LAMMPS and classical molecular dynamics for materials modeling

SAND2015-5705PE

Steve Plimpton
Sandia National Laboratories

OLCF User Meeting
June 2015 - Oak Ridge National Labs

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.
Presentation: SAND2015-4073C

LAMMPS from 10,000 meters

Large-scale Atomic/Molecular Massively Parallel Simulator
http://lammps.sandia.gov

- Materials modeling: soft matter, solids, mesoscale

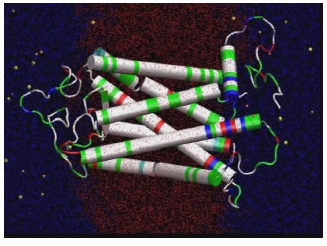
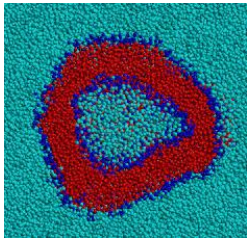
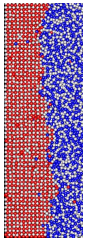
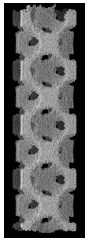
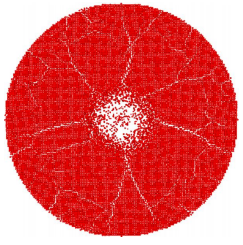

- Particle simulator at varying length and time scales
  electrons $\Rightarrow$ atomistic $\Rightarrow$ coarse-grained $\Rightarrow$ continuum
- Spatial-decomposition of simulation domain for parallelism
- OpenMP, GPU, Phi enhanced
- Can be coupled to other scales: QM, kMC, FE, CFD, ...

Research directions for MD

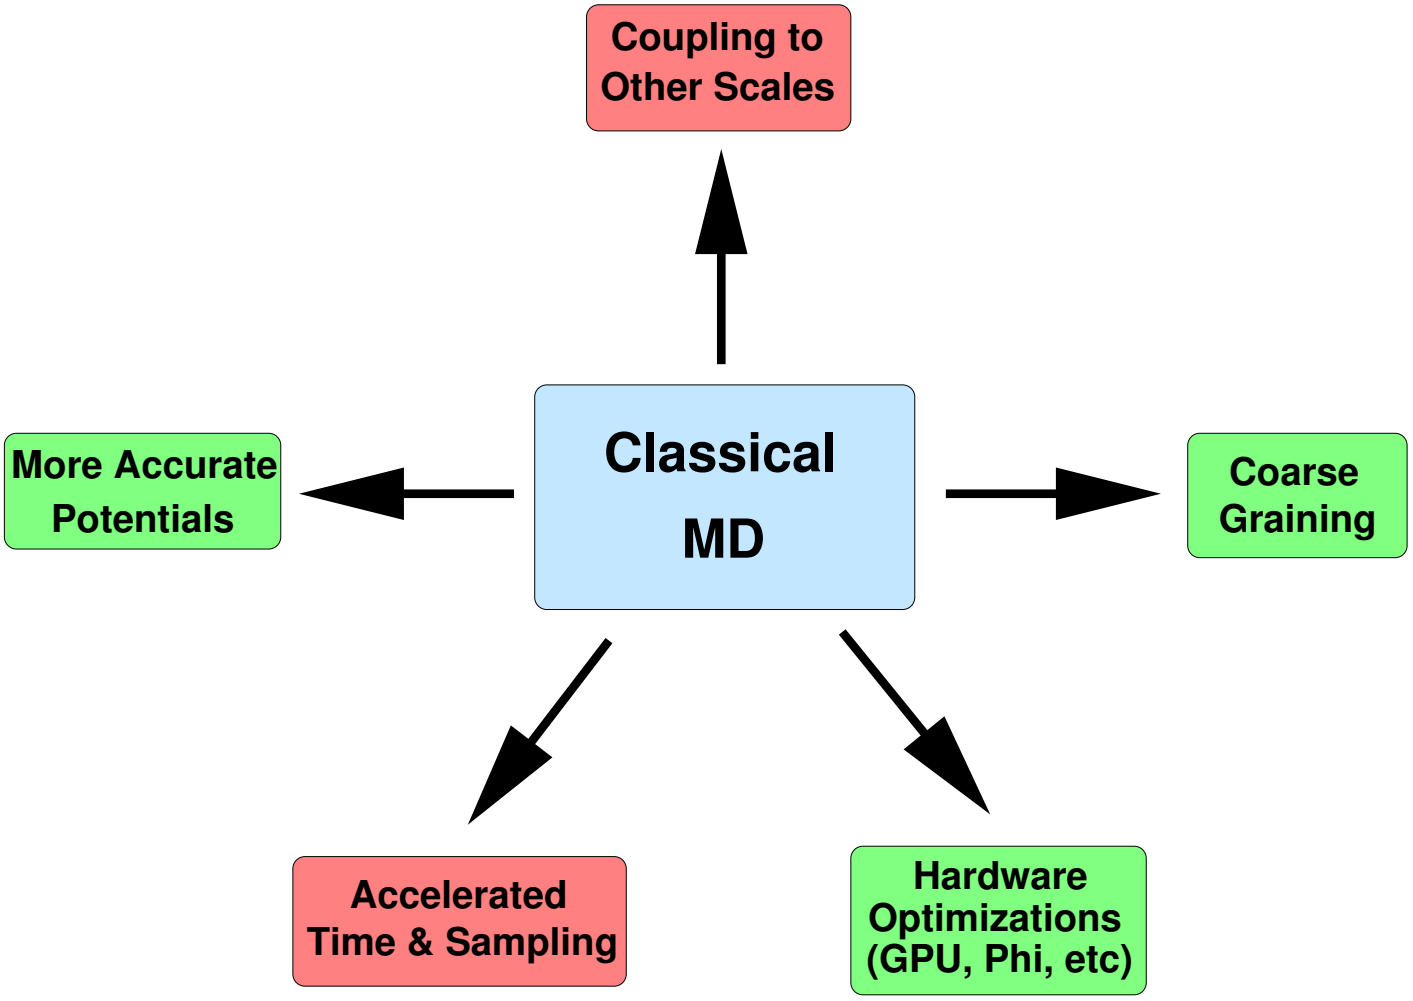

See you at the movies ...

CGI modeling advances by Pixar

Bug's Life (1998)
vegetation

Monsters, Inc (2001)
hair

Finding Nemo (2003)
water

Cars (2006)
painted surfaces

Ratatouille (2007)
food

Wall-E (2008)
rust & decay

Moore's Law for interatomic potentials (force fields)

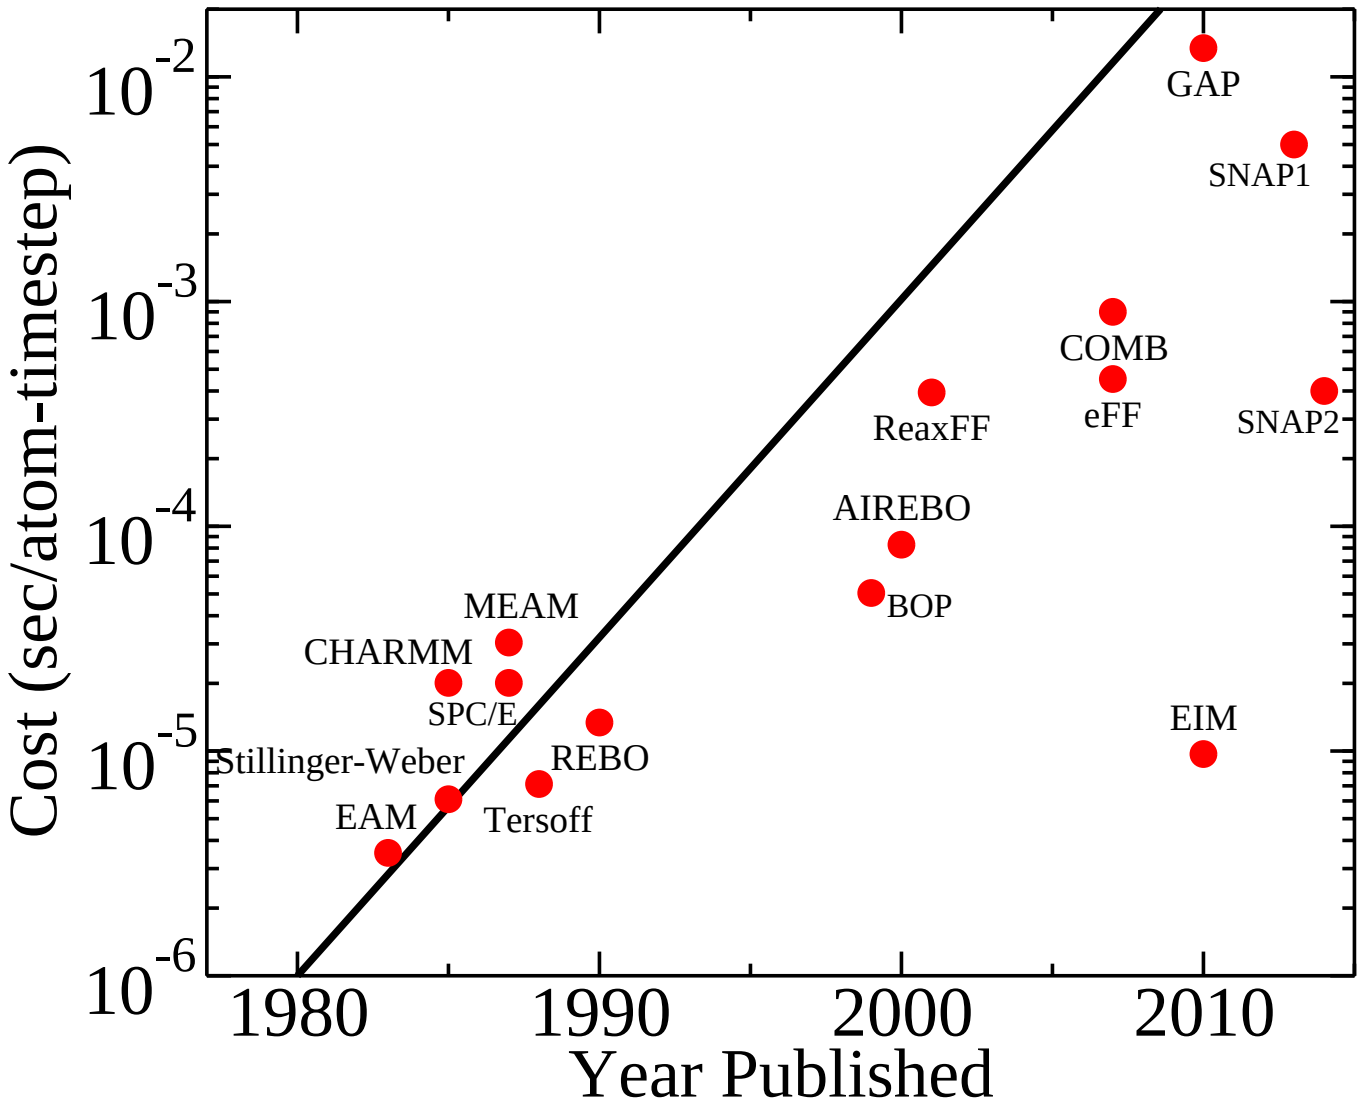

ReaxFF for shock-induced initiation of detonation

- Work by Ray Shan & Aidan Thompson (Sandia)
- **PETN** = explosive material made of organic molecules
- Simulate "slow" shock wave passing thru PETN crystal

