A Brief Overview of Molecular Dynamics, Statistical Mechanics, Atomic Potentials

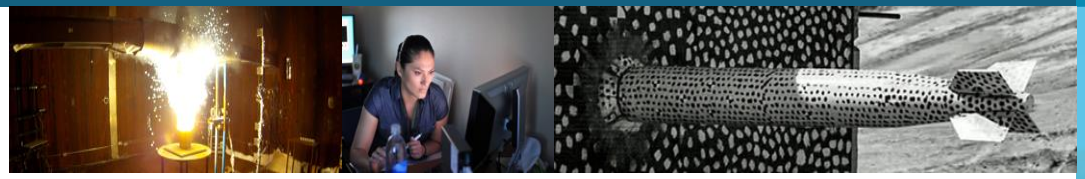

Stan Moore

2019 LAMMPS Workshop

Albuquerque, NM

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Molecular Dynamics: What is it?

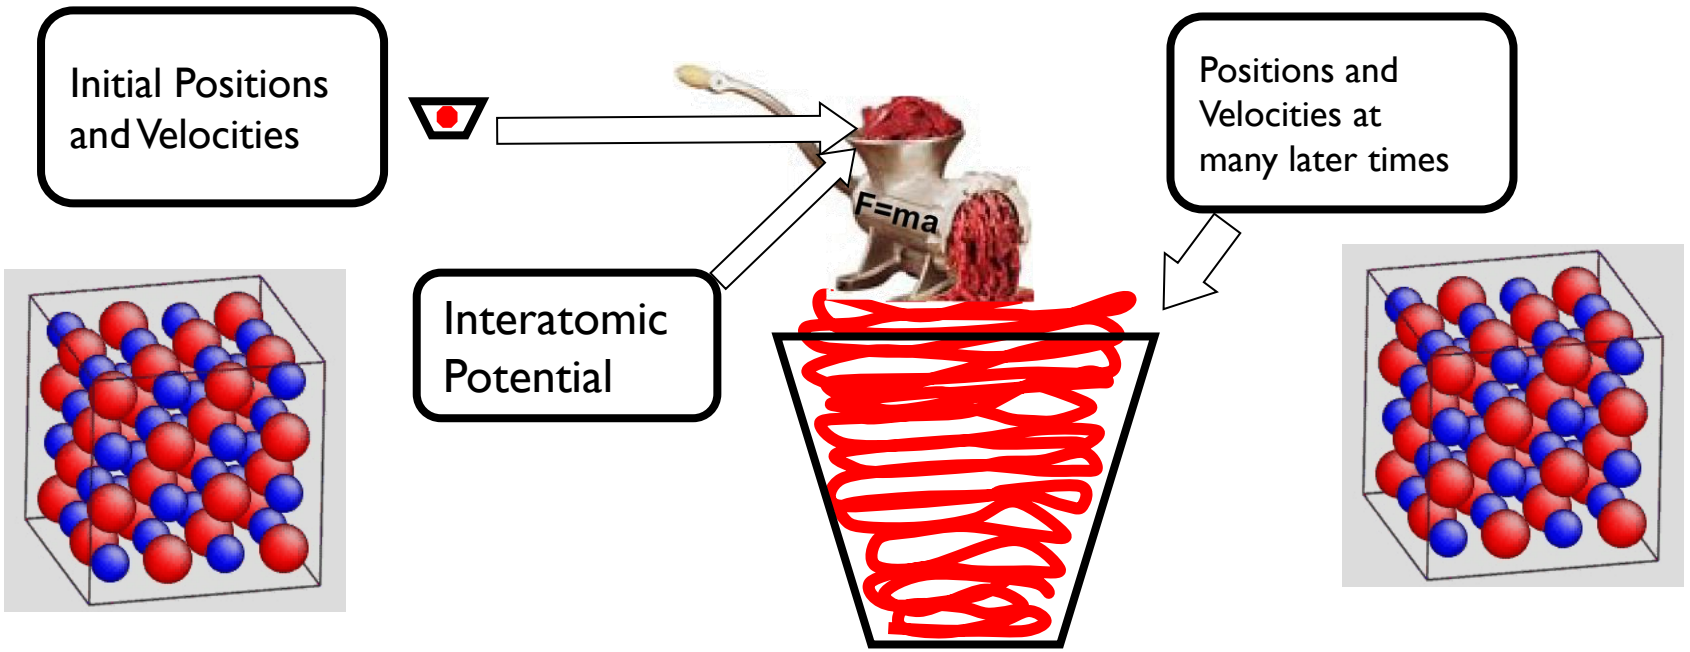

Mathematical Formulation

Classical Mechanics

Atoms are Point Masses: $r_1, r_2, \ldots . r_N$

Positions, Velocities, Forces: r_i, v_i, F_i

Potential Energy Function = $V(r^N)$

$6N$ coupled ODEs

$$\frac{d\mathbf{r}_i}{dt} = \mathbf{v}_i$$

$$\frac{d\mathbf{v}_i}{dt} = \frac{\mathbf{F}_i}{m_i}$$

$$\mathbf{F}_i = -\frac{d}{d\mathbf{r}_i} V\left(\mathbf{r}^N\right)$$

What is MD good for?

Quantum mechanical electronic structure calculations (QM) provide accurate description of mechanical and chemical changes on the atom-scale: 10x10x10~1000 atoms

Atom-scale phenomena drive a lot of interesting physics, chemistry, materials science, mechanics, biology…but it usually plays out on a much larger scale

Mesoscale: much bigger than an atom, much smaller than a glass of soda.

QM and continuum/mesoscale models (CM) can not be directly compared.

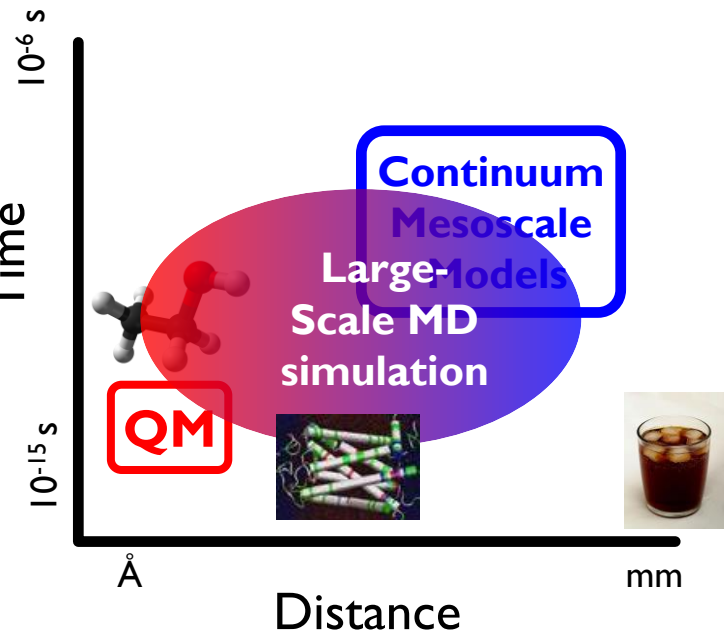

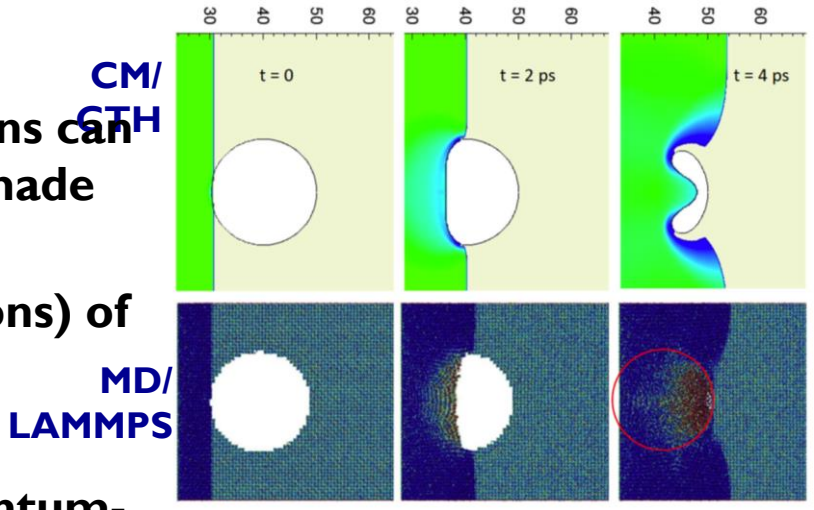

CTH images courtesy of David Damm, Sandia

- **Small molecular dynamics (MD) simulations can be directly compared to QM results, and made to reproduce them**
- **MD can also be scaled up to millions (billions) of atoms, overlapping the low-end of CM**
- **Limitations of MD orthogonal to CM**
- **Enables us to inform CM models with quantum-accurate results**

