High-Performance Computing for Atomistic Materials Modeling

Steve Plimpton
Sandia National Labs

HPC Symposium 2012
Lehigh University - March 2012

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.

High-Performance Computing for Atomistic Materials Modeling

Steve Plimpton
Sandia National Labs

HPC Symposium 2012
Lehigh University - March 2012

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

Pixar

- Faster computers enable better models
- "Law of Constancy of Pain"
- Adept at discarding low-level details
- Animations "look" right
- Success due to more than modeling, tell a good story
- Ultimate goal: fast & realistic modeling of any material with minimum of effort by user

Pixar

- Faster computers enable better models
- “Law of Constancy of Pain”
- Adept at discarding low-level details
- Animations “look” right
- Success due to more than modeling, tell a good story
- Ultimate goal: fast & realistic modeling of any material with minimum of effort by user

- Bottom line: 12 films, $600M/film, 21 Academy Awards

Pixar vs HPC materials modeling

- Faster computers enable better models
- "Law of Constancy of Pain"
- Adept at discarding low-level details
- Animations "look" right $\Rightarrow$ **Simulations match experiment**
- Success due to more than modeling, tell a good **science** story
- Ultimate goal: fast & realistic modeling of any material with minimum of effort by simulator

Pixar vs HPC materials modeling

- Faster computers enable better models
- “Law of Constancy of Pain”
- Adept at discarding low-level details
- Animations “look” right $\Rightarrow$ **Simulations match experiment**
- Success due to more than modeling, tell a good **science** story
- Ultimate goal: fast & realistic modeling of any material with minimum of effort by simulator

- Bottom line: $600M/**grant**, 21 Academy Awards **(sorry)**

Machines are getting faster (and harder to use)

1 Petaflop in 2008, now at **10 Petaflops**

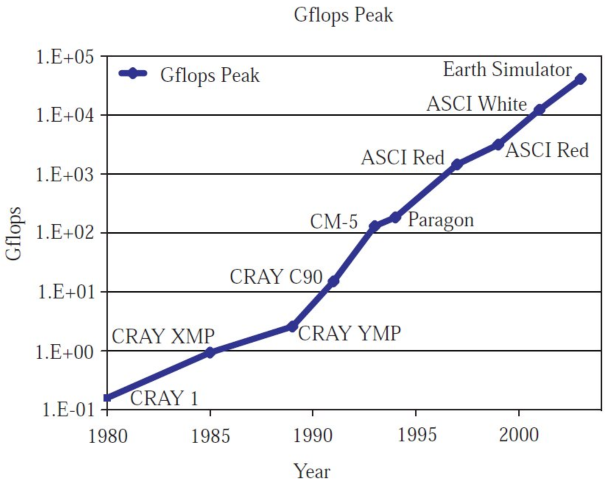

PDE matrix solvers are getting more clever

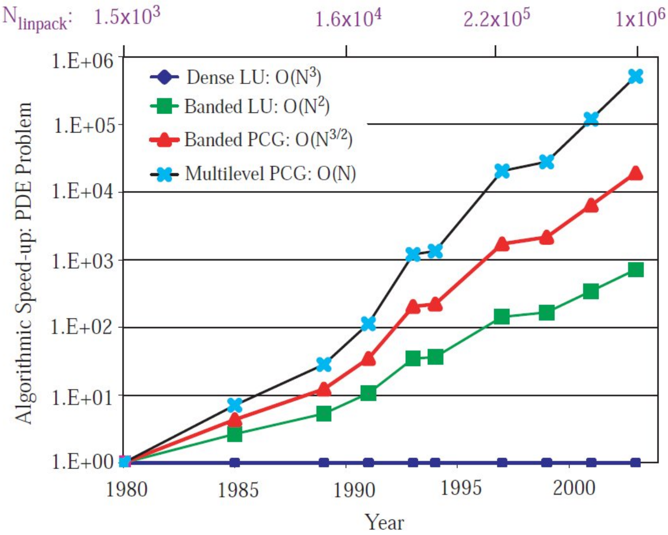

Empirical potentials are getting more clever and costly!

