Scalable Molecular Dynamics with NAMD

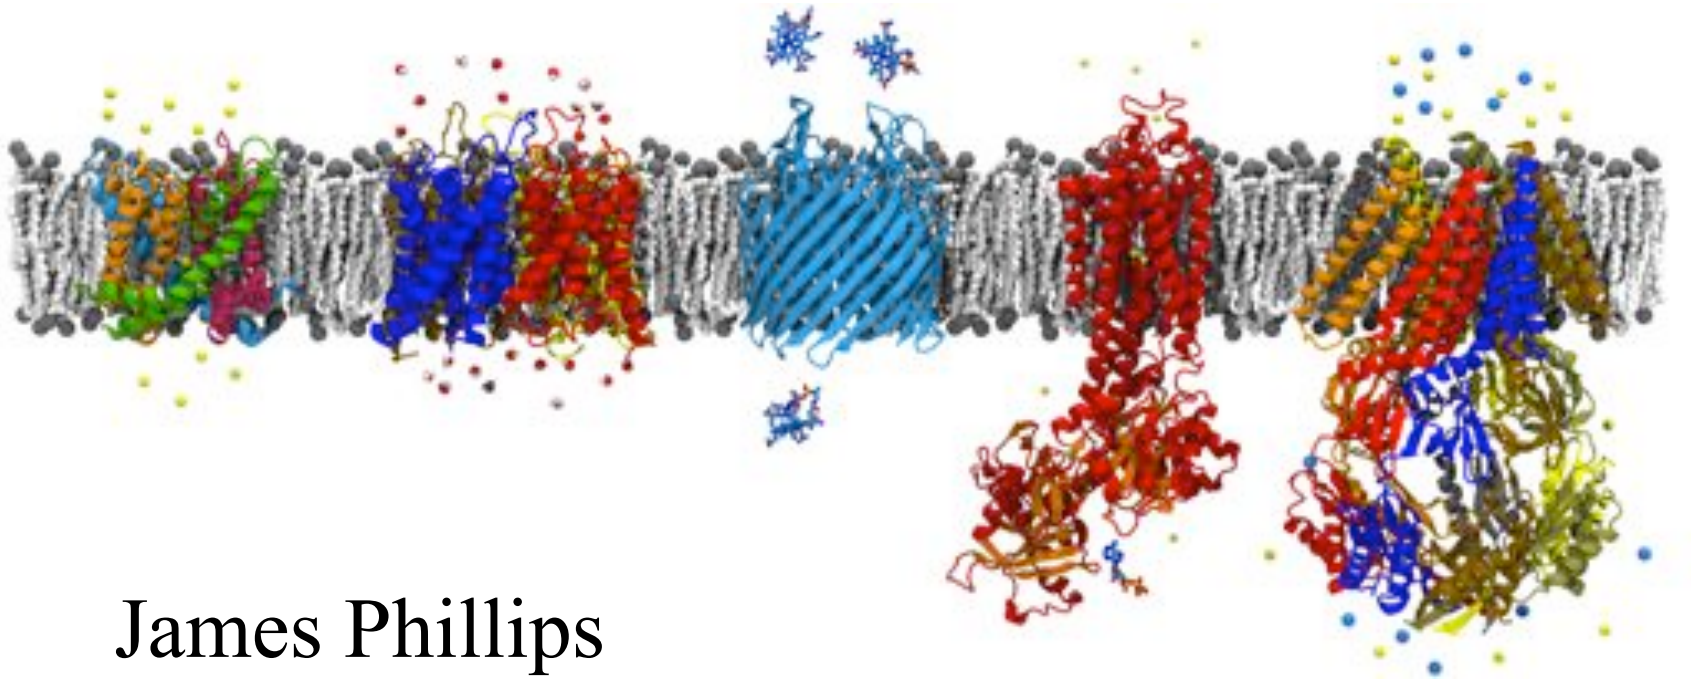

James Phillips
Beckman Institute, University of Illinois
http://www.ks.uiuc.edu/Research/namd/

NIH BTRC 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.
Center software provides multi-scale computational modeling.

femtoseconds → hours

Ångstrom → microns

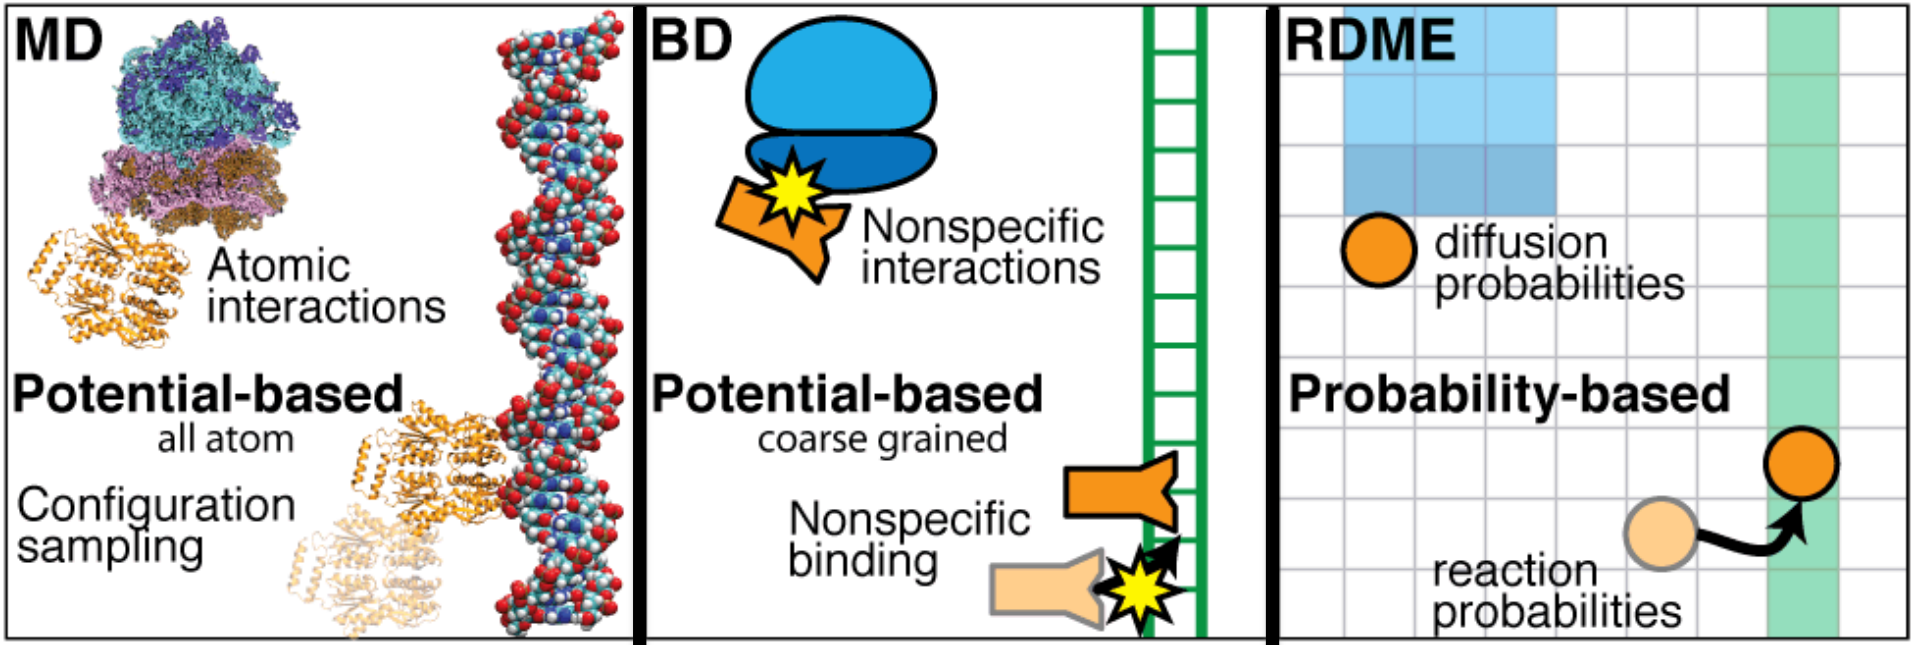

NAMD Scalable Molecular Dynamics
MDFF Molecular Dynamics Flexible Fitting
HMMM Highly Mobile Membrane Mimetic
VMD Visual Molecular Dynamics

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

Collaborative Driving Projects

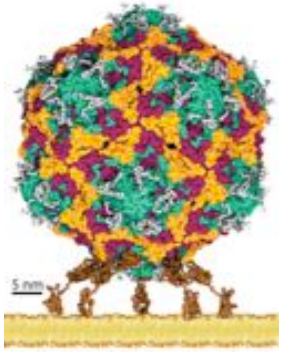

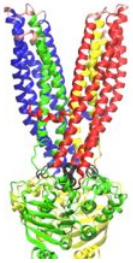

1. Ribosome	R. Beckmann (U. Munich) J. Frank (Columbia U.) T. Ha(UIUC) K. Fredrick (Ohio state U.) R. Gonzalez (Columbia U.)
2. Blood Coagulation Factors	J. Morrissey (UIUC) S. Sligar (UIUC) C. Rienstra (UIUC) G. Gilbert (Harvard)
3. Whole Cell Behavior	W. Baumeister (MPI Biochem.) J. Xiao (Johns Hopkins 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.)

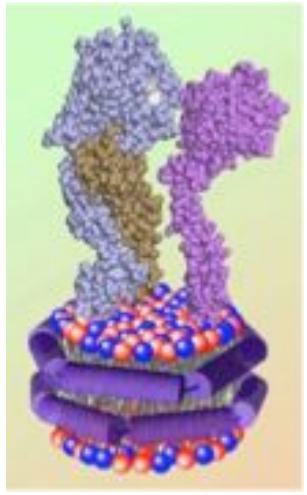

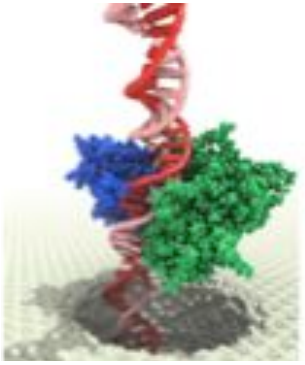

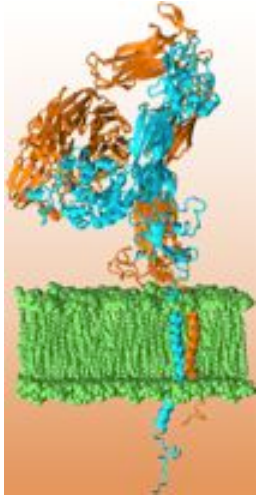

Ribosome Driving Project

Target of over 50%
of antibiotics

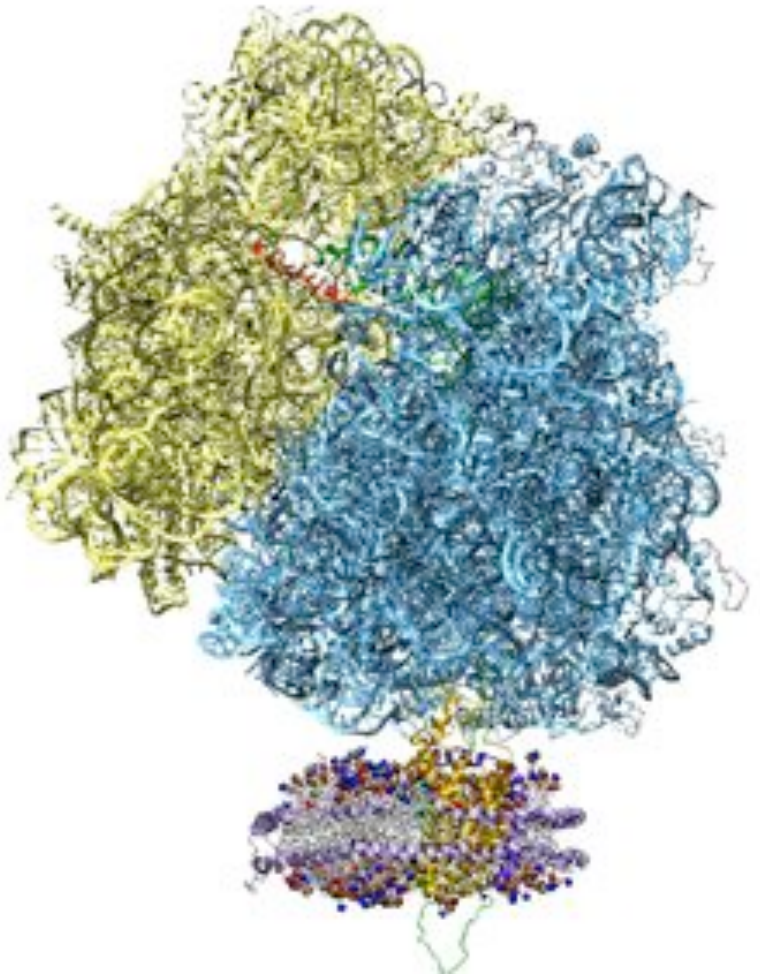

Many related diseases. e.g. Alzheimer's disease due to dysfunctional ribosome (J. Neuroscience 2005, 25:9171-9175)

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

Viral Infection Driving Projects

Poliovirus

Poliovirus is a model system for understanding how non-enveloped viruses bind to and enter a host cell.