MD Versatility

Coupling to Solid Mechanics

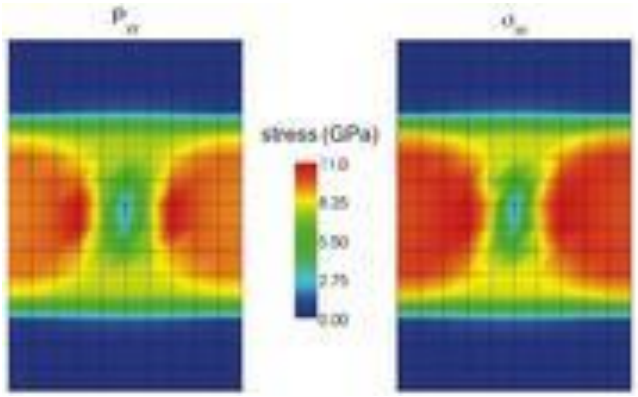

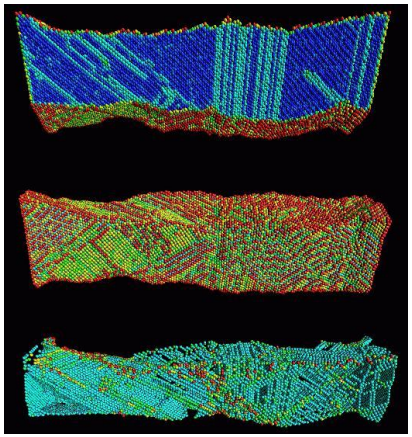

Materials Science

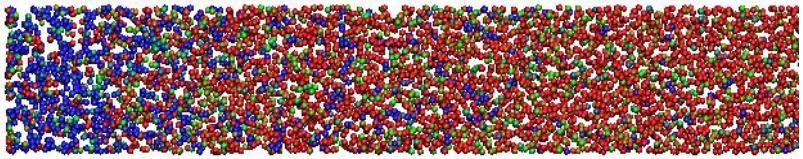

Chemistry

Biophysics

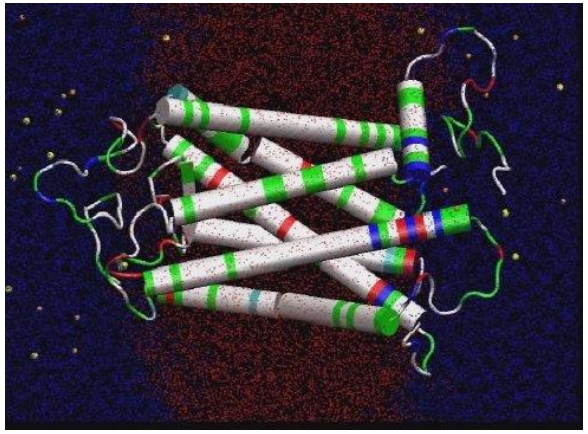

Granular Flow

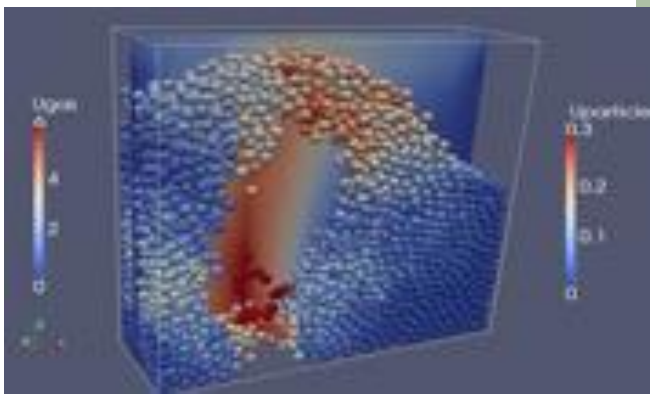

MD Basics

Particles are normally modeled as spheres, though other shapes (e.g. ellipsoids) are possible

Typically use a rectangular or triclinic simulation cell

Commonly use periodic boundary conditions: reduces finite size effects from boundaries and simulates bulk conditions

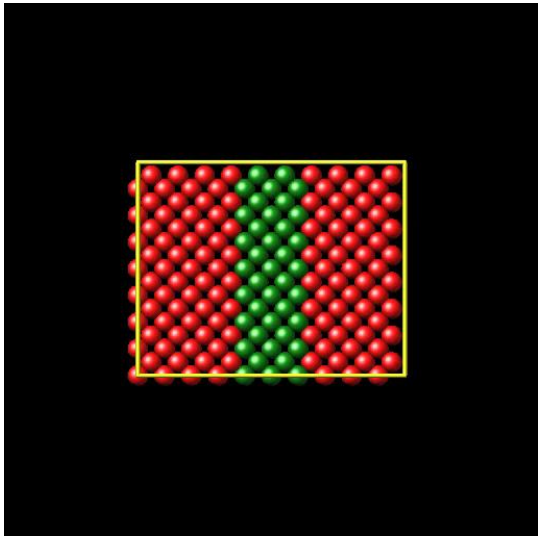

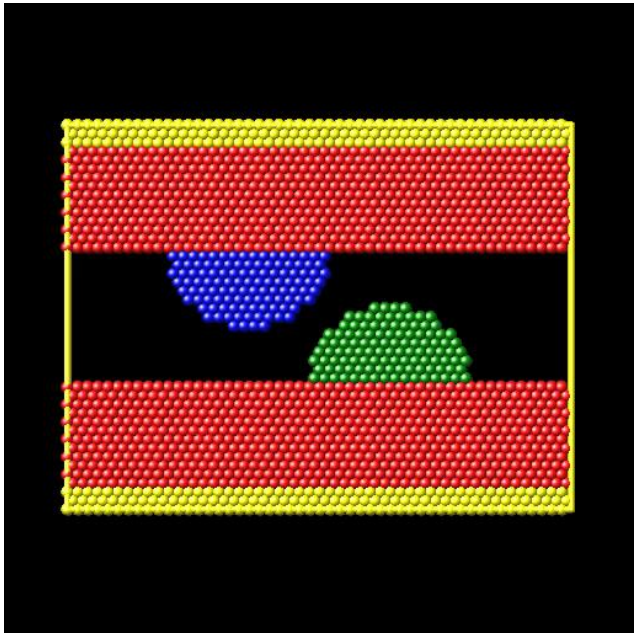

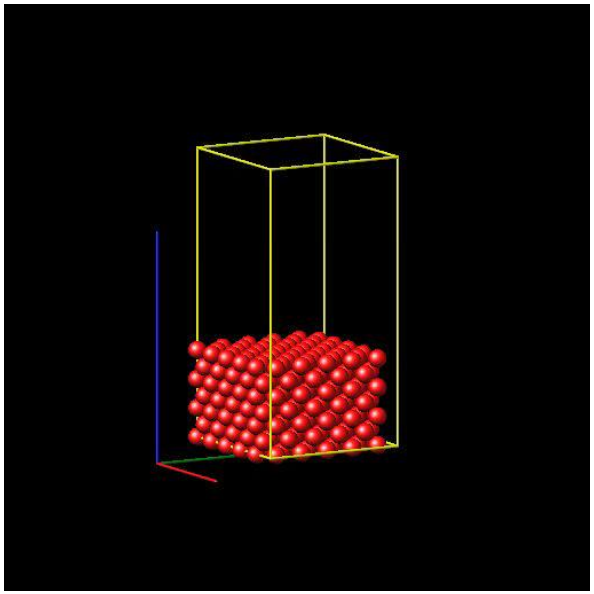

MD Time Integration Algorithm

- Most codes and applications use variations and extensions to the Størmer-Verlet explicit integrator:
- Only second-order : $\delta E = |\langle E \rangle - E_0| \sim \Delta t^2$, but….
- **time-reversible map**: switching sign of Δt takes you back to initial state
- **measure-preserving**:Volume of differential cube (δ**v**,δ**x**) is conserved (but not shape).
- **symplectic:** Conserves sum of areas of differential parallelogram (δ**v**,δ**x**) projected onto each particular (v_i,x_i) plane

Velocity form of Størmer-Verlet

$$\textit{For istep} < \textit{nsteps}:$$

$$\mathbf{v} \leftarrow \mathbf{v} + \frac{\Delta t}{2}\mathbf{F}$$

$$\mathbf{x} \leftarrow \mathbf{x} + \Delta t\,\mathbf{v}$$

$$\textit{Compute}\ \mathbf{F}(\mathbf{x})$$

$$\mathbf{v} \leftarrow \mathbf{v} + \frac{\Delta t}{2}\mathbf{F}$$

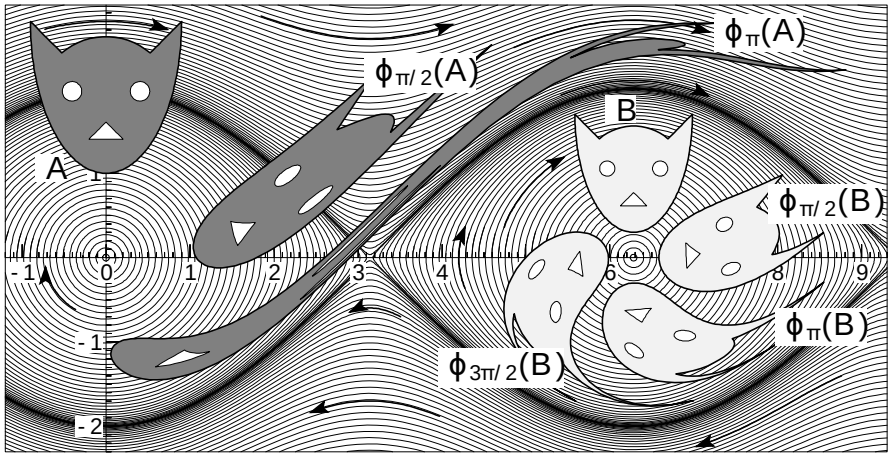

Figure 3: Area preservation of the flow of Hamiltonian systems

Ernst Hairer, Lubich, Wanner, *Geometric Numerical Integration* (2006)

MD Time Integration Algorithm

- **time-reversibility and symplecticity**: global stability of Verlet trumps local accuracy of high-order schemes
- More specifically, it can be shown that for Hamiltonian equations of motion, Størmer-Verlet exactly conserves a "shadow" Hamiltonian and $E-E_S \sim O(\Delta t^2)$
- For users: no energy drift over millions of timesteps
- For developers: easy to decouple integration scheme from efficient algorithms for force evaluation, parallelization.
- Symplectic high-order Runge-Kutta methods exist, but not widely adopted for MD

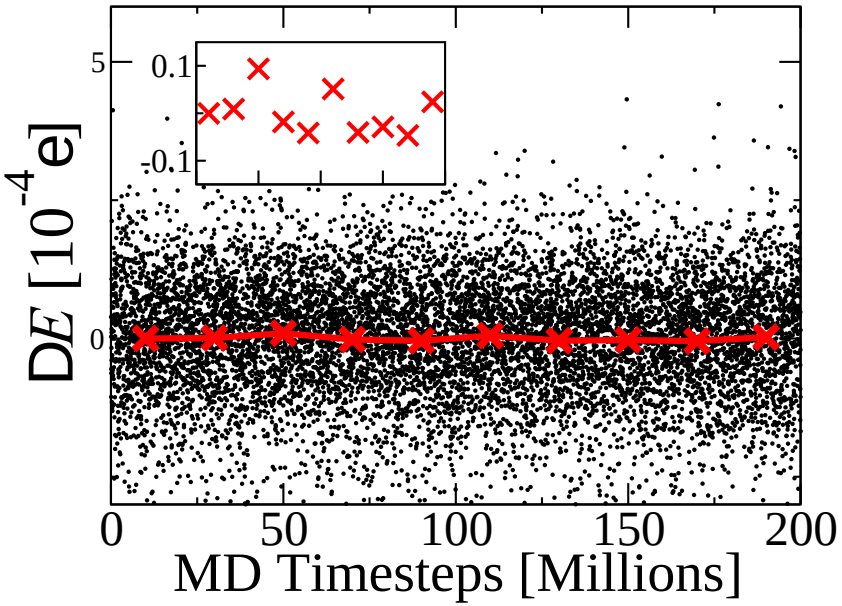

32 atom LJ cluster, 200 million MD steps, $\Delta t=0.005, T=0.4$

Statistical Mechanics Basics

Statistical Mechanics: relates macroscopic observations (such as temperature and pressure) to microscopic states (i.e. atoms)