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

Moore's Law for potentials

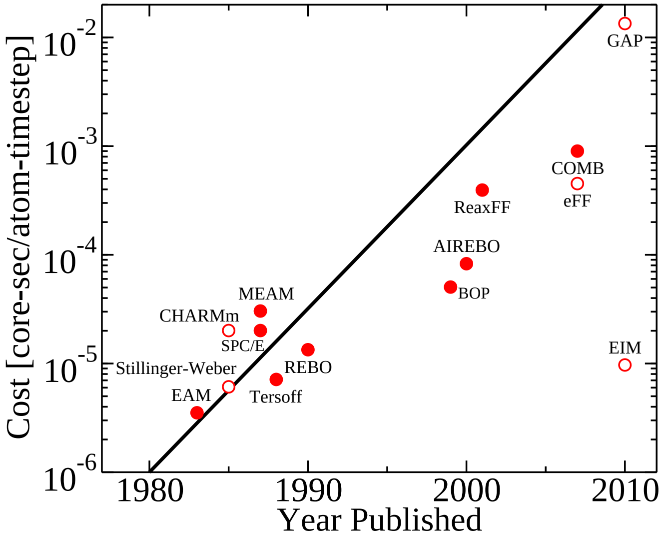

Research directions in classical MD, enabled by HPC

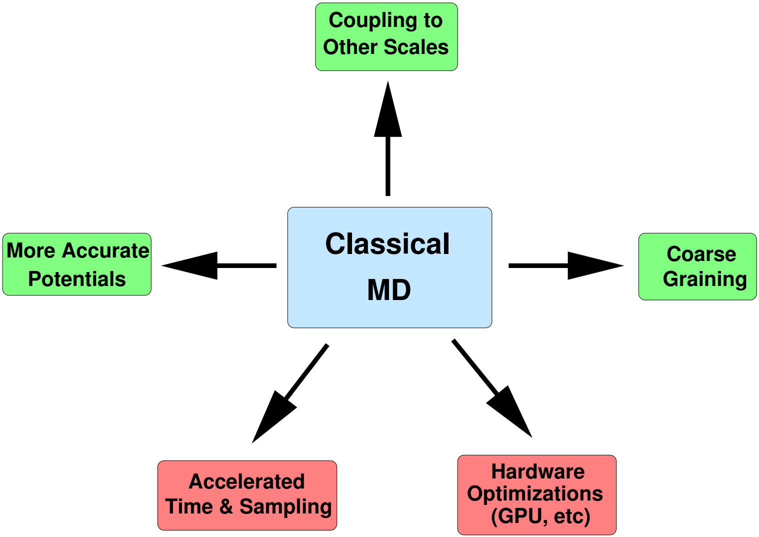

Particle suspensions

- **Colloids** = large particles, sub-micron to few microns in diameter, squishy or hard, bare or coated, spherical or irregular in shape
- **Nanoparticles** = small colloids, 2-100 nm in diameter
- **Suspension** = background fluid

Particle suspensions

- **Colloids** = large particles, sub-micron to few microns in diameter, squishy or hard, bare or coated, spherical or irregular in shape
- **Nanoparticles** = small colloids, 2-100 nm in diameter
- **Suspension** = background fluid

- Applications:
  - bulk: paints, gels, pastes
  - films: self-assembly of nano-structured materials
- Processing steps:
  - extrusion, flow, drying, self-assembly
- Measure/predict:
  - **diffusion**, **rheology**, aggregation/dispersion

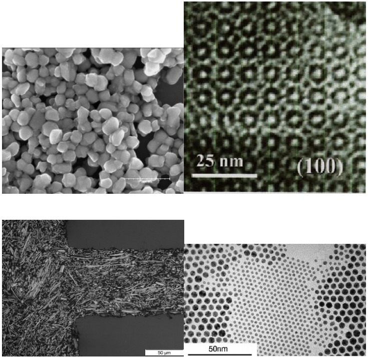

What we want to model

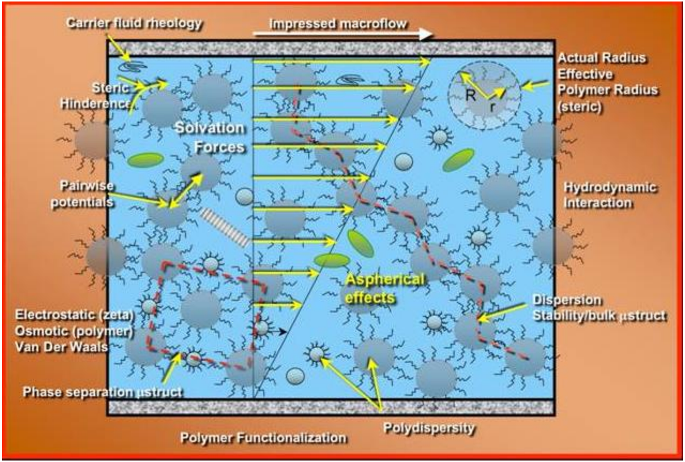

Spherical vs aspherical, bare vs coated, polydisperse, agglomeration, response to shear, ...

