The Origin of NAMD

Evolution and Intelligent Design

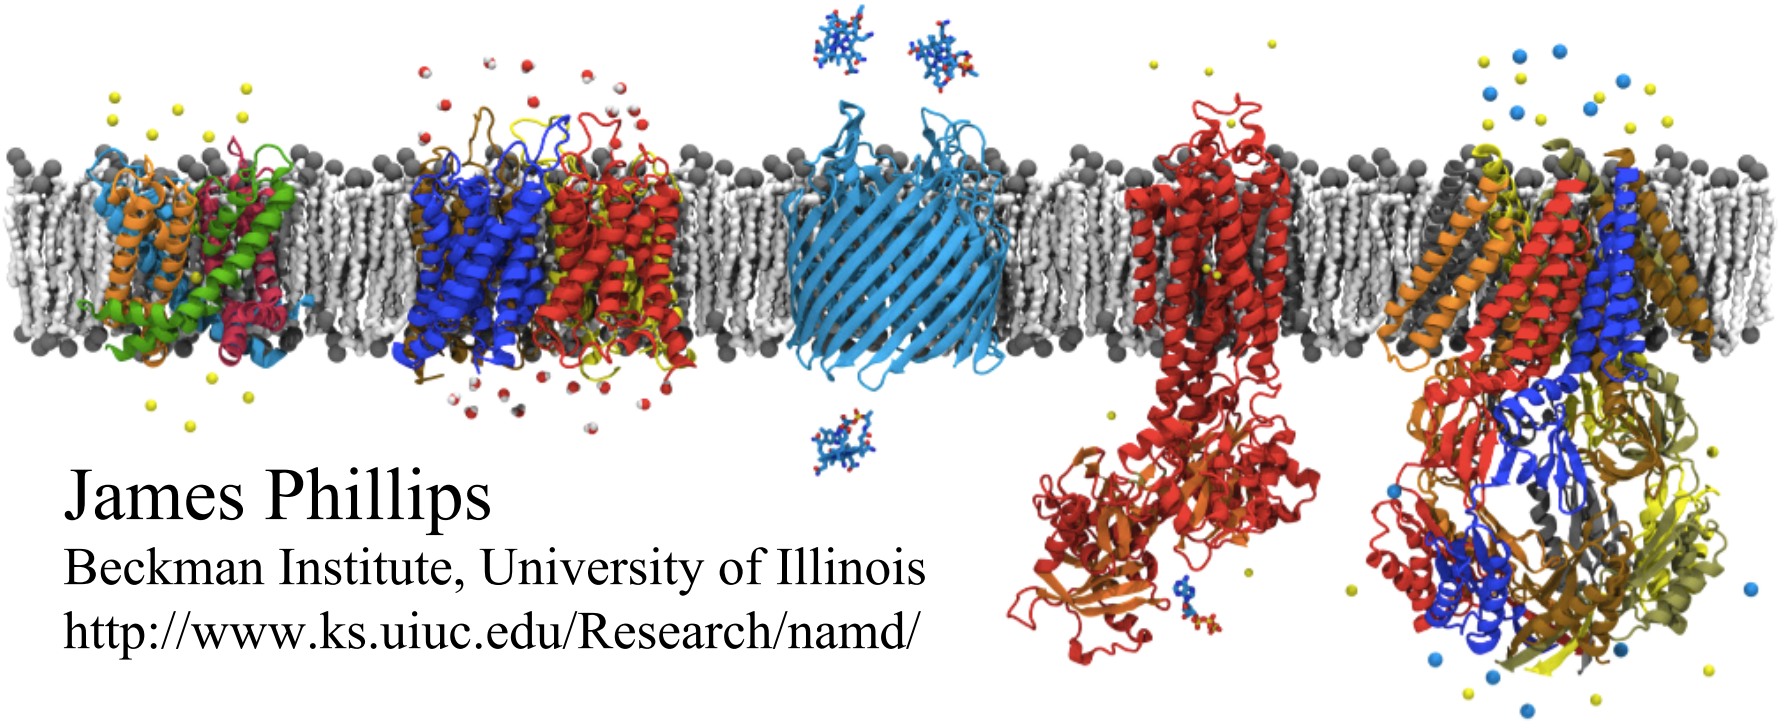

James Phillips
Beckman Institute, University of Illinois
http://www.ks.uiuc.edu/Research/namd/


NIH Center for Macromolecular Modeling and Bioinformatics
1990-2017
Beckman Institute
University of Illinois at
Urbana-Champaign

Physics of in vivo Molecular Systems

Biomolecular interactions span many orders of magnitude in space and time.

femtoseconds Center software provides multi-scale computational modeling. hours

Ångstrom microns

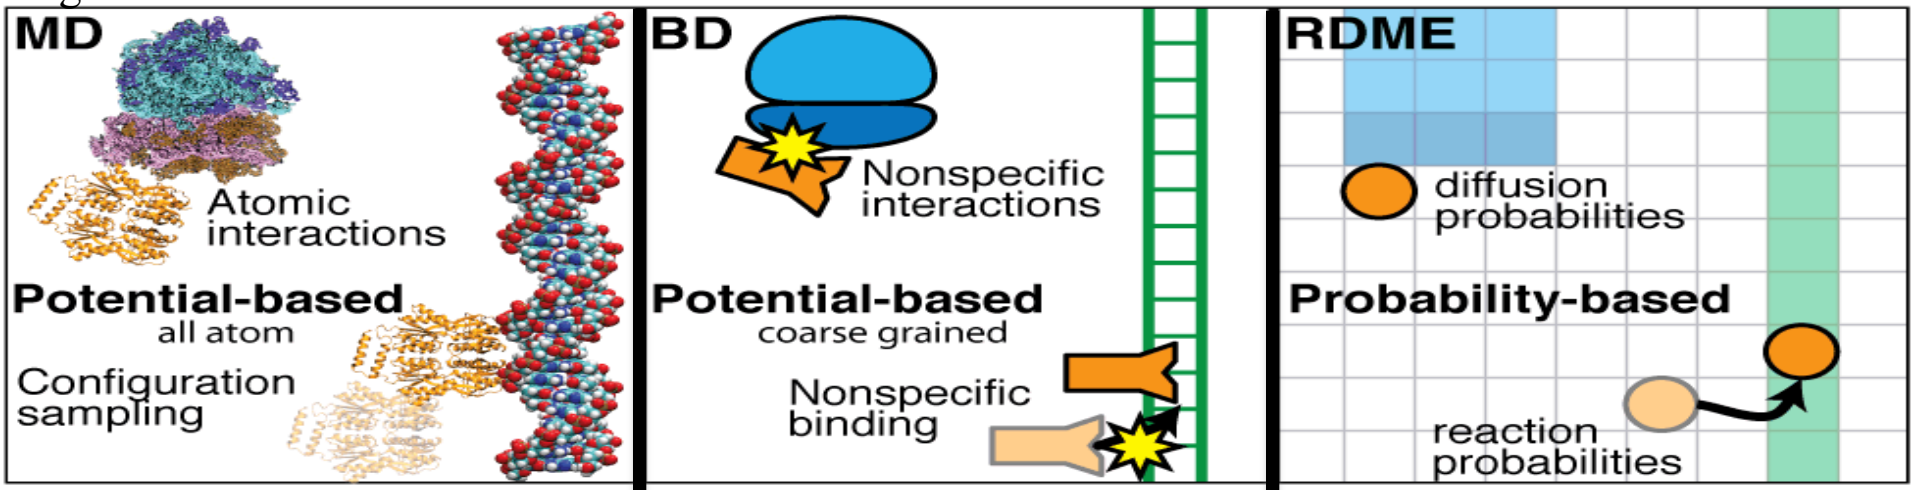

NAMD Scalable Molecular Dynamics

MDFF Molecular Dynamics Flexible Fitting

HMMM Highly Mobile Membrane Mimetic

VMD Visual Molecular Dynamics

BrownianMover Brownian Dynamics

VMD Visual Molecular Dynamics

NAMD Scalable Molecular Dynamics

BD: Brownian Dynamics

LatticeMicrobes Whole Cell Simulations

VMD Visual Molecular Dynamics

RDME: Reaction-diffusion master equation

NAMD: Scalable Molecular Dynamics

2002 Gordon Bell Award

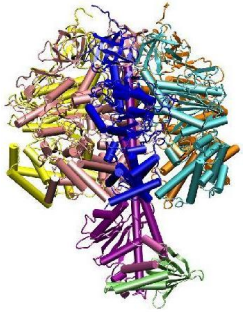

ATP synthase

PSC Lemieux

57,000 Users, 3000 Citations

Computational Biophysics Summer School

Blue Waters Target Application

Illinois Petascale Computing Facility

GPU Acceleration

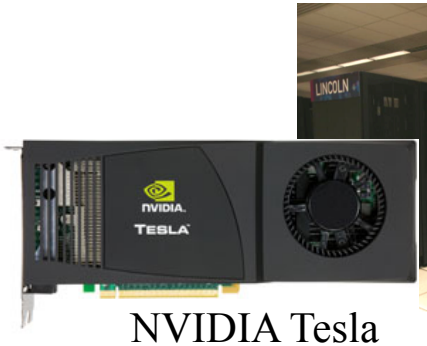

NVIDIA Tesla

NCSA Lincoln

Biomedical Technology Research Center for Macromolecular Modeling and Bioinformatics
http://www.ks.uiuc.edu/

Beckman Institute, UIUC

VMD – "Visual Molecular Dynamics"

- Visualization and analysis of molecular dynamics simulations, sequence data, volumetric data, quantum chemistry simulations, particle systems, …
- User extensible with scripting and plugins
- http://www.ks.uiuc.edu/Research/vmd/

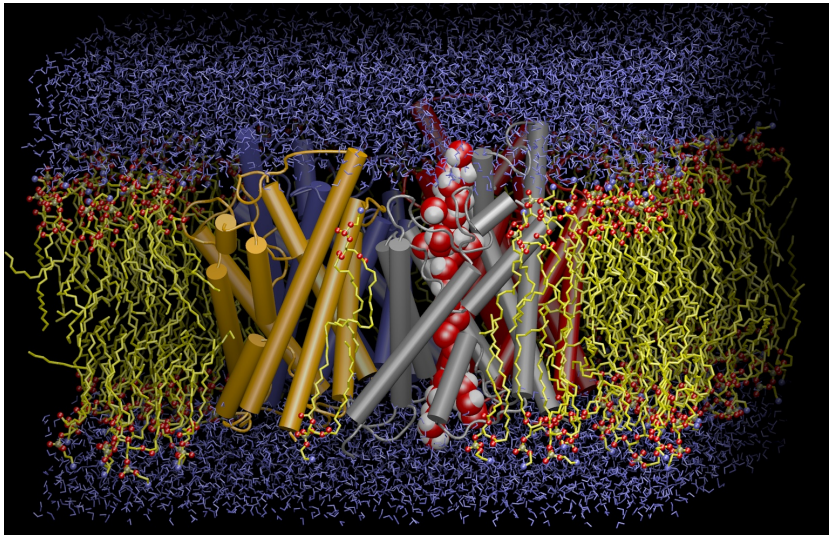

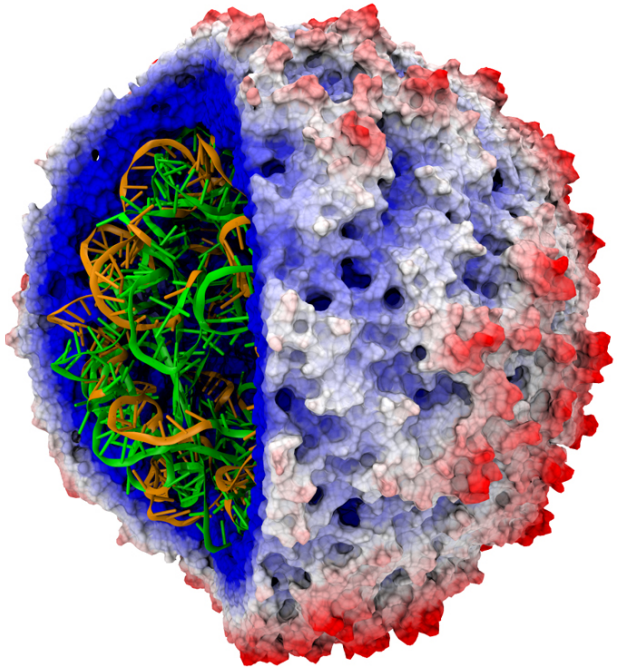

NAMD impact is broad and deep

- Comprehensive, industrial-quality software
  - Integrated with VMD for simulation setup and analysis
  - Portable extensibility through Tcl scripts (also used in VMD)
  - Consistent user experience from laptop to supercomputer
- Large user base – 57,000 users
  - 10,300 (18%) are NIH-funded; many in other countries
  - 16,600 have downloaded more than one version
- Leading-edge simulations
  - "most-used software" on NICS Cray XT5 (largest NSF machine)
  - "by far the most used MD package" at TACC (2nd and 3rd largest)
  - NCSA Blue Waters early science projects and acceptance test
  - Argonne Blue Gene/Q early science project

Outside researchers choose NAMD and succeed

3000 external citations since 2007

Corringer, et al., *Nature*, 2011

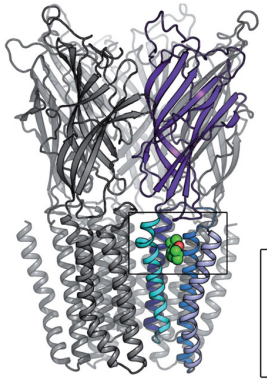

180K-atom 30 ns study of anesthetic binding to bacterial ligand-gated ion channel provided "complementary interpretations…that could not have been deduced from the static structure alone."

Bound Propofol Anesthetic

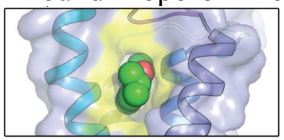

Voth, et al., *PNAS*, 2010

500K-atom 500 ns investigation of effect of actin depolymerization factor/cofilin on mechanical properties and conformational dynamics of actin filament.

