Transforming research in molecular biology through extreme acceleration of AMBER molecular dynamics simulations

Smart sampling for the 99%.

By Ross C. Walker

SDSC SAN DIEGO SUPERCOMPUTER CENTER

Walker Molecular Dynamics Lab

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

GPU Acceleration	Lipid Force Field Development	QM/MM MD	Biofuels

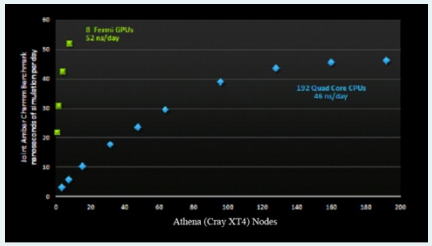

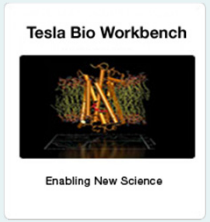

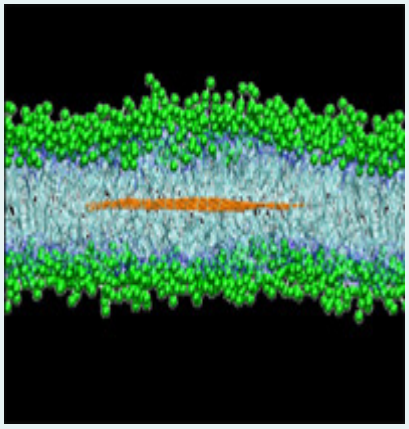

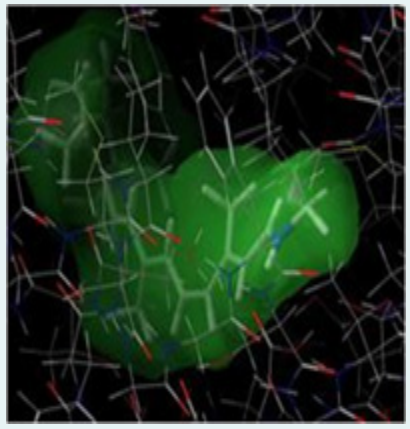

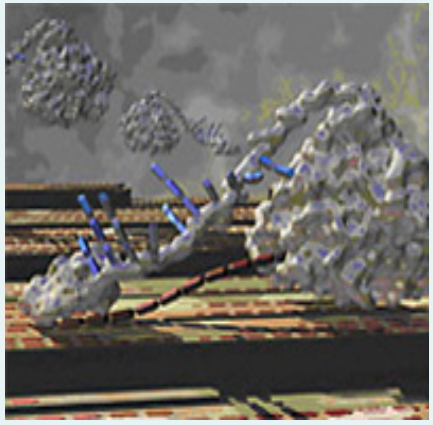

Postdocs: Andreas Goetz, Romelia Salomon

Graduate Students: Ben Madej (UCSD/SDSC), Jason Swails (UFL), Sarah Rosen (Imperial College), Age Skjevik (Bergen)

Undergraduate Researchers: Robin Betz, Matthew Clark

What is a GPU?

- **Graphics Processing Unit**

Tseng Labs ET4000/W32p 1991

S3 Graphics ViRGE 1992

Voodoo3 2000 AGP card 1999

- ‘Specialist’ processor for accelerating the rendering of computer graphics.
- Invented by 3DFX (Later bought by NVIDIA) in 1997.
- Originally fixed function pipelines
  - Invention of OpenGL added programmability.
  - Pixels can be programmed with specific textures.
  - Onboard memory for storing textures.
- Development driven by $150 billion gaming industry.

GeForce 6600 GT Personal Cinema 2004

NVIDIA GeForce GTX 280 2008

Why Interest GPUs?

Peak Flops

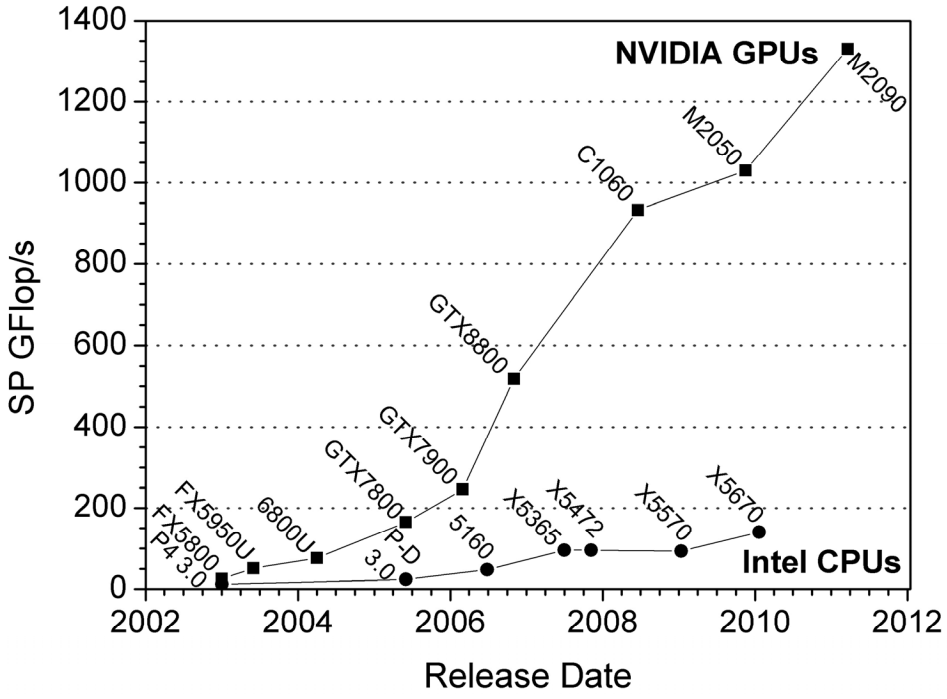

Memory Bandwidth

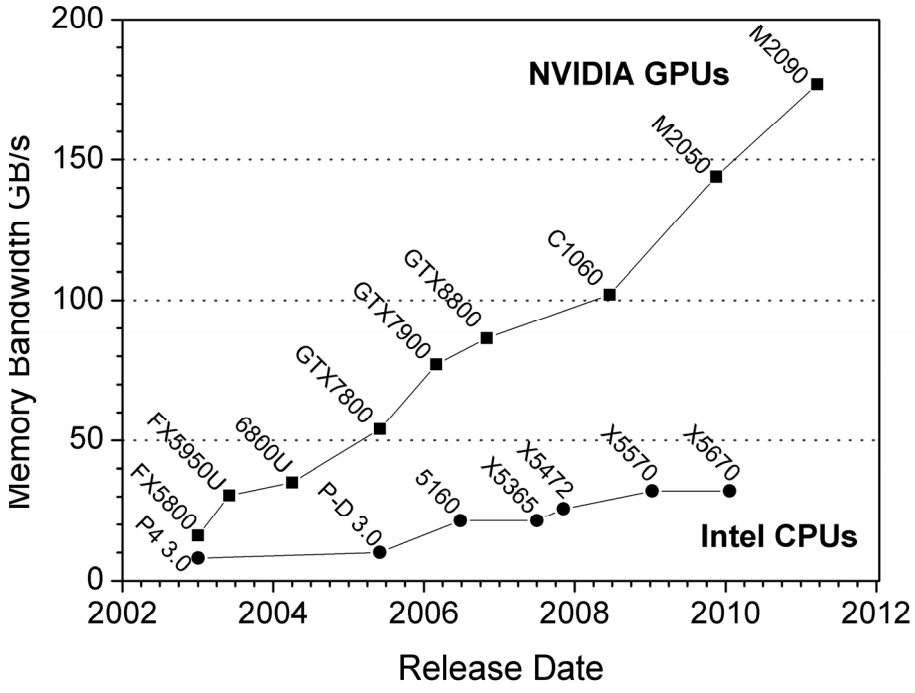

What's the Catch?

GPUs Require Rewriting Lots of Code (Is it worth it?)

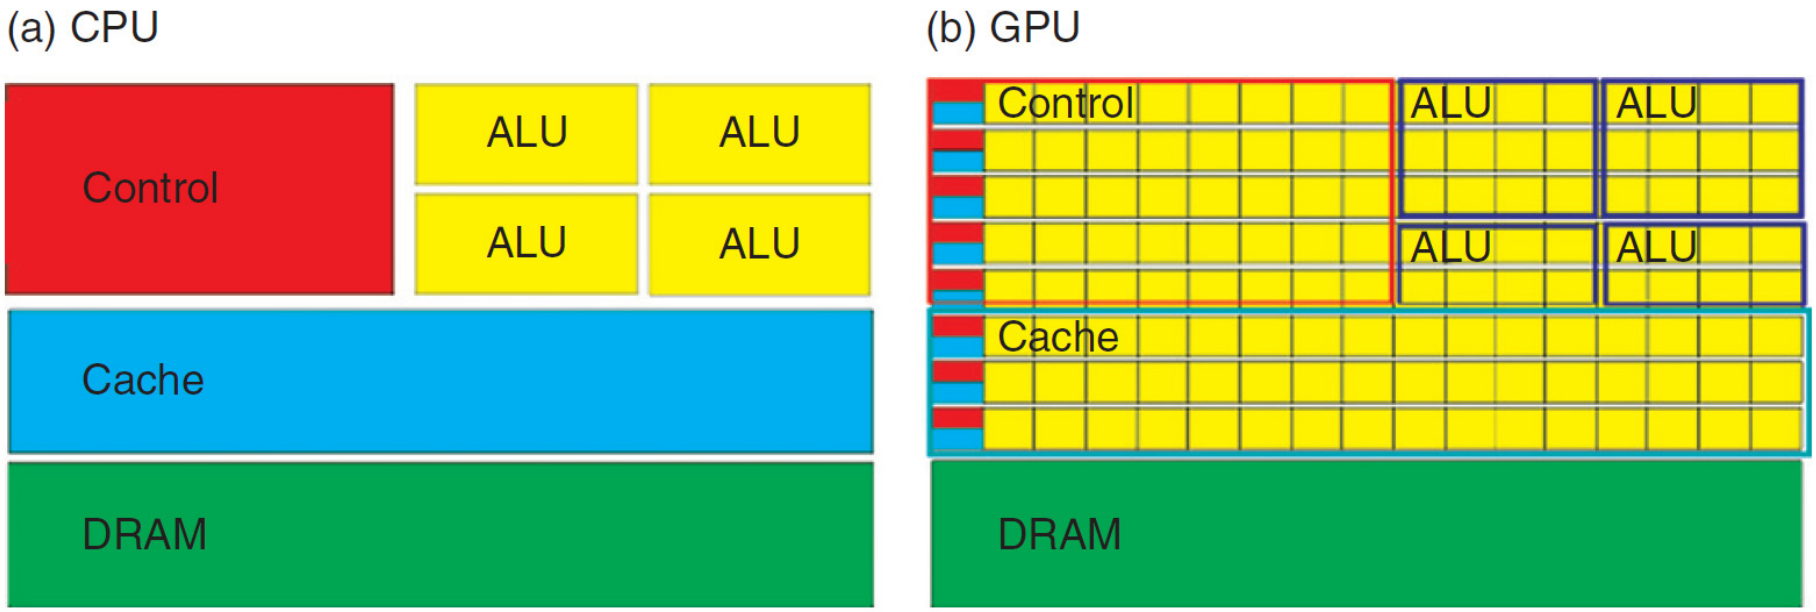

Spoiler: Yes. But LOTS of work.

‘Bio’ GPU Codes

http://www.nvidia.com/object/tesla_bio_workbench.html

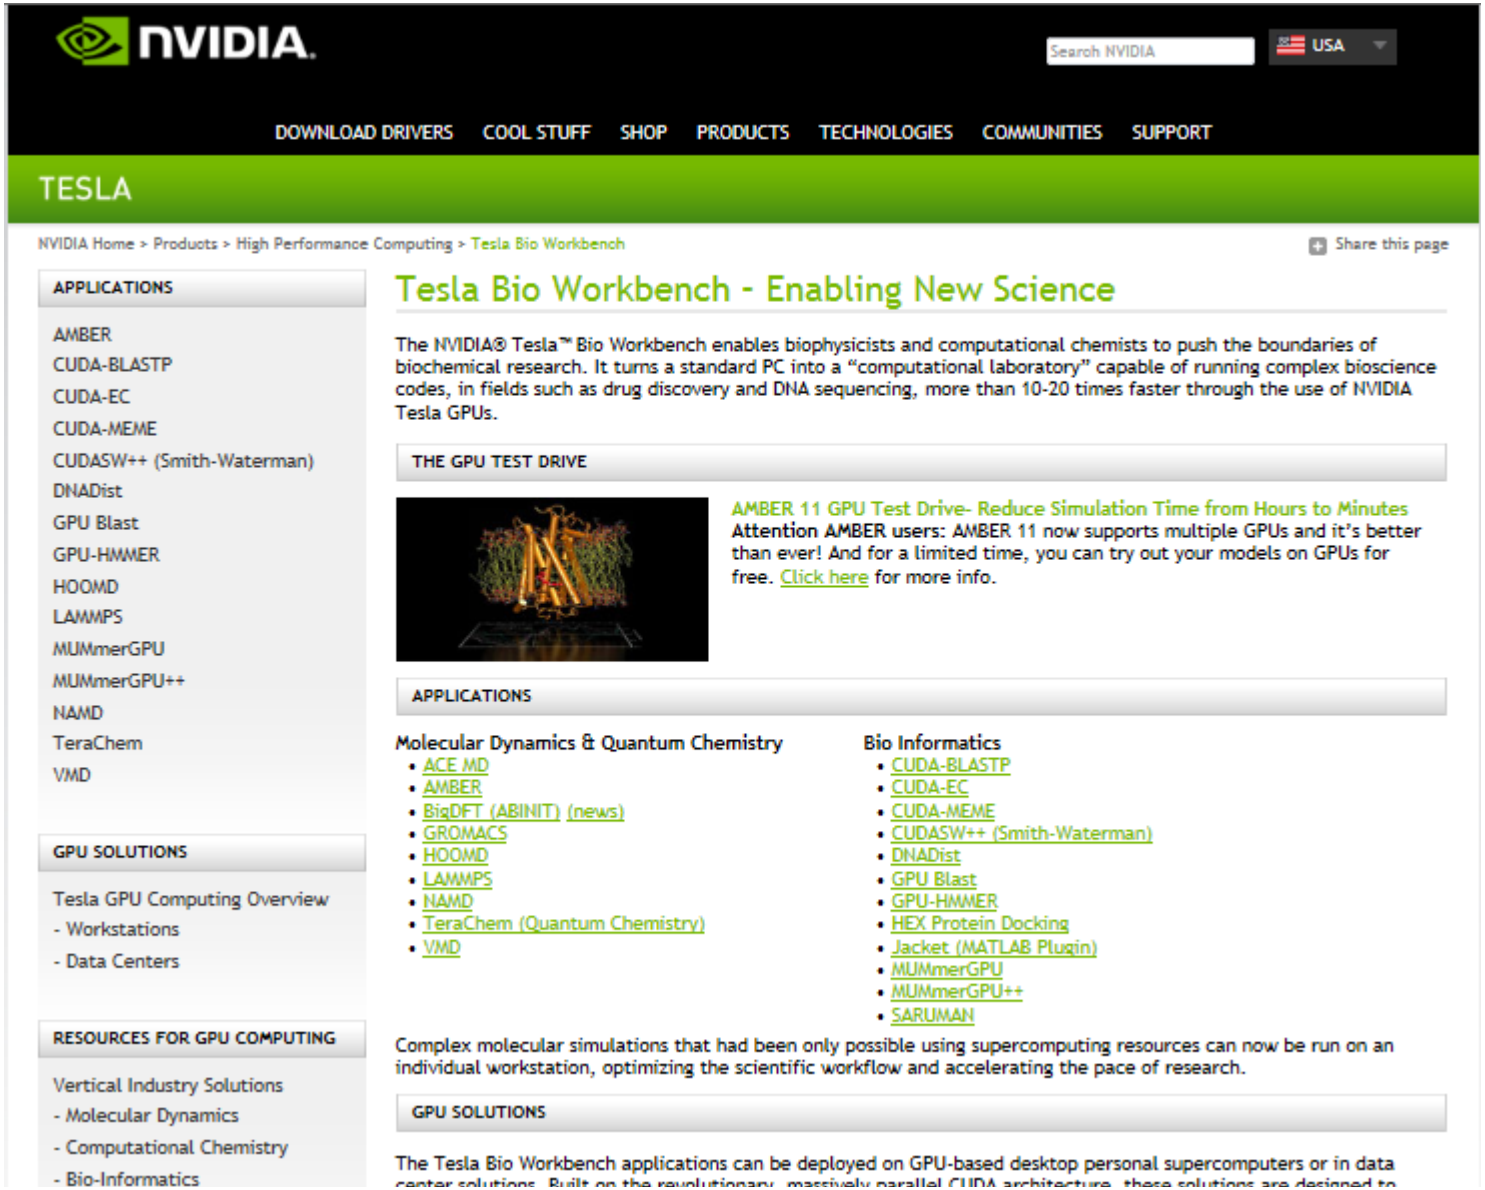

Review (Blatant Plug)

- Xu, D., Williamson, M.J., Walker, R.C., "***Advancements in Molecular Dynamics Simulations of Biomolecules on Graphical Processing Units.***", *Ann. Rep. Comp. Chem.*, vol 6. 2010, 2-19 (Chapter 1).

