AMBER and Kepler GPUs

Julia Levites, *Sr. Product Manager, NVIDIA*
Ross Walker, *Assistant Research Professor and NVIDIA CUDA Fellow*
San Diego Supercomputer and Department of Chemistry & Biochemistry

Walker Molecular Dynamics Lab

http://www.wmd-lab.org/

GPU Acceleration	Lipid Force Field Development	QM/MM MD	Automated Refinement

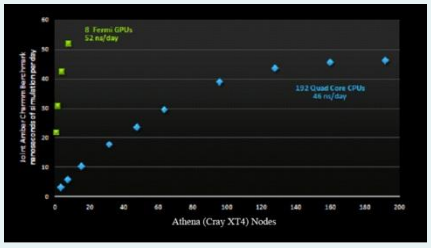

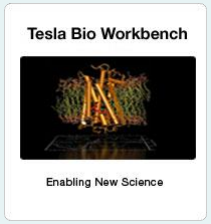

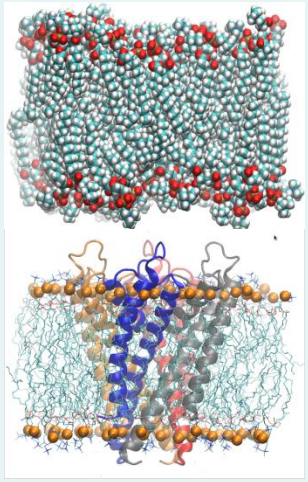

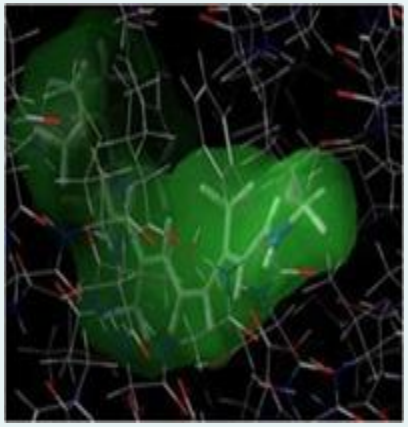

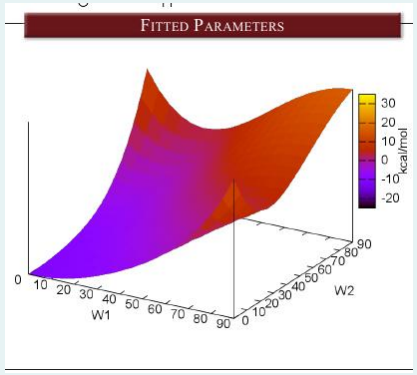

Researchers / Postdocs: Andreas Goetz, Romelia Salomon, JianYin
Graduate Students: Ben Madej (UCSD/SDSC), Justin McCallum (Imperial College), Age Skjevik (UCSD/SDSC/Bergen), Davide Sabbadin (SDSC)
Undergraduate Researchers: Robin Betz, Matthew Clark, Mike Wu

What is Molecular Dynamics?

- **In the context of this talk:**
  - The simulation of the dynamical properties of condensed phase biological systems.
    - Enzymes / Proteins
    - Drug Molecules
    - Biological Catalysts
  - Classical Energy Function
    - Force Fields
    - Parameterized (Bonds, Angles, Dihedrals, VDW, Charges…)
    - Integration of Newton's equations of motion.
    - Atoms modeled as points, electrons included implicitly within the parameterization.

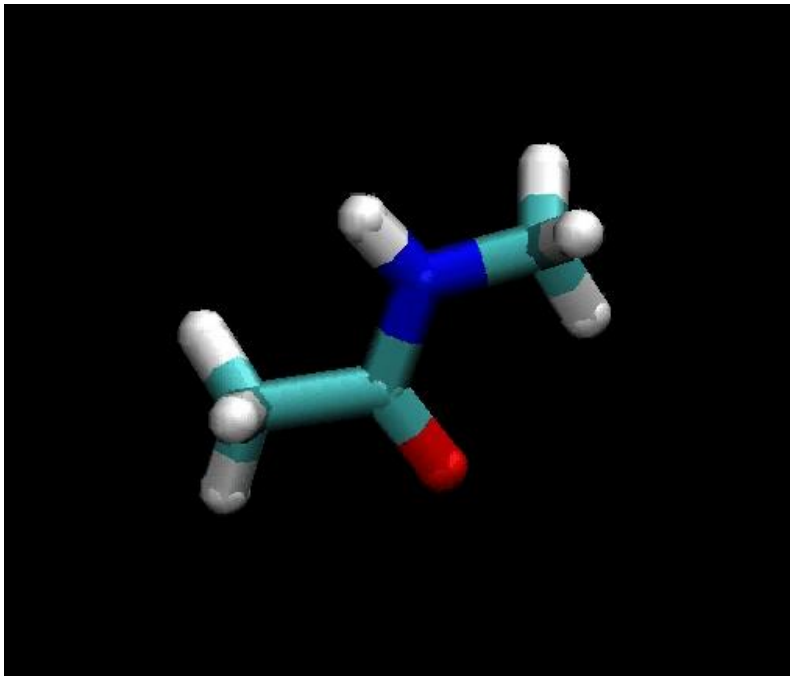

Why Molecular Dynamics?

- **Atoms move!**
  - Life does NOT exist at the global minimum.
  - We may be interested in studying time dependent phenomena, such as molecular vibrations, structural reorganization, diffusion, etc.
  - We may be interested in studying temperature dependant phenomena, such as free energies, anharmonic effects,
  - etc.
- **Ergodic Hypothesis**
  - Time average over trajectory is equivalent to an ensemble average.
  - Allows the use of MD for statistical mechanics studies.

Force Fields

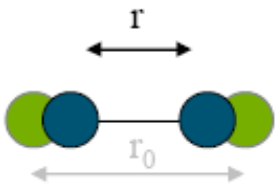

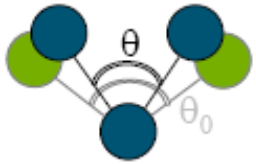

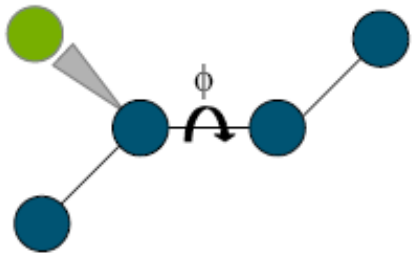

$$E_{stretching} = \sum_{1,2\,pairs} K_r (r - r_0)^2 \qquad E_{bending} = \sum_{1,2\,pairs} K_\theta (\theta - \theta_0)^2 \qquad E_{torsion} = \sum_{1,4\,pairs} K_\phi (1 - \cos(n\phi - \delta))$$

$$E = E_{stretch} + E_{bend} + E_{torsion} + E_{impropers} + E_{electrostatic} + E_{vanderWaals}$$

$$E_{impropers} = \sum_{impr.} K_\varphi (\varphi - \varphi_0)^2$$

$$E_{electrostatic} = \sum_{\substack{\text{nonbonded} \\ \text{i-k pairs}}} q_i \cdot q_k / D.r_{ik}$$

$$E_{van-der-Waals} = \sum_{nonbonded\ pairs} \left(\frac{A_{ik}}{r_{ik}^{12}} - \frac{C_{ik}}{r_{ik}^{6}} \right)$$

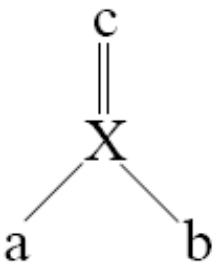

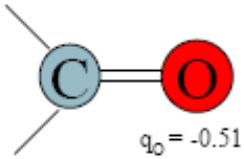

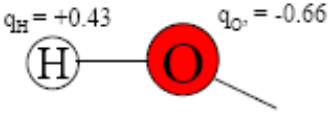

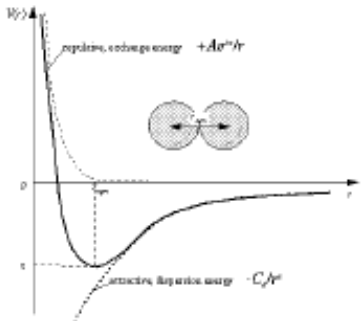

What is AMBER?

An MD simulation package

12 Versions as of 2012

distributed in two parts:

- *AmberTools*, preparatory and analysis programs, free under GPL

- *Amber* the main simulation programs, under academic licensing

independent from the accompanying forcefields

A set of MD forcefields

fixed charge, biomolecular forcefields: ff94, ff99, ff99SB, ff03, ff11, ff12

experimental polarizable forcefields e.g. ff02EP

parameters for general organic molecules, solvents, carbohydrates (Glycam), etc.

in the public domain

The AMBER Development Team
A Multi-Institutional Research Collaboration

Principal contributors to the current codes:

David A. Case (Rutgers University)
Tom Darden (NIEHS)
Thomas E. Cheatham III (University of Utah)
Carlos Simmerling (Stony Brook)
Junmei Wang (UT Southwest Medical Center)
Robert E. Duke (NIEHS and UNC-Chapel Hill)
Ray Luo (UC Irvine)
Mike Crowley (NREL)
Ross Walker (SDSC)
Wei Zhang (TSRI)
Kenneth M. Merz (Florida)
Bing Wang (Florida)
Seth Hayik (Florida)
Adrian Roitberg (Florida)
Gustavo Seabra (Florida)

Kim F. Wong (University of Utah)
Francesco Paesani (University of Utah)
Xiongwu Wu (NIH)
Scott Brozell (TSRI)
Thomas Steinbrecher (TSRI)
Holger Gohlke (J.W. Goethe-Universität)
Lijiang Yang (UC Irvine)
Chunhu Tan (UC Irvine)
John Mongan (UC San Diego)
Viktor Hornak (Stony Brook)
Guanglei Cui (Stony Brook)
David H. Mathews (Rochester)
Celeste Sagui (North Carolina State)
Volodymyr Babin (North Carolina State)
Peter A. Kollman (UC San Francisco)