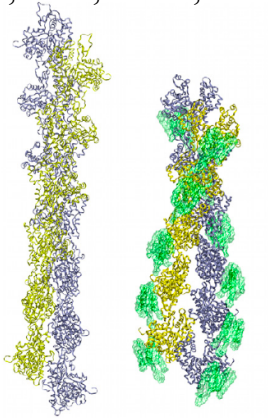

Bare actin Cofilactin

Recent NAMD Simulations in *Nature*

- **M. Koeksal, et al.,** *Taxadiene synthase structure and evolution of modular architecture in terpene biosynthesis*. (2011)
- **C.-C. Su, et al.**, *Crystal structure of the CusBA heavy-metal efflux complex of Escherichia coli.* (2011)
- **D. Slade, et al.**, *The structure and catalytic mechanism of a poly(ADP-ribose) glycohydrolase*. (2011)
- **F. Rose, et al.,** *Mechanism of copper(II)-induced misfolding of Parkinson's disease protein*. (2011)
- **L. G. Cuello, et al.**, *Structural basis for the coupling between activation and inactivation gates in K(+) channels*. (2010)
- **S. Dang, et al.,**, *Structure of a fucose transporter in an outward-open conformation.* (2010)
- **F. Long, et al.**, *Crystal structures of the CusA efflux pump suggest methionine-mediated metal transport*. (2010)
- **R. H. P. Law, et al.**, *The structural basis for membrane binding and pore formation by lymphocyte perforin.* (2010)
- **P. Dalhaimer and T. D. Pollard**, *Molecular Dynamics Simulations of Arp2/3 Complex Activation.* (2010)
- **J. A. Tainer, et al.**, *Recognition of the Ring-Opened State of Proliferating Cell Nuclear Antigen by Replication Factor C Promotes Eukaryotic Clamp-Loading*. (2010)

Computational Microscopy

Ribosome: synthesizes proteins from genetic information, target for antibiotics

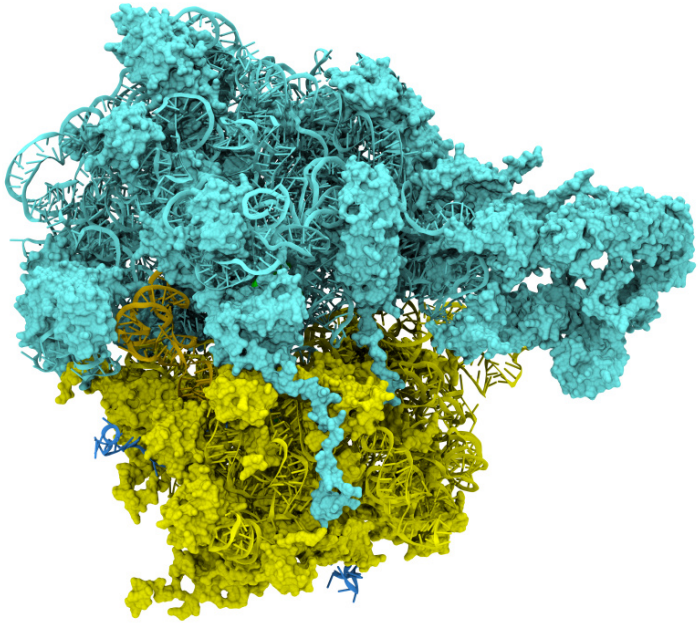

Silicon nanopore: bionanodevice for sequencing DNA efficiently

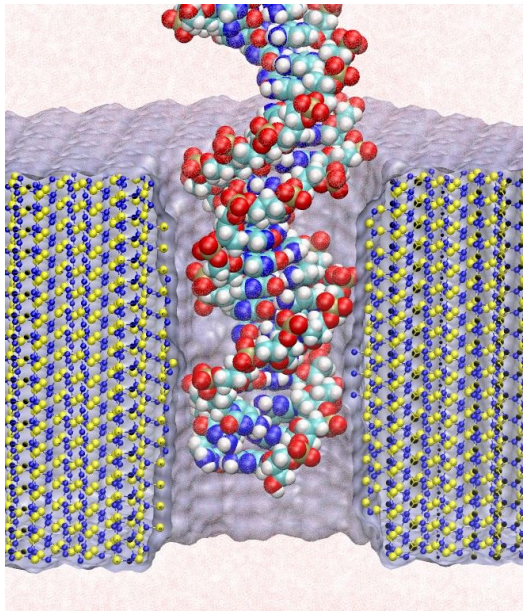

Molecular Mechanics Force Field

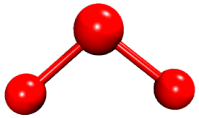

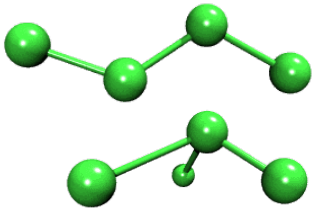

$$U(\vec{R}) = \underbrace{\sum_{bonds} k_i^{bond}(r_i - r_0)^2}_{U_{bond}} + \underbrace{\sum_{angles} k_i^{angle}(\theta_i - \theta_0)^2}_{U_{angle}} +$$

$$\underbrace{\sum_{dihedrals} k_i^{dihe}[1 + \cos(n_i\phi_i + \delta_i)]}_{U_{dihedral}} +$$

$$\underbrace{\sum_i \sum_{j \neq i} 4\epsilon_{ij}\left[\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right] + \sum_i \sum_{j \neq i} \frac{q_i q_j}{\epsilon r_{ij}}}_{U_{nonbond}}$$

Classical Molecular Dynamics

Energy function: $U(\vec{r}_1, \vec{r}_2, \cdots \vec{r}_N) = U(\vec{R})$

used to determine the force on each atom: $m_i \frac{d^2 \vec{r}_i}{dt^2} = \vec{F}_i = -\vec{\nabla} U(\vec{R})$

Newton's equation represents a set of N second order differential equations which are solved numerically via the Verlet integrator at discrete time steps to determine the trajectory of each atom. $\vec{r}_i(t + \Delta t) = 2\vec{r}_i(t) - \vec{r}_i(t - \Delta t) + \frac{\Delta t^2}{m_i} \vec{F}_i(t)$

Small terms added to control temperature and pressure.

Parallel Programming Lab

University of Illinois at Urbana-Champaign

Siebel Center for Computer Science

http://charm.cs.illinois.edu/

Develop abstractions in context of full-scale applications

Quantum Chemistry (QM/MM)

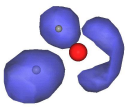

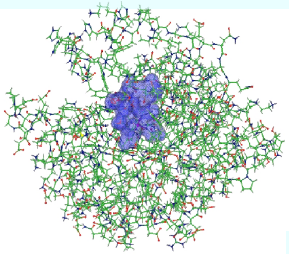

NAMD: Molecular Dynamics

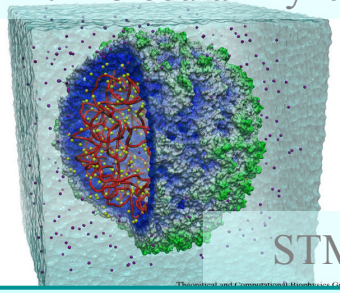

Protein Folding

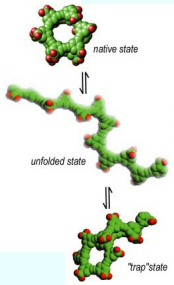

Computational Cosmology

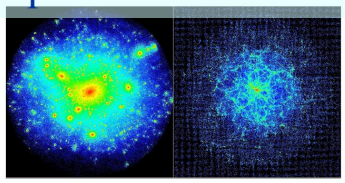

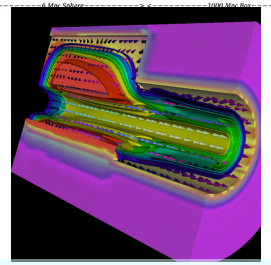

STM virus simulation

Parallel Objects,

Adaptive Runtime System

Libraries and Tools

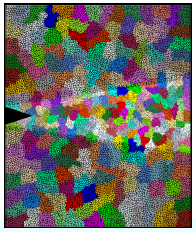

Crack Propagation

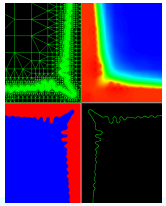

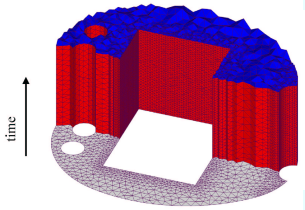

Rocket Simulation

Dendritic Growth

Space-time meshes

PPL Partnership

- NIH-funded collaboration since 1992
  - C++ was new and poorly understood
  - Charm++ was born (from Chare kernel)
  - PVM was popular (original NAMD target)
  - MPI was new
- Symbiotic relationship
  - PPL provides tools and expertise
  - NAMD provides feedback and use cases

1994: Our First Cluster

- 14 HP workstations
- 125 MHz processor
- 128 MB memory
- 100 Mbit network (optical ATM switch)
- ~ $20K per processor (cheaper than CM5!)

1994: Our First NAMD

- Written in C++
- Parallelized with PVM
- Spatial decomposition
- Message driven
- DPMTA electrostatics
- ~ 10,000 atom systems
- Up to ~ 8 processors

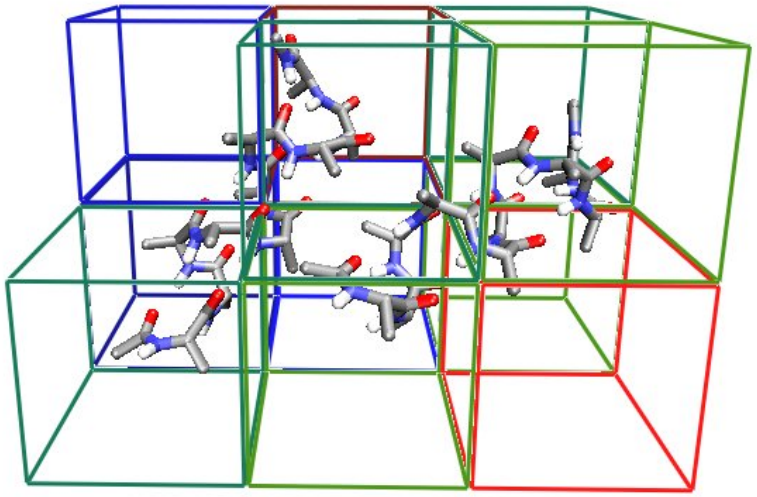

MDScope (1994)

- NAMD – dynamics on workstation clusters
  - PPL collaboration just started
  - Student-led revolt to write new code
  - Not A…'s Molecular Dynamics
- VMD – visualization
  - Originally VRChem for CAVE
  - NIH told us not to do visualization
  - Existing codes poor for trajectories
  - Scriptable interface, eventually Tcl
  - Andrew model: wrap other tools
  - Bill model: write code ourselves
- MDCOMM – steering communication
  - Collaboration with NCSA
  - Eventually replaced with raw sockets

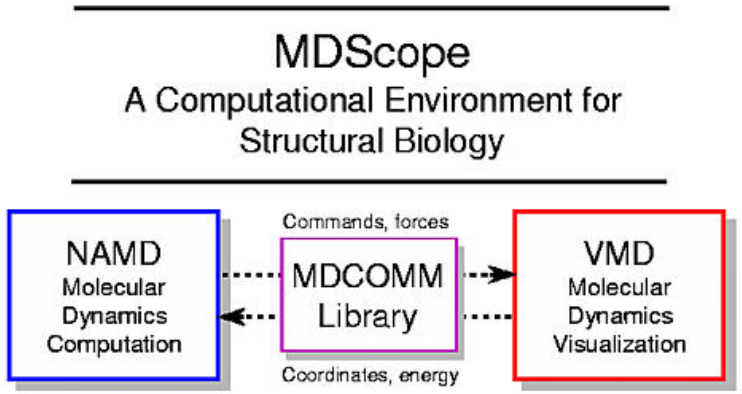

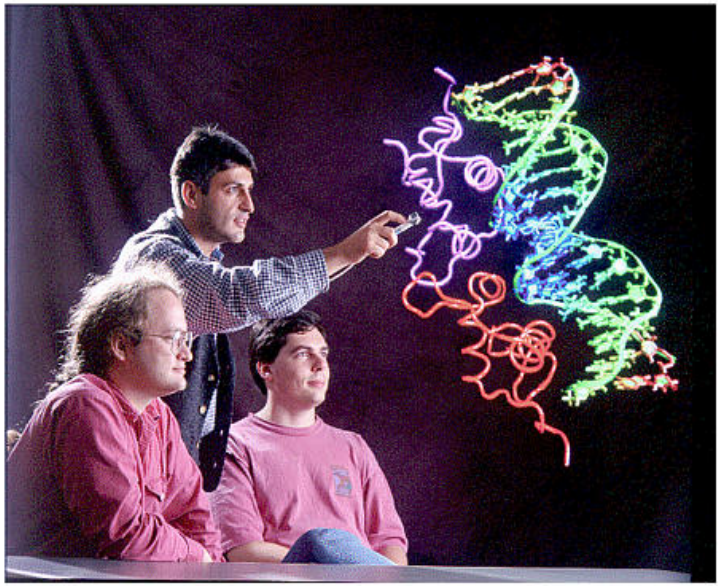

Nelson, M., Humphrey, W., Kufrin, R., Gursoy, A., Dalke, A., Kale, L., Skeel, R., and Schulten, K. *Comput. Phys. Commun.* **91**, 111-134.

