A GPU implementation for improved granular simulations with LAMMPS.

Emmanuel N. Millán[1,2,3], Christian Ringl[6], Carlos S. Bederián[5], María Fabiana Piccoli[4], Carlos García Garino[2], Herbert M. Urbassek[6] and Eduardo M. Bringa[1,3]

1. CONICET, Mendoza
2. ITIC, Universidad Nacional de Cuyo
3. Instituto de Ciencias Básicas, Universidad Nacional de Cuyo, Mendoza
4. Universidad Nacional de San Luis, San Luis, Argentina.
5. Instituto de Física Enrique Gaviola, CONICET
6. Fachbereich Physik und Forschungszentrum OPTIMAS, Universitat Kaiserslautern, Erwin-Schrodinger-Straße, D-67663 Kaiserslautern, Germany

http://sites.google.com/site/simafweb/

Funding: PICT2009-0092, SeCTyP 2011-2013

Interstellar dust plays an important role in astrophysical processes

Grain collisions matters for evolution, astro-chemistry, etc.

- Grain diameters: 0.005 μm – 5 μm
- Size distribution: $n(a) \propto a^{-3.5}$

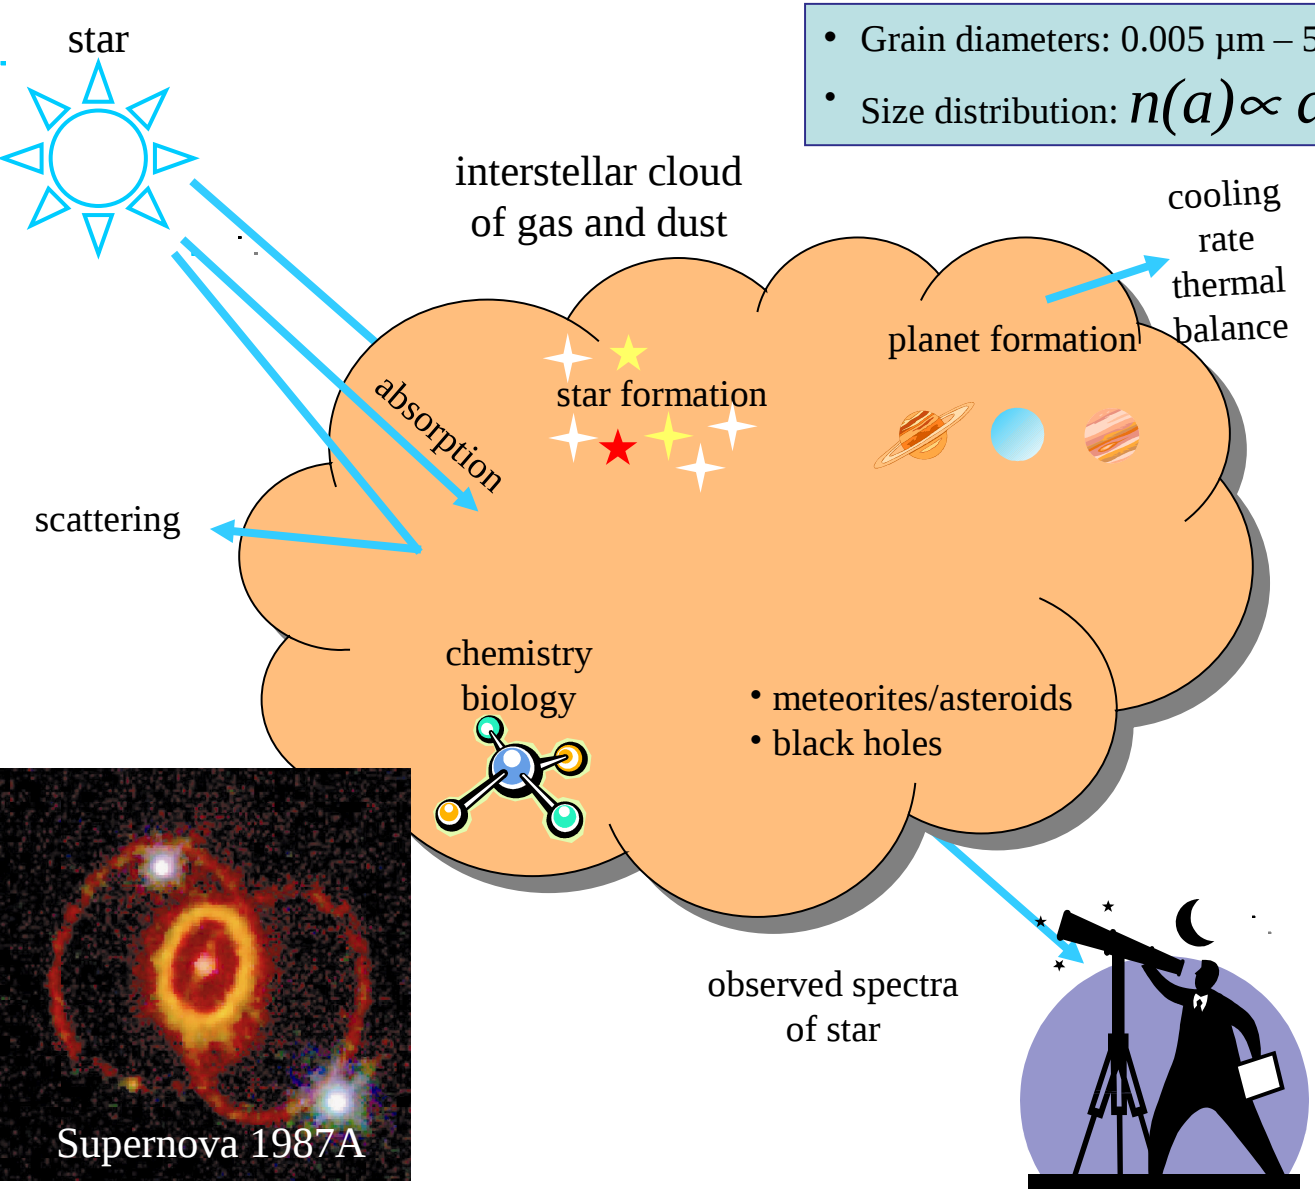

Modification of materials: phase transitions and damage to space hardware, and due to shrapnel (NIF)

Supernova 1987A

Large-scale MD links nano and microscales in damage induced by nanoprojectiles

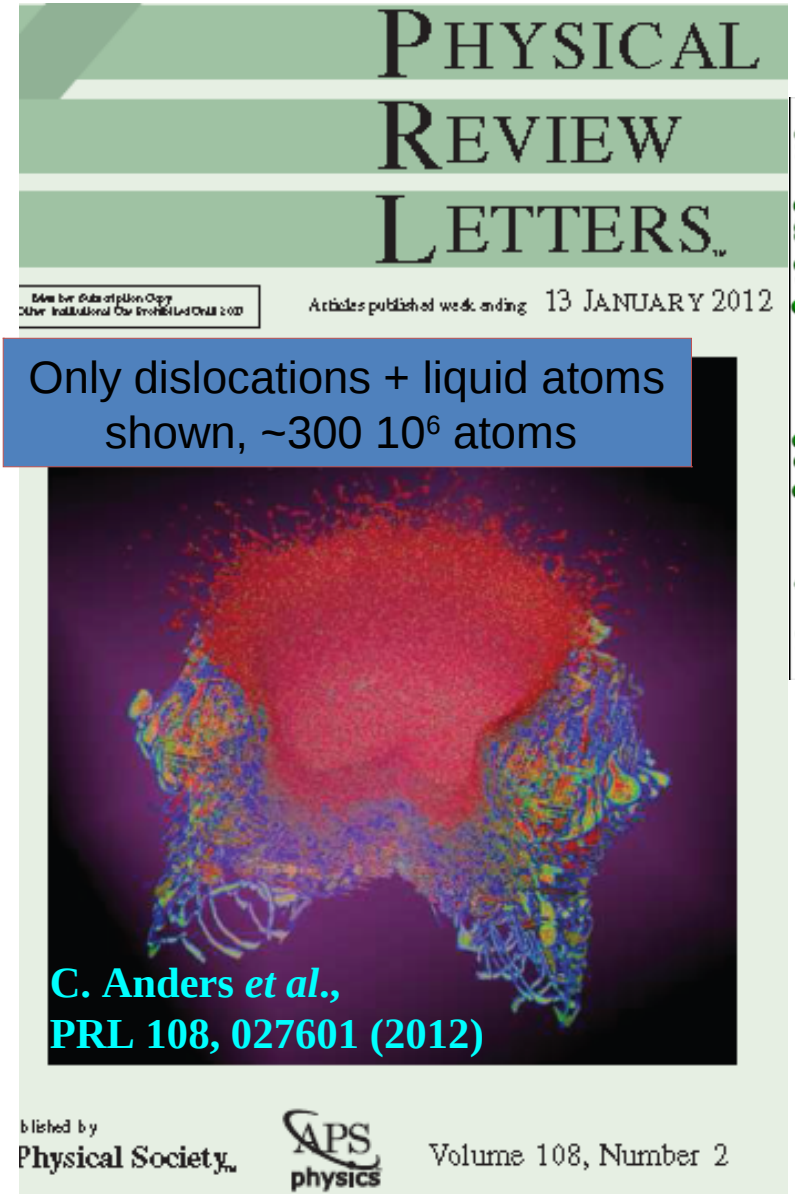

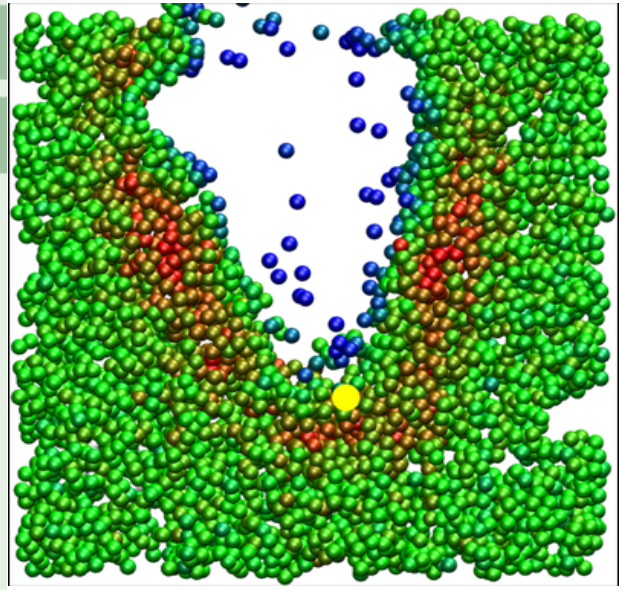

Granular mechanics of grain-surface collisions
Ringl *et al.*, PRE 86, 061313 (2012) **PRE KALEIDOSCOPE**

