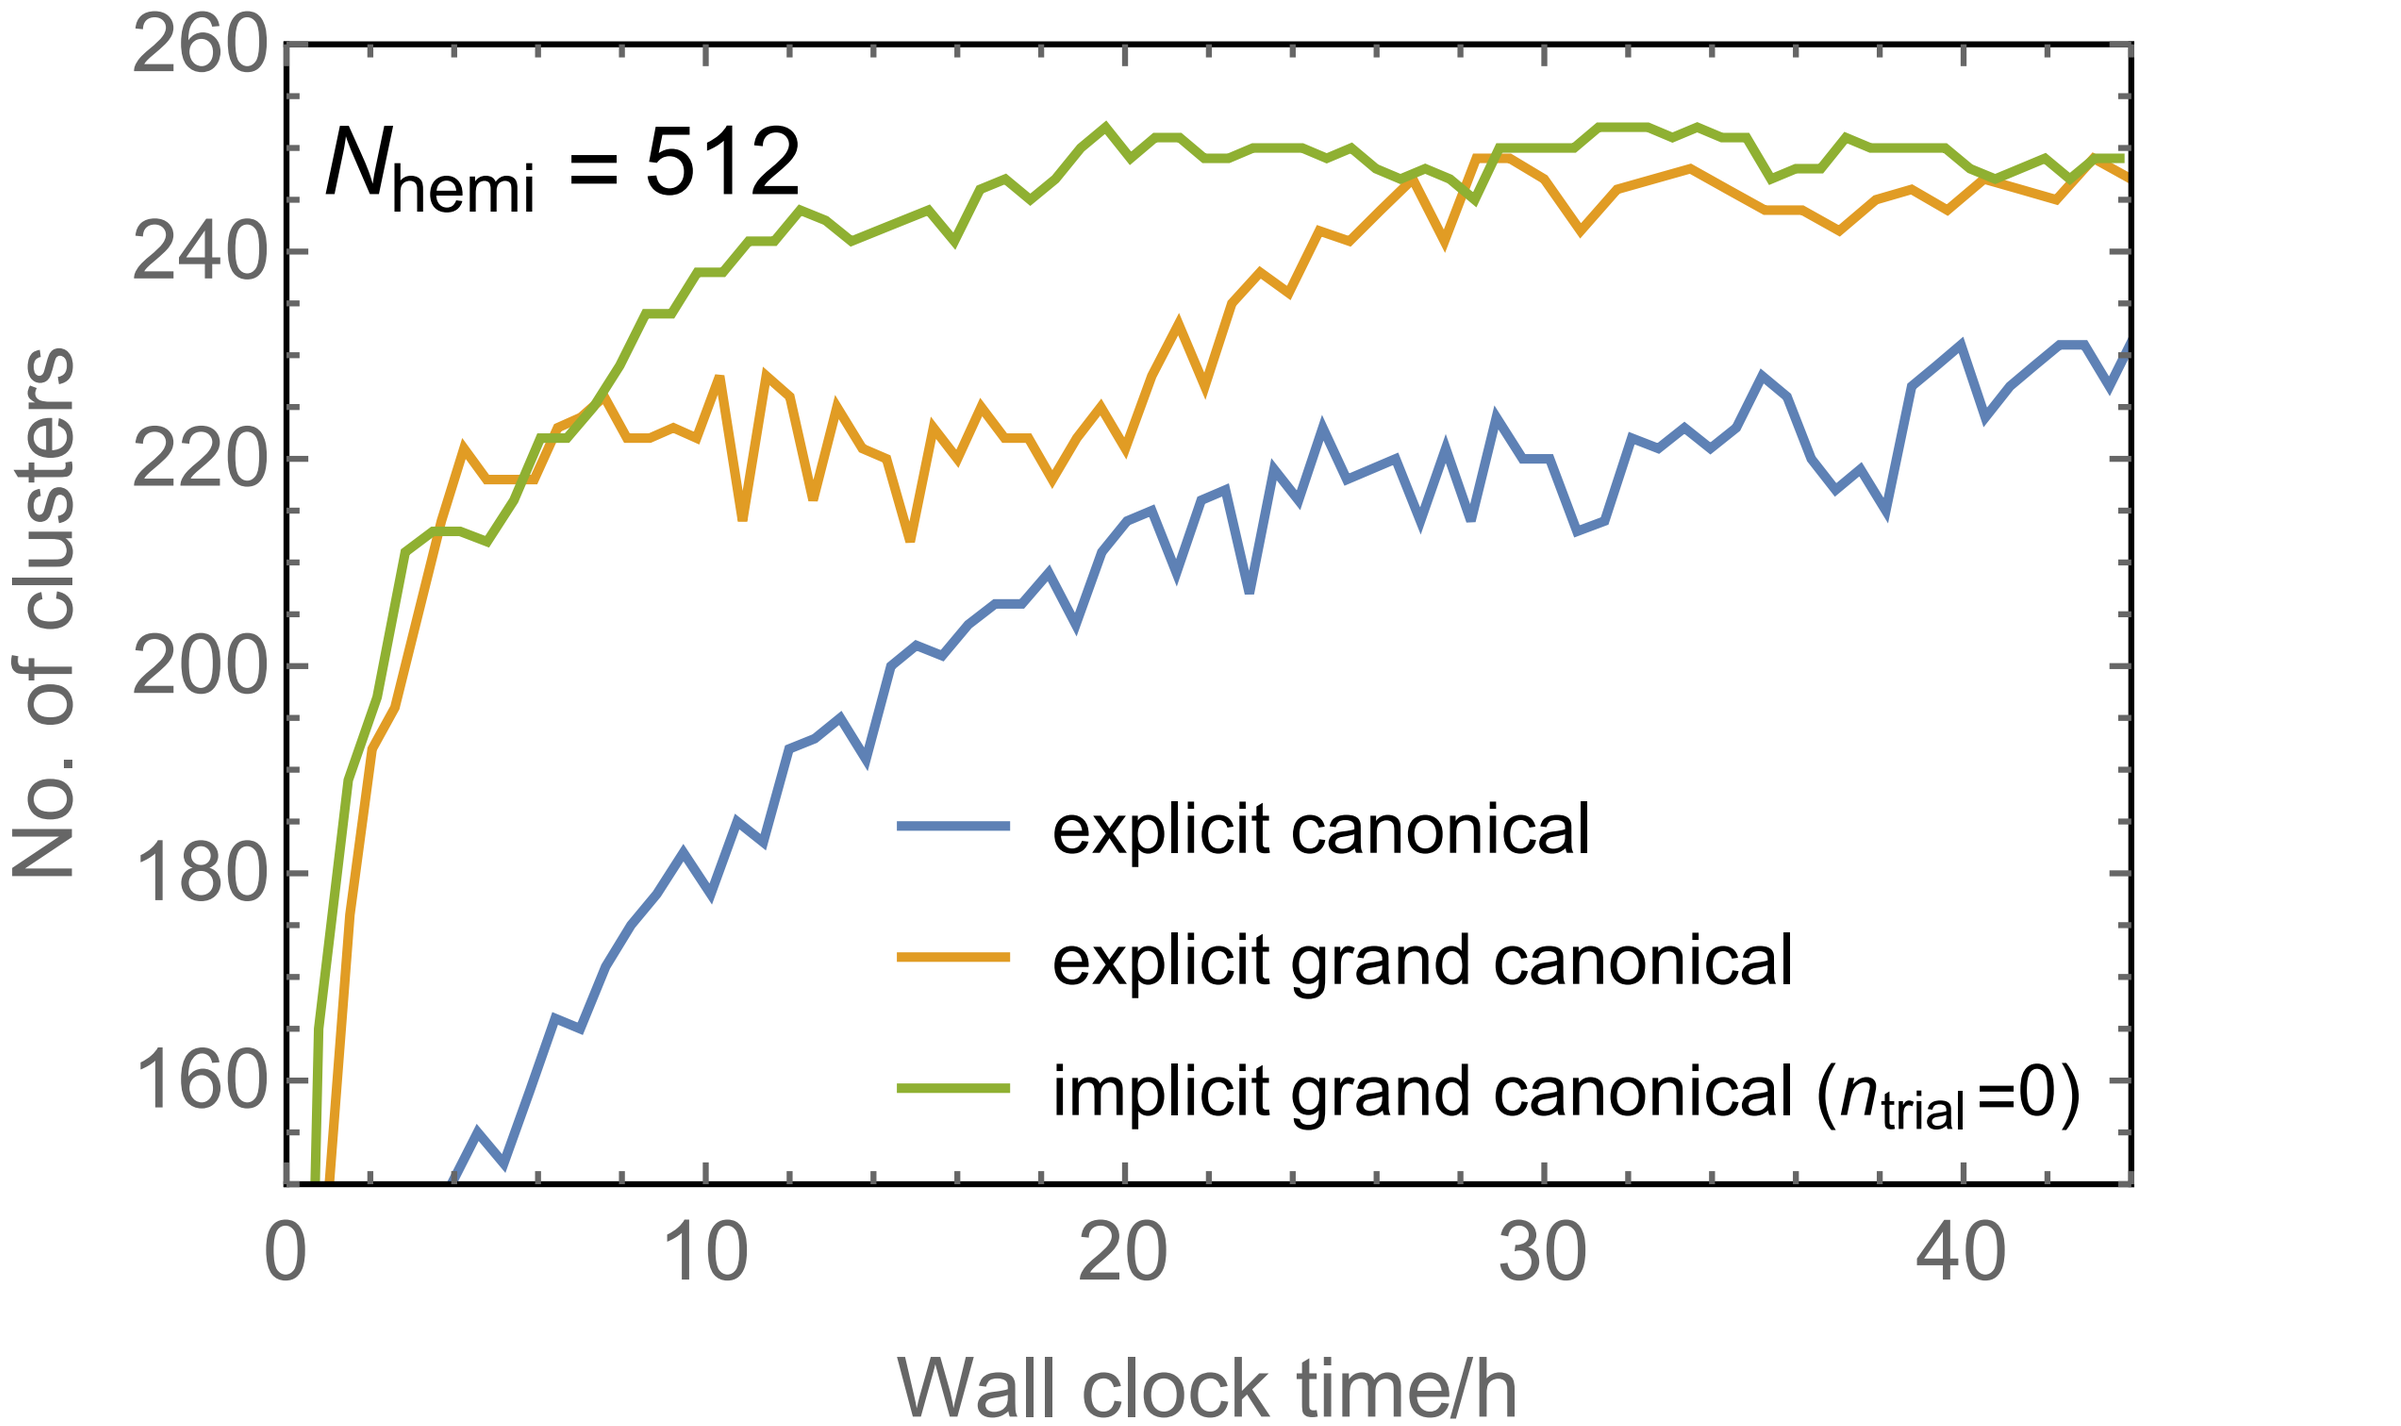

8 Haswell cores

Strong scaling on OLCF Cray XK7

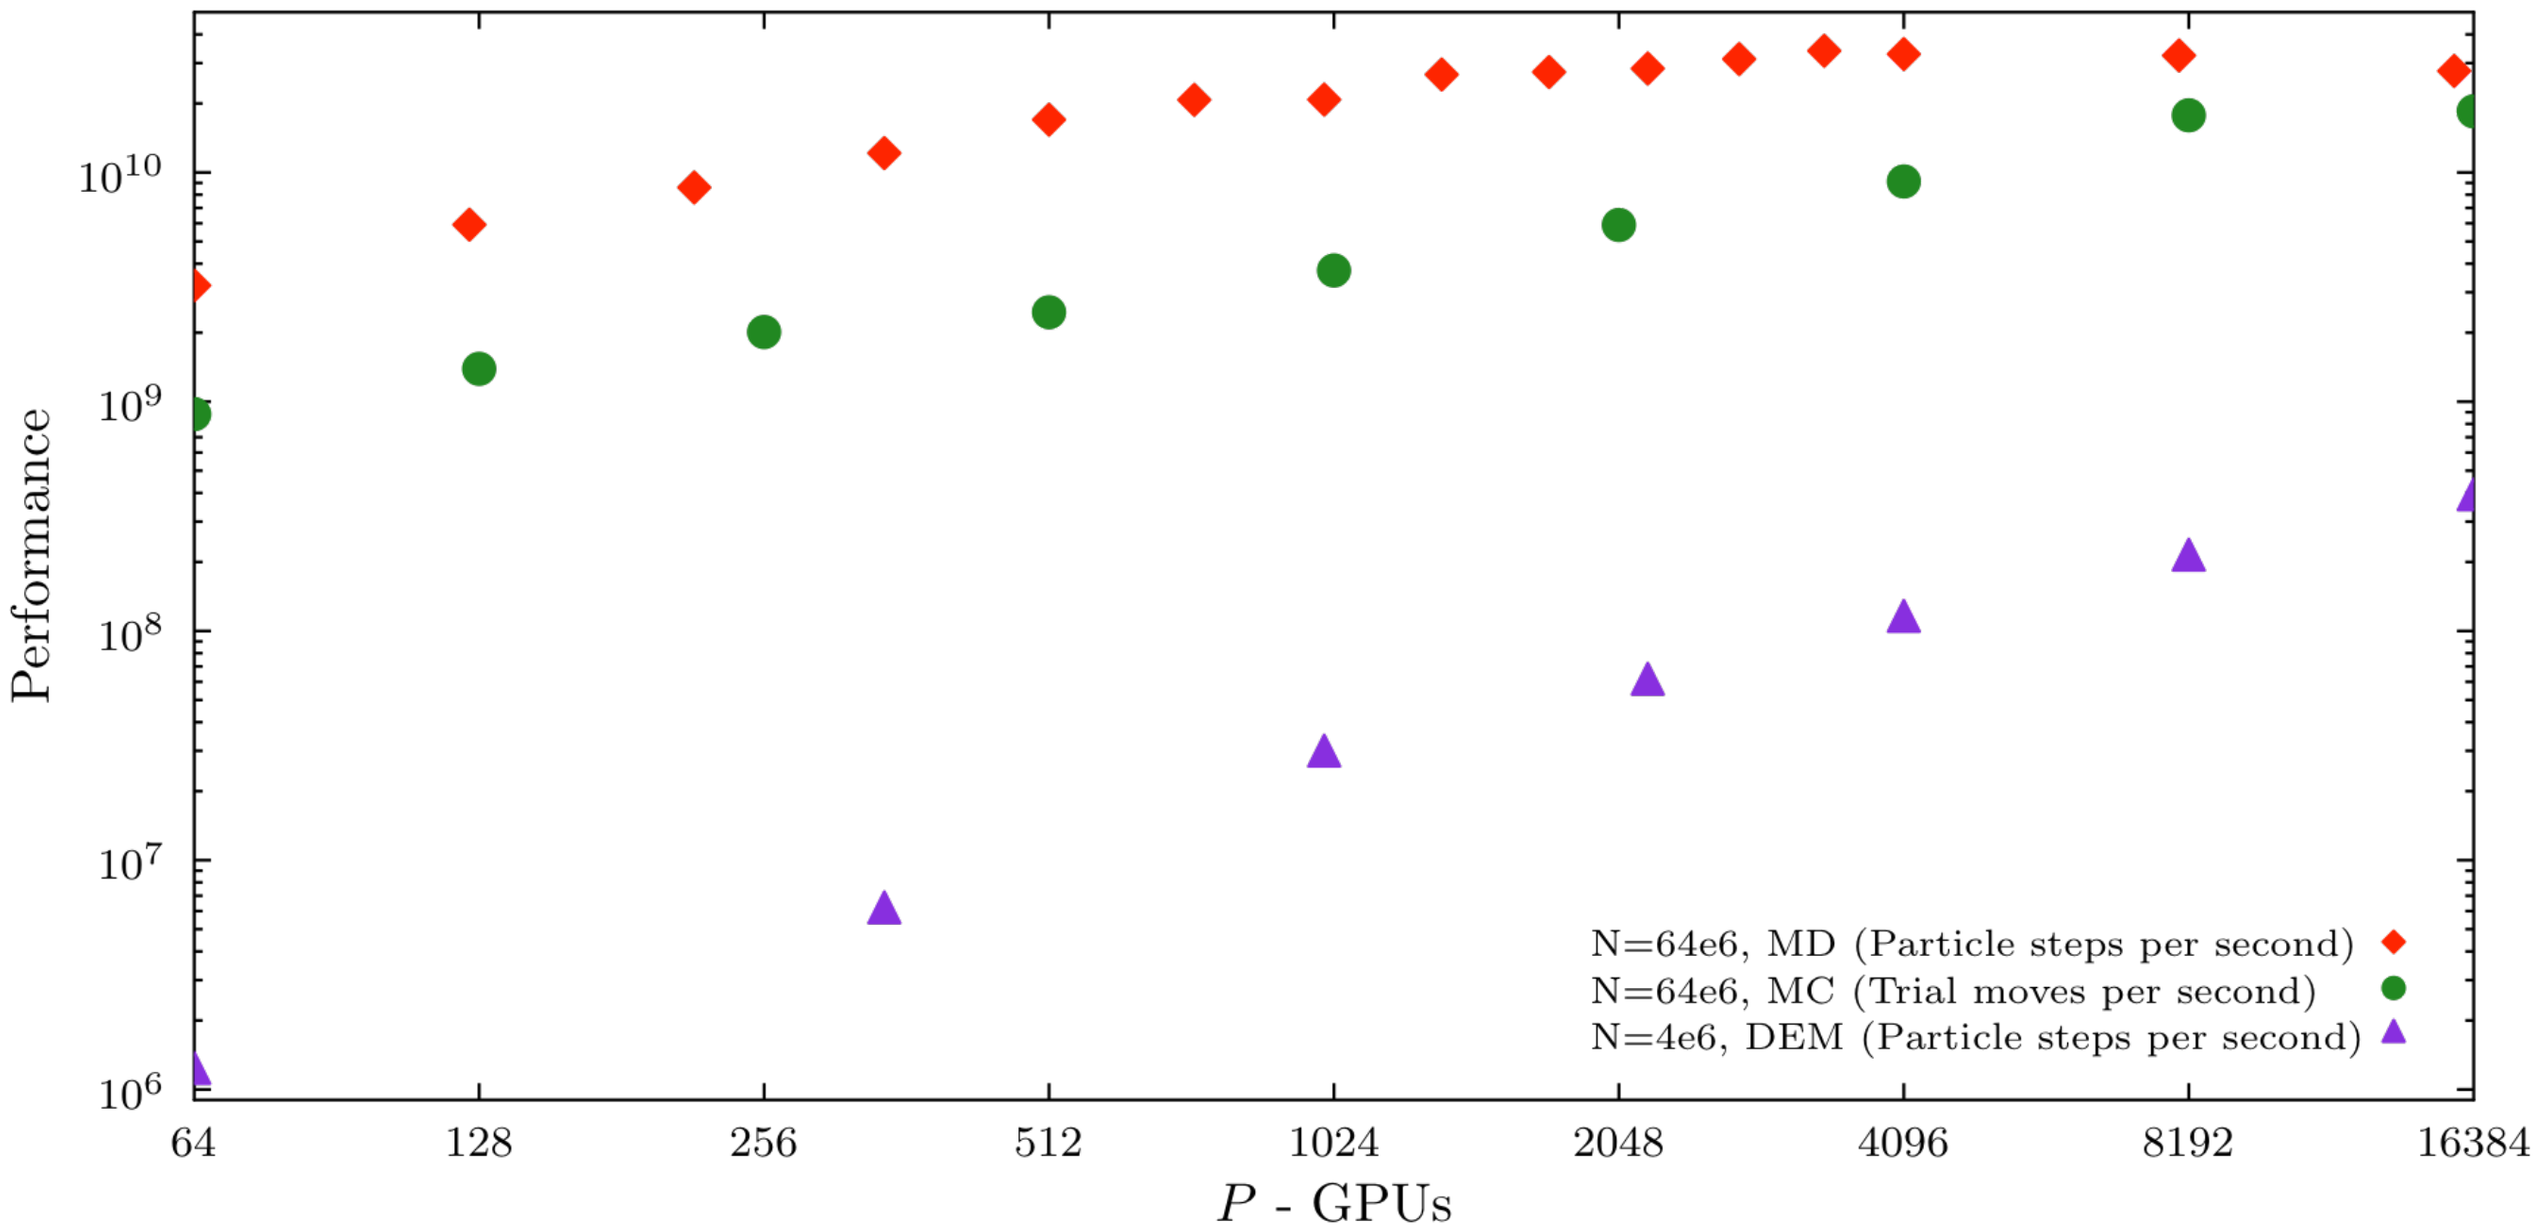

Summary

- Moore's law is slowing - we can't afford to wait a couple of years for a performance doubling anyways

- By extracting **parallelism** from GPUs, threads, vector instructions and MPI we build high performing simulation codes

- Parallel code = previously unsolvable problems become **solvable**

Acknowledgements

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

Strong scaling: J. Glaser, T. D. Nguyen, J. A. Anderson, P. Lui, F. Spiga, J. A. Millan, D. C. Morse, and S. C. Glotzer, "Strong scaling of general-purpose molecular dynamics simulations on GPUs," Comput. Phys. Commun., vol. 192, no. July, pp. 97–107, 2015.

Monte Carlo code not yet publicly available.

- It will eventually be released open-source as part of HOOMD-blue
- Paper on hard disks: Anderson, J. A. et al., JCP **254**, 27-38 (2013)
- Paper on 3D, anisotropic shapes, multi-GPU: *coming soon*

This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575; XSEDE award DMR 140129.

This work was partially supported by a Simons Investigator award from the Simons Foundation to Sharon Glotzer. This work is also partially supported by the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF-15-1-0185.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Lab- oratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC05-00OR22725. This research is part of the Blue Waters sustained petascale computing project, which is supported by the National Science Foundation (award number ACI 1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana–Champaign and its National Center for Supercomputing Applications. We thank the University of Cambridge for providing access to their Wilkes cluster.

EMail: jsglaser@umich.edu

Questions?