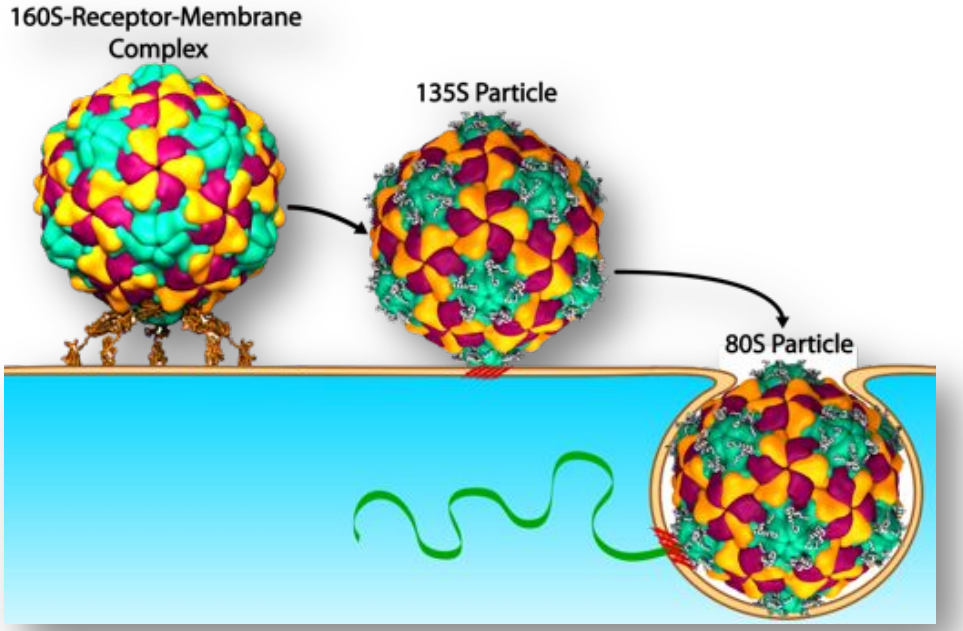

Human Immunodeficiency Virus 1

Knowledge of HIV capsid atomic structure may reveal disassembly mechanism and guide novel therapies.

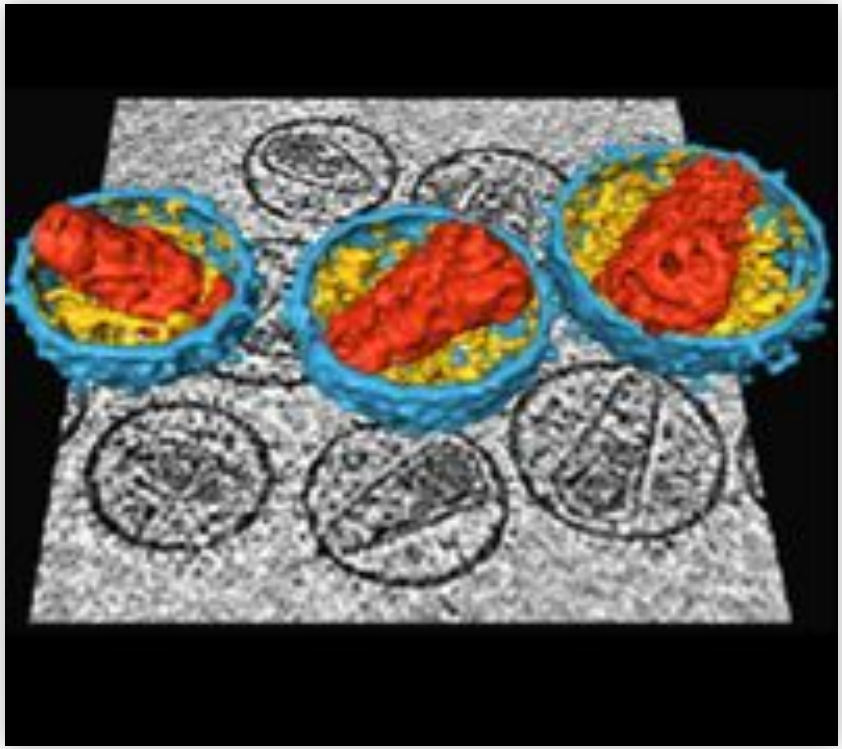

Briggs et al. Structure (2006) 14:15-20.

Blue Waters Early Science Project

"The first all-atom structure of an **HIV virus capsid** in its tubular form, courtesy Klaus Schulten, University of Illinois at Urbana-Champaign Theoretical and Computational Biophysics Group/ Beckman Institute; Angela Gronenborn and Peijun Zhang, University of Pittsburgh School of Medicine Center for HIV Protein Interactions/Department of Structural Biology."

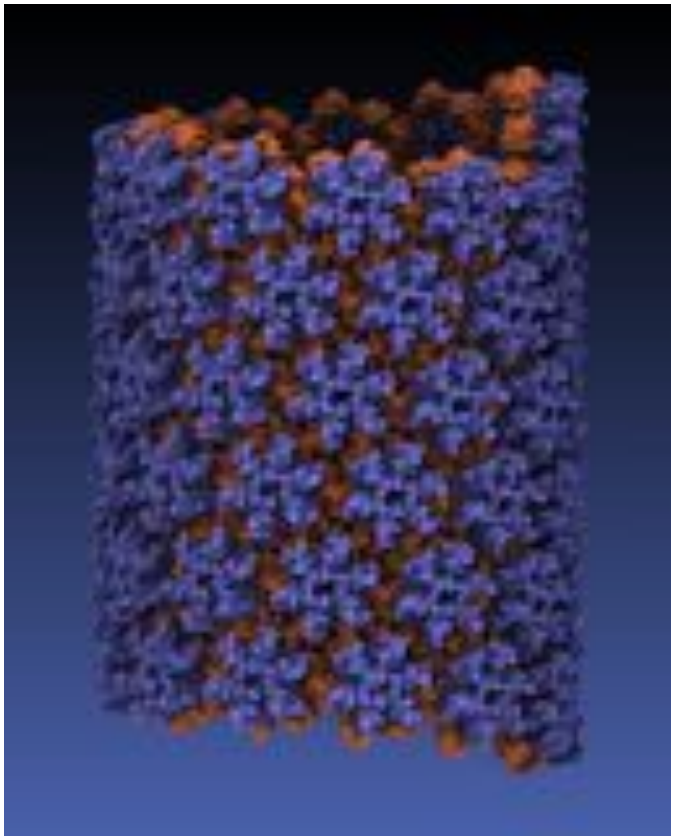

NAMD: Scalable Molecular Dynamics

2002 Gordon Bell Award

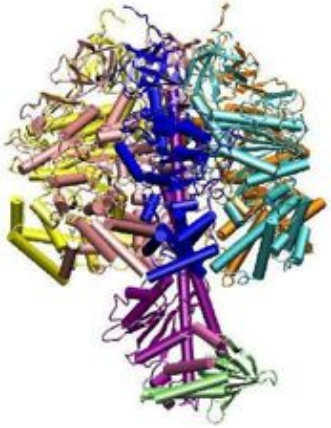

ATP synthase

PSC Lemieux

51,000 Users, 2900 Citations

Computational Biophysics Summer School

Blue Waters Target Application

Illinois Petascale Computing Facility

GPU Acceleration

NVIDIA Tesla

NCSA Lincoln

Larger machines enable larger simulations

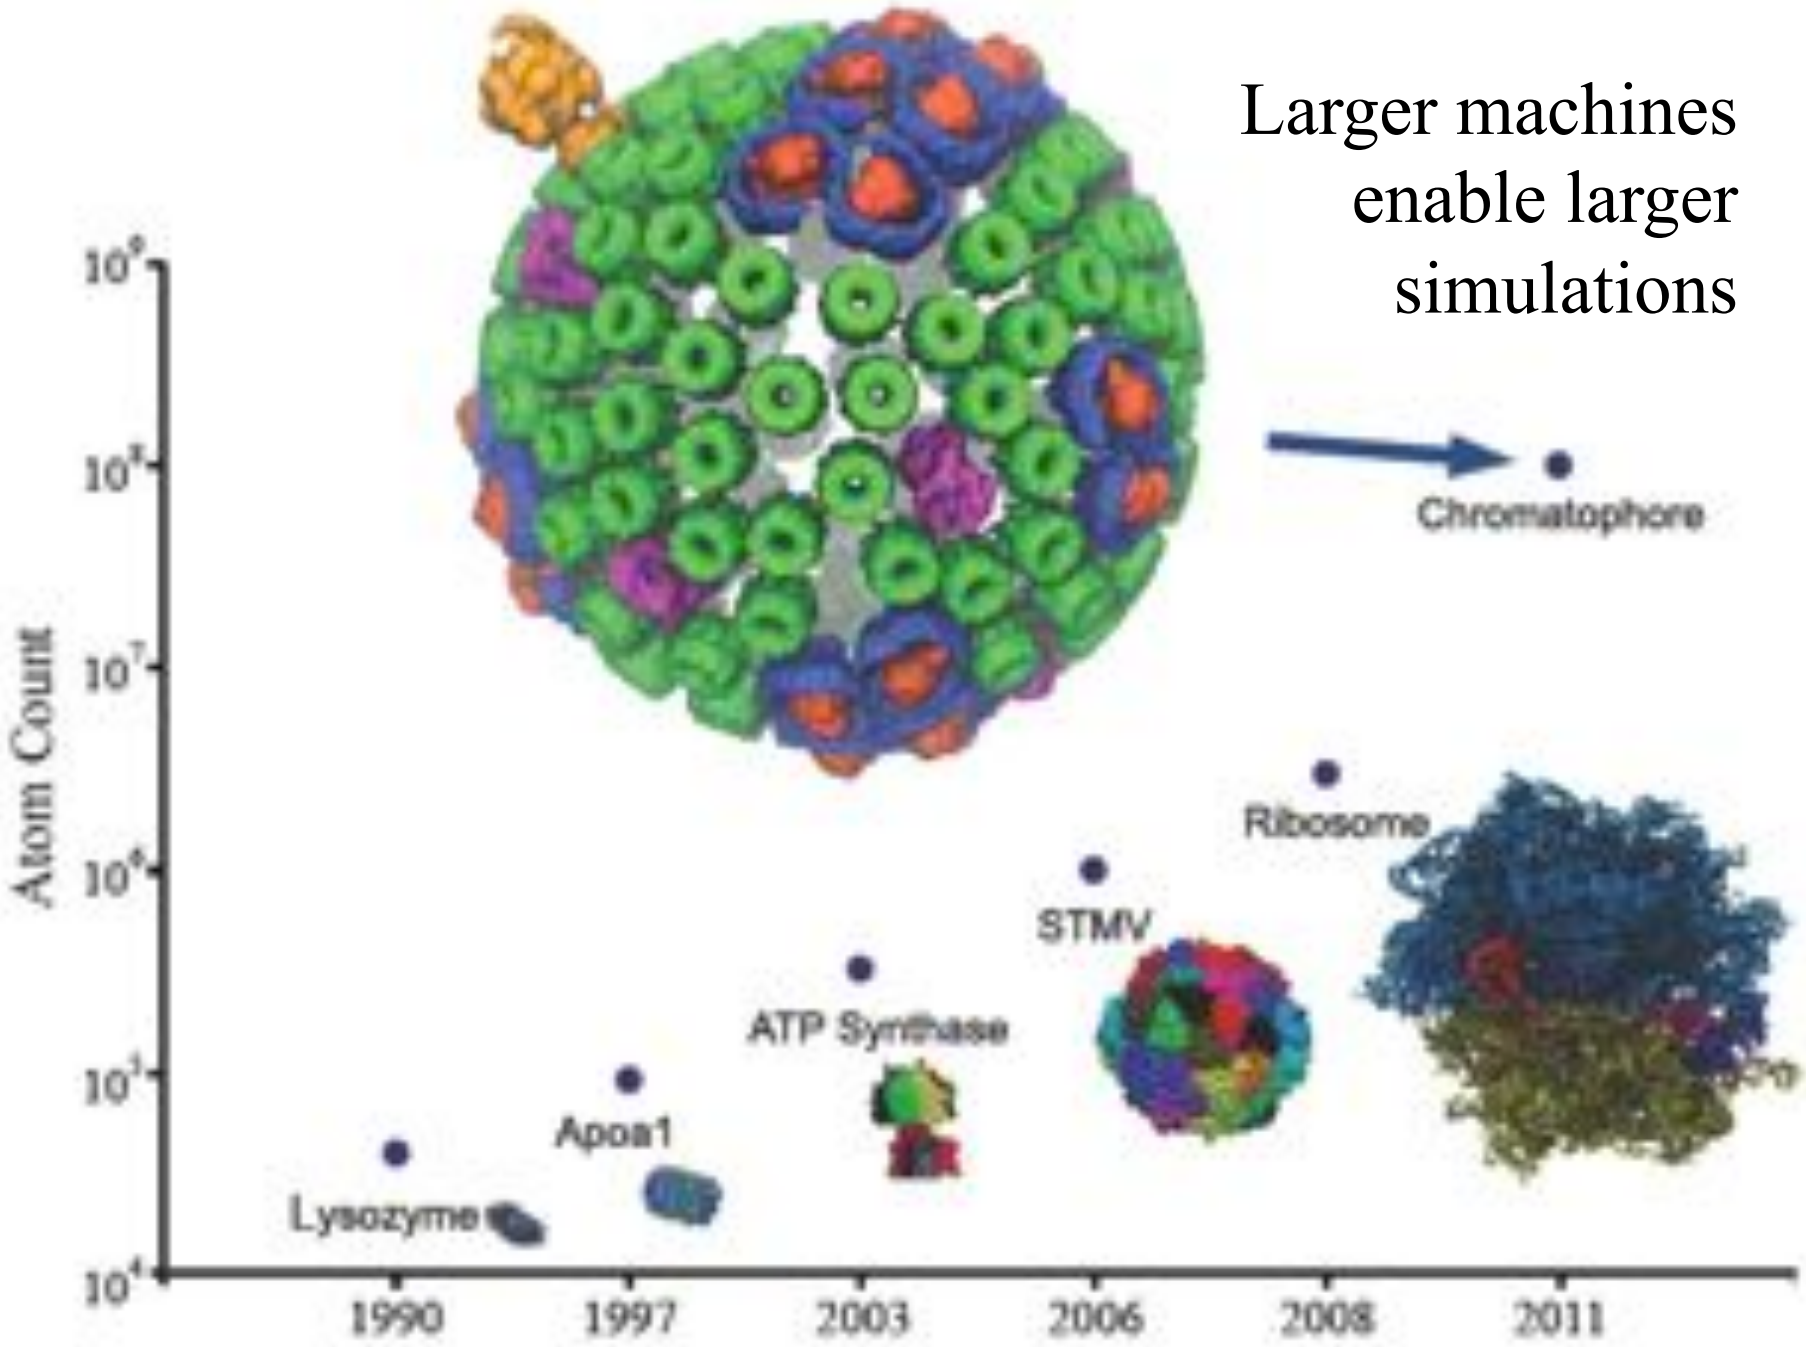

NAMD impact is broad and deep

- 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 – 51,000 users
  - 9,100 (18%) are NIH-funded; many in other countries
  - 14,100 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

2100 external citations since 2007

Corringer, et al., *Nature*, 2011

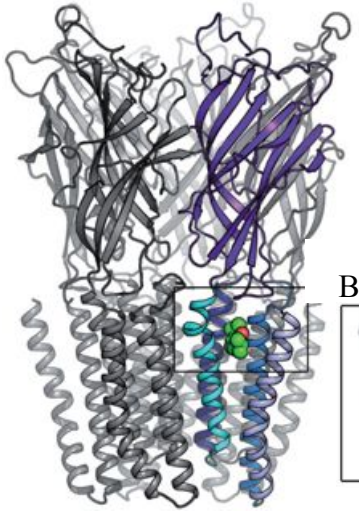

180K-atom 30 ns study of anesthetic binding to bacterial ligand-gated ion channel provided "complementary interpretations…that could not have been deduced from the static structure alone."

