Enhanced Charm++ Replica Facility For *The Computational Microscope*

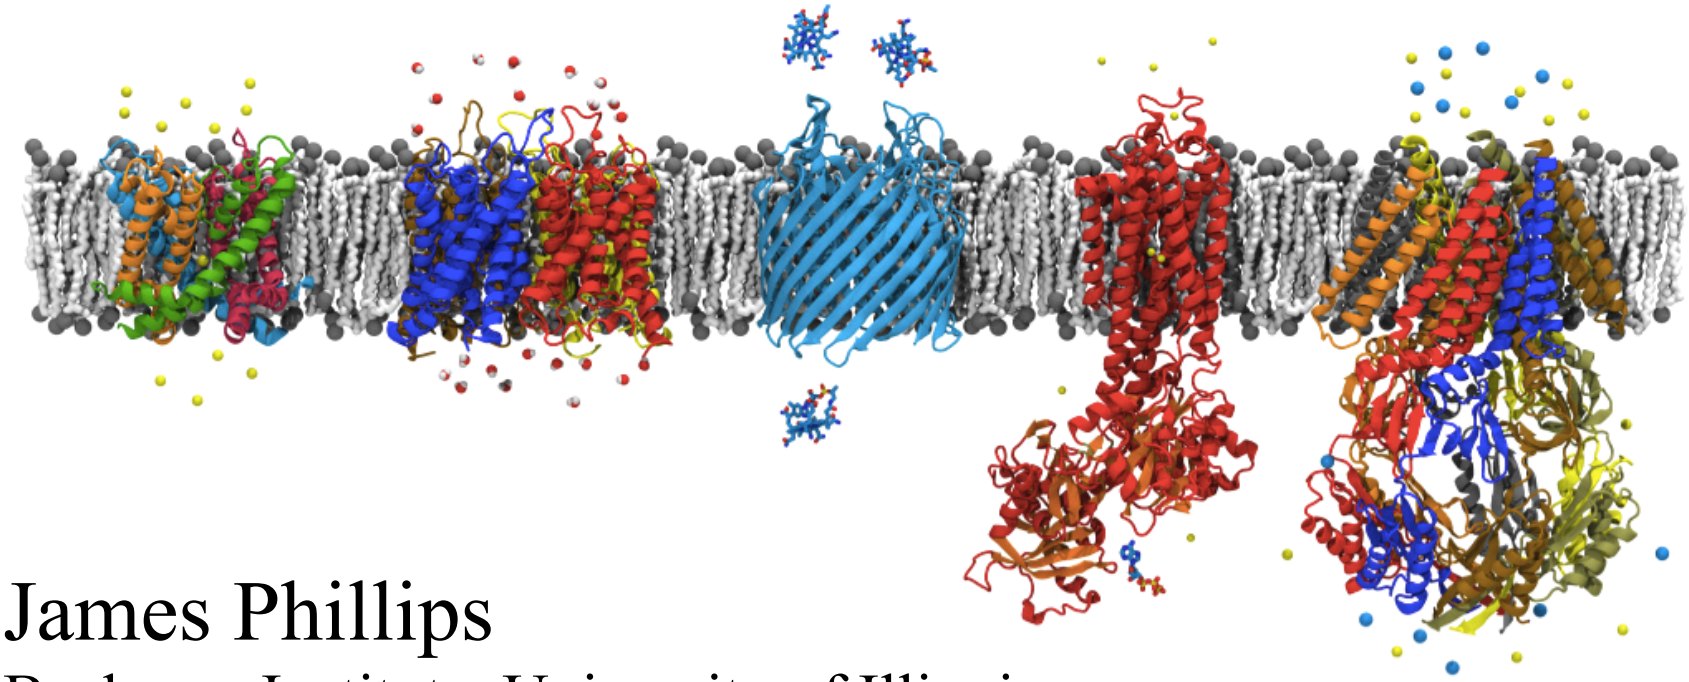

James Phillips
Beckman Institute, University of Illinois
http://www.ks.uiuc.edu/Research/namd/

NCSA/UIUC Enhanced Intellectual Services for Petascale Performance (NEIS-P^2) Program 2013 Blue Waters Symposium

NAMD: Scalable Molecular Dynamics

2002 Gordon Bell Award

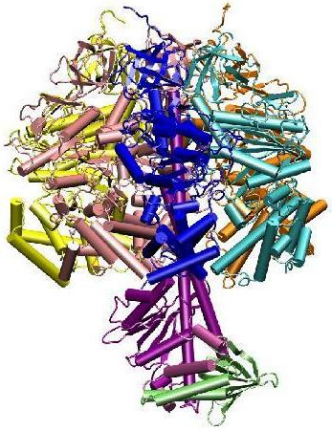

ATP synthase

PSC Lemieux

58,000 Users, 2900 Citations

Computational Biophysics Summer School

Blue Waters Target Application

Illinois Petascale Computing Facility

GPU Acceleration

NVIDIA Tesla

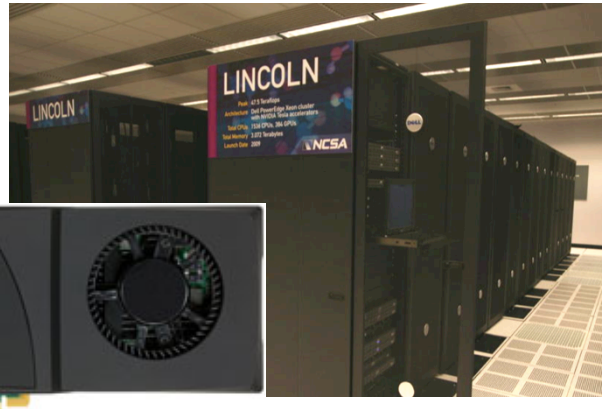

NCSA Lincoln

Computational Microscopy

Ribosome: synthesizes proteins from genetic information

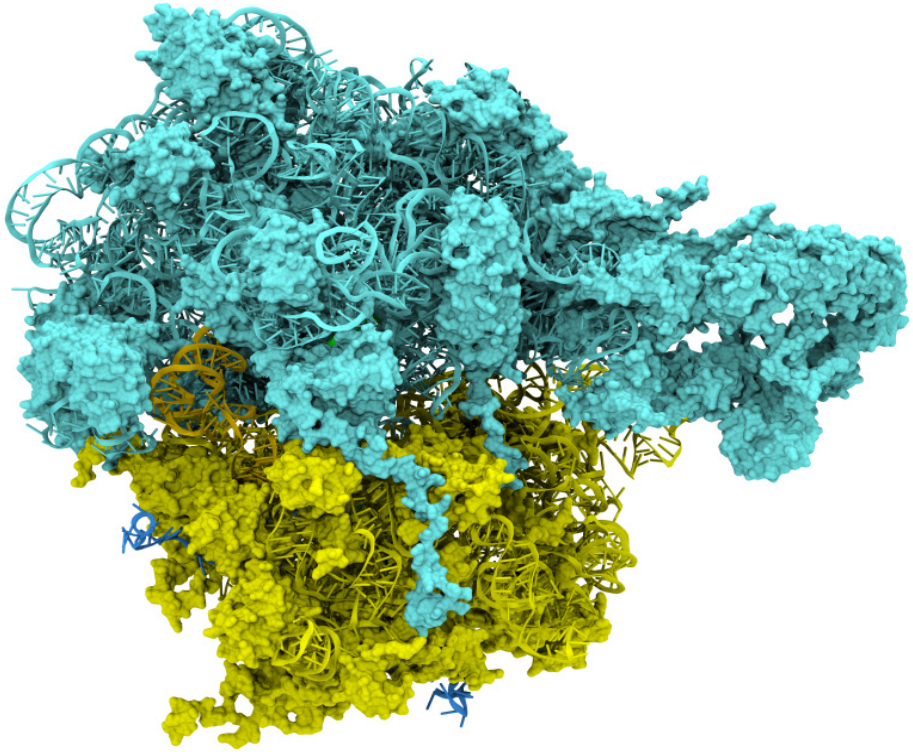

Silicon nanopore: bionanodevice for sequencing DNA

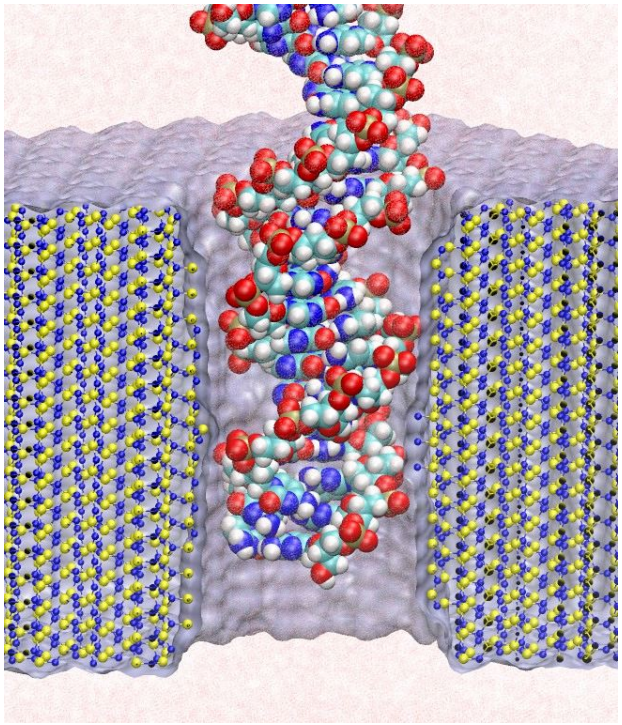

National Center for Research Resources

Molecular Mechanics Force Field

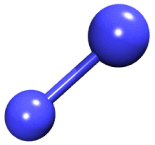

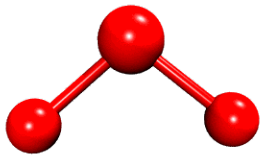

$$U(\vec{R}) = \underbrace{\sum_{bonds} k_i^{bond}(r_i - r_0)^2}_{U_{bond}} + \underbrace{\sum_{angles} k_i^{angle}(\theta_i - \theta_0)^2}_{U_{angle}} +$$

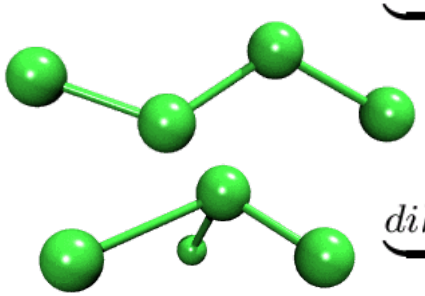

$$\underbrace{\sum_{dihedrals} k_i^{dihe}[1 + \cos(n_i\phi_i + \delta_i)]}_{U_{dihedral}} +$$

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

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.

Biomolecular Time Scales

Motion	Time Scale (sec)
Bond stretching	10^{-14} to 10^{-13}
Elastic vibrations	10^{-12} to 10^{-11}
Rotations of surface sidechains	10^{-11} to 10^{-10}
Hinge bending	10^{-11} to 10^{-7}
Rotation of buried side chains	10^{-4} to 1 sec
Allosteric transistions	10^{-5} to 1 sec
Local denaturations	10^{-5} to 10 sec

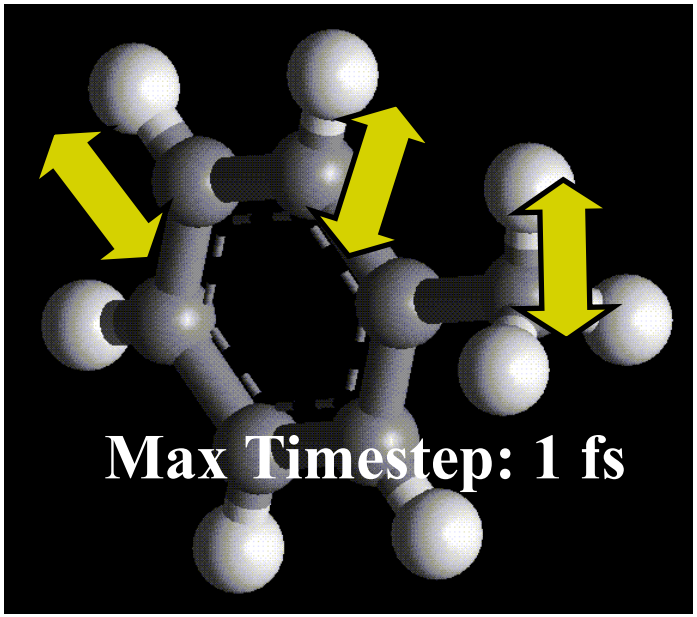

Parallel Programming Lab
University of Illinois at Urbana-Champaign

Siebel Center for Computer Science

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

Historical perspective

- 1993: First paper with Charm++ as a name for our system appeared in OOPSLA '93
  - Prior to that we had a C-based system called Charm
  - So, it has been (gulp) 20 years!
- NAMD development in Charm++ : 94-95
- 1994: Articulated Design Philosophy

Object based over-decomposition

- Let the programmer decompose computation into objects
  - Work units, data-units, composites
- Let an intelligent runtime system assign objects to processors
  - RTS can change this assignment during execution
- This empowers the RTS
  - The research agenda started with the simple precept above, just before NAMD,
  - Continued until now!

Charm++

- Multiple indexed collection of C++ objects
- Programmer expresses communication between objects