Length and time scales

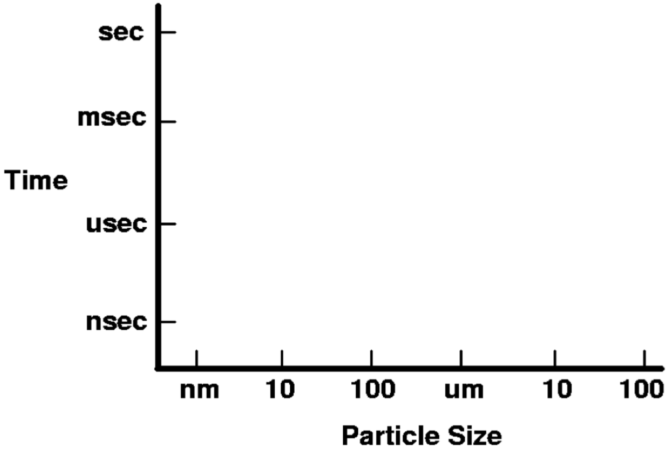

Length and time scales

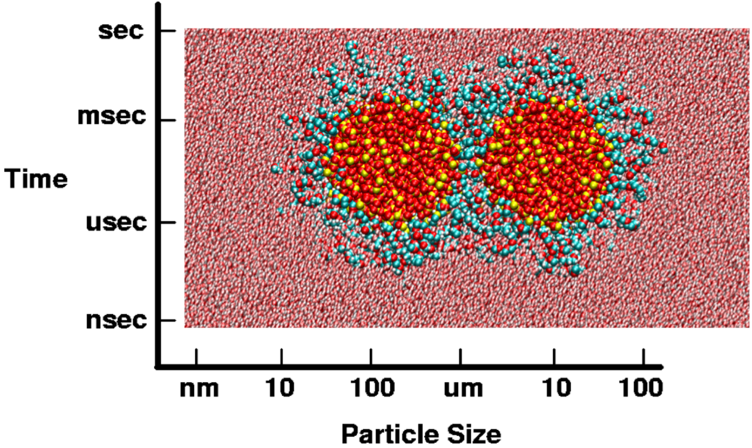

Length and time scales

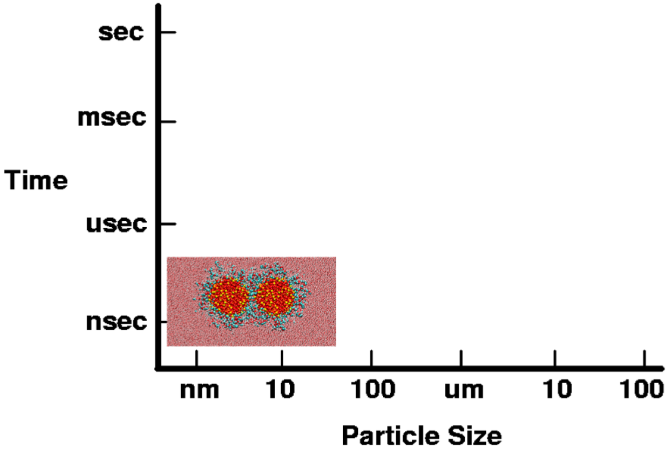

Length and time scales

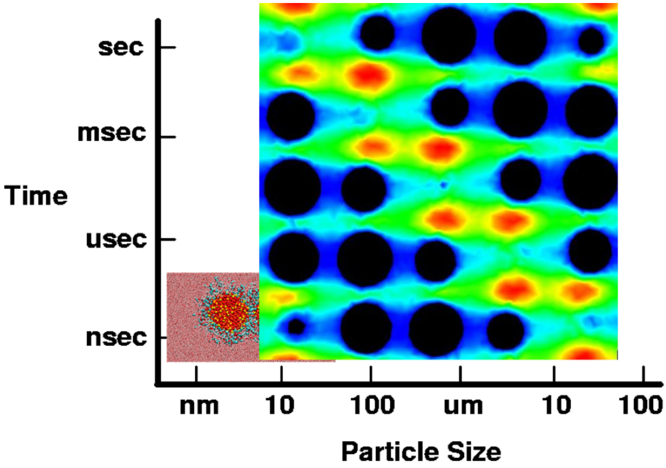

Length and time scales

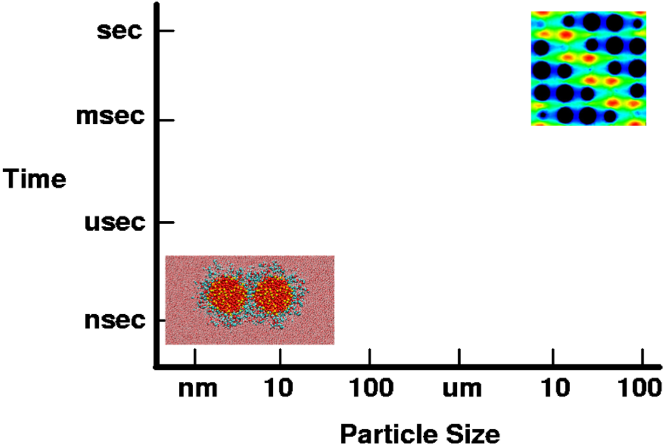

Length and time scales

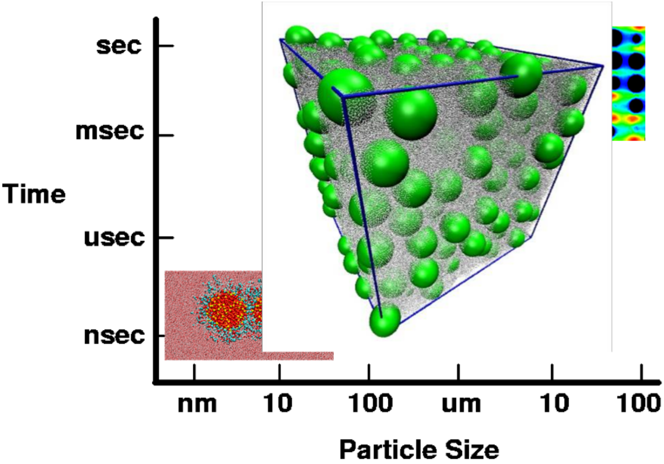

Length and time scales

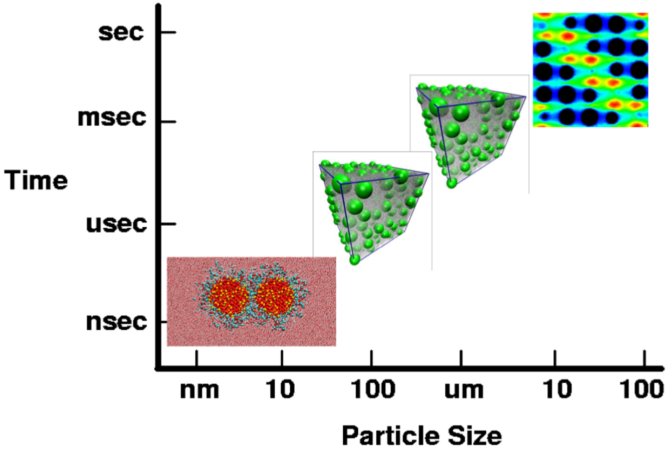

Coarse-grained nanoparticle interactions

- Integrated LJ potential over volume of nanoparticle
  Everaers (PRE 2003)