Bound Propofol Anesthetic

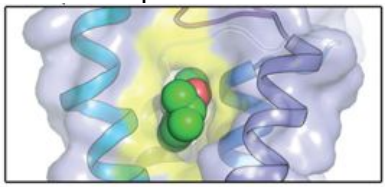

Voth, et al., *PNAS*, 2010

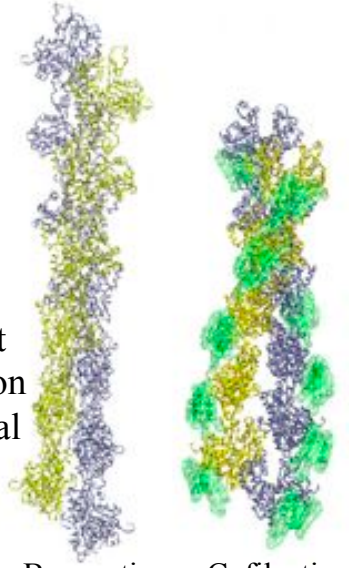

500K-atom 500 ns investigation of effect of actin depolymerization factor/cofilin on mechanical properties and conformational dynamics of actin filament.

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)
- **D. Krepkiy, et al.,**, *Structure and hydration of membranes embedded with voltage-sensing domains*. (2009)
- **N. Yeung, et al.,**, *Rational design of a structural and functional nitric oxide reductase*. (2009)
- **Z. Xia, et al.**, *Recognition Mechanism of siRNA by Viral p19 Suppressor of RNA Silencing: A Molecular Dynamics Study*. (2009)

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)

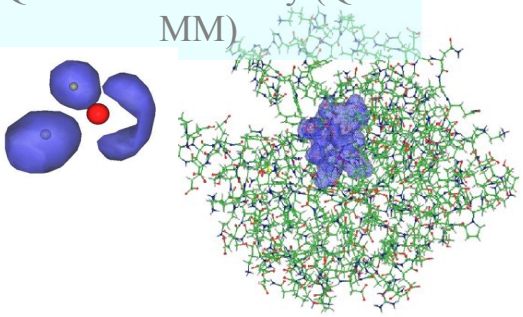

NAMD: Molecular Dynamics

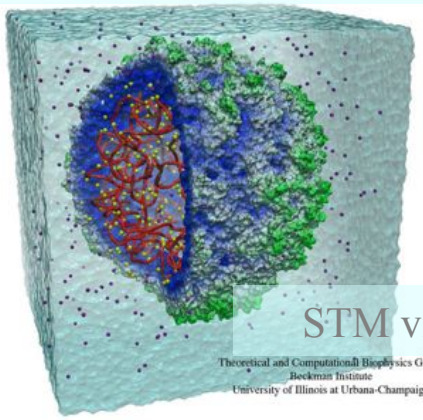

Protein Folding

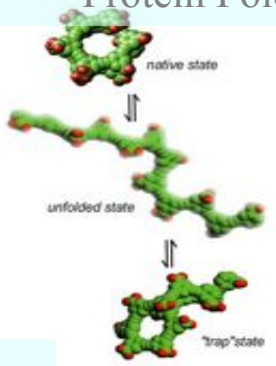

STM virus simulation

Computational Cosmology

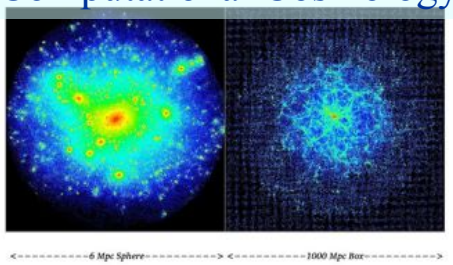

Parallel Objects,

Adaptive Runtime System

Libraries and Tools

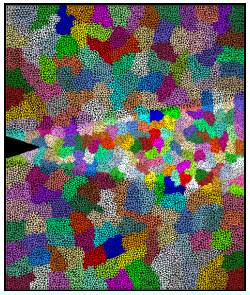

Crack Propagation

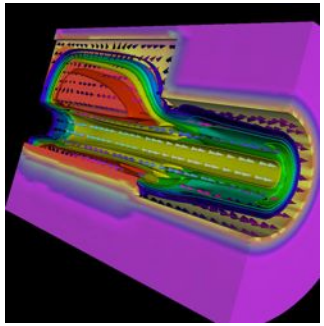

Rocket Simulation

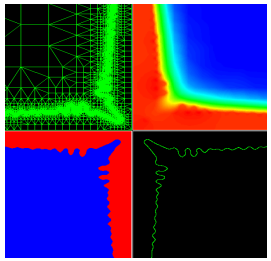

Dendritic Growth

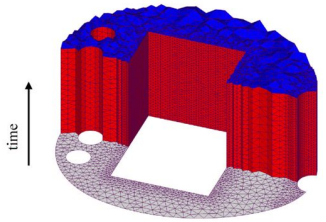

Space-time meshes

The enabling CS technology of parallel objects and intelligent Runtime systems has led to several collaborative applications in CSE

Computing research drives NAMD

- Parallel Programming Lab – (Laxmikant Kale)
  - Charm++ is an Adaptive Parallel Runtime System
    - Gordon Bell Prize 2002
    - Three publications at Supercomputing 2011
    - Four panels discussing the future necessity of our ideas
- 20 years of co-design for NAMD performance, portability, and productivity, adaptivity
    - Recent example: Implicit Solvent deployed in NAMD by 1 RA in 6 months. 4x more scalable than similar codes
- Yesterday's supercomputer is tomorrow's desktop

NAMD 2.8 Highly Scalable Implicit Solvent Model

NAMD Implicit Solvent is 4x more scalable than Traditional Implicit Solvent for all system sizes, implemented by one GRA in 6 months.

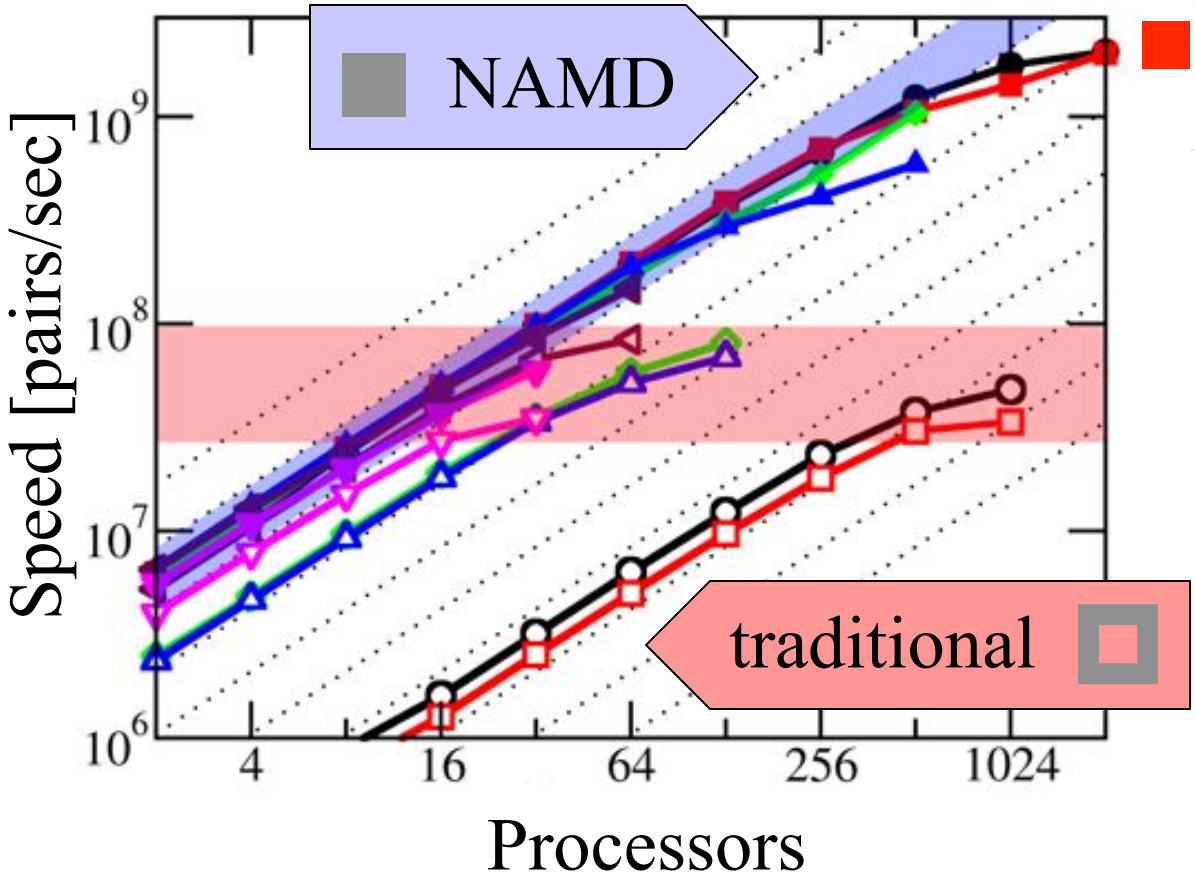

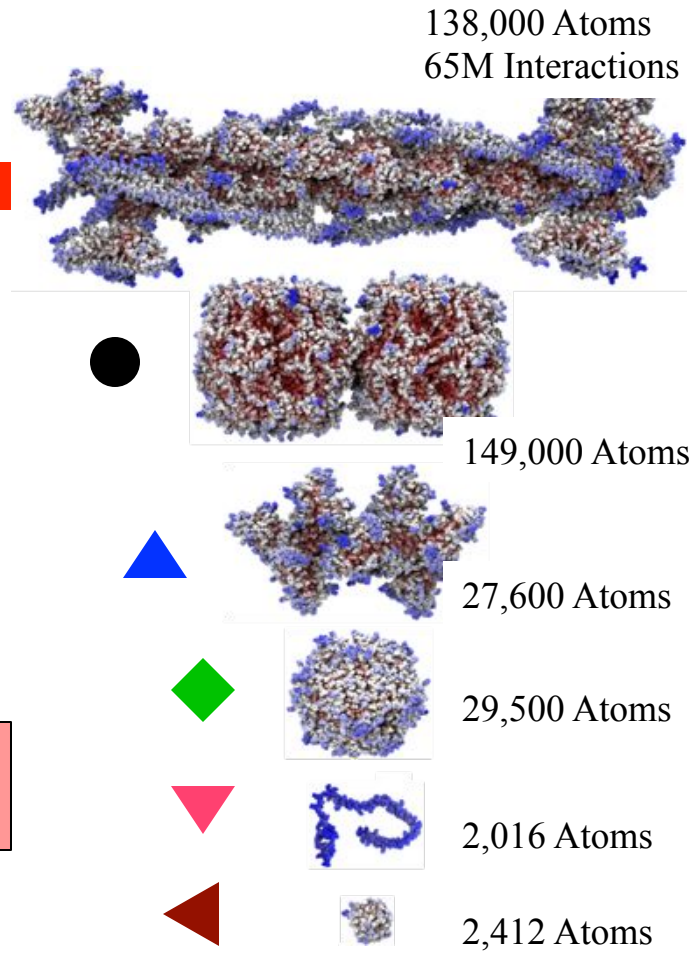

Tanner et al., J. Chem. Theory and Comp., 7:3635-3642, 2011

Charm++ Used by NAMD

- Parallel C++ with *data driven* objects.
- Asynchronous method invocation.
- Prioritized scheduling of messages/execution.
- Measurement-based load balancing.
- Portable messaging layer.

NAMD Overlapping Execution

Phillips *et al., SC2002*.