Phase space: a space in which all possible states of a system are represented. For N particles: $6N$-dimensional phase space (3 position variables and 3 momentum variables for each particle)

Ensemble: an idealization consisting of a large number of virtual copies of a system, considered all at once, each of which represents a possible state that the real system might be in, i.e. a probability distribution for the state of the system

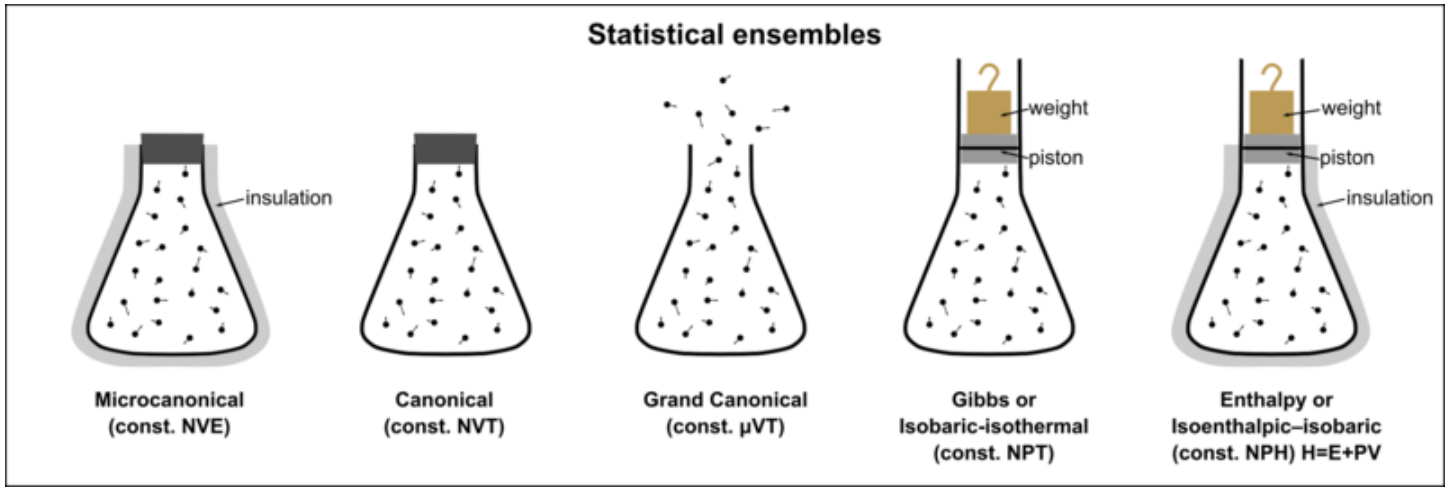

Thermostats and Barostats

Using the velocity-verlet time integrator gives the microcanonical ensemble (NVE). How to simulate canonical (NVT) or isothermal-isobaric (NPT) ensembles?

Temperature is related to atom velocities through statistical mechanics, pressure is related to volume of the simulation cell

Could just scale velocities and volume to the exact desired values, but this does not allow for fluctuations with a distribution typical for the ensemble

Instead Nose-Hoover style integrators are commonly used: dynamic variables are coupled to the particle velocities (thermostatting) and simulation box dimensions (barostatting)

Nose-Hoover uses a *damping* parameter specified in time units which determines how rapidly the temperature or pressure is relaxed. If the damping parameter is too small, the temperature/pressure can fluctuate wildly; if it is too large, the temperature/pressure will take too long to equilibrate

Intermolecular Potentials

Quantum chemistry: solves Schrödinger equation to get forces on atoms. Accurate but very computationally expensive and only feasible for small systems

Molecular dynamics: uses empirical force fields, sometimes fit to quantum data. Not as accurate but much faster.

Typically only interact with atoms in a spherical cutoff and only consider pair-wise or three-body interactions

Lennard-Jones Potential

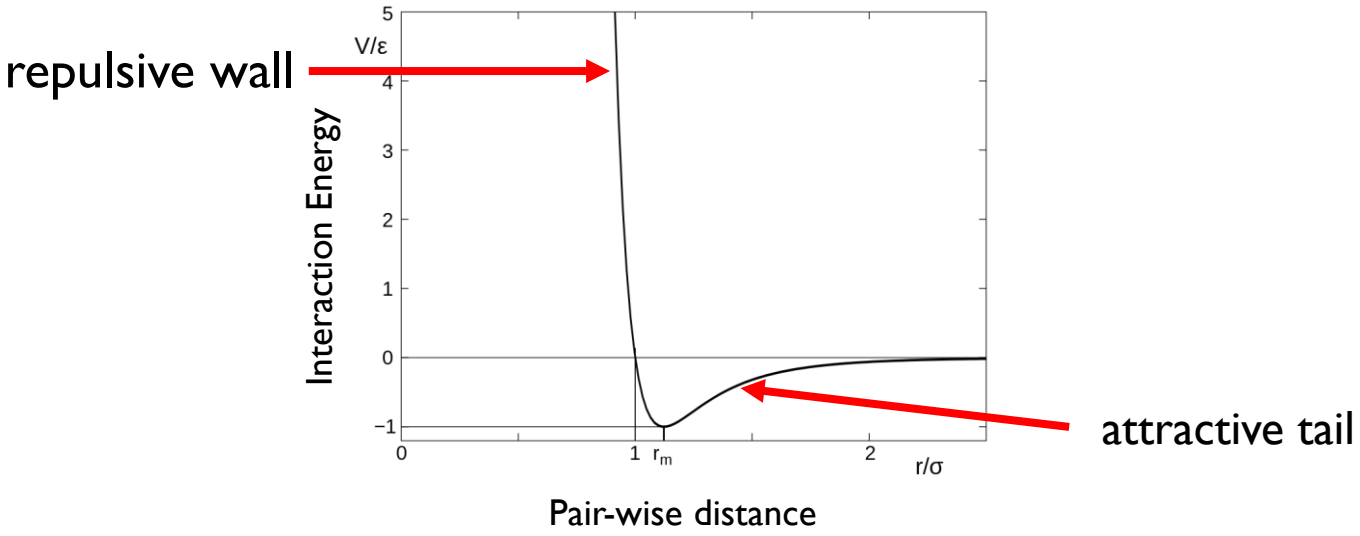

Interatomic Potentials

LAMMPS Potentials by Material

Biomolecules: CHARMM, AMBER, OPLS, COMPASS (class 2), long-range Coulombics via PPPM, point dipoles, ...

Polymers: all-atom, united-atom, coarse-grain (bead-spring FENE), bond-breaking, …

Materials: EAM and MEAM for metals, Buckingham, Morse, Yukawa, Stillinger-Weber, Tersoff, COMB, SNAP, ...

Chemistry: AI-REBO, REBO, ReaxFF, eFF

Mesoscale: granular, DPD, Gay-Berne, colloidal, peridynamics, DSMC...

Hybrid: can use combinations of potentials for hybrid systems: water on metal, polymers/semiconductor interface, colloids in solution, …

More Interatomic Potentials

LAMMPS Potentials by Functional Form
pairwise potentials: Lennard-Jones, Buckingham, ...

charged pairwise potentials: Coulombic, point-dipole

manybody potentials: EAM, Finnis/Sinclair, modified EAM (MEAM), embedded ion (EIM), Stillinger-Weber, Tersoff, AI-REBO, ReaxFF, COMB

coarse-grained potentials: DPD, GayBerne, ...

mesoscopic potentials: granular, peridynamics

long-range electrostatics: Ewald, PPPM, MSM

implicit solvent potentials: hydrodynamic lubrication, Debye
force-field compatibility with common CHARMM, AMBER, OPLS, GROMACS options

LAMMPS Potentials

See lammps.sandia.gov/bench.html#potentials

Potential	System	Atoms	Timestep	CPU	LJ Ratio
Granular	chute flow	32000	0.0001 tau	5.08e-7	0.34x
FENE bead/spring	polymer melt	32000	0.012 tau	5.32e-7	0.36x
Lennard-Jones	LJ liquid	32000	0.005 tau	1.48e-6	1.0x
DPD	pure solvent	32000	0.04 tau	2.16e-6	1.46x
EAM	bulk Cu	32000	5 fmsec	3.59e-6	2.4x
Tersoff	bulk Si	32000	1 fmsec	6.01e-6	4.1x
Stillinger-Weber	bulk Si	32000	1 fmsec	6.10e-6	4.1x
EIM	crystalline NaCl	32000	0.5 fmsec	9.69e-6	6.5x
SPC/E	liquid water	36000	2 fmsec	1.43e-5	9.7x
CHARMM + PPPM	solvated protein	32000	2 fmsec	2.01e-5	13.6x
MEAM	bulk Ni	32000	5 fmsec	2.31e-5	15.6x
Peridynamics	glass fracture	32000	22.2 nsec	2.42e-5	16.4x
Gay-Berne	ellipsoid mixture	32768	0.002 tau	4.09e-5	28.3x
AIREBO	polyethylene	32640	0.5 fmsec	8.09e-5	54.7x
COMB	crystalline SiO2	32400	0.2 fmsec	4.19e-4	284x
eFF	H plasma	32000	0.001 fmsec	4.52e-4	306x
ReaxFF	PETN crystal	16240	0.1 fmsec	4.99e-4	337x
ReaxFF/C	PETN crystal	32480	0.1 fmsec	2.73e-4	185x
VASP/small	water	192/512	0.3 fmsec	26.2	17.7e6
VASP/medium	CO2	192/1024	0.8 fmsec	252	170e6
VASP/large	Xe	432/3456	2.0 fmsec	1344	908e6