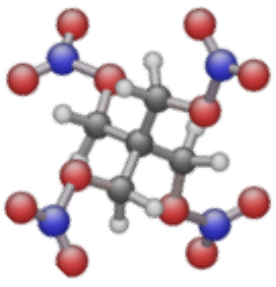

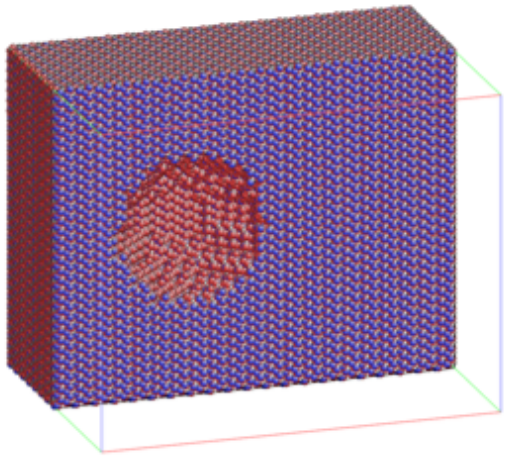

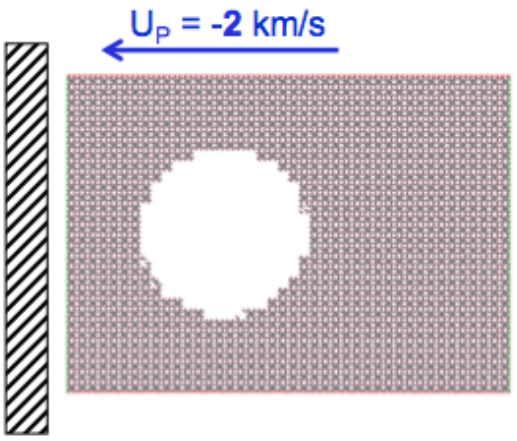

ReaxFF for shock-induced initiation of detonation

- Work by Ray Shan & Aidan Thompson (Sandia)
- **PETN** = explosive material made of organic molecules
- Simulate "slow" shock wave passing thru PETN crystal

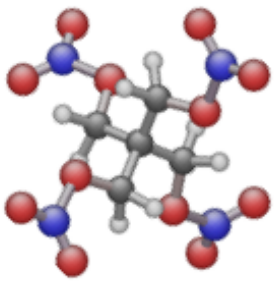

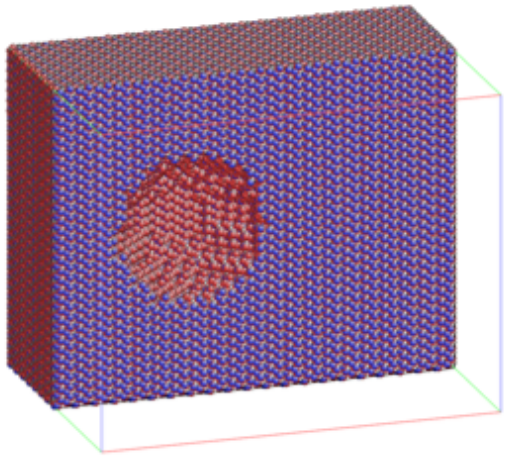

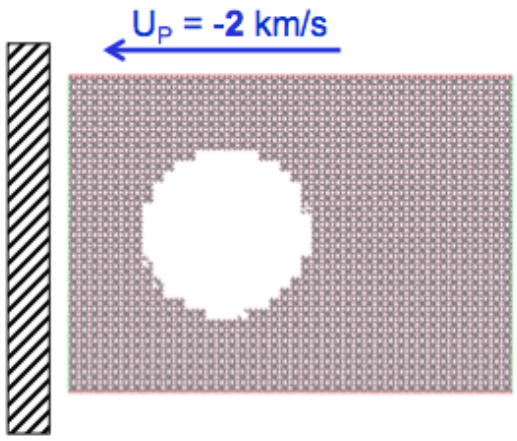

- Use a **reactive force field** (ReaxFF)
  - detonation triggered by initiation of exothermic reactions
- Quantify detonation **sensitivity** to
  orientation, defects, impurities ... a safety issue

Large-scale 20 nm void simulations

8.5M atoms (300x200x1.3 nm), 500 psec
5M steps, 1500 hours on 64K cores (4 t/c) of BQ/Q

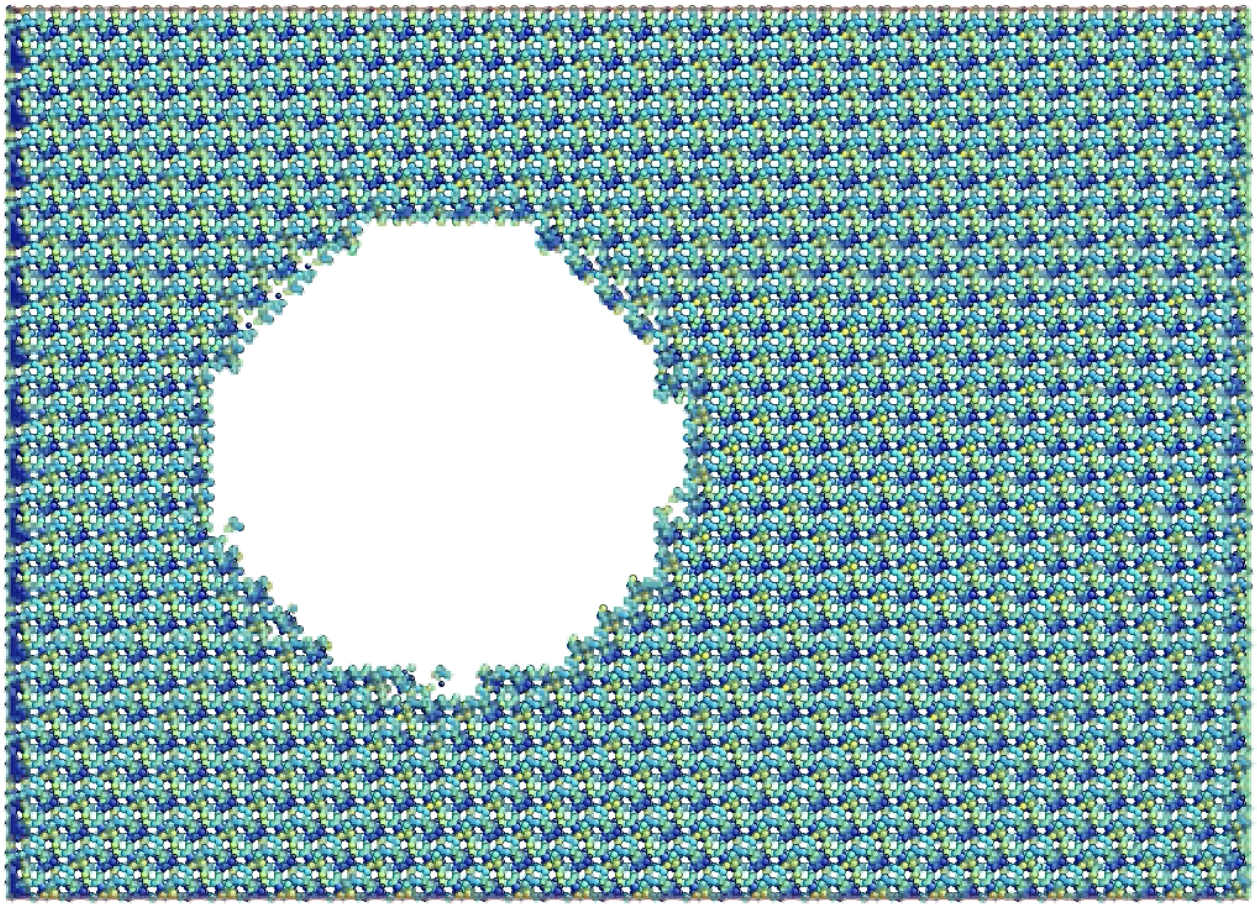

Shan, et al, Proc 15th International Detonation Symp (2015)

Sensitivity of PETN to shock stress

- **Heat and stress** trigger chemical reactions
- Hot-spot lowers initiation threshold stress by 30%
- Increased **sensitivity** with increasing void size

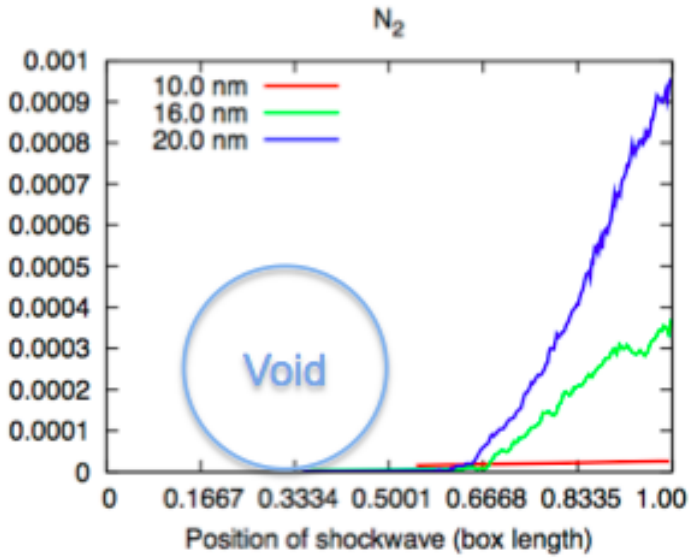

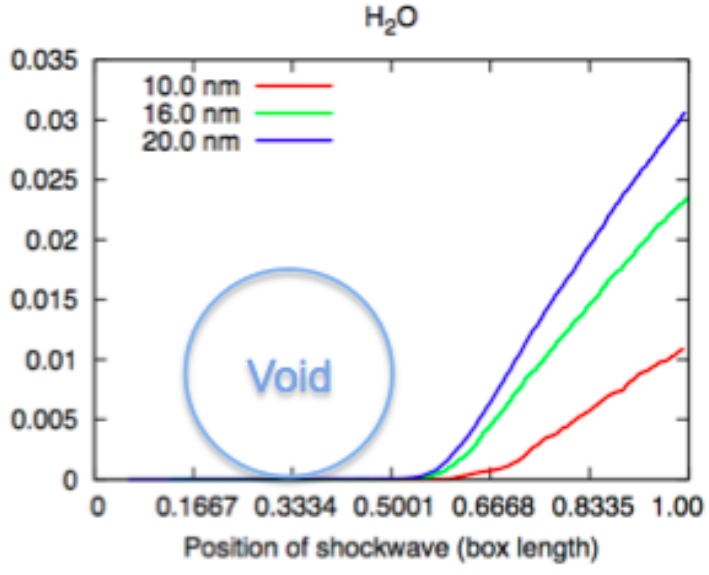

Sensitivity of PETN to shock stress

- **Heat and stress** trigger chemical reactions
- Hot-spot lowers initiation threshold stress by 30%
- Increased **sensitivity** with increasing void size

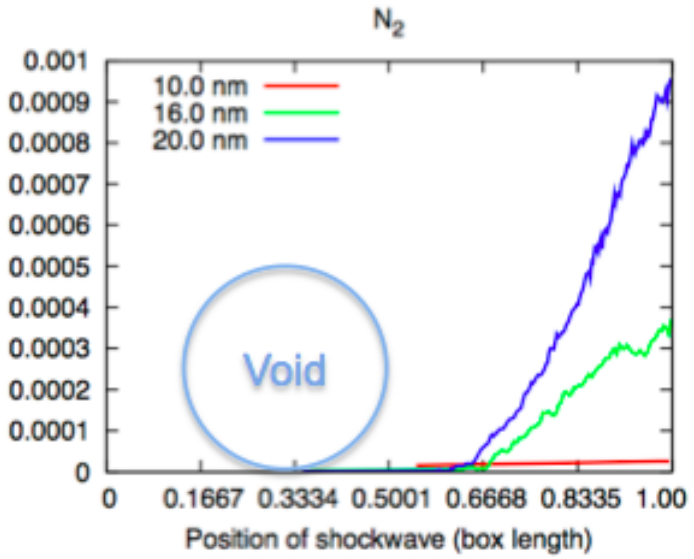

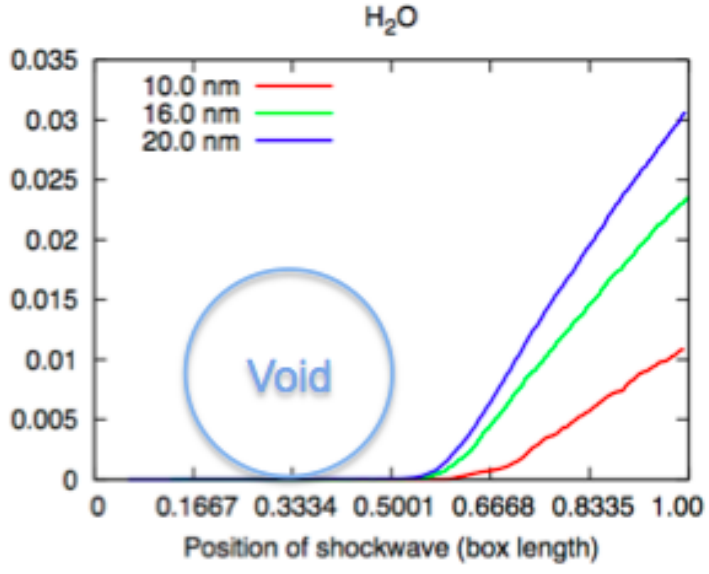

- Agrees with **experiment** except:
  - real PETN can have micron-scale voids
  - experiments performed at lower shock speeds
  - bigger/longer/CG simulations needed to bridge that gap

Quantum-accuracy with empirical potentials?