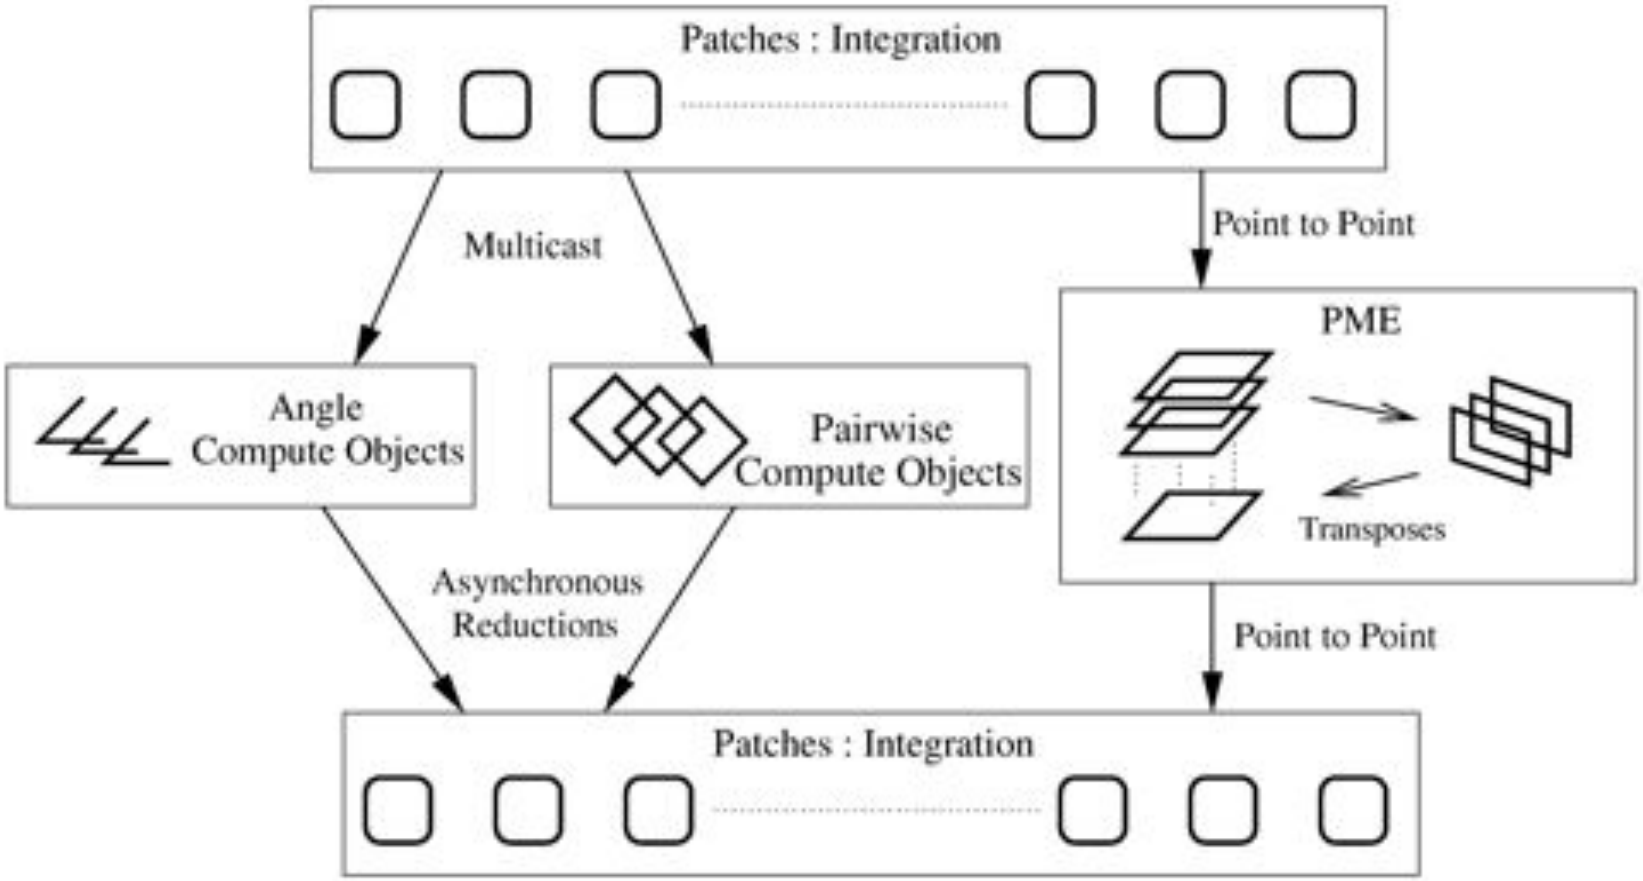

Objects are assigned to processors and queued as data arrives.

Tsubame (Tokyo) Application of GPU Accelerated NAMD

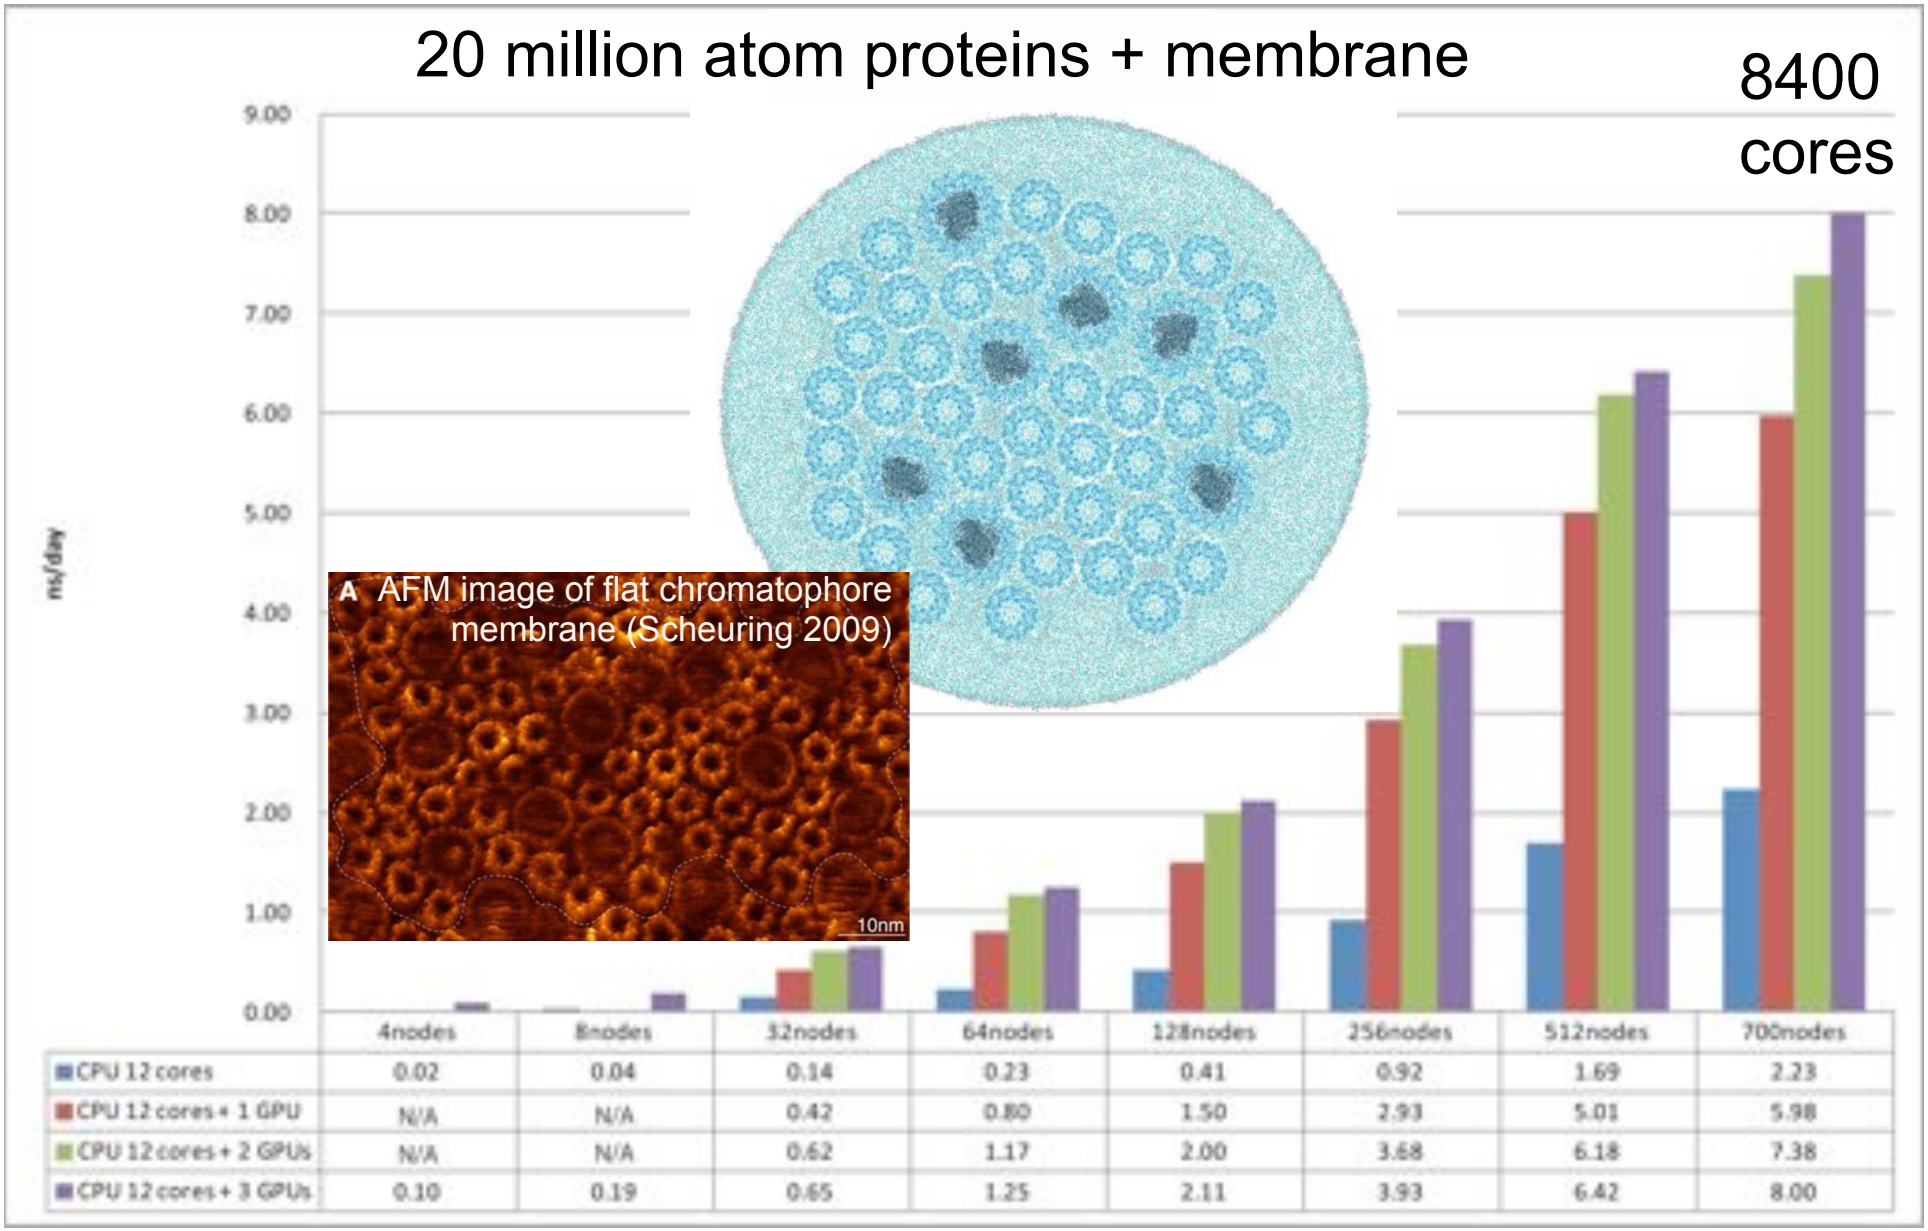

	4nodes	8nodes	32nodes	64nodes	128nodes	256nodes	512nodes	700nodes
CPU 12 cores	0.02	0.04	0.14	0.23	0.41	0.92	1.69	2.23
CPU 12 cores + 1 GPU	N/A	N/A	0.42	0.80	1.50	2.93	5.01	5.98
CPU 12 cores + 2 GPUs	N/A	N/A	0.62	1.17	2.00	3.68	6.18	7.38
CPU 12 cores + 3 GPUs	0.10	0.19	0.65	1.25	2.11	3.93	6.42	8.00

Cray Gemini Optimization

- The new Cray machine has a better network (called Gemini)
- MPI-based NAMD scaled poorly
- BTRC implemented direct port of Charm++ to Cray
  - *uGNI* is the lowest level interface for the Cray Gemini network
  - Removes MPI from NAMD call stack