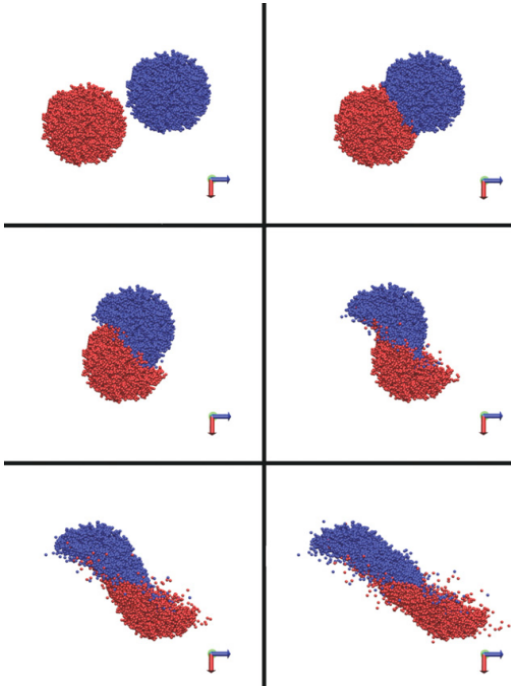

Granular mechanics of nano-grain collisions
Ringl *et al.*, Ap.J. **752** (2012) 151

Motivation

- Improve granular simulation within LAMMPS in GPUs
- Study performance of CPU and GPU code for granular simulations.
- Execute simulations in different hardware configurations.
- Compare results of Granular Simulations with Molecular Dynamics using Lennard-Jones potential.

Granular Simulations cont.

• Granular mechanics plays an important role in many branches of science and engineering, from astrophysics applications in planetary and interstellar dust clouds, to processing of industrial mixtures and powders

• Granular simulations [Silbert(2001), Duran (2000)] follow (spherical) grains and the short-range forces between two grains are classified as normal or tangential forces.

- Normal forces are divided in repulsive and adhesive forces.
- Tangential forces include several friction forces due to gliding, rolling and torsional motion.

• Each grain interacts only with other grains within an extremely short distance, leading to around 3-5 neighbors per grain in these particular simulations.

• The improved granular potential developed in Ringl 2012 is introduced in this work in the GPU version of LAMMPS.

• The potential is called “granularEasy”.

Ringl, C., Urbassek, H.M.: A LAMMPS implementation of granular mechanics: Inclusion of adhesive and microscopic friction forces. Computer Physics Communications 183 (2012) 986–992
Silbert, L.E., Ertas, D., Grest, G.S., Halsey, T.C., Levine, D., Plimpton, S.J.: Granular flow down an inclined plane: Bagnold scaling and rheology. Physical Review E 64 (2001) 051302
Duran, J.: Sands, powders, and grains: an introduction to the physics of granular materials (Springer Verlag, 2000).

Physical scenario

• Large prismatic cell containing 70,000 silica grains arranged in a disordered, amorphous solid structure with high porosity.

• The box is elongated, with its length along the z axis twice its length along x and y axis.

• The top half of the cell is empty, to mimic empty space above a surface. This is a typical starting configuration to model impact processes, where material ejected from the surface moves towards the top of the box.

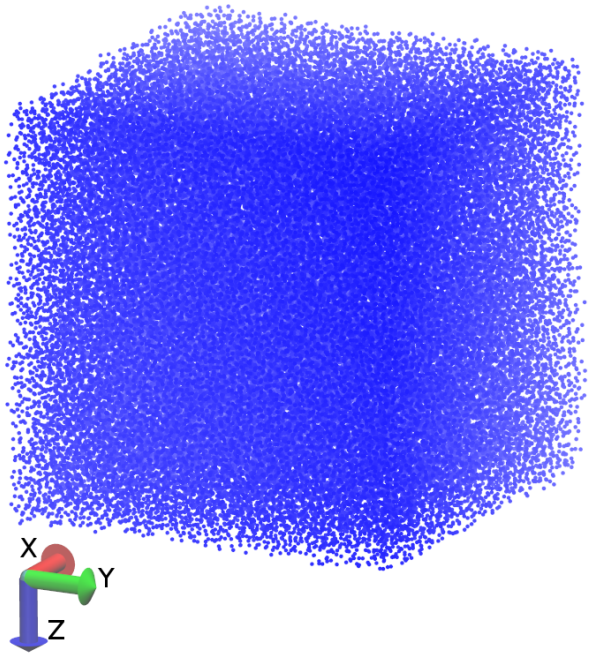

Molecular Dynamics: LAMMPS

LAMMPS characteristics:

- Mature.
- Open source code.
- In active development and with a fairly large user community.
- LAMMPS can work with MPI, OpenMP, CUDA and recently with OpenCL.