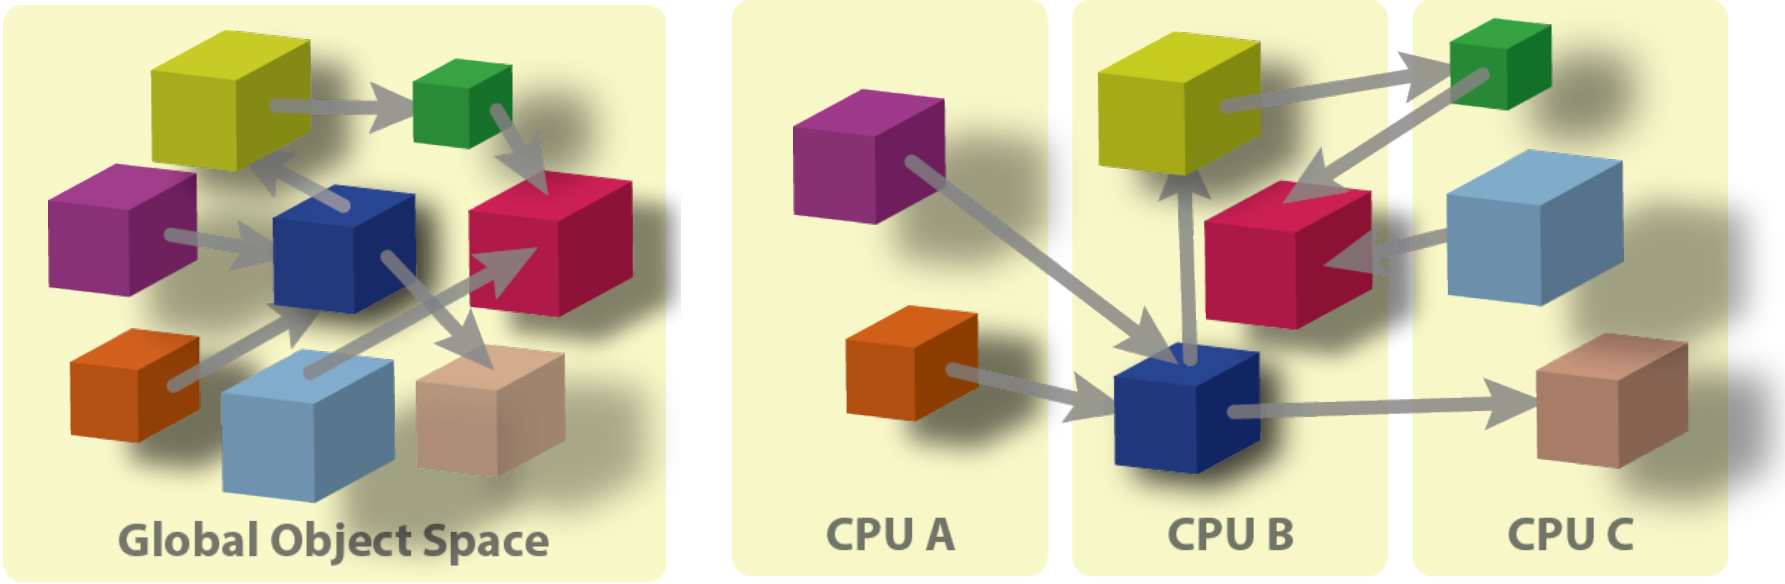

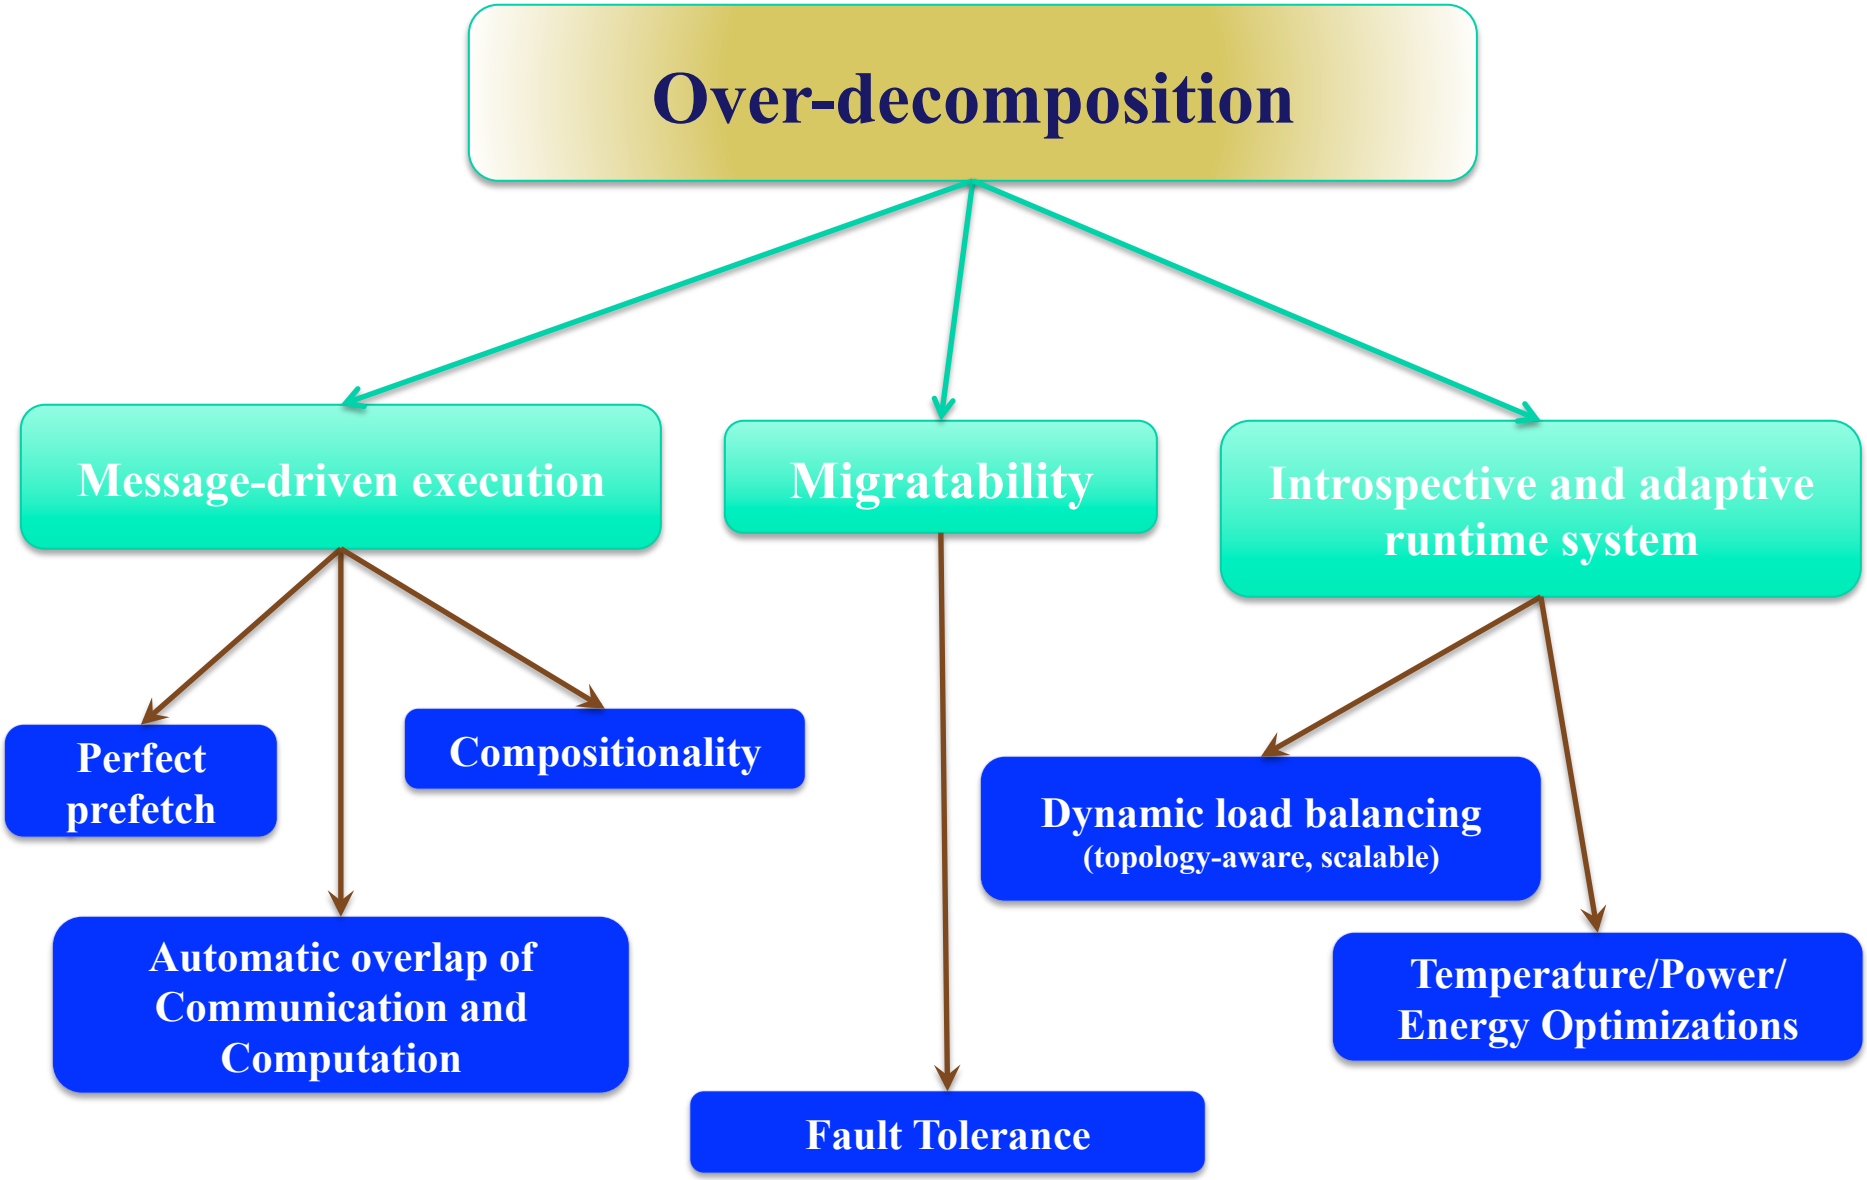
Over-decomposition
Message-driven execution
Migratability
Introspective and adaptive runtime system
Perfect prefetch
Compositionality
Automatic overlap of Communication and Computation
Fault Tolerance
Dynamic load balancing
(topology-aware, scalable)
Temperature/Power/ Energy Optimizations

Charm++ is a robust system

- Nightly build system
  - Dozens of machine, OS, Network, compiler combinations
  - 50,000 NAMD users depend on it
- Git repository, in addition to regular releases
- Online manuals, tutorials, ..
- Recent addition of Redmine (publicly accessible) to track bug-reports, bug-fixes, and feature development

Charm++ is scalable & efficient

- Demonstrated on many min-apps of our own
  - Barnes-Hut (scaled to 32k+ cores)
  - LU (HPL)
  - LeanMD
  - Structured Adaptive Mesh Refinement (AMR)
  - FFT
  - Sparse-Triangular solver
  - ….
  - HPC Challenge award in 2011 (shared w Chapel)

Some Charm++ Applications

NAMD

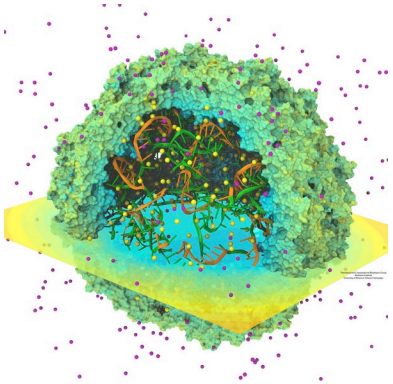

OpenAtom

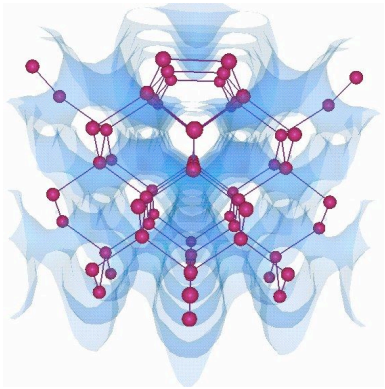

ChaNGa

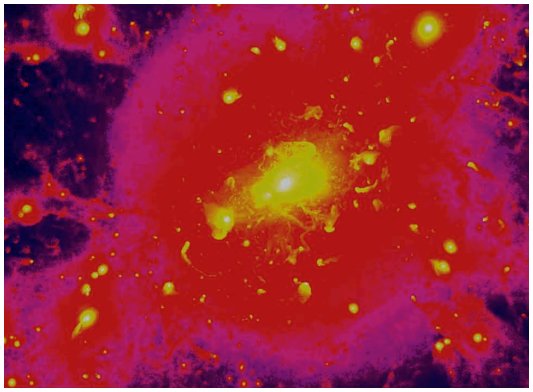

Rocstar

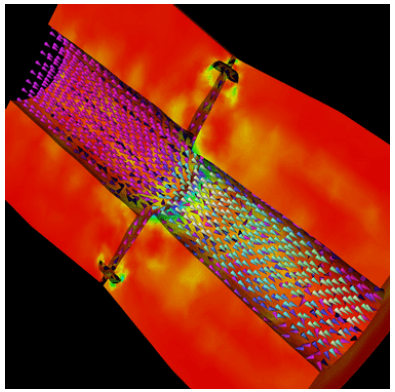

Others:

- **BRAMS**
- **ISAM**
- **Stochastic Optimization**
- **…**

Episimdemics

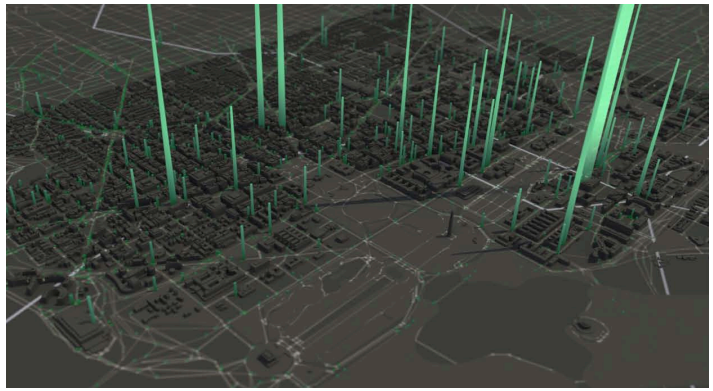

An upcoming book
Surveys seven major applications developed using Charm++

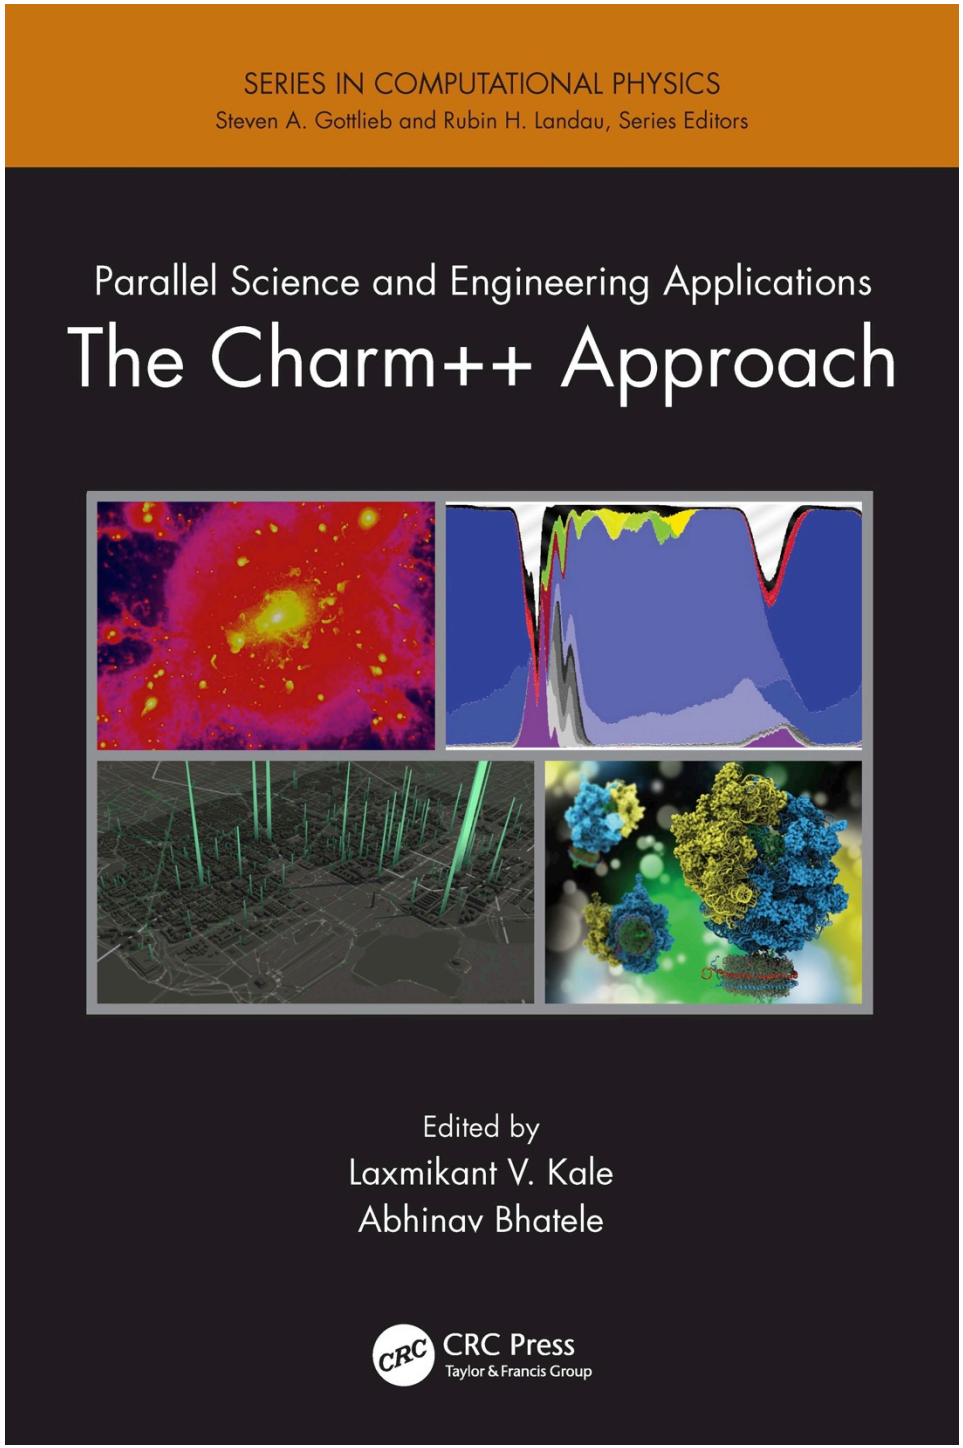

Charm++ Used by NAMD

- 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.*

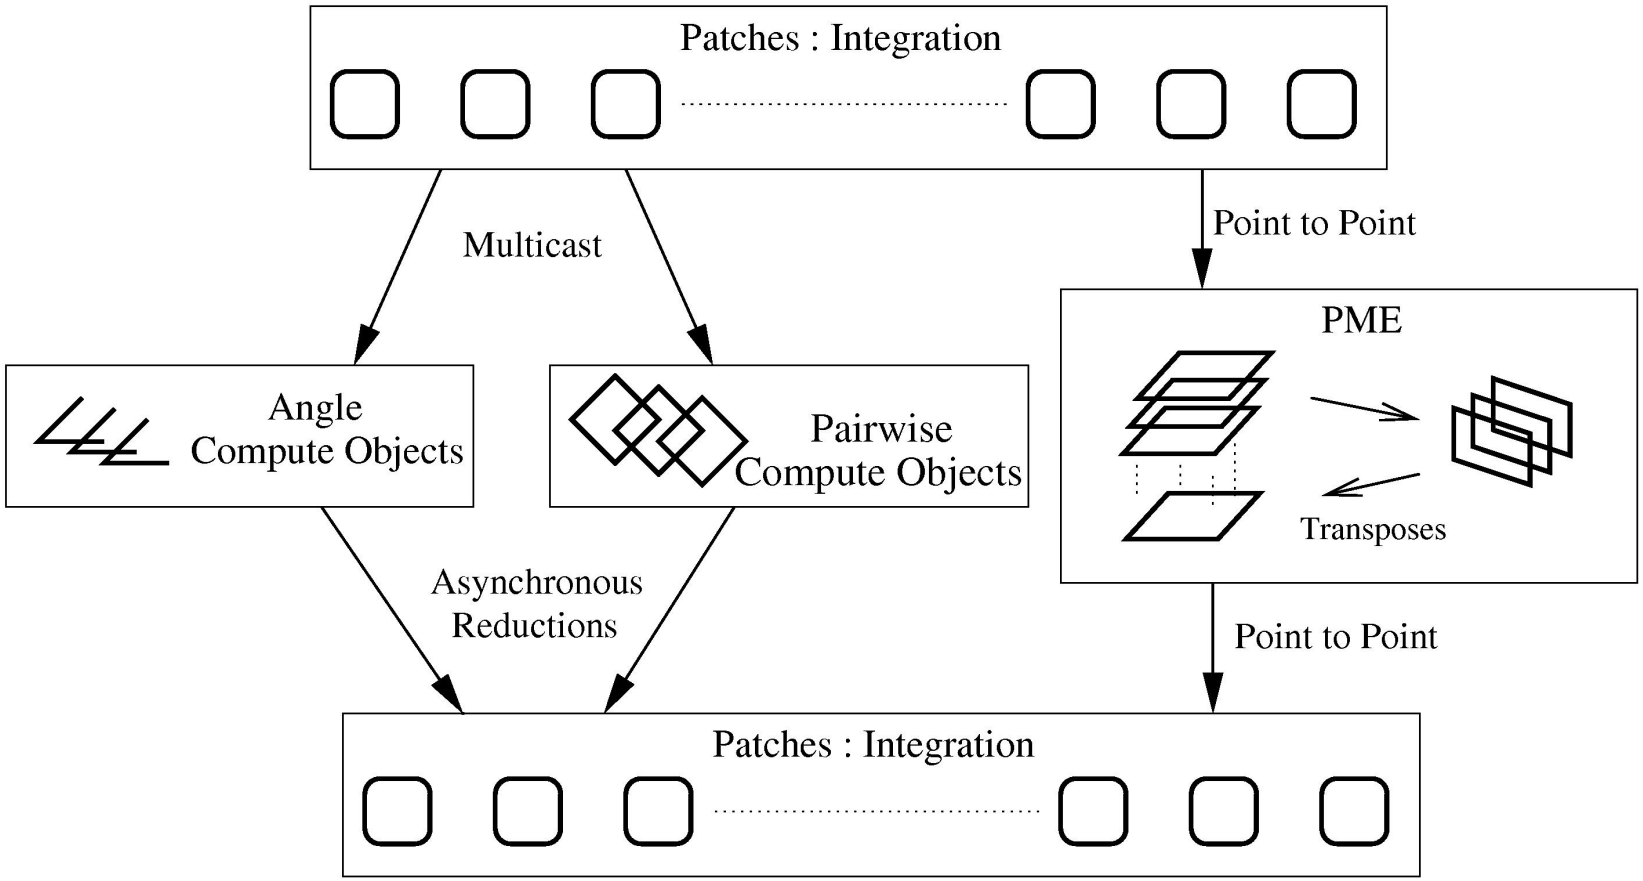

Objects are assigned to processors and queued as data arrives.

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.