- **GAP** = Gaussian approximation potentials
  - Gábor Csányi, Albert Bartók-Partay (U Cambridge)
- **SNAP** = spectral neighbor analysis potentials
  - Aidan Thompson and collaborators (Sandia)
- Aim for **quantum-level accuracy** in some cases:
  - interpolate to *ab initio* potential energy surface
  - Gaussian process: high-dimensional interpolation technique
  - trained on set of QM configurations, energy, forces
  - expensive, but cost still $O(N)$ in number of atoms

Quantum-accuracy with empirical potentials?

- **GAP** = Gaussian approximation potentials
  - Gábor Csányi, Albert Bartók-Partay (U Cambridge)
- **SNAP** = spectral neighbor analysis potentials
  - Aidan Thompson and collaborators (Sandia)
- Aim for **quantum-level accuracy** in some cases:
  - interpolate to *ab initio* potential energy surface
  - Gaussian process: high-dimensional interpolation technique
  - trained on set of QM configurations, energy, forces
  - expensive, but cost still $O(N)$ in number of atoms

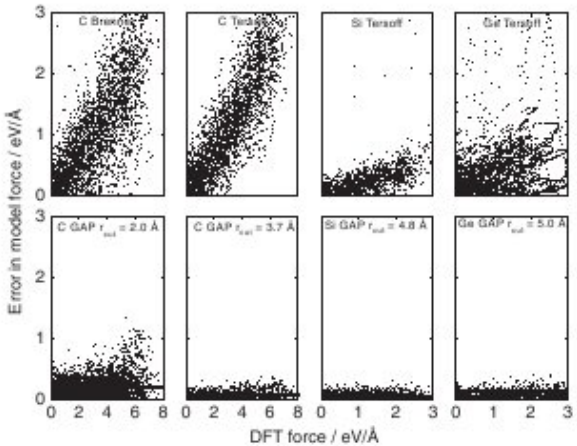

Reduces errors relative to DFT

Bartok, et al, PRL, 104, 136403 (2010).

- Our interest: semiconductors & metals like InP, Ta, Be

Success with SNAP potential for Tantalum

Energy barrier for screw dislocation migration in bcc Ta

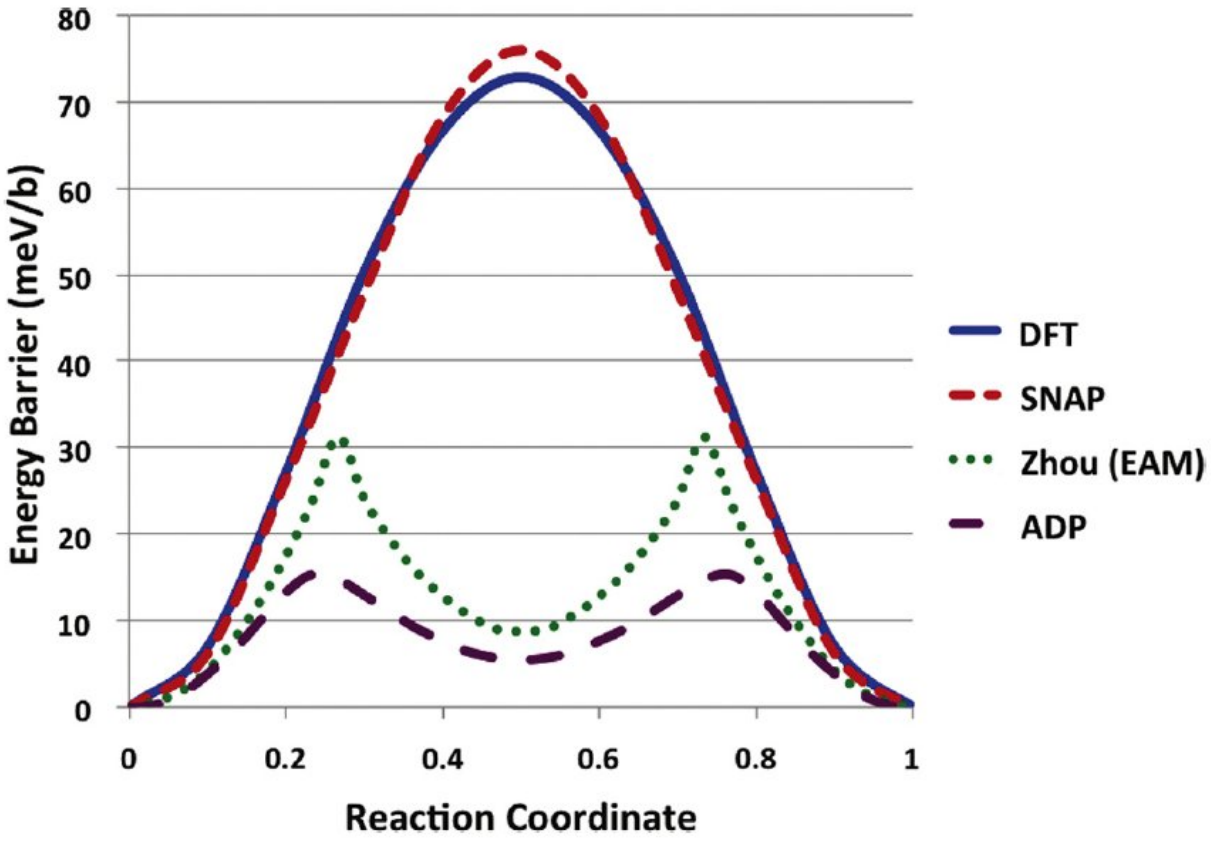

- **DAKOTA** optimization package used to iteratively fit
- *Thompson, et al, J Comp Phys, 285, 316-330 (2015).*

Strong scaling of SNAP on Titan

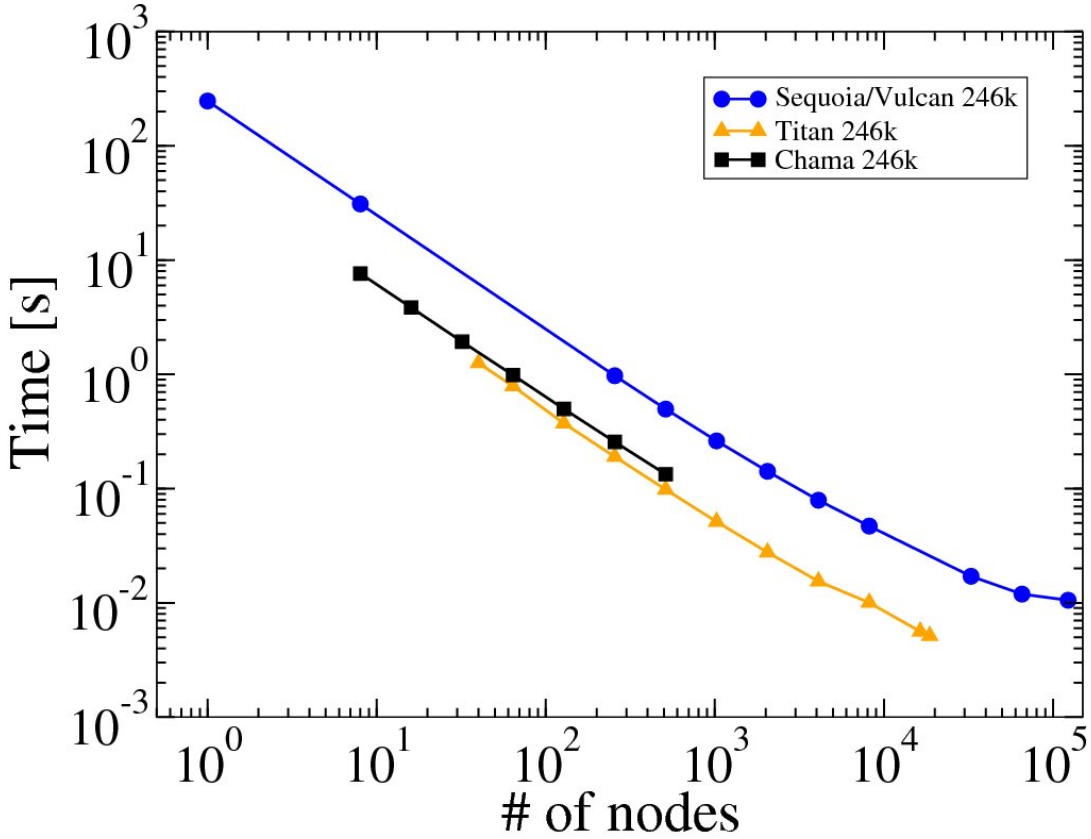

- 246K atoms on 18K nodes of Titan, /red13 atoms/GPU (!)
- SNAP optimized for GPU by Christian Trott (Sandia)
- *Trott, et al, Supercomputing,*
  in Lecture Notes Comp Sci, 8488, Springer, 1934 (2014).

Coarse graining to extend length & time scales

CRADA with companies interested in **solvated nanoparticles**

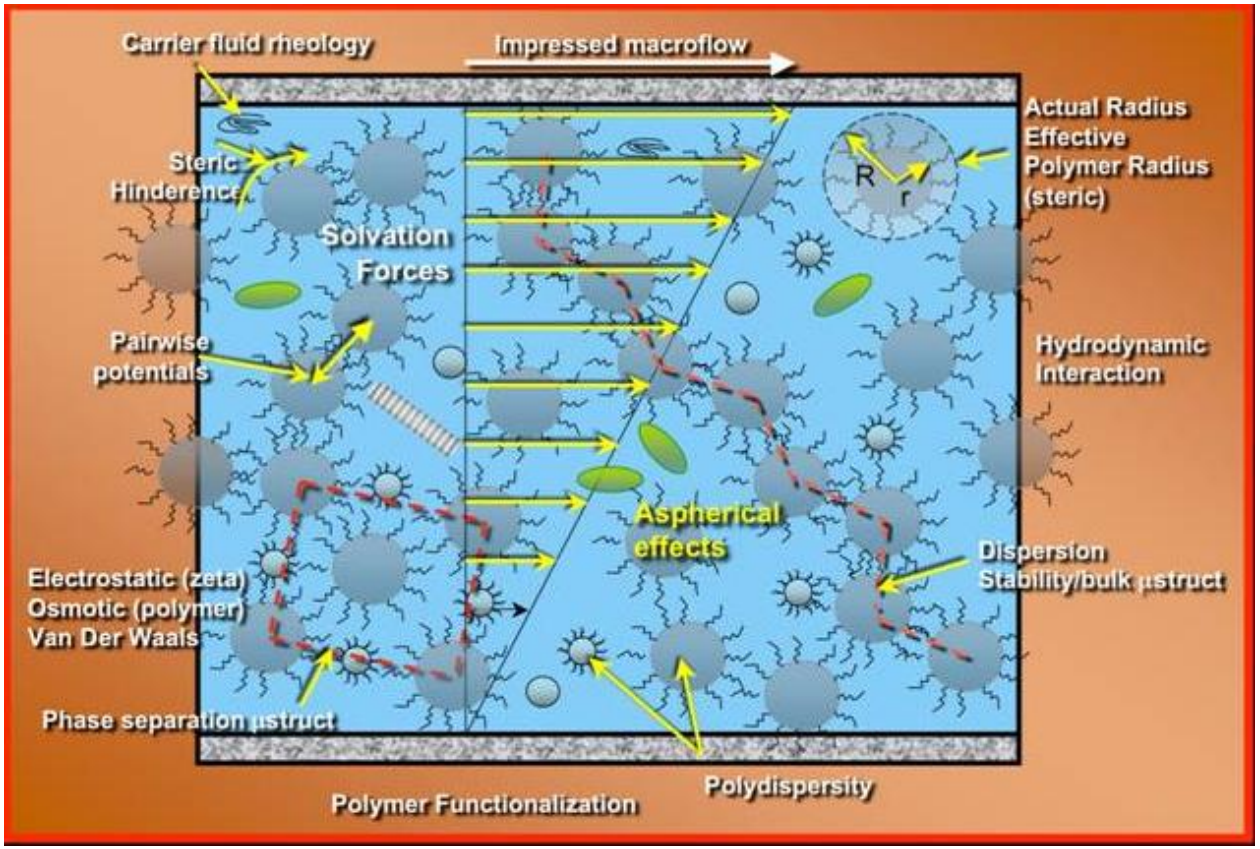

Spherical vs aspherical, bare vs coated, polydisperse, agglomeration, response to shear, ...