LAMMPS & Distributed Computing

- MPI library for different cores of a single workstation or nodes of a cluster.
- OpenMP for multicore workstations. *
- The complete workload can be executed in GPUs.
- The workload can be assigned to CPU cores and the GPUs available at a workstation. Dynamic or static load balancing can be considered.

* For more info on OpenMP support in LAMMPS see Axel Kohlmeyer web site: http://goo.gl/LQqAm

LAMMPS: GPU execution

- GPU Package:
  - Novel, most up to date GPU code. GNU GPL v2 license.
  - Main developers: Paul Crozier (Sandia), Mike Brown, Arnold Tharrington, Scott Hampton (Oak Ridge), Axel Kohlmeyer (Temple), Christian Trott, Lars Winterfeld (Ilmenau, Germany), Duncan Poole, Peng Wang (Nvidia), etc.
  - Multiple MPI processes can execute in the same GPU card.
  - Geryon Library: supports OpenCL and CUDA see http://users.nccs.gov/~wb8/geryon/index.htm
- USER-CUDA package:
  - Main developer: Christian Trott
  - Supports only CUDA.
  - In some type of simulations is faster than GPU package.
  - Only one MPI process can be running in a GPU.
  - The USER-CUDA package uses less memory than the GPU package.
  - Supports Granular simulations

For more information see: http://lammps.sandia.gov/doc/Section_accelerate.html#acc_8

HW & SW Infrastructure

- Phenom Workstation: AMD Phenom 1055T x6 at 2.8 GHz, with 12 GB of RAM and NVIDIA Tesla C2050 with 448 cores at 1.15 GHz and 3 GB of ECC memory. Slackware 13.37 64 bit with OpenMPI 1.4.2 and GCC 4.5.3.

- Cluster: two nodes with four AMD Opteron 6272 of 16 cores, 64 CPU cores with 128 GB of RAM each node. Rocks Linux 5.5 with OpenMPI 1.4.3 and GCC 4.4

- UNC Mendieta cluster: eight nodes with two Intel Xeon E5-2680 CPUs at 2.7GHz and 32 GB of RAM each, housing twelve NVIDIA Tesla M2090 GPUs with 6 GB GDDR5 memory (177 GBps) and 512 cuda cores at 1.3 GHz, and two NVIDIA Tesla C2075 GPUs with 6 GB GDDR5 memory (150 GBps) and 448 cuda cores at 1.15 GHz. 20 Gbps InfiniBand DDR. Linux CentOS 6.4 with MPICH 3.0.4.

- LAMMPS version dated 11-Mar-2013, compiled with OpenMPI 1.4.2 and gcc 4.5.3 with -O2 optimizations, USER-CUDA package with CUDA 5.0 and compute capability 2.0.

Benchmarks: Phenom workstation

Granular simulation with the GranularEasy pair style, 13 different amounts of grains, and 1000 steps, GPU Tesla C2050.
The 7.5e4 curve represents the results obtained in C. Ringl, Comp. Phys. **183**, 986 (2012).
The 0.6e4 curve is the performed obtained in this work using the GPU.