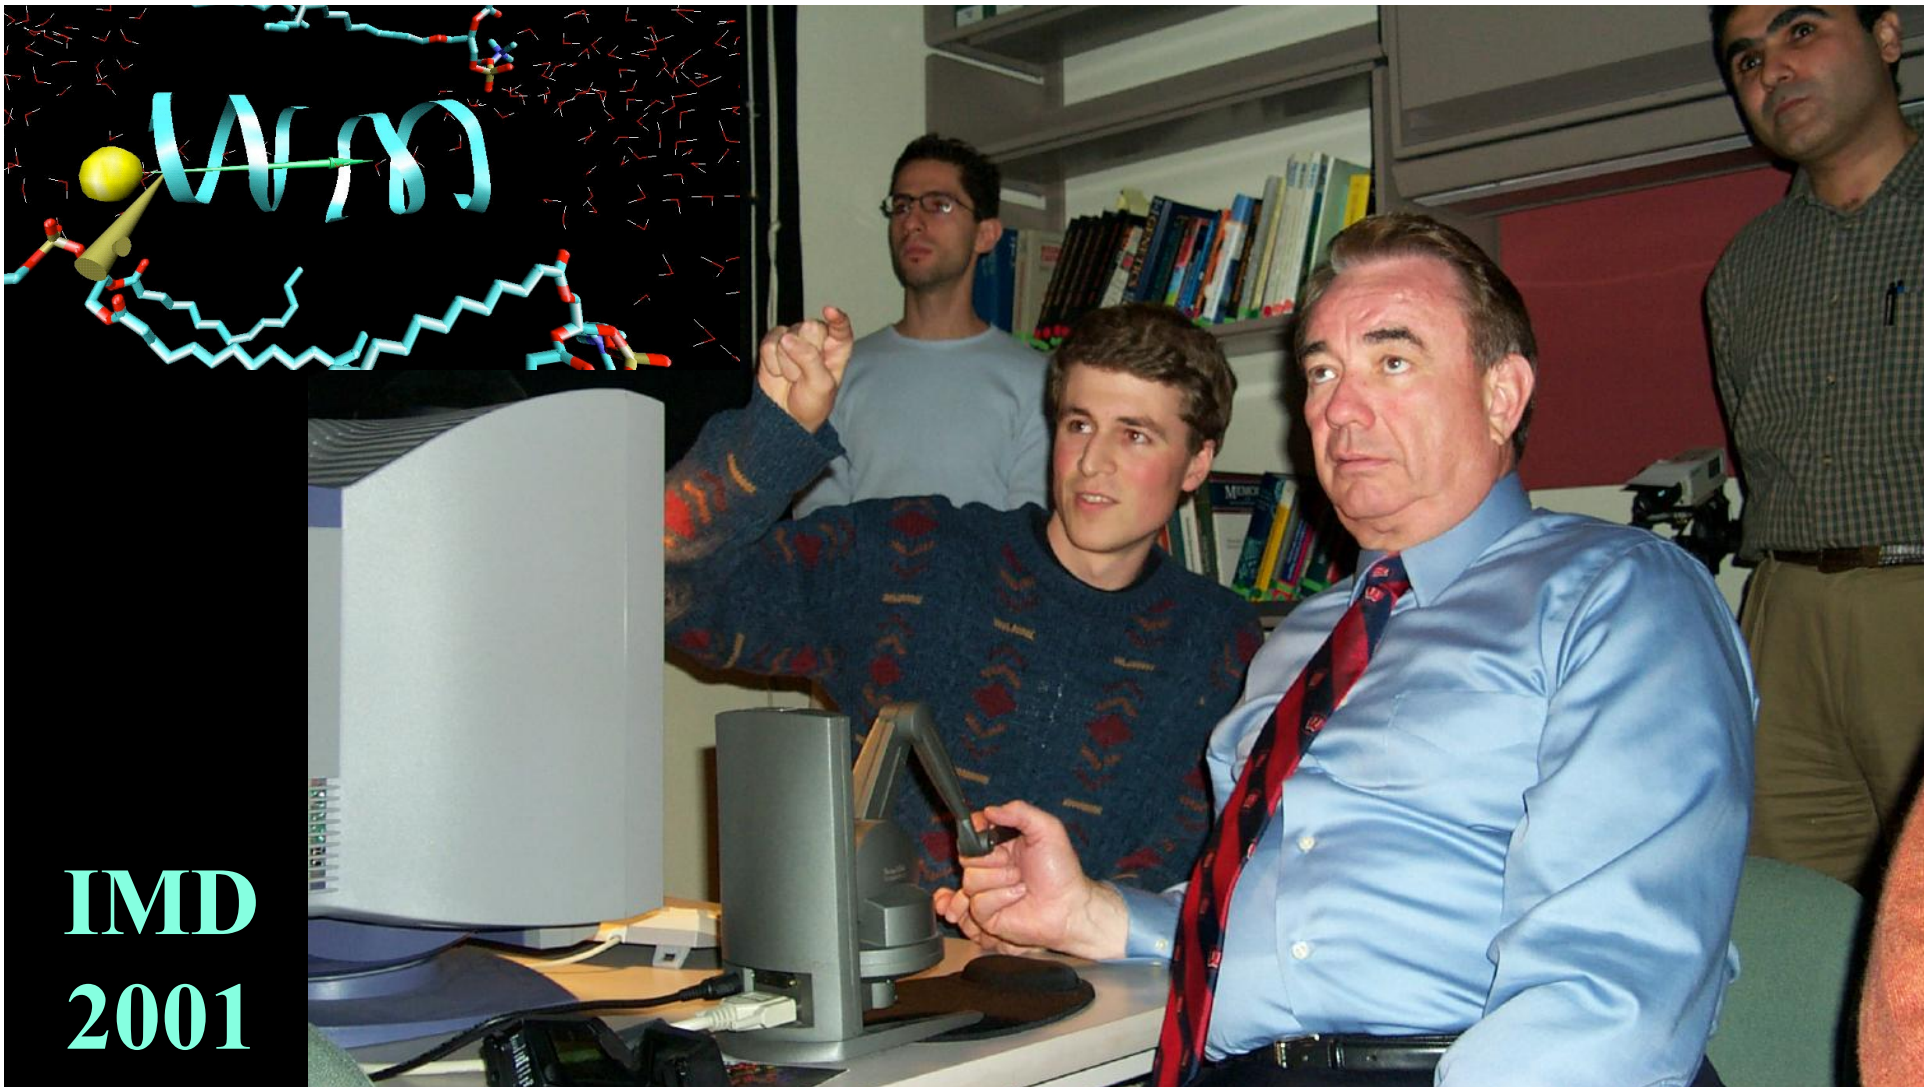
IMD
2001

Sizes of Simulations Over Time

BPTI
3K atoms

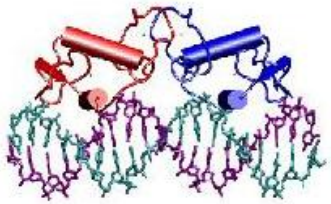

Estrogen Receptor
36K atoms (1996)

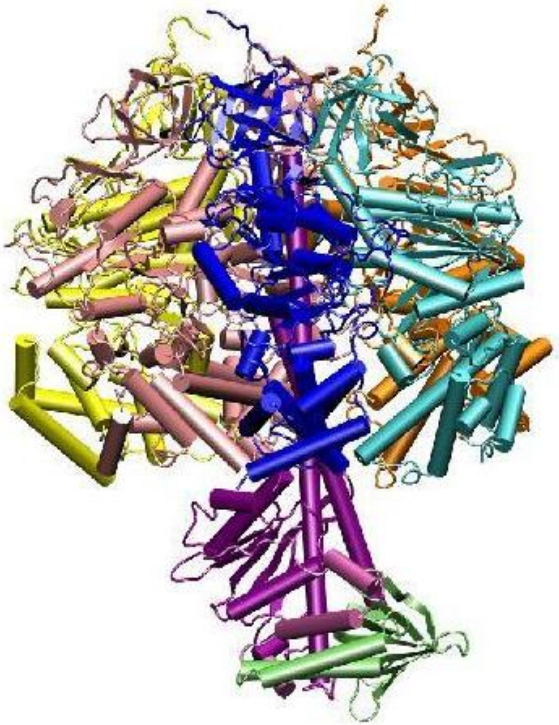

ATP Synthase
327K atoms
(2001)

Our Solution: Parallel Computing

HP 735 cluster
14 processors
(1994)

SGI Origin 2000
128 processors (1997)

PSC Lemieux AlphaServer SC
3000 processors (2002)

Poorly Scaling Approaches

- Replicated data
  - All atom coordinates stored on each processor
  - Communication/Computation ratio: O(P log P)
- Partition the atom array across processors
  - Nearby atoms may not be on the same processor
  - C/C ratio: O(P)
- Distribute force matrix to processors
  - Matrix is sparse, non uniform
  - C/C Ratio: O(sqrt P)

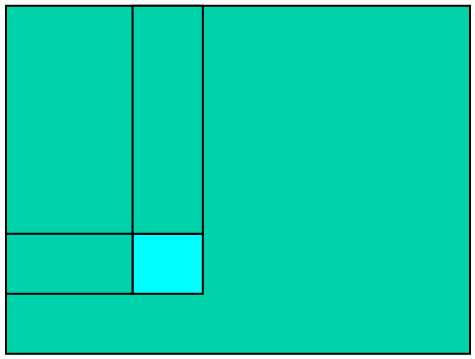

NAMD 1 Spatial Decomposition

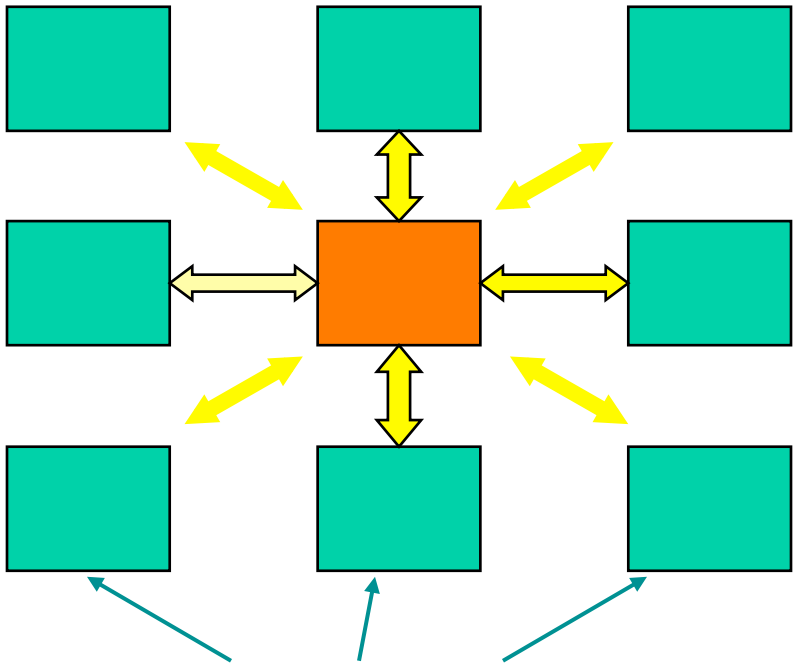

- Atoms spatially distributed to cubes
  - Multiple cubes per processor.
  - Early example of virtualization!
- Size of each cube :
  - Just a larger than cut-off radius
  - Communicate only w/ neighbors
  - Work for each pair of neighbors
- C/C ratio: O(1)
- However:
  - Load Imbalance
  - Limited Parallelism

NAMD 2 Hybrid Decomposition

Kale *et al., J. Comp. Phys.* **151**:283-312, 1999.

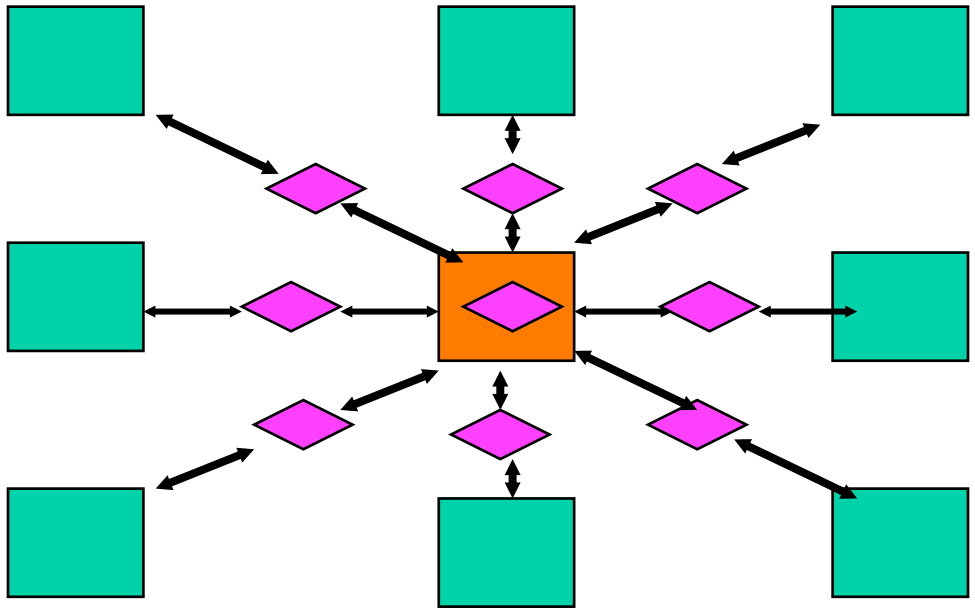

- Spatially decompose data and communication.
- Separate but related work decomposition.
- "Compute objects" facilitate iterative, measurement-based load balancing system.

NAMD Overlapping Execution

Phillips *et al., SC2002.*

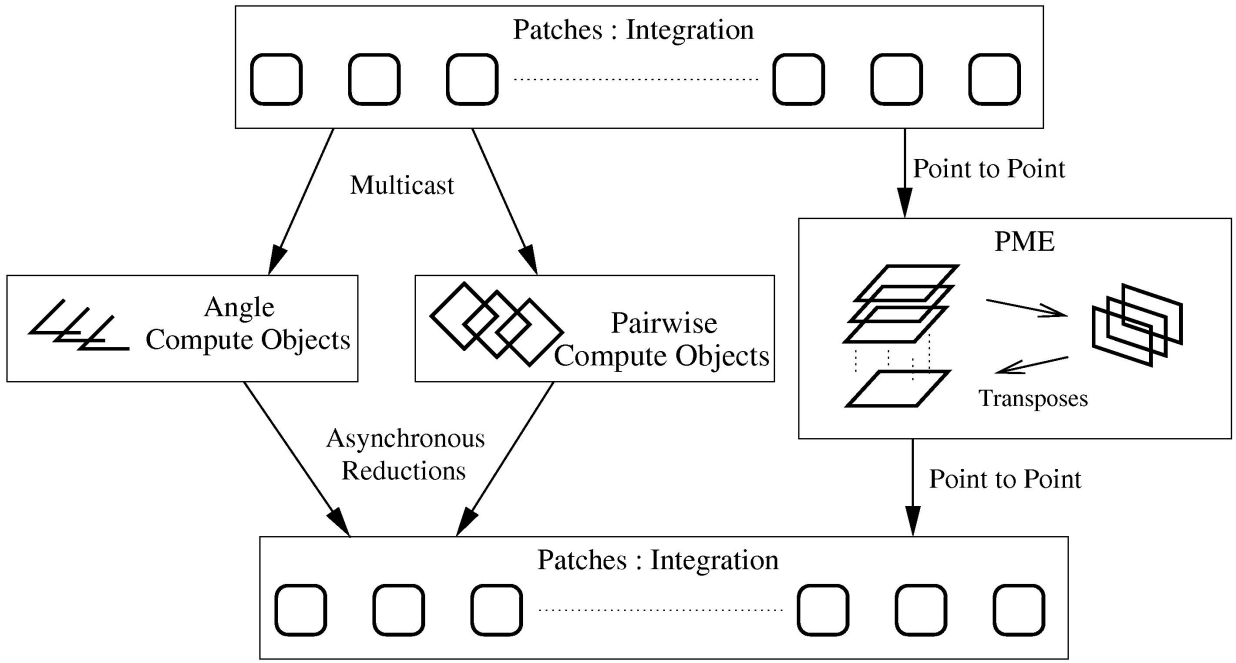

Objects are assigned to processors and queued as data arrive

Implemention in 1997 Charm++

- Parallel C++ with *data driven* objects.
- Object groups:
  - Global object with a "representative" on each PE.
- Asynchronous method invocation.
- Prioritized scheduling of messages/execution.
- Measurement-based load balancing.
- Portable messaging layer.

