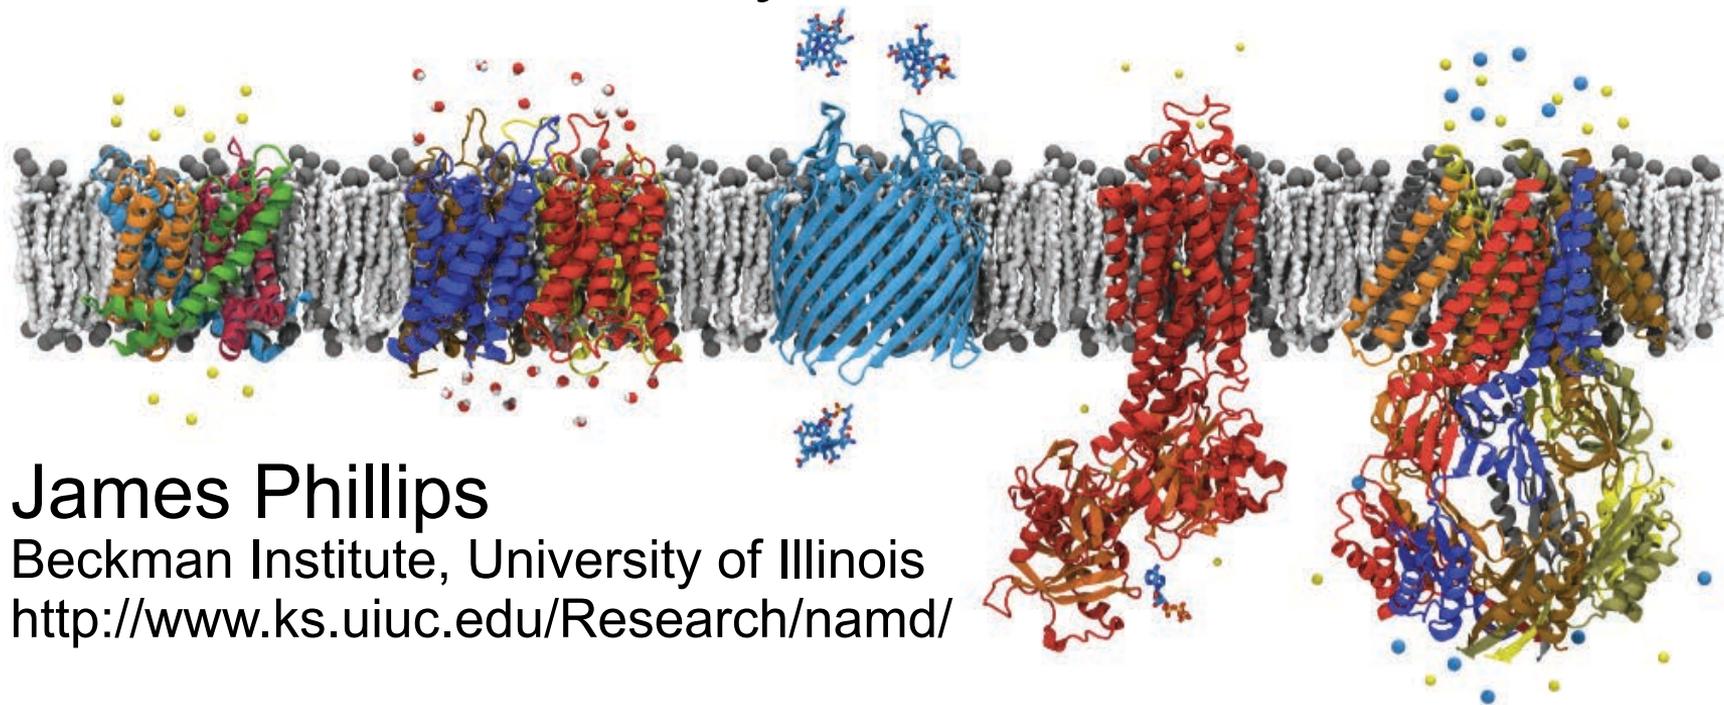


# Plenary: Real Life Experience 3 - NAMD

## *DOE HPCOR on Scientific Software Architecture for Portability and Performance*



James Phillips

Beckman Institute, University of Illinois

<http://www.ks.uiuc.edu/Research/namd/>

# NIH Biomedical Technology Research Center for Macromolecular Modeling and Bioinformatics

Developers of the widely used computational biology software VMD and NAMD

290,000 registered VMD users  
72,000 registered NAMD users

600 publications (since 1972)  
over 54,000 citations

5 faculty members  
8 developers  
1 systems  
administrator  
17 postdocs  
46 graduate students  
3 administrative staff