AMBER Usage

- **Approximately 850 site licenses (per version) across most major countries.**

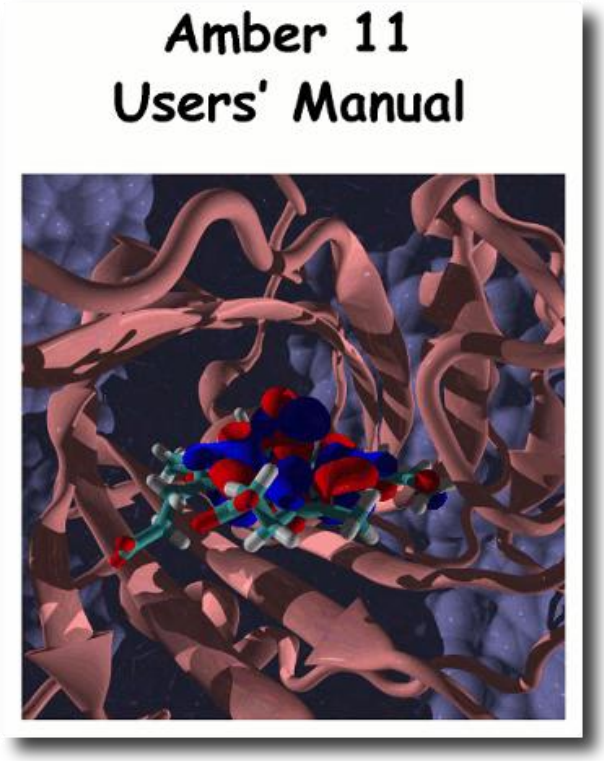

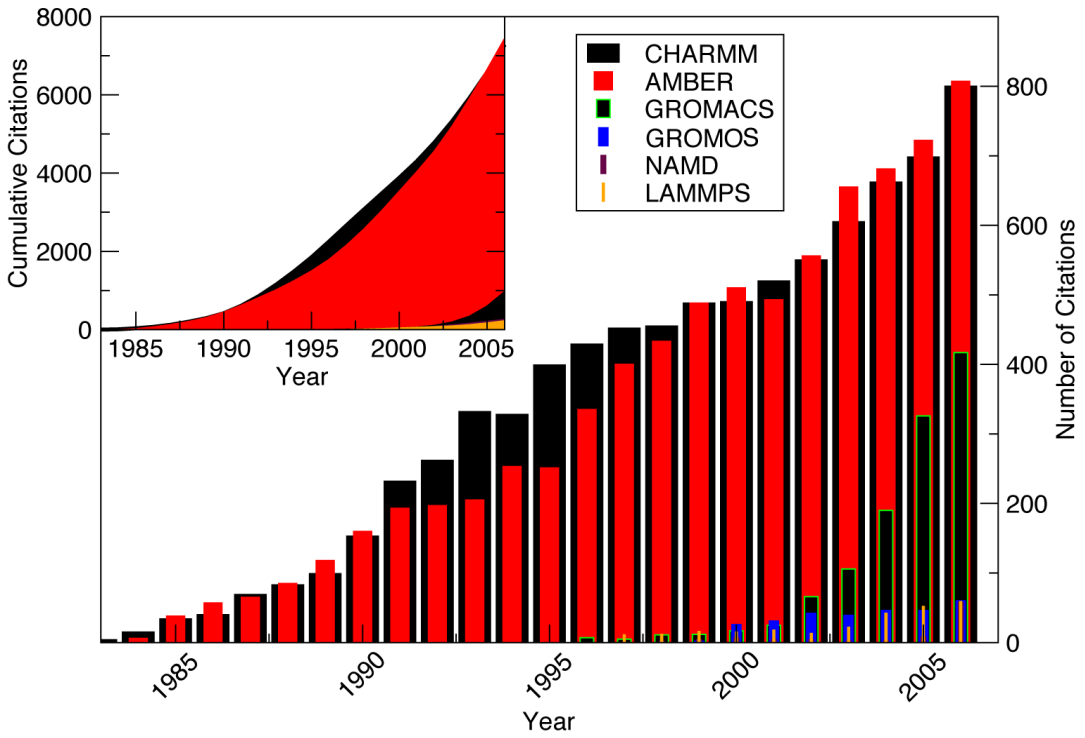

What can we do with Molecular Dynamics?

- **Can simulate time dependent properties.**
  - Protein domain motions.
  - Small Protein Folds.
  - Spectroscopic Properties.

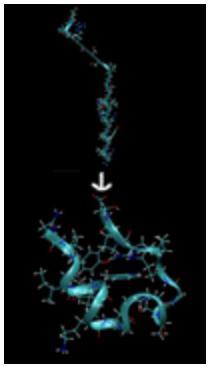

- **Can simulate ensemble properties.**
  - Binding free energies.
    - Drug Design
    - Biocatalyst Design
  - Reaction Pathways
  - Free Energy Surfaces.

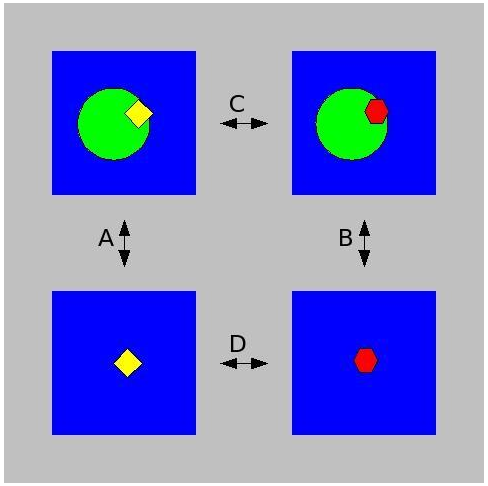

Why do we need Supercomputers? (Complex Equations)

$$U(R) = \sum_{bonds} K_r (r - r_{eq})^2 + \sum_{angles} K_\theta (\theta - \theta_{eq})^2$$

$$+ \sum_{dihedrals} \frac{V_n}{2} \left[1 + \cos(n\phi - \gamma)\right] + \sum_{i<j}^{atoms} \frac{A_{ij}}{R_{ij}^{12}} - \frac{B_{ij}}{R_{ij}^{6}}$$

$$+ \sum_{i<j}^{atoms} \frac{q_i q_j}{\varepsilon R_{ij}}$$

Why do we need Supercomputers?
Lots of Atoms

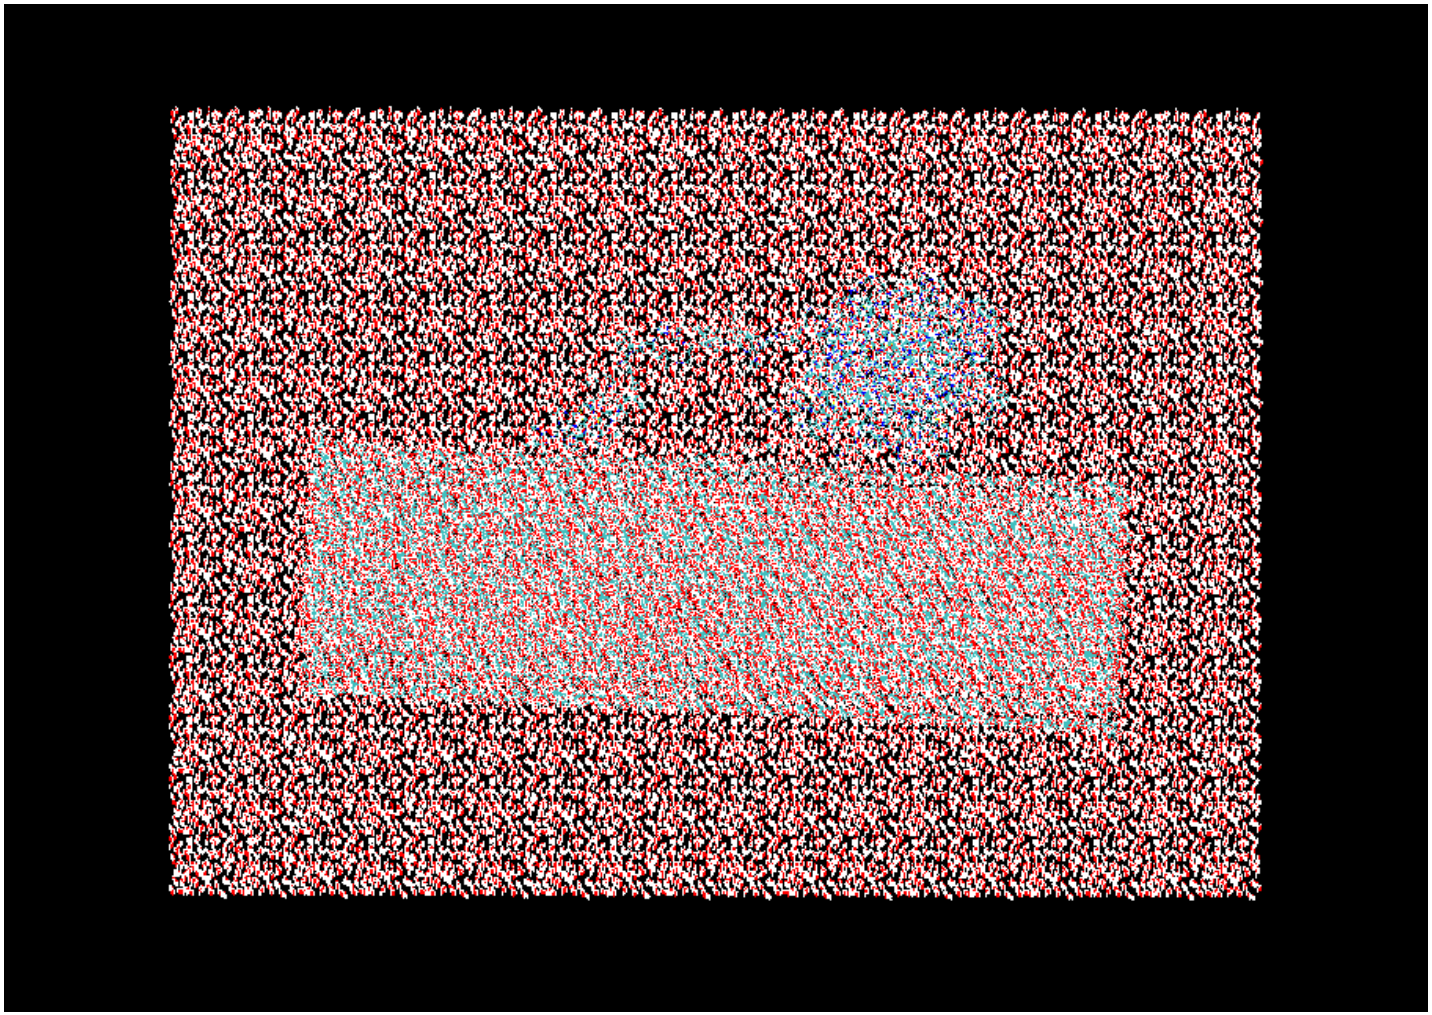

Why do we need Supercomputers? Lots of Time Steps

- ***Maximum time per step is limited by fastest motion in system (vibration of bonds)***

  = 2 femto seconds (0.000000000000002 seconds)

  (Light travels 0.006mm in 2 fs)

- ***Biological activity occurs on the nano-second to micro-second timescale.***

  1 micro second = 0.000001 seconds

SO WE NEED

500 million steps to reach 1 microsecond!!!