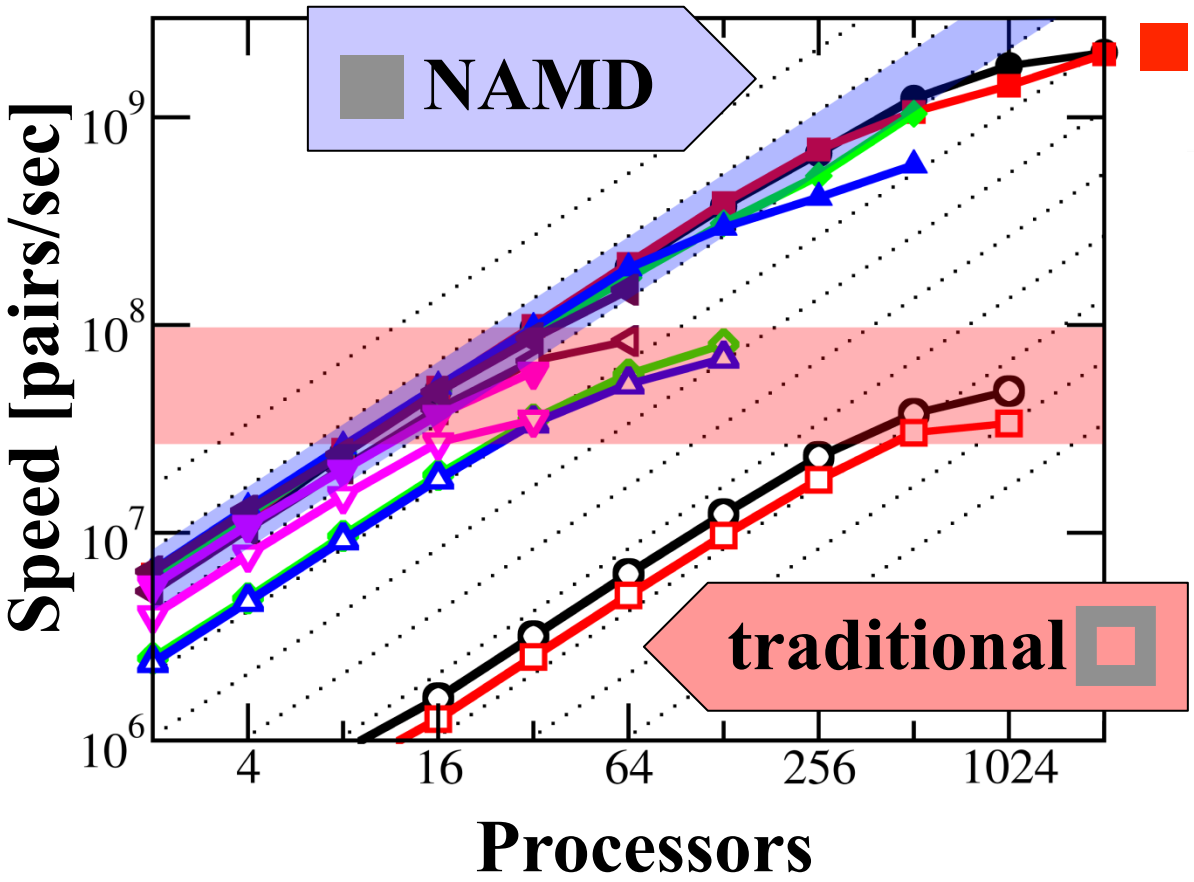

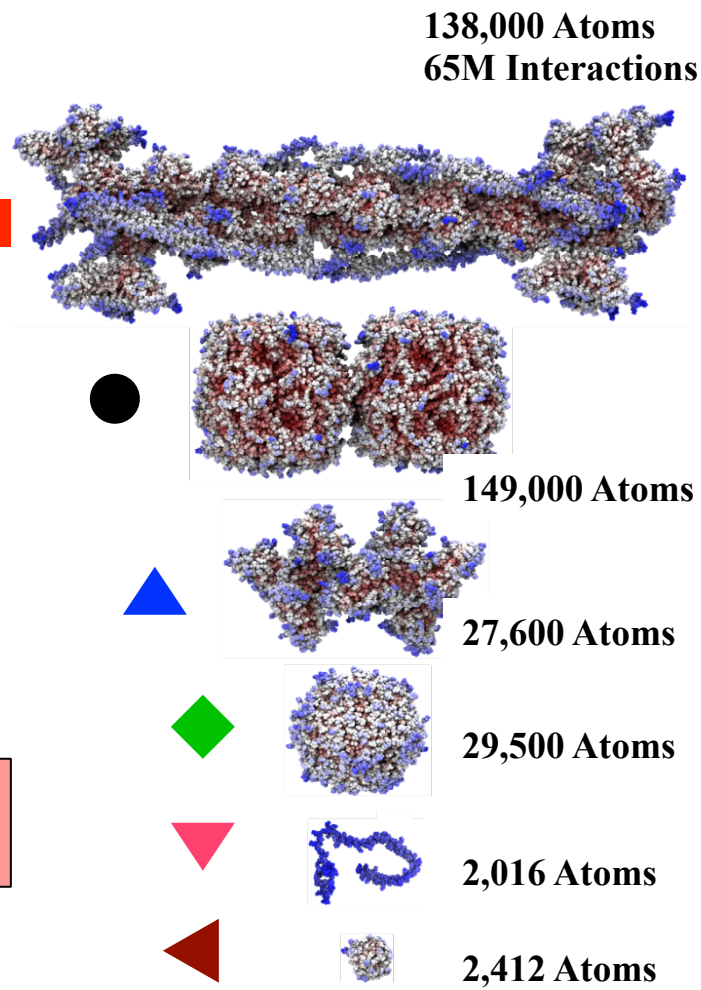

Tanner et al., J. Chem. Theory and Comp., 7:3635-3642, 2011

100M Atoms on Titan vs Jaguar

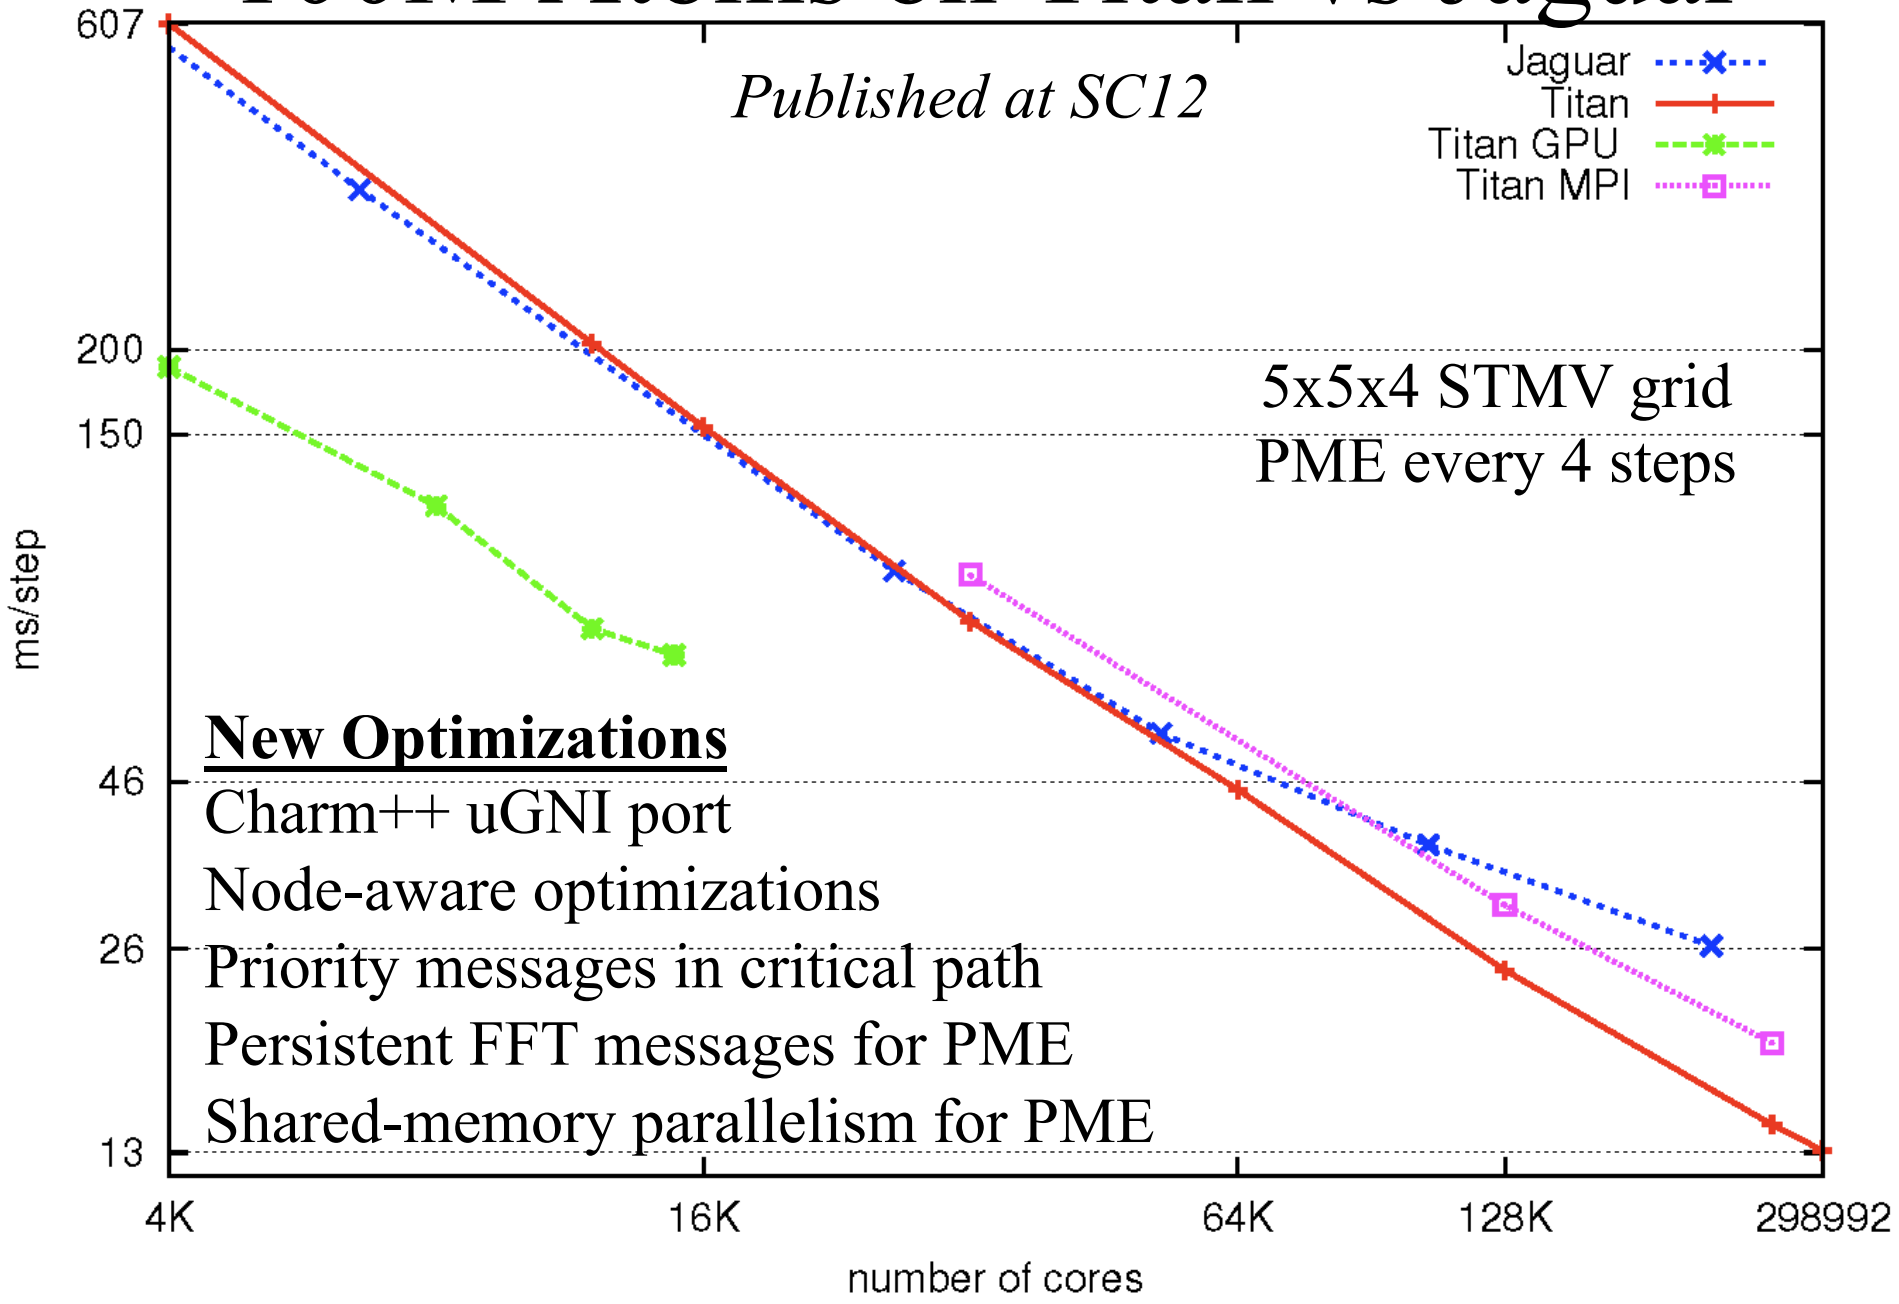

Larger machines enable larger simulations

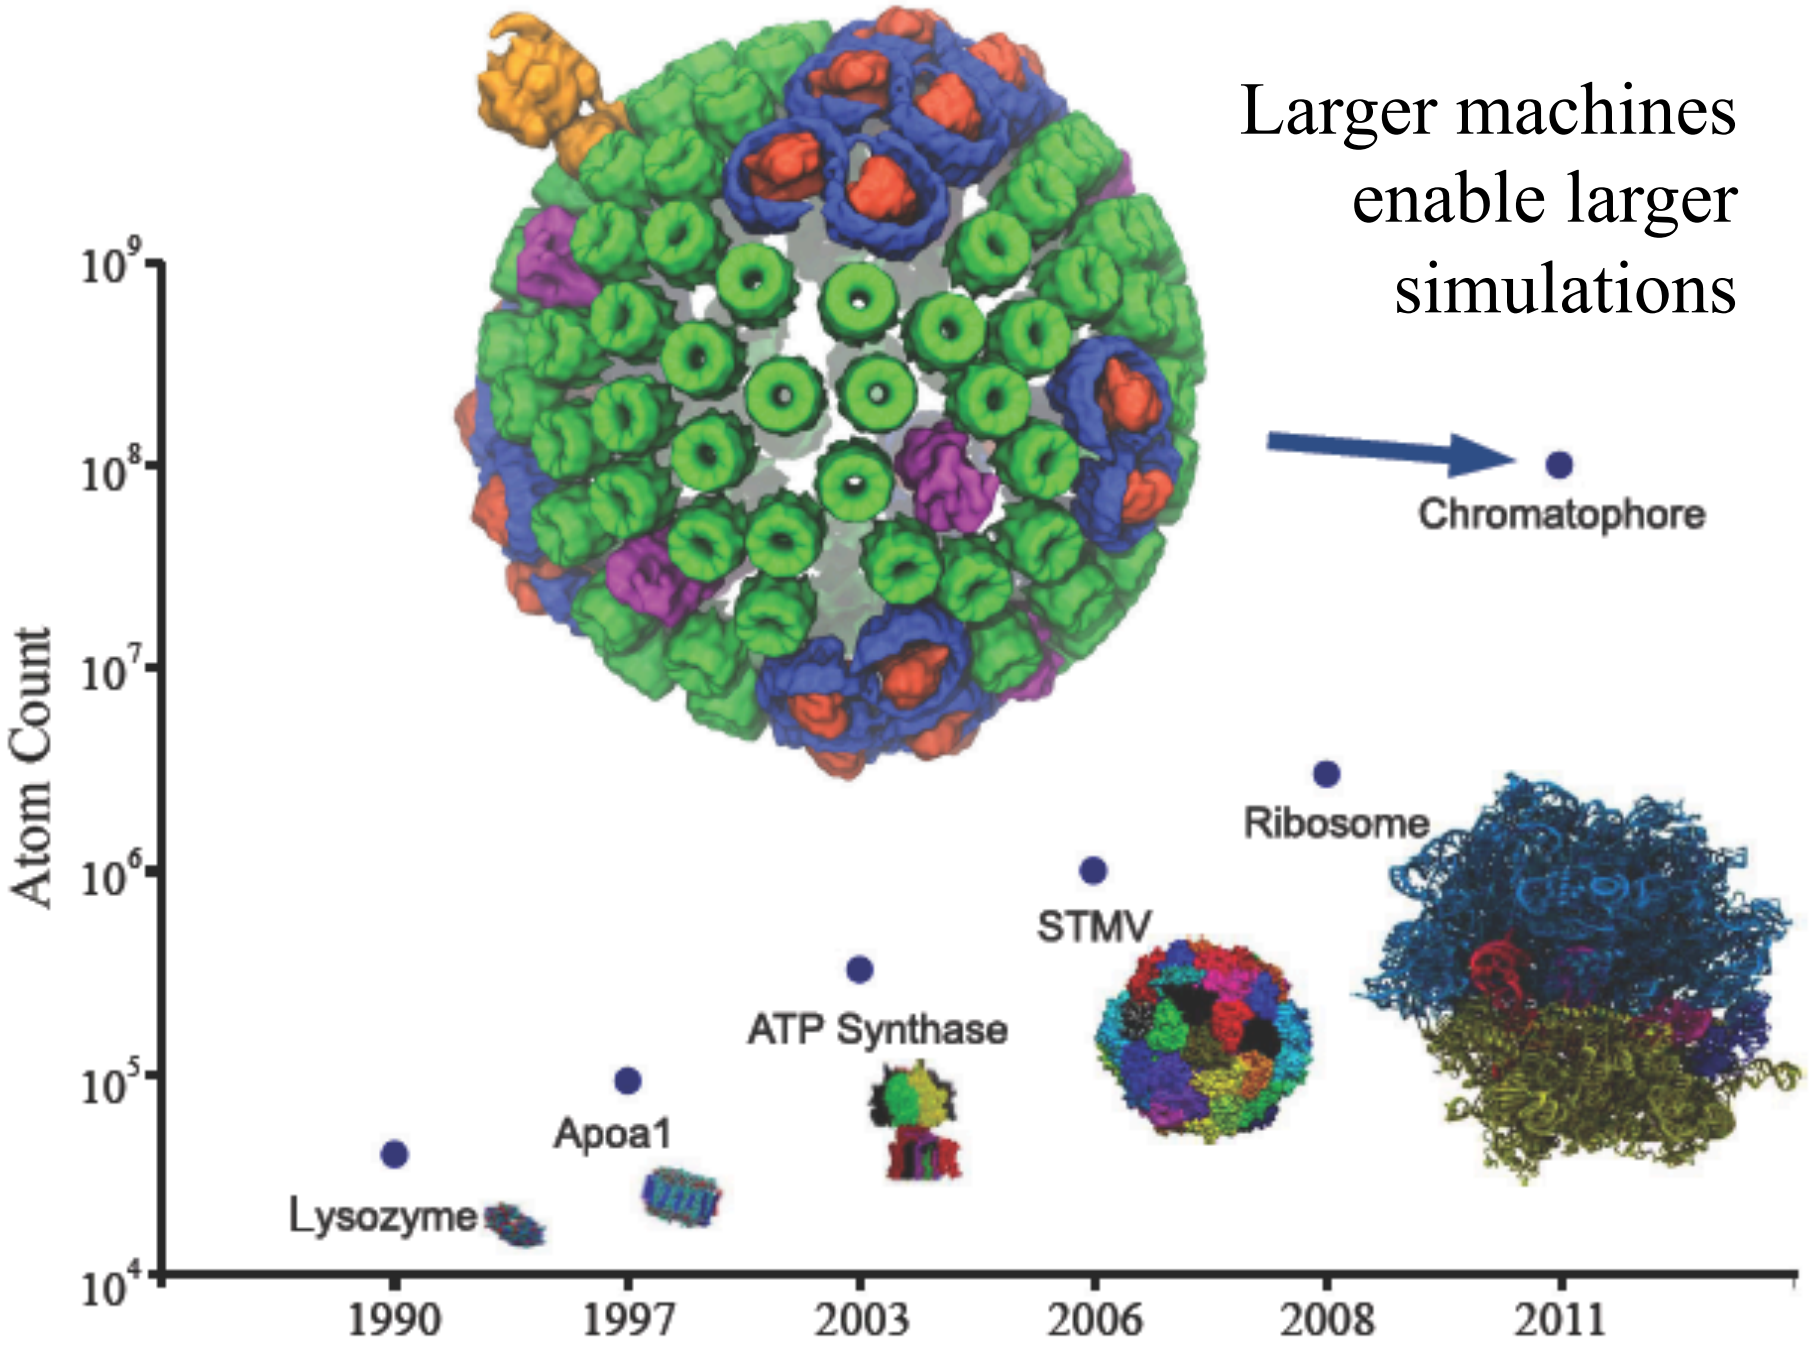

NIH BTRC for Macromolecular Modeling and Bioinformatics

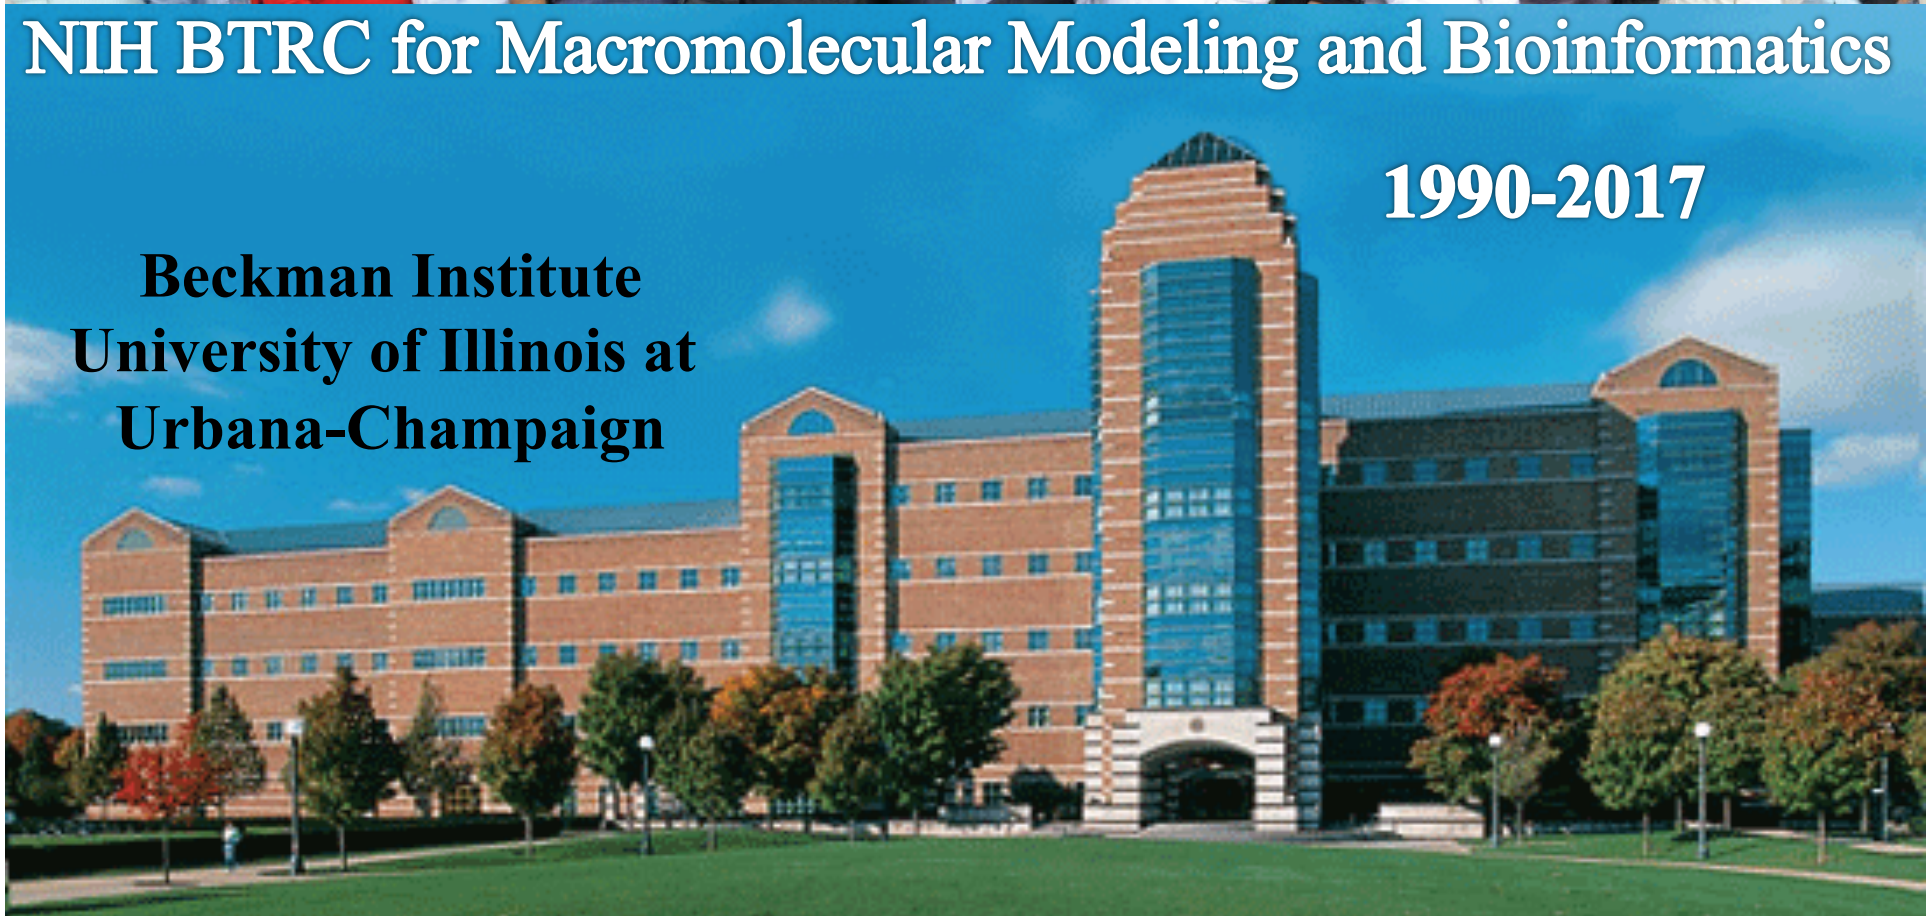

Physics of in vivo Molecular Systems

Biomolecular interactions span many orders of magnitude in space and time.
Center software provides multi-scale computational modeling.