*Renewed 2012-2017  
with 10.0 score (NIH)*

research projects include: virus  
capsids, ribosome, photosynthesis,  
protein folding, membrane reshaping,  
animal magnetoreception

## Achievements Built on People

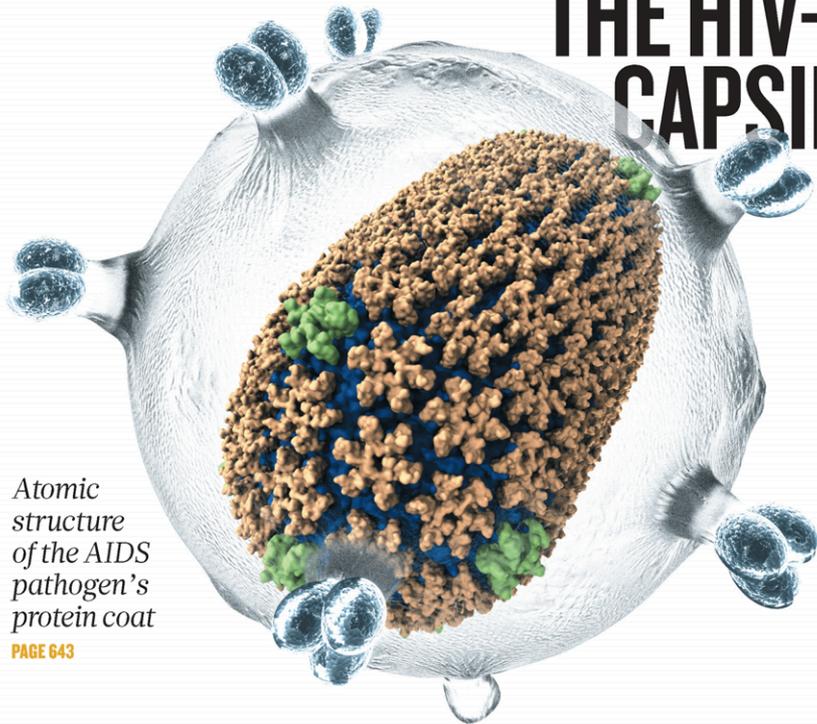


Tajkorshid, Luthey-Schulten, Stone, Schulten, Phillips, Kale, Mallon

# nature

THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE

## THE HIV-1 CAPSID



Atomic  
structure  
of the AIDS  
pathogen's  
protein coat

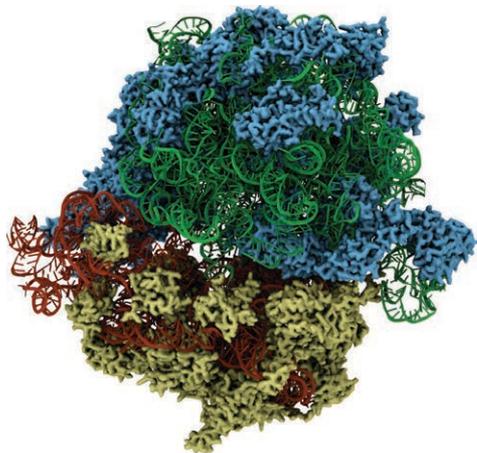
PAGE 643

## 2013 *HPCwire* Editors' Choice Award for Best Use of HPC in Life Sciences



# Other Petascale Projects Using NAMD

From cellular machines  
to the pharmacy...