Moving Molecular Dynamics Forward?

(at the speed of memory / interconnect?)

Just build bigger supercomputers
Beliefs in moving to the exascale

Atheist

+++R, -P

Heretic

++R, P

Believer

+R, +P

R = Node Speed
P = No. Nodes

(Today)

R, P

True Believer

R, ++P

Luddites

-R, -P

Fanatic

-R, +++P

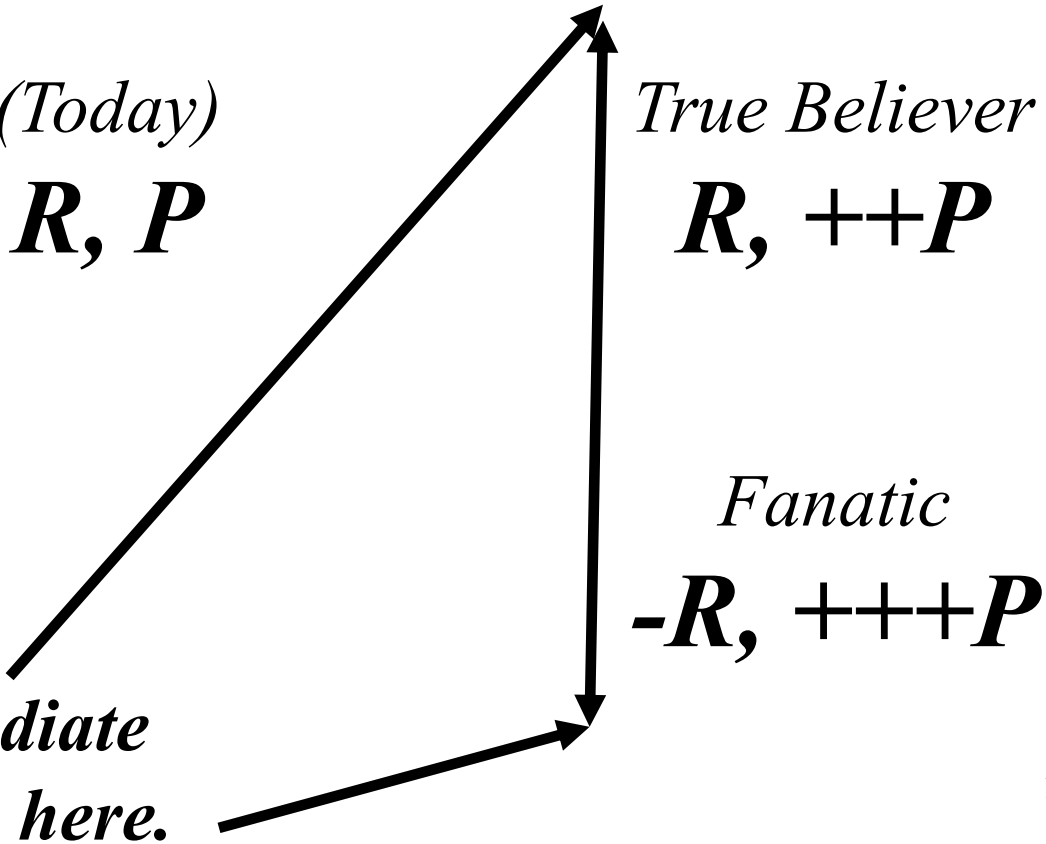

The Problem: Immediate Future is somewhere here.

The Problem(s)

- **Molecular Dynamics is inherently serial.**
  - To compute t+1 we must have computed all previous steps.
  - We cannot simply make the system bigger since these need more sampling (although many people conveniently forget this).
    - 100M atoms = 300M degrees of freedom (d.o.f)
    - 10ns = 5,000,000 time steps = 60x less time steps than d.o.f.
  - We can run ensembles of calculations but these present their own challenges (both practical and political).

Better Science?

- **Bringing the tools the researcher needs into his own lab.**
  - Can we make a researcher's desktop look like a small cluster (remove the queue wait)?
  - Can we make MD truly interactive (real time feedback / experimentation?)
  - Can we find a way for a researcher to cheaply increase the power of all his graduate students workstations?
    - Without having to worry about available power (power cost?).
    - Without having to worry about applying for cluster time.
    - Without having to have a full time 'student' to maintain the group's clusters?
- **GPU's offer a possible cost effective solution.**

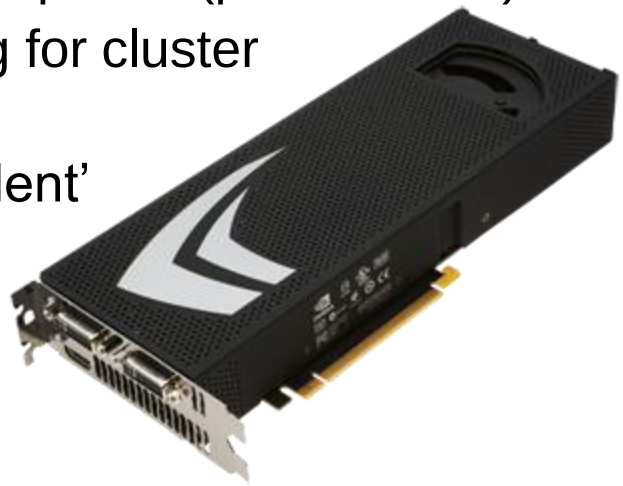

Requirements

- **Any implementation that expects to gain widespread support must be:**
  - Simple / transparent to use.
    - Scientists want science first.
    - Technology is the enabler, NOT the science.
    - Whichever path is the easiest will be the one which is taken.
  - Not make additional approximations.
  - Have broad support.
  - Have longevity (5+ years minimum).

The Project

- **Develop a GPU accelerated version of AMBER's PMEMD.**

San Diego Supercomputer Center

Ross C. Walker

NVIDIA

Scott Le Grand

Duncan Poole

Funded as a pilot project (1 year) under NSF SSE Program & renewed for 3 more years. ☺

SDSC SAN DIEGO SUPERCOMPUTER CENTER

UCSD

NVIDIA

Project Info

- **AMBER Website: http://ambermd.org/gpus/**

Publications

1. Goetz, A.W., Williamson, M.J., Xu, D., Poole, D.,Grand, S.L., Walker, R.C. "Routine microsecond molecular dynamics simulations with amber - part i: Generalized born", Journal of Chemical Theory and Computation, 2012, 8 (5), pp 1542-1555, DOI:10.1021/ct200909j
2. Pierce, L.C.T., Salomon-Ferrer, R. de Oliveira, C.A.F. McCammon, J.A. Walker, R.C., "Routine access to millisecond timescale events with accelerated molecular dynamics.", Journal of Chemical Theory and Computation, 2012, 8 (9), pp 2997-3002, DOI: 10.1021/ct300284c
3. Salomon-Ferrer, R.; Case, D.A.; Walker, R.C.; "An overview of the Amber biomolecular simulation package", WIREs Comput. Mol. Sci., 2012, *in press*, DOI: 10.1002/wcms.1121
4. Grand, S.L.; Goetz, A.W.; Walker, R.C.; "SPFP: Speed without compromise - a mixed precision model for GPU accelerated molecular dynamics simulations", Chem. Phys. Comm., 2013, *184, pp374-380*, DOI: 10.1016/j.cpc.2012.09.022
5. Salomon-Ferrer, R.; Goetz, A.W.; Poole, D.; Le Grand, S.; Walker, R.C.* "*Routine microsecond molecular dynamics simulations with AMBER - Part II: Particle Mesh Ewald*" , *J. Chem. Theory Comput.*, (in review), 2013

Original Design Goals

- **Transparent to the user.**
  - Easy to compile / install.
  - AMBER Input, AMBER Output
  - Simply requires a change in executable name.
- **Cost effective performance.**
  - C2050 should be equivalent to 4 or 6 standard IB nodes.
- **Focus on accuracy.**
  - Should NOT make any additional approximations we cannot rigorously defend.
  - Accuracy / Precision should be directly comparable to the standard CPU implementation.

Version History

- **AMBER 10 – Released Apr 2008**
  - Implicit Solvent GB GPU support released as patch Sept 2009.
- **AMBER 11 – Released Apr 2010**
  - Implicit and Explicit solvent supported internally on single GPU.
  - Oct 2010 – Bugfix.9 doubled performance on single GPU, added multi-GPU support.
- **AMBER 12 – Released Apr 2012**
  - Added Umbrella Sampling Support, REMD, Simulated Annealing, aMD, IPS and Extra Points.
  - Aug 2012 – Bugfix.9 new SPFP precision model, support for Kepler I, GPU accelerate NMR restraints, improved performance.
  - Jan 2013 – Bugfix.14 support CUDA 5.0, Jarzynski on GPU, GBSA. Kepler II support.

Supported Features Summary

- **Supports 'standard' MD**
  - Explicit Solvent (PME)
    - NVE/NVT/NPT
  - Implicit Solvent (Generalized Born)
  - AMBER and CHARMM classical force fields

- **Thermostats**
  - Berendsen, Langevin, Anderson

- **Restraints / Constraints**
  - Standard harmonic restraints
  - Shake on hydrogen atoms

New in AMBER 12

- Umbrella Sampling
- REMD
- Simulated Annealing
- Accelerated MD
- Isotropic Periodic Sum
- Extra Points

Precision Models

***SPSP* - Use single precision for the entire calculation with the exception of SHAKE which is always done in double precision.**

***SPDP* - Use a combination of single precision for calculation and double precision for accumulation (default < AMBER 12.9)**

***DPDP* – Use double precision for the entire calculation.**

SPFP – New![1] – Single / Double / Fixed precision hybrid. Designed for optimum performance on Kepler I. Uses fire and forget atomic ops. Fully deterministic, faster and more precise than SPDP, minimal memory overhead. (default >= AMBER 12.9)

Q24.40 for Forces, Q34.30 for Energies / Virials

[1] Scott Le Grand, Andreas W. Goetz, Ross C. Walker, "*SPFP: Speed without compromise - a mixed precision model for GPU accelerated molecular dynamics simulations*", Comp. Phys. Comm., 2012, in review.