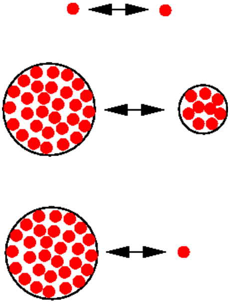

$$
\begin{aligned}
U_A &= -\frac{A_{cc}}{6}\left[\frac{2a_1a_2}{r^2-(a_1+a_2)^2}+\frac{2a_1a_2}{r^2-(a_1-a_2)^2}+\ln\left(\frac{\mathrm{r}^2-(\mathrm{a}_1+\mathrm{a}_2)^2}{\mathrm{r}^2-(\mathrm{a}_1-\mathrm{a}_2)^2}\right)\right]\\
U_R &= \frac{A_{cc}}{37800}\frac{\sigma^6}{r}\Bigg[\\
&\quad \frac{r^2-7r(a_1+a_2)+6(a_1^2+7a_1a_2+a_2^2)}{(r-a_1-a_2)^7}\\
&\quad +\frac{r^2+7r(a_1+a_2)+6(a_1^2+7a_1a_2+a_2^2)}{(r+a_1+a_2)^7}\\
&\quad -\frac{r^2+7r(a_1-a_2)+6(a_1^2-7a_1a_2+a_2^2)}{(r+a_1-a_2)^7}\\
&\quad -\frac{r^2-7r(a_1-a_2)+6(a_1^2-7a_1a_2+a_2^2)}{(r-a_1+a_2)^7}\Bigg]\\
U &= U_A+U_R, \qquad r<r_c
\end{aligned}
$$

Coarse-grained nanoparticle interactions

- Integrated LJ potential over volume of nanoparticle
  Everaers (PRE 2003)

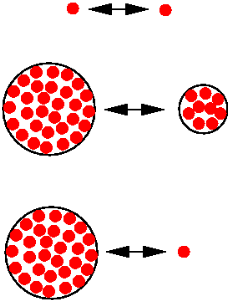

$$
\begin{aligned}
U_A &= -\frac{A_{cc}}{6}\left[\frac{2a_1a_2}{r^2-(a_1+a_2)^2}+\frac{2a_1a_2}{r^2-(a_1-a_2)^2}+\ln\left(\frac{r^2-(a_1+a_2)^2}{r^2-(a_1-a_2)^2}\right)\right] \\
U_R &= \frac{A_{cc}}{37800}\frac{\sigma^6}{r}\Bigg[\\
&\qquad \frac{r^2-7r(a_1+a_2)+6(a_1^2+7a_1a_2+a_2^2)}{(r-a_1-a_2)^7} \\
&\qquad +\frac{r^2+7r(a_1+a_2)+6(a_1^2+7a_1a_2+a_2^2)}{(r+a_1+a_2)^7} \\
&\qquad -\frac{r^2+7r(a_1-a_2)+6(a_1^2-7a_1a_2+a_2^2)}{(r+a_1-a_2)^7} \\
&\qquad -\frac{r^2-7r(a_1-a_2)+6(a_1^2-7a_1a_2+a_2^2)}{(r-a_1+a_2)^7}\Bigg] \\
U &= U_A+U_R, \qquad r < r_c
\end{aligned}
$$

- Pairwise and analytic (albeit expensive), **Hamaker constant**,
  similar formulation for ellipsoids

Colloid/colloid forces are now long-range

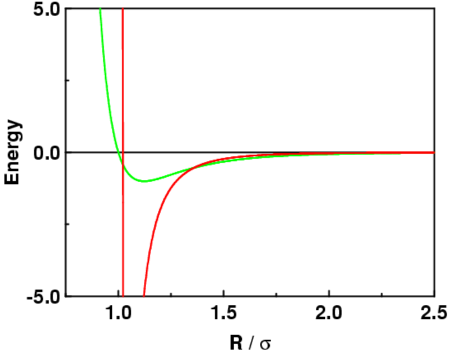

- Distance is plotted in particle diameters
- For colloid/solvent size ratio of 20:1, LJ cutoff of **2.5 sigma** translates to: colloid/colloid cutoff = **50 sigma**, colloid/solvent cutoff = **25 sigma**

Finding neighbors in mixtures

- Goal: efficient Verlet lists
  - SS cutoff at 2.5
  - CS cutoff at 12.5
  - CC cutoff at 25
- Bin size = 1/2 of max cutoff
- Problem: compute R_{ij} for too many pairs
- Performance:
  0.5M particles, 100 S/C
  - 4.7 secs (pair)
  - **242 secs (neighbor)**
  - **36 secs (communication)**
- Exacerbated as size ratio grows

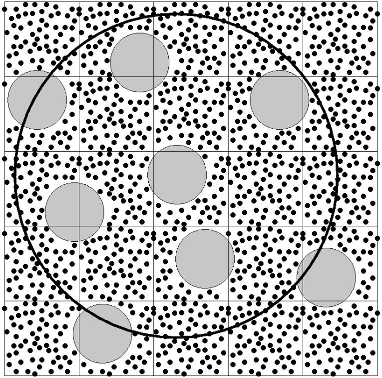

Faster neighbor finding

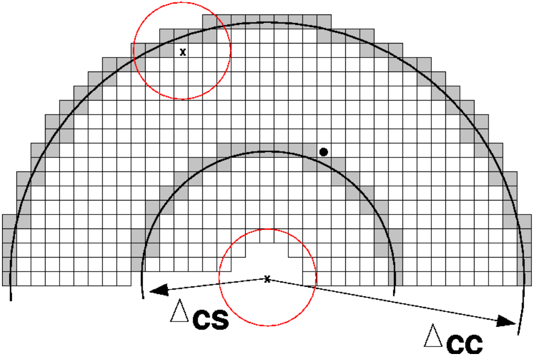

- Bin based on smallest cutoff
- Store 2 values with each stencil bin (closest/furthest corners)
- Only need compute R_{ij} in shaded ring/shell of bins
- Improved performance:

  4.7 secs (pair), **2.4 secs (neigh)**, **2.1 secs (comm)**

Performance improvement

3 size ratios, 3 volume fractions, CPU time for 100 steps

Ratio	P	N	old/old	new/old	new/new
5:1	27	200K	17.8	6.8	5.4
		143K	7.6	4.0	3.1
		87K	3.0	2.0	1.6
10:1	64	672K	471	19.2	11.2
		558K	284	15.9	8.2
		445K	165	11.3	5.6
20:1	125	2.3M	5630	88.6	38.6
		2.0M	4390	73.7	31.2
		1.7M	2750	56.9	22.0

Bottom line: new algorithms as much as **100x faster**

Shear of polymer chains

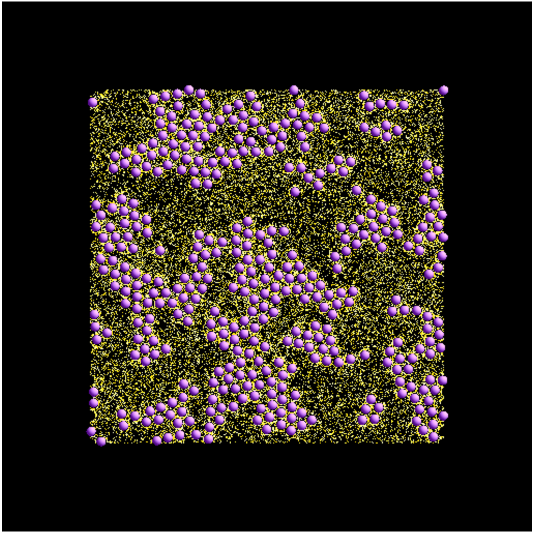

Further coarse graining via SRD particles

- SRD = **stochastic rotation dynamics**
  - *Malevanets (J Chem Phys 99), Padding (PRL 04), Hecht (PRE 05)*
  - intermediate Peclet numbers ~ 1
    Pe = ratio of advection to diffusion
- Basic idea:
  - SRDs are background solvent for colloids
  - two timescales: $t_{colloid}$ and t_{SRD}
  - **small** $t_{colloid}$:
    - colloids move, C/C interactions
    - SRDs advect (no self-interactions)
    - detect colloid-SRD collisions
  - **large** t_{SRD}:
    - bin SRDs, v is randomly rotated

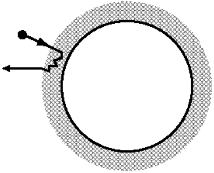

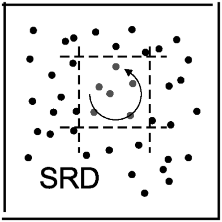

Parallel performance

- Key point: SRDs are **cheap** due to no S/S interactions
- Challenge: Keep them cheap when 1000s per colloid