Accuracy = Higher Cost

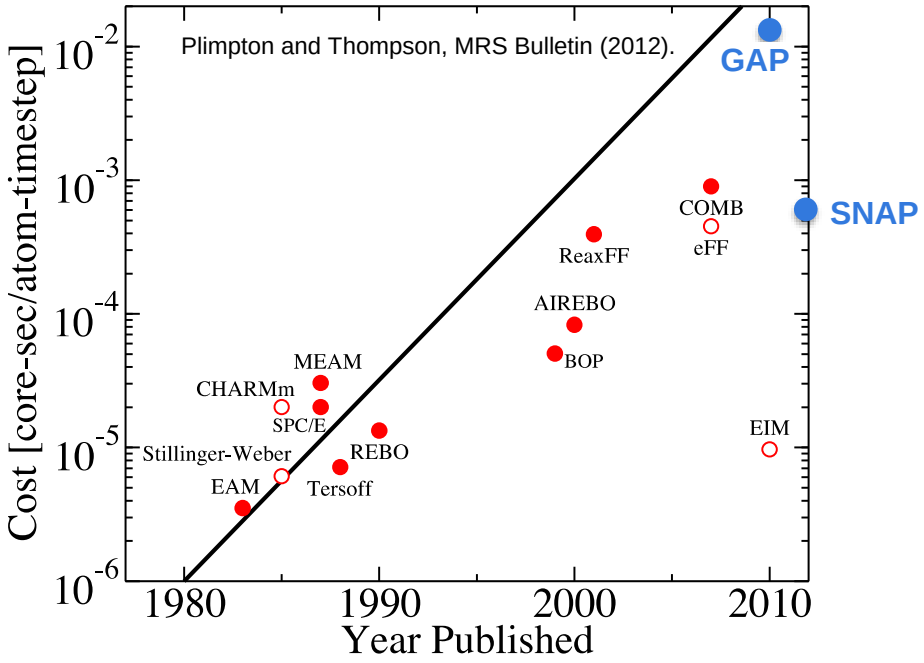

Moore's Law for Interatomic Potentials

Plimpton and Thompson, MRS Bulletin (2012).

Hybrid force fields

Metal droplet on an LJ surface

- metal → metal atoms interact with Embedded-Atom potential
- surface → surface atoms interact with the Lennard-Jones potential
- metal/surface interaction is also computed via the Lennard-Jones potential

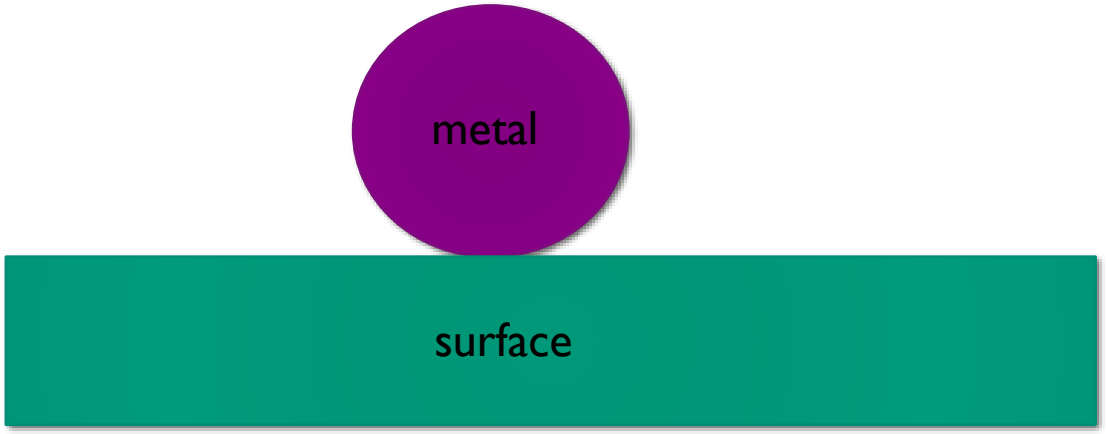

Neighbor Lists

Neighbor lists are a list of neighboring atoms within the interaction cutoff + skin for each central atom

Extra skin allows lists to be built less often

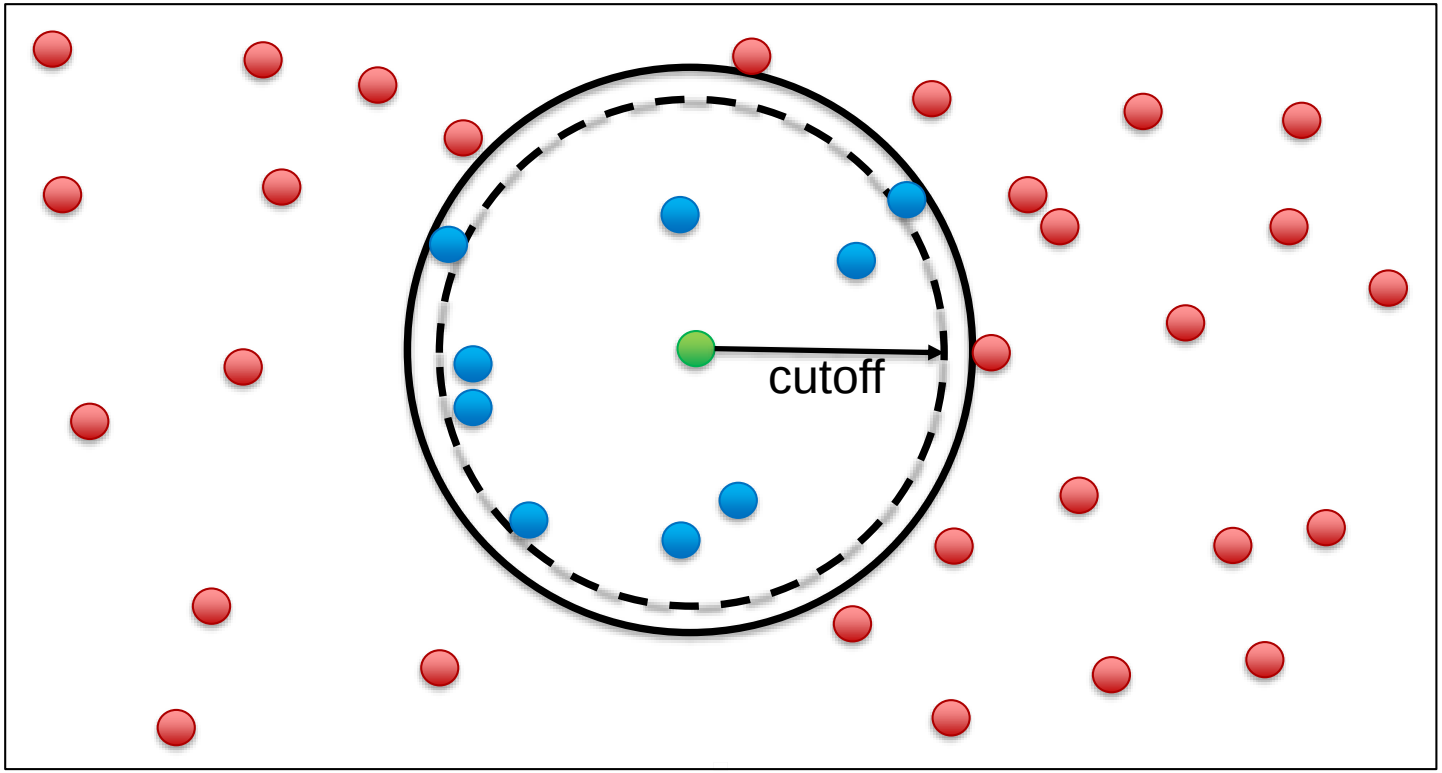

Half Neighbor List

For pair-wise interactions, the force on one atom is equal in magnitude, but opposite in sign, as on the other atom

Can use this principle to reduce computation and the size of the neighbor list (at the cost of more inter-processor communication)

Each pair is stored only once

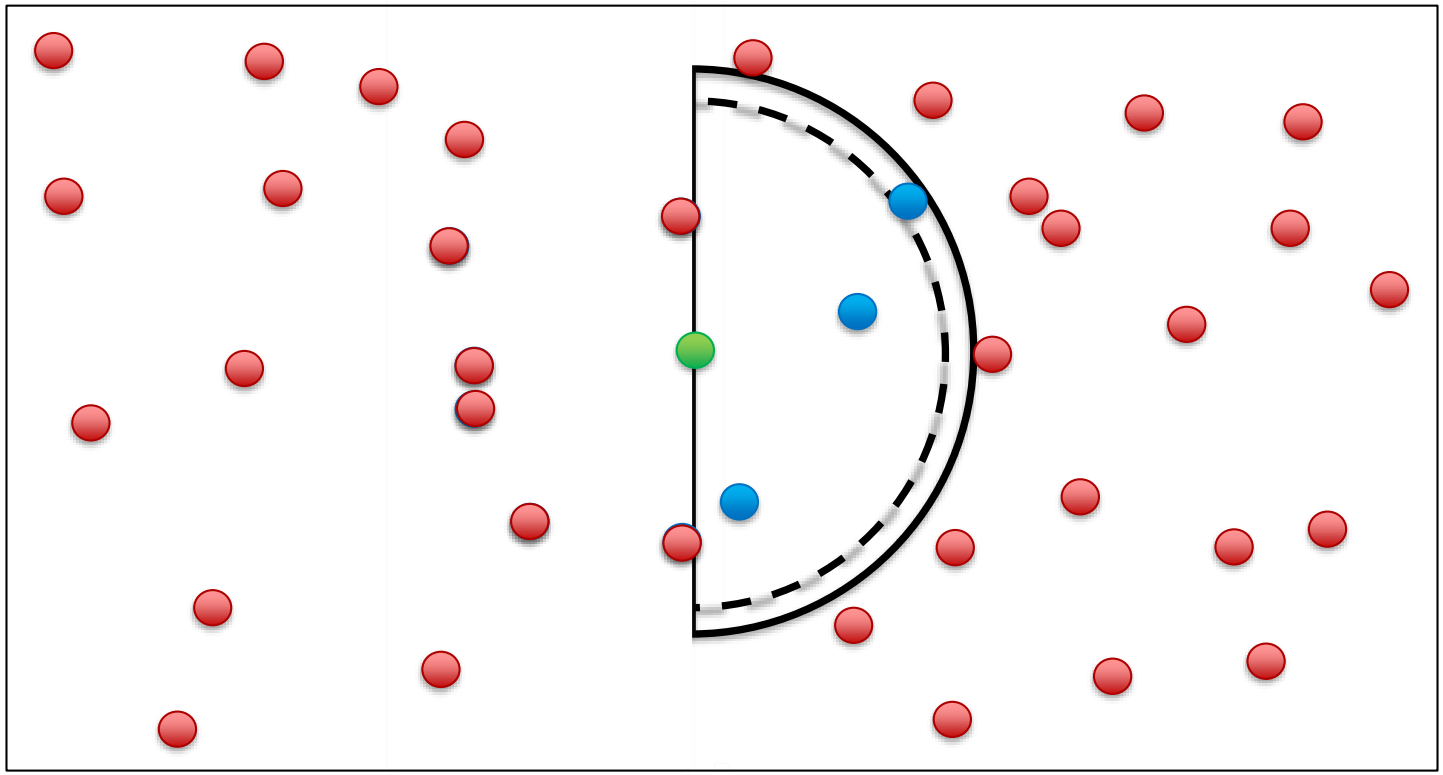

Full Neighbor List

Each pair stored twice, doubles computation but reduces inter-processor communication (can be faster on GPUs)