Sequence of coarse-grained models in LAMMPS

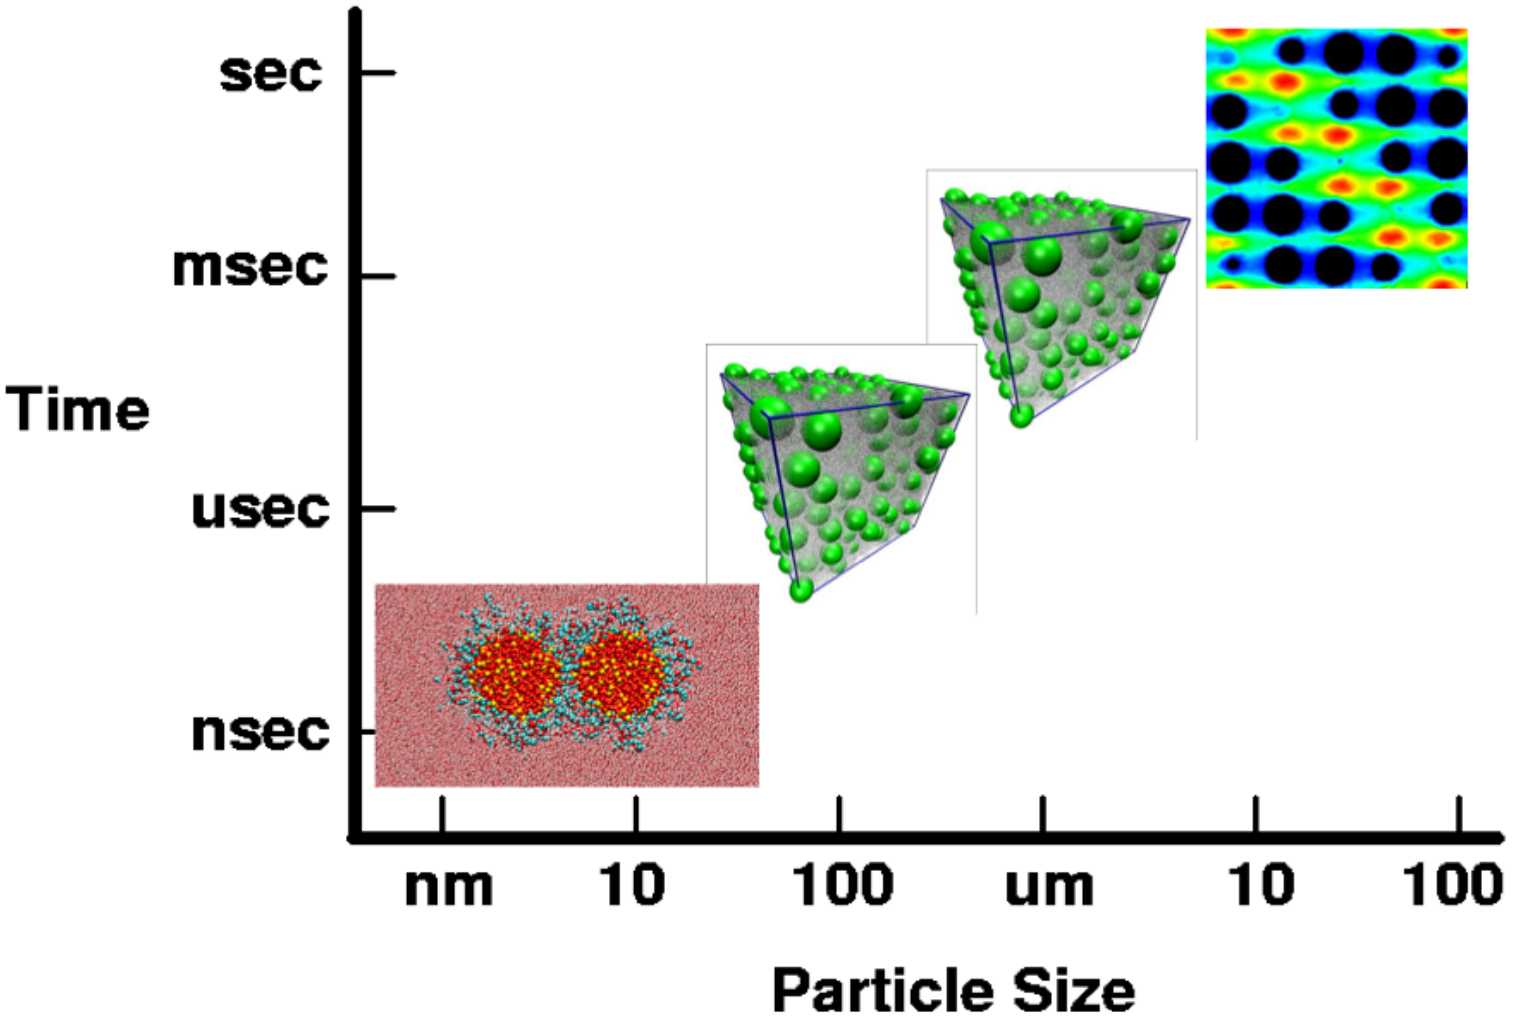

Coarse-graining of nanoparticles and solvent

Integrated LJ potential over NP volume: *Everaers (PRE 2003)*

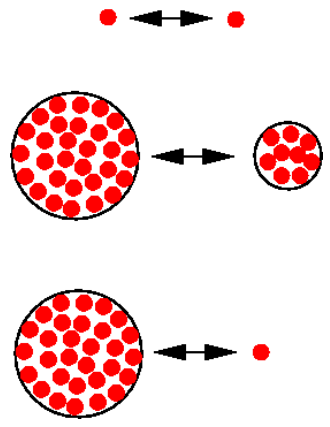

$$U_A = -\frac{A_{cc}}{6}\left[\frac{2a_1a_2}{r^2-(a_1+a_2)^2}+\frac{2a_1a_2}{r^2-(a_1-a_2)^2}+\ln\left(\frac{r^2-(a_1+a_2)^2}{r^2-(a_1-a_2)^2}\right)\right]$$

$$U_R = \frac{A_{cc}}{37800}\frac{\sigma^6}{r}\Bigg[\frac{r^2-7r(a_1+a_2)+6(a_1^2+7a_1a_2+a_2^2)}{(r-a_1-a_2)^7}+\frac{r^2+7r(a_1+a_2)+6(a_1^2+7a_1a_2+a_2^2)}{(r+a_1+a_2)^7}-\frac{r^2+7r(a_1-a_2)+6(a_1^2-7a_1a_2+a_2^2)}{(r+a_1-a_2)^7}-\frac{r^2-7r(a_1-a_2)+6(a_1^2-7a_1a_2+a_2^2)}{(r-a_1+a_2)^7}\Bigg]$$

$$U = U_A + U_R, \qquad r < r_c$$

SRD = **stochastic rotation dynamics** for solvent, then FLD:
Padding (PRL 04), Kumar and Higdon, (PRE 2010)

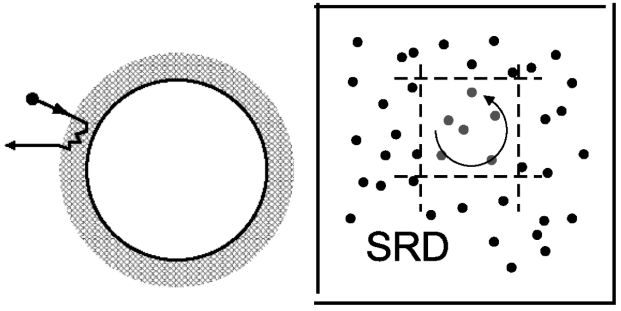

Diffusion across time scales and volume fractions

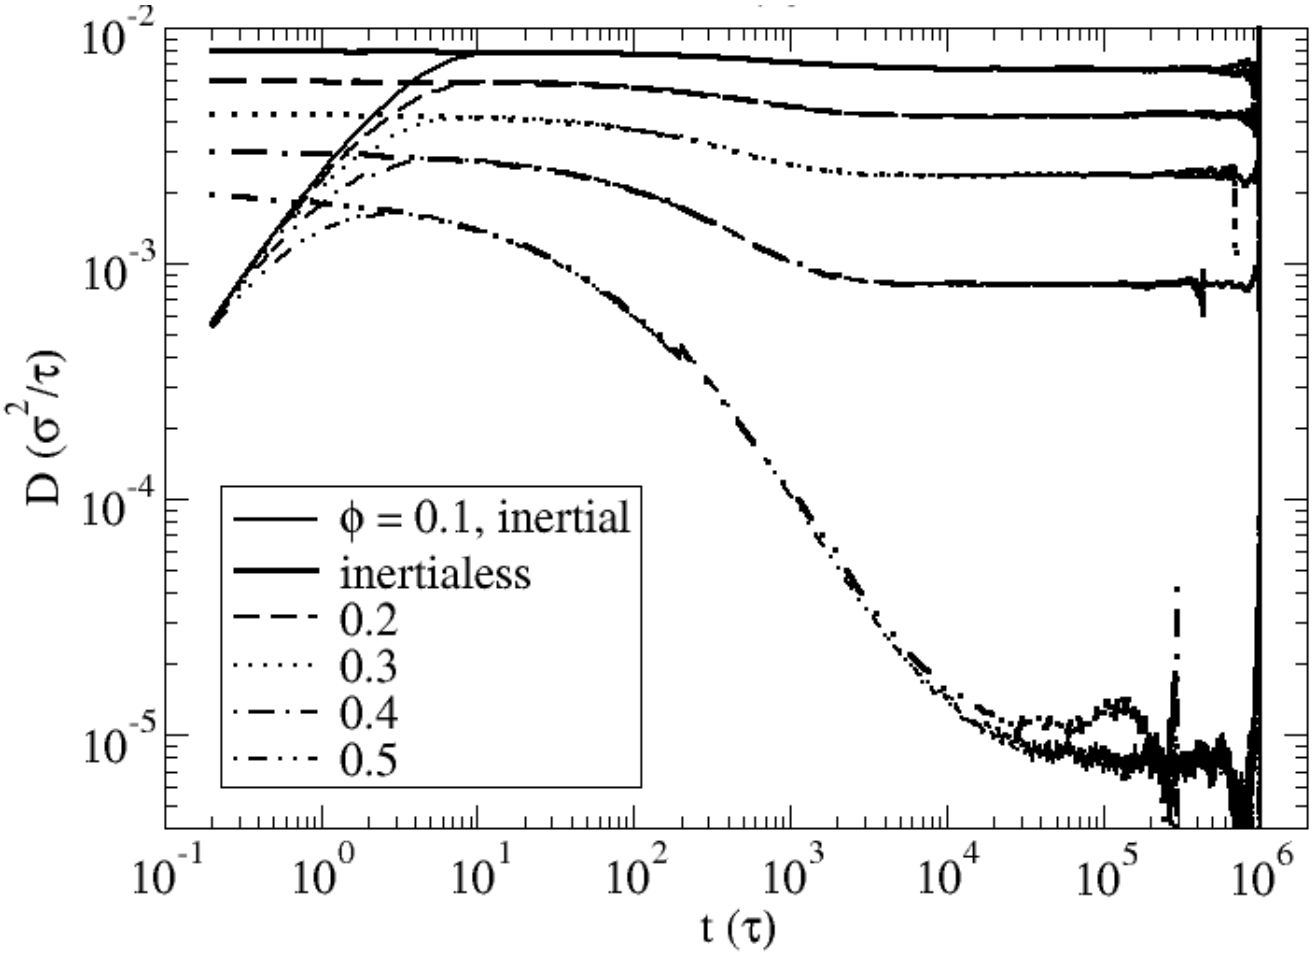

Bolintineanu, et al, Comp Particle Mechanics, 1, 321-356 (2014).