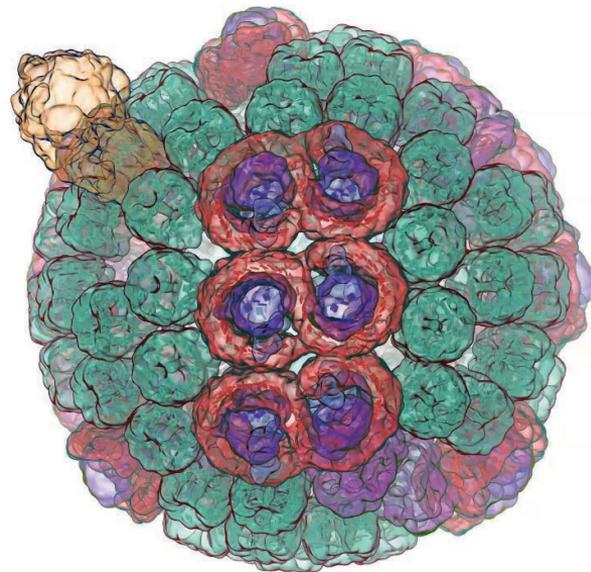
ribosome  
3 M atoms, multiple copies

From woodchips to gasoline...



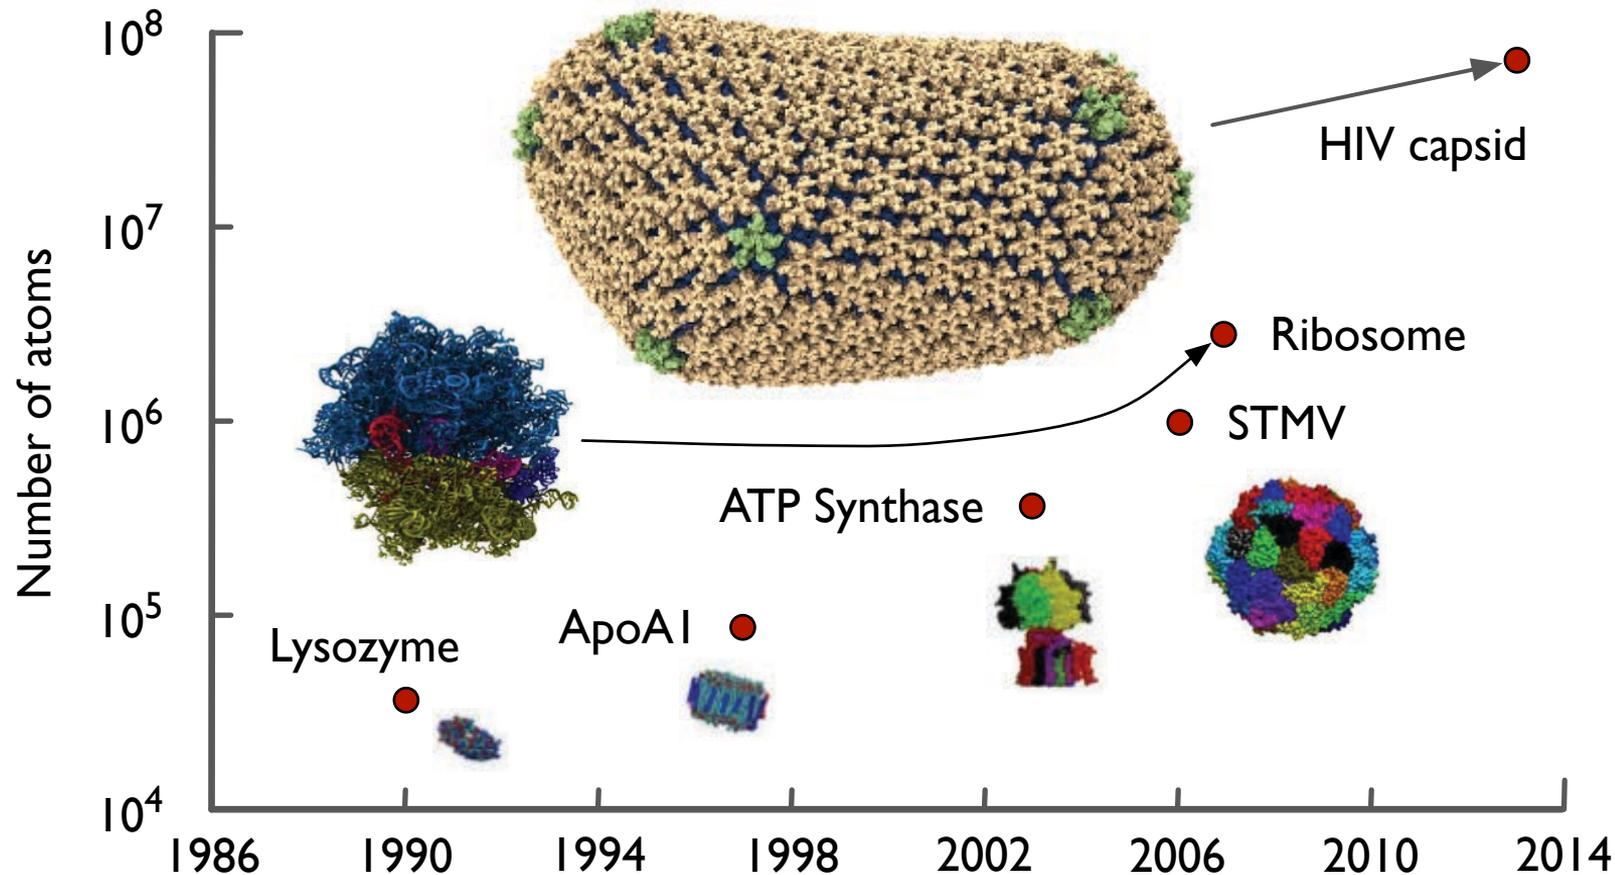