Startup Phases

- In parallel after reading input data on pe 0
- Used to load data, construct and link objects
- Phases separated by quiescence detections
- Simplest way to ensure side effects done
- Allows diagnosis of:
  - What operation is crashing or hanging
  - What operation is using memory or time

Coordinated Message Priorities

- Computes enqueued as messages
- Priorities based on critical path
  - Earlier step before later step
  - Earlier PME stages before later stages
  - Computes for remote patches before local
- Modules must be coordinated!
  - All priorities defined in Priorities.h
  - Diagnose and confirm using Projections

2006 NAMD Performance

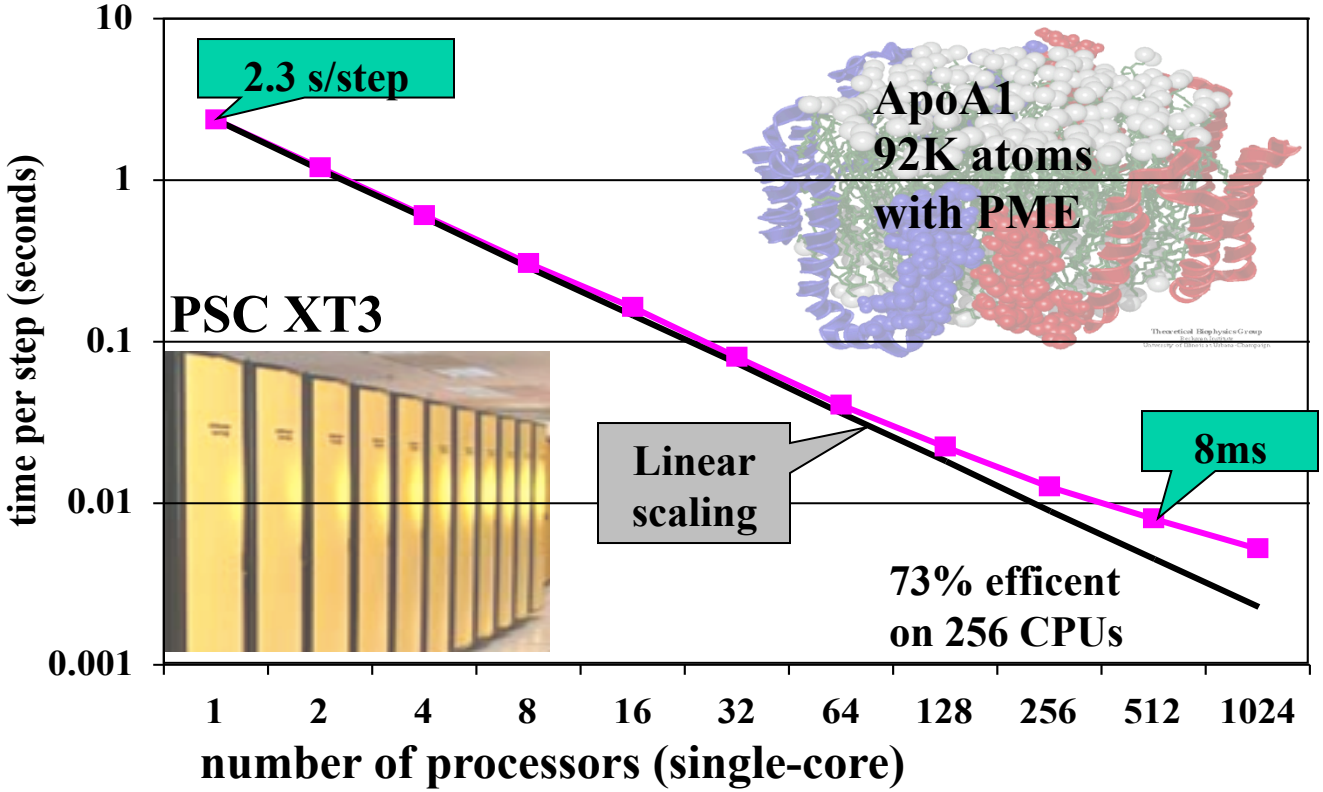

The Grid (2006)

- One driving project needed to run a “large” number of similar simulations
- Collaboration with NCSA
- Utilize Globus tools
- Result “maybe 30% more efficient…after getting everything set up”
- Never used again

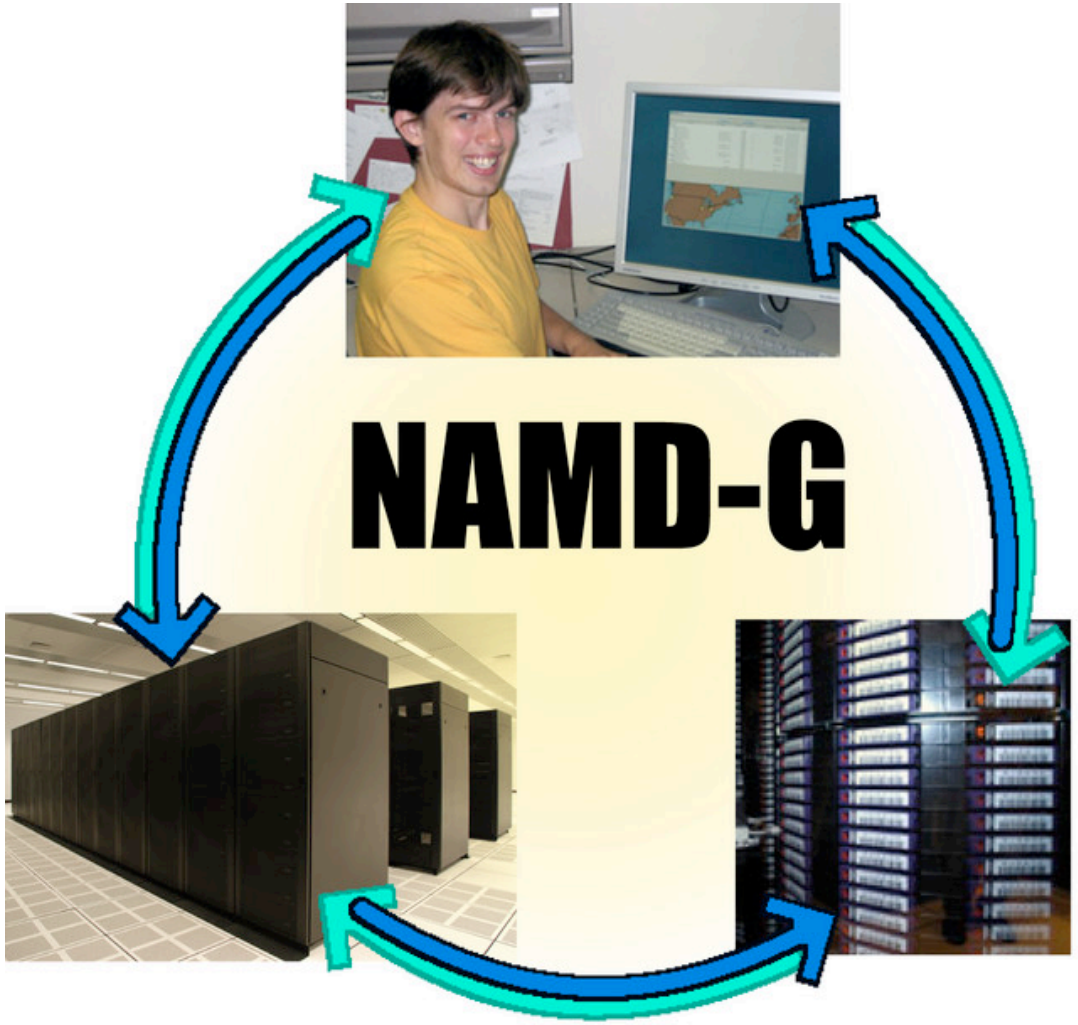

Early Acceleration Options

- Outlook in 2005-2006:
  - FPGA reconfigurable computing (with NCSA)
    - Difficult to program, slow floating point, expensive
  - Cell processor (NCSA hardware)
    - Relatively easy to program, expensive
  - ClearSpeed (direct contact with company)
    - Limited memory and memory bandwidth, expensive
  - MDGRAPE
    - Inflexible and expensive
  - Graphics processor (GPU)
    - Program must be expressed as graphics operations


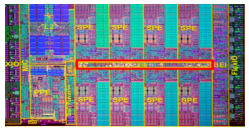


CUDA: Practical Performance

November 2006: NVIDIA announces CUDA for G80 GPU.

- CUDA makes GPU acceleration usable:
  - Developed and supported by NVIDIA.
  - No masquerading as graphics rendering.
  - New shared memory and synchronization.
  - No OpenGL or display device hassles.
  - Multiple processes per card (or vice versa).
- Center and collaborators make it useful:
  - Experience from VMD development
  - David Kirk (Chief Scientist, NVIDIA)
  - Wen-mei Hwu (ECE Professor, UIUC)

Stone *et al.*, *J. Comp. Chem.* **28**:2618-2640, 2007.

2006 “Large” Simulation Performance

- 2.7 M atoms
- 2048-core XT3
- 35 ms per step
- 2.1 teraflop/s
  - 1 gigaflop/s per core
- 1300 atoms/core
  - 2-3 patches per core
  - near-linear scaling

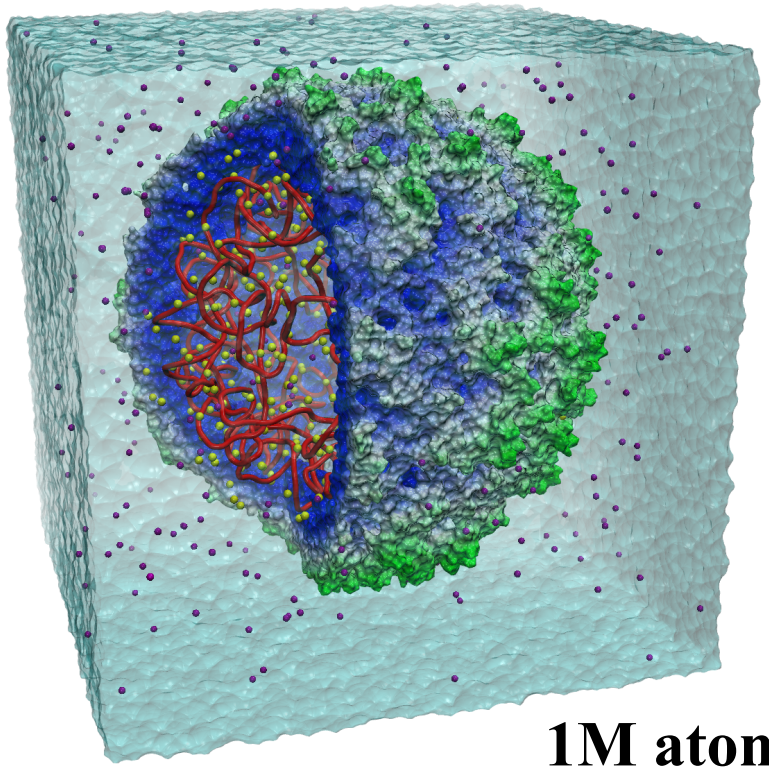

1M atom virus

2006 NSF Model Problem

- 100 M atoms
  - 1000A cell
  - 238,000 patches
  - 1024x1024x1024 FFT
- 10 ns per 25 hours
  - 2 fs timestep
  - 18 ms per step
- 151 teraflop/s
  - 150,000 cores?
- half time per step
  - 200,000 cores?
- 1 patch per core (max)
  - 250,000 cores
  - 15 ms per step

Larger machines enable larger simulations

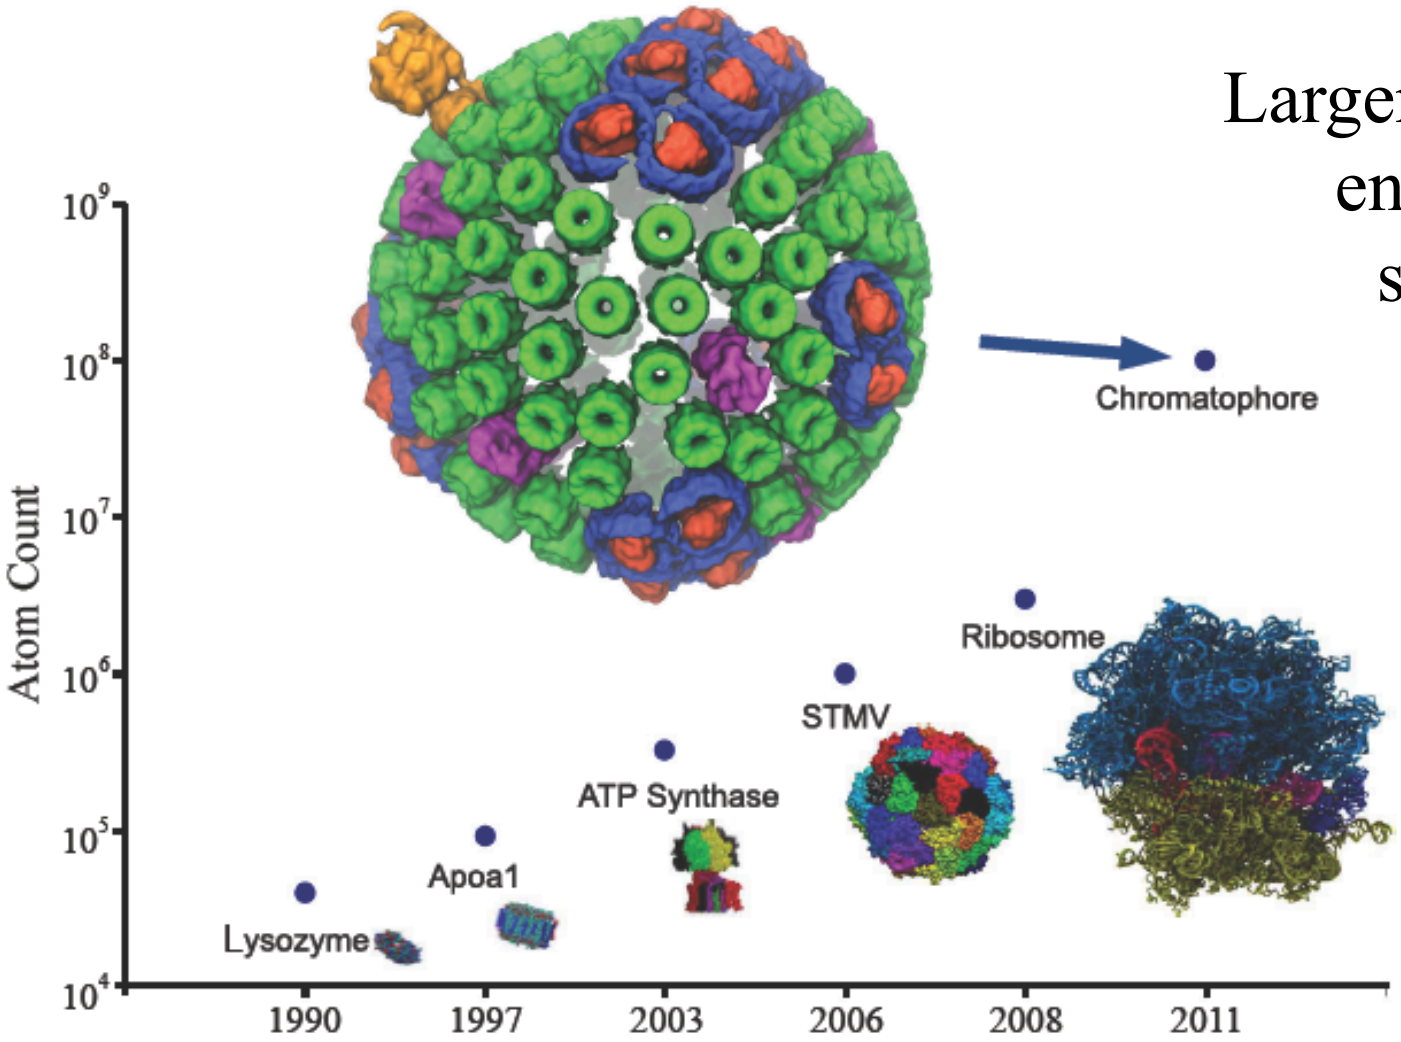

Beckman Institute, UIUC

Gearing Up for Petascale

- 2006 - NSF calls for 100 million atom simulation
  - Had just published million-atom virus simulation
- Issues to address:
  - Find scientific questions worthy of resources
  - Build model and initial coordinates
  - Store output trajectory
  - Analyze output trajectory
  - Scale NAMD to 100 million atoms
  - Scale NAMD to petascale machine(s)