Supported System Sizes

Table 1: Approximate maximum atom counts that can be treated with the GPU implementation of PME explicit solvent simulations in the released version of AMBER 12 using the DPDP, SPDP and SPFP precision models. Test systems are cubic boxes of TIP3P water molecules (for details on the simulations see supplemental information). Error-correction code (ECC) was switched off on the Tesla cards (M2090, K10, K20X).

GPU		PME (max atoms)		
type	memory	DPDP	SPDP	SPFP
GTX580	3.0 GB	870,000	1,060,000	1,240,000
Tesla M2090	6.0 GB	1,820,000	2,230,000	2,680,000
GTX680	2.0 GB	460,000	710,000	920,000
K10	4.0 GB	1,270,000	1,520,000	1,810,000
K20X	6.0 GB	1,890,000	2,270,000	2,710,000

Running on GPUs

- **Details provided on: http://ambermd.org/gpus/**

- **Compile (assuming nvcc >= 4.2 installed)**
  - cd $AMBERHOME
  - ./configure –cuda gnu
  - make install
  - make test

- **Running on GPU**
  - Just replace executable pmemd with pmemd.cuda
    - $AMBERHOME/bin/pmemd.cuda –O –i mdin …

If set process is exclusive mode is on for each GPU, pmemd just 'Does the right thing'™

Performance

AMBER 12

Implicit Solvent Performance (SPFP)

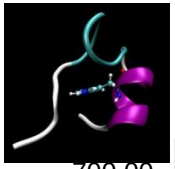

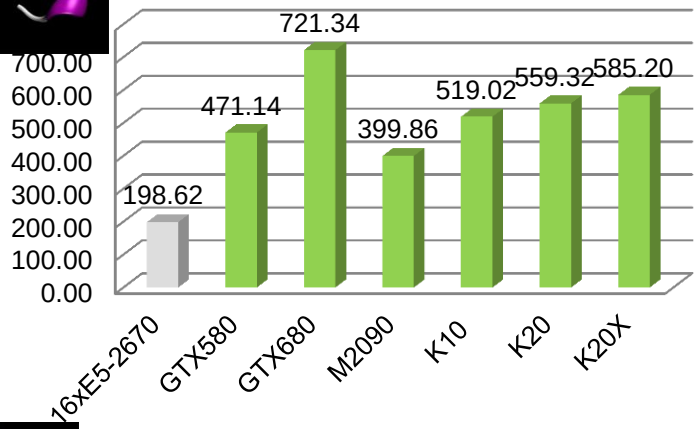

As expected the performance differential is larger for bigger systems.

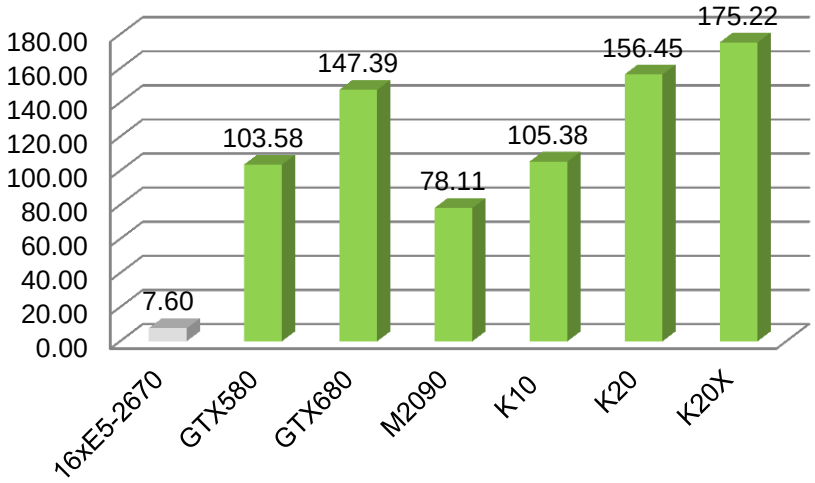

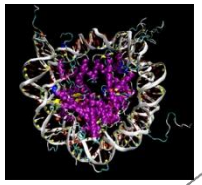

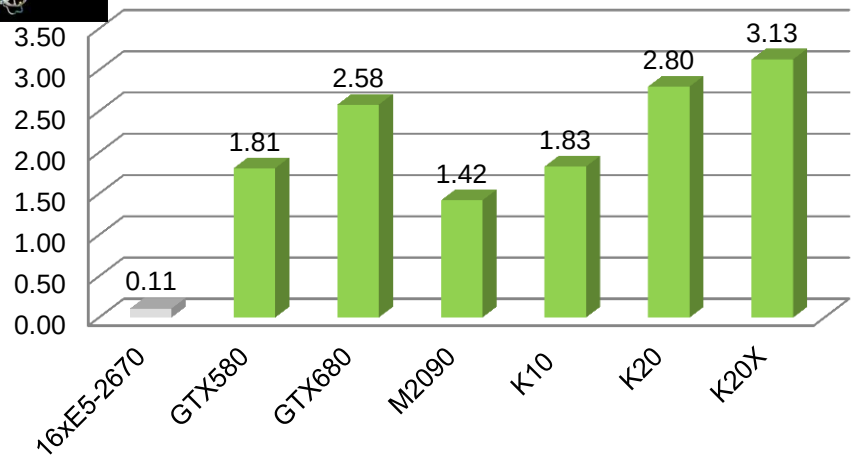

Throughput in ns/day

Explicit Solvent Performance (JAC-DHFR NVE Production Benchmark)

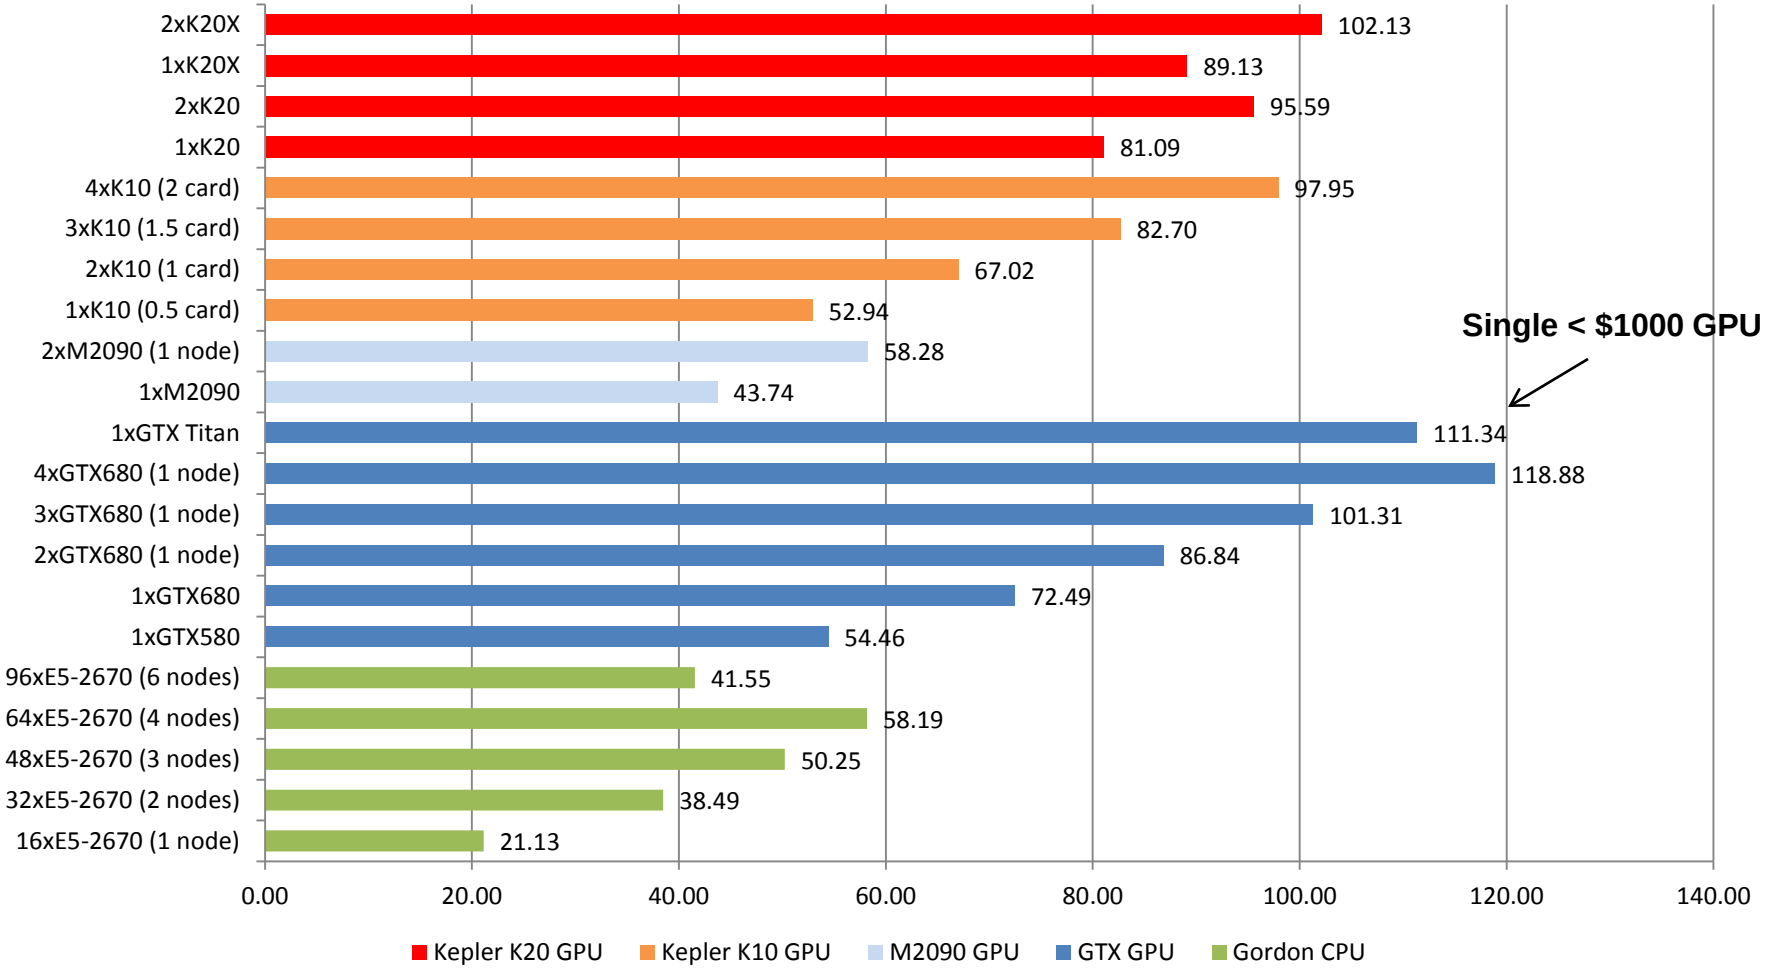

Throughput in ns/day

Explicit Solvent Benchmark (Cellulose NVE Production Benchmark)