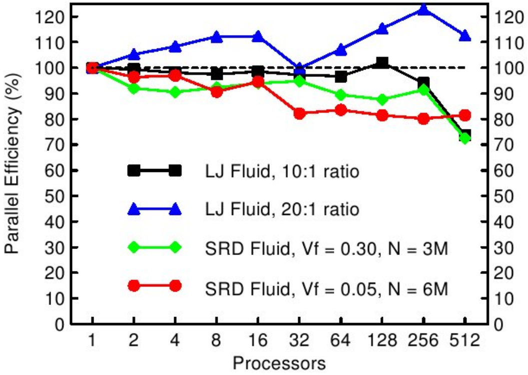

Savings vs explicit LJ solvent:
30x in speed (per particle), **20x** in timestep

Diffusion coefficient versus volume fraction

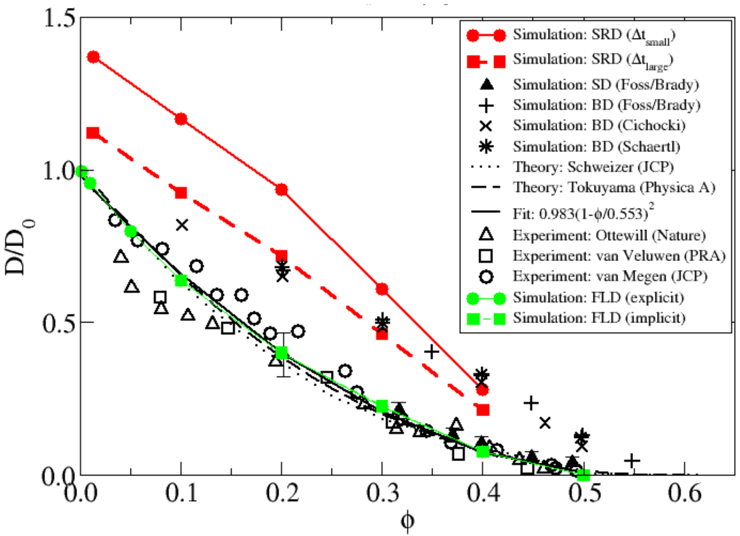

Shear viscosity for varying shear rate and volume fraction

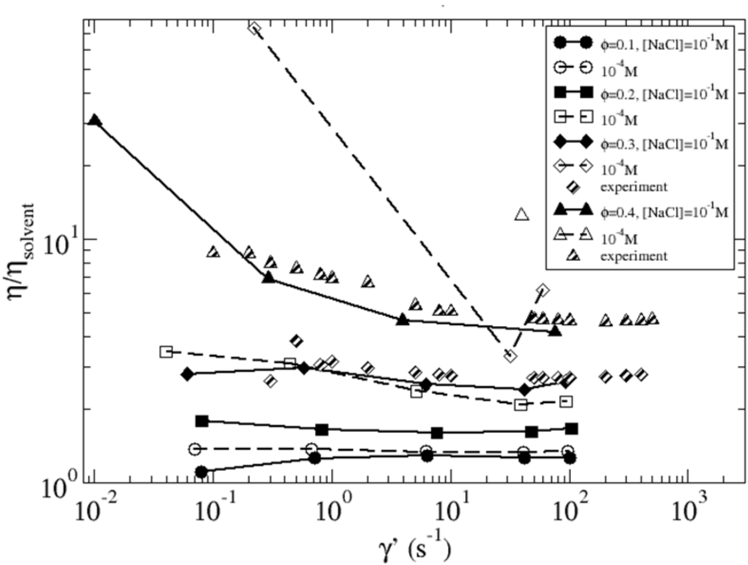

Composite aspherical particles

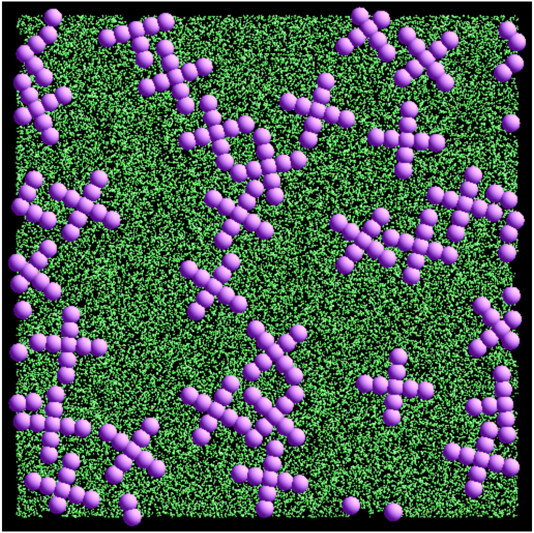

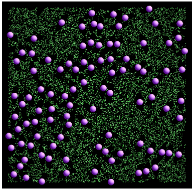

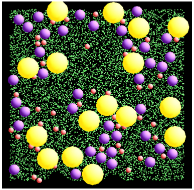

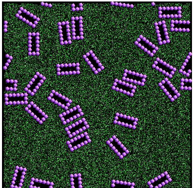

Generalized aspherical particles (2d and 3d)

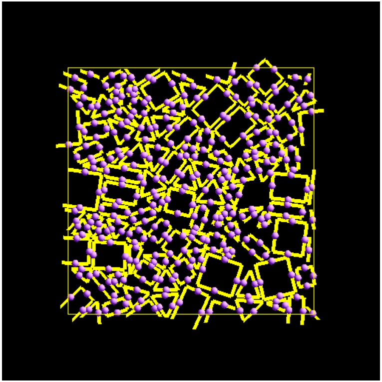

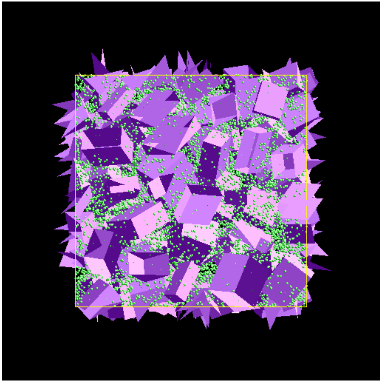

Examples of large-scale materials modeling enabled by HPC

- **Fracture of CNT bundles** (all-atom)
- **Shock compression of low-density polymer foam** (all-atom)
- **Brain blood flow** (multi-scale)
  - coarse-grained DPD particles + Navier-Stokes fluid flow
- **Atom-to-Continuum coupling** (multi-scale)
  - all-atom + finite elements
- **Microstructural grain evolution** (mesoscale)
  - kinetic Monte Carlo, not molecular dynamics

Fracture of CNT fibers

Charles Cornwell (US Army ERDC): *JCP 134, 204708 (2011)*

- Interested in **strong, light-weight materials**
- Fiber = bundle of CNTs of varying length arrayed end-to-end
- **Cross-linking** atoms and bonds added randomly

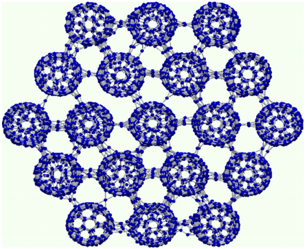

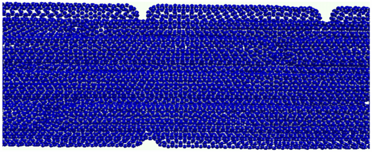

Fracture of CNT fibers

800 nm length fiber bundle, 1.2M atoms, 6 ns simulation

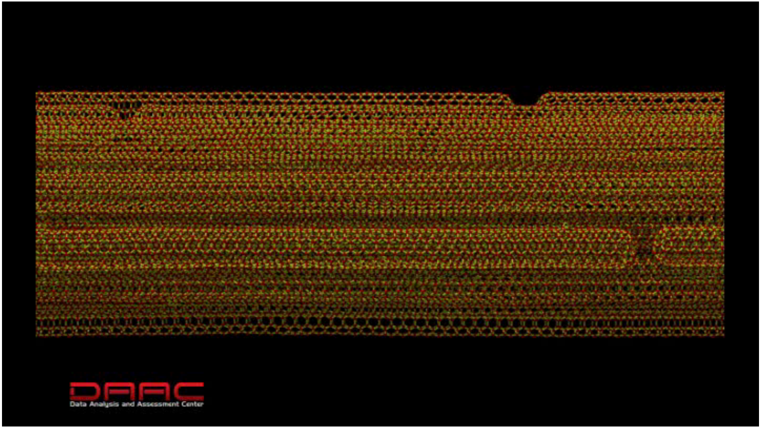

Fracture of CNT fibers

- **AIREBO** bond-order potential, tuned for C and H
- Strain applied incrementally, with relaxation in between
- **Tensile strength** at breakage increased with fiber length and cross-link density
- Maximum strength = **60 GPa** (800 nm fibers, 0.75% density)
  - high-strength steel = 2 GPa
  - Kevlar = 3.5 GPa
  - single CNT = 110 GPa