2010 Schulten PRAC Simulations

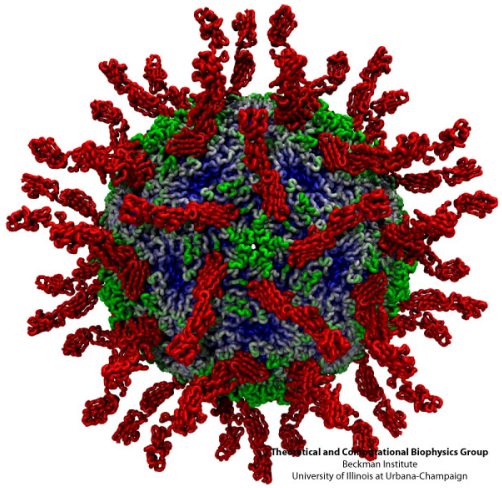

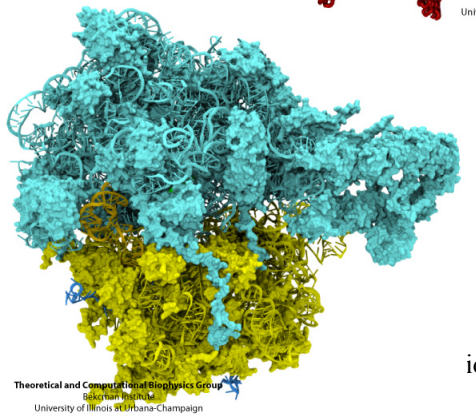

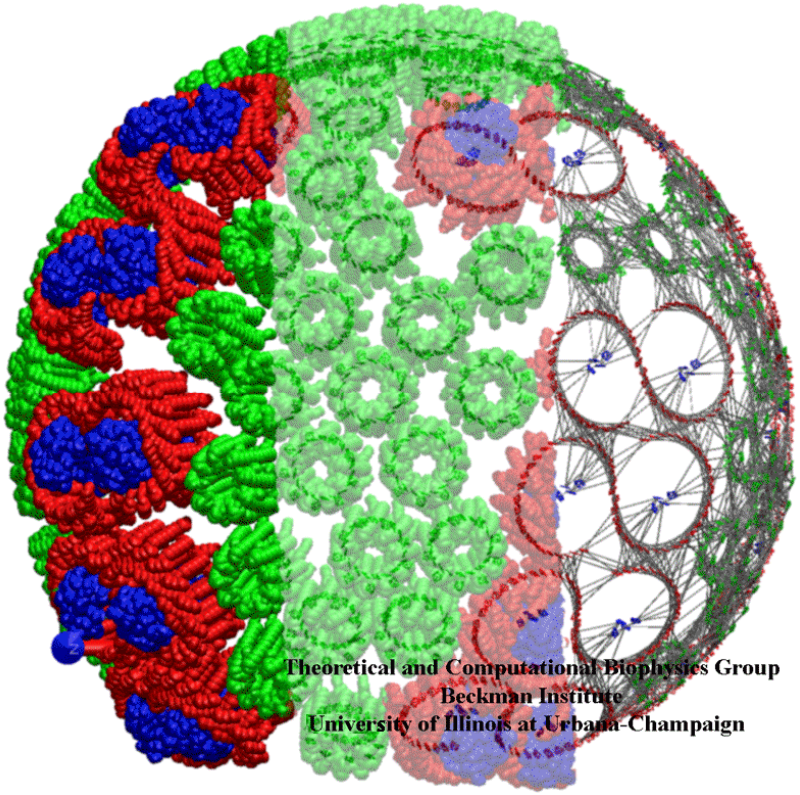

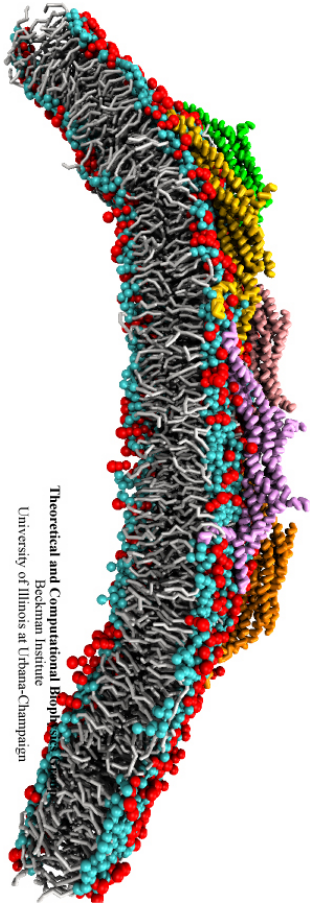

iomedical Technology Research Center for Macromolecular Modeling and Bioinformatics
http://www.ks.uiuc.edu/

Petascale Gateway Facility

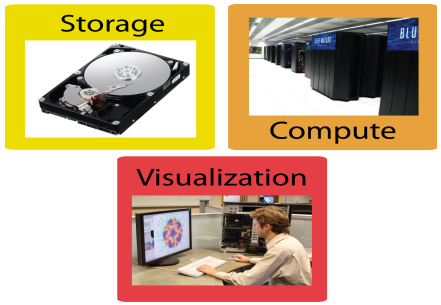

Petascale computers will create ~100x larger datasets. The Petascale Gateway provides the necessary storage and analysis facilities all located in the Beckman Institute.

Stereo & High Resolution Visualization Facility
3D & High-Resolution Projection

Storage & Archive
500 TB
120 TB Archive System

1-10 Gigabit Network

External Resources, 90% of our Computer Power

Simulation Preparation & Analysis Clusters
4+6=10 Million SUs/year & 28 GPUs

Visitor, Researcher & Developer Workstations

VMD 1.9.1: Improved Support for Large Datasets

- New structure building tools, file formats, and data structures enable VMD to operate efficiently up to 100M atoms
  - Up to 30% more memory efficient
  - Analysis routines optimized for large structures, up to 20x faster for calculations on 100M atom complexes
  - New and revised graphical representations support smooth trajectory animation for multi-million atom complexes; VMD remains interactive even when displaying surface reps for 20M atom membrane patch
- Uses multi-core CPUs and GPUs for the most demanding computations

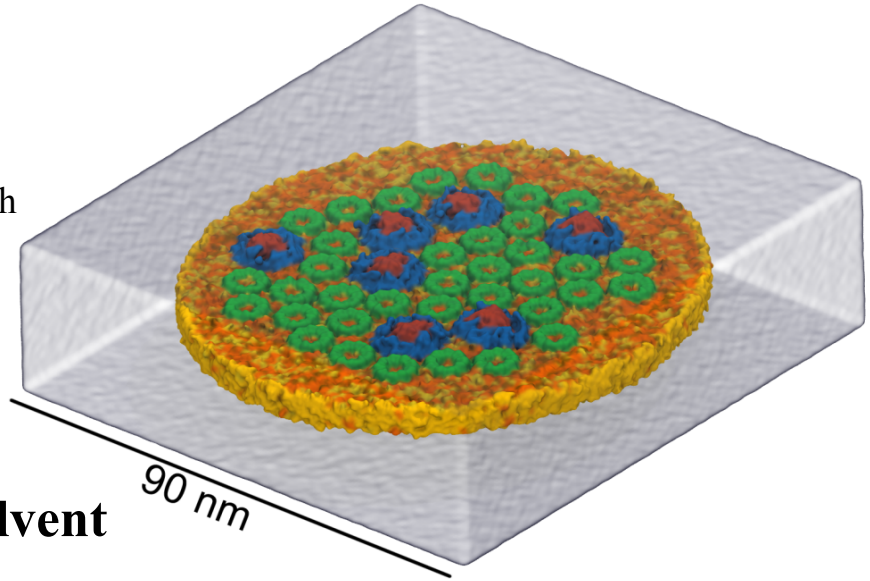

20M atoms:
membrane patch and solvent

VMD Interactive Display & Analysis of Terabytes of Data:

Out-of-Core Trajectory I/O w/ Solid State Disks

450MB/sec to 4GB/sec

A DVD movie per second!

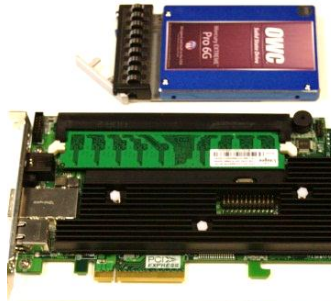

Commodity SSD, SSD RAID

- Timesteps loaded on-the-fly (out-of-core)
  - Eliminates memory capacity limitations, even for multi-terabyte trajectory files
  - High performance achieved by new trajectory file formats, optimized data structures, and efficient I/O
- Analyze long trajectories significantly faster using just a personal computer

Immersive out-of-core visualization of large-size and long-timescale molecular dynamics trajectories. J. Stone, K. Vandivort, and K. Schulten.
***Lecture Notes in Computer Science*, 6939:1-12, 2011.**

VMD Out-of-Core Trajectory I/O Performance: SSD-Optimized Trajectory Format, 8-SSD RAID

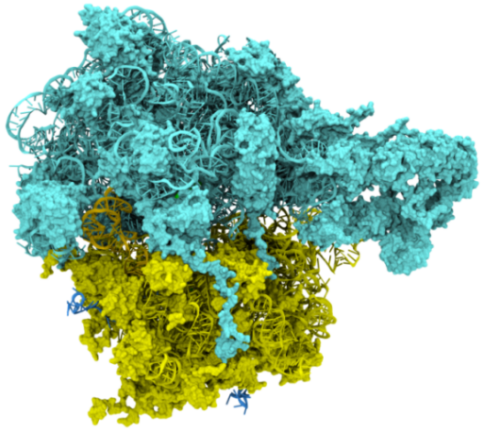

Ribosome w/ solvent
3M atoms
3 frames/sec w/ HD
60 frames/sec w/ SSDs

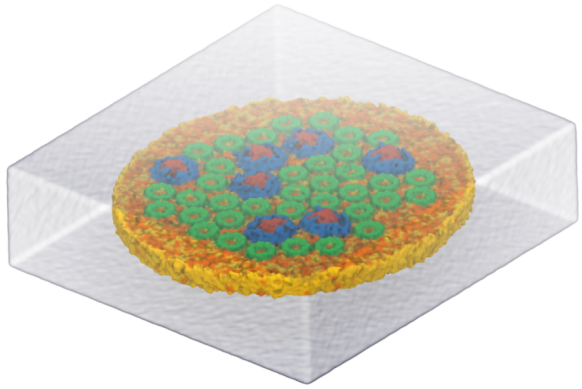

Membrane patch w/ solvent
20M atoms
0.4 frames/sec w/ HD
8 frames/sec w/ SSDs

New SSD Trajectory File Format 2x Faster vs. Existing Formats
VMD I/O rate ~2.1 GB/sec w/ 8 SSDs

VMD for Demanding Analysis Tasks
Parallel VMD Analysis w/ MPI

- Analyze trajectory frames, structures, or sequences in parallel on clusters and supercomputers:
  - Compute hydrogen bonding during lambda repressor folding
  - Parallel rendering, movie making
- Addresses computing requirements beyond desktop
- User-defined parallel reduction operations, data types
- Dynamic load balancing:
  - Tested with up to 15,360 CPU cores
- **Supports GPU-accelerated clusters and supercomputers**

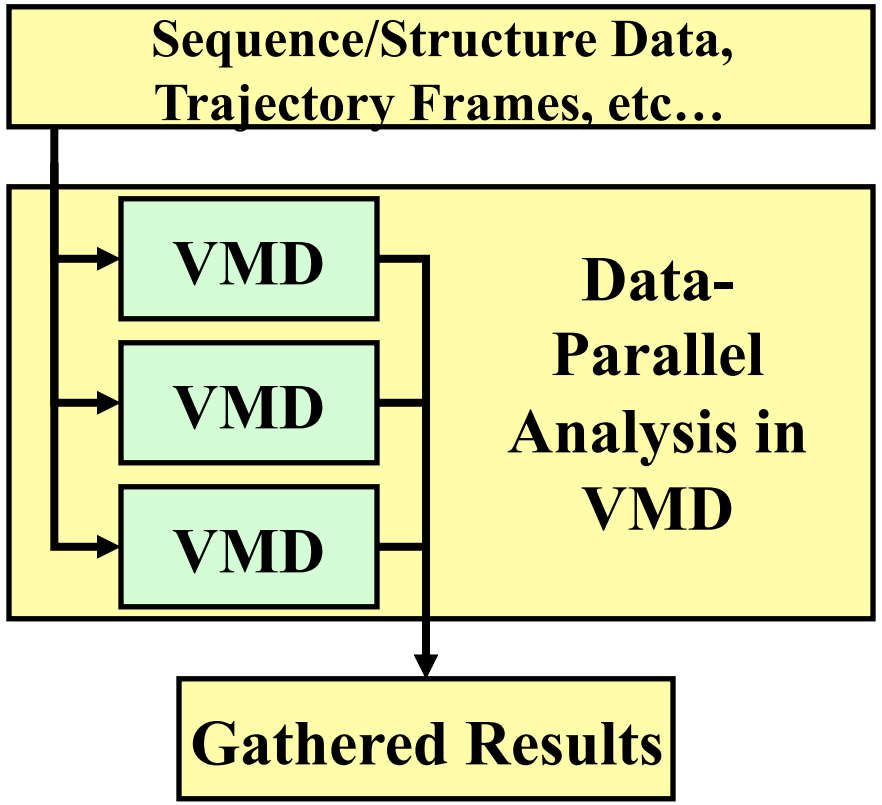

NAMD for Large Systems

- Per-node memory usage
  - Exploit redundant structure
  - Pre-compressed static data
  - Distributed per-atom data
  - Special “memopt” build
  - Not all features supported
  - NAMD-only file formats
    - May change between versions
    - No VMD reader/writer - ever
- I/O performance
  - Data is relatively small
  - Parallelized POSIX I/O
  - Performance is just OK
  - New Charm++ I/O library being co-developed
- Parallelize load balancer
  - Local load balancing only

Cray Gemini Optimization

- The new Cray machine has a better network (called Gemini)
- MPI-based NAMD scaled poorly
- Direct port of Charm++ to Cray uGNI layer
  - Removes MPI from NAMD call stack