CPU: AMD Phenom x6 1055t 2.8GHz
GPU: NVIDIA Tesla C2050

AVG speedup
GPU vs **1** CPU core = 7x
GPU vs **6** CPU core = 2.95x

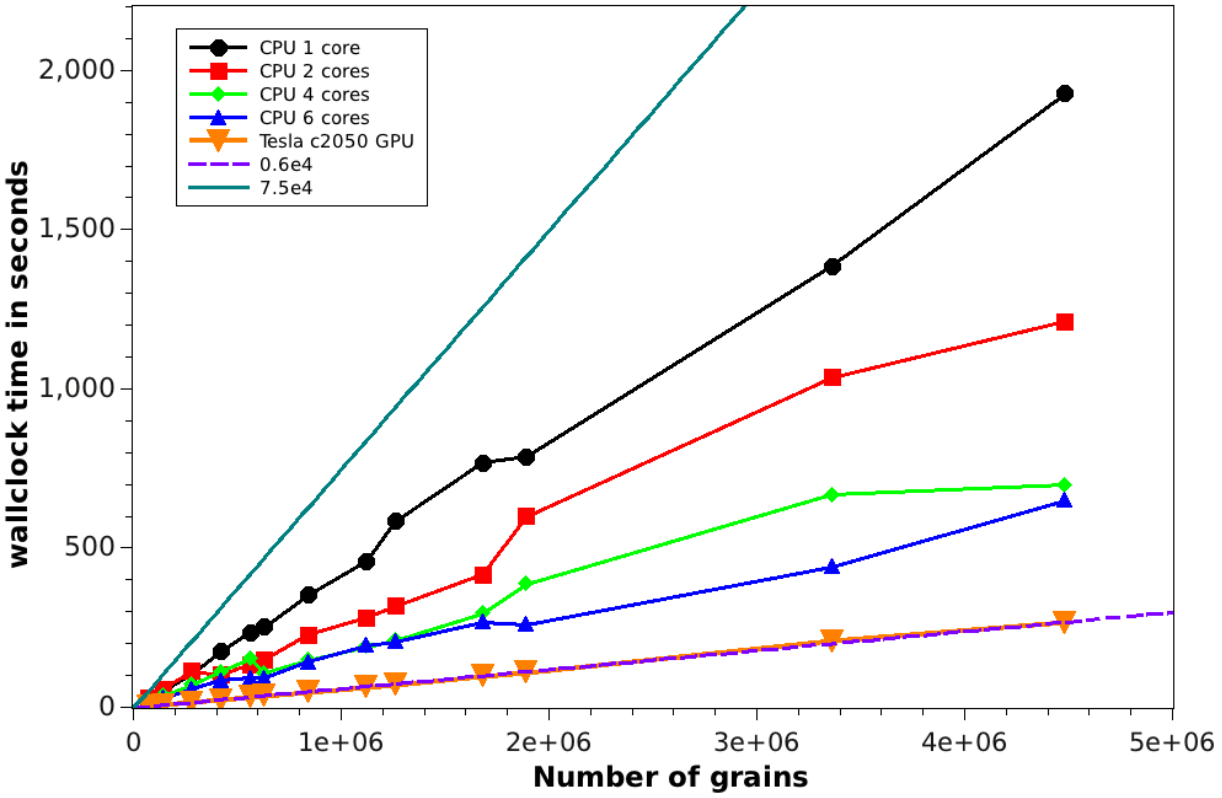

Benchmarks: Clusters

Granular simulation with the GranularEasy pair style, with 4.48e6 grains and 1000 steps, for 1 through 64 processes, in Mendieta and ICB-ITIC clusters. Various GPUs are tested: C2050, C2075 and M2090.

Tesla c2050 GPU ~= 16 CPU cores ICB-ITIC cluster.

Mendieta Tesla M2090 GPUs best performance using 4 GPUs in two cluster nodes speedup of ~4.2 x against the best CPU result (ICB-ITIC cluster with 16 CPU cores).

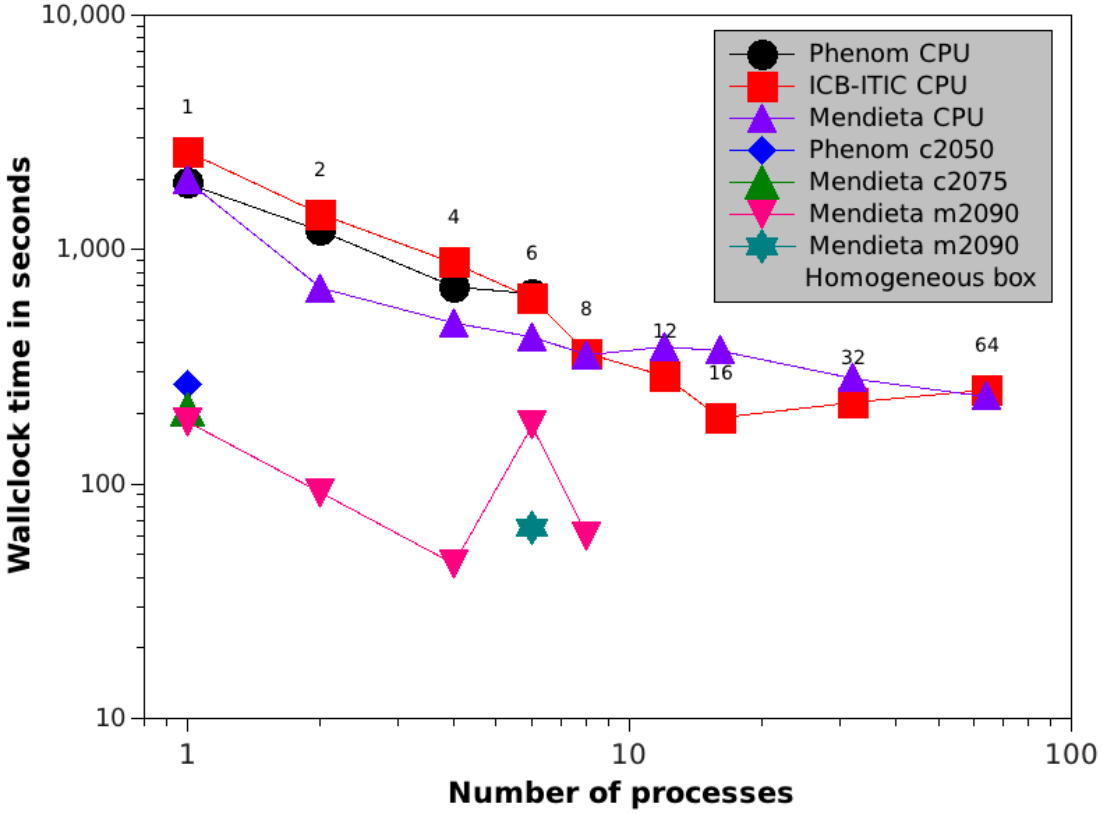

Benchmarks: Communication

GPU Melt simulation, Lennard Jones potential, 1000 steps with 256e3 atoms, each atom has ~70 neighbors with 2.5 cutoff and ~500 neighbors with 5.0 cutoff