femtoseconds → hours

Ångstrom → microns

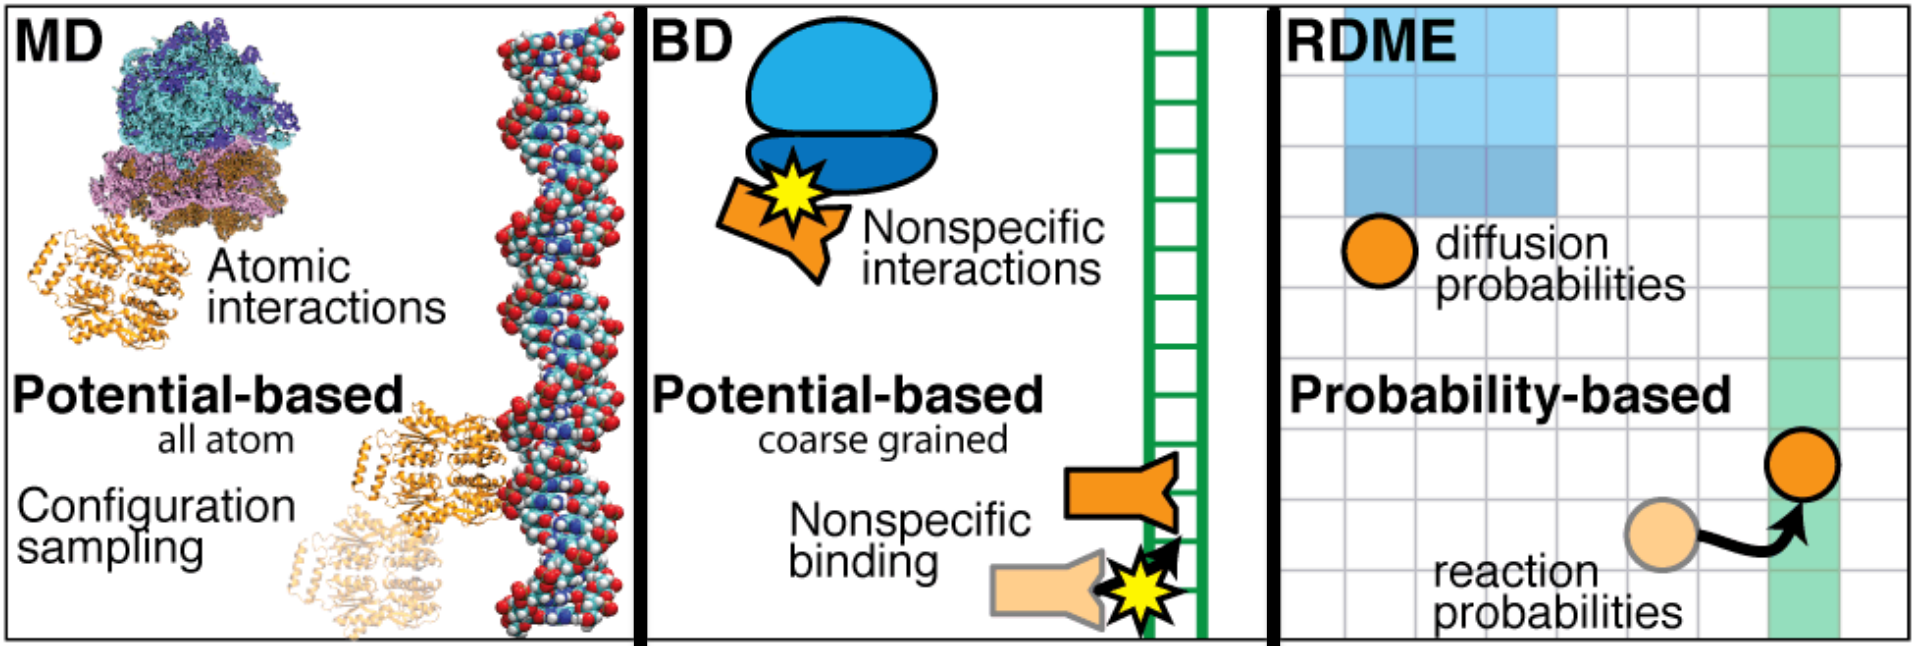

NAMD Scalable Molecular Dynamics
MDFF Molecular Dynamics Flexible Fitting
HMMM Highly Mobile Membrane Mimetic
VMD Visual Molecular Dynamics

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

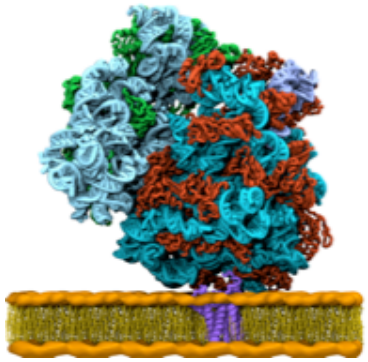

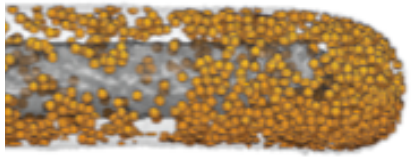

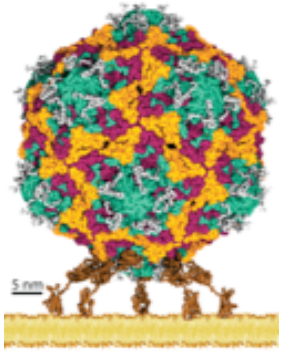

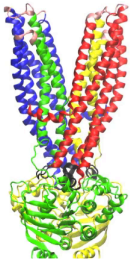

1. Ribosome	R. Beckmann (U. Munich) J. Frank (Columbia U.) T. Ha(UIUC) K. Fredrick (Ohio state U.) R. Gonzalez (Columbia U.)
2. Blood Coagulation Factors	J. Morrissey (UIUC) S. Sligar (UIUC) C. Rienstra (UIUC) G. Gilbert (Harvard)
3. Whole Cell Behavior	W. Baumeister (MPI Biochem.) J. Xiao (Johns Hopkins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

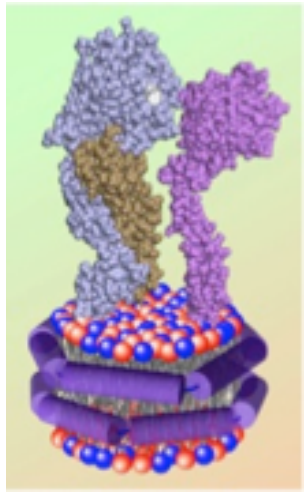

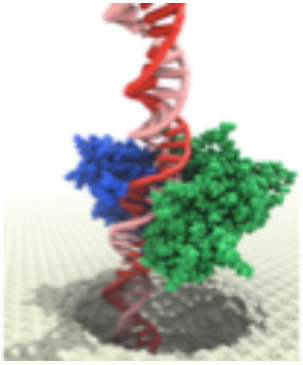

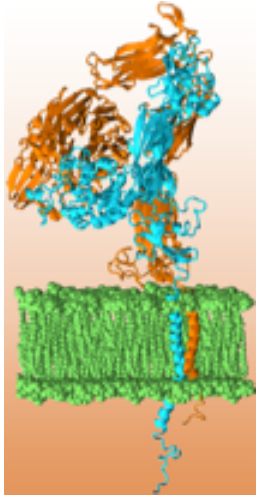

Viral Infection Driving Projects

Poliovirus

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

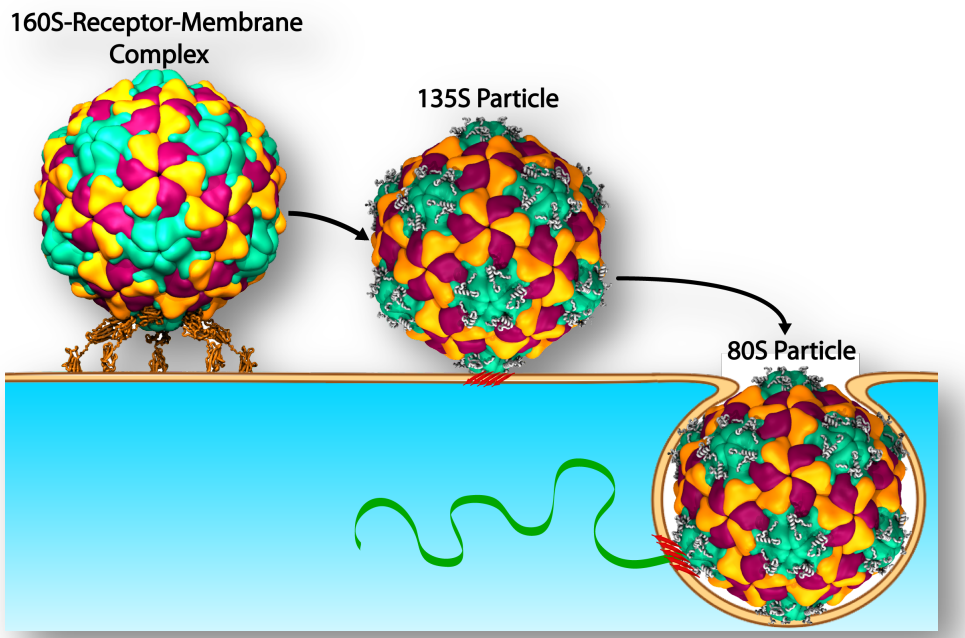

Human Immunodeficiency Virus 1

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

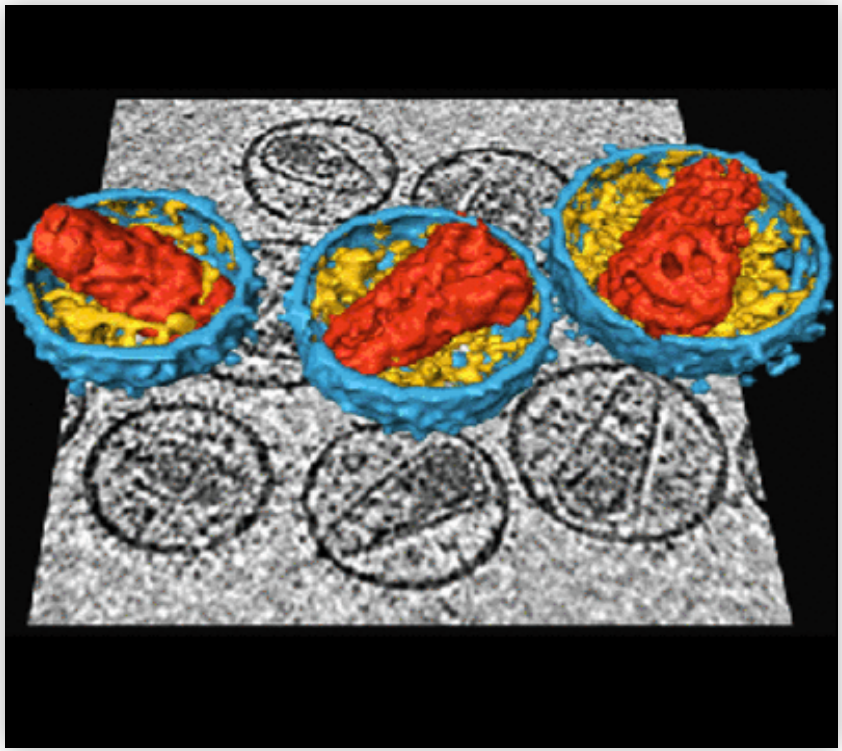

Briggs et al. Structure (2006) 14:15-20.

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. Biol. of the Cell 2005, 16:279-291)

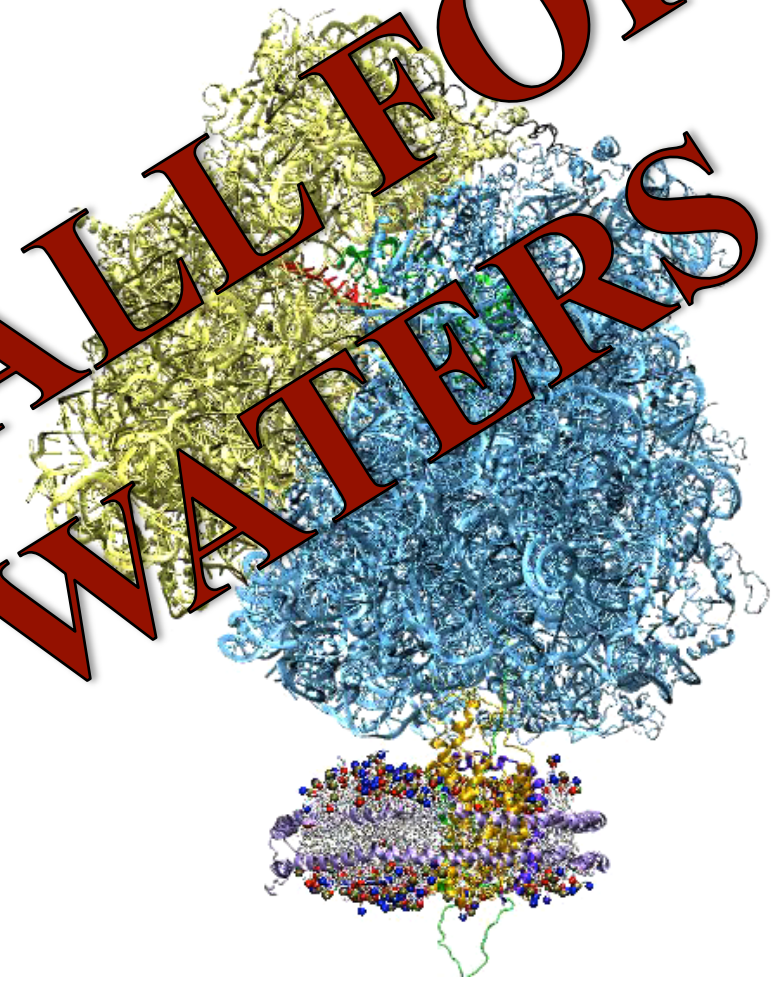

TOO SMALL FOR BLUE WATERS

NAMD 2.6 Replica Exchange (2006)

- Implemented entirely in Tcl:
  - Rapid development
  - User-modifiable
  - Portable
- Master Tcl interpreter:
  - Splits allocated nodes and launches NAMD slaves
  - Communicates with slaves via Unix network sockets
  - Issues commands to slave Tcl interpreters

} Not portable

- Slave Tcl interpreter in NAMD slaves:
  - Listen to master socket
  - Run commands
  - Return results
- Supports parallel tempering

NAMD 2.9 MPI-Based Replica Exchange

- Small patch for Charm++ MPI machine layer startup code:
  - Call MPI_Comm_split() to create “row” and “column” communicators
  - Charm++ uses “row” communicators instead of MPI_COMM_WORLD
  - Tcl interface to MPI_Send(), MPI_Recv(), MPI_Sendrecv() on “columns”
- 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)

  } 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

NEIS-P^2 Project: Bypass MPI

- Modify Charm++ "LRTS" generic machine layer
- Enable fine-grained inter-replica communications
- Optimize replica partitions for machine topology

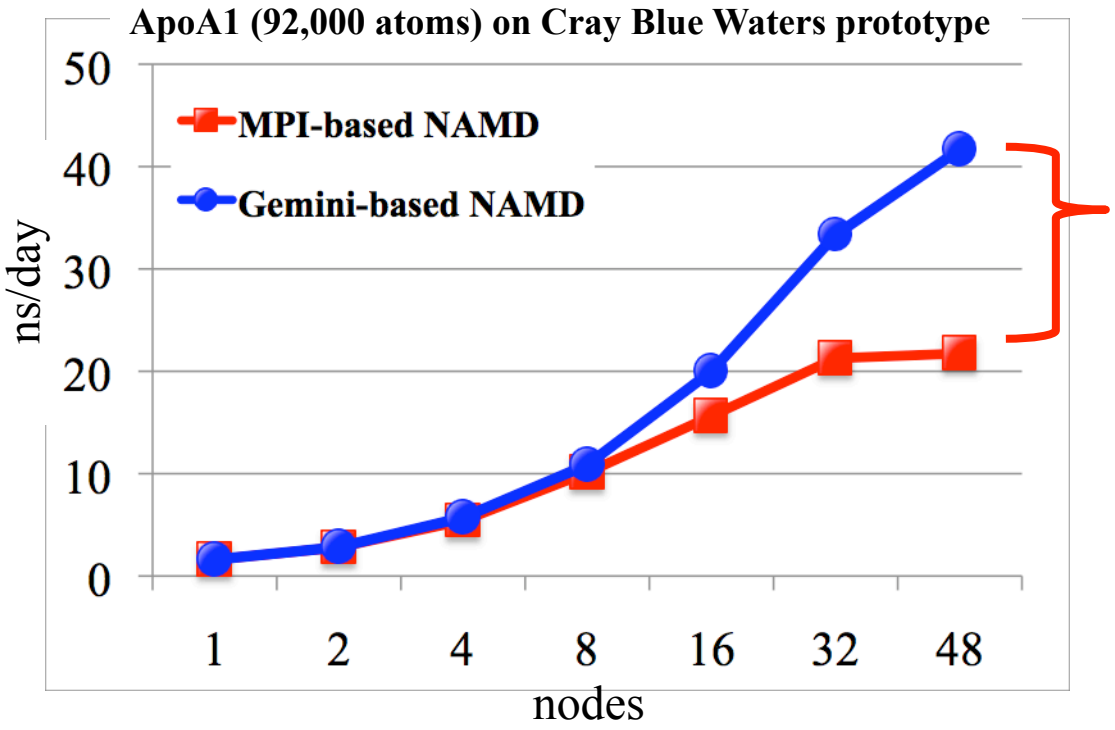

LRTS Gemini doubles usable XE nodes for strong scaling and is essential for use of multiple GPU-accelerated nodes per replica.