second-generation biofuels  
> 10 M atoms

From solar energy to cellular fuel...



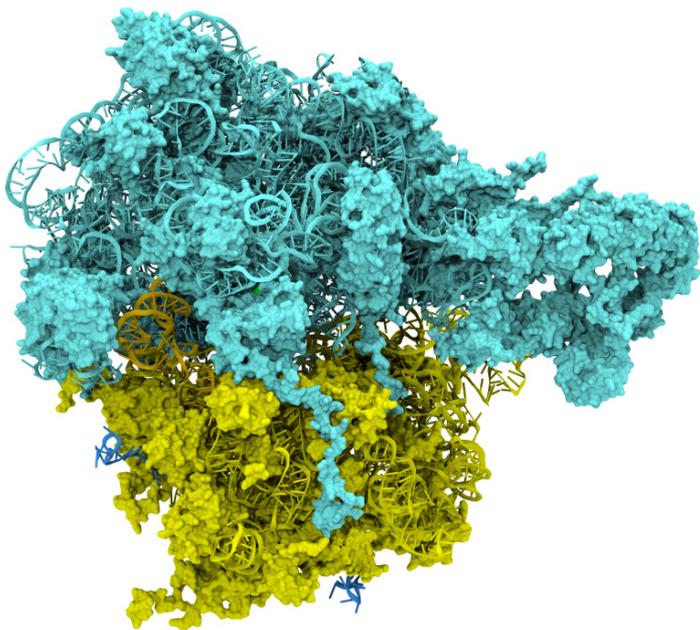
photosynthetic  
chromatophore  
100 M atoms