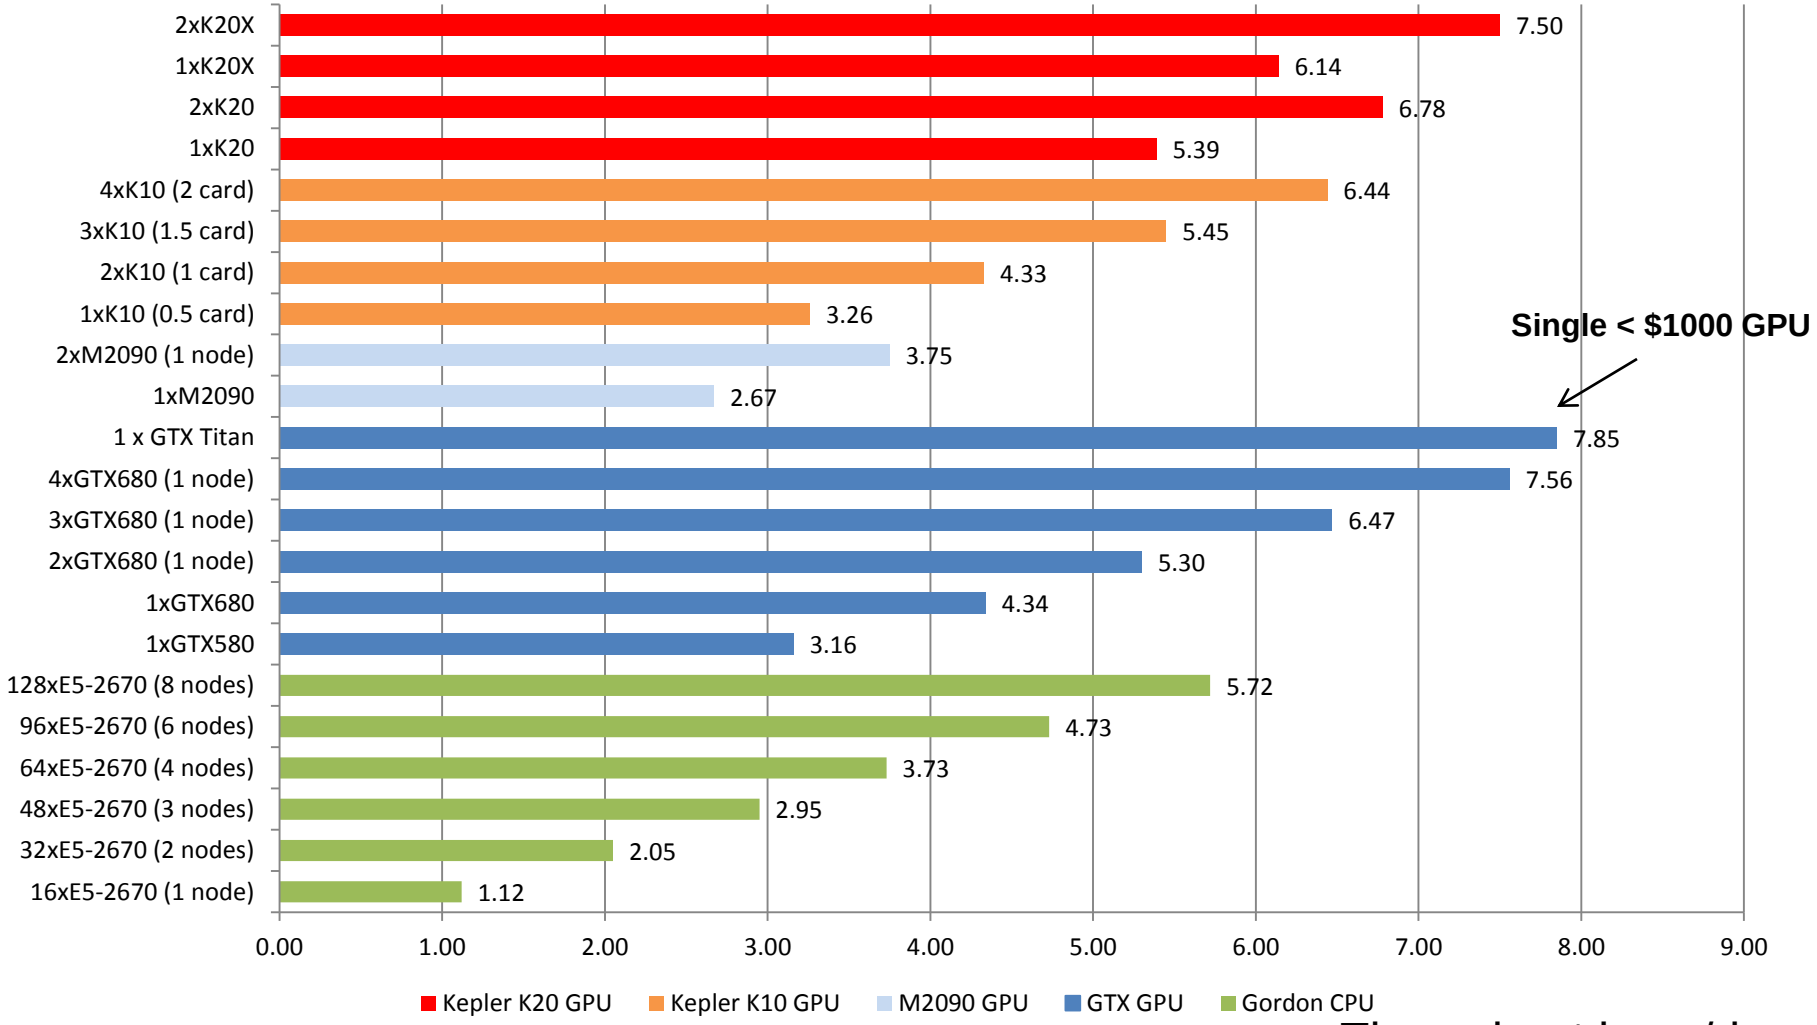

Performance Example (TRP Cage)

CPU 8xE5462

```
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:50][caffeine:0.00][rcw:2009_11_13_Benchmark_for_SC09_Talk]$ tail -f mdout
```

GPU C2050

```
[10:35][caffeine:0.98][rcw:2009_11_13_Benchmark_for_SC09_Talk]$ rm -f mdout
[10:35][caffeine:0.83][rcw:2009_11_13_Benchmark_for_SC09_Talk]$ touch mdout
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$
[10:35][caffeine:0.76][rcw:2009_11_13_Benchmark_for_SC09_Talk]$ tail -f mdout
```