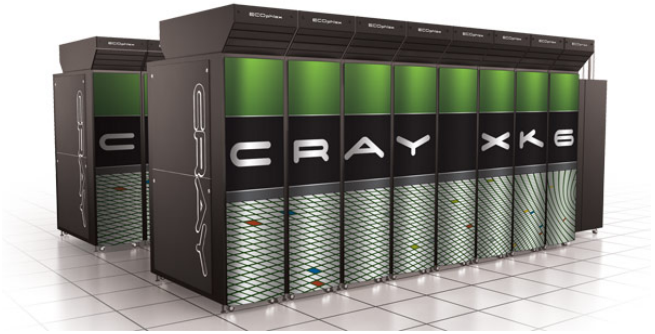

NAMD 2.10 Scalable Replica Exchange

- More general Charm++ integration:
  - NAMD 2.9 used MPI communicator splitting
  - NAMD 2.10 splits replicas in Charm++ 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)

  } Same Tcl scripts as NAMD 2.9

  - Finite Temperature String method
  - Nudged elastic band

  } 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

NEIS-P^2 Collateral Benefits

- Adds a new dimension of parallelism to Charm++
  - Existing code runs unmodified within a partition
  - Enables evolution from petascale to exascale
- Supports partitions of different sizes
  - Single-node master partition
  - Small output-buffering partition
  - Varying performance requirements
- Enables pruning of undesirable nodes from job
  - Extract compact set of nodes from allocation
  - Eliminate Gemini shared with other jobs
- Enables replication-based soft fault tolerance

Free Energy Landscapes of GlpT Transporter Obtained from **Bias-Exchange** (Umbrella Sampling) Simulations

Mahmoud Moradi
Tajkhorshid Lab, May 2013

Inward-Facing↔Outward-Facing (IF-OF) Transition of GlpT Transporter

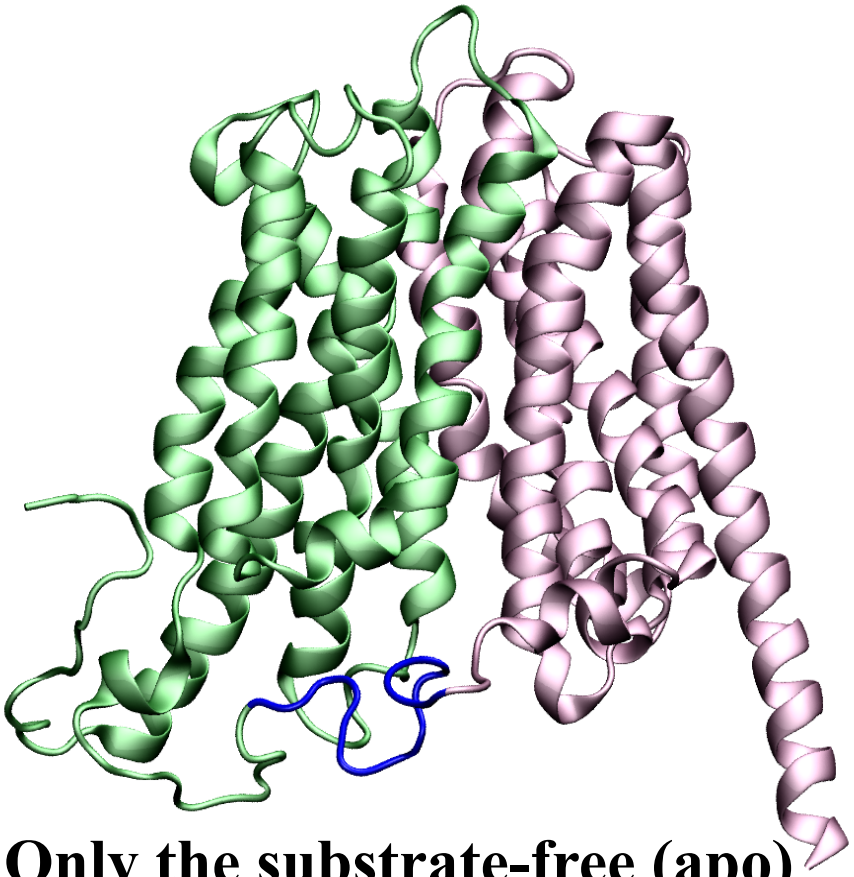

Only the substrate-free (apo) IF structure is available.

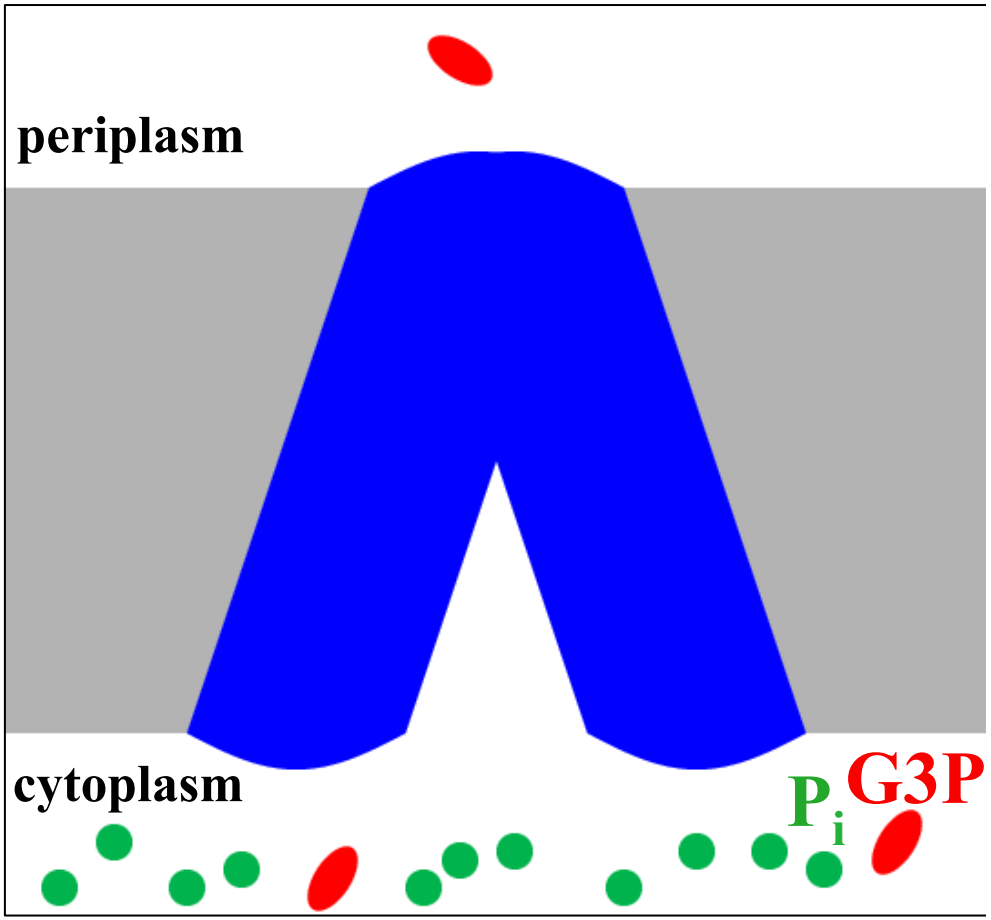

The proposed mechanism.

Huang, *et al.*, *Science* 301, 616 (2003).

Molecular Dynamics Simulation Protocols

1. Equilibrating GlpT in explicit solvent/membrane in the apo IF state (~120,000 atoms).
2. Finding/optimizing a biasing protocol capable of inducing the IF➔OF transition in a reliable way through advanced *colvar* (collective variables) techniques.
3. Free energy calculations using bias-exchange umbrella sampling simulations based on the optimized biasing protocol.
4. Repeating the simulations for a substrate-bound GlpT and comparing the results to those of the apo GlpT.

IF↔OF transition induced by imposing rotational change on helices TM1 and TM7 (using orientation quaternions)

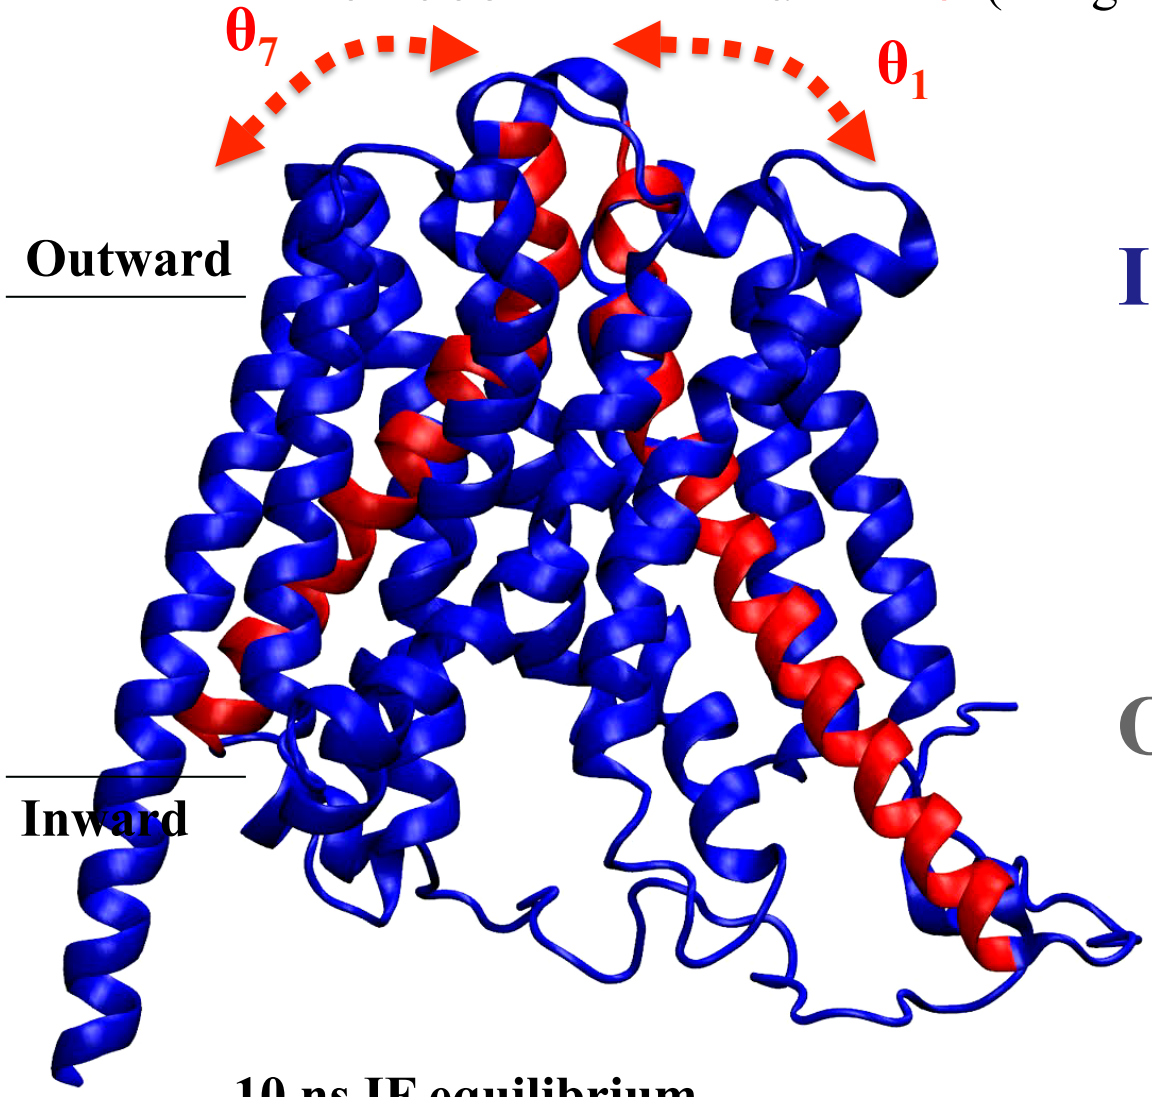

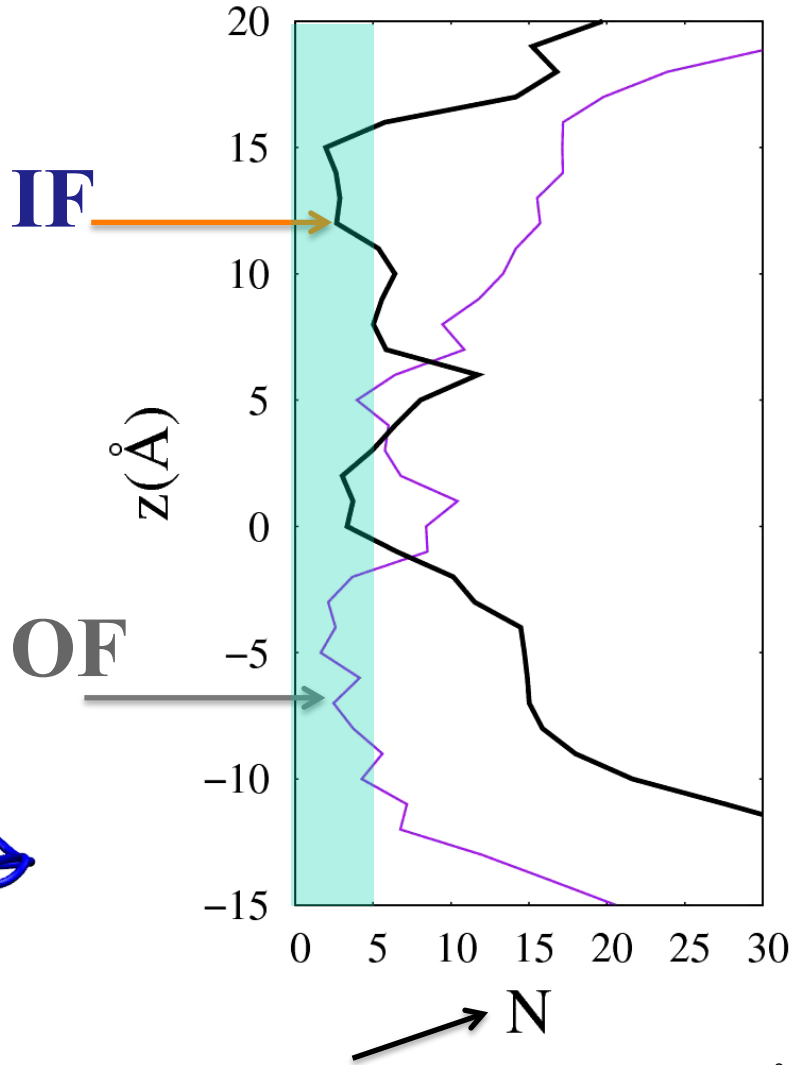

10 ns IF equilibrium
20 ns nonequilibrium (IF➔OF)
10 ns OF equilibrium

Number of water molecules along the pore (per Å)
(verifying the transition from IF to OF)

Free Energy Calculations Using Bias-Exchange Umbrella Sampling

- *Colvars*: orientation quaternions associated with helices 1 and 7 (two quaternion-based colvars).
- Initial conformations: 12 conformations selected from the optimized transition pathway (nonequilibrium simulations).
- Colvar centers and force constants associated with different replicas: Optimized (empirically) to get similar exchange rates for all neighboring replicas (28-35%).