Viscosity of nanoparticles in SRD fluid

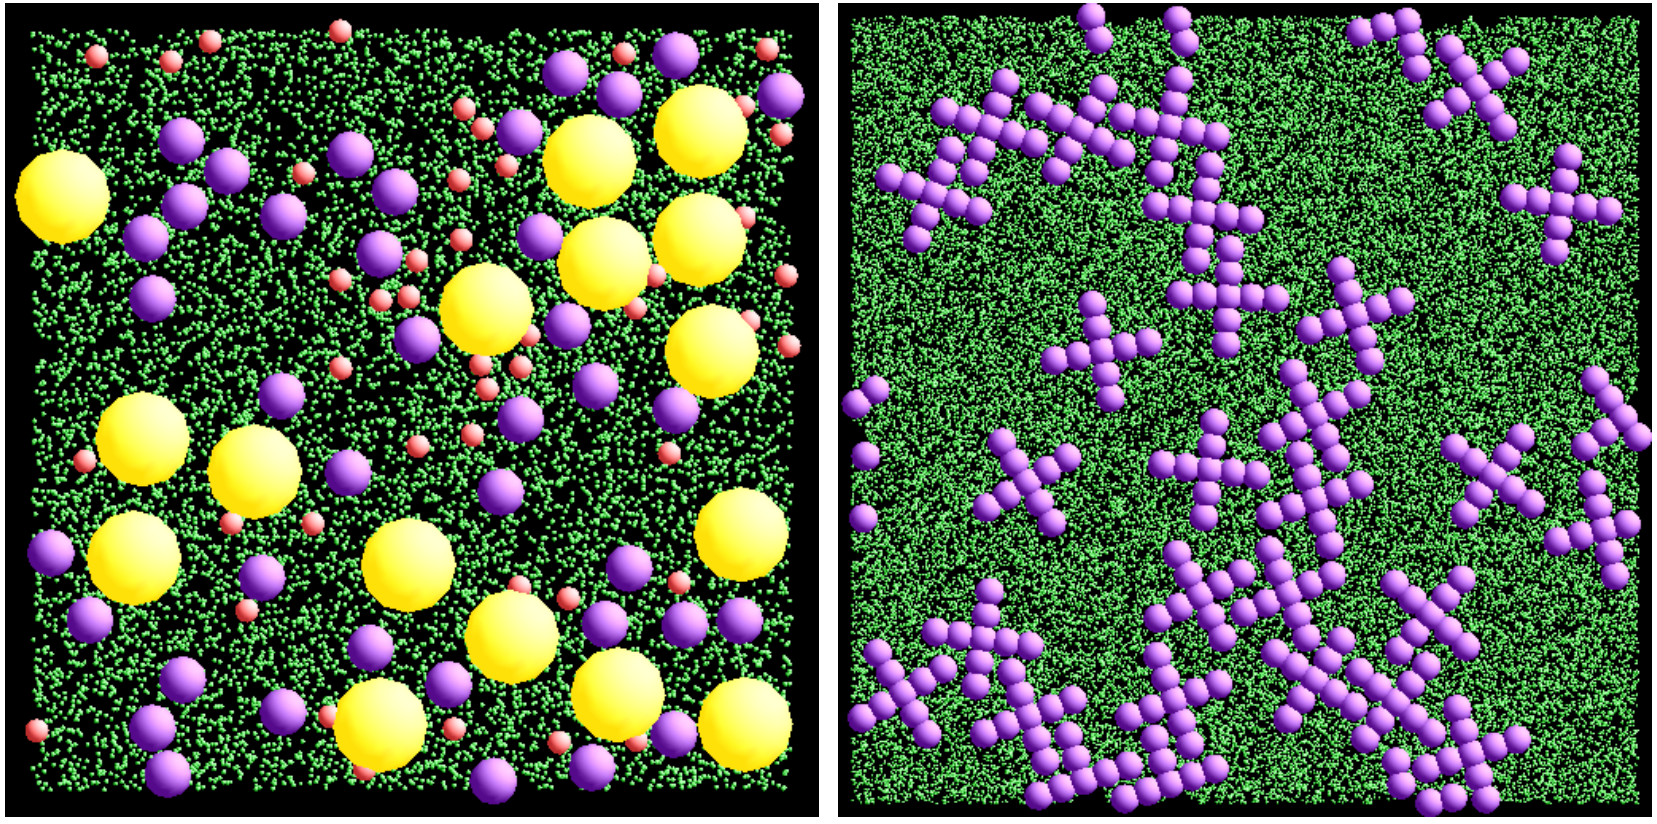

Muller-Plathe algorithm induces V-shaped velocity profile

Arbitrary-shape aspherical NPs with and w/out solvent

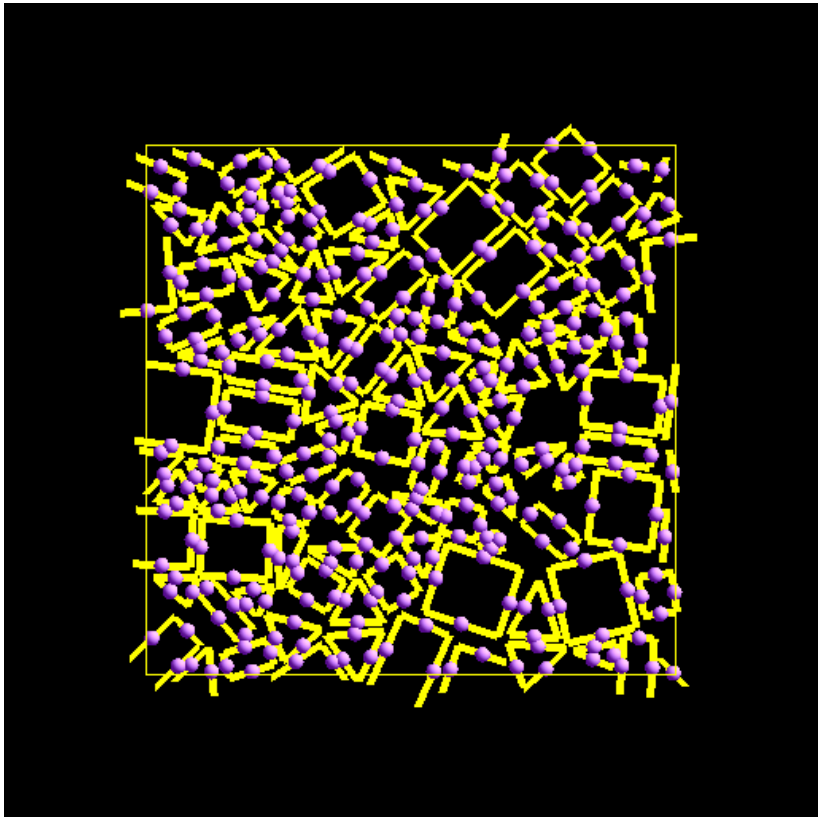

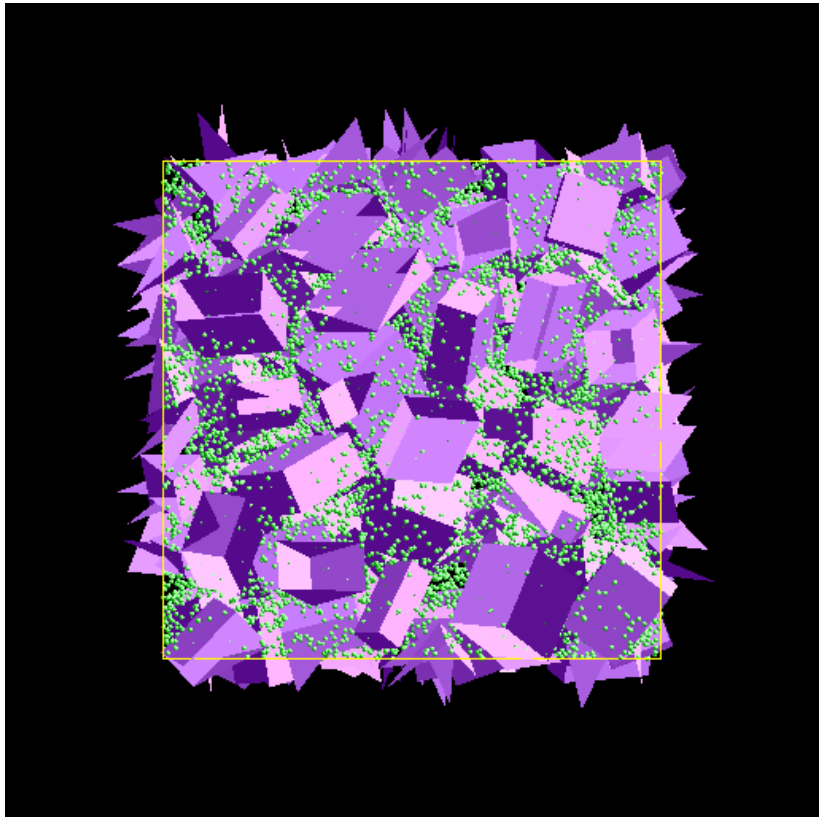

Rigid bodies

2d particles are line-segmented surfaces

3d particles triangulated surfaces

Liquid crystal thin film rupture on Titan

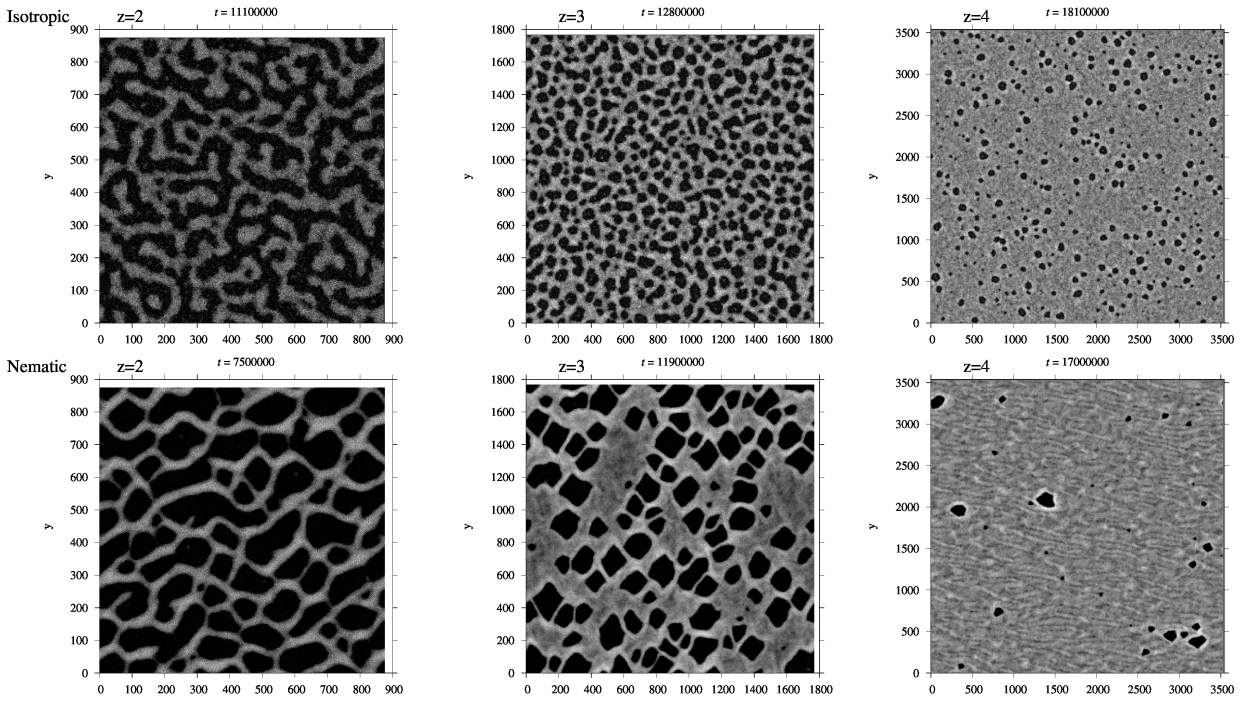

- Mike Brown and collaborators (ORNL)
- Liquid crystal molecule = coarse-grained ellipsoid
- GayBerne potential on GPU = ∼100x faster than CPU core
- Titan GPU/CPU node is ∼7x vs two multi-core CPUs
- *Nguyen, et al, Nanoscale, 6, 3083-96 (2014).*