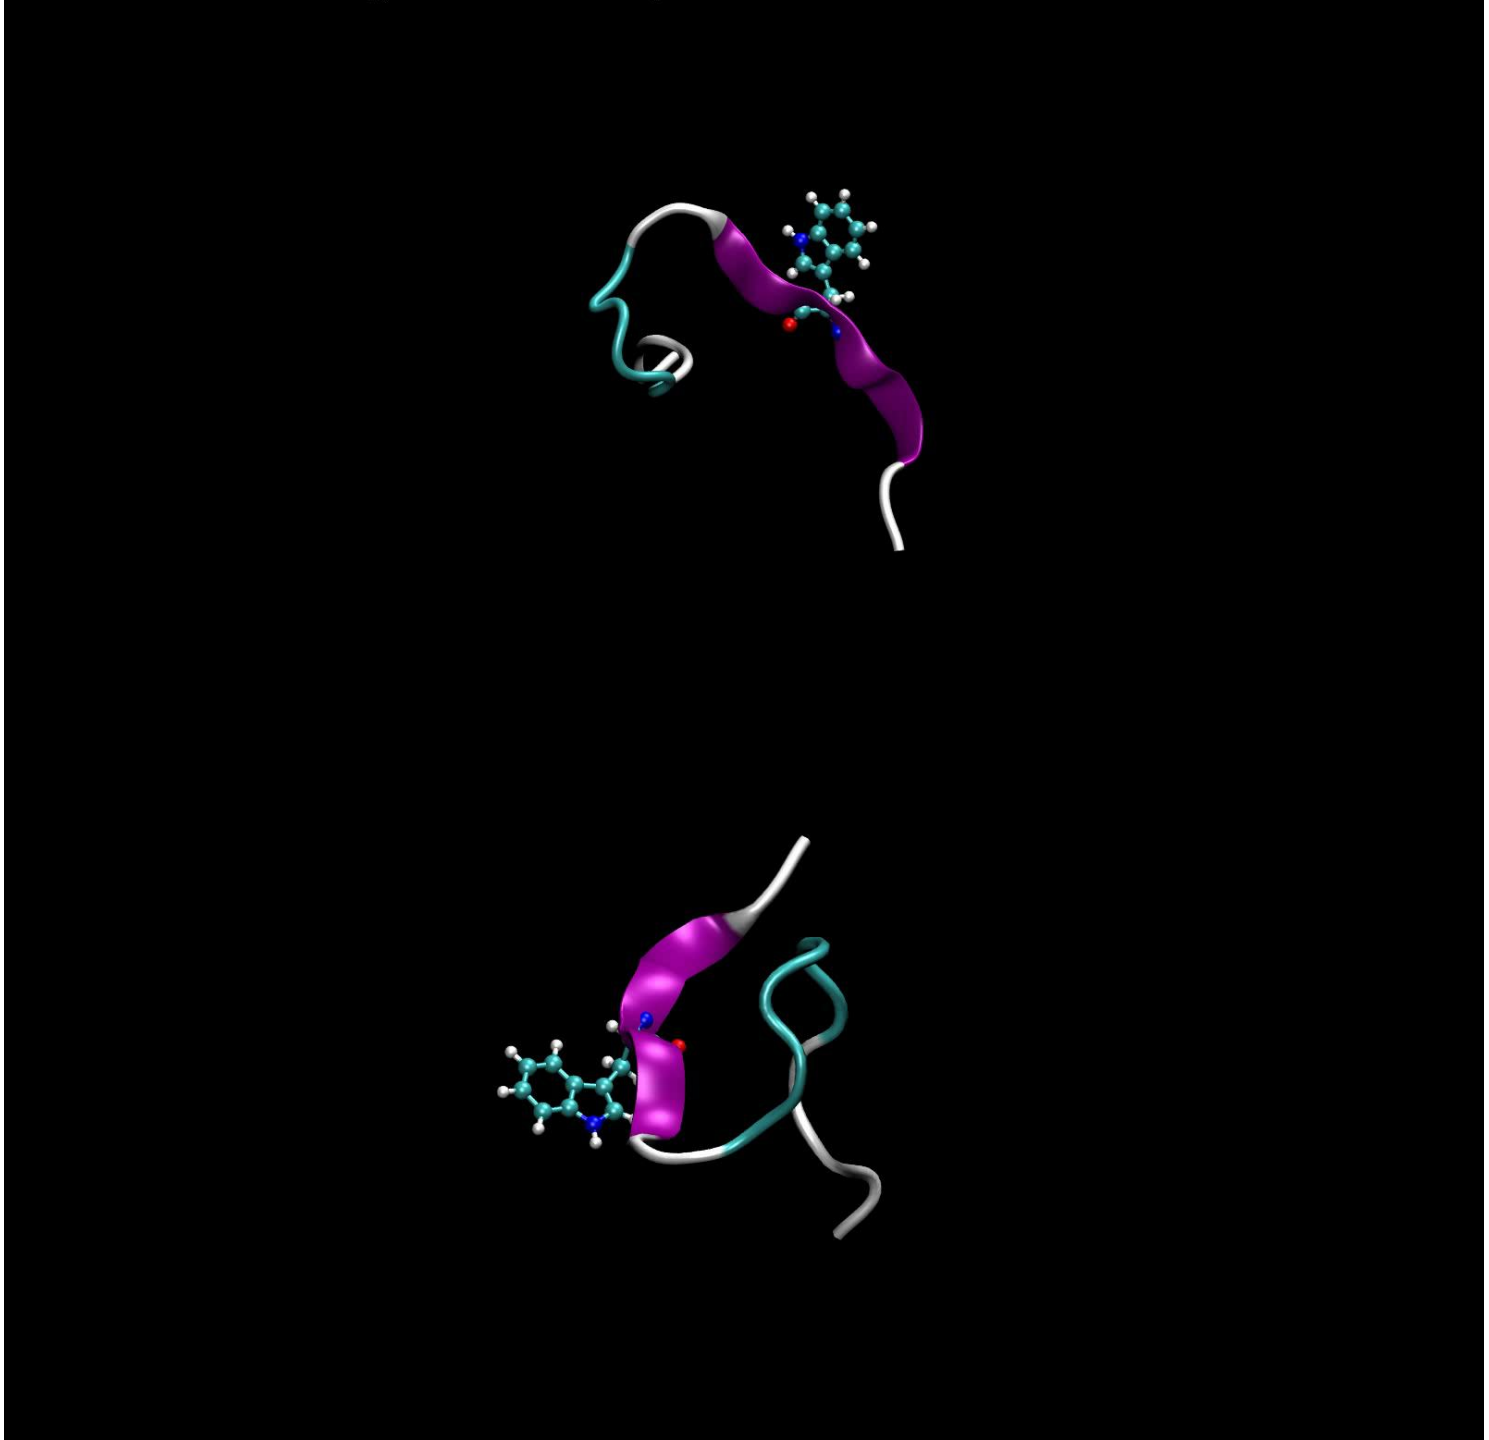

Interactive MD?

- **Single nodes are now fast enough, GPU enabled cloud nodes actually make sense as a back end now.**