Only 12 replicas (each running for 40 ns) are enough for a good sampling.

Following the trajectory of one of the 12 shows that it visits most of the windows in 40 ns

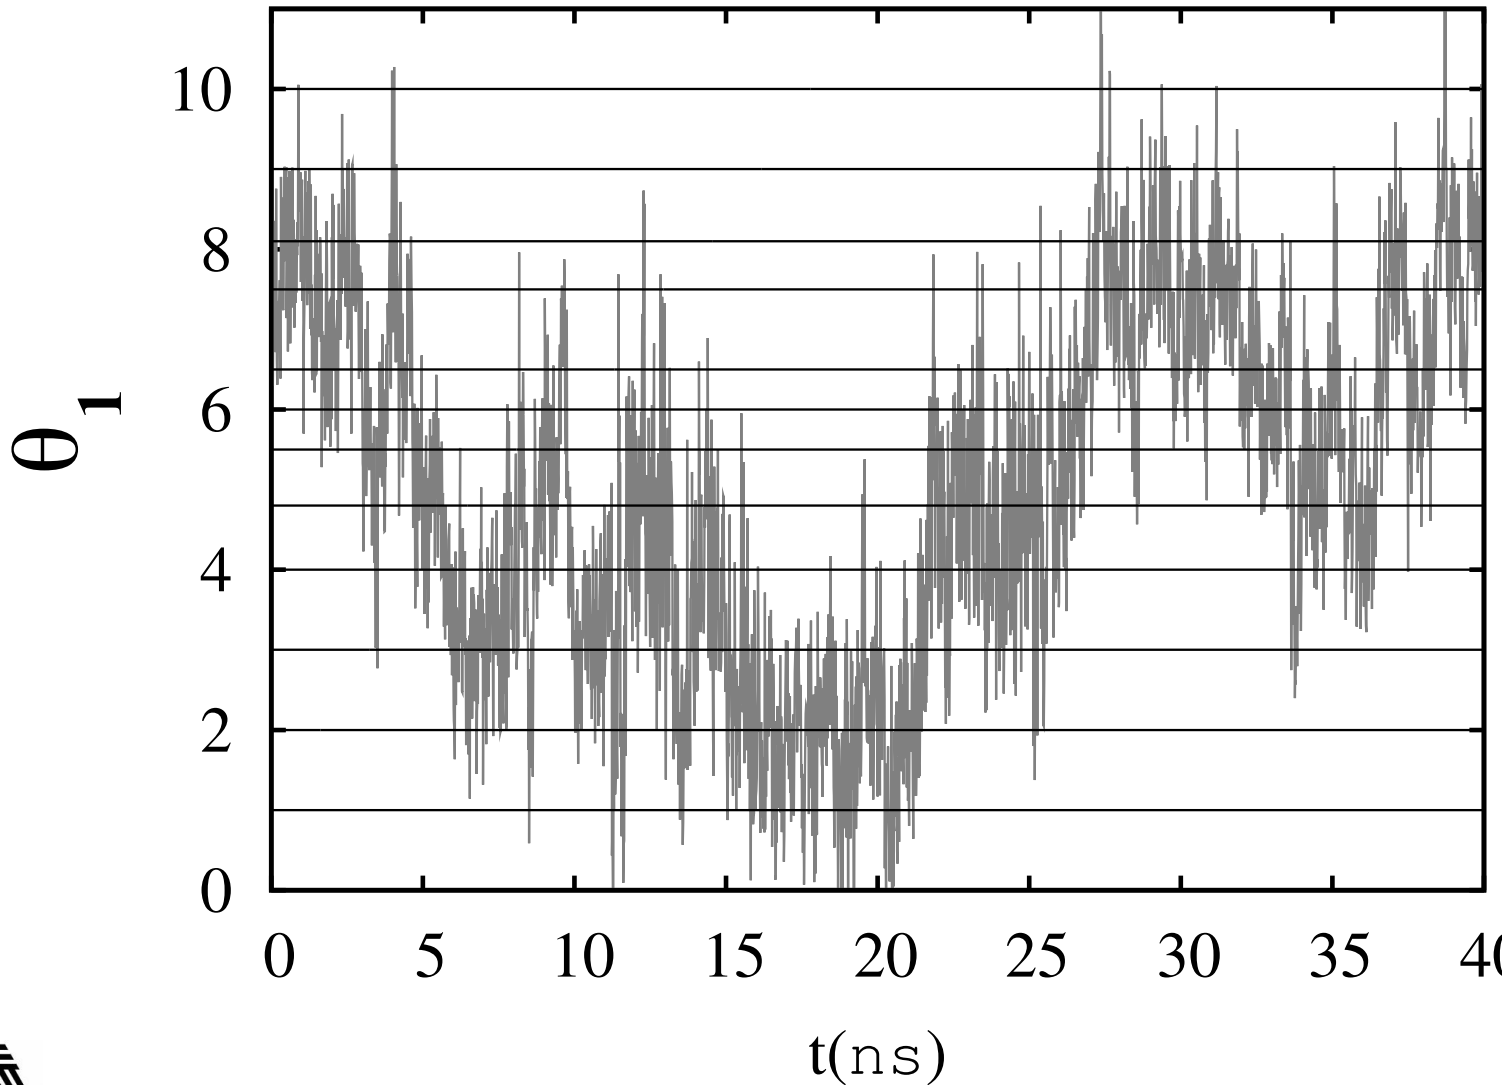

Centers of the windows in the θ_1 space.