# A brief history of NAMD (and VMD)

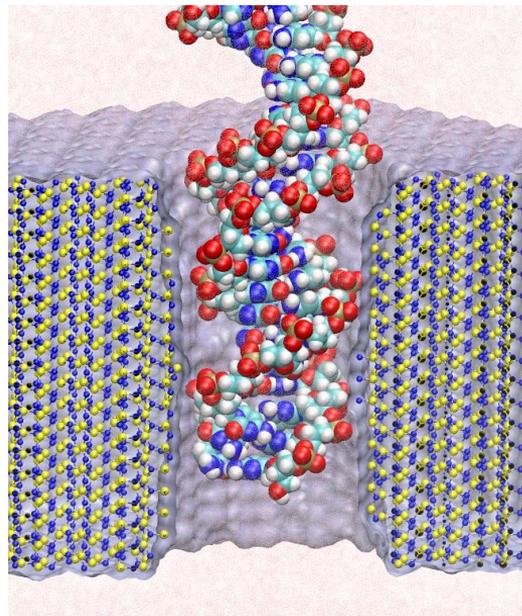


# Computational Microscopy

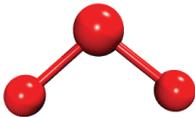
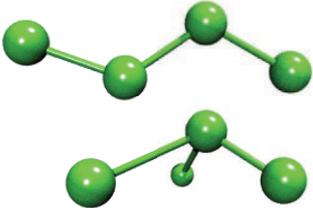
Ribosome: synthesizes proteins from genetic information, target for antibiotics



Silicon nanopore: bionanodevice for sequencing DNA efficiently



# Molecular Mechanics Force Field

$$\begin{aligned}
 U(\vec{R}) = & \underbrace{\sum_{\text{bonds}} k_i^{\text{bond}} (r_i - r_0)^2}_{U_{\text{bond}}} + \underbrace{\sum_{\text{angles}} k_i^{\text{angle}} (\theta_i - \theta_0)^2}_{U_{\text{angle}}} + \\
 & \underbrace{\sum_{\text{dihedrals}} k_i^{\text{dihe}} [1 + \cos(n_i \phi_i + \delta_i)]}_{U_{\text{dihedral}}} + \\
 & \underbrace{\sum_i \sum_{j \neq i} 4\epsilon_{ij} \left[ \left( \frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left( \frac{\sigma_{ij}}{r_{ij}} \right)^6 \right]}_{U_{\text{nonbond}}} + \sum_i \sum_{j \neq i} \frac{q_i q_j}{\epsilon r_{ij}}
 \end{aligned}$$

# Classical Molecular Dynamics

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

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.