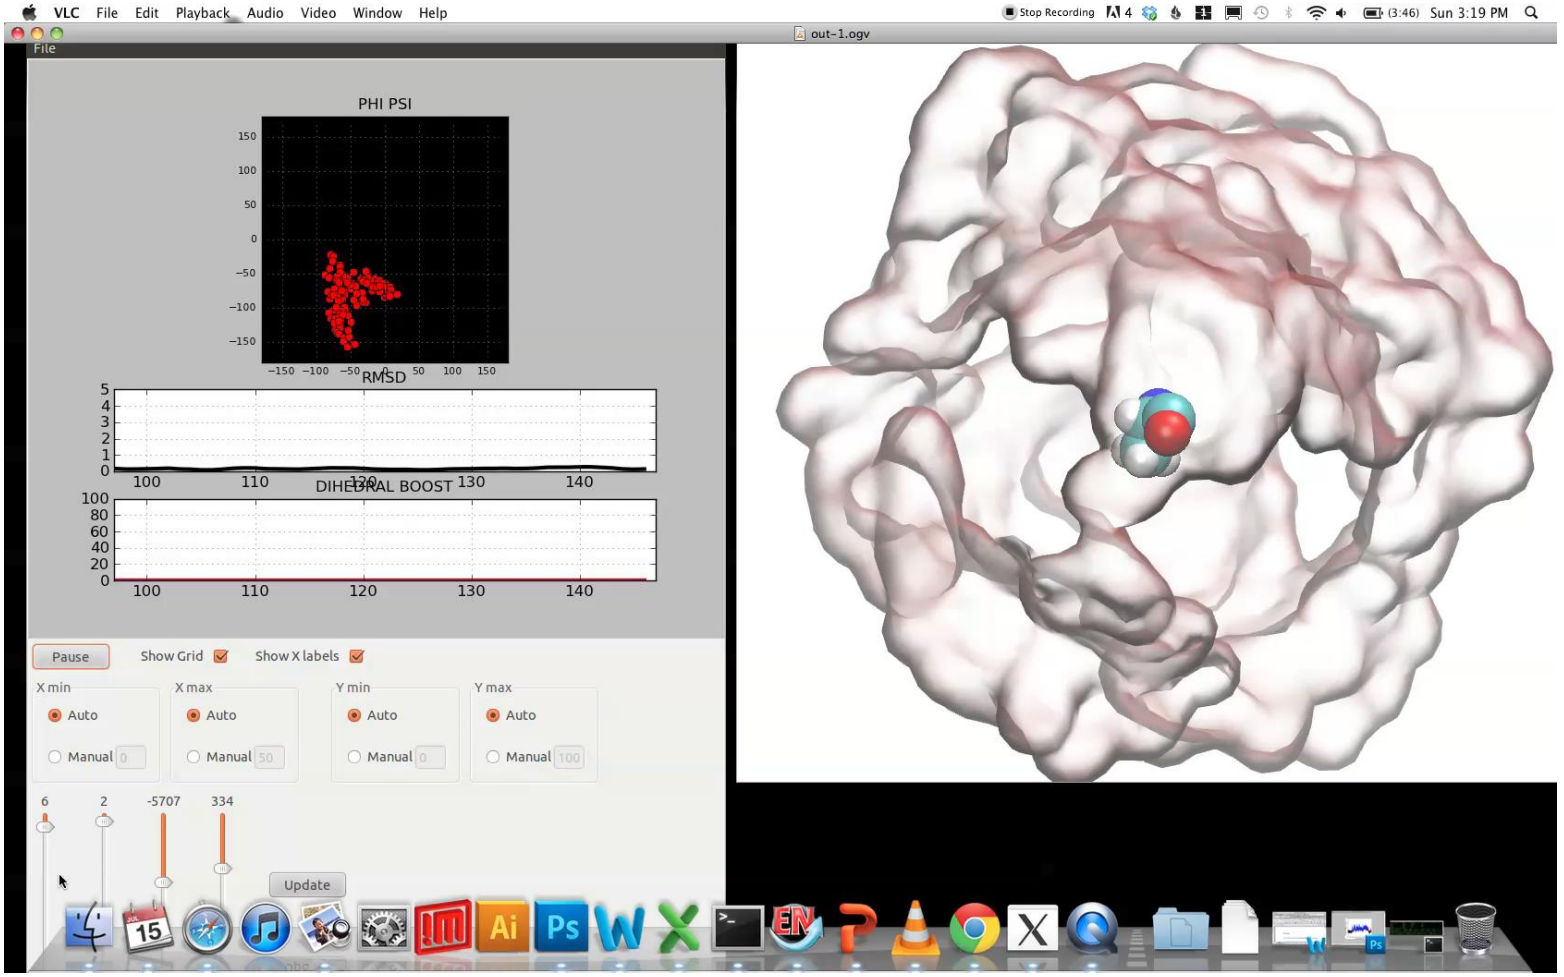

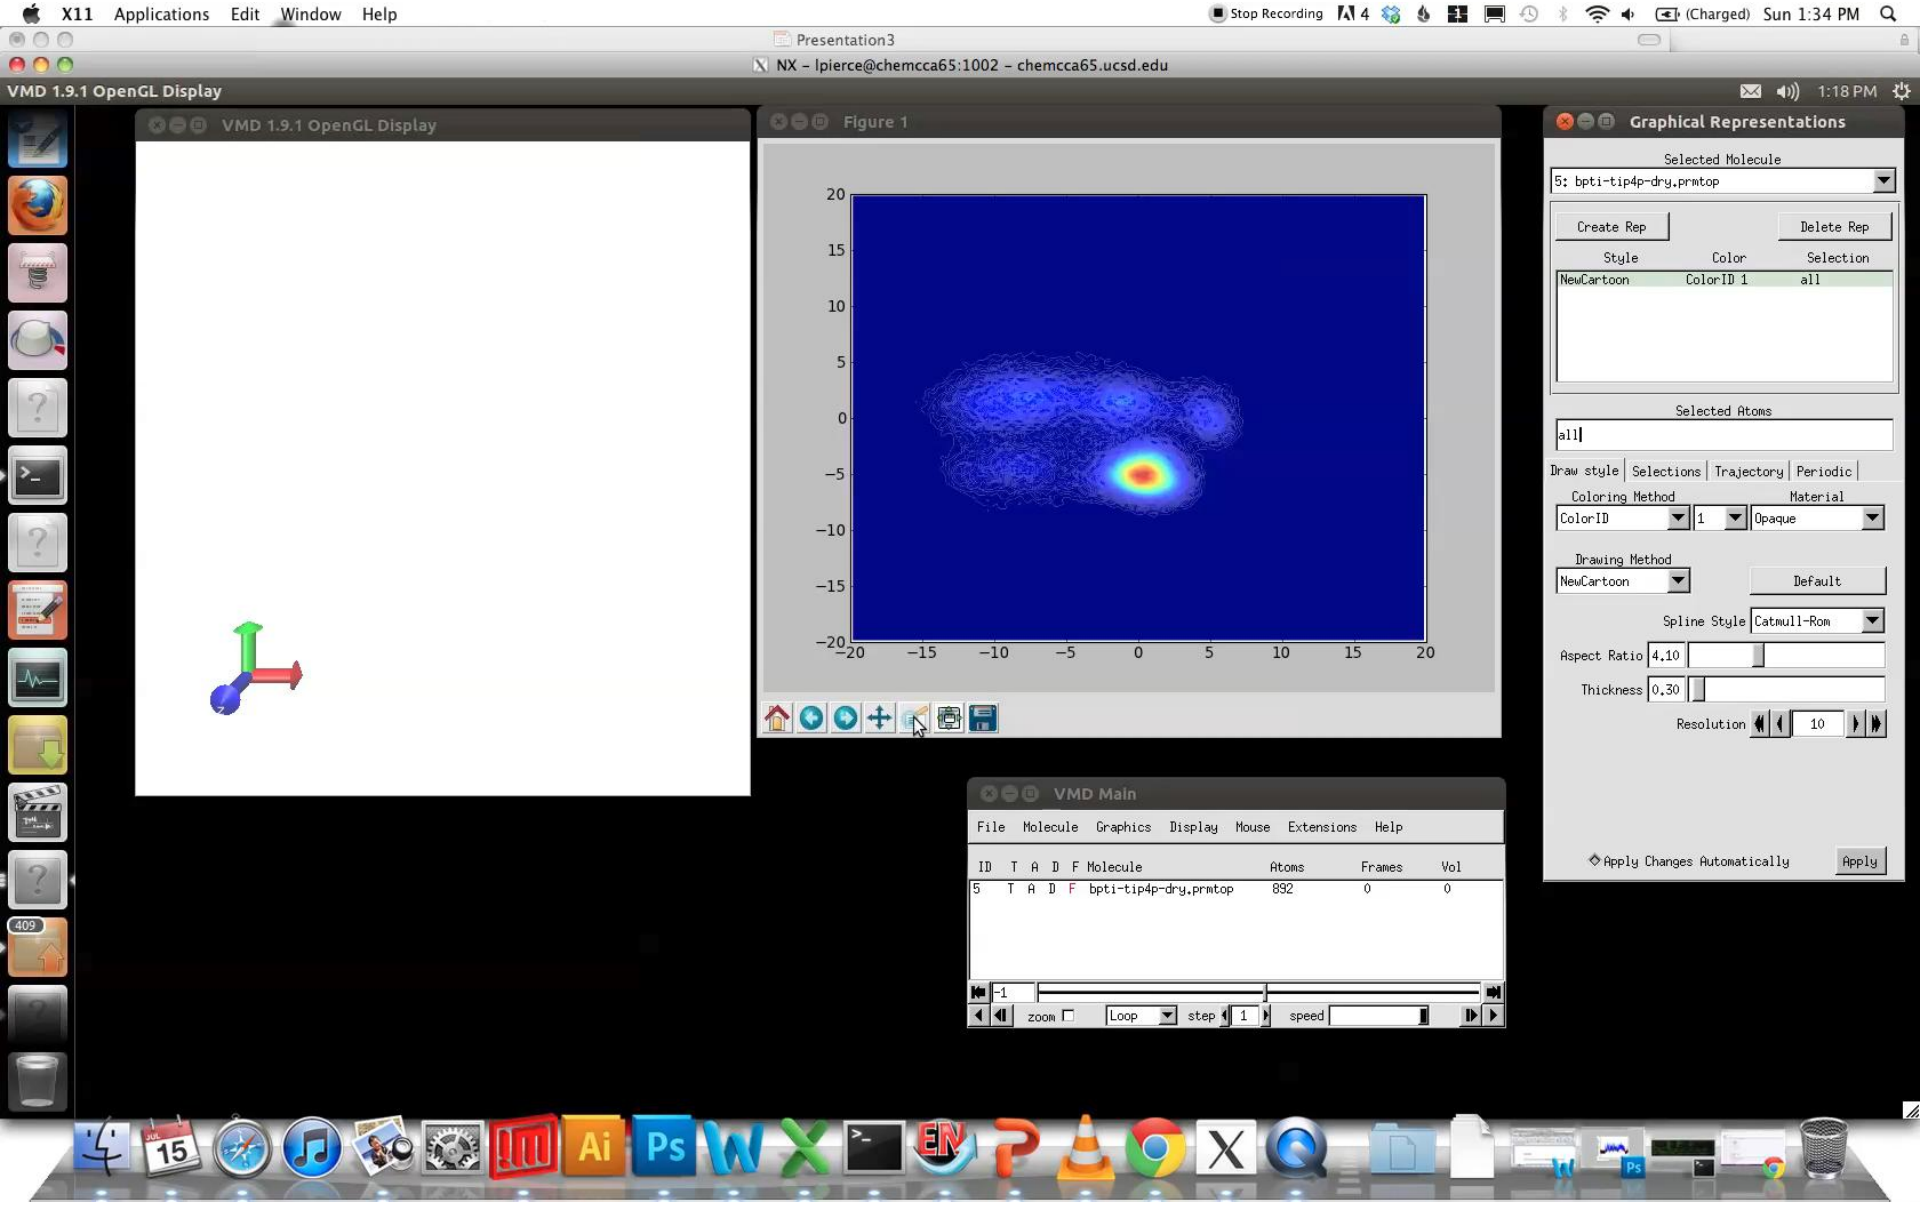
X11 Applications Edit Window Help
Stop Recording
(Charged) Sun 1:34 PM
Presentation3
NX – lpierce@chemcca65:1002 – chemcca65.ucsd.edu
VMD 1.9.1 OpenGL Display
1:18 PM
VMD 1.9.1 OpenGL Display
Figure 1
Graphical Representations
Selected Molecule
5: bpti-tip4p-dry.prmtop
Create Rep
Delete Rep
Style
Color
Selection
NewCartoon
ColorID 1
all
Selected Atoms
all
Draw style
Selections
Trajectory
Periodic
Coloring Method
Material
ColorID
1
Opaque
Drawing Method
NewCartoon
Default
Spline Style
Catmull-Rom
Aspect Ratio
4.10
Thickness
0.30
Resolution
10
Apply Changes Automatically
Apply
VMD Main
File
Molecule
Graphics
Display
Mouse
Extensions
Help
ID T A D F Molecule
Atoms
Frames
Vol
5 T A D F bpti-tip4p-dry.prmtop
892
0
0
zoom
Loop
step
1
speed

Recommended Hardware

Supported GPUs (examples, not exhaustive)

Hardware Version 3.5

- Kepler K20 / K20X
- GTX Titan

Hardware Version 3.0

- Kepler K10
- GTX670 / 680 / 690

Hardware Version 2.0

- Tesla M2050/M2070/M2075/M2090 (and C variants)
- GTX560 / 570 / 580 / 590
- GTX465 / 470 / 480

C1060/S1070 / GTX280 etc also supported

Graphics card

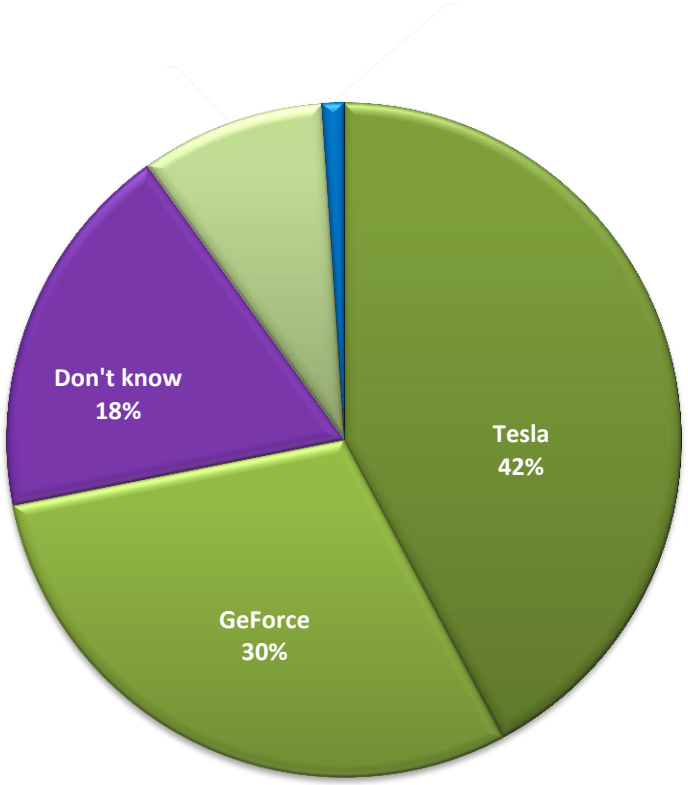

For the research group you belong to, please write in the specific card models if you know them.