Load-balancing via recursive coordinate bisectioning

- Often needed for coarse-grained models
  - DPD, SPH, Peridynamics, granular, etc

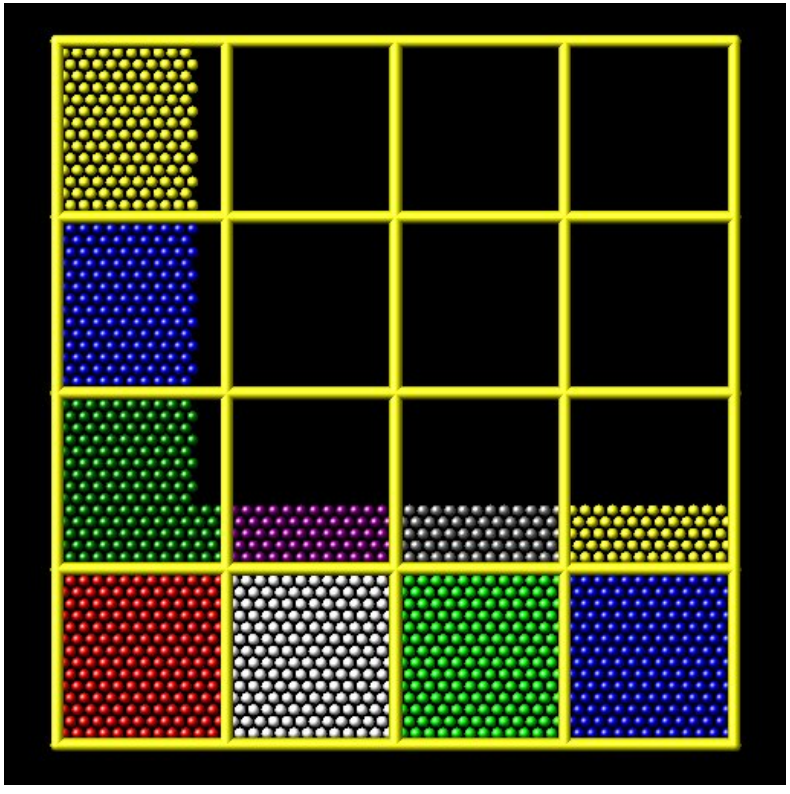

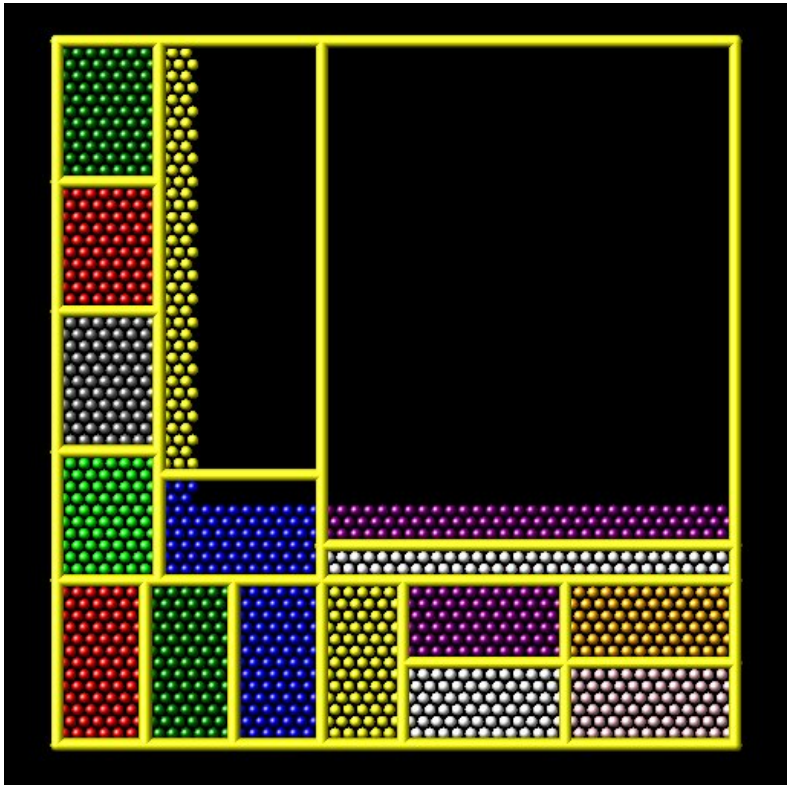

- Worked to reduce comm with 26 neighbors to 6+ (for 3d)

Load-balancing examples for soft and hard materials

2d SPH "water" flowing over a dam
Georg Ganzenmueller, Fraunhofer EMI

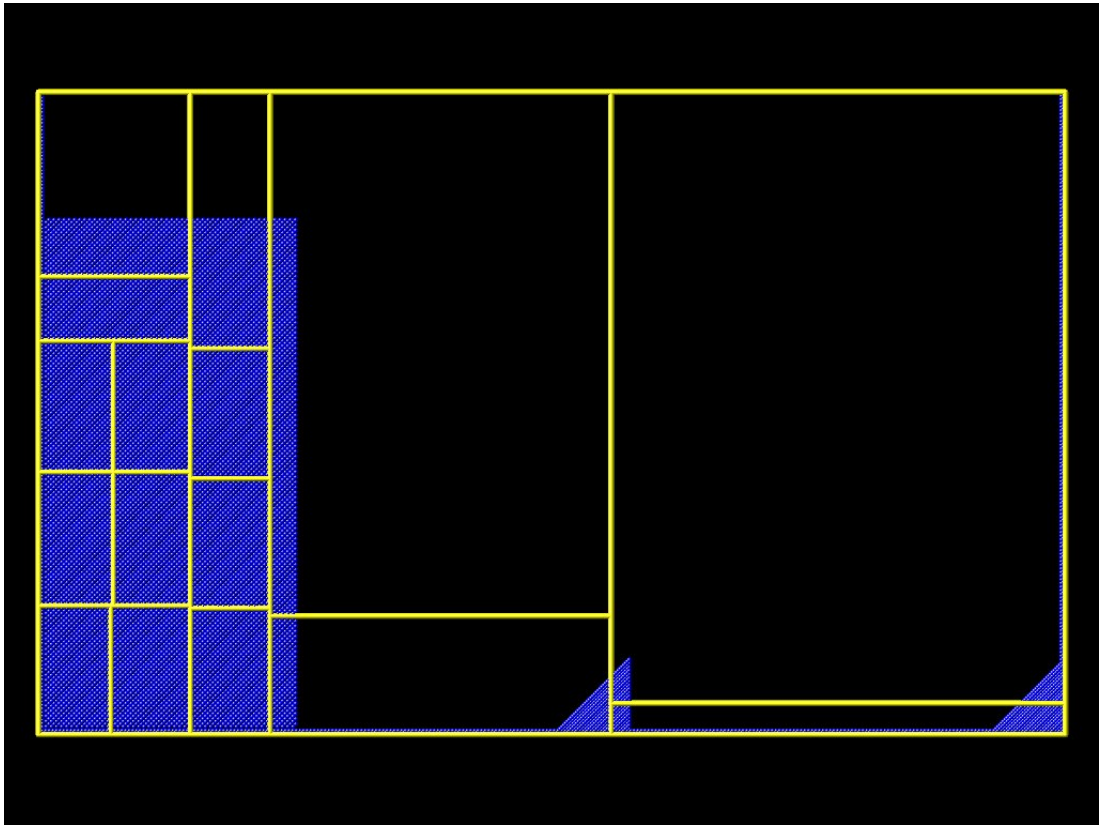

Load-balancing examples for soft and hard materials

2d SPH "water" flowing over a dam
Georg Ganzenmueller, Fraunhofer EMI

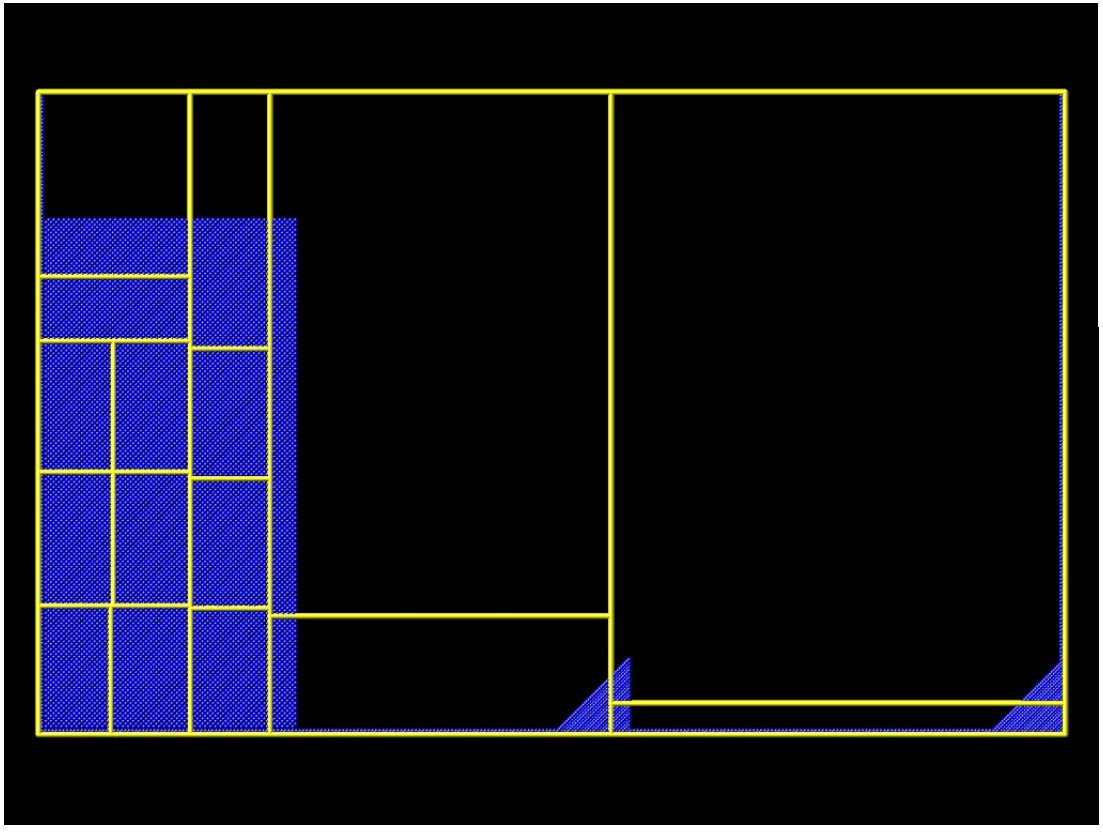

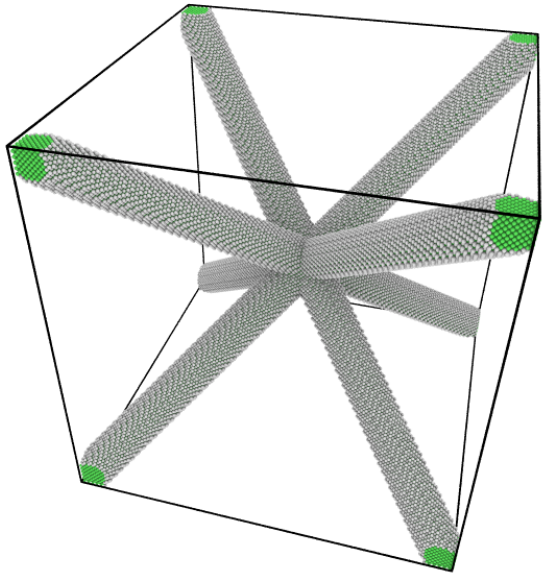

Atomic microlattice of metal struts
Alex Stukowski
(Tech Univ Darmstadt)