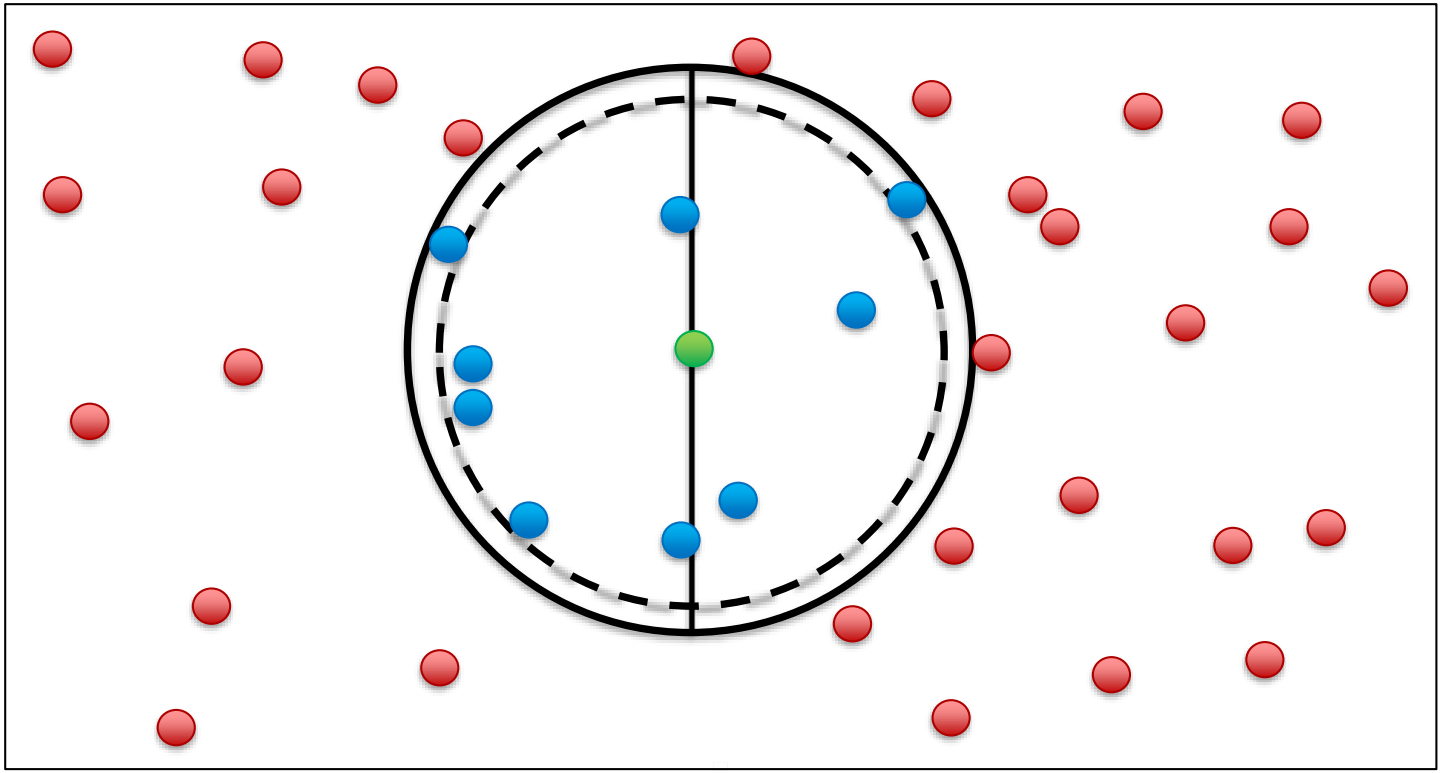

Basic MD Timestep

During each timestep (without neighborlist build):

1. initial integrate
2. compute forces (pair, bonds, etc.)
3. final integrate
4. output (if requested on this timestep)

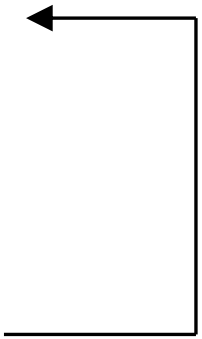

*Computation of diagnostics (i.e. thermodynamic properties) can be scattered throughout the timestep

MPI Parallelization Approach

Domain decomposition: each processor owns a portion of the simulation domain and atoms therein

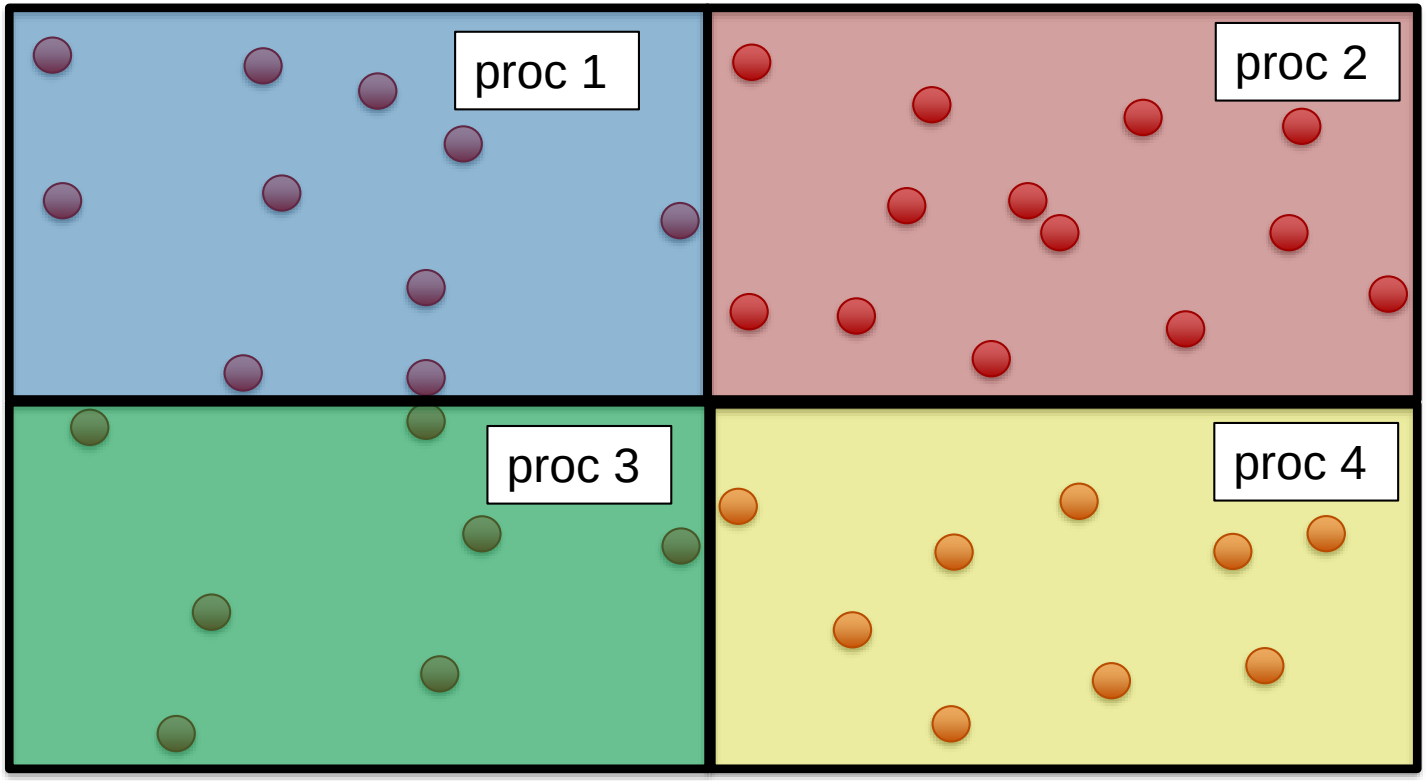

Ghost Atoms

The processor domain is also extended include needed ghost atoms (copies of atoms located on other processors)