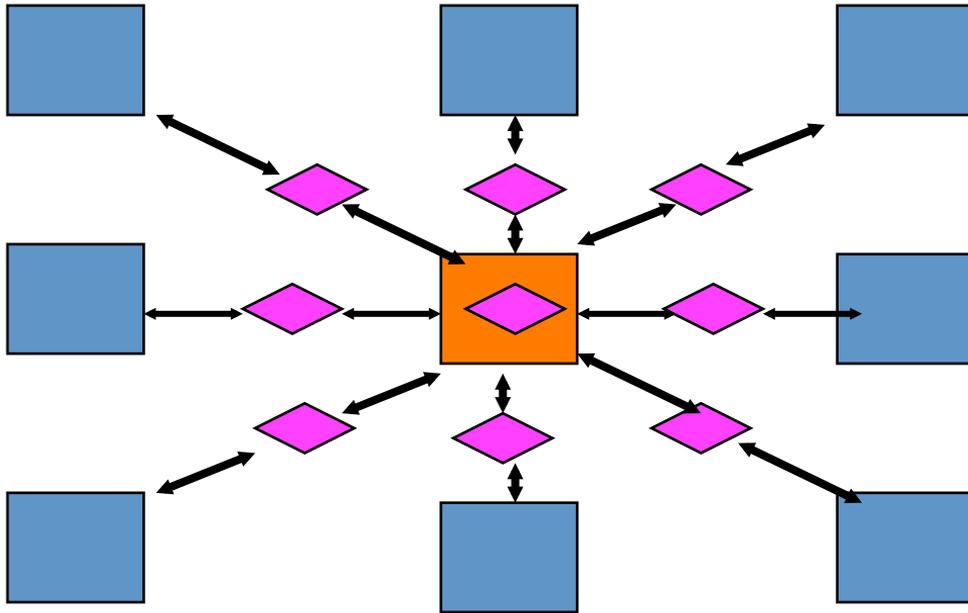
# Long-term Charm++ Collaboration

- Illinois Parallel Programming Lab
  - Prof. Laxmikant Kale
  - [charm.cs.illinois.edu](http://charm.cs.illinois.edu)
- Long standing collaboration
  - Since start of Center in 1992
  - Gordon Bell award at SC2002
  - Joint Fernbach award at SC12
- Synergistic research
  - NAMD requirements drive and validate CS work
  - Charm++ software provides unique capabilities
  - Enhances NAMD performance in many ways



# 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.

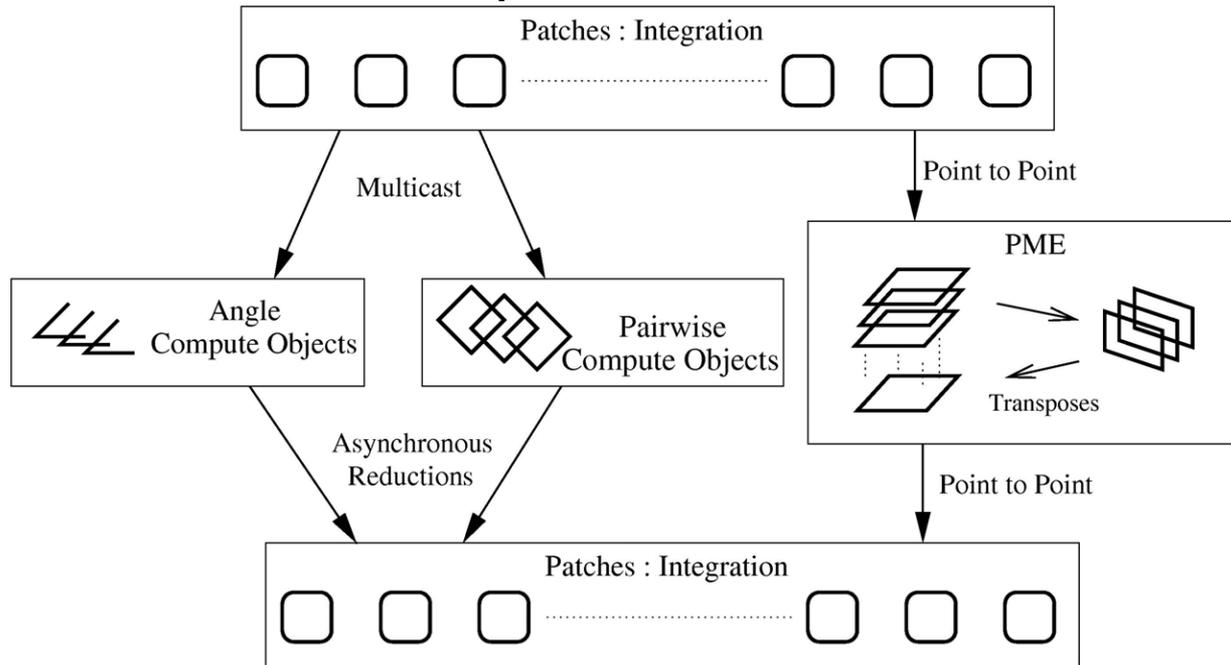
# Implementation in 1997 Charm++

- 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.