- Goetz, A.W., Woelfe, T., Walker, R.C., "***Quantum Chemistry on Graphics Processing Units***", *Ann. Rep. Comp. Chem.*, vol 6. 2010, 21-35 (Chapter 2).

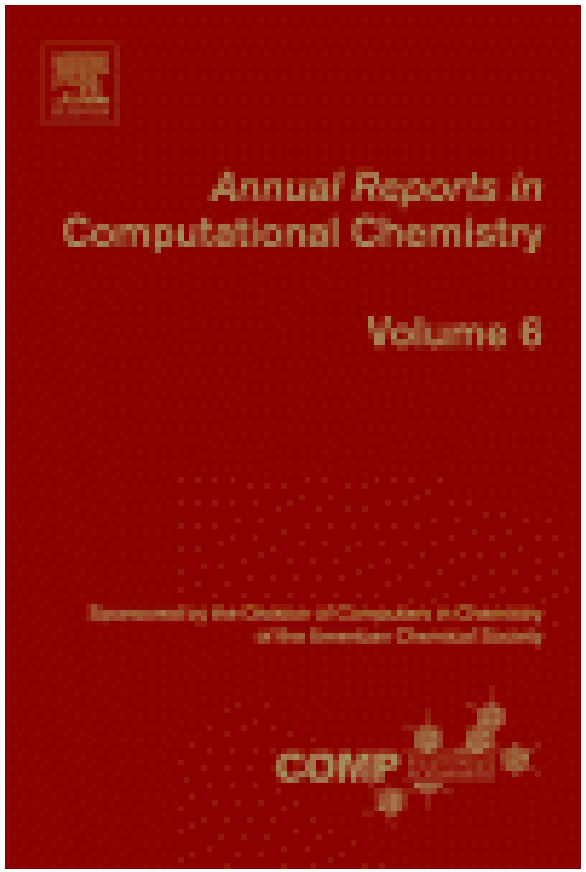

Transforming research in molecular biology through extreme acceleration of AMBER molecular dynamics simulations

The Project

- **A collaboration between NVIDIA and the AMBER Development Team.**

San Diego Supercomputer Center

Ross C. Walker

NVIDIA

Scott Le Grand

Duncan Poole

Funded as a pilot project (1 year) under NSF SSE Program

SDSC
SAN DIEGO SUPERCOMPUTER CENTER

UCSD

NVIDIA

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)

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.

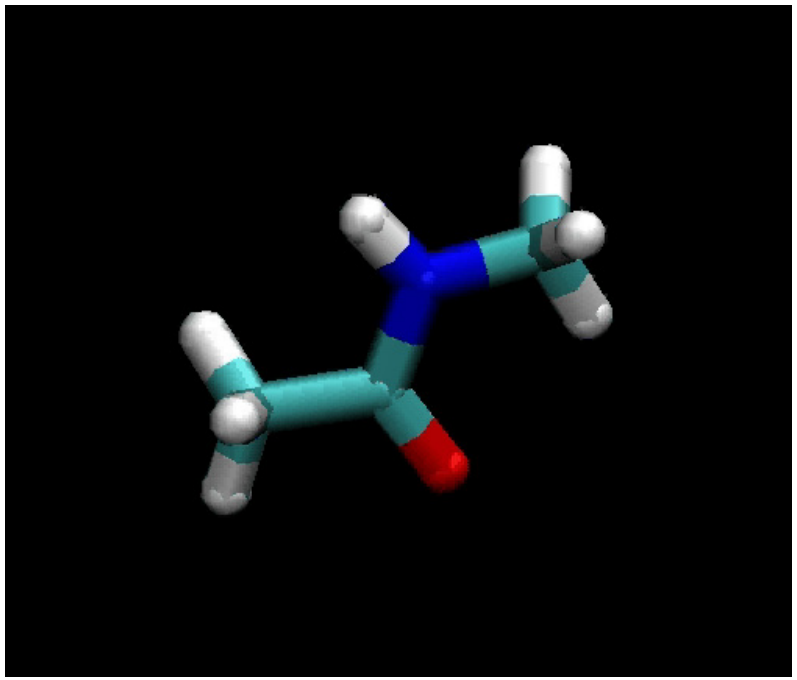

Why Molecular Dynamics?

- **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

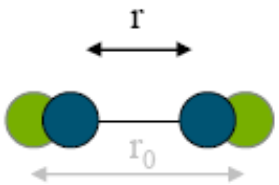

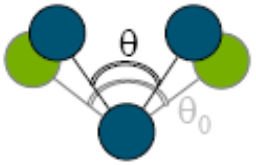

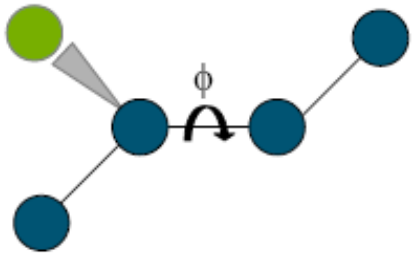

$$E_{stretching} = \sum_{1,2\,pairs} K_r (r - r_0)^2 \qquad E_{bending} = \sum_{1,2\,pairs} K_\theta (\theta - \theta_0)^2 \qquad E_{torsion} = \sum_{1,4\,pairs} K_\phi (1 - \cos(n\phi - \delta))$$

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

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

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

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

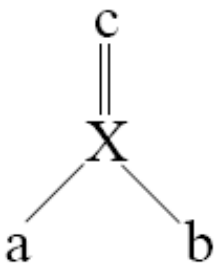

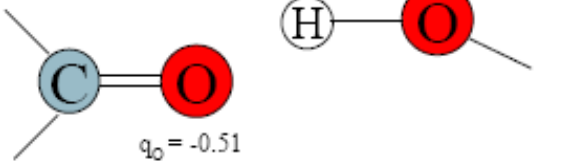

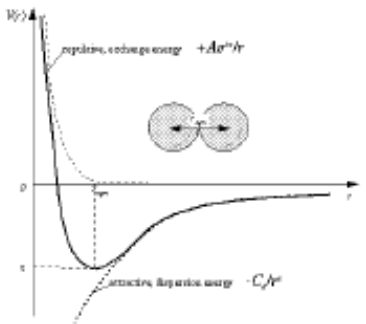

What is AMBER?

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

experimental polarizable forcefields e.g. ff02EP

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

in the public domain

AMBER Usage

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

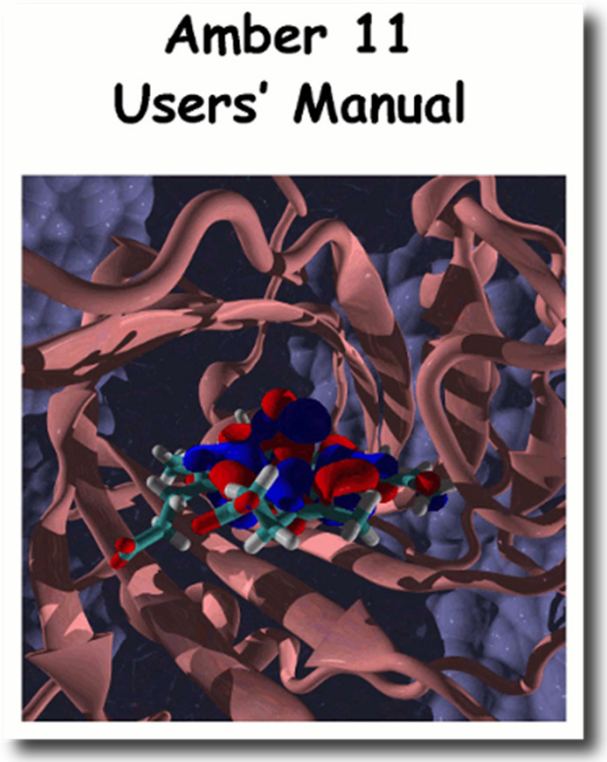

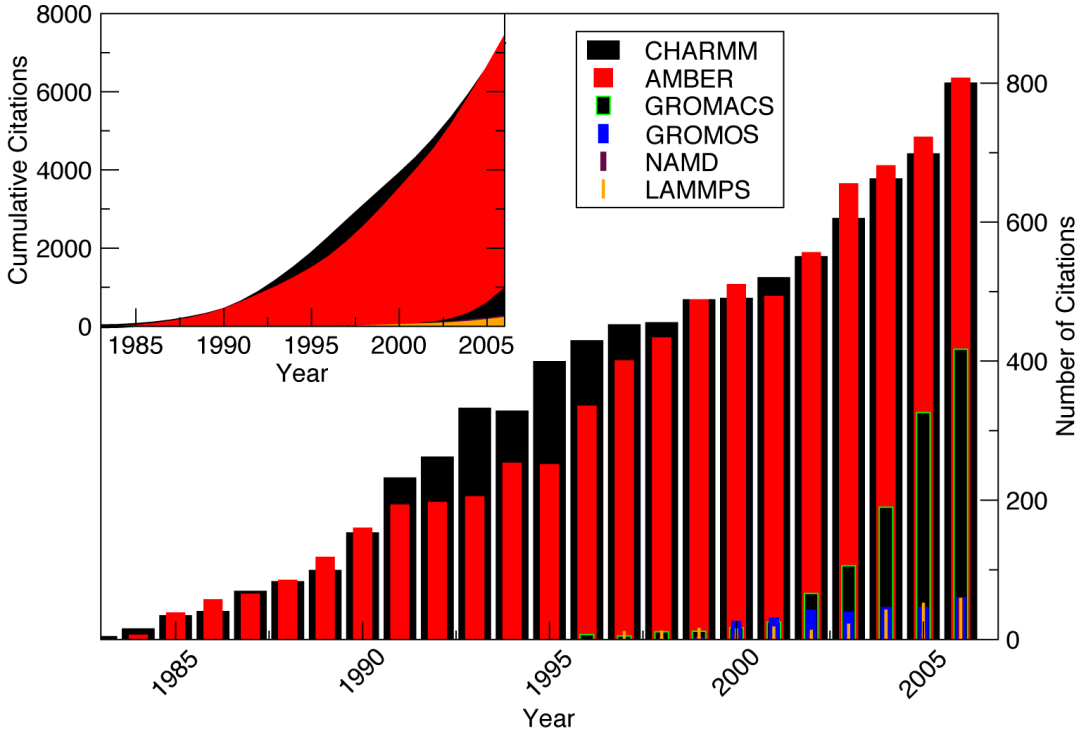

What can we do with Molecular Dynamics?

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

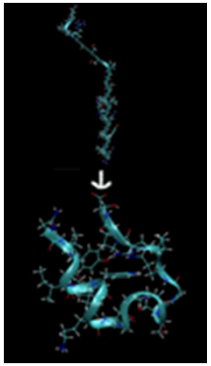

- **Can simulate ensemble properties.**
  - Binding free energies.
    - Drug Design
    - Biocatalyst Design
  - Reaction Pathways
  - Free Energy Surfaces.

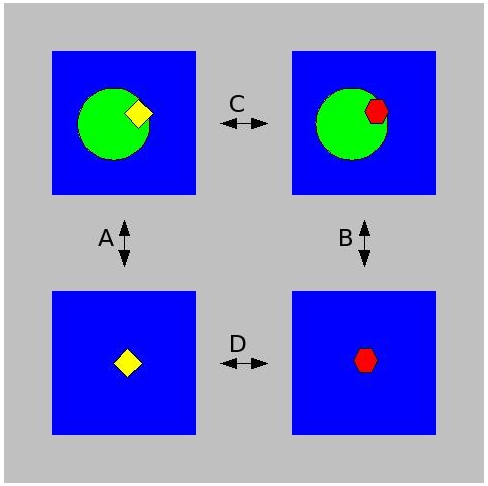

Why do we need Supercomputers? (Complex Equations)