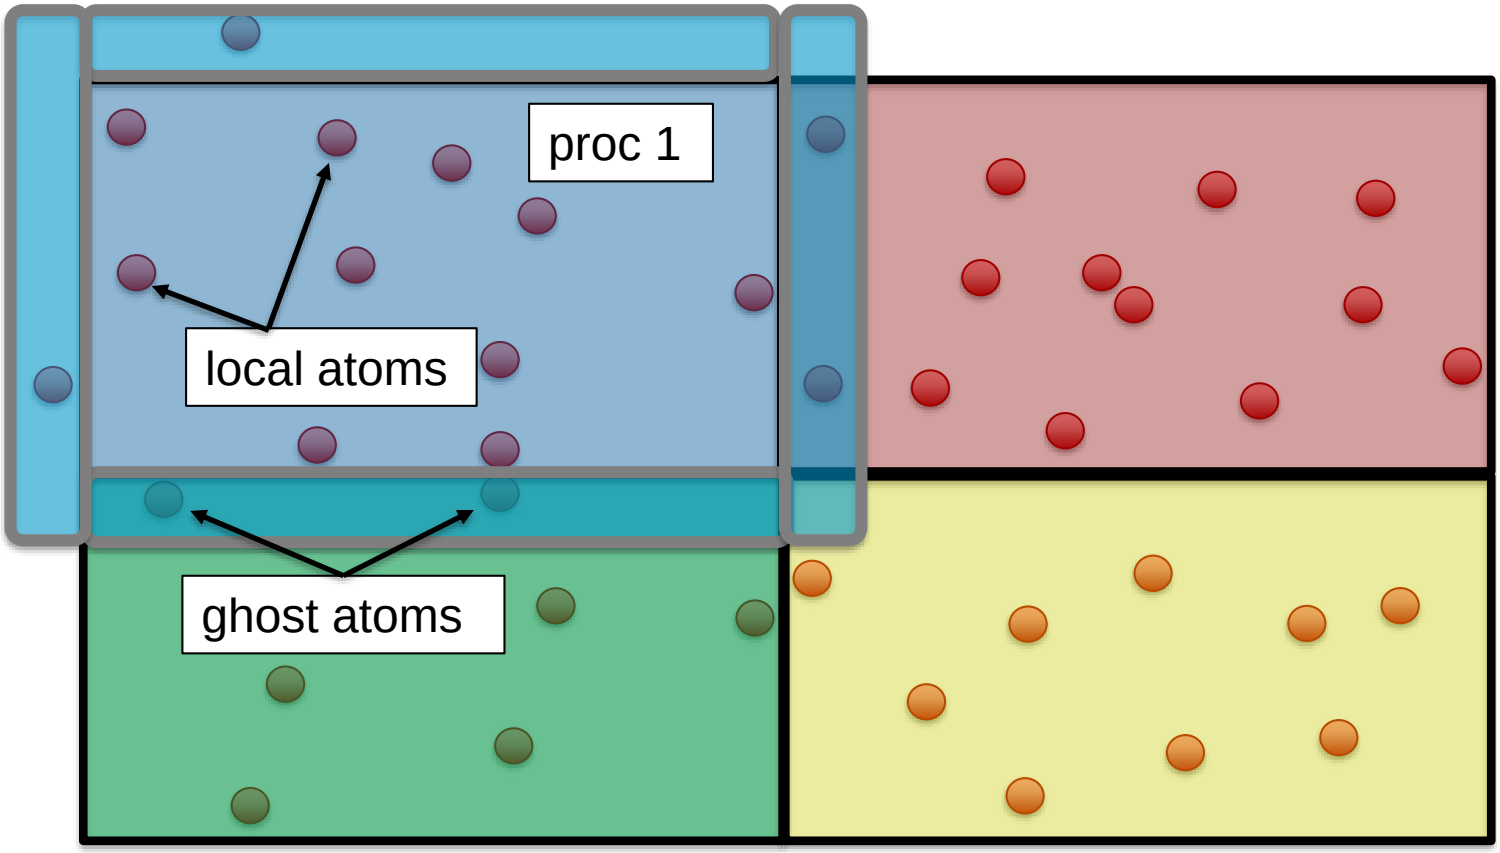

Communication Patterns

Forward communication updates ghost atoms properties (such as positions) from the corresponding real atoms on a different processor

Reverse communication takes properties (such as forces) accumulated on ghost atoms and updates them on the corresponding real atoms on a different processor

Exchange communication migrates real atoms from one processor to another

Border communication creates new ghost atoms

LAMMPS tries to minimize the number of MPI calls required between subdomains

Basic MD Timestep with MPI comm

During each timestep (without neighborlist build):

1. initial integrate
2. MPI forward communication
3. compute forces (pair, bonds, etc.)
4. MPI reverse communication (if newton flag on)
5. final integrate
6. output (if requested on this timestep)

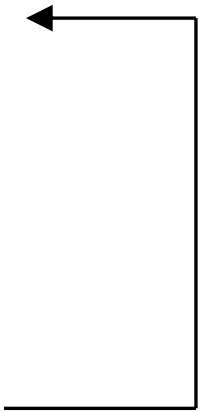

*Computation of diagnostics (fixes or computes) can be scattered throughout the timestep

Strong vs Weak Scaling

Strong scaling: hold system size fixed while increasing processor count (# of atoms/processor decreases)

Weak scaling: increase system size in proportion to increasing processor count (# of atoms/processor remains constant)

For perfect strong scaling, doubling the processor count cuts the simulation time in half

For perfect weak scaling, the simulation time stays exactly the same when doubling the processor count

Harder to maintain parallel efficiency with strong scaling because the compute time decreases relative to the communication time

Molecular Topology

Bonds: constrained length between two atoms

Angles: constrained angle between three atoms

Dihedrals: interactions between quadruplets of atoms

Impropers: “improper” interactions between quadruplets of atoms

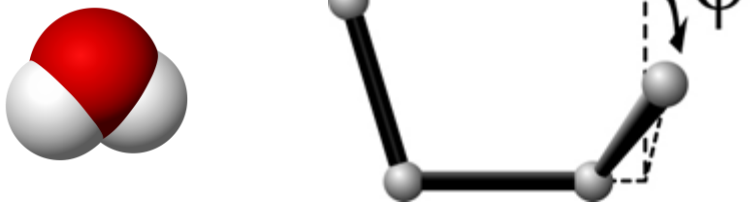

Long-Range Electrostatics

Truncation doesn't work well for charged systems due to long-ranged nature of Coulombic interactions

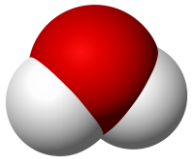

Use Kspace style to add long-range electrostatics:

- PPPM—usually fastest, uses FFTs
- Ewald—potentially most accurate, but slow for large systems
- MSM—multigrid method that also works for non-periodic systems

Usually specify a relative accuracy (1e-4 or 1e-5 typically used)

2D Slab Geometry with Kspace

The *slab* keyword allows a Kspace solver to be used for a systems that are periodic in x,y but non-periodic in z

Must use a boundary setting of "`boundary p p f`"

Actually treats the system as if it were periodic in z, but inserts empty volume between atom slabs and removing dipole inter-slab interactions so that slab-slab interactions are effectively turned off

May need to use reflecting walls in the z-dimension

Run vs Minimize

"run" command updates velocities and positions based on forces. System may blow up and crash LAMMPS if atoms overlap!

"minimize" command minimizes energy of the system by iteratively adjusting atom coordinates

- Good to minimize first if you built your system using an external tool
- Prevents LAMMPS crashing from overlapping atoms

Can use "run" after "minimize"

Load Balancing

Adjusts the size and shape of processor sub-domains within the simulation box

Attempts to balance the number of atoms or particles and thus indirectly the computational cost (load) more evenly across processors

Can be static or dynamic

shift

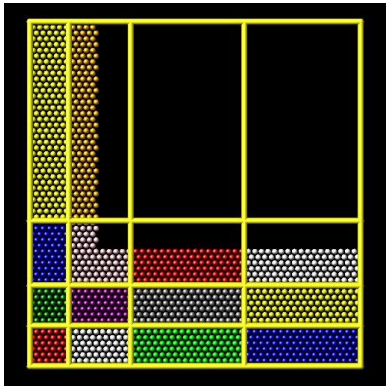

RCB

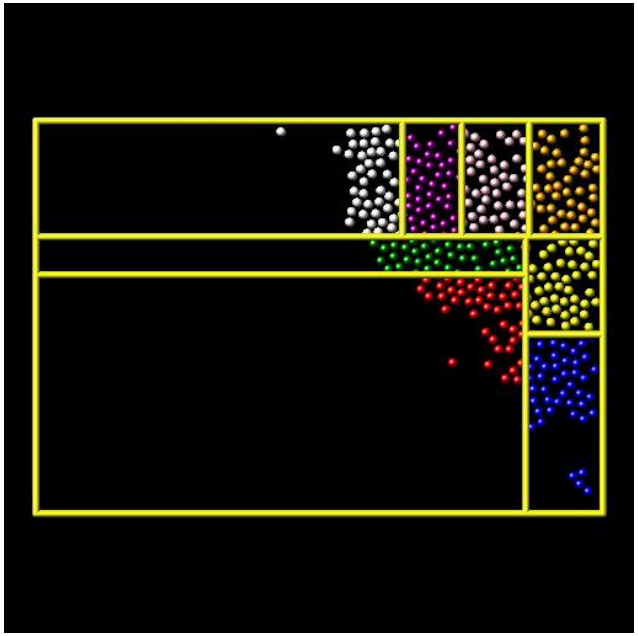

Example: Energetic Material

APS SCCM Meeting 2017 – Wood, Kittell, Yarrington and Thompson

Shock to Initiation, Deflagration

- Detailed chemistry is incorporated in these MD potentials, hot spot evolution is captured naturally.
- **Current capabilities for ReaxFF within LAMMPS is ~100M atoms running routinely on ~100k processors**
- KOKKOS-Reax/c package circumvents memory overflow errors and makes the code portable to modern architectures (GPU, KNL, HSW)