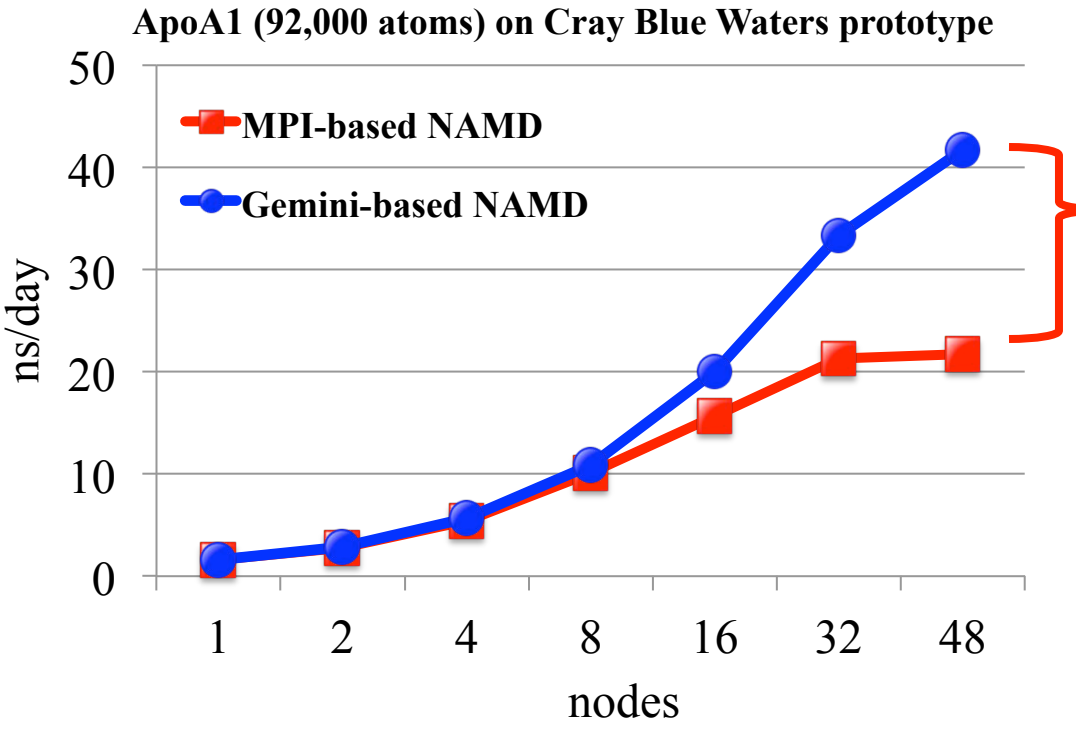

Gemini provides at least 2x increase in usable nodes for strong scaling

100M-atom Benchmark Performance

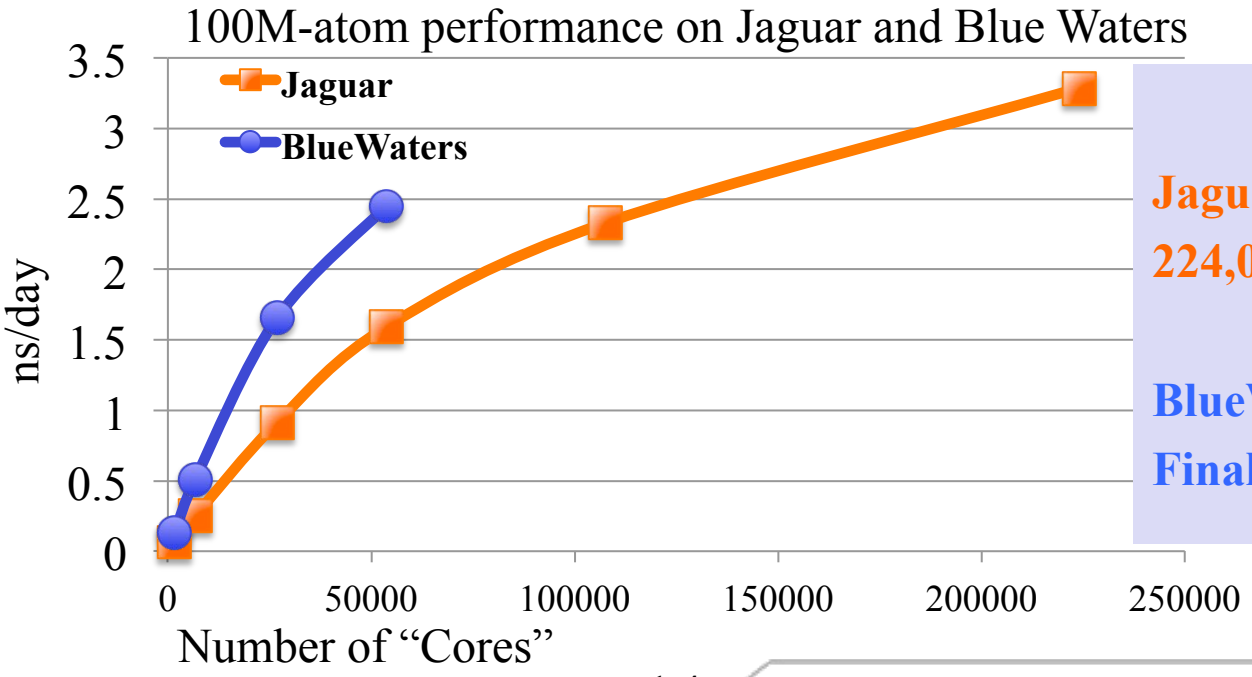

Cray

Jaguar XT5 (ORNL)

224,000 cores

BlueWaters ESS XE6 (UIUC)

Final: 380,000 cores, 3000 GPUs

Tsubame

Tokyo Institute of Tech.

4224 GPUs

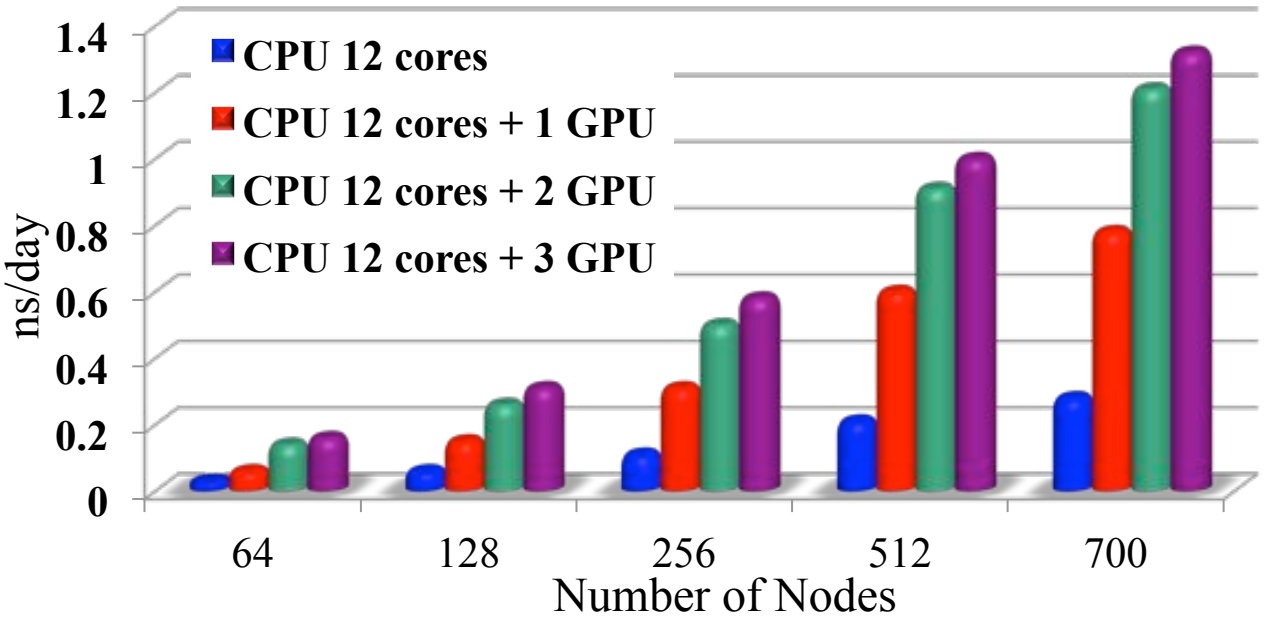

TitanDev Cray XK6 Results

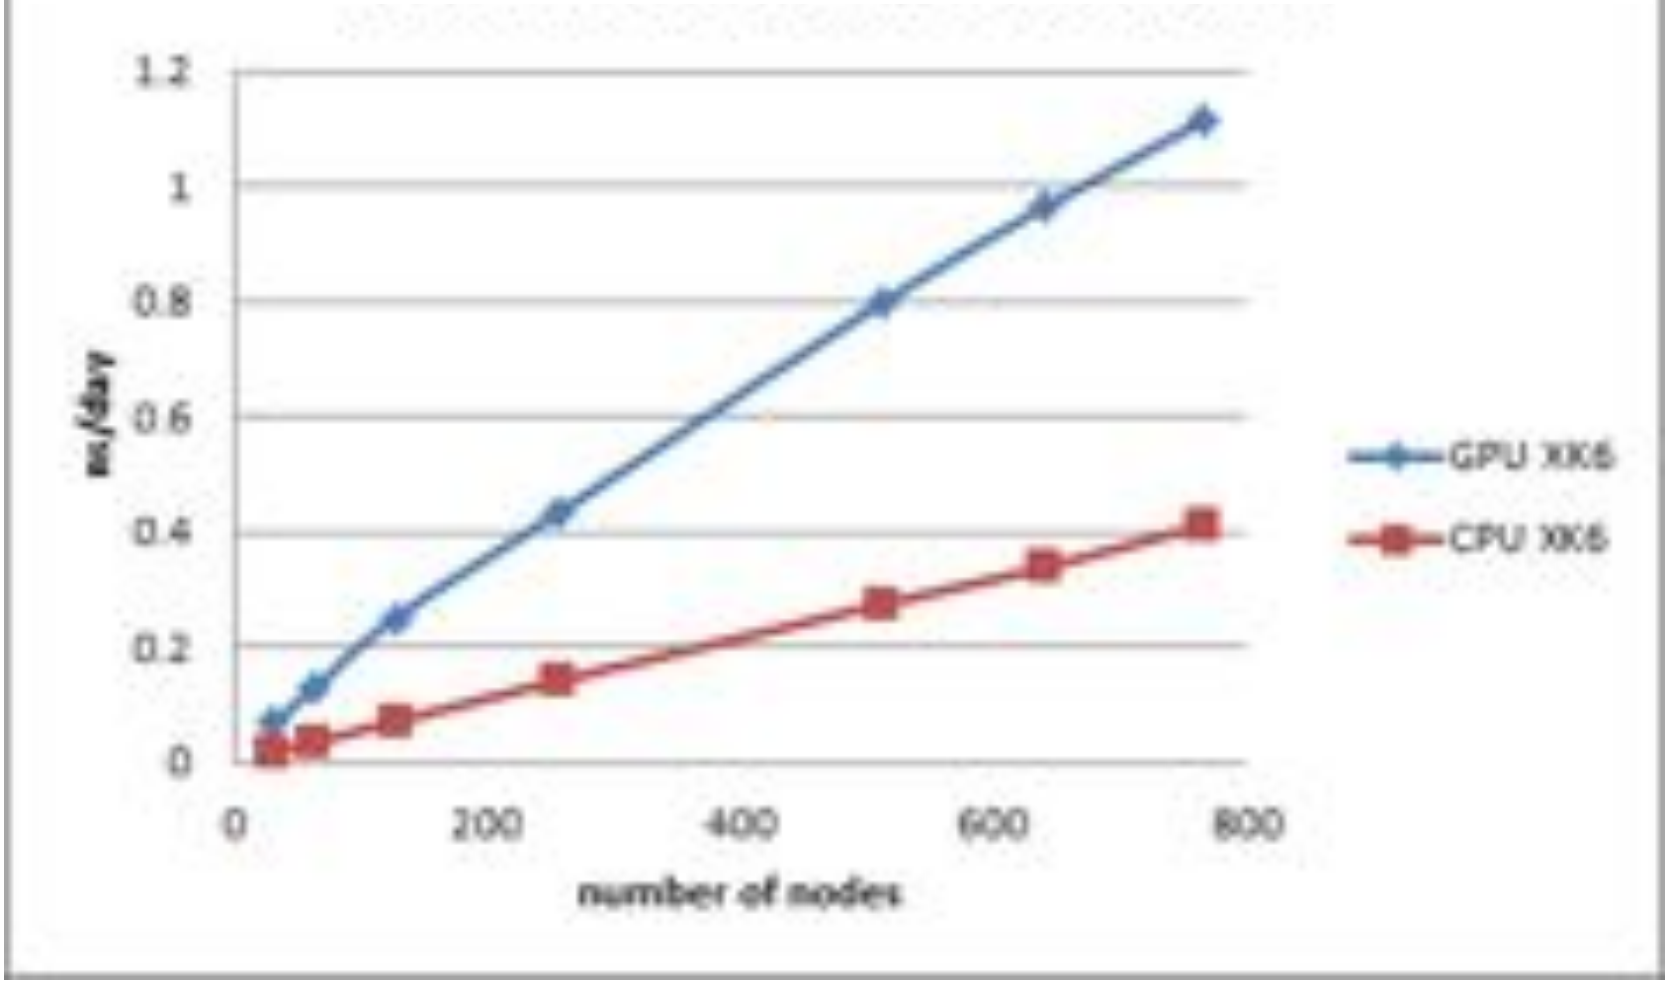

TitanDev Cray XK6 Results

	100 stmv s/step 32 nodes	100 stmv s/step 768 nodes
XK6, Fermi	1.23065 s	0.07657 s
XK6, without Fermi	4.62633 s	0.19688 s
XE6	2.318 s	0.10602 s
XK6 Fermi vs XK6 without Fermi	3.8 x	2.6 x
XK6 Fermi vs XE6	1.8 x	1.4 x