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

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

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

Why do we need Supercomputers?
Lots of Atoms

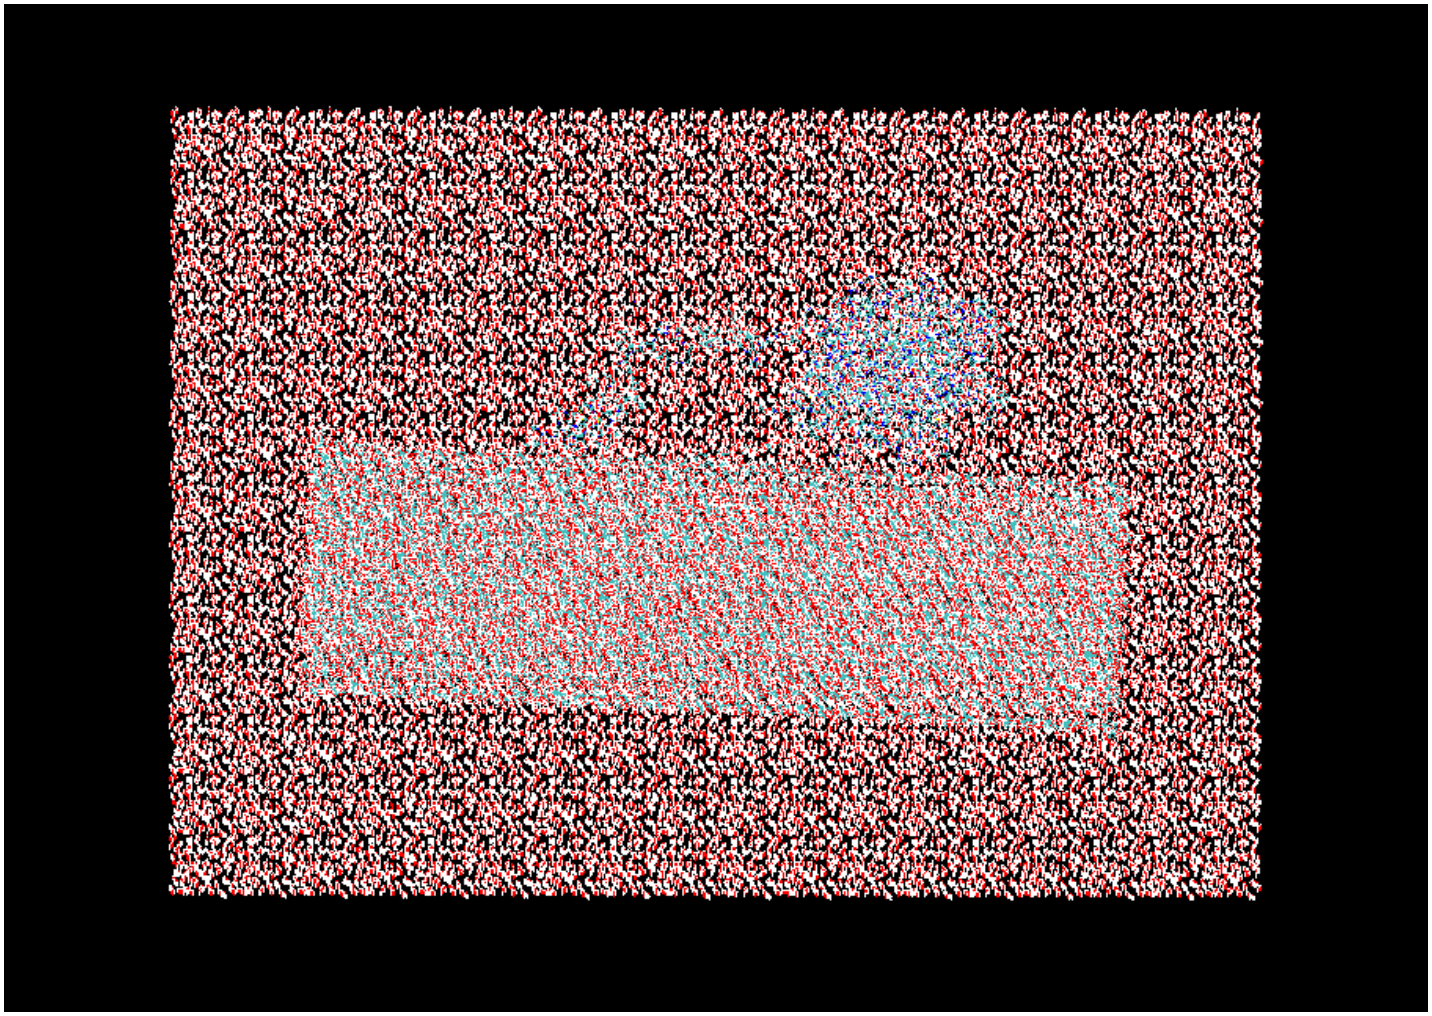

Why do we need Supercomputers?
Lots of Time Steps

- ***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!!!

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).

Better Science?

- **Bringing the tools the researcher needs into his own lab.**
  - Can we make a researcher's desktop look like a small cluster?
  - 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).

GPU Support in AMBER

GPU Support

- **Collaboration with NVIDIA to produce CUDA version of AMBER.**
  - PMEMD Engine
  - Implicit Solvent GB
  - Explicit Solvent PME

NVIDIA

- **Serial version released as part of AMBER v11.**
- **Parallel version is available as a patch against the released version.**

JCTC Journal of Chemical Theory and Computation

Article

pubs.acs.org/JCTC

Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 1. Generalized Born

Andreas W. Götz,[†] Mark J. Williamson,[†,∥] Dong Xu,[†,⊥] Duncan Poole,[‡] Scott Le Grand,[‡] and Ross C. Walker*[,†,§]

[†]San Diego Supercomputer Center, University of California San Diego, 9500 Gilman Drive MC0505, La Jolla, California 92093, United States

[‡]NVIDIA Corporation, 2701 San Tomas Expressway, Santa Clara, California 95050, United States

[§]Department of Chemistry and Biochemistry, University of California San Diego, 9500 Gilman Drive MC0505, La Jolla, California 92093, United States

Supporting Information

ABSTRACT: We present an implementation of generalized Born implicit solvent all-atom classical molecular dynamics (MD) within the AMBER program package that runs entirely on CUDA enabled NVIDIA graphics processing units (GPUs). We discuss the algorithms that are used to exploit the processing power of the GPUs and show the performance that can be achieved in comparison to simulations on conventional CPU clusters. The implementation supports three different precision models in which the contributions to the forces are calculated in single precision floating point arithmetic but accumulated in double precision (SPDP), or everything is computed in single precision (SPSP) or double precision (DPDP). In addition to performance, we have focused on understanding the implications of the different precision models on the outcome of implicit solvent MD simulations. We show results for a range of tests including the accuracy of single point force evaluations and energy conservation as well as structural properties pertaining to protein dynamics. The numerical noise due to rounding errors within the SPSP precision model is sufficiently large to lead to an accumulation of errors which can result in unphysical trajectories for long time scale simulations. We recommend the use of the mixed-precision SPDP model since the numerical results obtained are comparable with those of the full double precision DPDP model and the reference double precision CPU implementation but at significantly reduced computational cost. Our implementation provides performance for GB simulations on a single desktop that is on par with, and in some cases exceeds, that of traditional supercomputers.

1. INTRODUCTION

Since the first simulation of an enzyme using molecular dynamics (MD) was reported by McCammon et al.[1] in 1977, MD simulations have evolved to become important tools in rationalizing the behavior of biomolecules. The field has grown from that first 10-ps-long simulation of a mere 500 atoms to the point where small enzymes can be simulated on the microsecond time scale[2−4] and simulations containing millions of atoms can be considered routine.[5,6] However, such simulations are numerically very intensive, and using traditional CPU-centric hardware requires access to large-scale supercomputers or well-designed clusters with expensive interconnects that are beyond the reach of many research groups.

Numerous attempts have been made over the years to accelerate classical MD simulations by exploiting alternative hardware technologies. Some notable examples include ATOMS by AT&T Bell Laboratories,[7] FASTRUN by Columbia University and Brookhaven National Laboratory,[8] MDGRAPE by RIKEN,[9] and most recently Anton by DE Shaw Research LLC.[10] All of these approaches have, however, failed to make an impact on mainstream research because of their excessive cost. Additionally, these technologies have been based on custom hardware and do not form part of what would be considered a standard workstation specification. This has made it difficult to experiment with such technologies, leading to a lack of sustained development or innovation and ultimately their failure to mature into ubiquitous community-maintained research tools.

Graphics processing units (GPUs), on the other hand, have been an integral part of personal computers for decades, and a strong demand from the consumer electronics industry has resulted in significant sustained industrial investment in the stable, long-term development of GPU technology. In addition to low prices for GPUs, this has led to a continuous increase in the computational power and memory bandwidth of GPUs, significantly outstripping the improvements in CPUs. As a consequence, high-end GPUs can be considered standard equipment in scientific workstations, which means that they either already exist in many research laboratories or can be purchased easily with new equipment. This makes them readily available to researchers and thus attractive targets for acceleration of many scientific applications including MD simulations.

The nature of GPU hardware, however, has until recently made their use in general purpose computing challenging to all but those with extensive three-dimensional (3D) graphics programming experience. However, the development of application programming interfaces (APIs) targeted at general purpose scientific computing has reduced this complexity substantially such that

Received: December 20, 2011
Published: March 26, 2012

Goetz, A. W.; Williamson, M. J.; Xu, D.; Poole, D.; Grand, S.; 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

Supported Features

- **Supports 'standard' MD**
  - Explicit Solvent (PME)
    - NVE/NVT/NPT
  - Implicit Solvent (Generalized Born)
- **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

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.
  - Accuracy should be directly comparable to the standard CPU implementation.

Precision Models

- **Multiple codes have simply used single precision without any ‘REAL’ consideration of accuracy implications.**
  - Validation is now the ‘worst’ part of programming.
- **We have focused on accuracy first.**
  - Get the answers correct and validate!
  - Then improve performance.
- **We have implemented several precision models for testing.**

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*)**

DPDP **- Use double precision for the entire calculation**

Force Accuracy

Table 5: Deviations of forces (in kcal/(molÅ)) of the AMBER PMEMD GPU implementation using different precision models as compared to reference values obtained with the CPU implementation.

Precision model	TRPCage (304 atoms)	ubiquitin (1,231 atoms)	apo-myoglobin (2,492 atoms)	nucleosome (25,095 atoms)
max deviation				
SPSP	3.0×10^{-3}	4.8×10^{-3}	4.2×10^{-3}	2.7×10^{-2}
SPDP	5.6×10^{-5}	3.7×10^{-4}	1.6×10^{-4}	1.1×10^{-3}
DPDP	1.1×10^{-8}	7.3×10^{-8}	3.4×10^{-8}	8.0×10^{-8}
RMS deviation				
SPSP	5.0×10^{-4}	6.1×10^{-4}	4.1×10^{-4}	1.5×10^{-3}
SPDP	7.0×10^{-6}	1.5×10^{-5}	8.1×10^{-6}	3.0×10^{-5}
DPDP	1.5×10^{-9}	3.6×10^{-9}	2.6×10^{-9}	3.2×10^{-9}

Energy Conservation: Implicit Solvent

(kT/ns/d.o.f)

UBIQUITIN GB	dt = 0.5fs	dt = 1.0fs
CPU	-0.000008	-0.000835
DPDP	-0.000001	-0.000780
SPDP	-0.000008	-0.000631
SPSP	0.000589	0.001139
OpenMM	---	0.005411

0.000008 kT/ns/dof = 0.01798 kcal/mol/ns.

Energy Conservation: Explicit Solvent

(kT/ns/d.o.f)

DHFR	dt = 0.5fs	dt = 1.0fs	dt = 2.0fs*
CPU	0.000000	0.000001	-0.000047
DPDP	0.000007	0.000024	-0.000101
SPDP	-0.000052	0.000050	-0.000066
SPSP	0.001969	0.001171	3.954495
Gromacs 4	---	0.011xxx	0.005xxx
Desmond	---	0.017xxx	0.001xxx
NAMD	---	0.023xxx	---

+ = J. Chem. Theory Comput., Vol. 4, No. 3, 2008

* = shake on H bonds. 0.000008 kT/ns/dof = 0.01798 kcal/mol/ns.

Energy Conservation

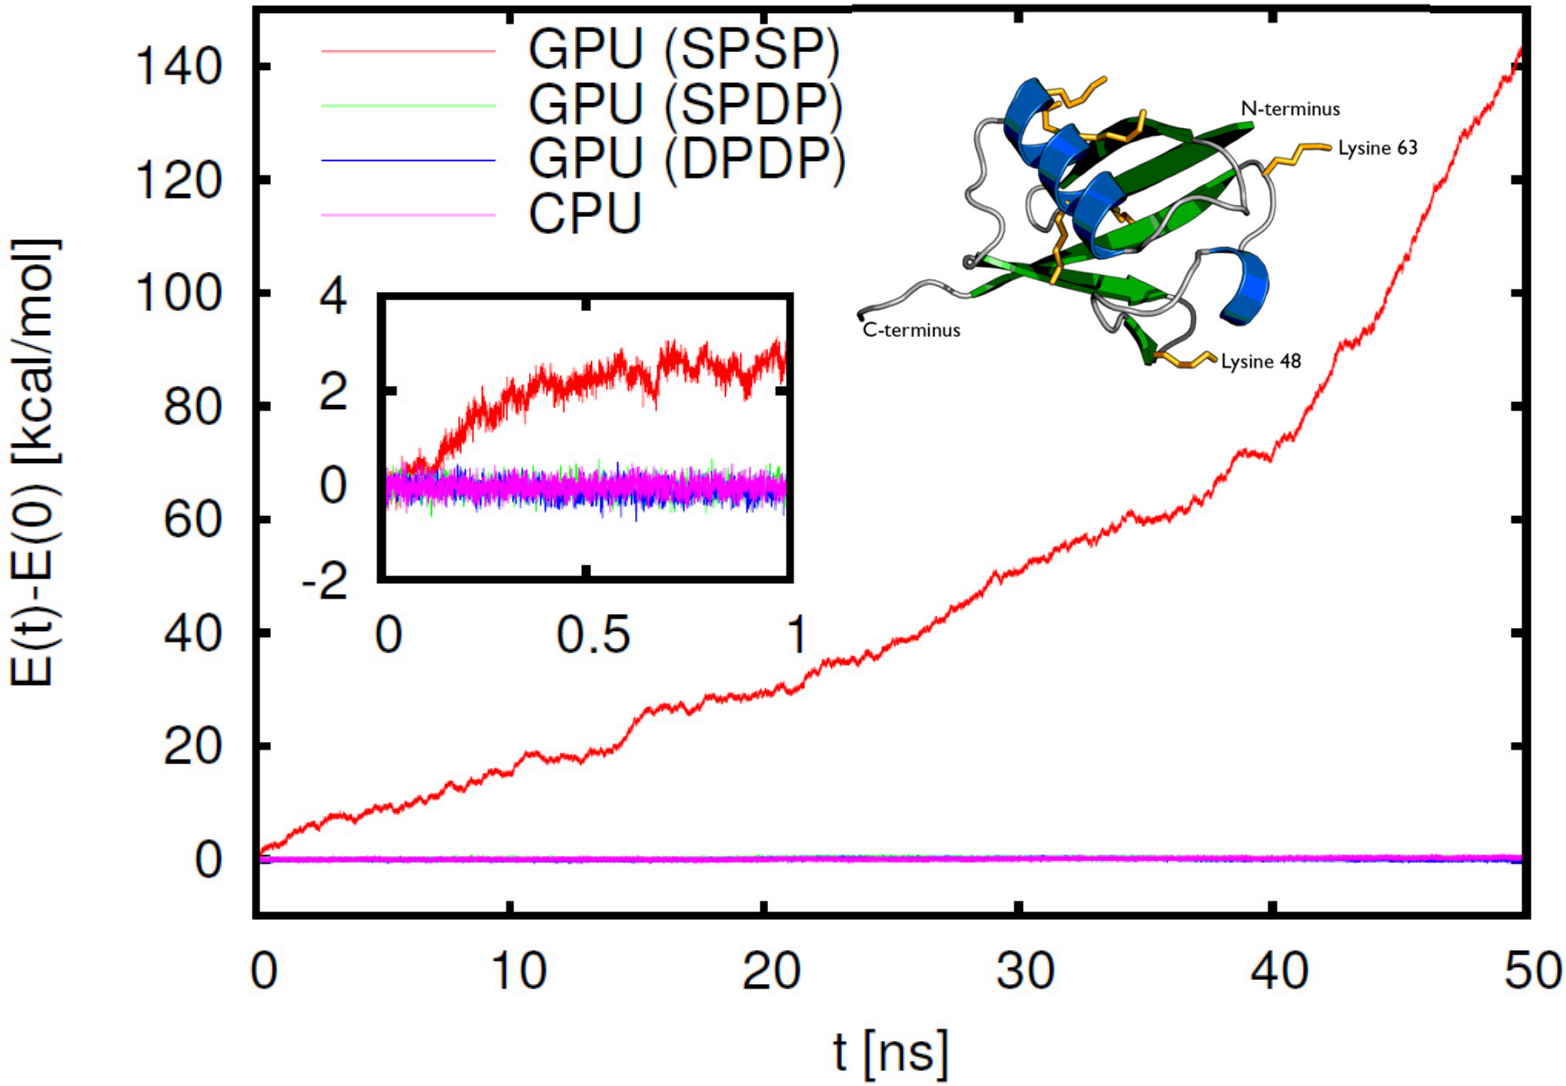

Single Precision Issues?