- star imbalance = 18x
- 13x speed-up for 21M atoms on 16K cores

Accelerator hardware: Aiming for MPI+X via Kokkos

- **Kokkos** = programming model developed at Sandia
  - minimize impact of new chip designs on applications
  - https://github.com/kokkos
- Goal: write application kernels only once,
  run efficiently on variety of current/future hardware

Accelerator hardware: Aiming for MPI+X via Kokkos

- **Kokkos** = programming model developed at Sandia
  - minimize impact of new chip designs on applications
  - https://github.com/kokkos
- Goal: write application kernels only once,
  run efficiently on variety of current/future hardware
- Two major components:
  1. Data access abstraction via **Kokkos arrays**
     - optimal layout & access pattern for each device
       GPU, Xeon Phi, etc
  2. **Parallel dispatch** of small chunks of work
     - auto-mapped onto back-end languages
       CUDA, OpenMP, etc

Accelerator hardware: Aiming for MPI+X via Kokkos

- **Kokkos** = programming model developed at Sandia
  - minimize impact of new chip designs on applications
  - https://github.com/kokkos
- Goal: write application kernels only once,
  run efficiently on variety of current/future hardware
- Two major components:
  1. Data access abstraction via **Kokkos arrays**
     - optimal layout & access pattern for each device
       GPU, Xeon Phi, etc
  2. **Parallel dispatch** of small chunks of work
     - auto-mapped onto back-end languages
       CUDA, OpenMP, etc
- Key task for application is to write **kernels** so they:
  - operate at fine granularity and are thread-safe
  - use Kokkos data structures (dual-view concept)
  - unfortunately LAMMPS has ~1000 "kernels"

Kokkos performance for Lennard-Jones pairwise kernel

One-node performance :

dual 8-core Intel Sandy Bridge Xeon CPUs

two NVIDIA Kepler GPUs

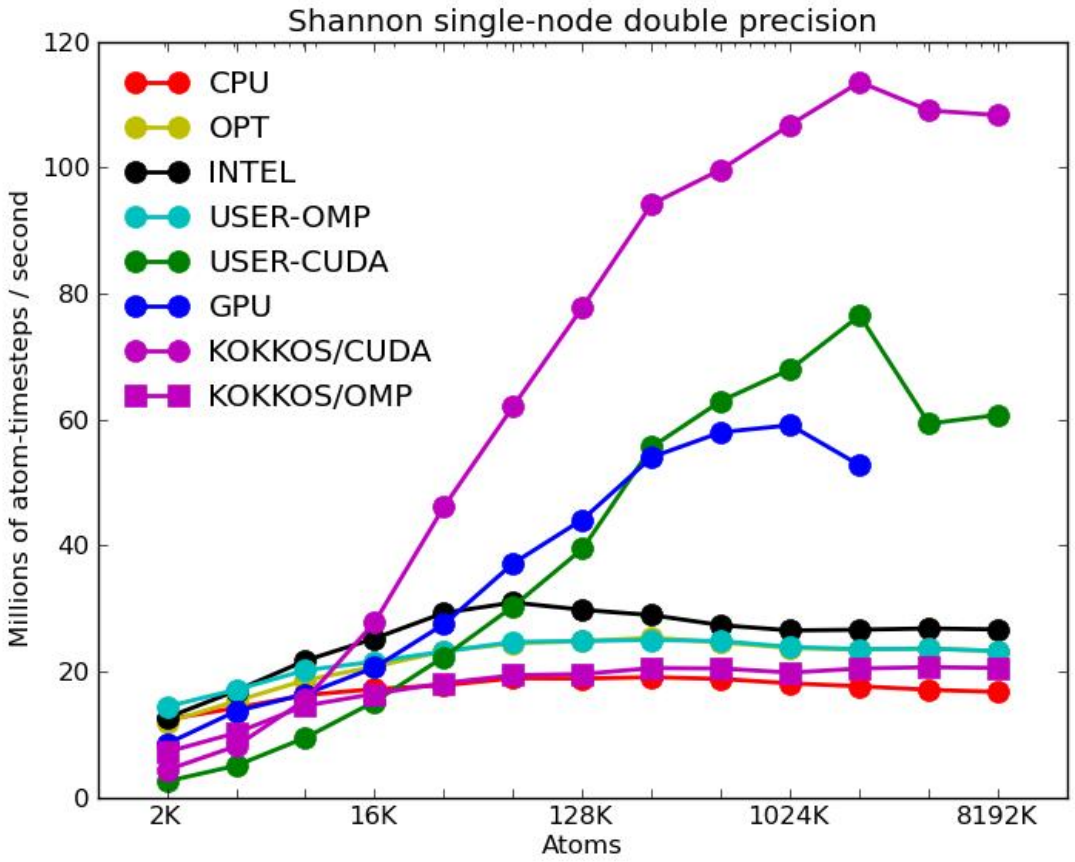

Kokkos for manybody potentials

- Completed: EAM, Tersoff, Stillinger-Weber
- Working on ReaxFF = reactive bond-order potential

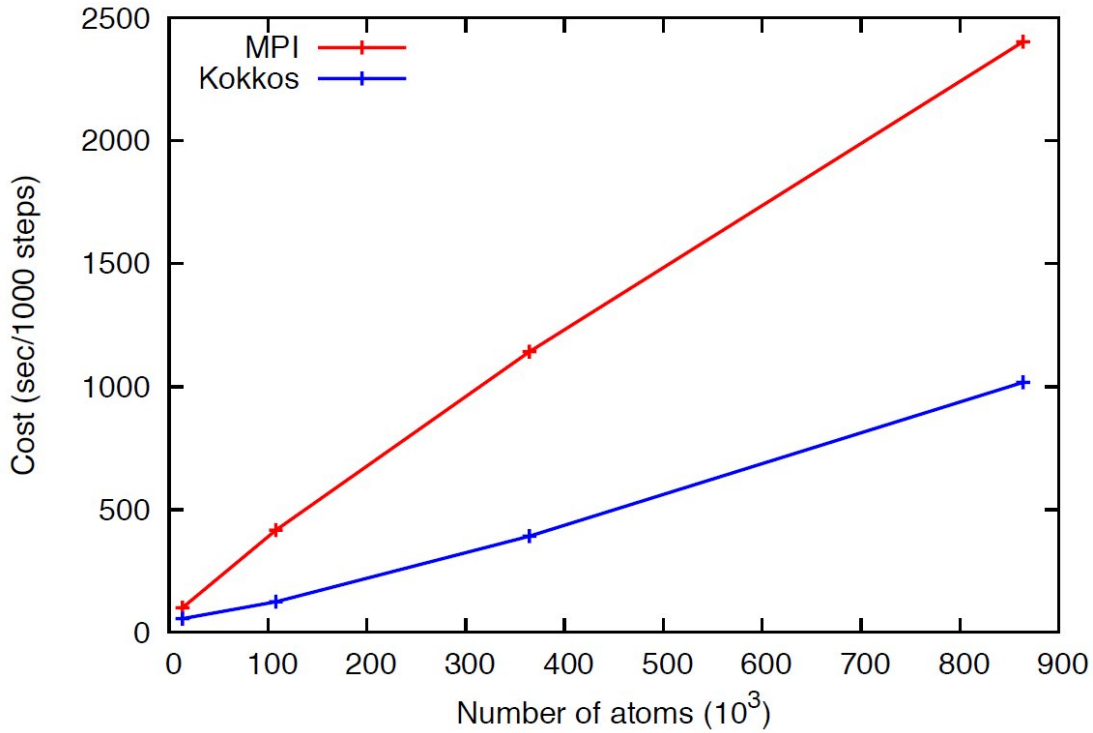

- About 50% complete: bond-order terms, Coulomb, LJ
- Todo: many-body, QEq (matrix solve or damped dynamics)

Multiscale & multiphysics via coupling to other codes

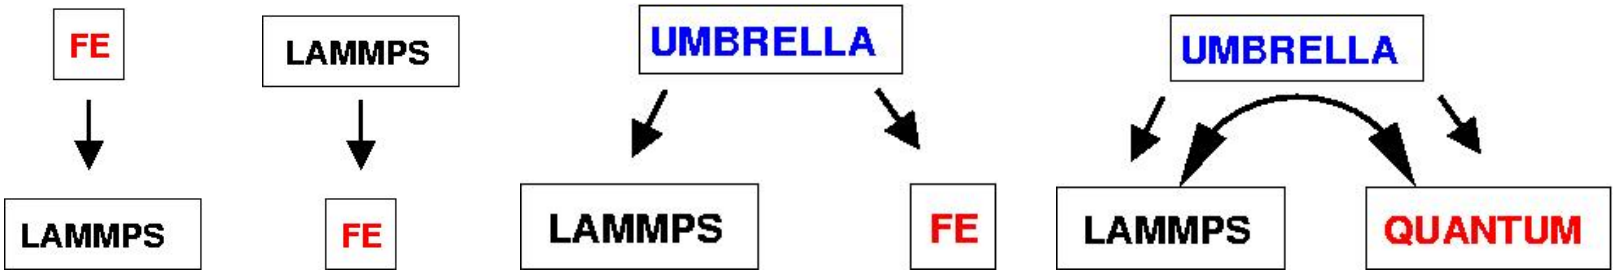

Goal: enable "easy" coupling of MD to QM, kMC, FE ...
via Python or C-style lib interface (C/C++/Fortran/etc)

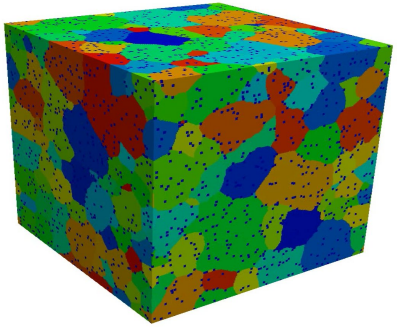

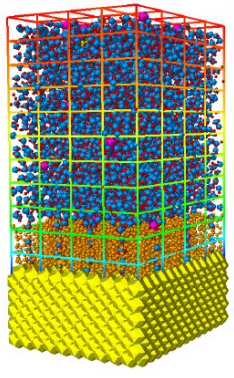

Saltwater-electrode-CNT system: mesh overlaps exactly with water-CNT atom region

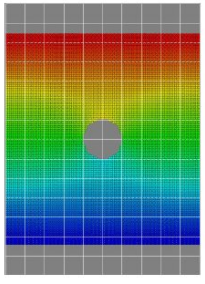

Circular hole in plate: mesh overlaps exactly with box, but atom region is subset

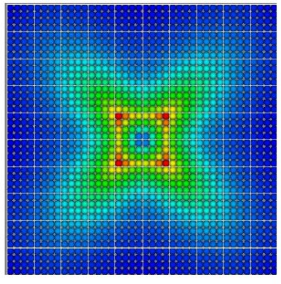

Elastic inclusion problem: mesh overlaps exactly with box and atoms

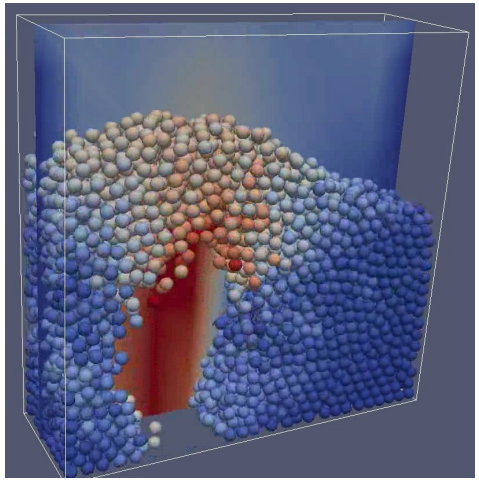

Thermal coupling with AtC package

Reese Jones, Jon Zimmerman, Jeremy Templeton,
Greg Wagner (Sandia)