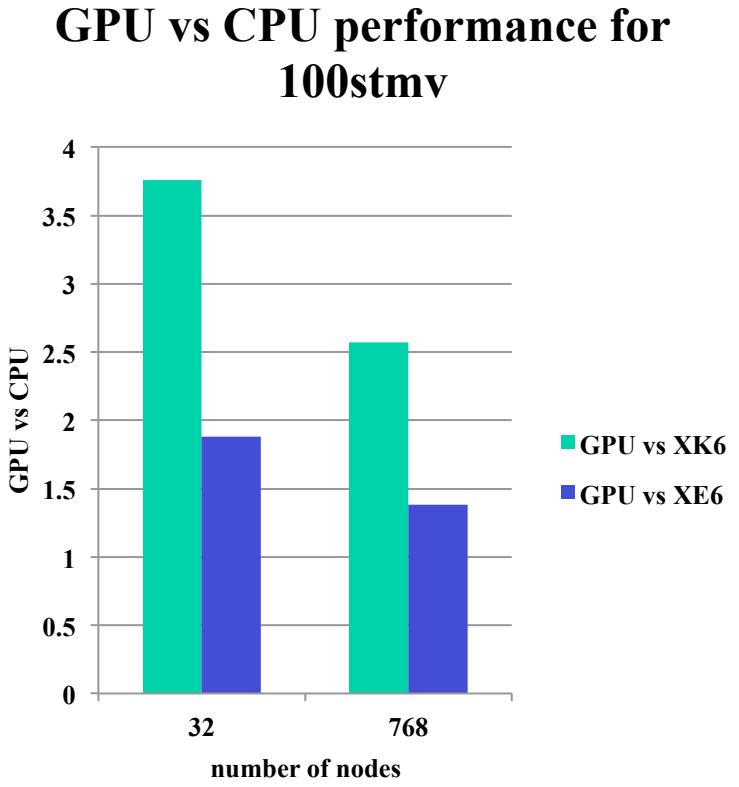

TitanDev Cray XK6 Results

	time step	non PME step	PME step	sum PME computation	kernel
GPU, 100 stmv, 32 nodes	1.249 s	1.009 s	1.833 s	0.796 s	0.964 s
GPU, 4 stmv, 30 nodes	0.057 s	0.049 s	0.065 s	0.028 s	0.042 s
GPU, 100 stmv, 768 nodes	0.085 s	0.046 s	0.176 s	0.034 s	0.042 s
strong scaling	0.616	0.905	0.435	0.988	0.966
weak scaling	0.680	1.061	0.368	0.845	1.004