- **NVE does not work, but we can use a thermostat no?**

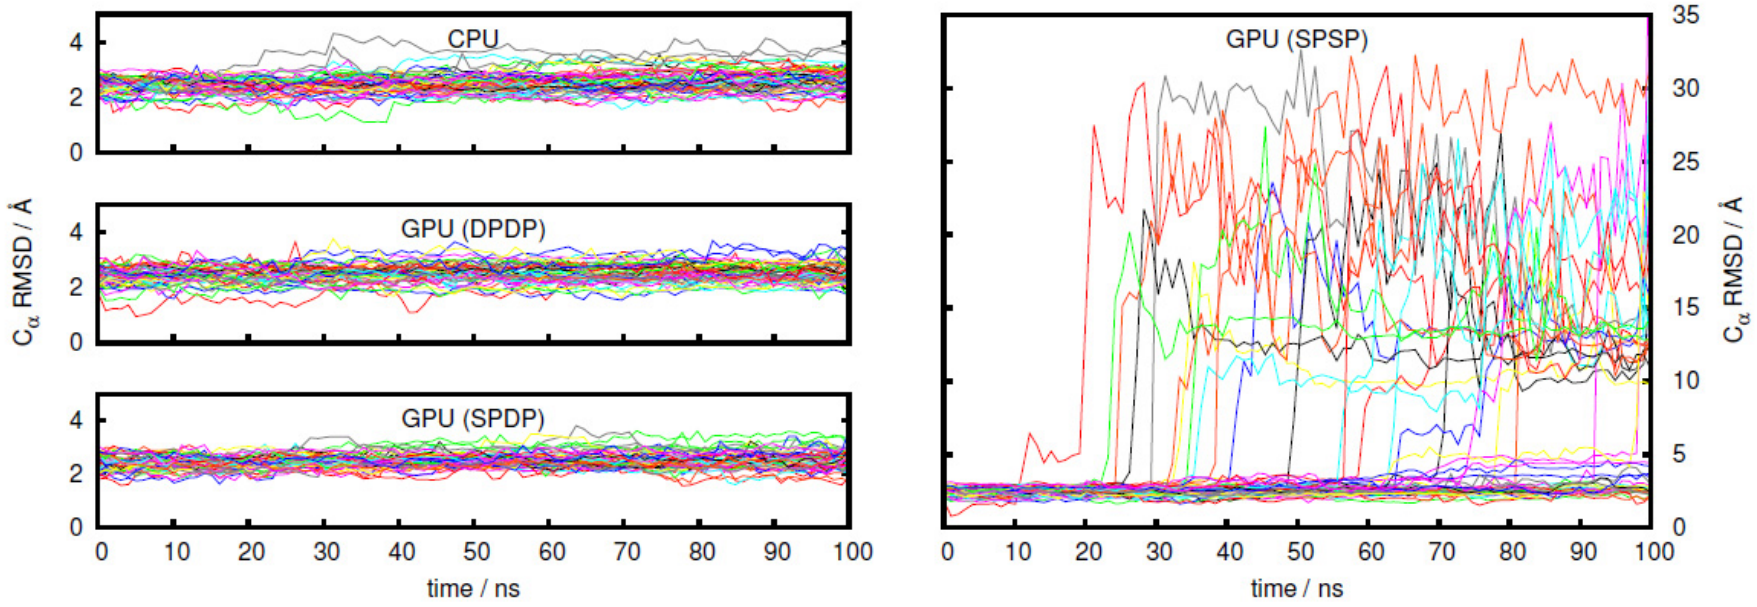

Figure 4: Root-mean-square deviations (RMSDs) of the C_α backbone carbon atoms of ubiquitin (excluding the flexible tail, residues 71 to 76) with respect to the crystal structure for 50 independent trajectories as obtained with the CPU implementation and the GPU implementation of PMEMD using different precision models.

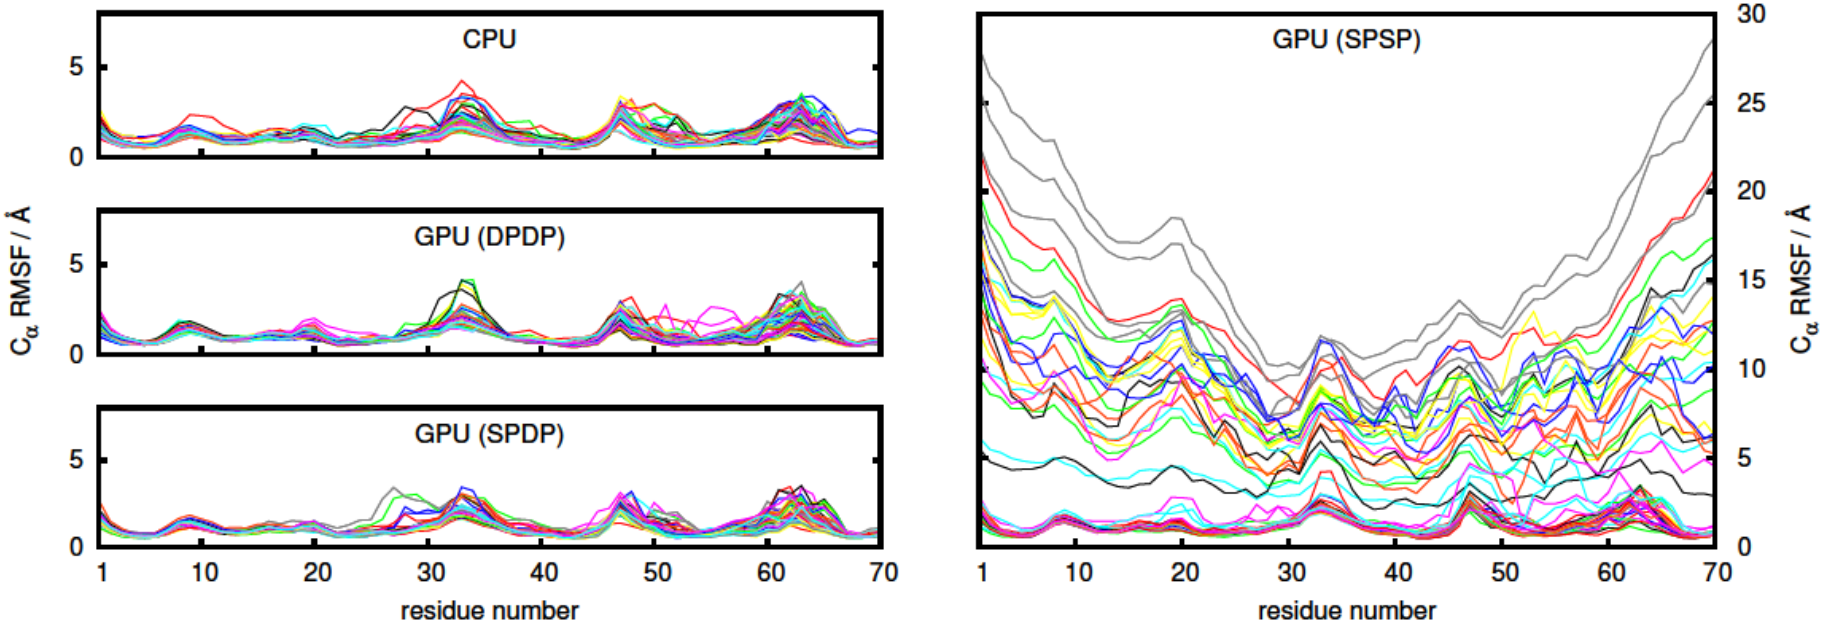

Figure 5: Root-mean-square fluctuations (RMSFs) of the C_α backbone carbon atoms of ubiquitin residues 71 to 76 with respect to the crystal structure for 50 independent trajectories of 100 ns length as obtained with the CPU implementation and the GPU implementation of PMEMD using different precision models.

Performance

AMBER 11

Implicit Solvent Performance

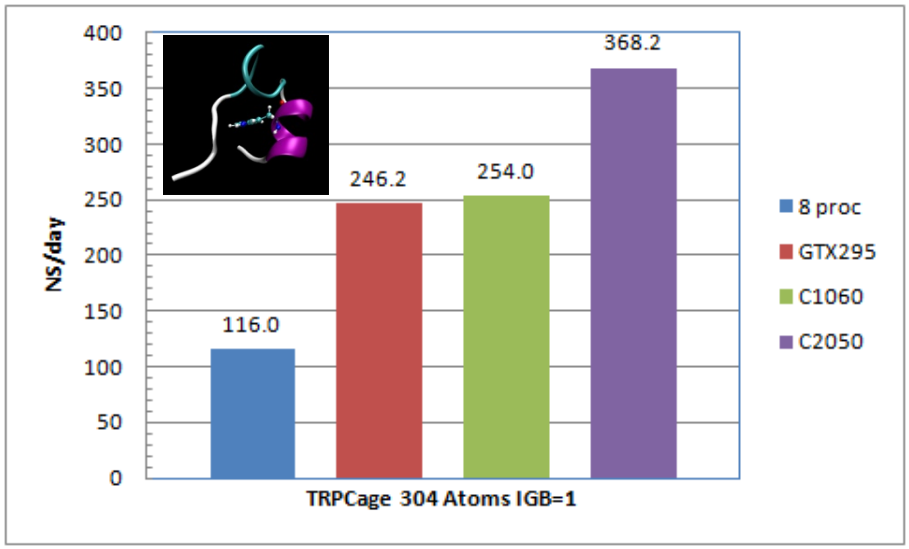

As expected the performance differential is larger for bigger systems.

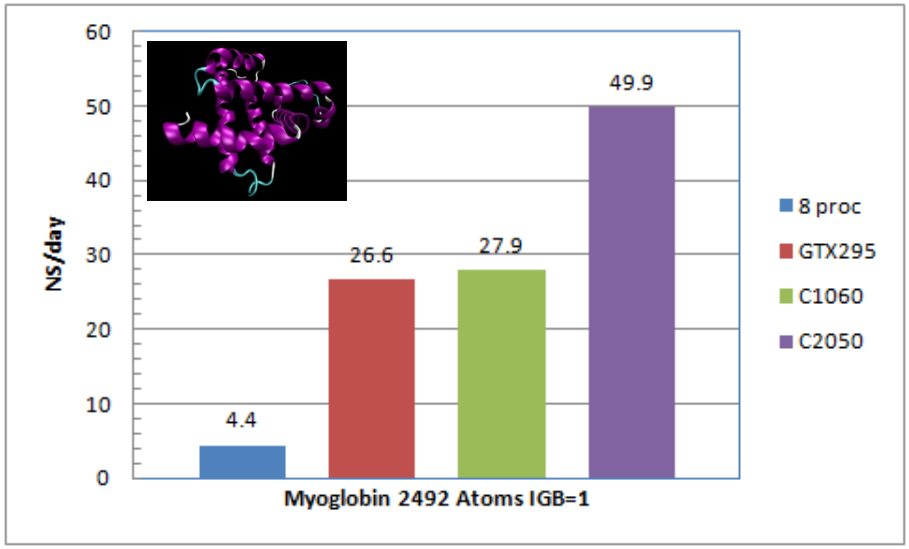

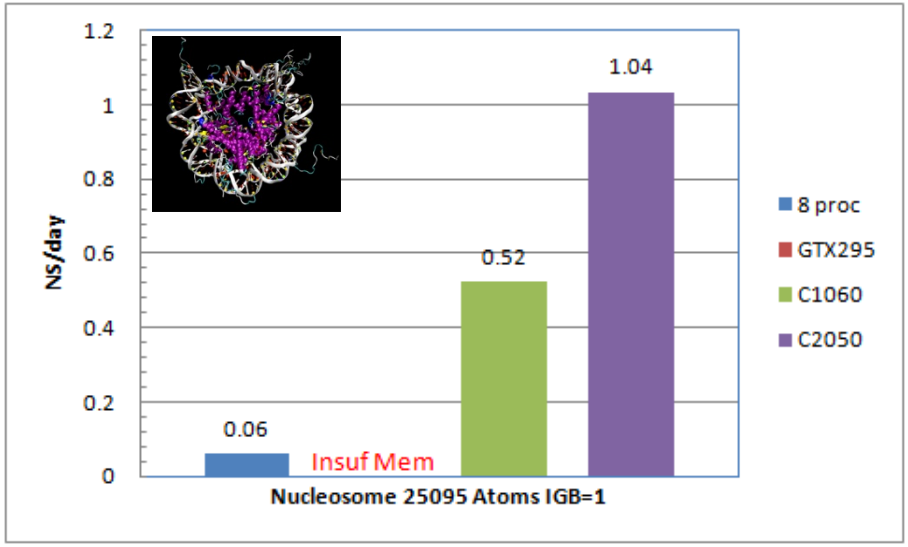

8 proc = Intel E5462 @ 2.8GHz

Explicit Solvent Performance (JAC DHFR Production Benchmark)