ApoA1 (92,000 atoms) on Cray Blue Waters prototype

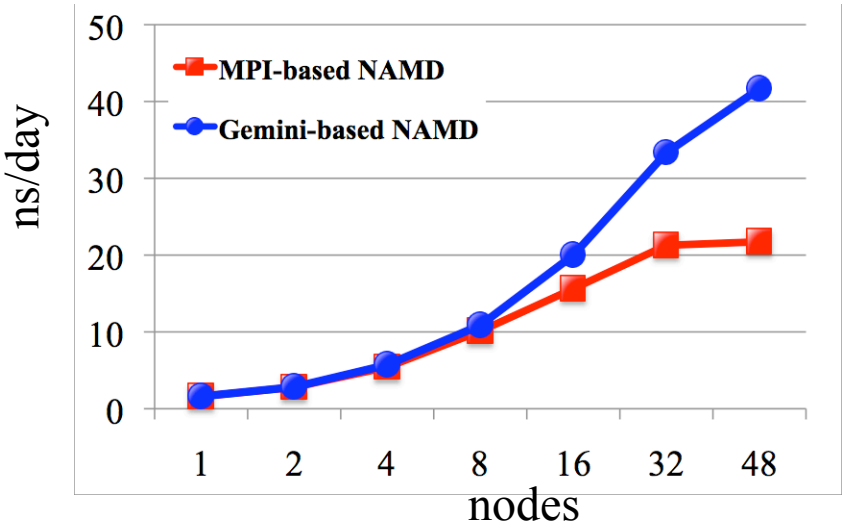

Gemini provides at least 2x increase in usable nodes for strong scaling

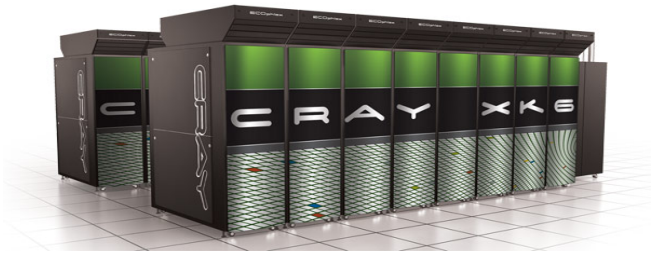

100M Atoms on Titan vs Jaguar

Sun *et al.*,
SC2012

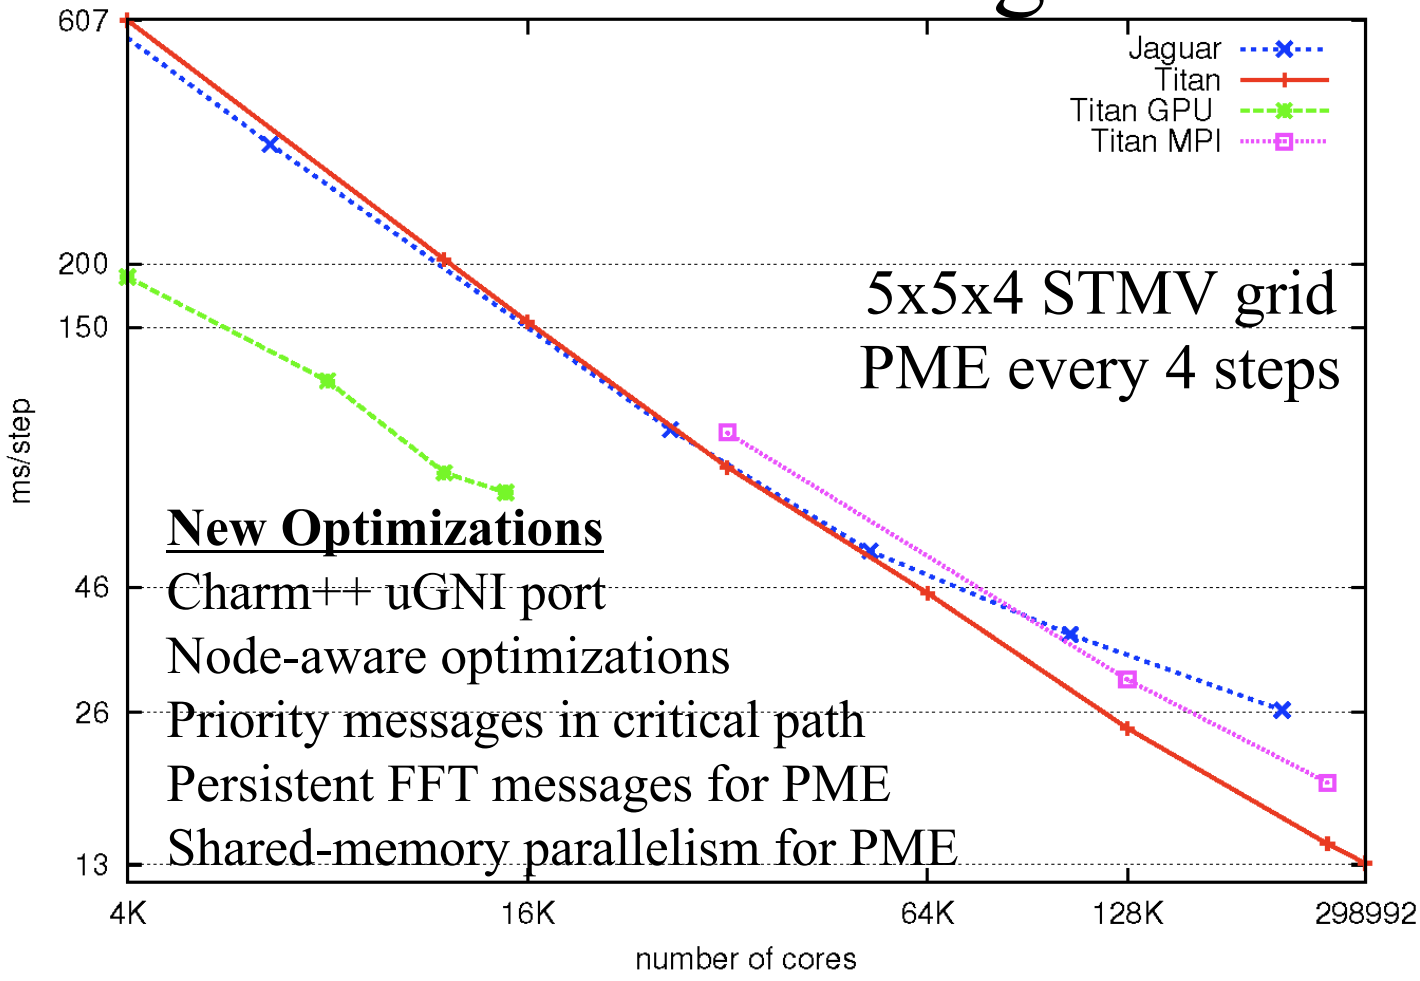

June 2013

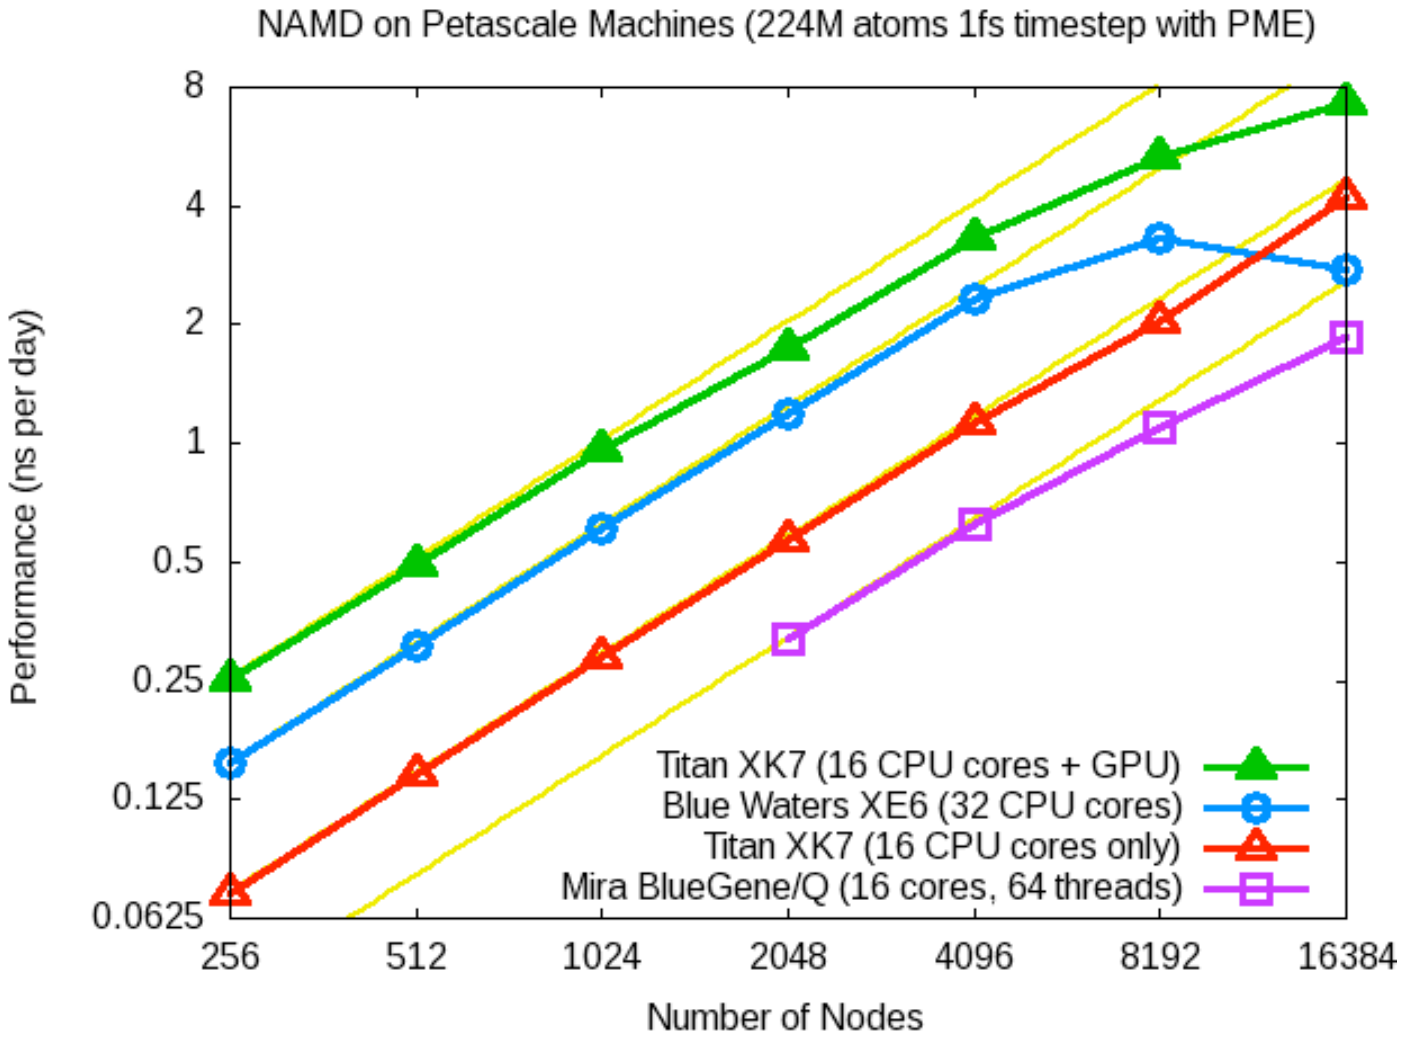

ATPESC 2013

NAMD: Practical Supercomputing

- 57,000 users can't all be computer experts.
  - 16,600 have downloaded more than one version.
  - 3000 citations of NAMD reference papers.
- One program for all platforms.
  - Desktops and laptops – setup and testing
  - Linux clusters – affordable local workhorses
  - Supercomputers – free allocations on XSEDE
  - Blue Waters – sustained petaflop/s performance
- User knowledge is preserved.
  - No change in input or output files.
  - Run any simulation on **any number of cores.**
- Available free of charge to all.

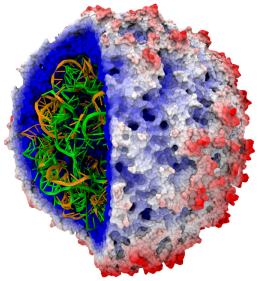

Phillips *et al.*, *J. Comp. Chem.* **26**:1781-1802, 2005.

Distribution and Licensing

- Binaries and source code
  - Charm++ included
- Annual releases
- Nightly builds
- Registration required
- Public CVS access available
- Installed on supercomputers

- No redistribution
- Citation required
- Registration required
- Use for any purpose
- Combine up to 10% of source with at least 50% original code without restriction
- VMD plugins use BSD license

Support and Training

- Public mailing list
  - Other scientists know best
  - Archived and searchable
  - Social conventions apply
- Bug report emails
- Personal support
  - Driving projects
  - New capabilities
- Tutorials and Case Studies
  - Written by scientists
  - Focus on science problems
- Hands-on workshops
  - Taught by scientists
  - Several per year
  - Various locations
  - Requires only laptop

Advertising

- Demonstrate science that can be done with the code
- Make sure all contributors can share credit:
  - NCSA
  - NSF
  - DOE
  - NVIDIA
  - PPL
  - Collaborators

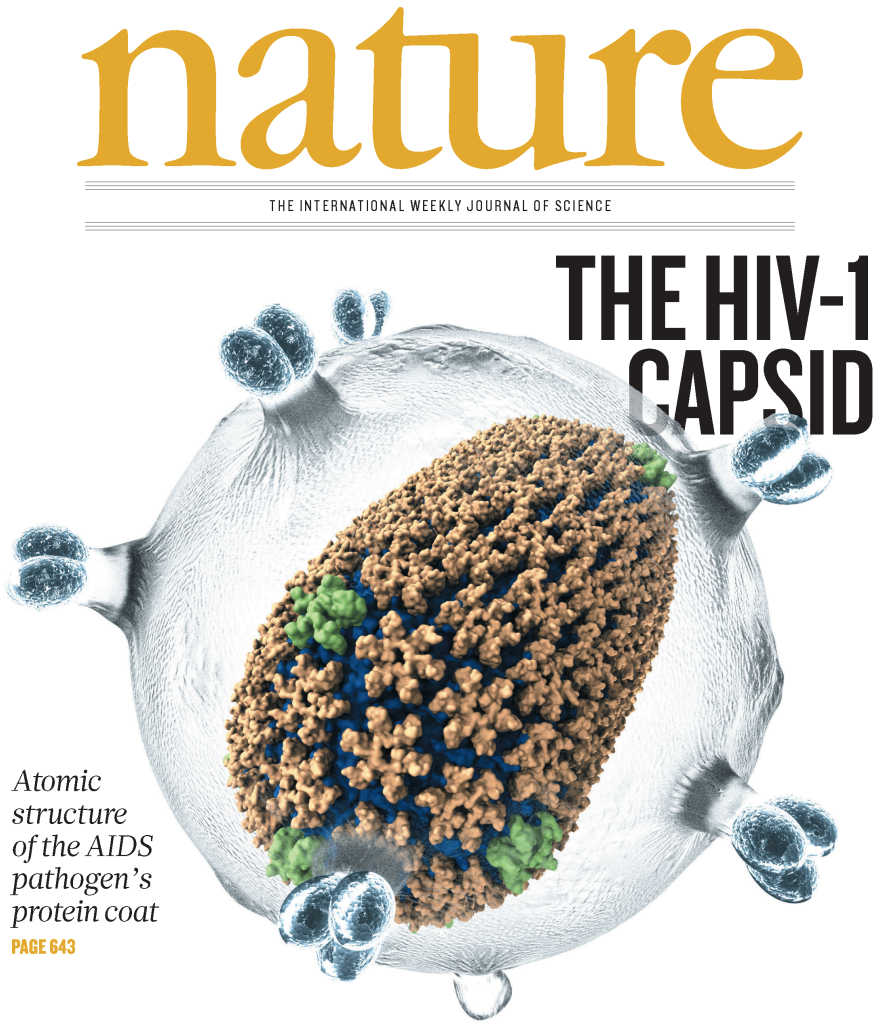

Development Process/Philosophy

- Five-year funding cycle
  - Code, science, publish, proposal
- Evolutionary development
  - Fully functional code at all times
  - No stable/development branches
  - Large changes by refactoring only
- Simplify – don't manage
  - Separation of responsibilities
  - Alignment of incentives
  - Low coupling between people
- No code without an eager user
- No single-user features
- No schedules, no promises
- No design/code documentation
  - Source code must be **discoverable**
  - Use sandboxes to hide complexity
- Priorities and opportunities
  - Enabling new science
  - Supporting outside developers