TitanDev Strong Scaling

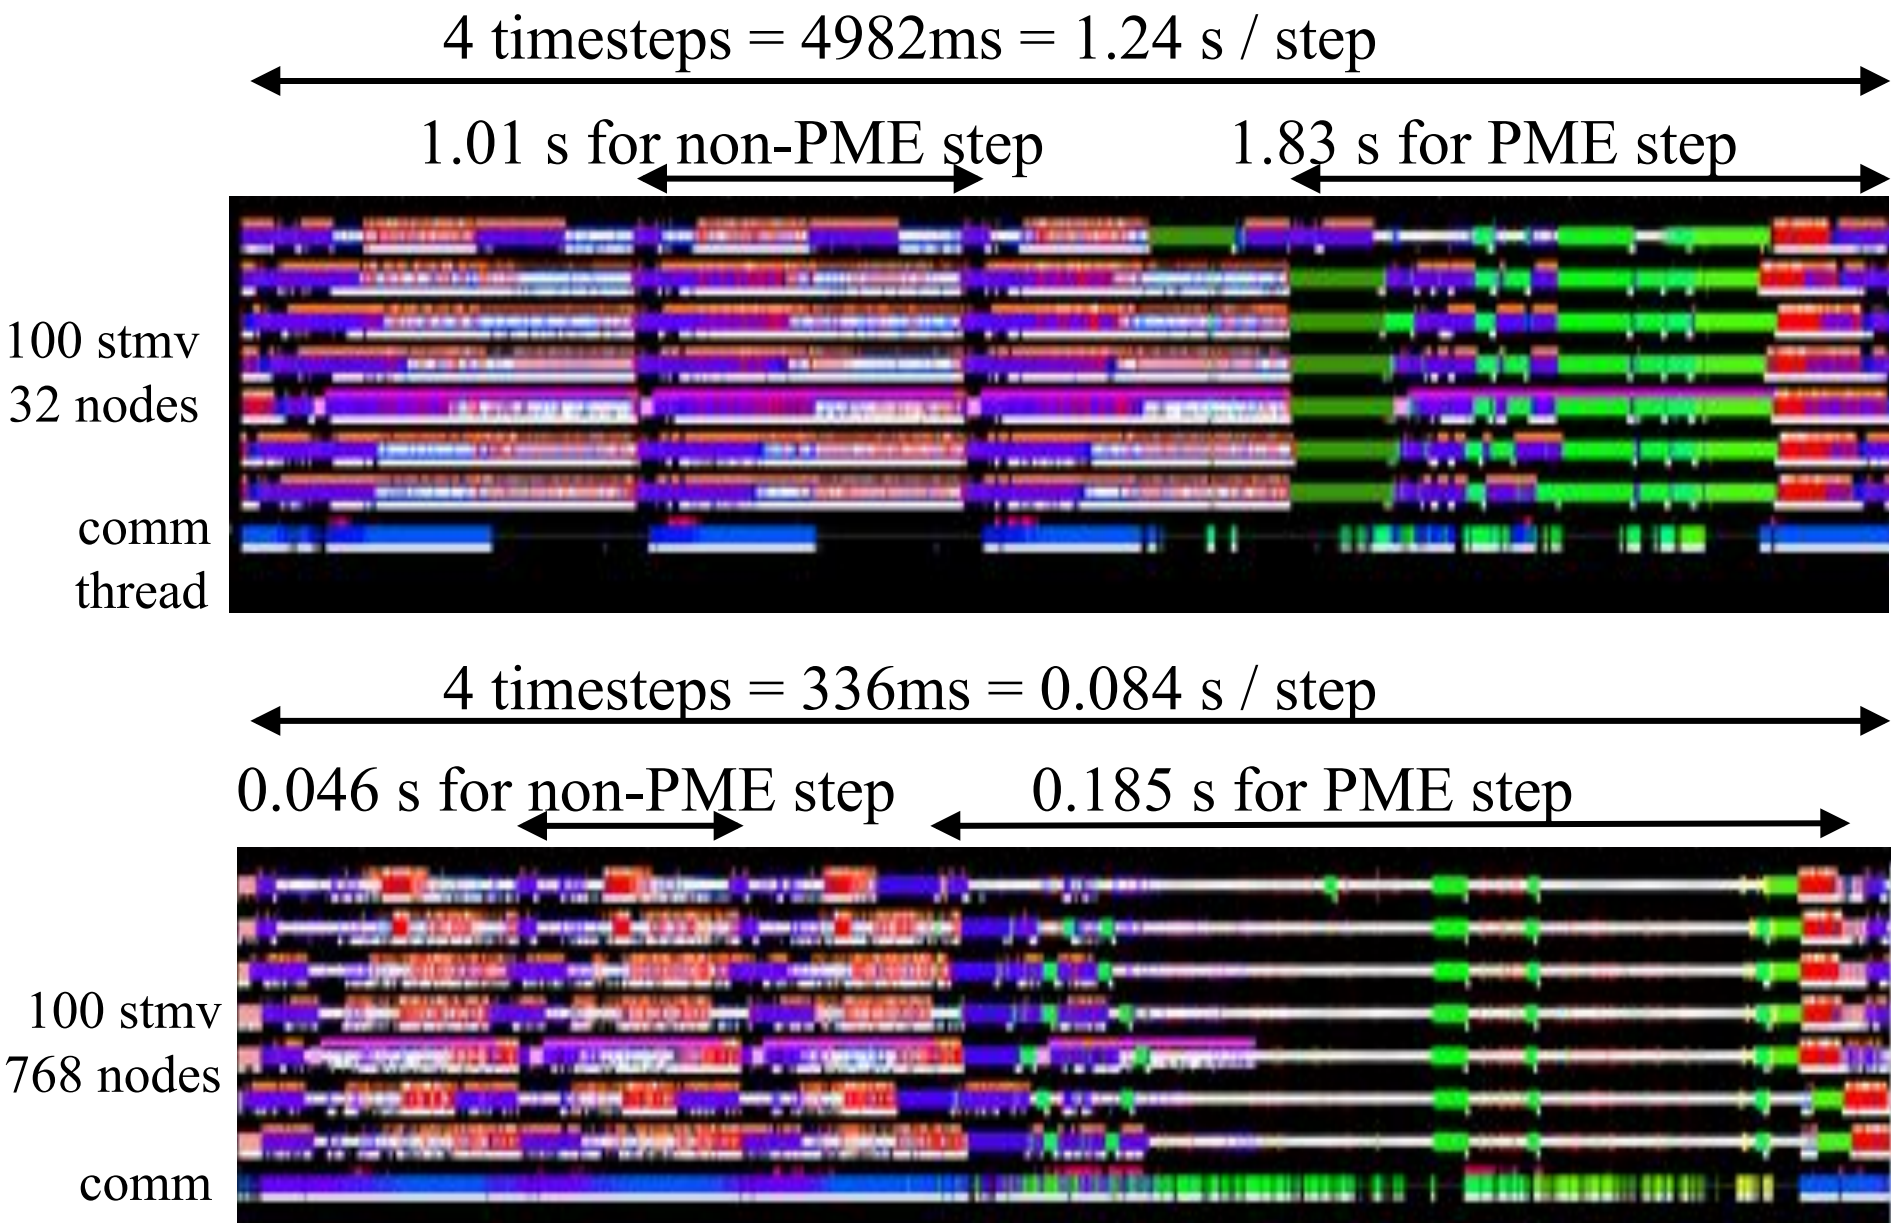

TitanDev Weak Scaling

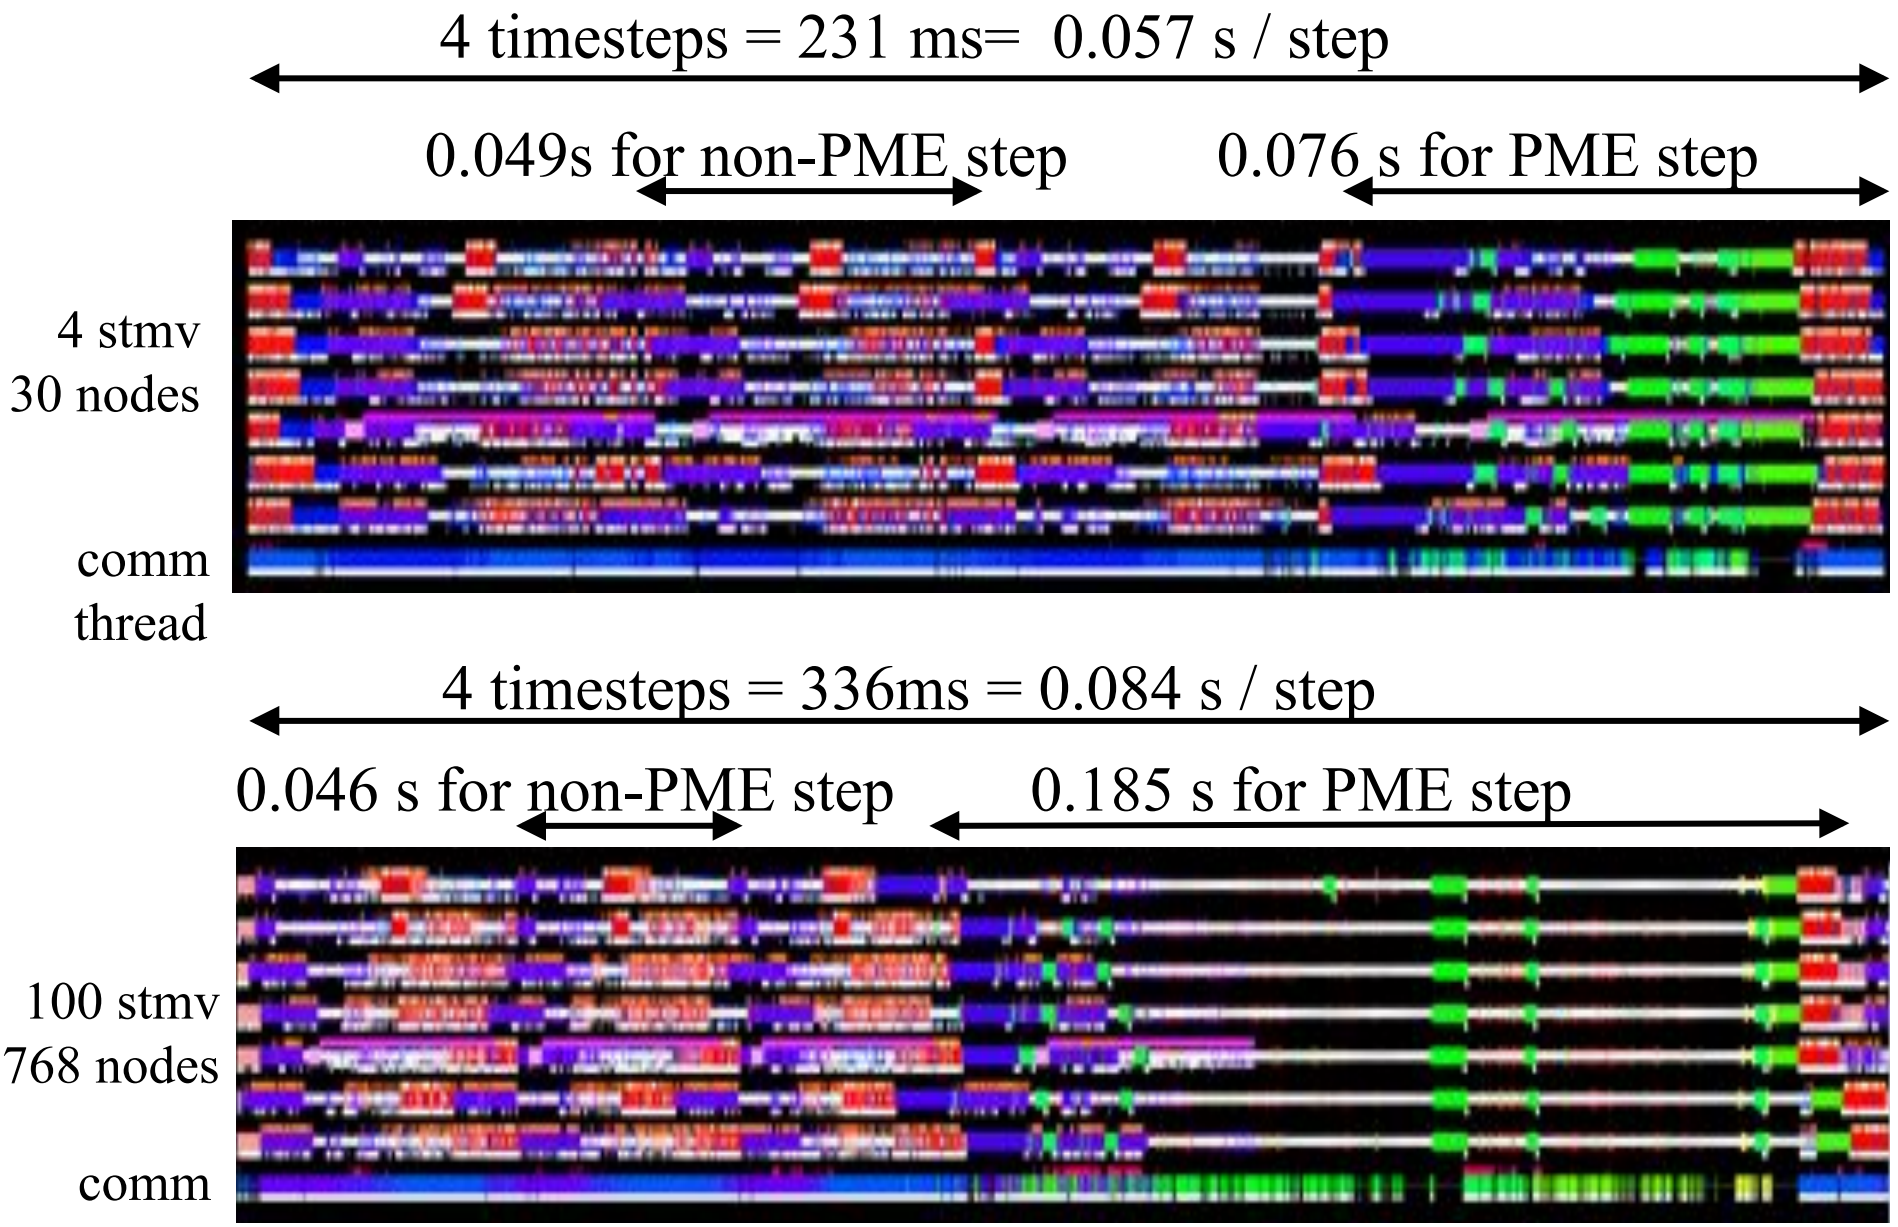

PME delays – tracing data needed for one ungrid calculation

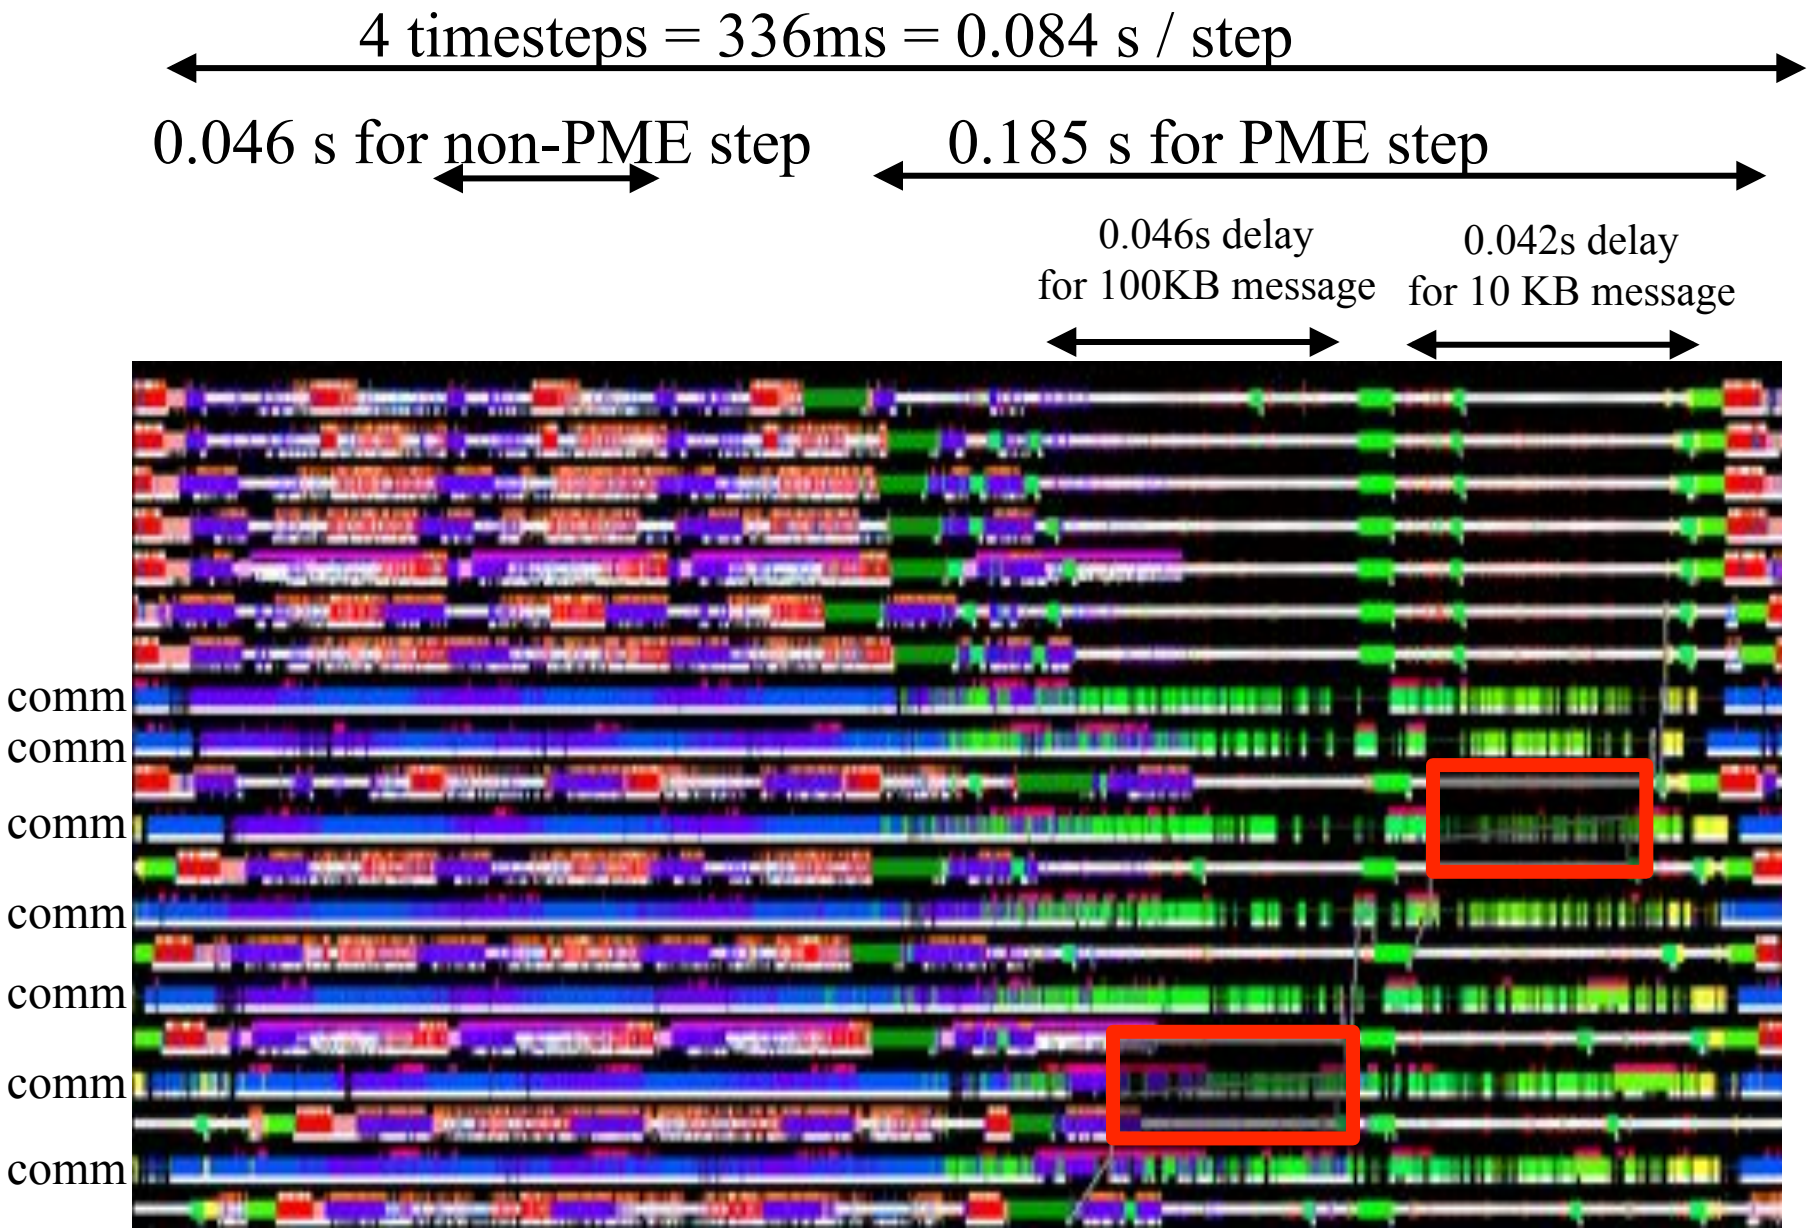

Strategy to improve scalability

- Fix issues with communication
  - 23x16x2 topology limits bisection bandwidth
- Push PME work to the GPU
  - Charge gridding overlaps coordinate receive
- Start GPU work sooner
  - Currently waiting for all coordinate receives
  - Use streams to launch work as data arrives

NAMD 2.9 Scalable Replica Exchange

- 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

  } Released in NAMD 2.9

  - Hamiltonian exchange (alchemical or conformational)
  - Finite Temperature String method
  - Nudged elastic band

  } Enabled for Roux group

- 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

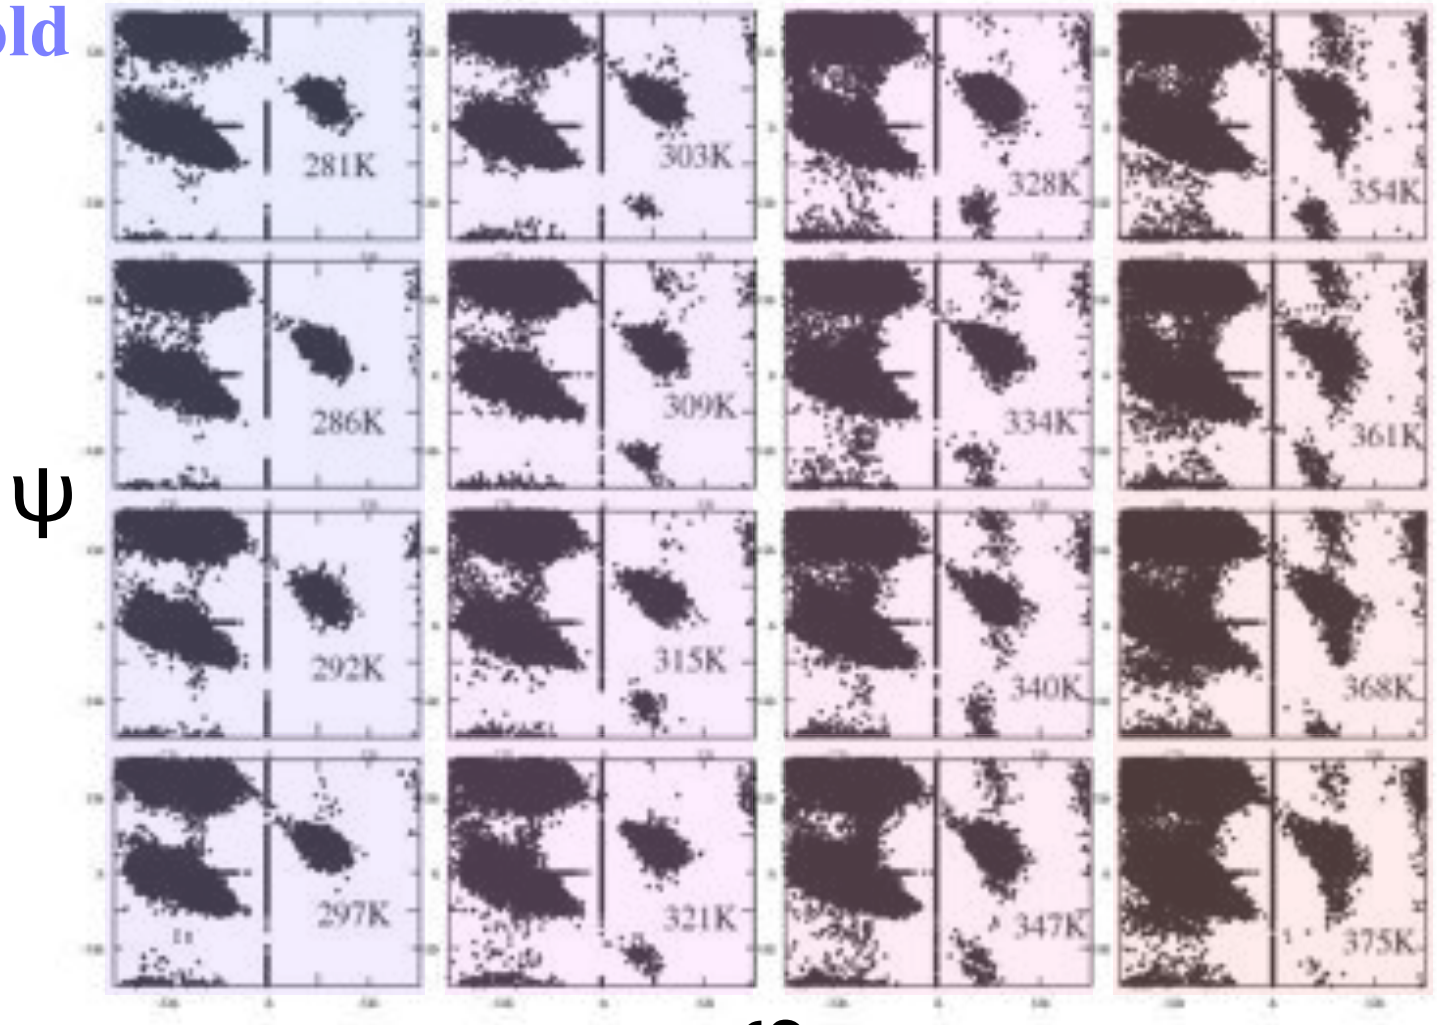

Data from Luo & Roux, ANL/UC.