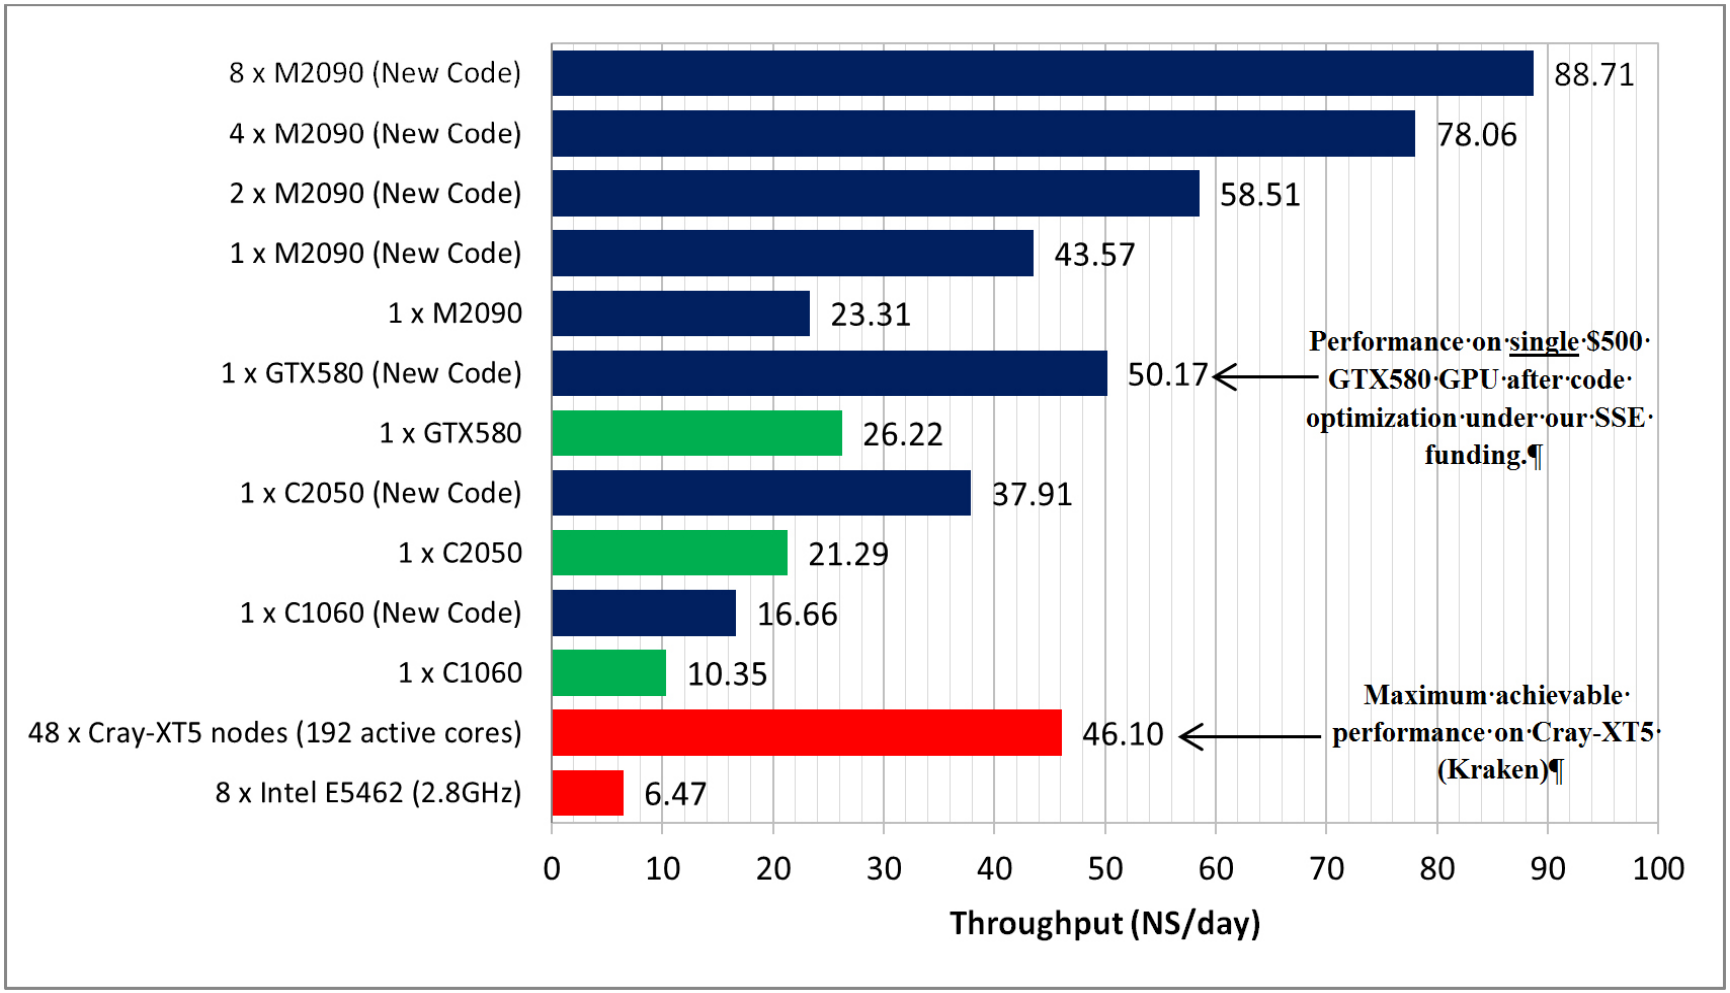

Performance Example (TRP Cage)

CPU 4xE5462

```
[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]$ tail -f mdout
```

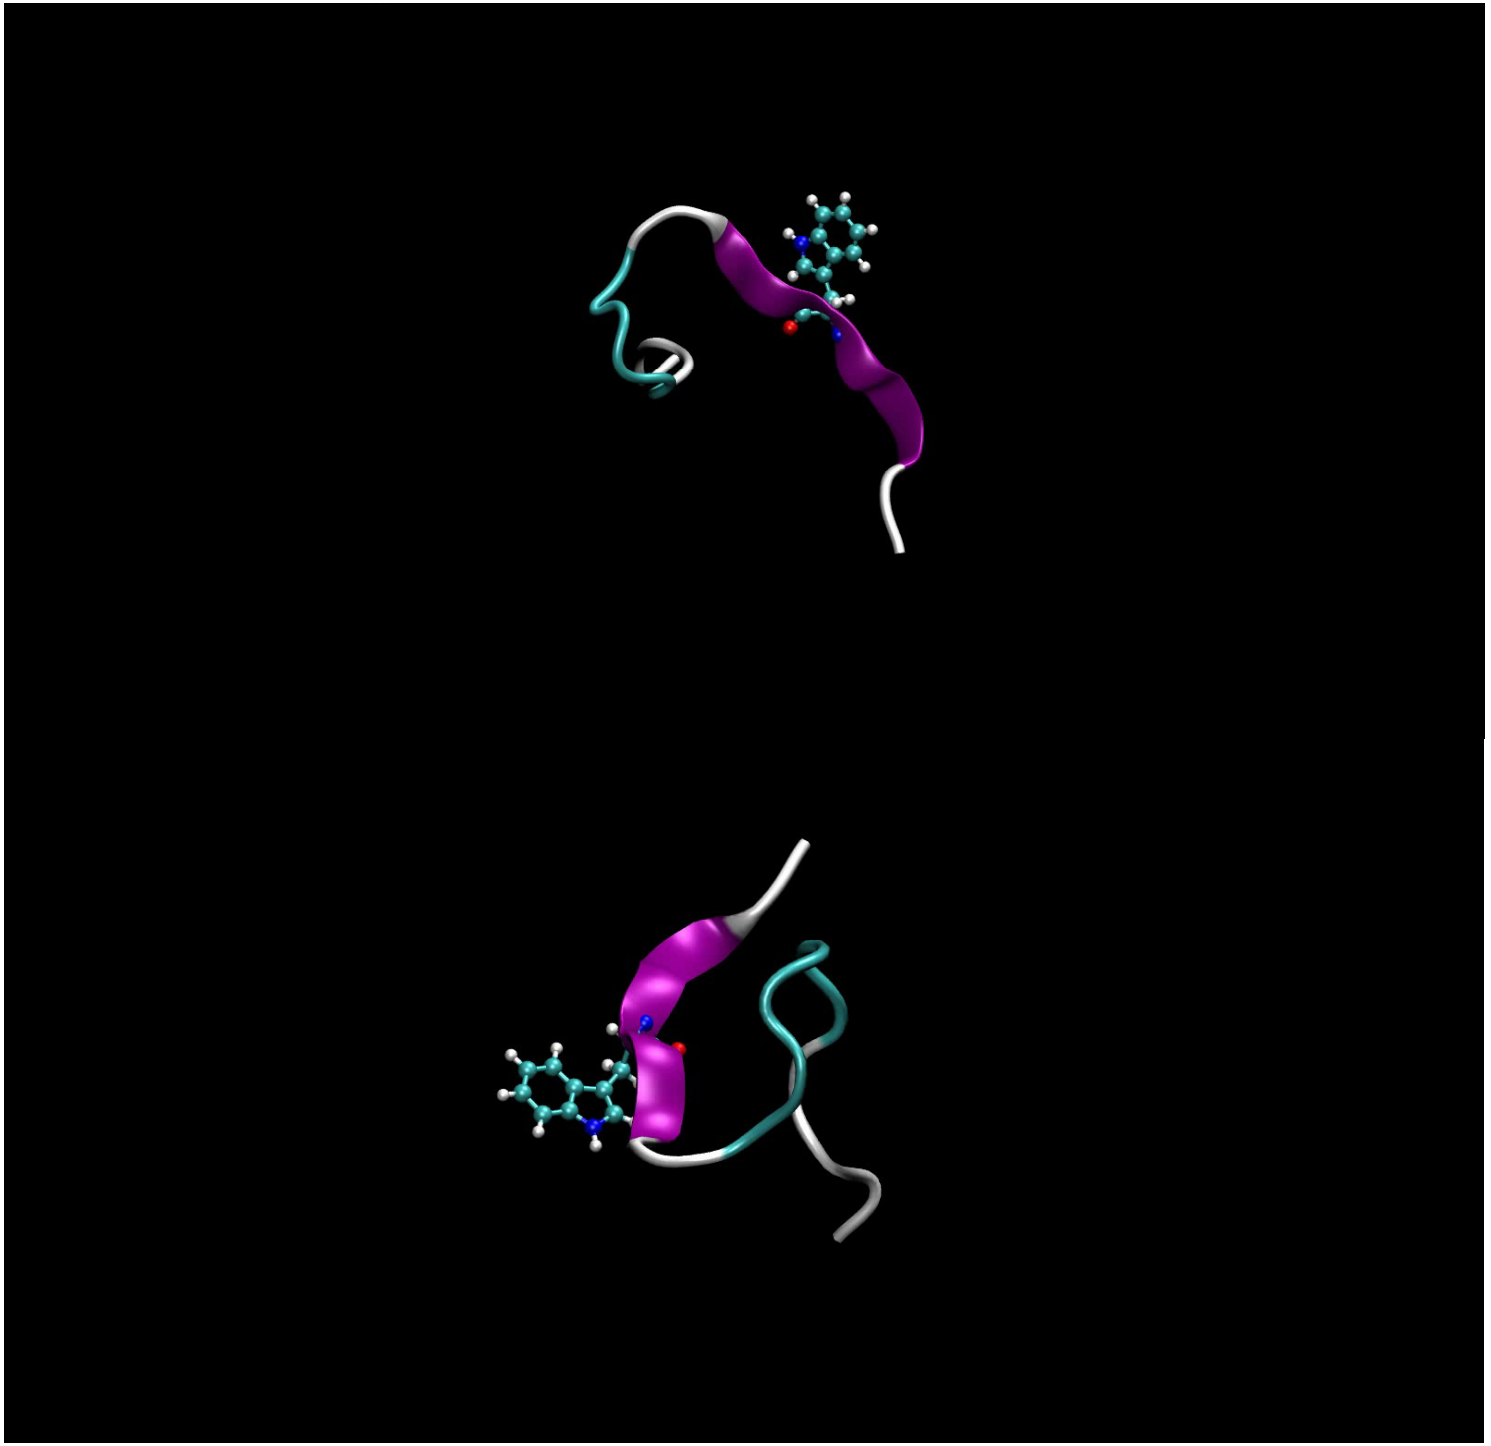

Smart sampling for the 99%

Outline

- **Enhanced Sampling**
  - AMD = Accelerated Molecular Dynamics
- **MD on the GPU**
- **Combining enhanced sampling with the power of the GPU**
- **Scaling of the new code**
- **Application to BPTI**

Why do we need enhanced sampling?

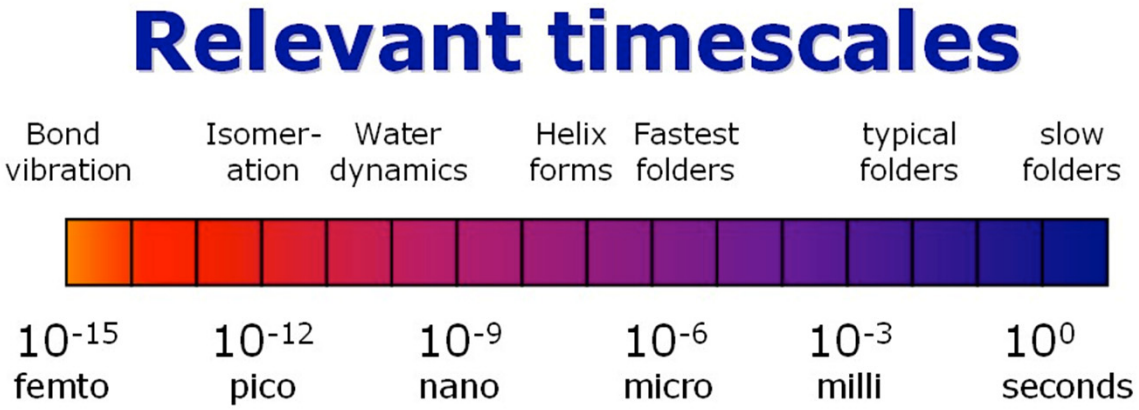

- **16 order of magnitude range**
  - Femtosecond timesteps
  - Need to simulate micro to milliseconds

Motivation: Rare events

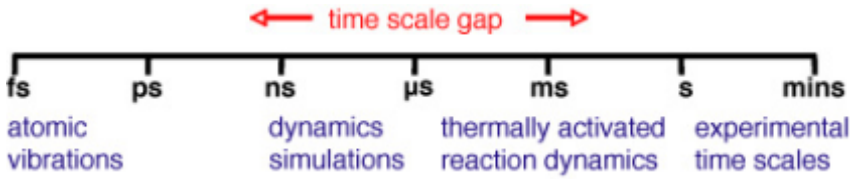

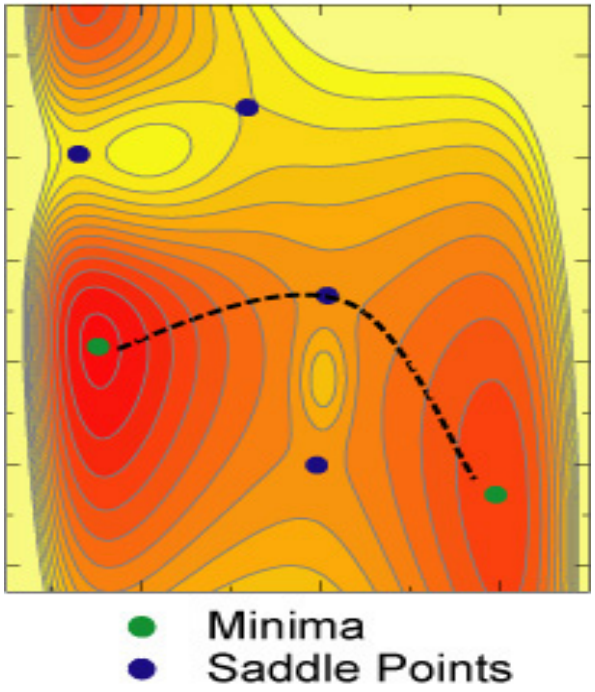

- Many processes in nature involve transition through rare events: reactions, conformational changes.
- Low probability, therefore, long simulation times
- Most time is spent in vibrations in stable basins
- Need more (GPUs, Anton, MPI, etc) and more efficient sampling (Metadynamics, AMD, replica exchange, etc)