Collaborative Driving Projects

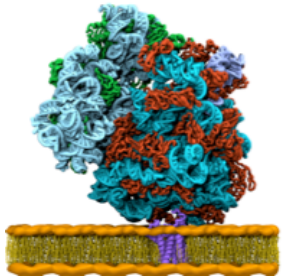

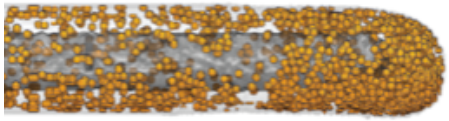

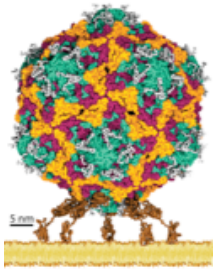

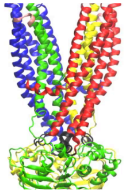

1. Ribosome	R. Beckmann (U. Munich) J. Frank (Columbia U.) T. Ha(UIUC) K. Fredrick (Ohio state U.) R. Gonzalez (Columbia U.)
2. Blood Coagulation Factors	J. Morrissey (UIUC) S. Sligar (UIUC) C. Rienstra (UIUC) G. Gilbert (Harvard)
3. Whole Cell Behavior	W. Baumeister (MPI Biochem.) J. Xiao (Johns Hopkins U.) C.N. Hunter (U. Sheffield) N. Price (U. Washington)
4. Biosensors	R. Bashir (UIUC) J. Gundlach (U. Washington) G. Timp (U. Notre Dame) M. Wanunu (Northeastern U.) L. Liu (UIUC)
5. Viral Infection Process	J. Hogle (Harvard U.) P. Ortoleva (Indiana U.) A. Gronenborn (U. Pittsburgh)
6. Integrin	T. Ha (UIUC) T. Springer (Harvard U.)
7. Membrane Transporters	H. Mchaourab (Vanderbilt U.) R. Nakamoto (U. Virginia) D.-N. Wang (New York U.) H. Weinstein (Cornell U.)

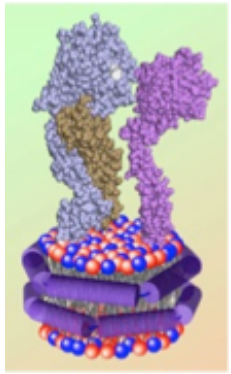

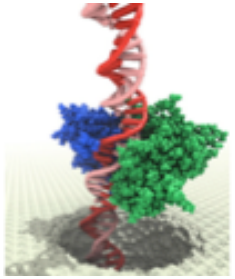

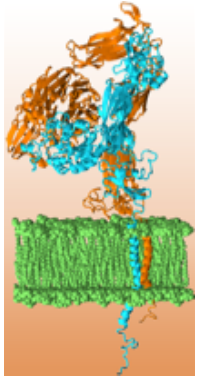

Collaborative Driving Projects

- Nearly every experimental collaboration relies on NAMD.
- High-end simulations push scaling efforts.
  - Try to anticipate needs: Million-atom virus just worked in 2006.
- Innovative simulations generate feature requests:

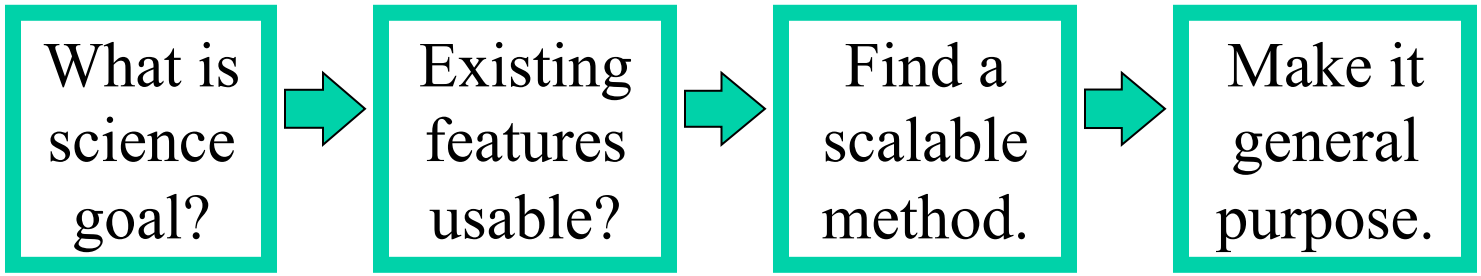

Adaptability Through Scripting

- Tcl customizations are **portable**
- Top-level protocols:
  - Minimize, heat, equilibrate
  - Simulated annealing
  - Replica exchange (originally via sockets)
- Long-range forces on selected atoms
  - Torques and other steering forces
  - Adaptive bias free energy perturbation
  - Coupling to external coarse-grain model
- Special boundary forces
  - Applies potentially to every atom
  - Several design iterations for efficiency
  - Shrinking phantom pore for DNA

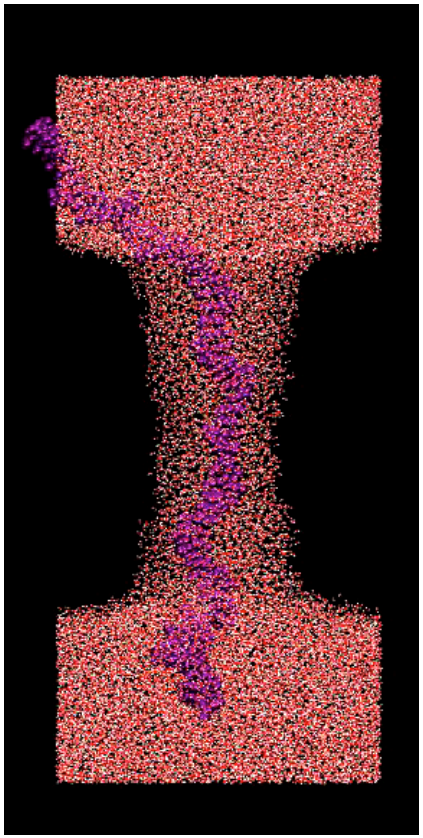

NAMD 2.9 Scalable Replica Exchange

- MPI-based partitions easier to use *and* more efficient:
  - Eliminates complex, machine-specific launch scripts
  - Scalable pair-wise communication between replicas
  - Fast communication via high-speed network
- Basis for many enhanced sampling methods:
  - Parallel tempering (temperature exchange)
  - Umbrella sampling for free-energy calculations
  - Hamiltonian exchange (alchemical or conformational)
  - Finite Temperature String method
  - Nudged elastic band

} Released in NAMD 2.9

- Great power *and* flexibility:
  - **Enables petascale simulations of modestly sized systems**
  - Leverages features of Collective Variables module
  - Tcl scripts can be highly customized and extended

First application of **parallel tempering** is CHARMM Drude-oscillator polarizable force field development by Alex MacKerell (U. Maryland)

Distribution of backbone dihedral angles at different temperatures from 64-replica simulation of Acetyl-(AAQAA)3-amide peptide on Blue Gene/**P**

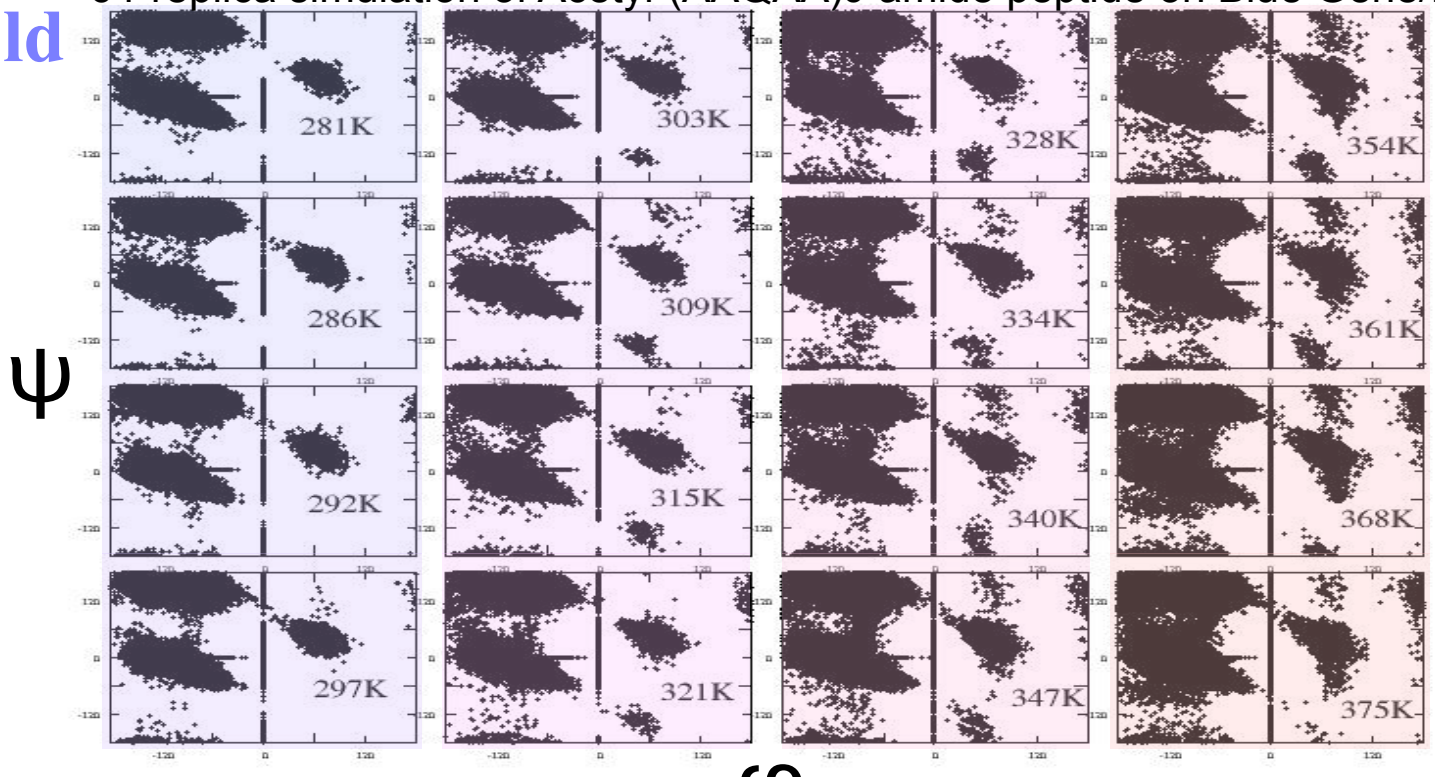

Data from Luo **&** Roux, ANL/UC.

Membrane Transporters – First application of replica exchange for umbrella sampling on collective variables

Quaternion-based order parameters from collective variables module

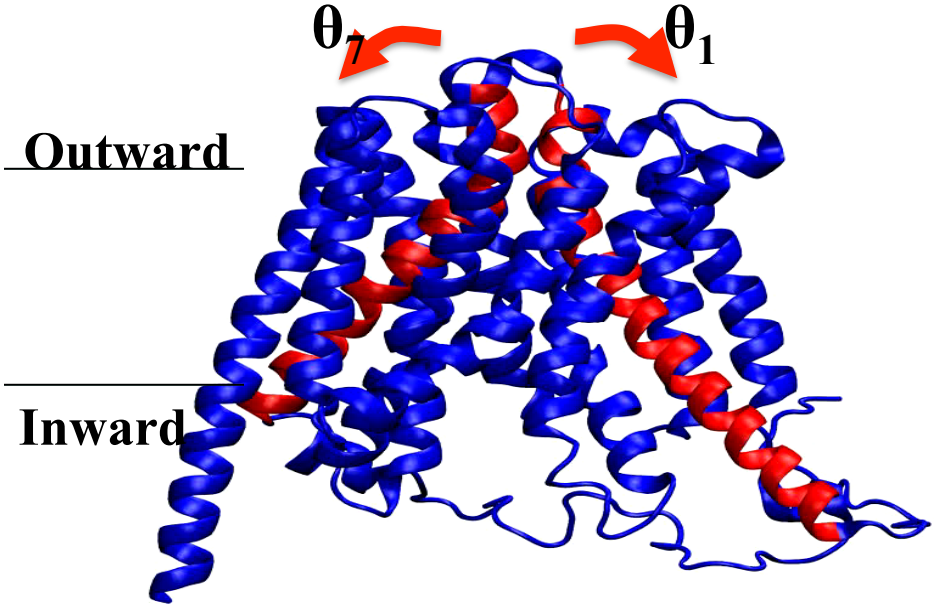

Inward-Facing↔Outward-Facing transition of GlpT transporter in explicit membrane/water environment (not shown)

