GPU Acceleration of Aspherical Particle Simulations

W. Michael Brown, Steve Plimpton

Thursday, February 25, 2010

Why Aspherical Particles?

- Particles in nature and manufacturing often have highly irregular shapes
- Liquid crystal simulations
- Coarse Graining
- Majority of computational particle mechanics (CPM) simulators treat only spherical particles
- Need a parallel and scalable implementation to attack realistic problems (LAMMPS)

Gay-Berne Potential

- Single-site potential for aspherical
- **S** is the shape matrix
- The **E** matrix characterizes the relative well depths of side-to-side, face-to-face, and end-to end interactions
- ~30 times the cost of an LJ interaction

$$U = U_r\left(\mathbf{A}_1,\mathbf{A}_2,\mathbf{r}_{12}\right)\eta_{12}\left(\mathbf{A}_1,\mathbf{A}_2\right)\chi_{12}\left(\mathbf{A}_1,\mathbf{A}_2,\hat{r}_{12}\right)$$

$$U_r = 4\varepsilon\left[\left(\frac{\sigma}{h_{12}+\gamma\sigma}\right)^{12} - \left(\frac{\sigma}{h_{12}+\gamma\sigma}\right)^{6}\right]$$

$$\eta_{12} = \left[\frac{2s_1s_2}{\det\left[\mathbf{A}_1^T\mathbf{S}_1^2\mathbf{A}_1 + \mathbf{A}_2^T\mathbf{S}_2^2\mathbf{A}_2\right]}\right]^{\upsilon/2}$$

$$s = \left[a_ib_i + c_ic_i\right]\left[a_ib_i\right]^{1/2}$$

$$\chi_{12} = \left[2\hat{\mathbf{r}}_{12}^T\left(\mathbf{A}_1^T\mathbf{E}_1\mathbf{A}_1 + \mathbf{A}_2^T\mathbf{E}_2\mathbf{A}_2\right)^{-1}\hat{\mathbf{r}}_{12}\right]^{\mu}$$