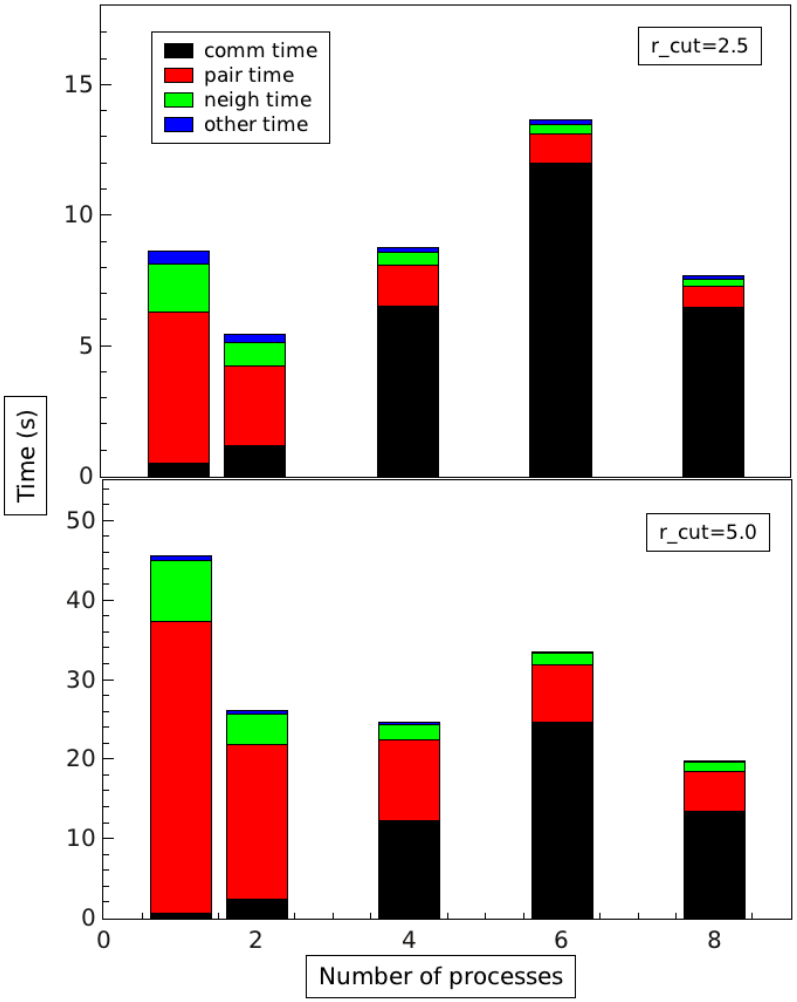

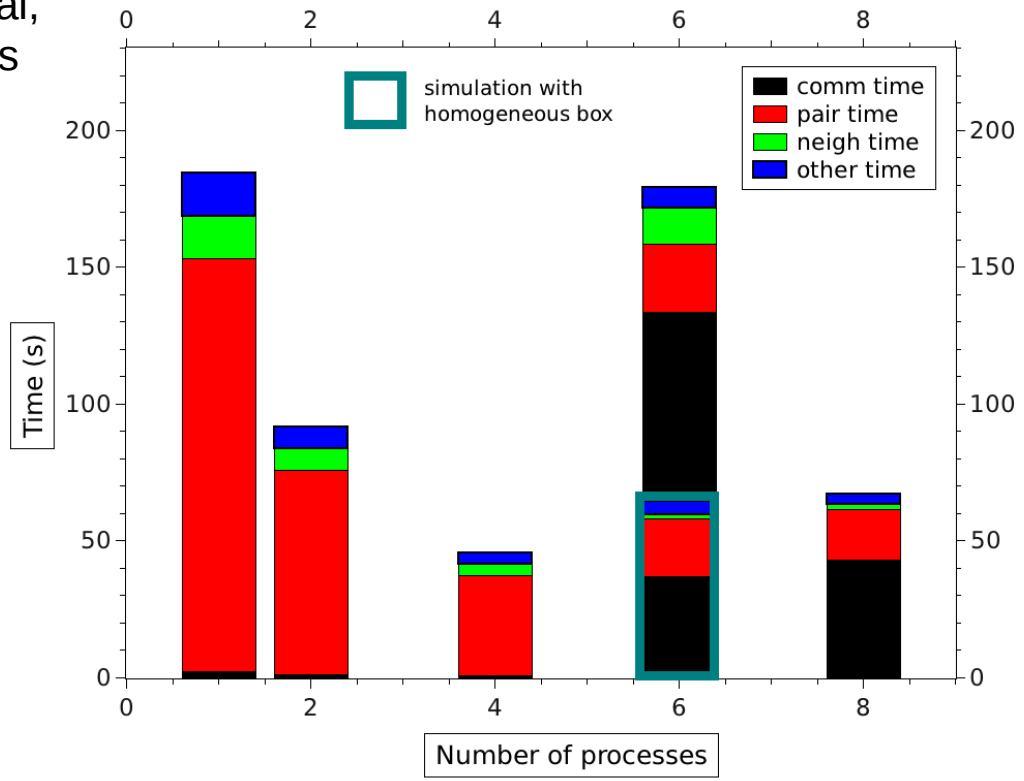

GPU Granular simulation, 1000 steps with 4.48e6 grains. For six processes, two simulations are shown, one has an elongated, half empty heterogeneous box, and the second has a cubic homogeneous box filled with grains. Each grain has 2-5 neighbors.