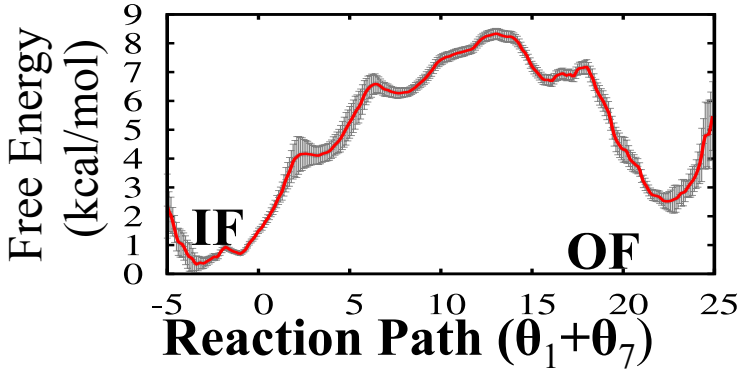

Efficient Reaction Path Sampling

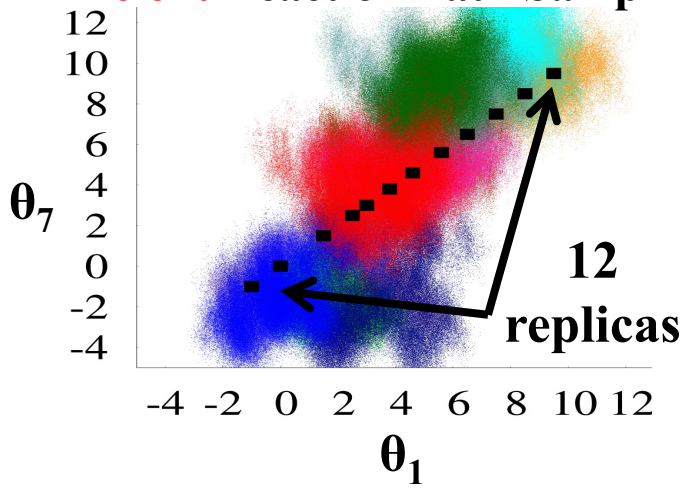

Collaborator Replica Exchange Applications

Wei Jiang, Yun Luo, Benoit Roux (Argonne Lab and U. Chicago)

2D umbrella sampling of EF-hand domain RMSD on 65,536 cores of BG/P

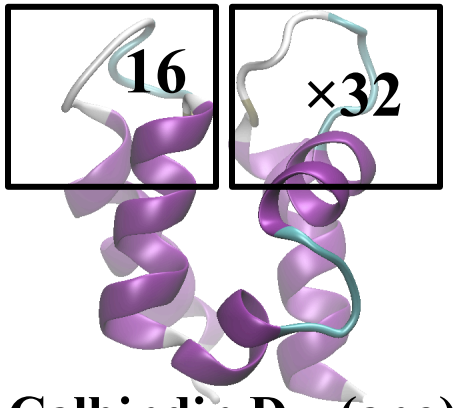

Calbindin D_{9k} (apo)
12,251 atoms/replica
16×32 = 512 replicas

Hamiltonian exchange method for alchemical free-energy perturbation "running well" on BG/Q

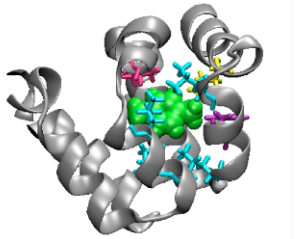

T4 Lysozyme/L99A

Future work (DOE INCITE award): Absolute binding free energy of antibiotics to

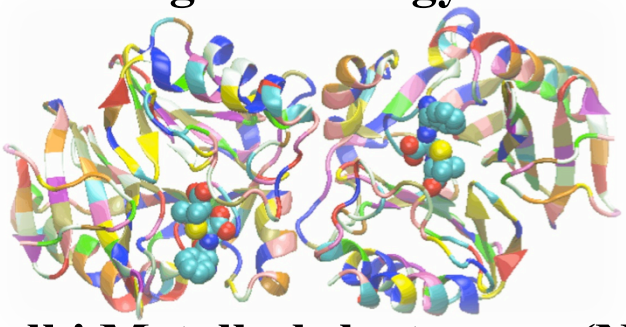

New Delhi Metallo-b-lactamase (NDM-1)

NAMD 2.10 Scalable Replica Exchange

- More general Charm++ integration:
  - NAMD 2.9 used MPI communicator splitting
  - NAMD 2.10 splits replicas in Converse low-level runtime (LRTS)
  - LRTS underlies MPI, Cray (uGNI), and BlueGene/Q (PAMI) implementations
- Basis for many enhanced sampling methods:
  - Parallel tempering (temperature exchange)
  - Umbrella sampling for free-energy calculations
  - Hamiltonian exchange (alchemical or conformational)
  - Finite Temperature String method
  - Nudged elastic band

Same Tcl scripts as NAMD 2.9

Future work enabled by Charm++ integration

- Better scaling for individual replicas:
  - **Cray uGNI layer essential for multi-node GPU replicas**
  - IBM BlueGene/Q will benefit similarly from PAMI layer
  - Porting native InfiniBand (ibverbs) layer to LRTS

Driving Project: The Ribosome

Target of over 50%
of antibiotics

Many related diseases. e.g. Alzheimer's disease due to dysfunctional ribosome (J. Neuroscience 2005, 25:9171-9175)

Localization failure of nascent chain lead to neurodegenerative disease (Mol. Bio. of the Cell 2005, 16:279-291)

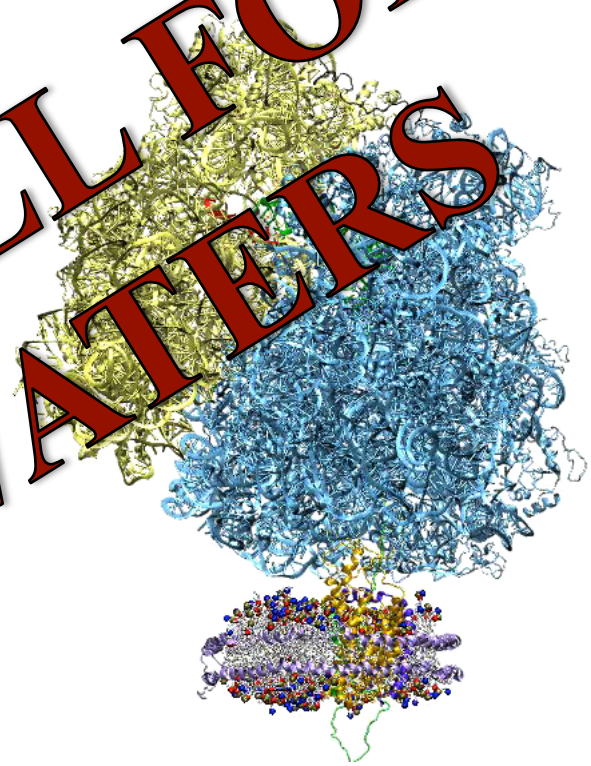

TOO SMALL FOR BLUE WATERS

Aug 2013

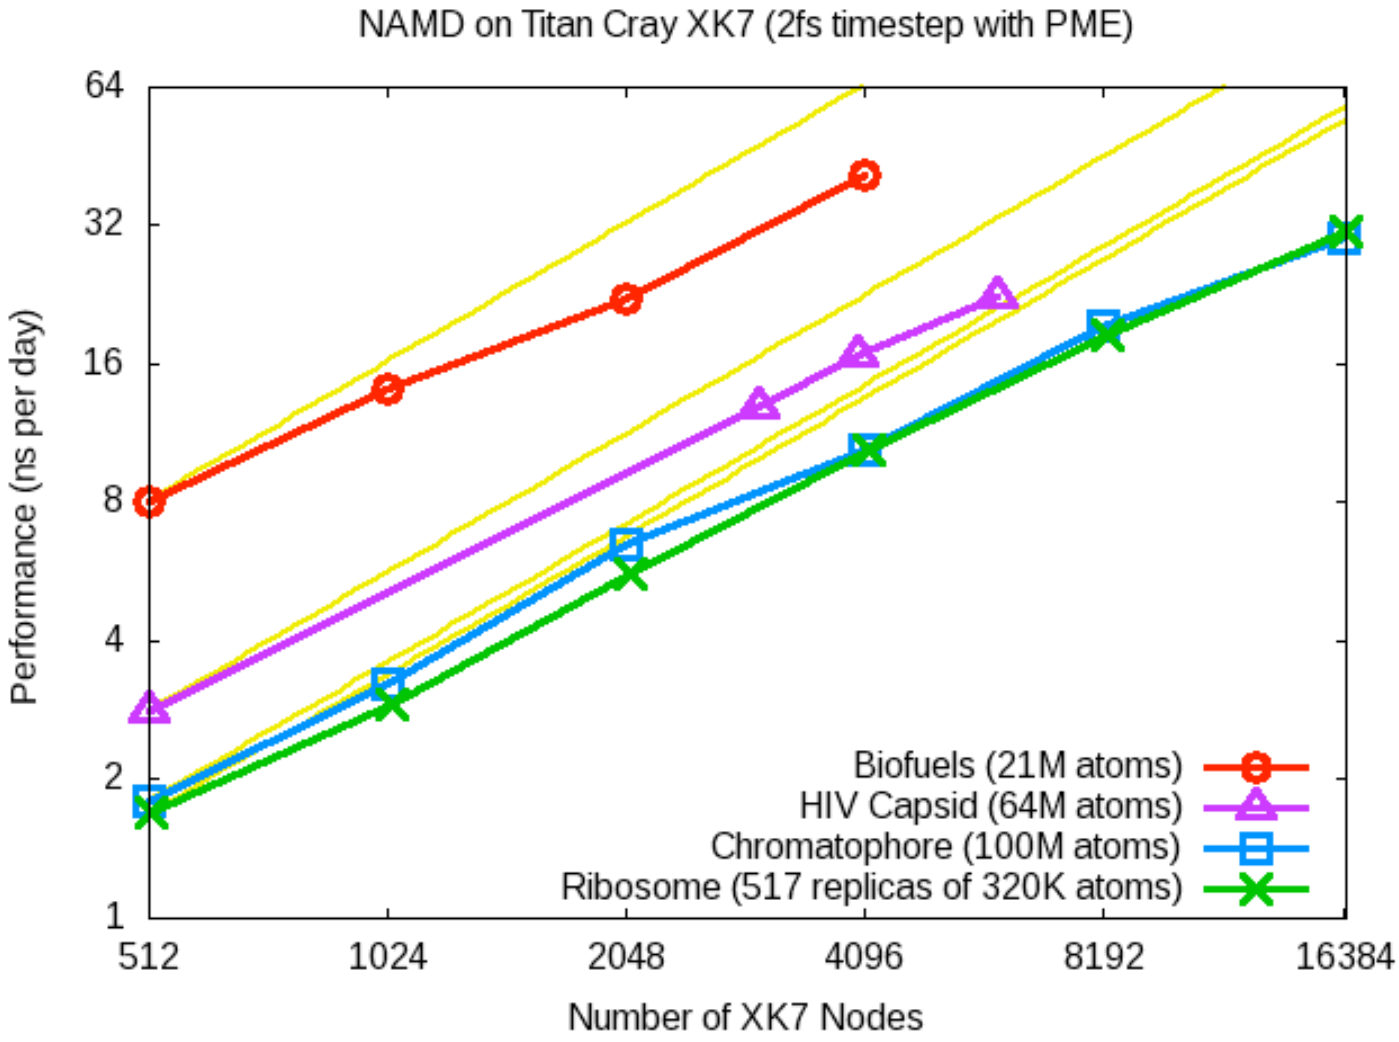

ATPESC 2013

Thanks to NIH, NSF, DOE, and 18 years of NAMD and Charm++ developers and users.

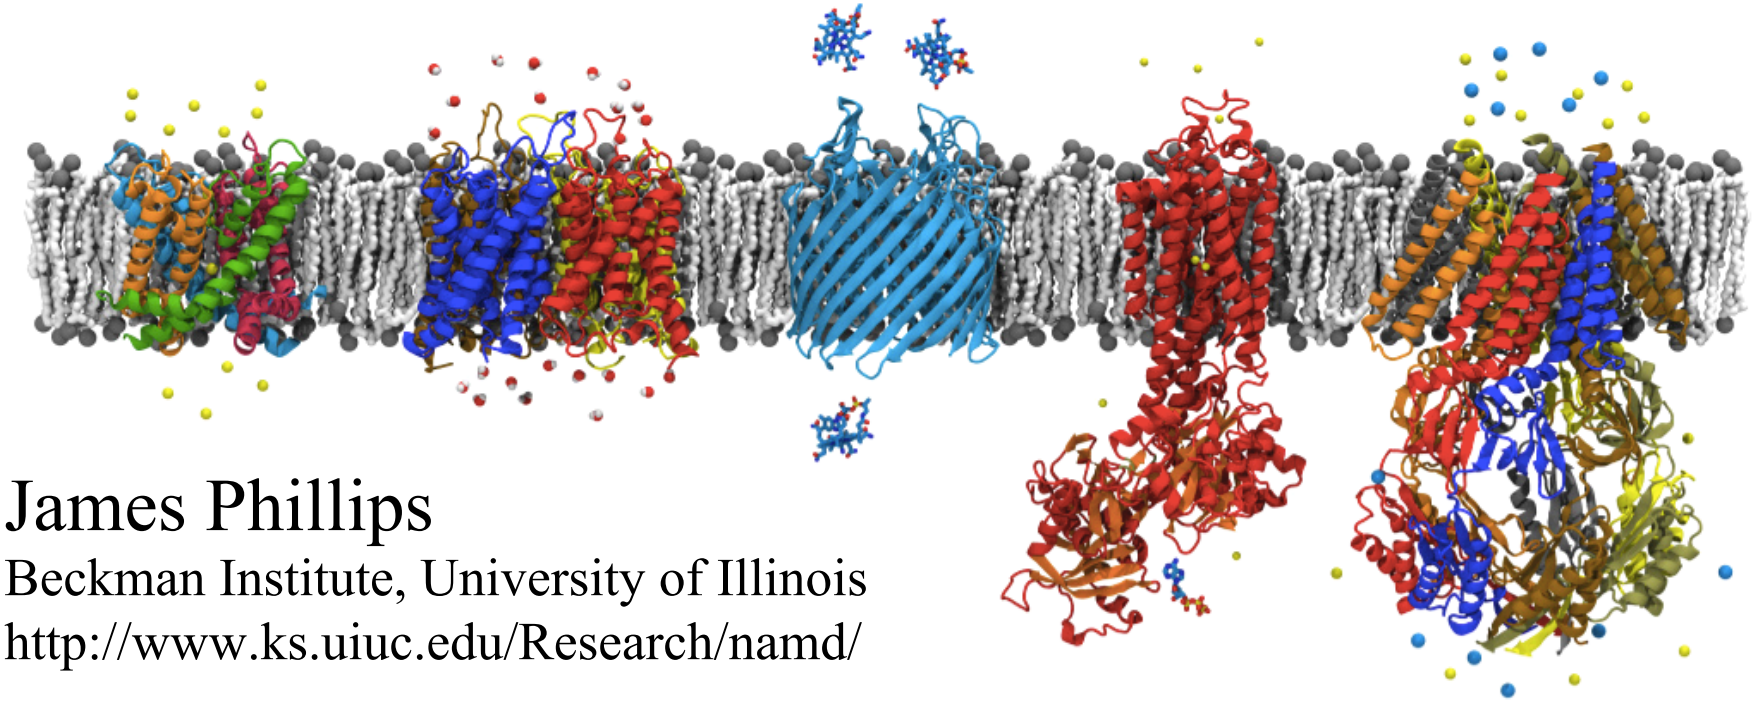

James Phillips
Beckman Institute, University of Illinois
http://www.ks.uiuc.edu/Research/namd/