Liquid Crystal Simulations

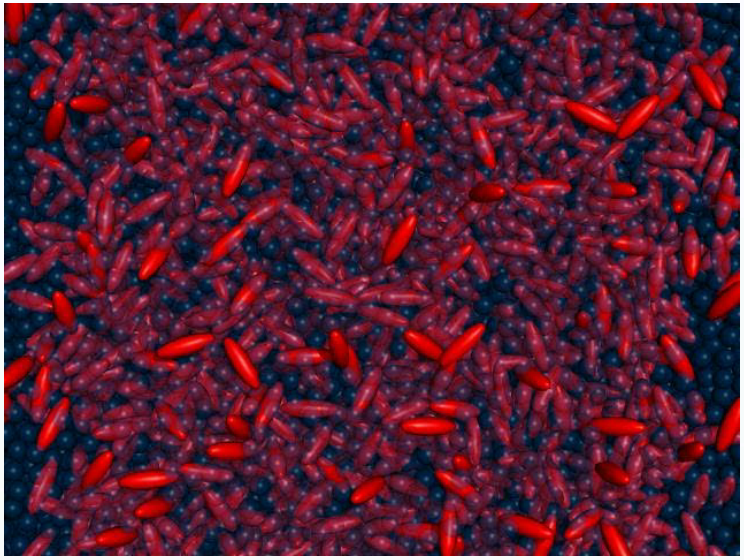

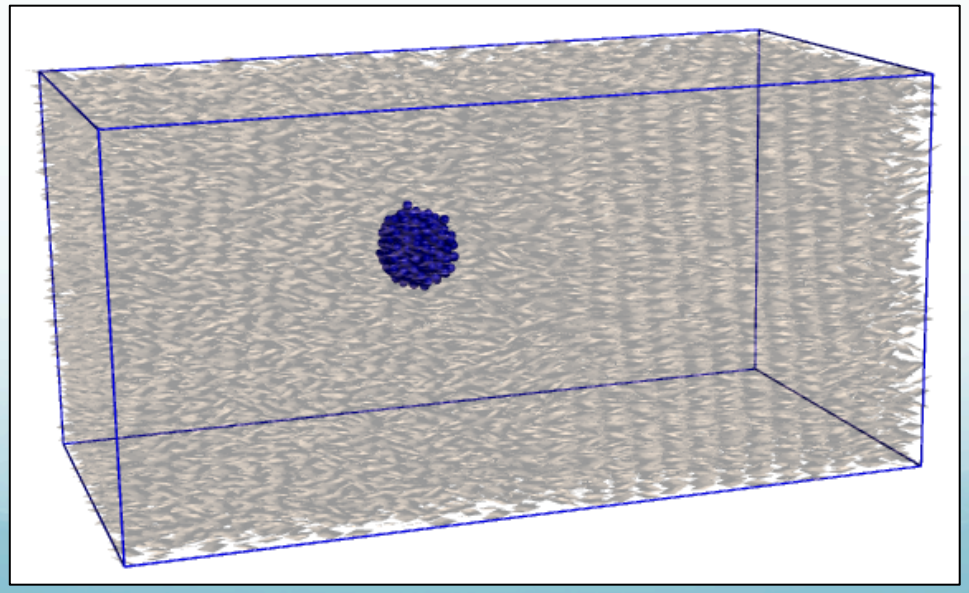

Accelerated Gay-Berne in LAMMPS

- Good candidate for GPU acceleration
  - *Very* expensive force calculation
- Available in the GPU package (make yes-asphere yes-gpu)
  - Can run on multiple GPUs on a single node or in a cluster
  - Multiple precision options: Single/Single, Single/Double, and Double/Double
  - Can simulate millions of particles per GPU

GPGPU Times Speedup vs 1 Core

(*c*=cutoff, *n*=particles)

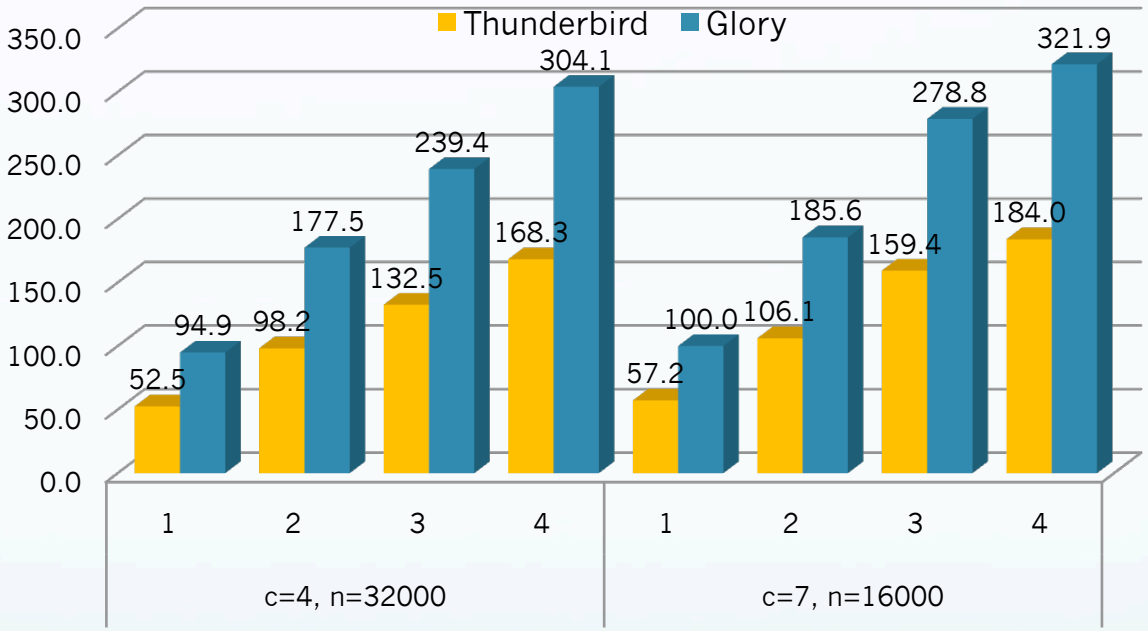

GPGPU: 1, 2, 3, or 4 NVIDIA, 240 core, 1.3 GHz GPGPUs
Thunderbird: *1 core of Dual 3.6 GHz Intel EM64T processors*
Glory: *1 core of Quad Socket/Quad Core 2.2 GHz AMD*

GPGPU Times Speedup vs 1 Node

(*c*=cutoff, *n*=particles)