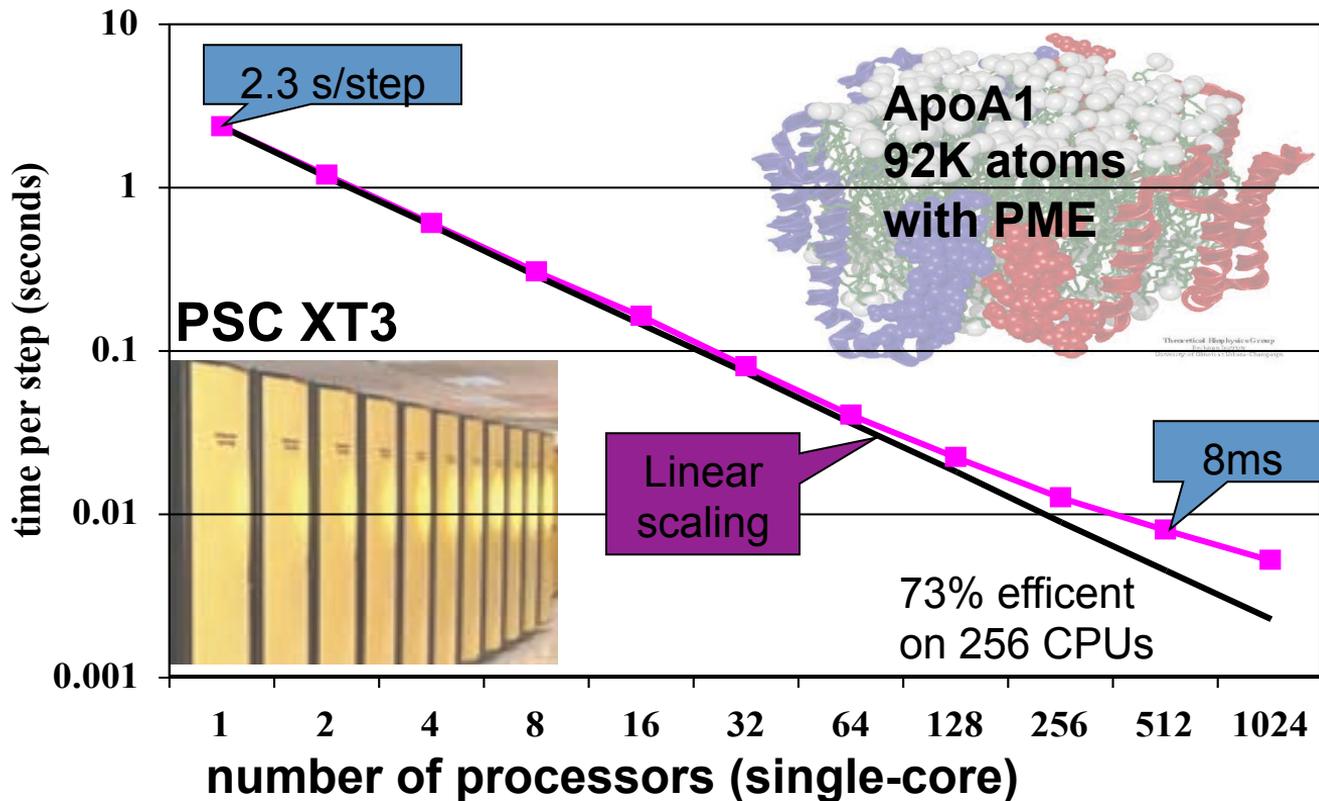
# NAMD Overlapping Execution

Phillips *et al.*, SC2002.



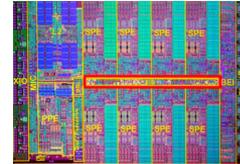
Objects are assigned to processors and queued as data arrives.