Temperature: Time = 0 ps

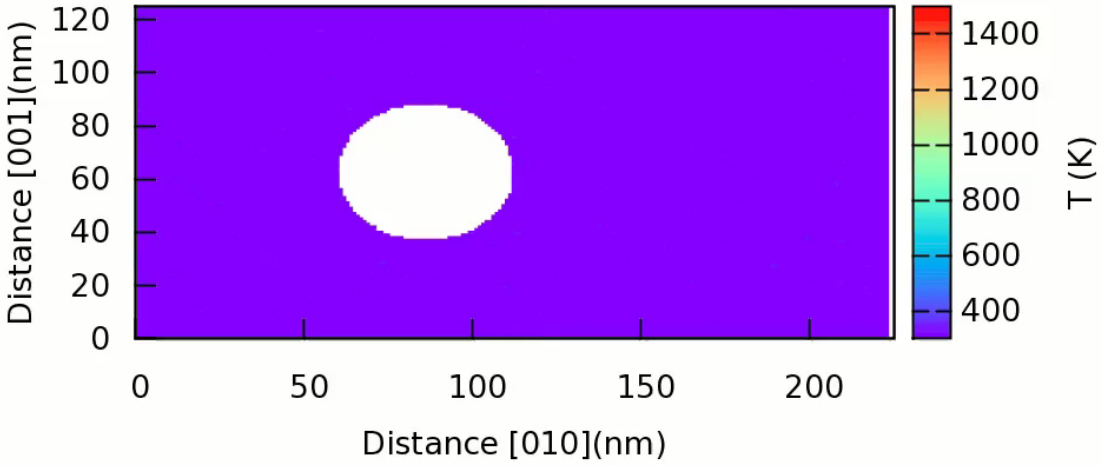

Strain Rate: Time = .100 ps

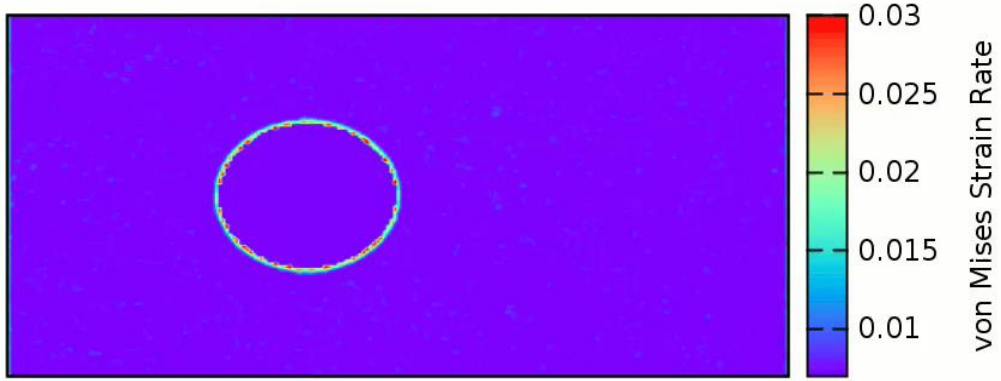

Example: Organic Nanowire

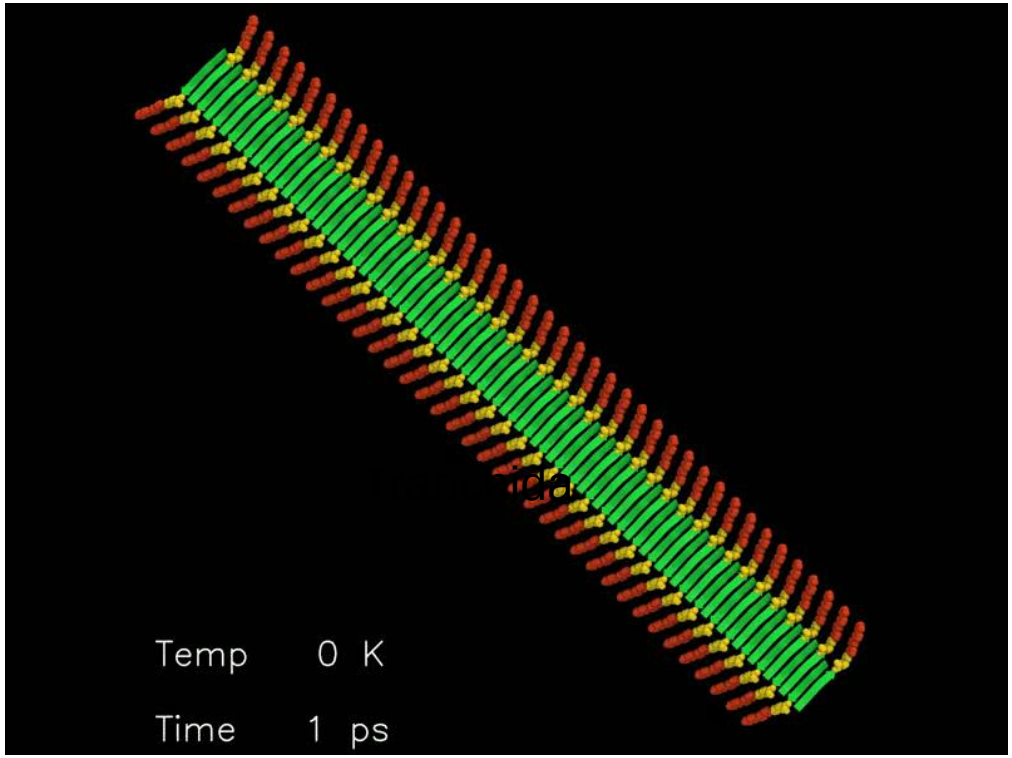

This is work by Alexey Shaytan et al. at the Dept of Energy-Related Nanomaterials (University of Ulm, Germany) on a large-scale fully atomistic MD simulation of the amyloid-like nanofibers formed by the conjugates of oligothiophenes and oligopeptides. Such compounds are very promising for applications in organic electronics (conductive organic nanowires).