Accelerated Molecular Dynamics

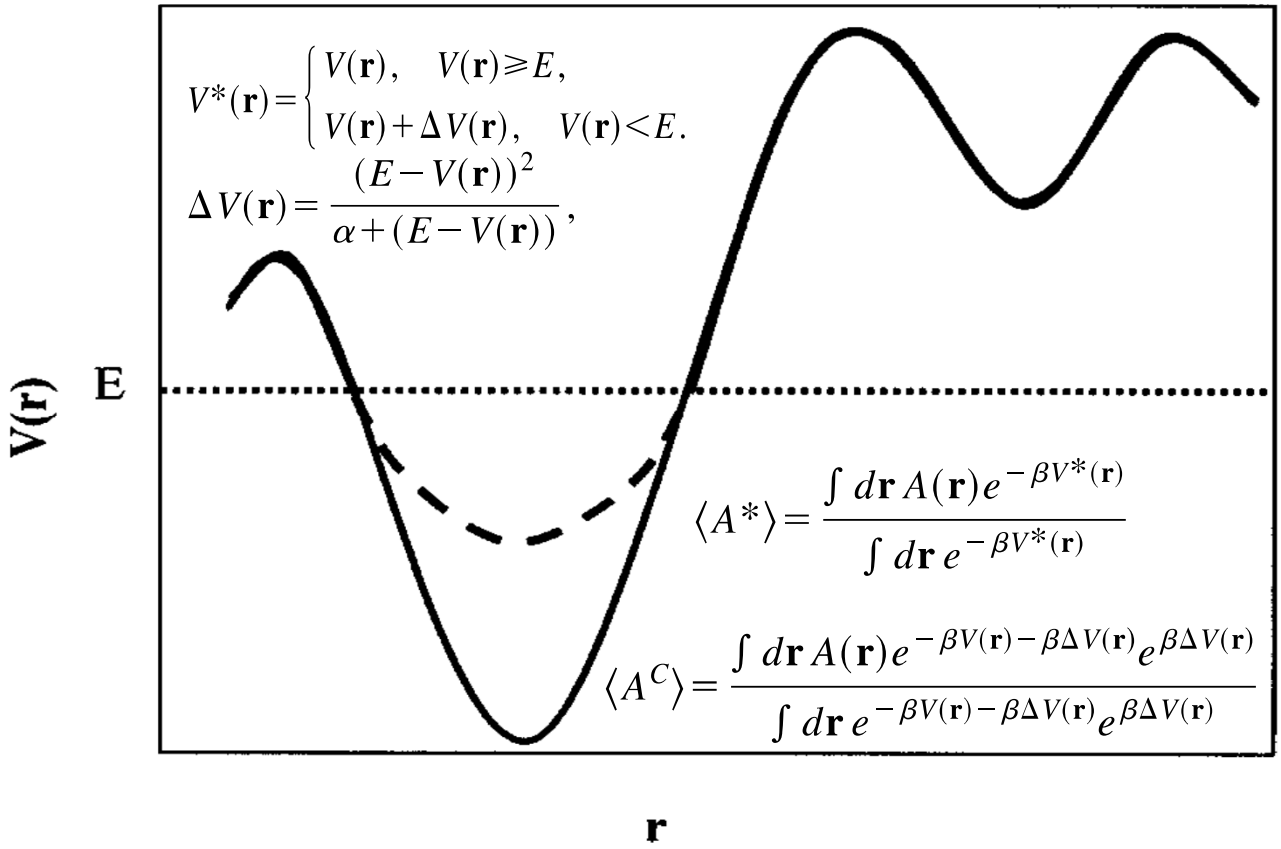

de Oliveira C.A.F., Hamelberg, D., McCammon, J.A., On the Application of Accelerated Molecular Dynamics to Liquid Water Simulations. J. Phys. Chem. B 2006.

Hamelberg, D., de Oliveira C.A.F., McCammon J.A. , *Sampling of slow diffusive conformational transitions with accelerated molecular dynamics.* The Journal of chemical physics, 2007.

Grant, B.J., Gorfe, A.A., and McCammon, J.A. , *Ras conformational switching: simulating nucleotide-dependent conformational transitions with accelerated molecular dynamics.* PLoS computational biology, 2009.

de Oliveira, C.A.F., et al., *Large-Scale Conformational Changes of Trypanosoma cruzi Proline Racemase Predicted by Accelerated Molecular Dynamics Simulation.* PLoS computational biology, 2011.

AMD effect on the Potential

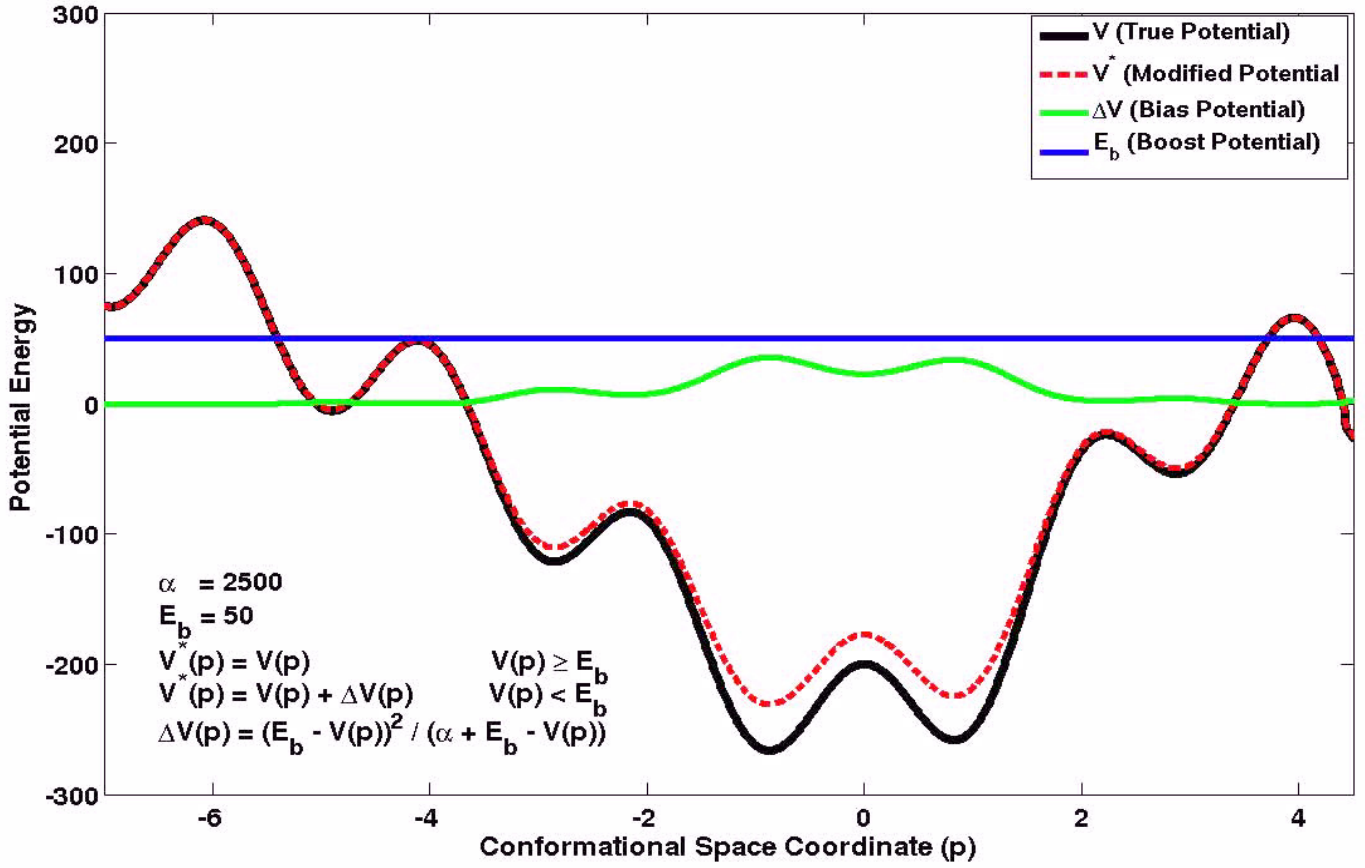

Implementing aMD on the GPU

- **Ported sander aMD to PMEMD**
  - Good performance improvement vs sander aMD
  - Scaling on the cluster vs NAMD AMD
- **Ross Walker + Romelia Salomon programmed the method into the GPU code**
  - GTX 580 43.2 ns/day MD
  - GTX 580 41.3 ns/day aMD

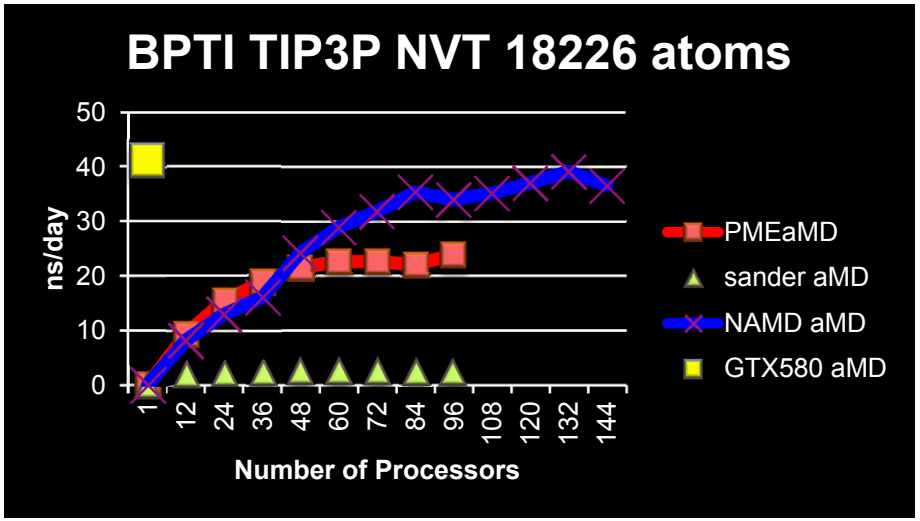

Wang, Y., et al., *Implementation of Accelerated Molecular Dynamics in NAMD.* Computational science & discovery, 2011. **4**(1).

The Benchmark

- **1-millisecond BPTI brute force MD Anton simulation**
- **Long lived conformations were observed ranging from 6 to 26 micro seconds**
- **Representative PDB Structures from Kinetic Clustering**
  - 1 Extracted at time 0.13562525 ms RED CLUSTER
  - 2 Extracted at time 0.93412525 ms BLUE CLUSTER
  - 3 Extracted at time 0.82912525 ms GREEN CLUSTER
  - 4 Extracted at time 0.43162525 ms PURPLE CLUSTER
  - 5 Extracted at time 0.66225025 ms BLACK CLUSTER

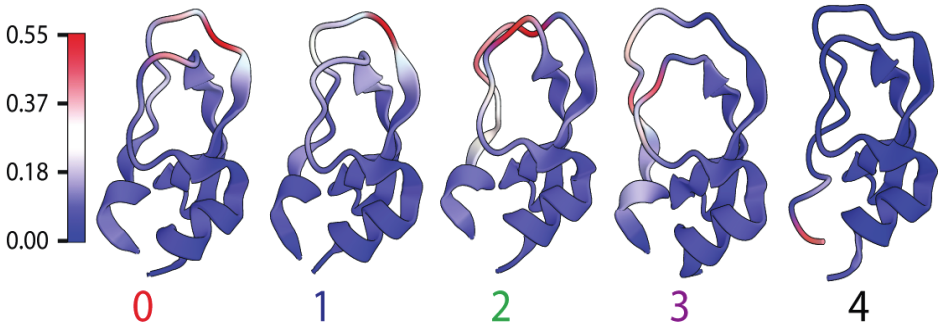

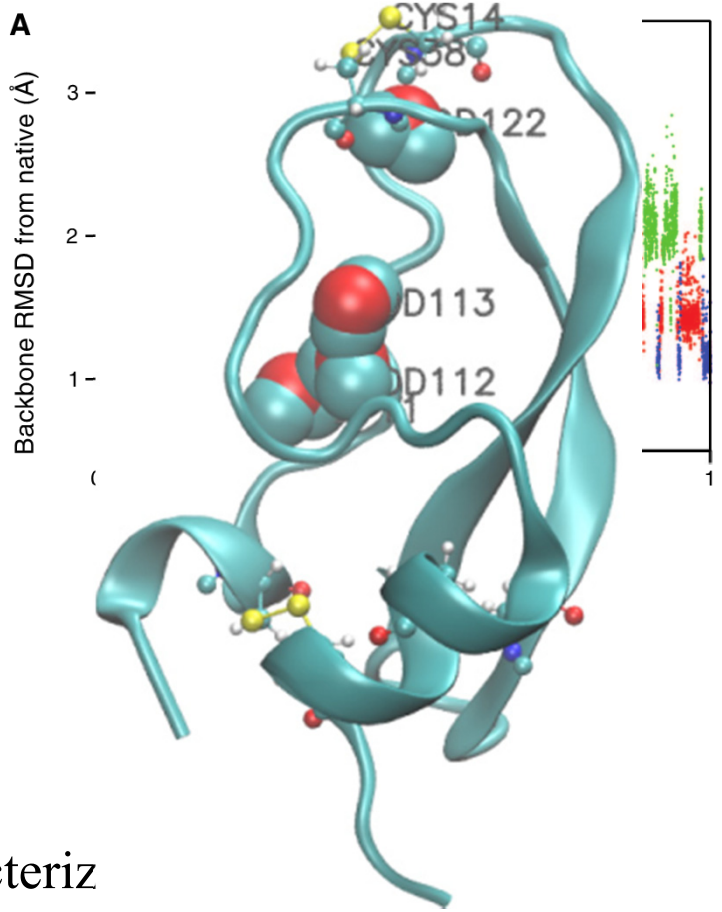

Shaw, D.E., et *al.* Atomic-Level Characteriz
Structural Dynamics of Proteins. Science 2010.

aMD Simulation Protocol

- System was setup as closely to that of Shaw's
- Same number atoms and box size
- Same force field Parm99SB-ILDN
- Same water model TIP4PEW
- 50ns of cMD was run to determine average dihedral and total energy this cMD run was extended to 500ns as a reference
- Acceleration levels for the GPU aMD run were chosen based on earlier work★✚

★ Grant, B.J., Gorfe, A.A., and McCammon, J.A. , *Ras conformational switching: simulating nucleotide-dependent conformational transitions with accelerated molecular dynamics.* PLoS computational biology, 2009.
✚ de Oliveira, C.A.F., et al., *Large-Scale Conformational Changes of Trypanosoma cruzi Proline Racemase Predicted by Accelerated Molecular Dynamics Simulation.* PLoS computational biology, 2011.