2D diffusion problem

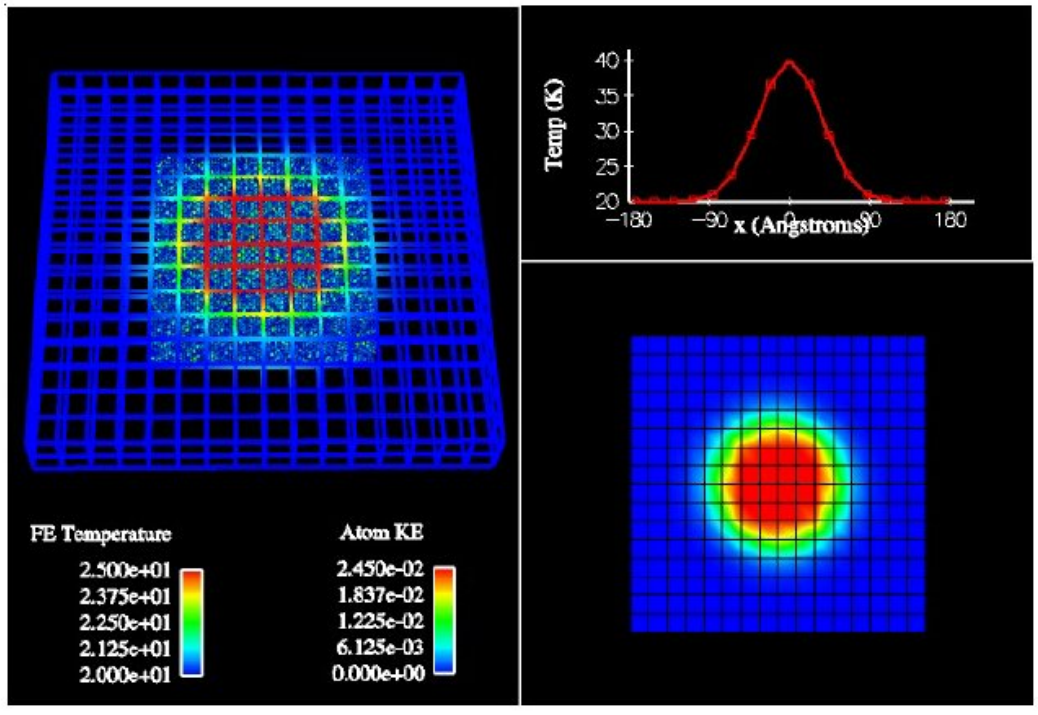

- Plate with embedded MD region (~33,000 atoms)
- Initialized to temperature field with gaussian profile
- Adiabatic boundary conditions at edges

Sandia National Laboratories

Mechanical coupling with AtC package

Elasto-dynamic response:

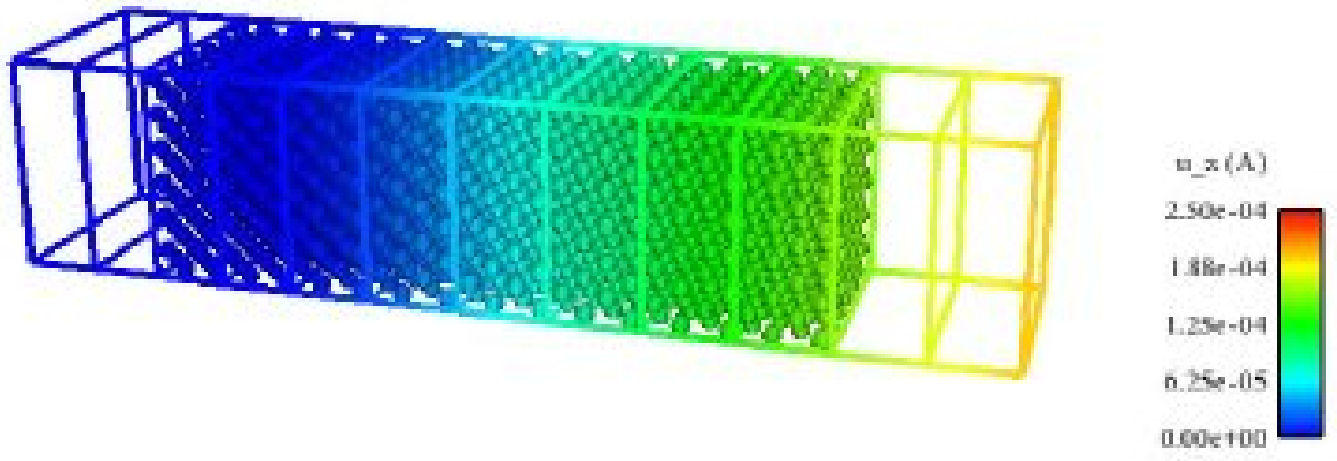

Granular + fluids modeling via OpenFOAM

Christoph Kloss (JKU) and add-on LIGGGHTS package
www.liggghts.com/www.cfdem.com

- particles + CAD mesh + **fluid**

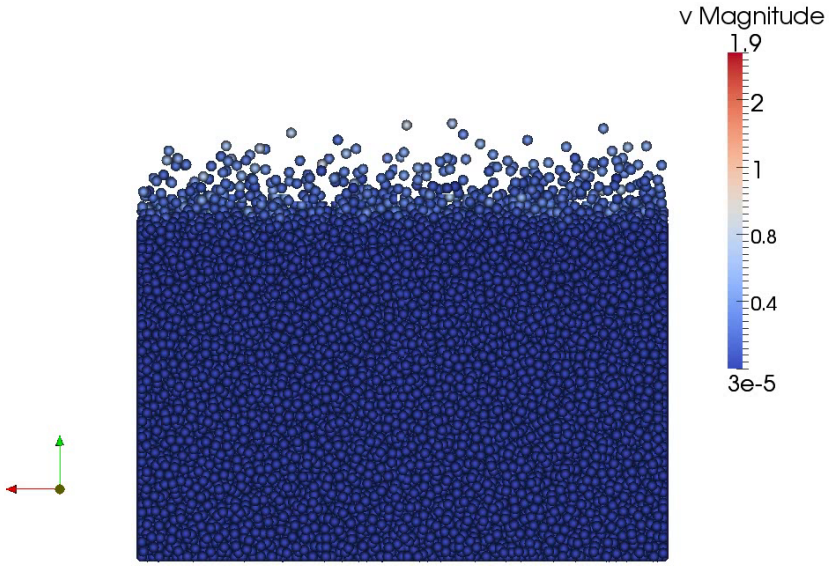

MD + kMC for stress-driven grain growth

- SPPARKS (kMC) runs Potts model for grain growth
  - Hamiltonian includes stress term
  - send **grain structure** to LAMMPS
- LAMMPS (MD) treats particles at grain boundary as larger
  - relaxes system
  - send per-particle **stress** to SPPARKS

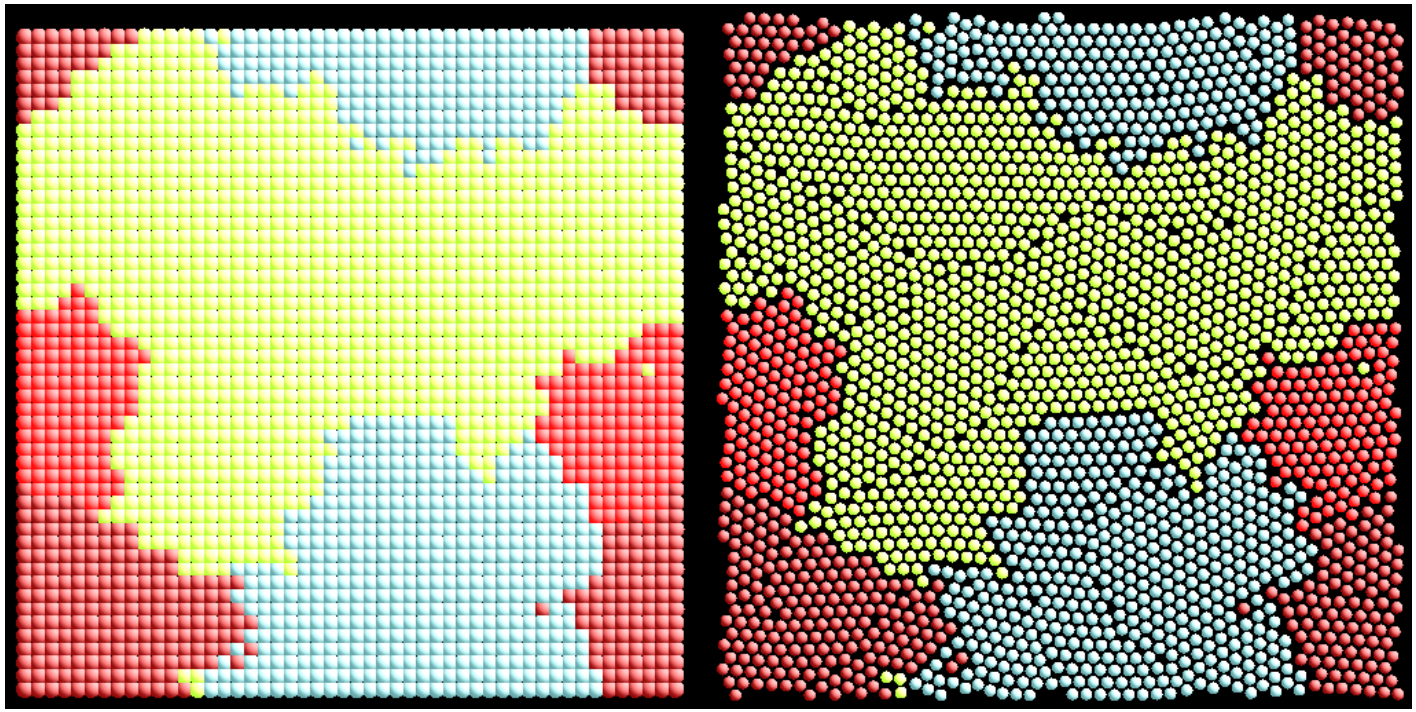

A think-outside-the-box example ...

LIGGGHTS package extension to LAMMPS for **granular models** and
FMI (Functional Mock-up Interface) for **mesh dynamics**
LIGGGHTS by C Kloss (DCS Computing)
Wheelloader model by C Schubert & T Dresden (Dresden Tech U)
Simulation by C. Richter & A. Katterfeld (U Magdeburg OV Guericke)


```
simtime :   0.000
throttle:   0.0%
brake   :   0.0%
gear    :       1
KiZyIn  :       0
HuZyIn  :      -1
LeZyIn  :       0
```

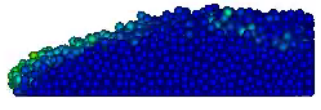

Thanks and links

- LAMMPS: **http://lammps.sandia.gov**
- **Funding**:
  - DOE (BES,BER), Sandia (ASC,LDRD)
  - NINE (university), CRADA with Corning, 3M, BASF
- **Joint work** with:
  - LAMMPS: Aidan Thompson, Paul Crozier, Stan Moore, Ray Shan, Axel Kohlmeyer (Temple U)
  - Kokkos: Carter Edwards & Christian Trott
- Two papers with more info:
  - *S. J. Plimpton and A. P. Thompson, "Computational Aspects of Many-body Potentials", MRS Bulletin, 37, 513-521 (2012).*
  - *D. S. Bolintineanu, et al, "Particle dynamics modeling methods for colloid suspensions", Comp Particle Mechanics, 1, 321-356 (2014).*