Viscosity for rigid-bodies in SRD fluid

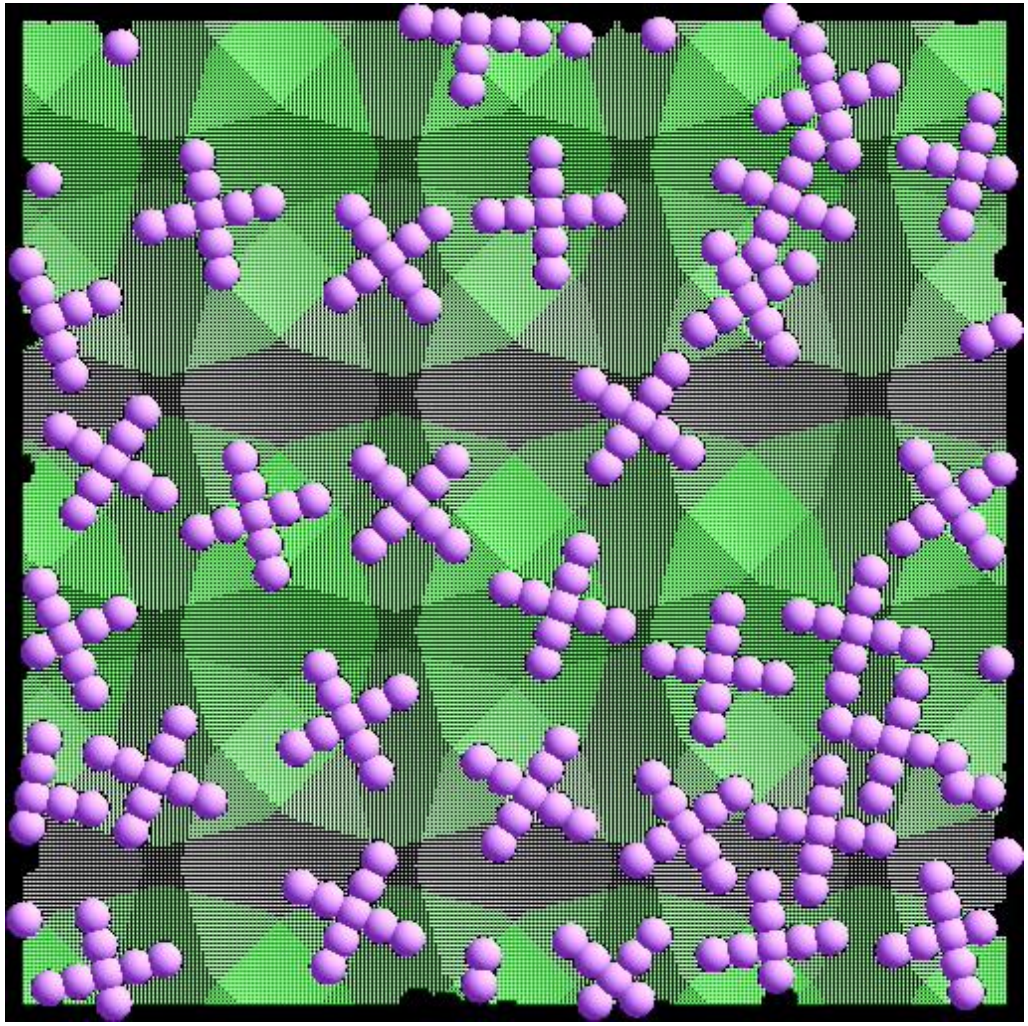

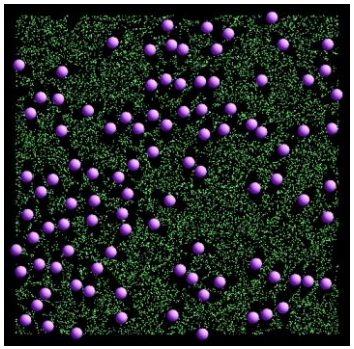

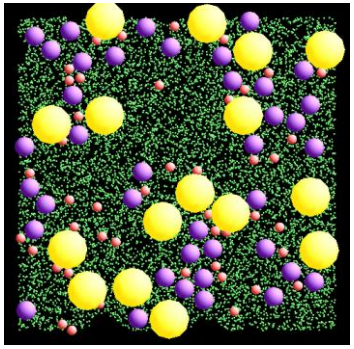

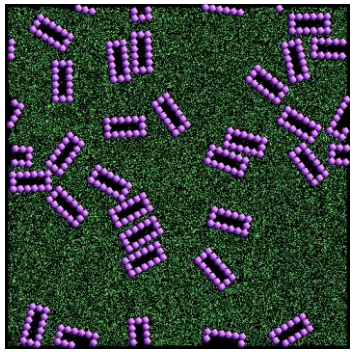

Triangle and line particle examples

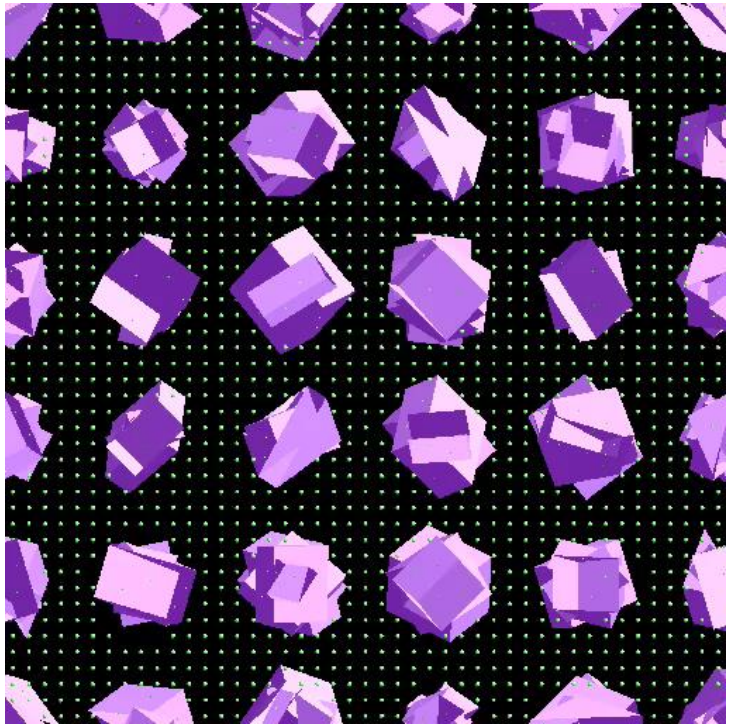

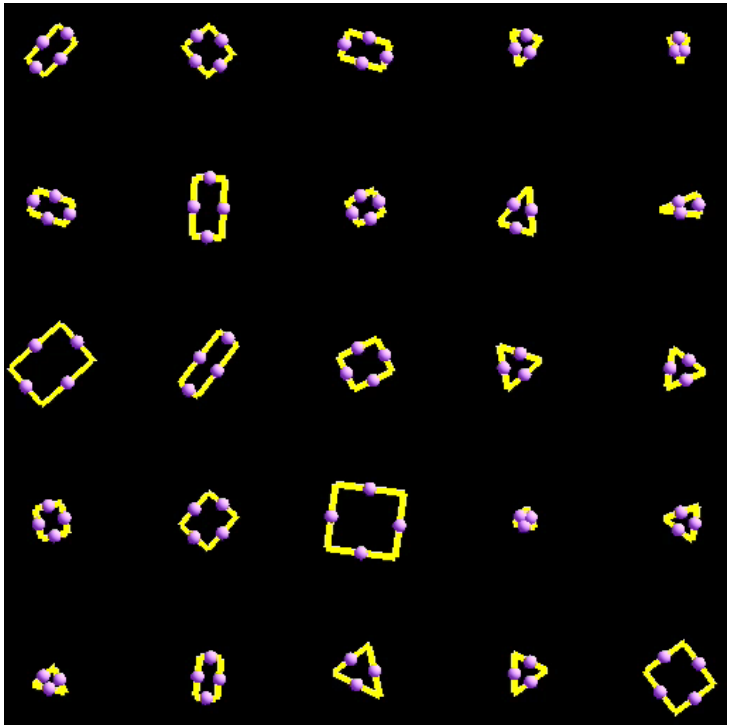

Peridynamics

PERI package, Mike Parks and Stewart Silling (Sandia)

Particle-based meshless non-local continuum model

Hi-deformation impact & fracture

Constitutive models encoded in pairwise interactions & bonding

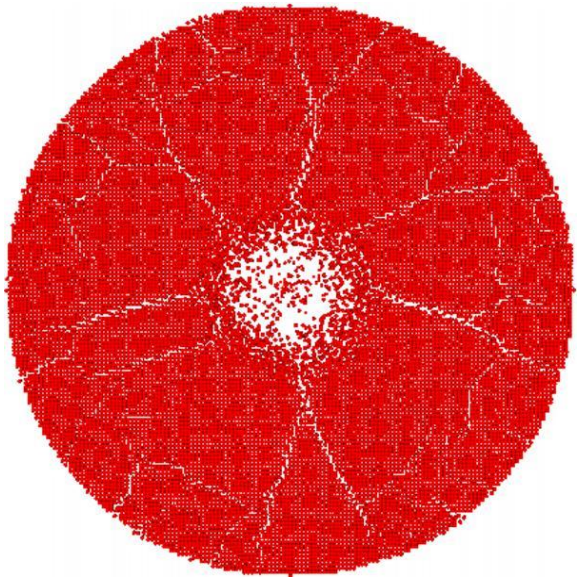

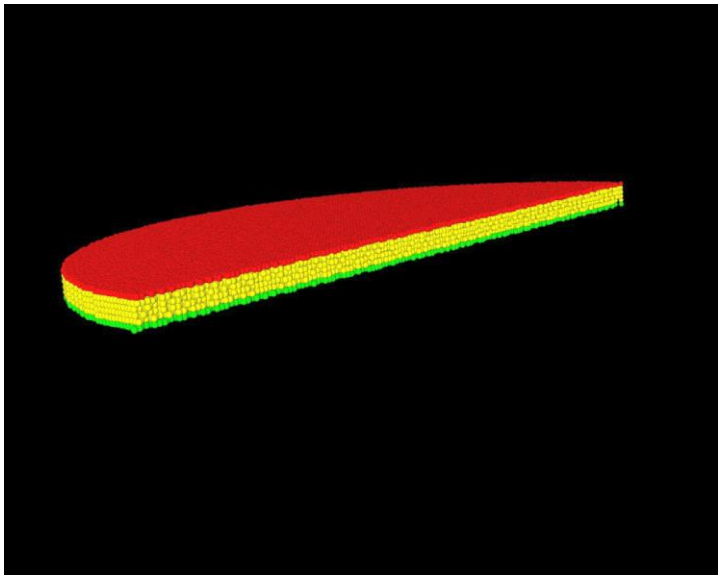

Smoothed particle hydrodynamics

USER-SPH package

Georg Ganzenmüller (Franhofer-Institute, EMI, Germany)

collapse of a water column

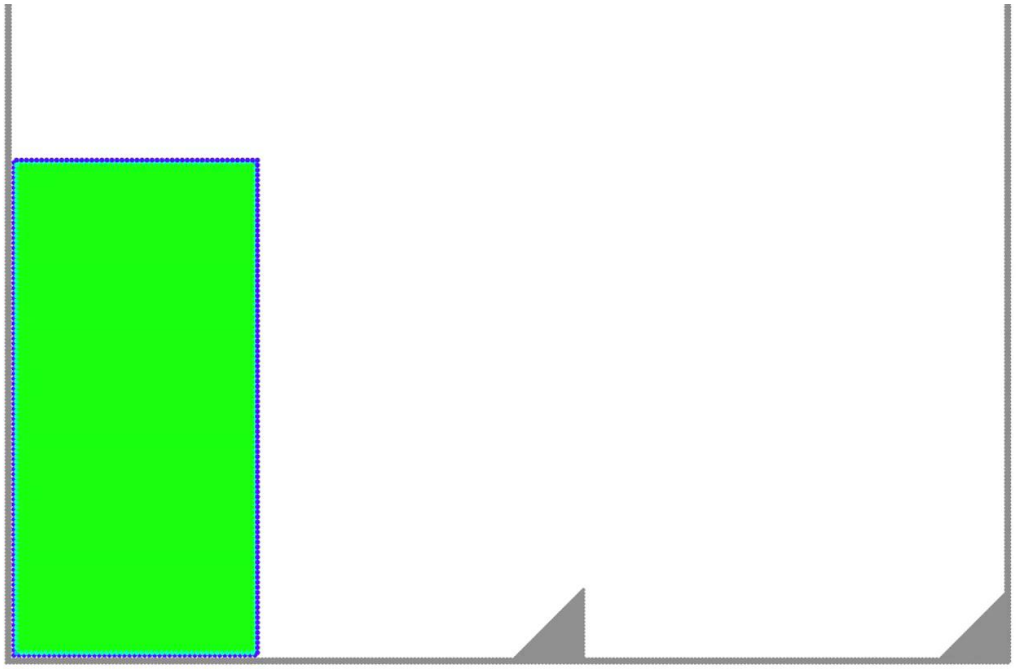

Granular modeling

GRANULAR package

Christoph Kloss group (JKU) created add-on LIGGGHTS code, (Christoph Kloss now at DCS Computing)

www.cfdem.com

particles + CAD mesh

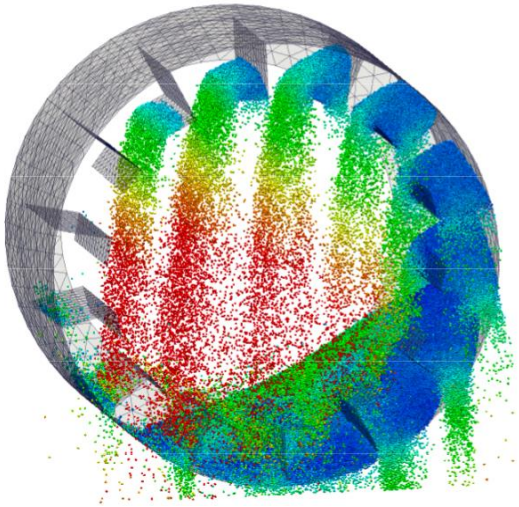

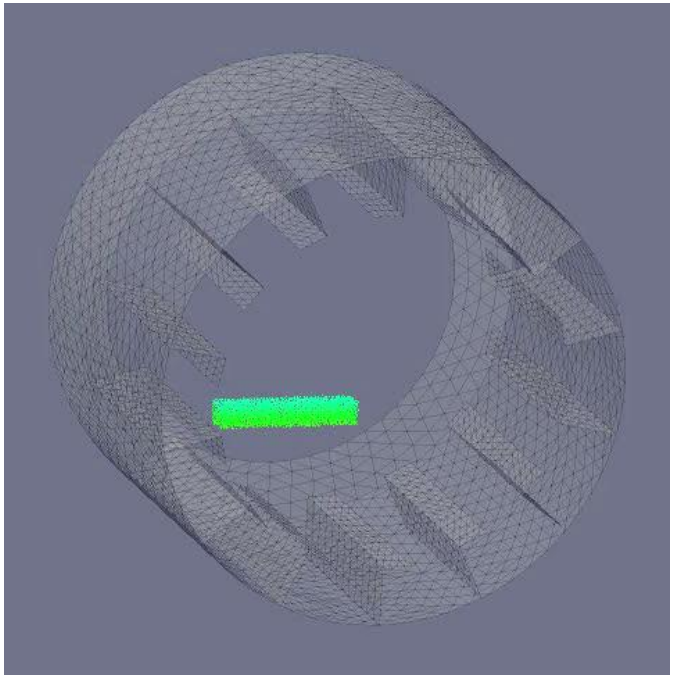

A bit beyond the mesoscale ...

LIGGGHTS code and

FMI (Functional Mock-up Interface) for mesh dynamics

Wheelloader model by C Schubert & T Dresden, (Dresden Tech U)

Simulation by C. Richter & A. Katterfeld, (U Magdeburg OV Guericke)