RMSD to "Long Lived" States

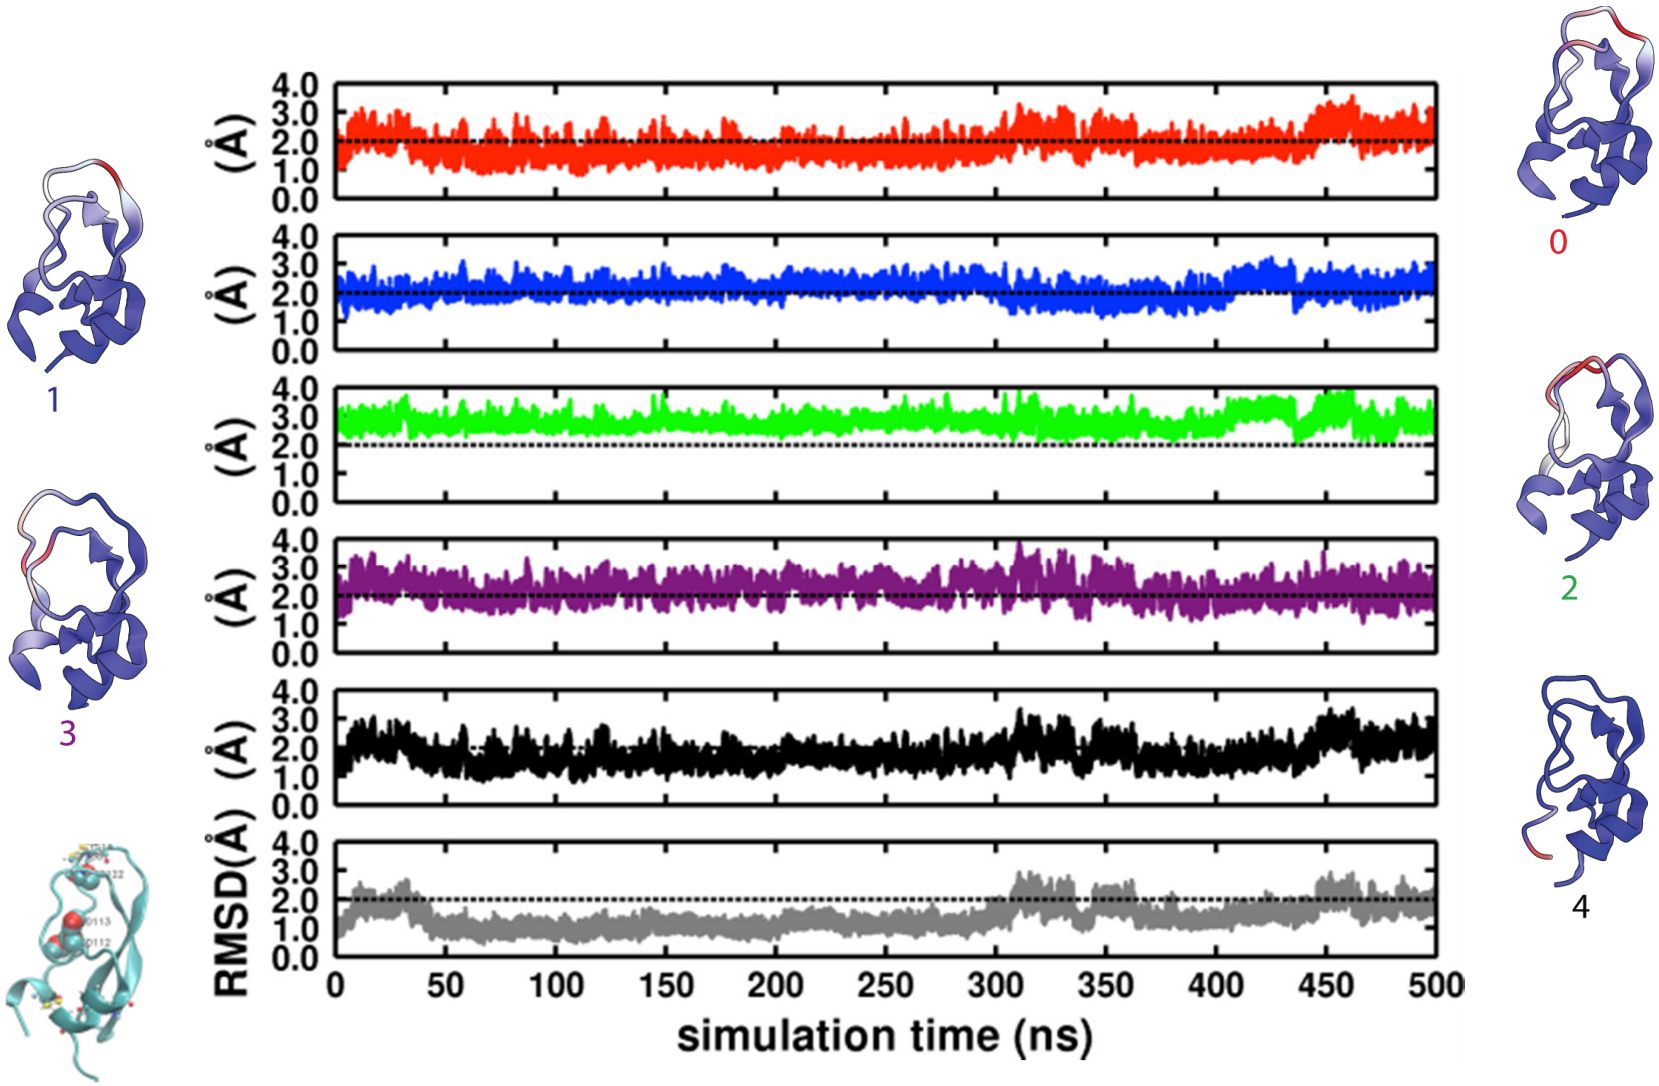

Principal Component Analysis

- **Build PC space based on the 500ns aMD simulation**
- **Project X-ray, Shaw structures, 500ns cMD, and 1ms cMD**

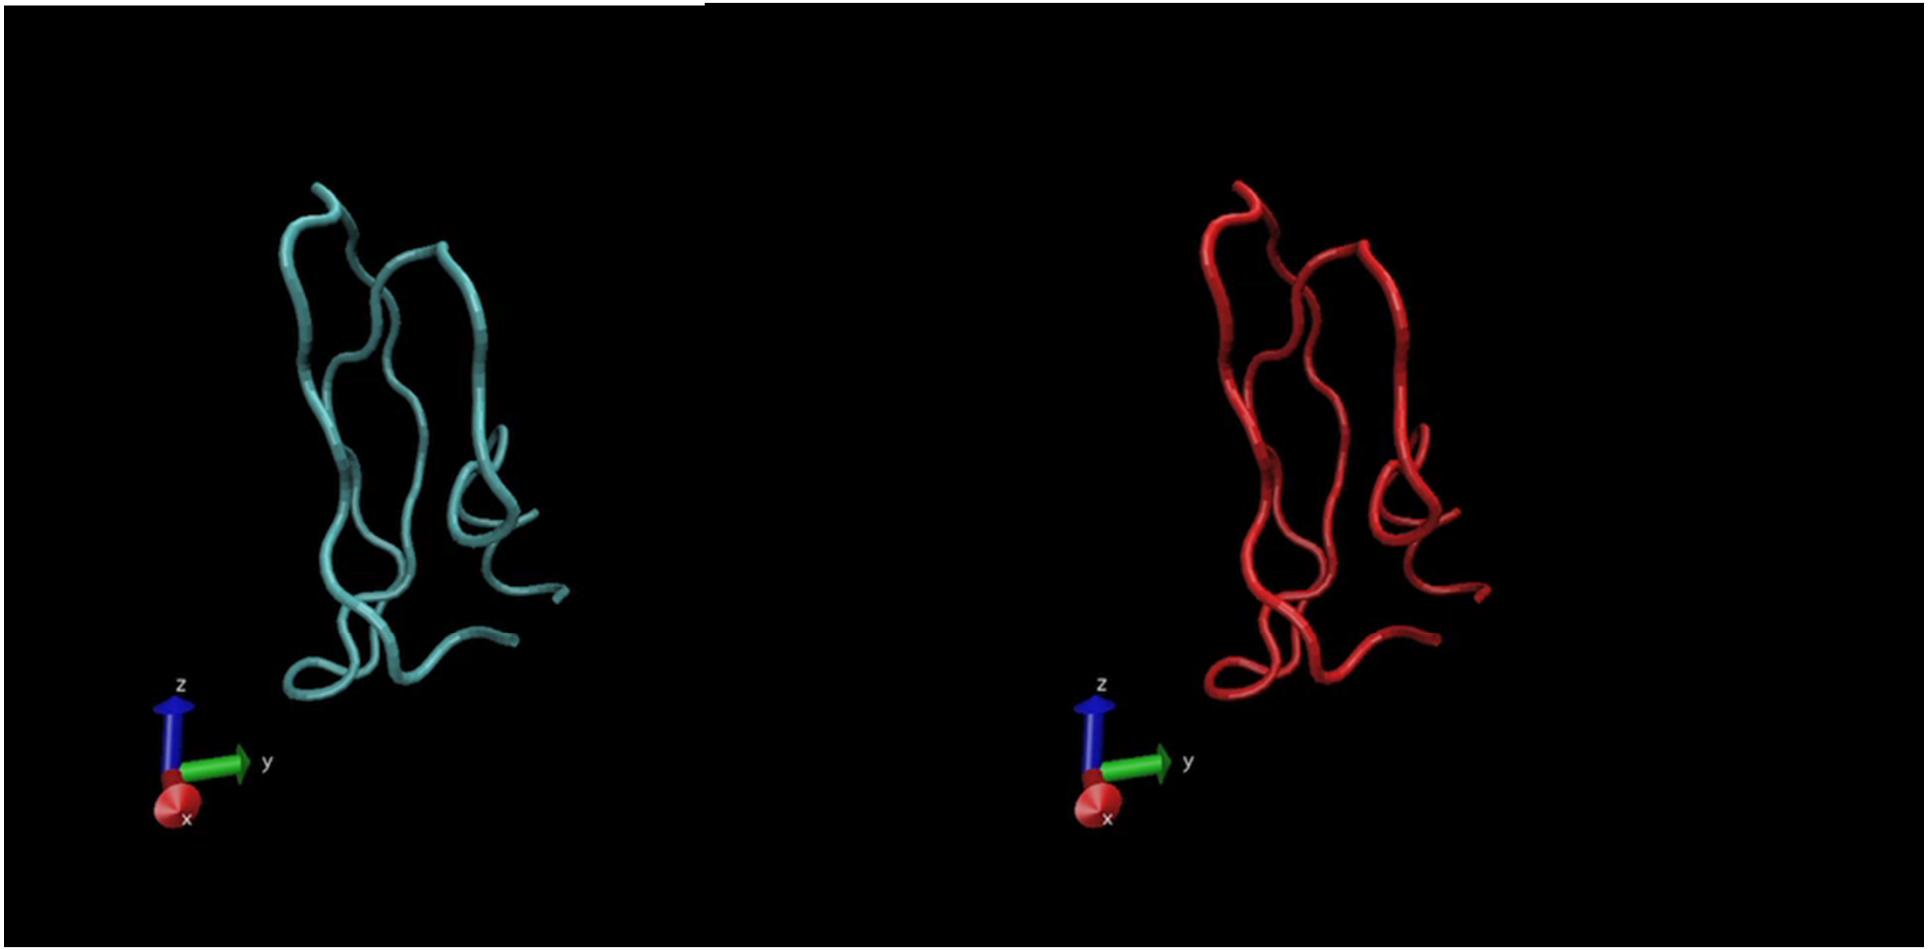

CYS14 CYS38 Chi 1 Analysis

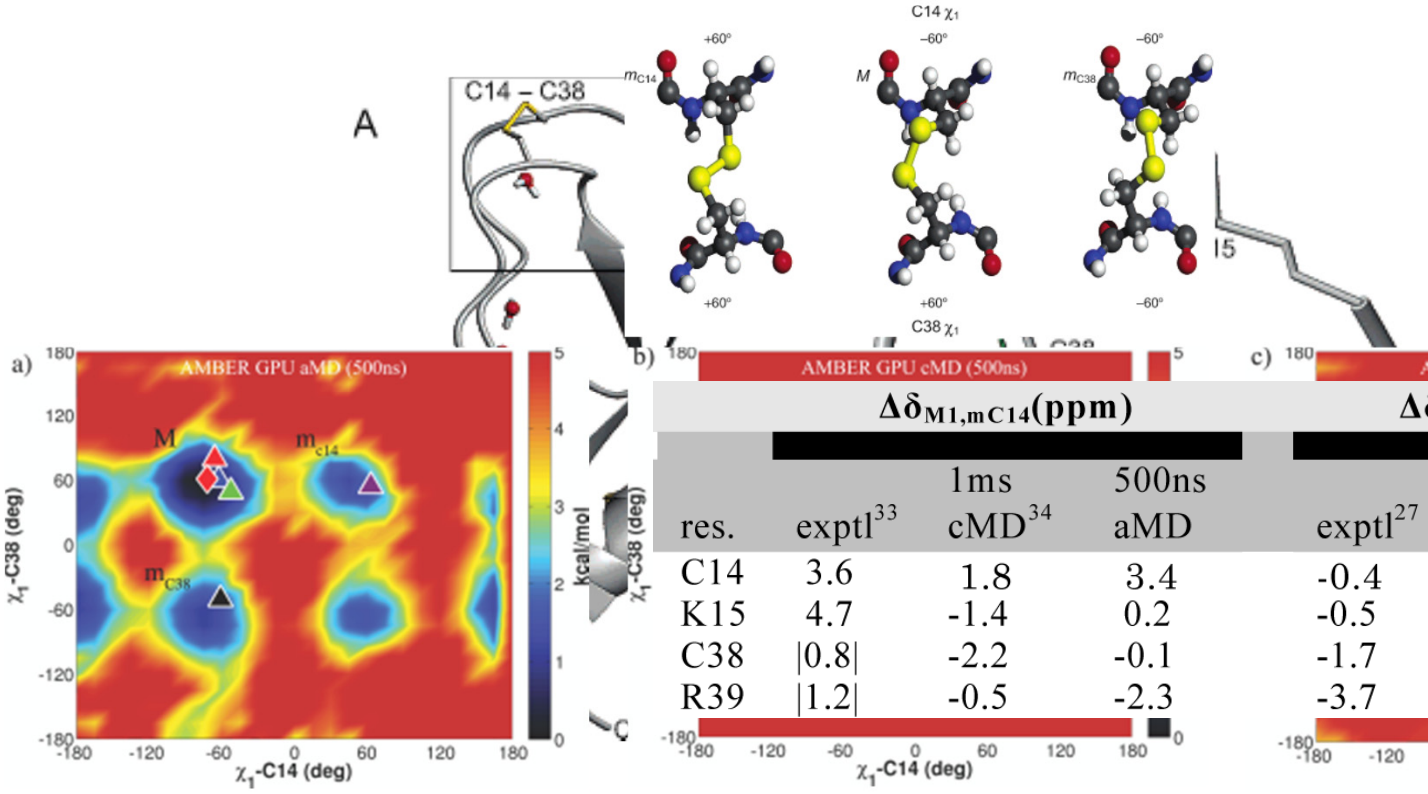

aMD = 1.8 PPM RMS
cMD = 2.7 PPM RMS

	$\Delta\delta_{M1,mC14}$(ppm)			$\Delta\delta_{M1,mC38}$(ppm)		
res.	exptl[33]	1ms cMD[34]	500ns aMD	exptl[27]	1ms cMD[34]	500ns aMD
C14	3.6	1.8	3.4	-0.4	0.6	0.1
K15	4.7	-1.4	0.2	-0.5	-1.5	-0.1
C38	|0.8|	-2.2	-0.1	-1.7	-1.9	-3.2
R39	|1.2|	-0.5	-2.3	-3.7	-2.6	-3.7

Grey, M.J, et *al.* Disulfide Bond Isomerization in Basic Pancreatic Trypsin Inhibitor: Multisite Chemical Exchange Quantified by CPMG Relaxation Dispersion and Chemical Shift Modeling. JACS, 2003.

Otting, G.; Liepinsh, E.; Wuethrich, K. *Journal of the American Chemical Society* 1991, *113*, 4363.