Shock compression of low-density polymer foam

Matt Lane and Aidan Thompson (Sandia): *APS SCCM, (2011)*

- Dynamic response of **porous materials** is complex
- Foam = 50-mer chains of poly(4-methyl-1-pentene) (PMP)
- Lattice of voids ⇒ low density of 0.3 g/cc
- Piston strike from left at 10-30 km/sec

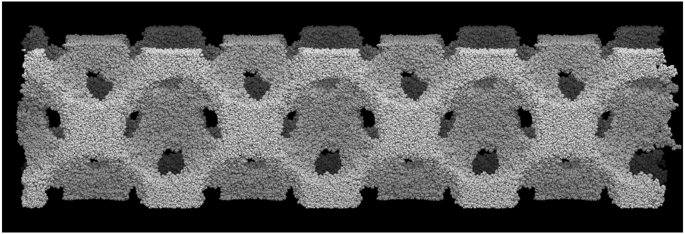

Shock compression of low-density polymer foam

80 nm thickness sample, 1.44M atoms, tens of picoseconds

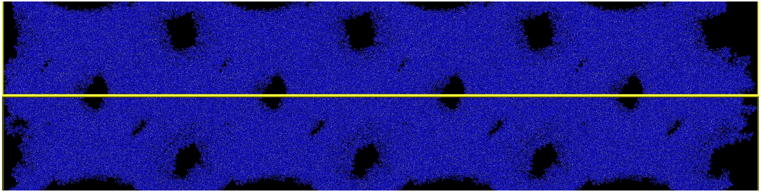

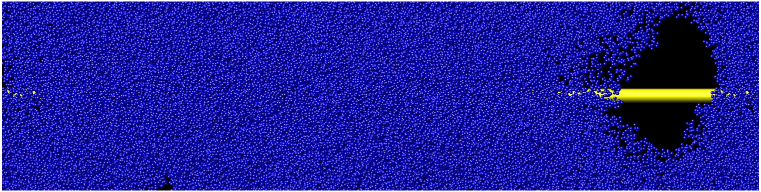

Shock compression of low-density polymer foam

- **ReaxFF** potential to allow bond dissociation
- Requires small 0.025 fm timestep
- Observe **jetting** of polymer fragments into voids ahead of shock front
- In contrast, dense polymers exhibit little dissociation
- Quantitative capture of Hugoniot response (dense), agrees with expt

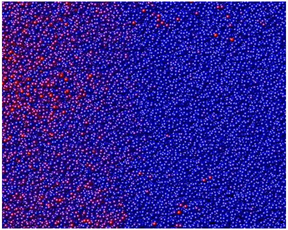

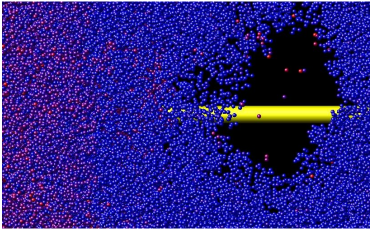

Brain blood flow

Dmitry Fedosov (FZ Juelich) and collaborators:
SC11 Gordon Bell competition finalist paper

- Interested in blood flow, clot formation, **aneurysm rupture**
- Patient-specific cerebrovascalature via imaging
- **DPD particles** to resolve individual blood cells
- Continuum **Navier-Stokes** via NekTor package and spectral element solution

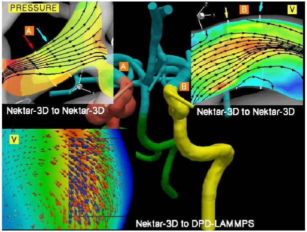

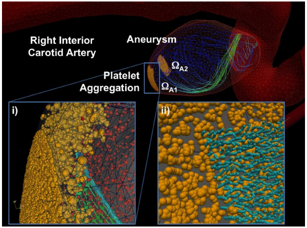

Brain blood flow

up to 823M particles, 2B N-S unknowns, 131,072 cores of BG/P

Atom-to-Continuum coupling

Reese Jones, Jon Zimmerman, Jeremy Templeton, Greg Wagner (Sandia): *CMAME 197, 3351 (2008)*