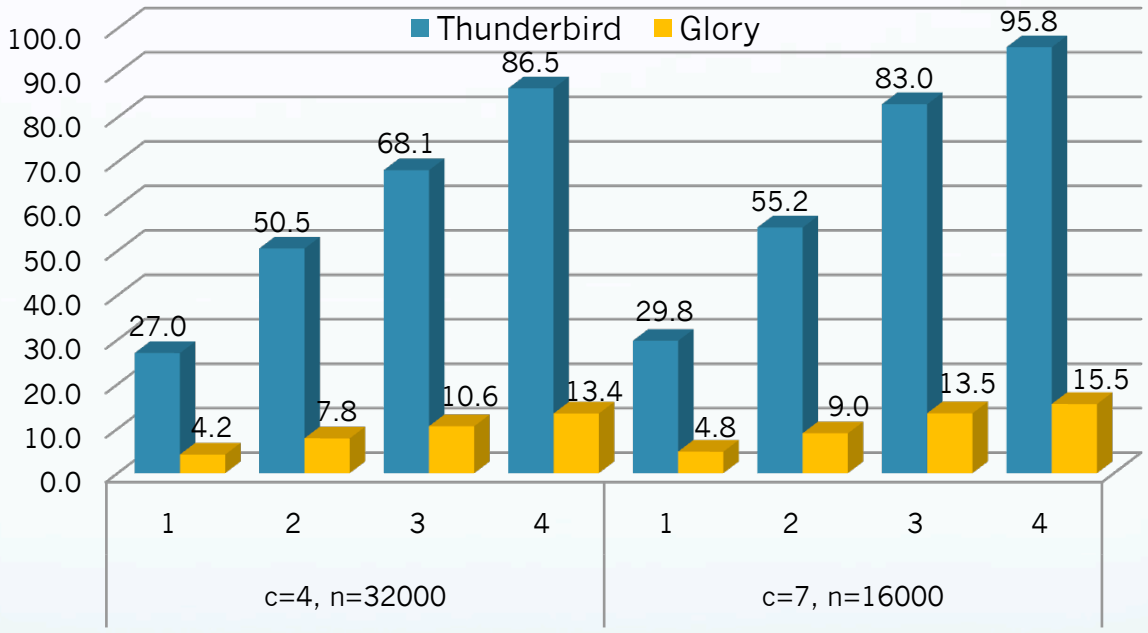

GPGPU: 1, 2, 3, or NVIDIA, 240 core, 1.3 GHz GPGPUs
Thunderbird: *2 procs, Dual 3.6 GHz Intel EM64T processors*
Glory: *16 procs, Quad Socket/Quad Core 2.2 GHz AMD*