Conclusions

- An improved granular model was implemented to run in GPUs within the USER-CUDA package of the LAMMPS Molecular Dynamics (MD) software.

- The granular model was validated by detailed comparison to results from the CPU version, and the code is being submitted for inclusion in the LAMMPS distribution.

- Speedups of ~7 x (compared to one CPU core) and ~3 x (compared to six CPU cores) were obtained.

- Benchmarks performed in different hardware configurations, including one CPU cluster and a hybrid CPU-GPU cluster, showed that GPU code is always faster than the CPU version.

Conclusions cont.

- For a single GPU, a linear scaling is observed with increasing system size.

- Beyond 4 GPUs processes, for the system tested, performance degrades due to communication increase.

- For granular simulations, a Tesla c2050 GPU is similar in performance to 16 CPU cores in the ICB-ITIC cluster (12th in LARTOP 50).

- The Mendieta cluster using Tesla M2090 GPUs give the best performance in our tests, using 4 GPUs in two cluster nodes gives a speedup of ~4.2 x against the best CPU result (ICB-ITIC cluster with 16 CPU cores). Similar results in GPU cluster in KTU (Germany).

Other GPU related work

- Biased Monte Carlo in GPU to describe settlements in arid environments (Millan et al., submitted (2013)

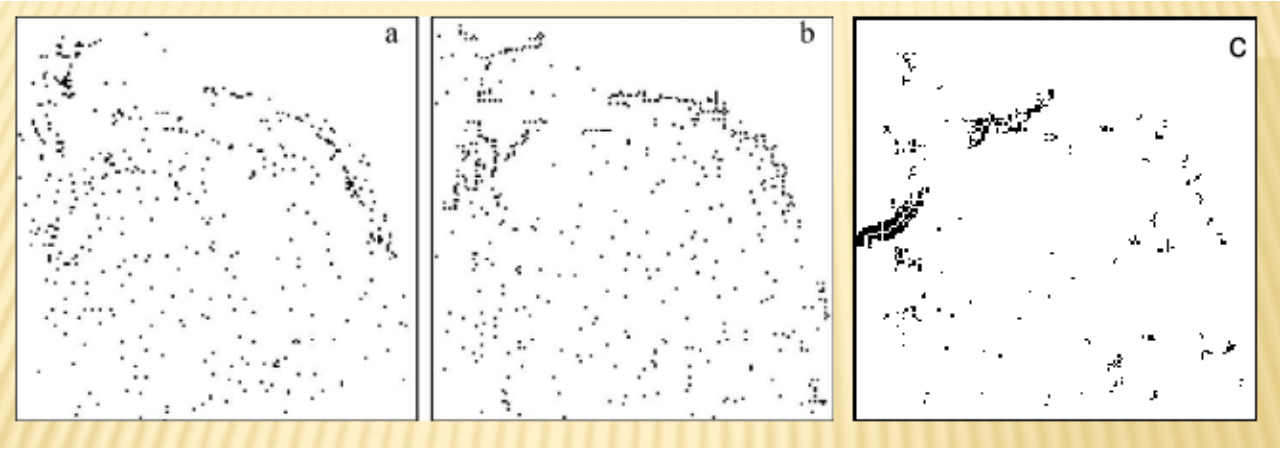

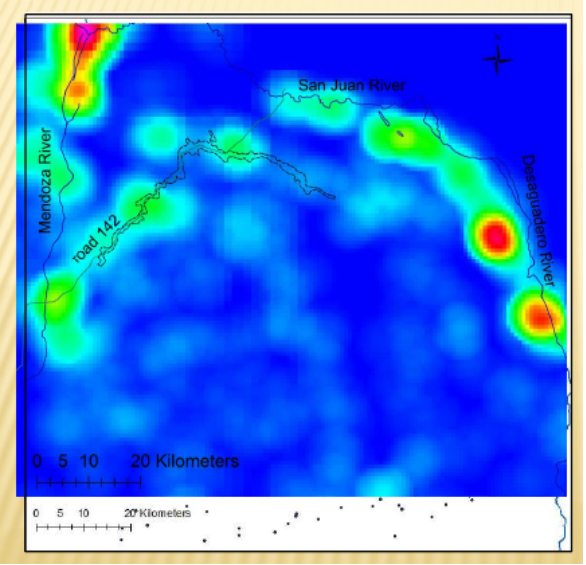

- Cellular Automata in GPU clusters (CACIC 2013, submitted)
- Medical Imaging processing (Fioretti et al., HPC2012)
- Plasticity in granular collisions (Millan et al., CACIC2012, MACI2013)
- Reaction-diffusion solvers (Schwen et al.)
- X-ray diffraction simulation (Higginbotham et al.)