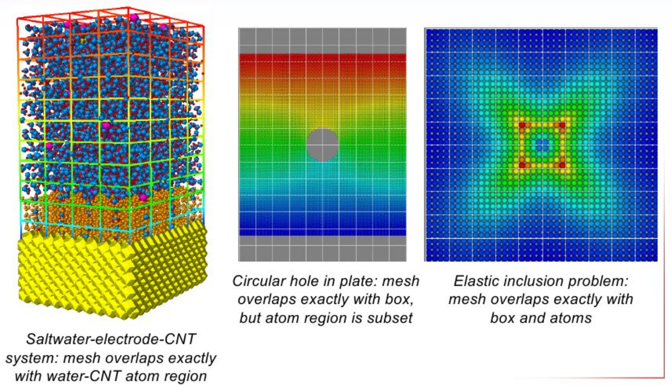

Atoms in parallel, **FE solution** in serial

Thermal coupling with AtC package

2D diffusion problem

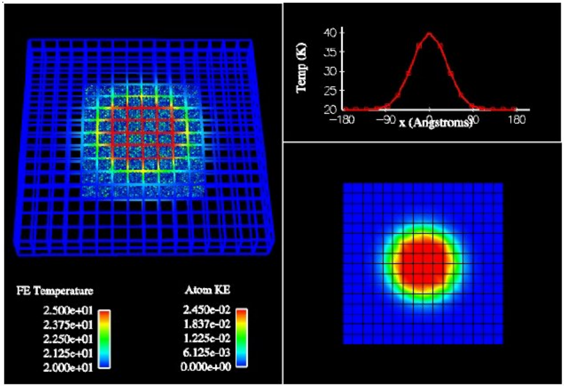

- Plate with embedded MD region (~33,000 atoms)
- Initialized to temperature field with gaussian profile
- Adiabatic boundary conditions at edges

Sandia National Laboratories

Mechanical coupling with AtC package

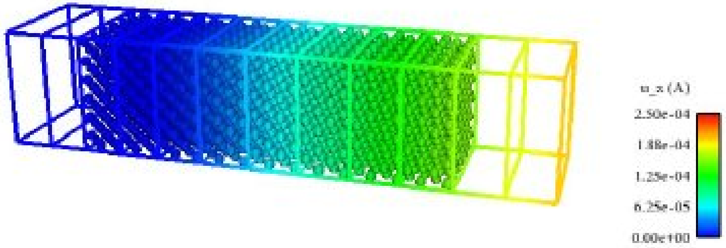

Materials modeling via on-lattice kMC

Work with Aidan Thompson, Ed Webb, Liz Holm, Corbett Battaile (Sandia)

- Variety of models for **time evolution** of processed materials
- Length scale = atomic to mesoscale (fcc or cubic lattices)
- Time scale = **microseconds to days**
- Hamiltonian encodes physics of interest via local interactions

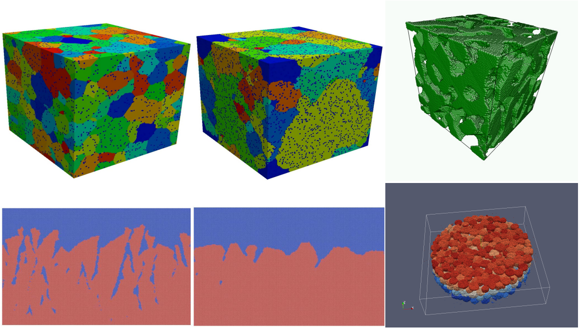

Microstructural grain evolution for normal grain growth

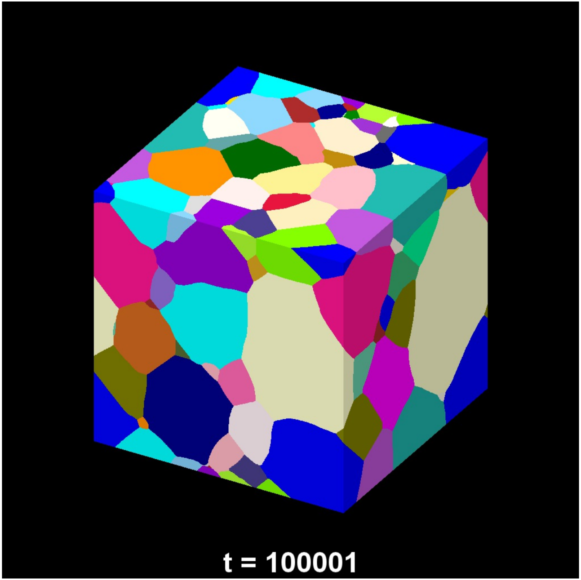

Summary

Exciting time to be involved in materials modeling ...

- Advances in **HPC**:
  - petaflops $\Rightarrow$ exaflops in next few years
  - GPUs, hybrid CPU/GPU nodes
- Advances in **empirical potentials**:
  - pushing towards quantum-level accuracy
- Advances in **numerical and parallel algorithms**:
  - longer length and time scales
  - multi-scale coupling

Summary

Exciting time to be involved in materials modeling ...

- Advances in **HPC**:
  - petaflops $\Rightarrow$ exaflops in next few years
  - GPUs, hybrid CPU/GPU nodes
- Advances in **empirical potentials**:
  - pushing towards quantum-level accuracy
- Advances in **numerical and parallel algorithms**:
  - longer length and time scales
  - multi-scale coupling

Progressing towards Pixar-like **ultimate goal**:
fast & realistic modeling of any material,
with minimum of effort by user

Thanks

- **Funding**:
  DOE and NINE funding
  CRADA with Corning & 3M & BASF

- **Collaborators**:
  Aidan Thompson, Flint Pierce, Jeremy Lechman, Gary Grest, Randy Schunk (Sandia), Pieter in't Veld (BASF)

Thanks

- **Funding**:
  DOE and NINE funding
  CRADA with Corning & 3M & BASF

- **Collaborators**:
  Aidan Thompson, Flint Pierce, Jeremy Lechman, Gary Grest,
  Randy Schunk (Sandia), Pieter in't Veld (BASF)

- LAMMPS MD package: **lammps.sandia.gov**
- Plimpton & Thompson, **many-body potentials**,
  MRS Bulletin, May 2012
- Petersen et al, **SRD method in MD**, J Chem Phys, 2010
- in't Veld et al, **colloid mixture methods**,
  Comp Phys Comm, 2008