64K GB Particle Simulation on Stella

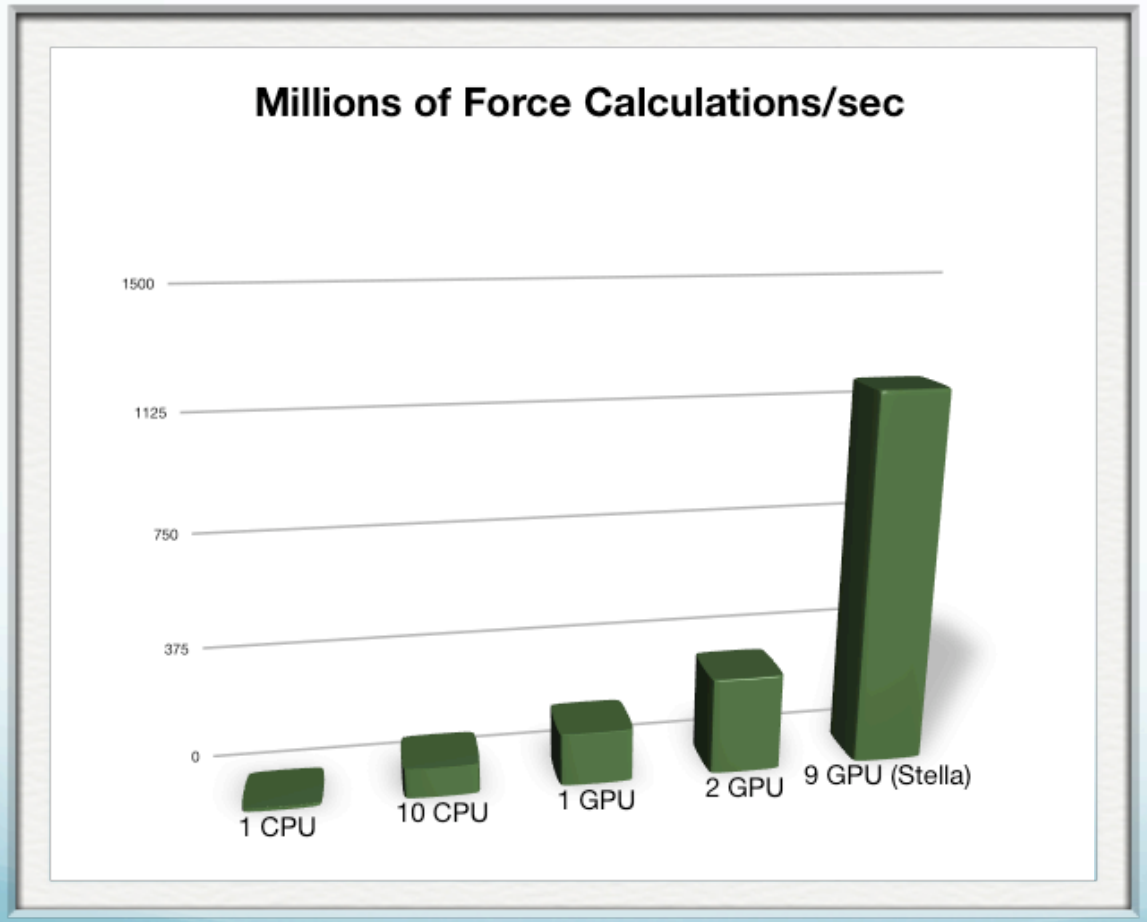

Creative Consultants LLC
Efficient Performance Computing

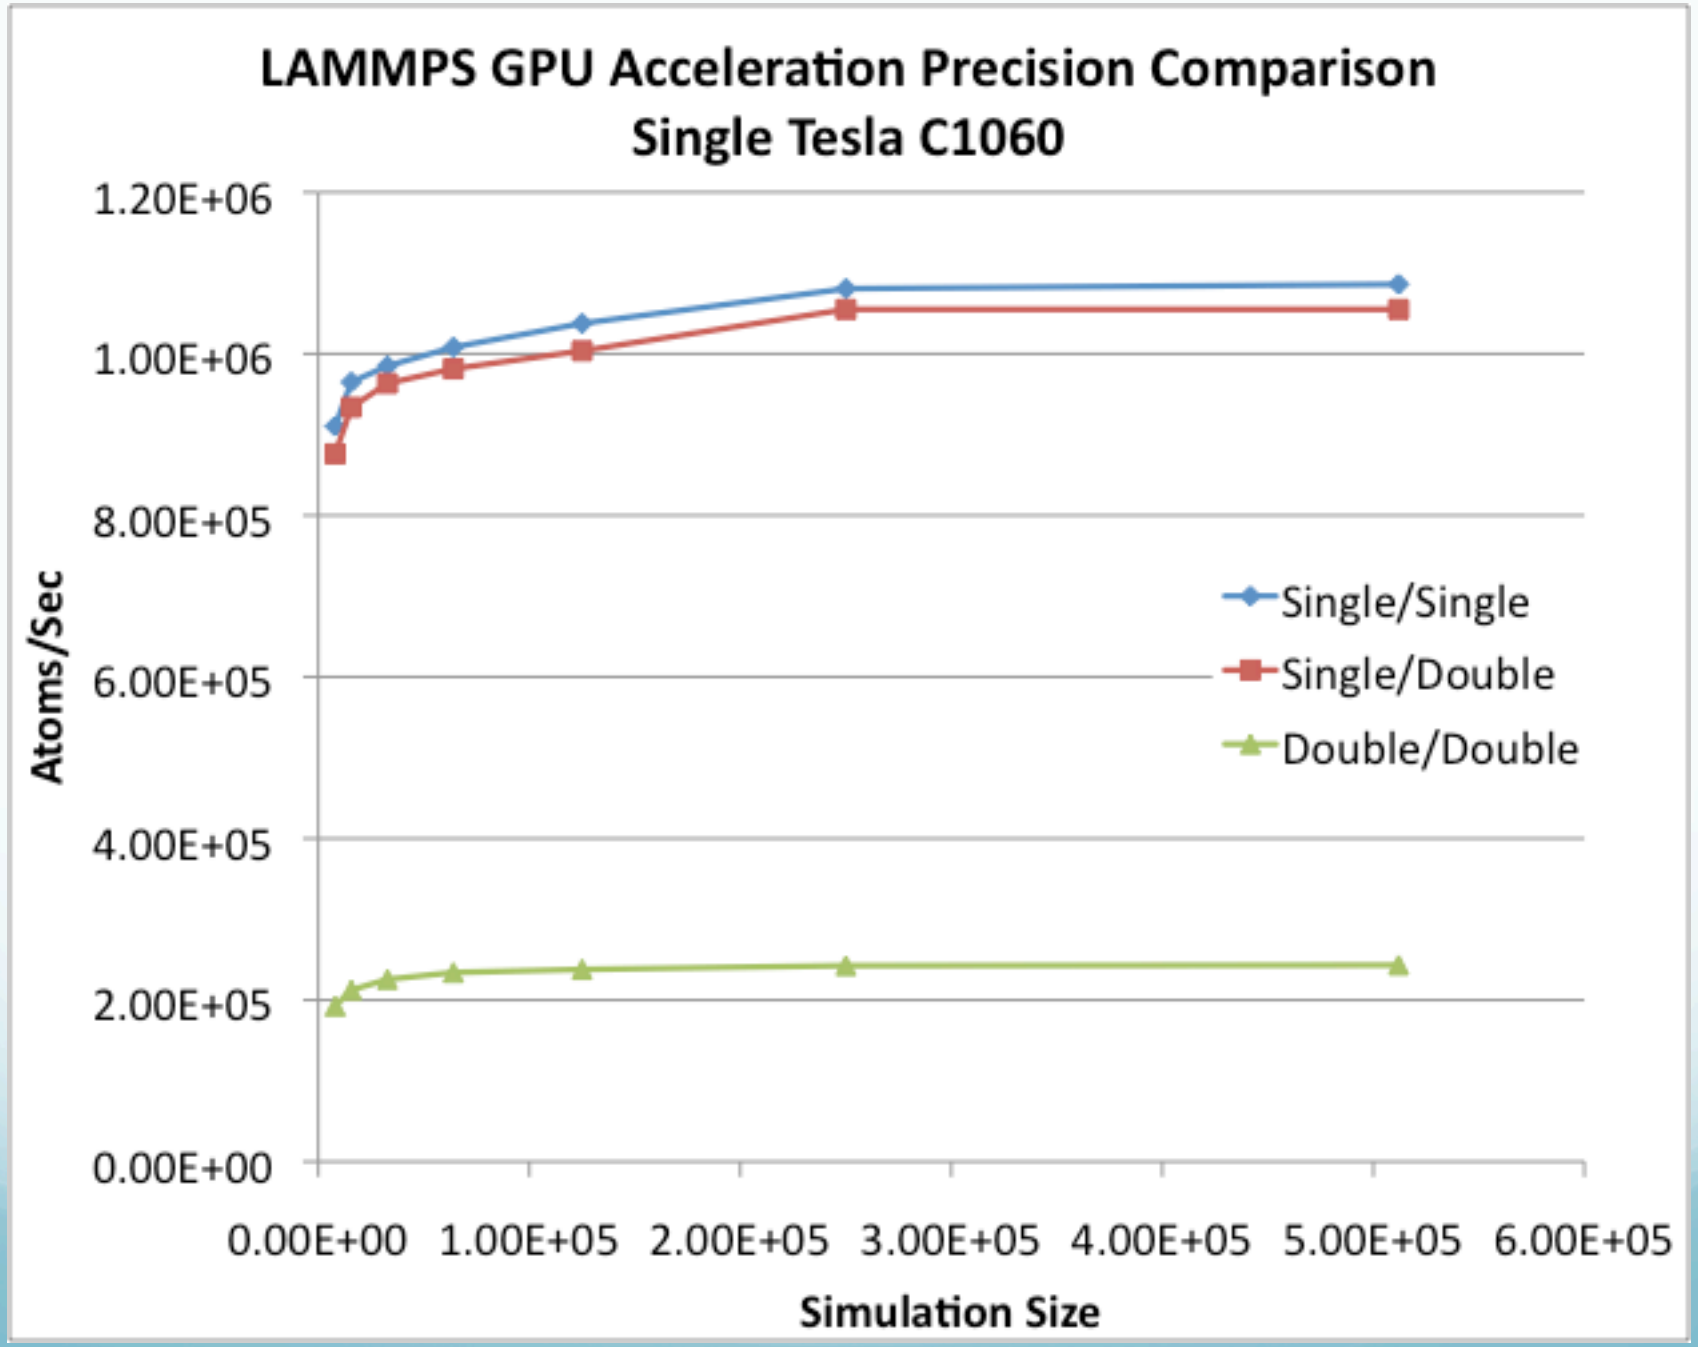
LAMMPS GPU Acceleration Precision Comparison
Single Tesla C1060
1.20E+06
1.00E+06
8.00E+05
6.00E+05
4.00E+05
2.00E+05
0.00E+00
Atoms/Sec
Single/Single
Single/Double
Double/Double
0.00E+00
1.00E+05
2.00E+05
3.00E+05
4.00E+05
5.00E+05
6.00E+05
Simulation Size

Conclusions

- A single 4-GPU accelerated node can run a simulation faster than a 256-core simulation on Thunderbird or Glory.
- The power requirements for the GPU accelerated run were <1.2 kW versus 11.2 kW on Glory or 44.8 kW on Thunderbird

Questions