Free Helices - Involved in the cytoplasmic opening/closing

GlpT Free Energy Maps

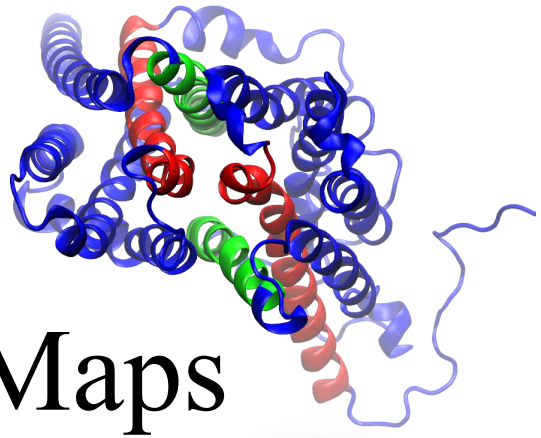

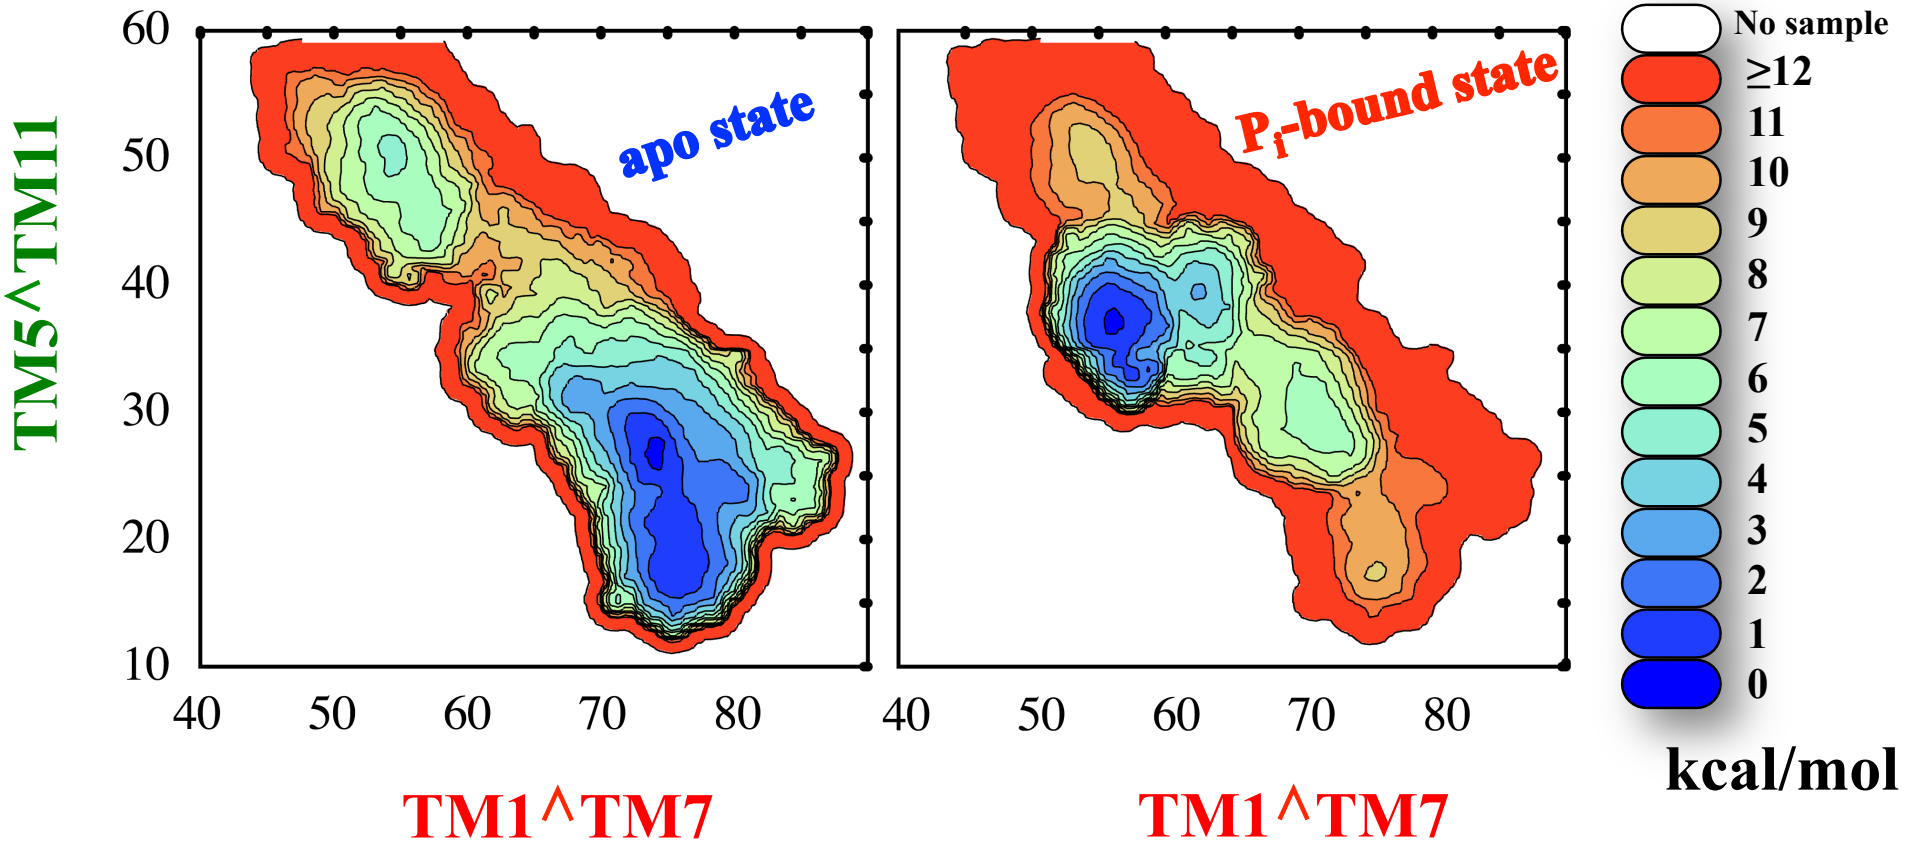

Steered Helices - Involved in the periplasmic opening/closing

GlpT Free Energy Minima

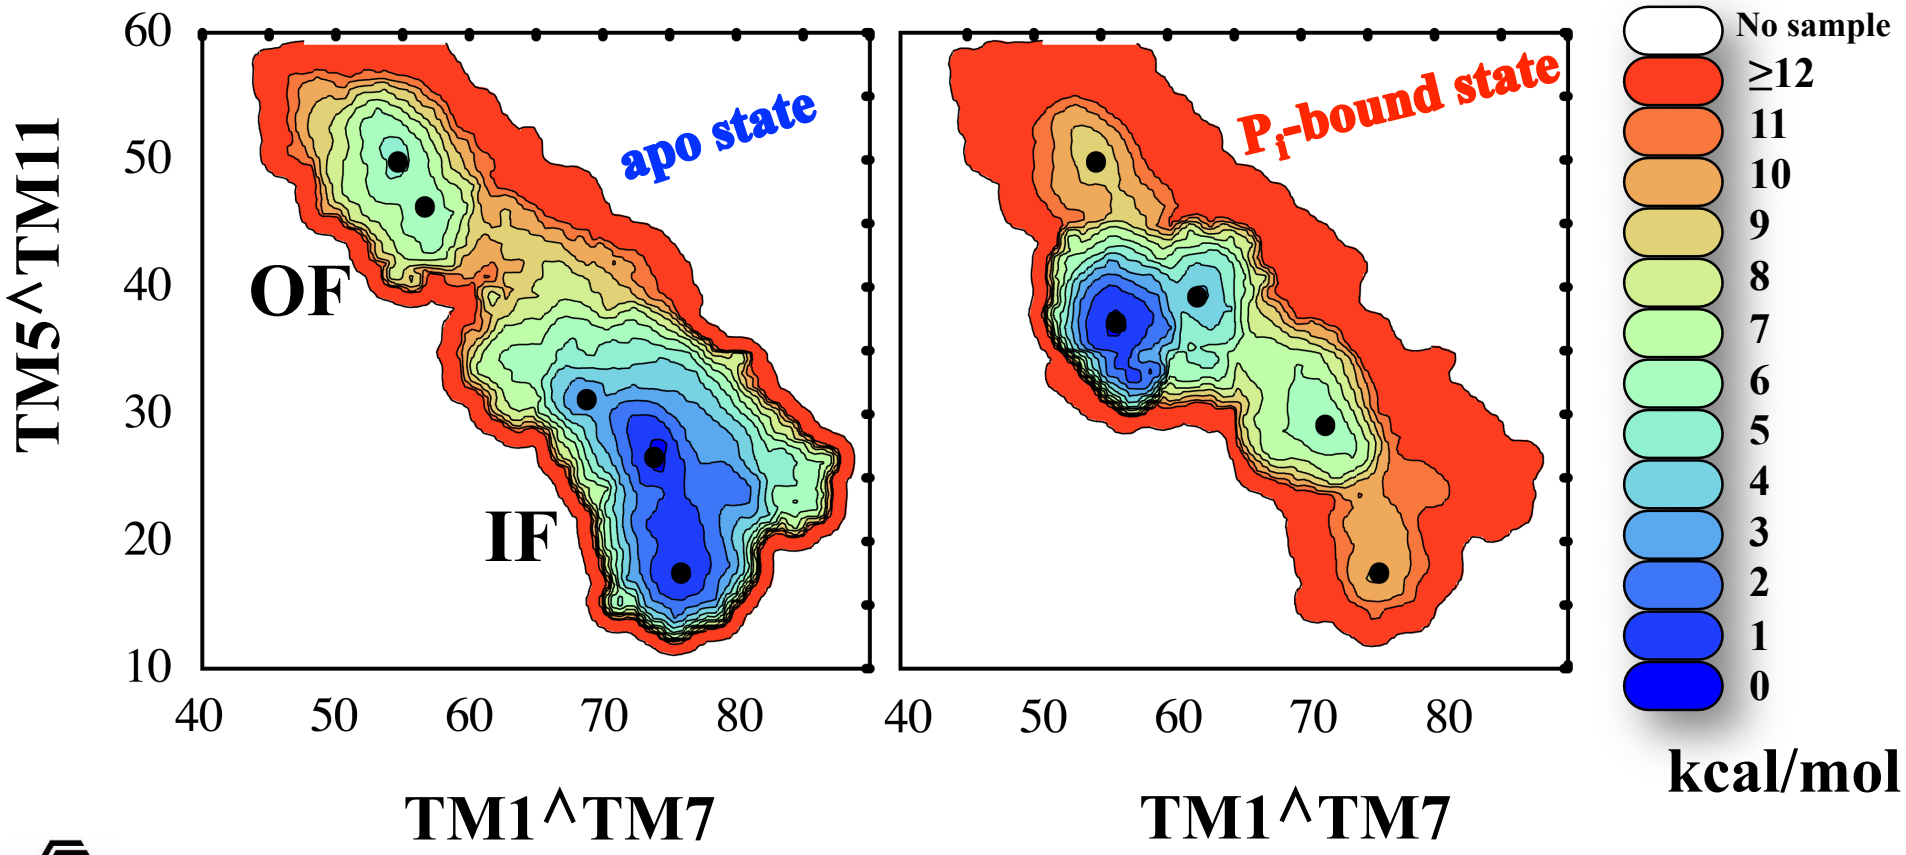

National Center for Research Resources

NIH BTRC for Macromolecular Modeling and Bioinformatics
http://www.ks.uiuc.edu/

Beckman Institute, UIUC

GlpT Least Free Energy Path

Method of Ensing *et al.*, *JPCB* 109 6676 (2005).

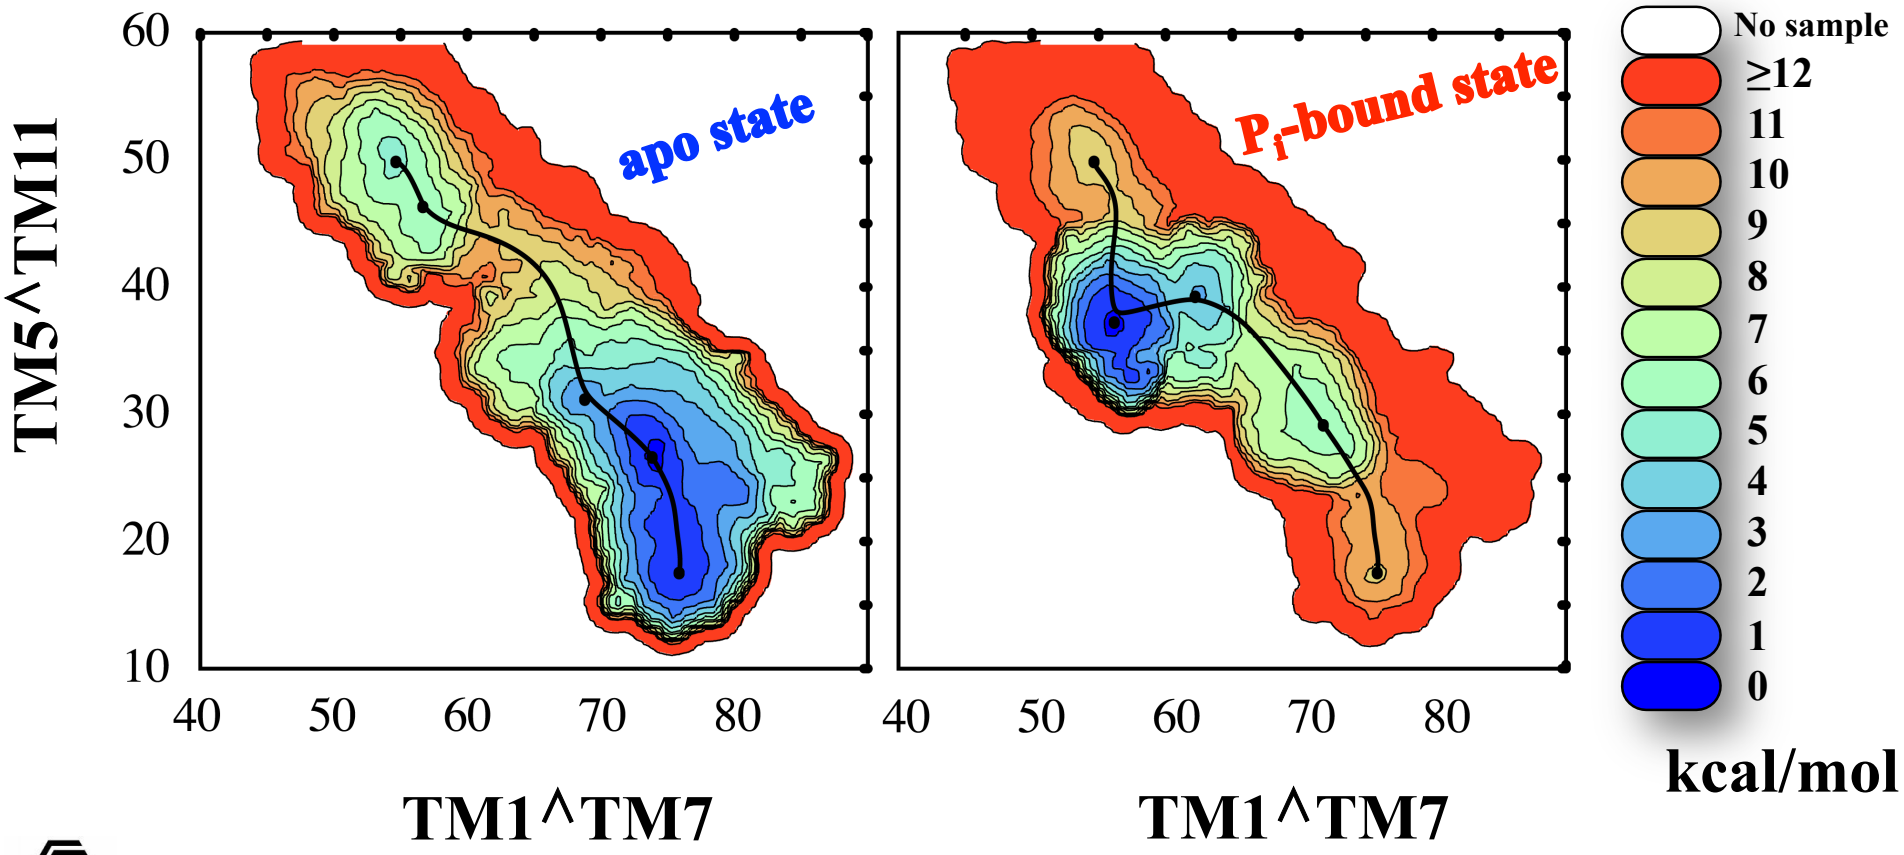

National Center for Research Resources

Free Energy Along GlpT Least Free Energy Path

It turns out the P_i substrate lowers the barrier and shifts the global minimum from the IF to the OF state.

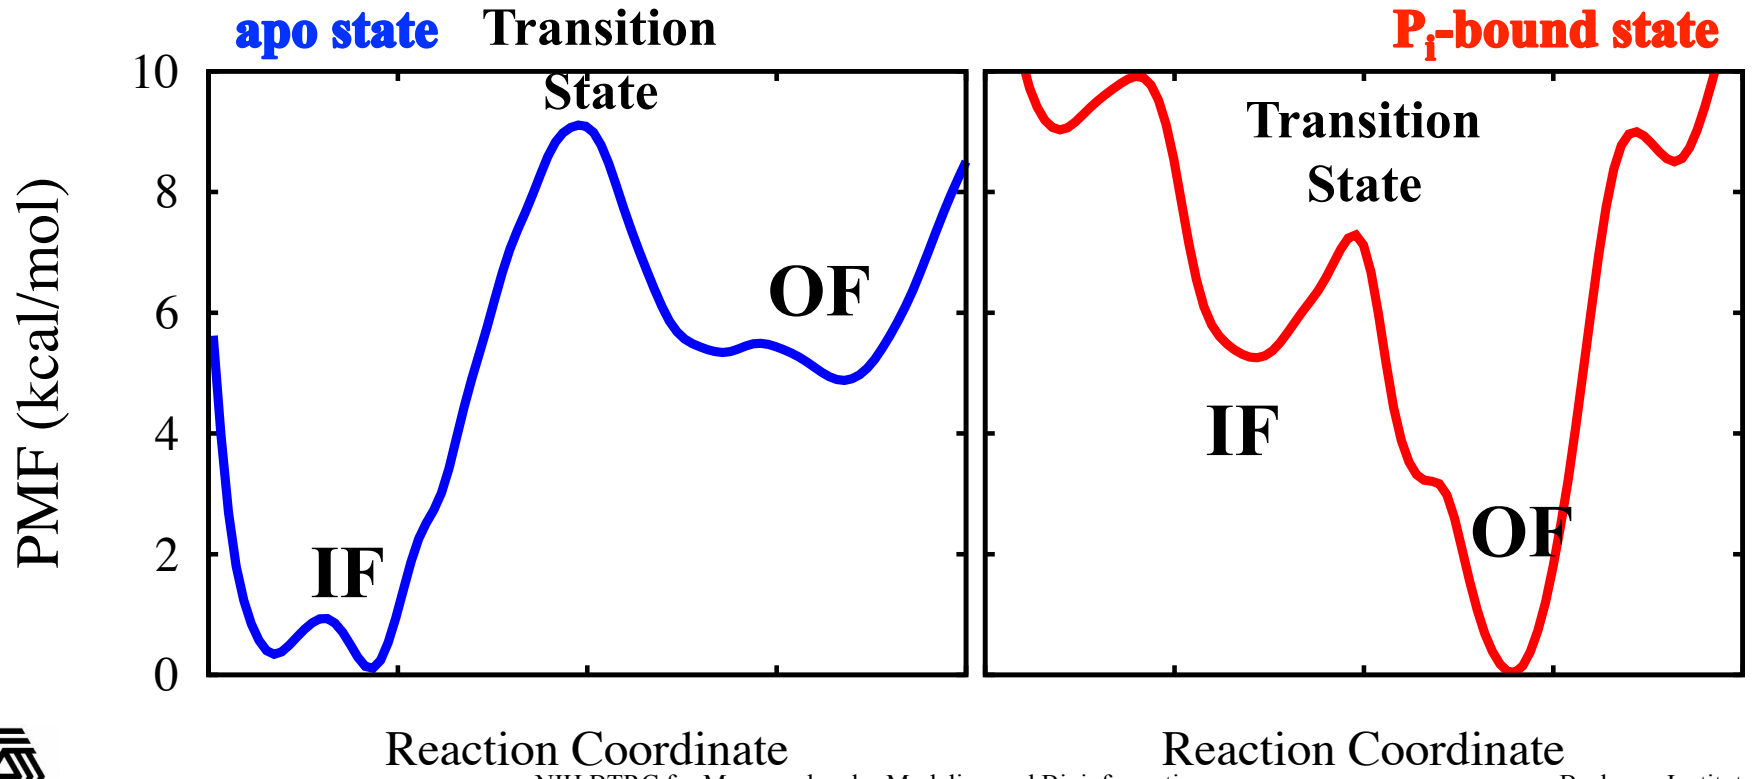

Thanks to: NIH, NSF, DOE,
Nikhil Jain, Yanhua Sun, Gengbin Zheng, Eric Bohm,
Mahmoud Moradi

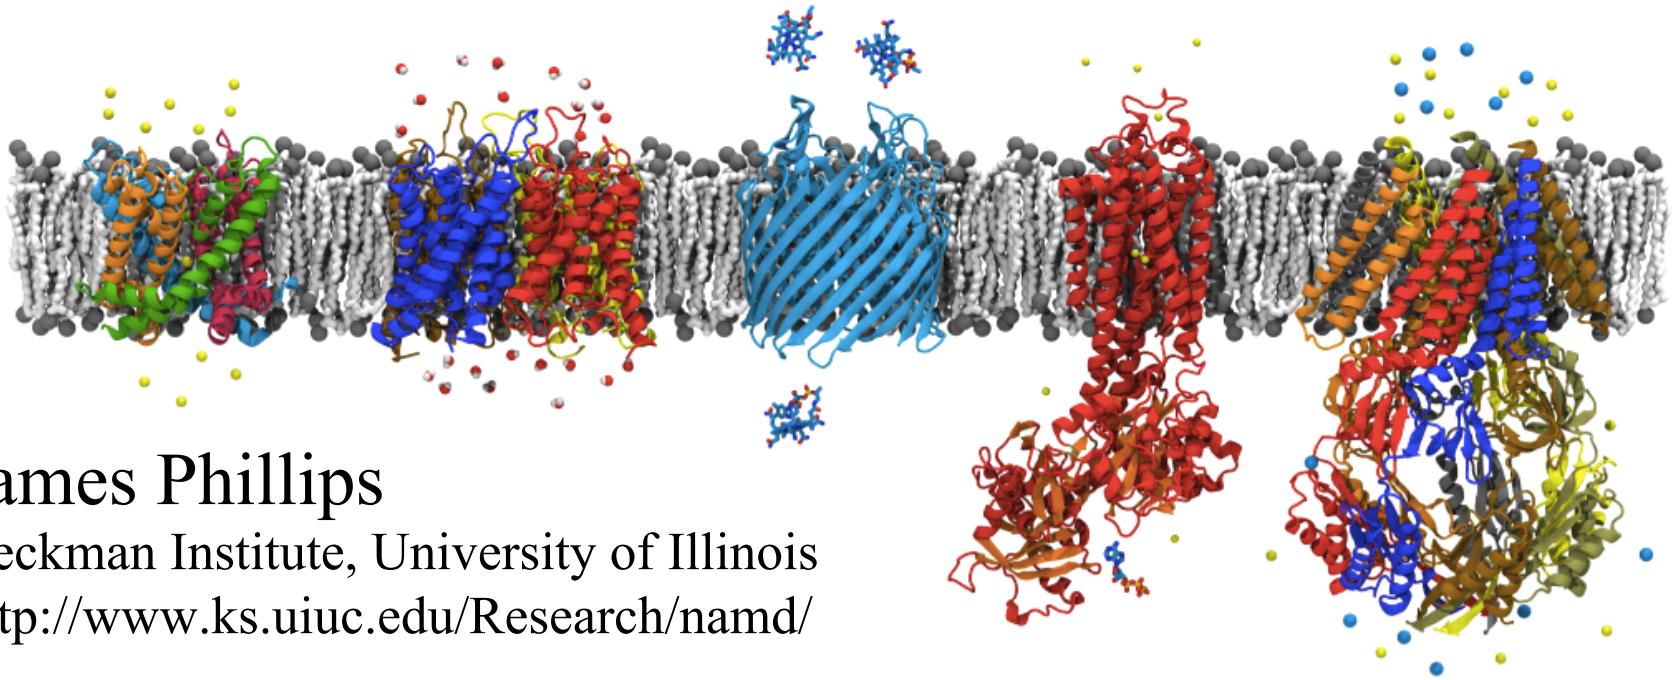

James Phillips
Beckman Institute, University of Illinois
http://www.ks.uiuc.edu/Research/namd/