Xue, Yi, et al. µs time-scale conformational exchange in proteins: using long MD trajectory to simulate NMR relaxation dispersion data. JACS, 2011

Population Analysis

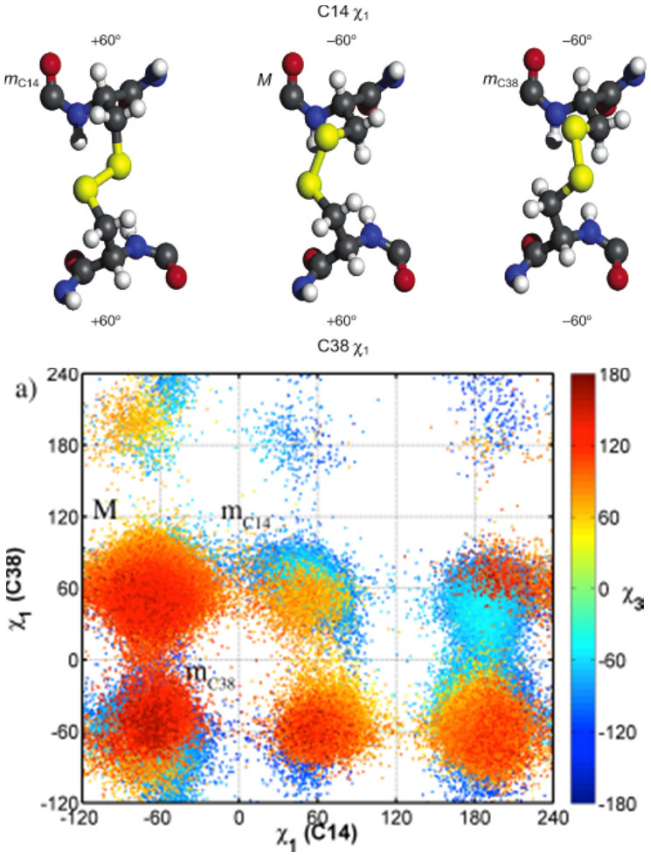

Table S1. Conformational Species of BPTI Classified According the C14-C38 Disulfide Bridge

Conformational States	Exptl pop(%)	1ms cMD pop(%)	500ns aMD pop(%)
M(Grey et al. [16])	~95	34	61
m_{C14}(Grey et al. [16])	~1	50	2.6
m_{C38}(Grey et al. [16])	~4	6	7.9

Grey, M.J, et *al.* Disulfide Bond Isomerization in Basic Pancreatic Trypsin Inhibitor: Multisite Chemical Exchange Quantified by CPMG Relaxation Dispersion and Chemical Shift Modeling. JACS, 2003.

Otting, G.; Liepinsh, E.; Wuethrich, K. *Journal of the American Chemical Society* 1991, *113*, 4363.

Xue, Yi, et al. μs time-scale conformational exchange in proteins: using long MD trajectory to simulate NMR relaxation dispersion data. JACS, 2011

Putting it all together

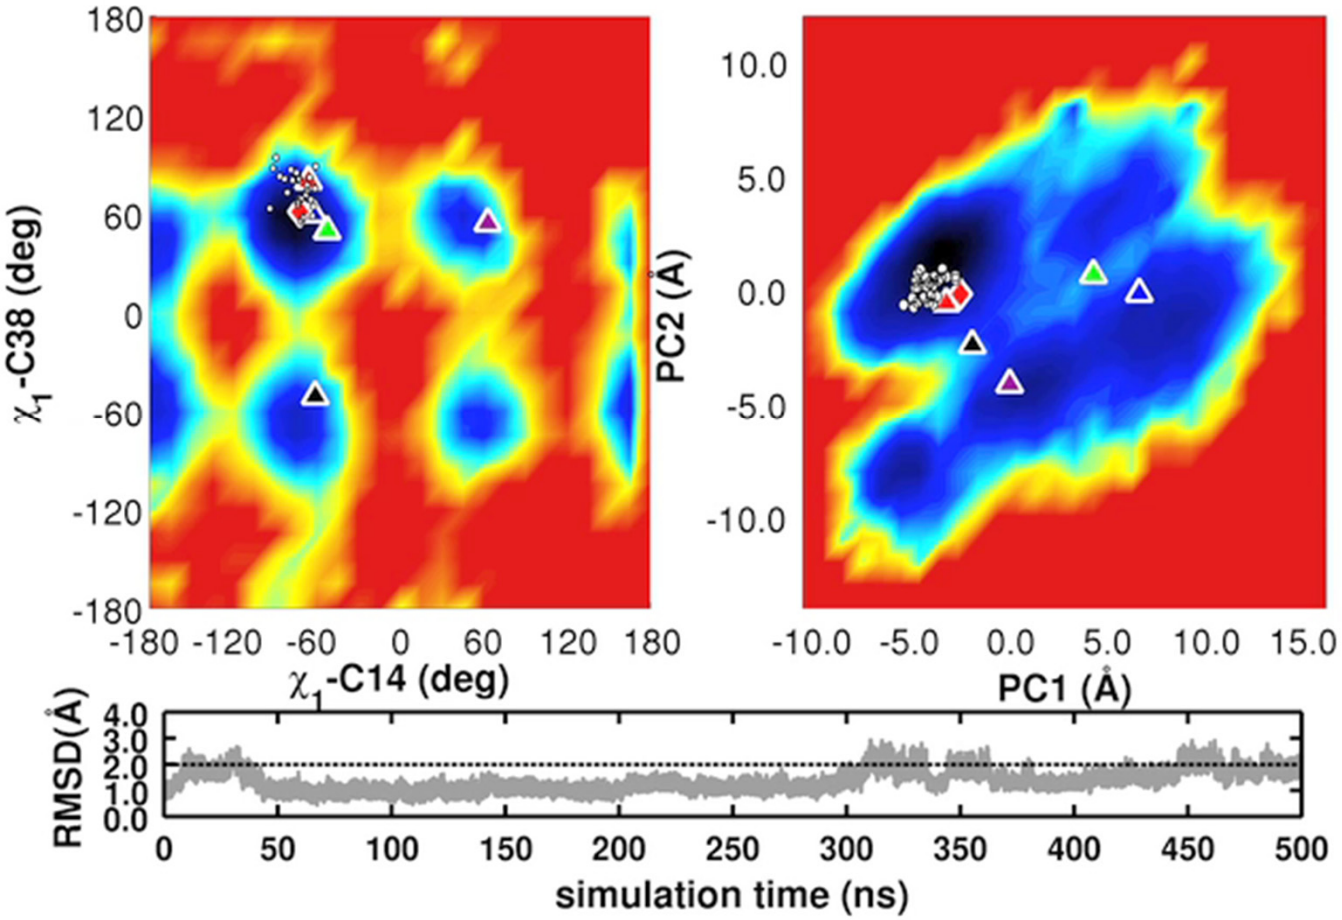

Water Occupancy

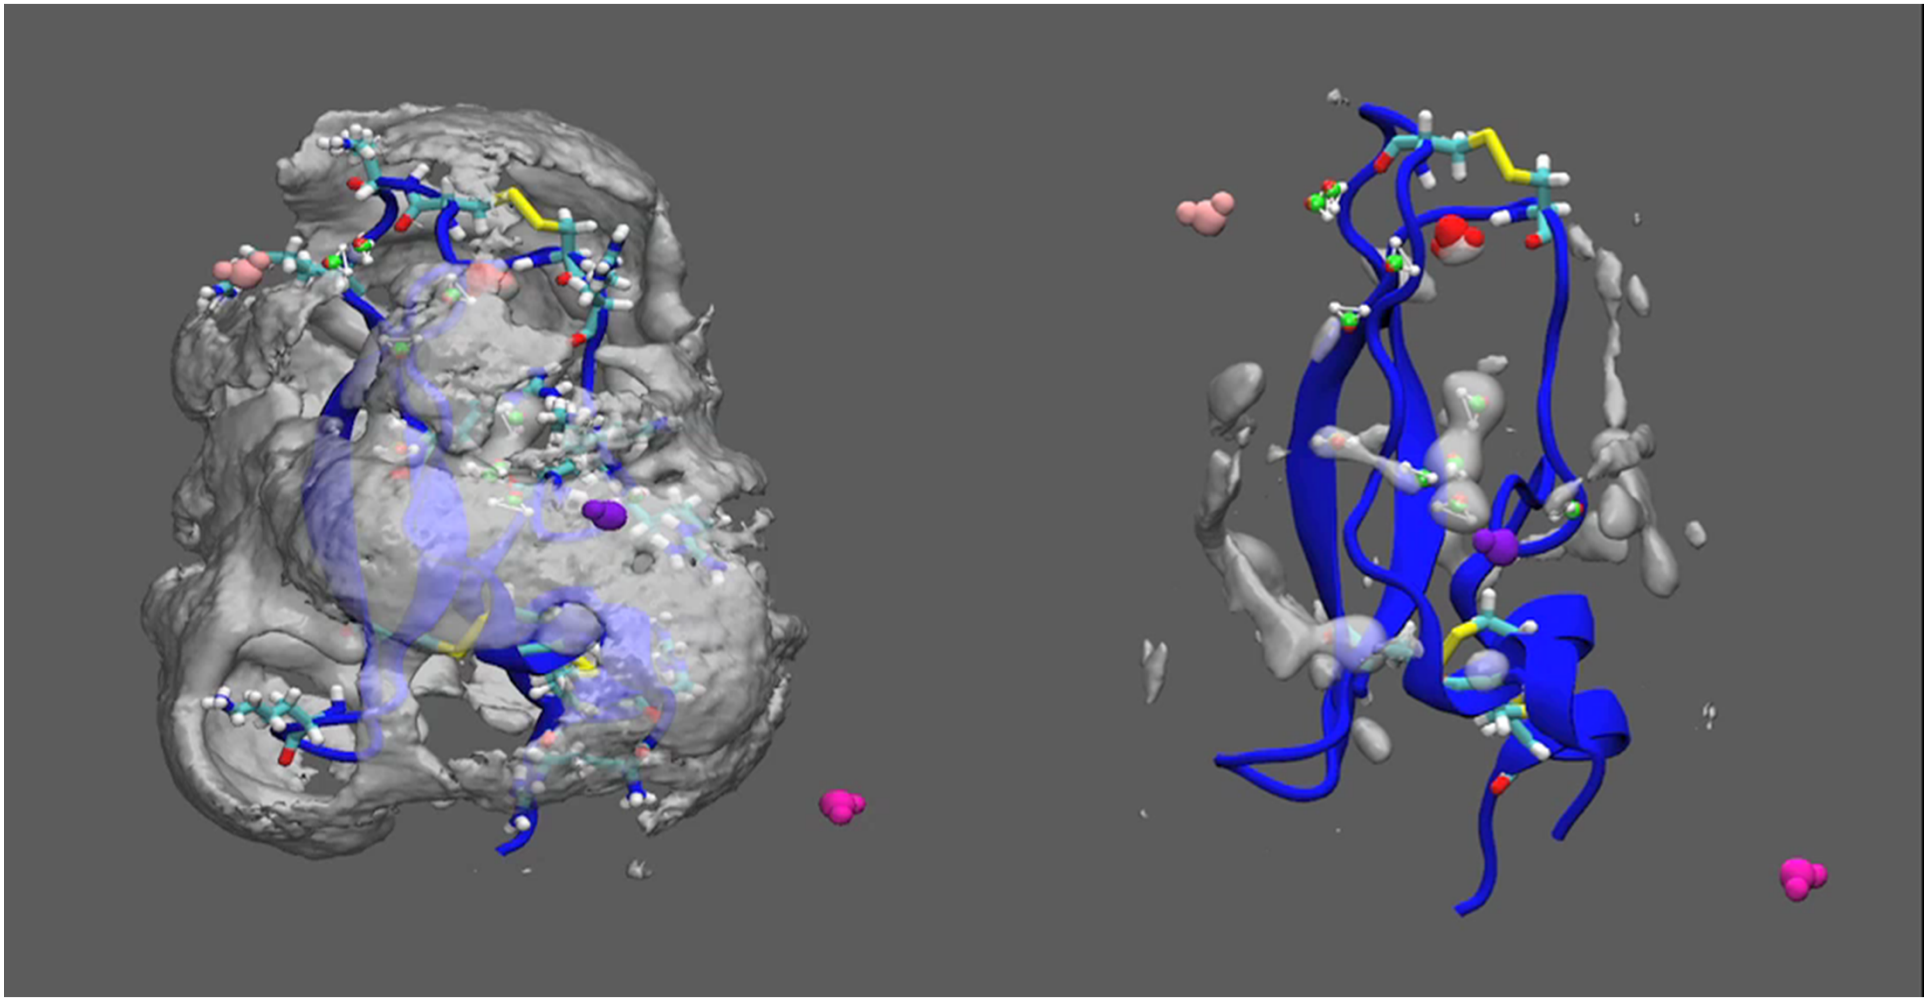

Anton vs Amber

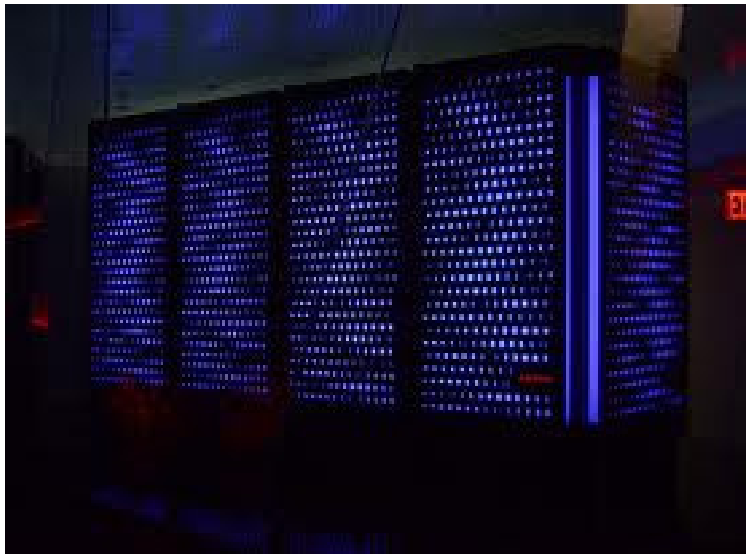

- **Cost: several million dollars**
- **Energy consumption ~116.5KW**
- **Team: 100 people involved**
- **Code hardwired to the hardware**

Cost: $30,000

Energy consumption: ~4.5 KW

Team: about 8 people involved

Flexible code

Conclusions

- **aMD has been successfully ported to the GPU**
  - To obtain the speed of one GPU 144 processors was needed
- **Sampled the phase space explored by the 1-millisecond DE Shaw simulation**
  - RMSD improvement compared to similar length MD simulation
  - “Long lived” waters were observed in the basins
  - All Chi1 states were sampled
  - We see improvements in NMR Observables
  - Observe correct isomerization populations for disulphide bonds

Acknowledgements

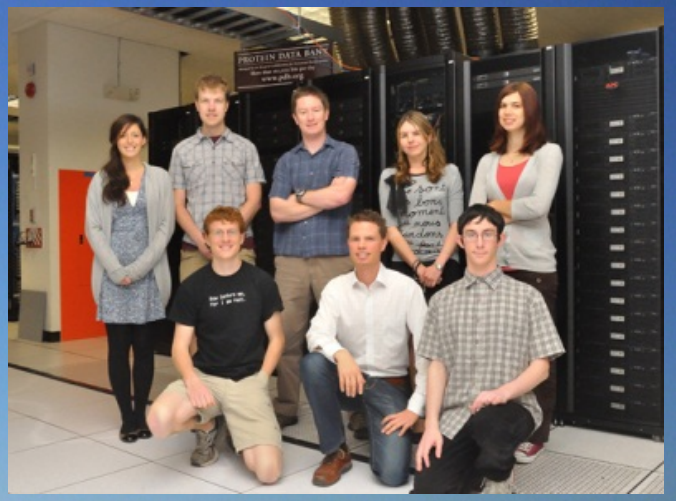

San Diego Supercomputer Center

University of California San Diego

National Science Foundation

NSF Strategic Applications Collaboration (AUS/ASTA) Program

NSF SI2-SSE Program

NVIDIA Corporation

Hardware + People

People

Mark Williamson

Andreas Goetz

Romelia Salomon

Scott Le Grand

Duncan Poole

Levi Pierce