LAMMPS (http://lammps.sandia.gov/) on GPUs

Package GPU (former GPU-LAMMPS)

•Need **CUDA GPU driver** and **CUDA toolkit**, but CUDA SDK is not needed.

•**Sample compilation (**compile gpu library, add files to main src dir, then compile whole code**)**:

cd ~/lammps/lib/gpu

emacs Makefile.linux (compute capability 1.3/2.0, single/mix/double precision)

make -f Makefile.linux (to obtain libgpu.a)

cd ../../src

emacs ./MAKE/Makefile.linux (here need to change options and paths)

make yes-asphere (for granular materials)

Make yes-manybody (for EAM in Package GPU)

make yes-kspace (for electrostatics, needs fftw.2.x)

make yes-gpu (to copy /src/gpu files to /src)

make linux (to obtain executable)

Modified Granular simulations on LAMMPS

Can simulate more than 1 million particles in workstation.
Improved adhesive forces and friction (gliding, rolling, and torsional).

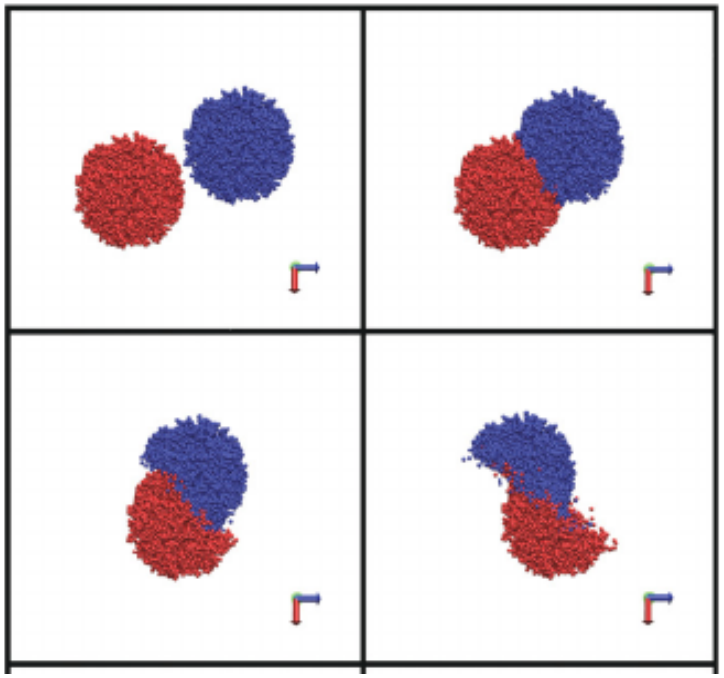

Velocity v = 5 m/s (v = 29_{vfrag}) & impact parameter b = 0.8 R

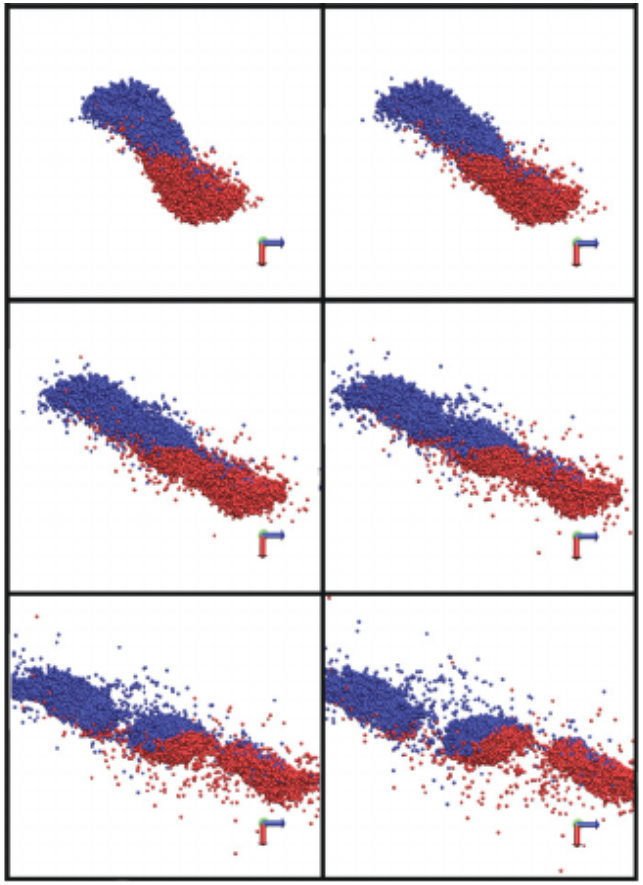

C. RINGL AND H.M. URBASSEK, A LAMMPS implementation of granular mechanics: Inclusion of adhesive and microscopic friction forces, Computer Physics Communications 183 (2012), pp 986-992.

Granular Simulations

• Granular mechanics plays an important role in many branches of science and engineering, from astrophysics applications in planetary and interstellar dust clouds, to processing of industrial mixtures and powders

• Granular simulations follow trajectories of particles with a finite volume which cannot be penetrated by other particles.

• Interactions amongst particles are typically short-ranged, and include contact and adhesive forces.

• Typical MD codes for point particles can also be adapted to run granular simulations efficiently, LAMMPS includes the possibility of running a few different granular models both in CPUs and GPUs.