Recommended Hardware

- **See the following page for continuous updates:**

http://ambermd.org/gpus/recommended_hardware.htm#hardware

AMBER MD Workstations		AMBER MD SimCluster
Kepler K10 MD Workstation	Kepler GTX680 MD Workstation	Kepler MD SimCluster (K20, K10 or GTX680)
		
• 2 Intel Xeon E5-2660 8-Core CPUs • 4 Tesla K10 8GB GPUs (8 x 4GB GPUS!) • 64 GB 1600MHz system memory • AMBER12 preinstalled, tested and optimized • CentOS • 3 year warranty	• 2 Intel Xeon E5-2660 8-Core CPUs • 4 GeForce GTX 680 4GB GPUs • 64 GB 1600MHz system memory • AMBER12 preinstalled, tested and optimized • CentOS • 3 year warranty	• 8U cluster (4 x 2U) • 2 Intel Xeon E5-2660 8-Core CPUs per node • 4 Tesla Kepler GPUs (K20, K10 or GTX680) per node • 64 GB 1600MHz system memory per node • AMBER12 preinstalled, tested and optimized • Optional InfiniBand QDR Switch & HCAs • Bright Computing Cluster Manager • full warranty (even with GTX680s)
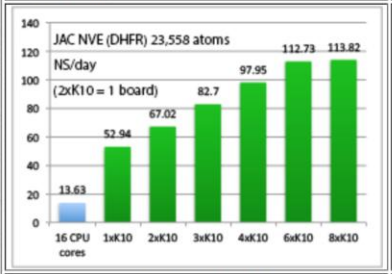	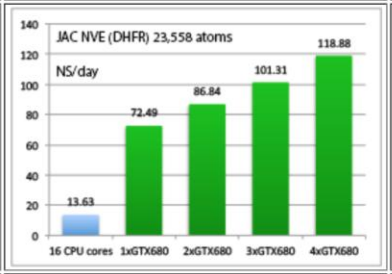	
($5000 - $18000) *As shown ~$15,400*	*($3000 - $8000)* *As shown ~$7,600*	*Price varies*

DIY 4 GPU System

Antec P280 Black ATX Mid Tower Case
http://tinyurl.com/a2wtkfr
$126.46

SILVERSTONE ST1500 1500W
ATX12V/EPS12V Power Supply
http://tinyurl.com/alj9w93
$299.99

AMD FX-8350 Eight Core CPU
http://tinyurl.com/b9teunj
$189.15

Corsair Vengeance 16GB
(2x8GB) DDR3 1600 MHz
Desktop Memory
http://tinyurl.com/amh4jyu
$96.87

GIGABYTE GA-990FXA-UD7
AM3+ AMD Motherboard
http://tinyurl.com/b8yvykv
$216.00

Seagate Barracuda 7200 3 TB
7200RPM Internal Bare Drive
ST3000DM001
http://tinyurl.com/a4ccfvj
$139.99

4of EVGA GeForce GTX 680
4096 MB GDDR5
http://tinyurl.com/d82lq8d
$534.59 each
Or K20X or GTX Titan

Total Price : $3206.82

Note cards in this system run at x8 so you can only run single GPU AMBER runs (but you can run 4 simultaneously at full speed) – If you want to be able to run MPI 2xGPU runs then only place 2 cards in the x16 slots.

Single Workstation

Based on Exxact Model Quantum TXR410-512R
(Available as Amber MD Workstation with AMBER 12 preinstalled)

(A) SuperServer Tower / 4U Convertible Chassis, Supports Up To 3x 5.25 Inch Bays, 8x 3.5 Inch Hot-Swap HDD Bays, Up To 4x Double-Width GPU, 1620W Redundant Power Supplies
(B) SuperServer Intel Patsburg Based Motherboard, Supports Up To 2x Sandy Bridge EP (Socket R) Series CPU, 2x 10/100/1000 NIC, Dedicated IPMI Port, 4x PCIE 3.0 x16 Slots, 2x PCIE 3.0 x8 Slots, Up To 512GB DDR3 1600MHz ECC/REG Memory
(C) Intel Xeon E5-2620 2.00 Ghz 15MB Cache 7.20GT/sec LGA 2011 6-Core Processor (2)
(D) Certified 4GB 240-Pin DDR3 SDRAM ECC Registered DDR3 1600 MHz Server Memory (8)
(E) Certified 2TB 7200RPM 64MB Cache 3.5 Inch SATA Enterprise Class HDD in a RAID 1 Configuration (2)
(G) GeForce GTX 680 4GB or GTX Titan or K20X 6GB 384-bit GDDR5 PCI Express 3.0 Accelerator (4)
(H) CentOS 6

Price ~ $6,500 (GTX 680) $8,500 (GTX Titan)
$20,000 (K20X)

Clusters

# of CPU sockets	2
Cores per CPU socket	4+ (1 CPU core drives 1 GPU)
CPU speed (Ghz)	2.0+
System memory per node (GB)	16 to 32
GPUs	Kepler K10, K20, K20X Fermi M2090, M2075, C2075
# of GPUs per CPU socket	1-2 **(4 GPUs on 1 socket is good to do 4 fast serial GPU runs)**
GPU memory preference (GB)	6
GPU to CPU connection	PCIe 2.0 16x or higher
Server storage	2 TB+
Network configuration	**Infiniband QDR or better (optional)**

Scale to multiple nodes with same single node configuration

Acknowledgements

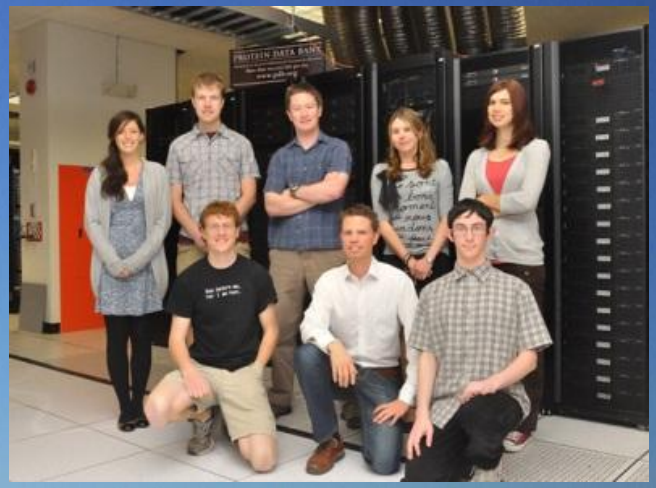

San Diego Supercomputer Center

University of California San Diego

National Science Foundation

NSF Strategic Applications Collaboration (AUS/ASTA) Program

NSF SI2-SSE Program

NVIDIA Corporation

Hardware + People

People

Romelia Salomon

Mike Wu

Robin Betz

Duncan Poole

Andreas Goetz

Matthew Clark

Jason Swails

Mark Berger

Scott Le Grand

Romelia Salomon

Ben Madej

Sarah Tariq

AMBER User Survey - 2011
GPU Momentum is Growing!

AMBER machines

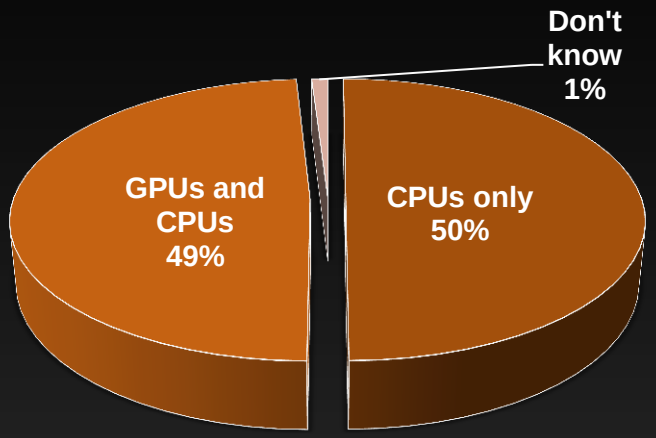

GPU Experience

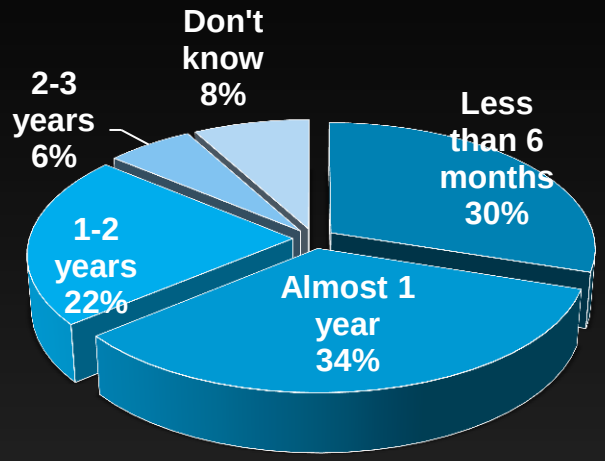

- 49% of AMBER machines have GPUs
- 85% of users have up to 2 years of GPU experience

Testimonials

> “The whole lab loves the GPU cluster. Students are now able to run AMBER simulations that would not have been feasible on our local CPU-based resources before. Research throughput the group has been enhanced significantly.”

Jodi Hadden
Chemistry Graduate Student
Woods Computing Lab
Complex Carbohydrate Research Center
University of Georgia

The University of Georgia

GPU Accelerated Apps Momentum

Key codes are GPU Accelerated!

Molecular Dynamics

- Abalone – GPU only code
- ACEMD – GPU only code
- AMBER
- CHARMM
- DL_POLY
- GROMACS
- HOOMD-Blue – GPU only code
- LAMMPS
- NAMD

Quantum Chemistry

- ABINIT
- BigDFT
- CP2K
- GAMESS
- Gaussian – in development
- NWChem – in development
- Quantum Espresso
- TeraChem – GPU only code
- VASP

Check many more apps at www.nvidia.com/teslaapps

Test Drive K20 GPUs!

Experience The Acceleration

- Run AMBER on Tesla K20 GPU today
- Sign up for FREE GPU Test Drive on remotely hosted clusters
  www.nvidia.com/GPUTestDrive

Test Drive K20 GPUs!

Experience The Acceleration

- Run AMBER on Tesla K20 GPU today
- Sign up for FREE GPU Test Drive on remotely hosted clusters
  www.nvidia.com/GPUTestDrive

Registration is Open!

March 18-21, 2013 | San Jose, CA

- **Four days**
- **Three keynotes**
- **400+ sessions**
- **One day of pre-conference developer tutorials**
- **150+ research posters**
- **Lots of networking events and opportunities**

Visit www.gputechconf.com for more info.