Replica exchange for **umbrella sampling** on collective variables

Quaternion-based order parameters from collective variables module

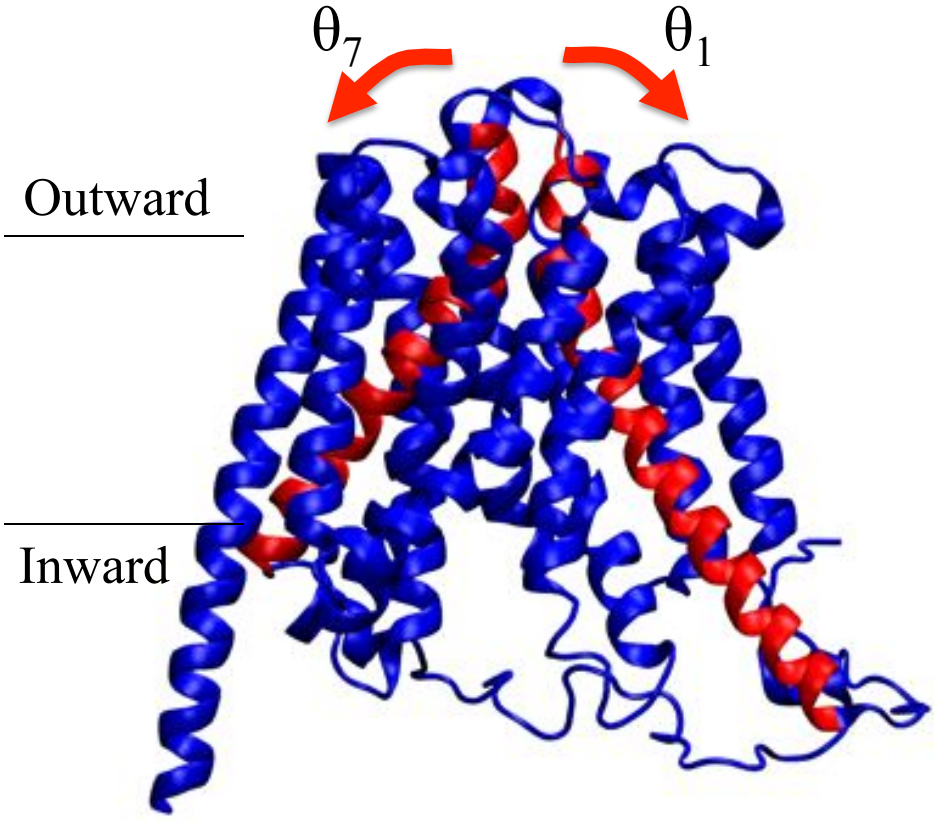

Inward-Facing↔Outward-Facing transition of GlpT transporter in explicit membrane/water environment (not shown)

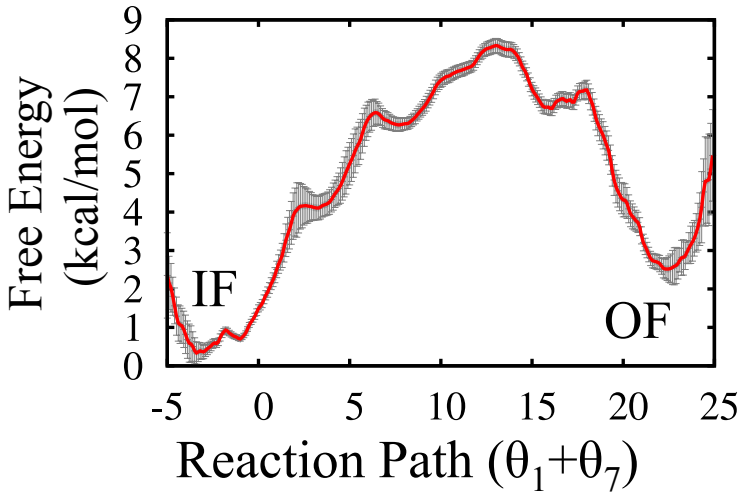

Efficient Reaction Path Sampling

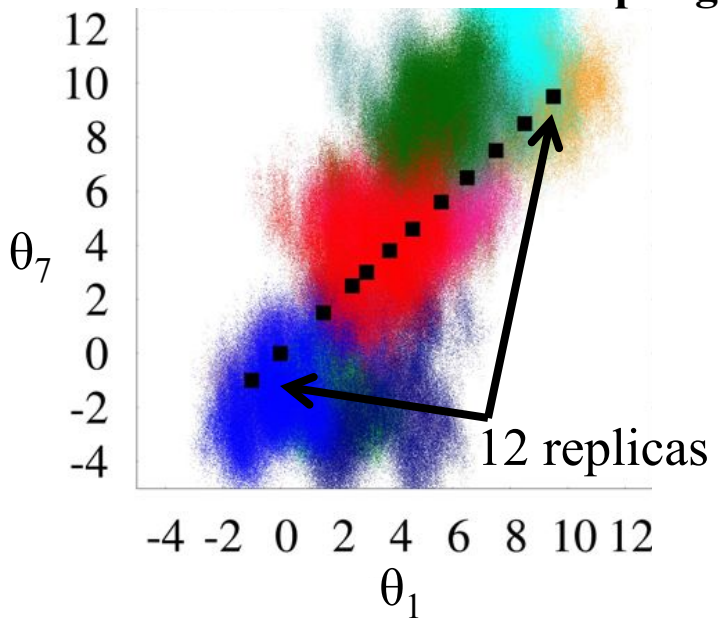

Collaborator Replica Exchange Applications

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

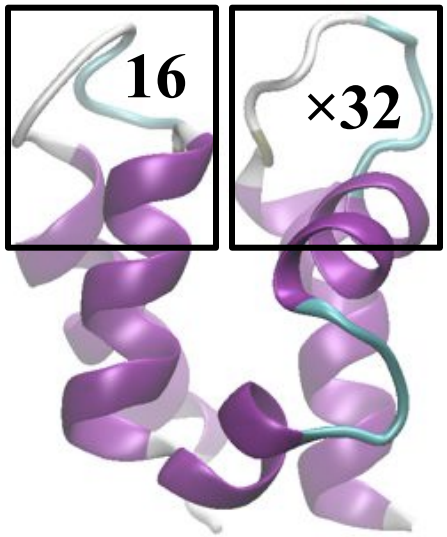

Calbindin D_{9k} (apo)
12,251 atoms/replica
16×32 = 512 replicas

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

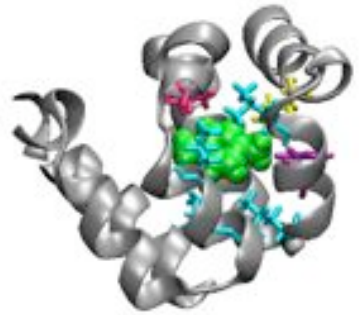

T4 Lysozyme/L99A

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

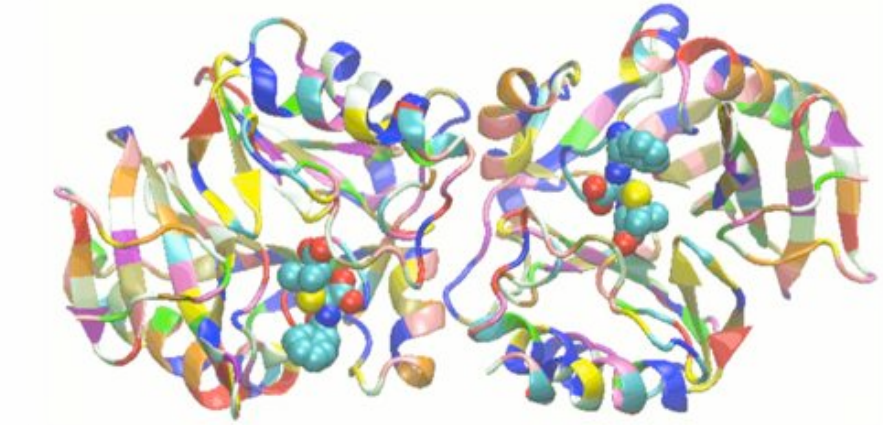

New Delhi Metallo-b-lactamase (NDM-1)

Thanks to: NIH, NSF, DOE,
NVIDIA (**Sarah Tariq**, Sky Wu,
Justin Luitjens, Nikolai Sakharnykh),
Cray (Sarah Anderson, Ryan Olson),
PPL (Eric Bohm, Yanhua Sun, Gengbin Zheng)
and 17 years of NAMD and Charm++
developers and users.

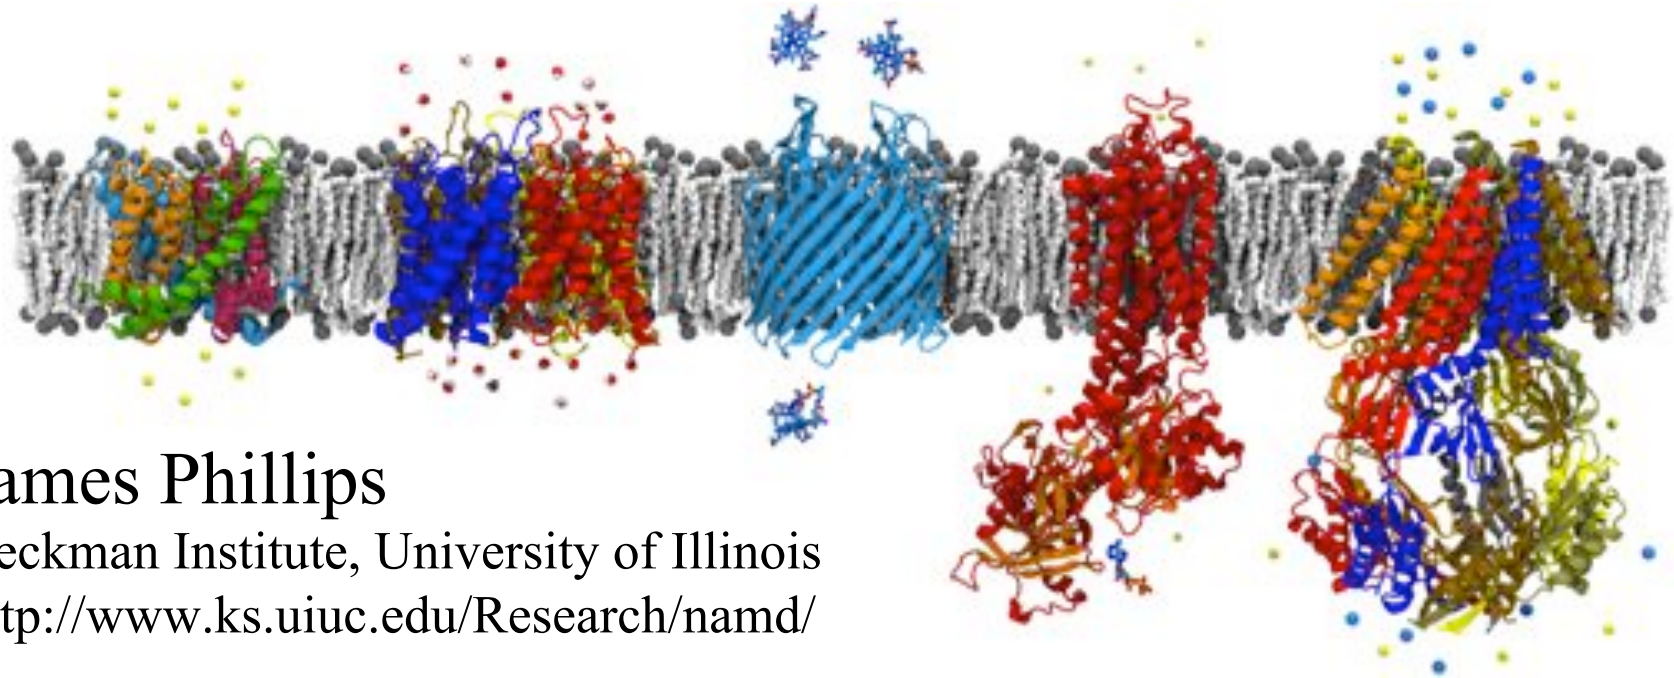

James Phillips
Beckman Institute, University of Illinois
http://www.ks.uiuc.edu/Research/namd/