# 2006 NAMD Performance



# Early Acceleration Options

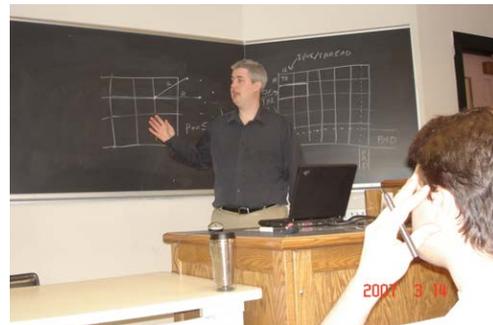
- Outlook in 2005-2006:
  - FPGA reconfigurable computing (with NCSA)
    - Difficult to program, slow floating point, expensive
  - Cell processor (NCSA hardware)
    - Relatively easy to program, expensive
  - ClearSpeed (direct contact with company)
    - Limited memory and memory bandwidth, expensive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



# CUDA: Practical Performance

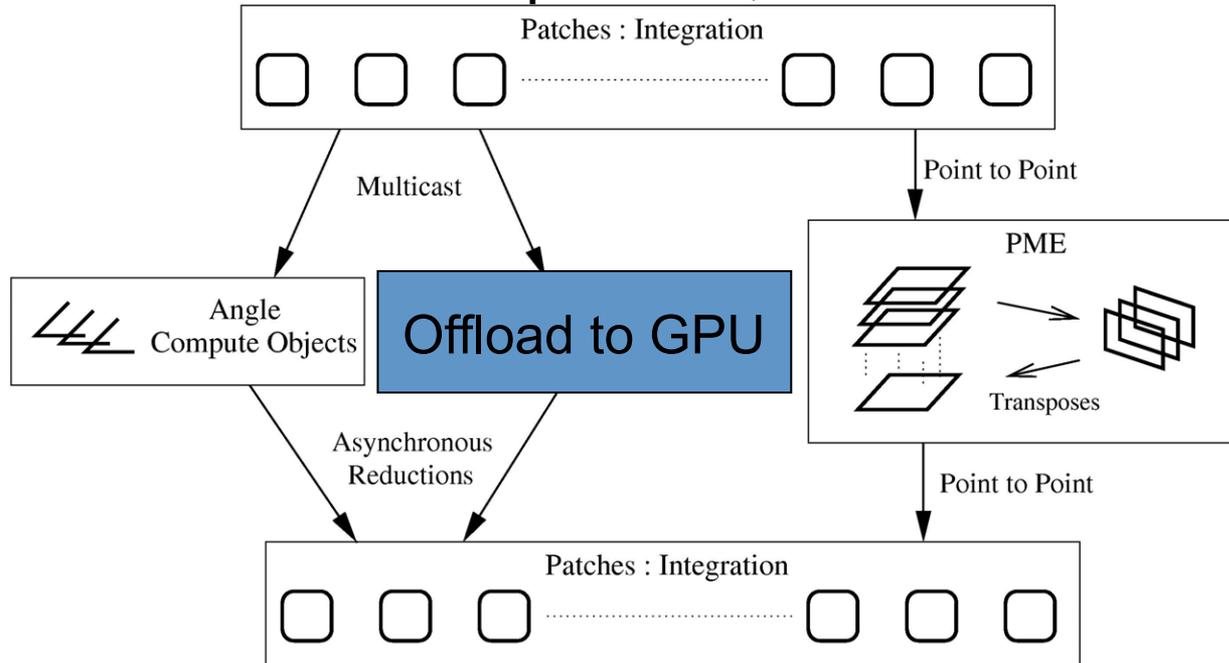
*November 2006: NVIDIA announces CUDA for G80 GPU.*

- 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)



# NAMD Overlapping Execution

Phillips *et al.*, SC2002.



Objects are assigned to processors and queued as data arrives.

# 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)

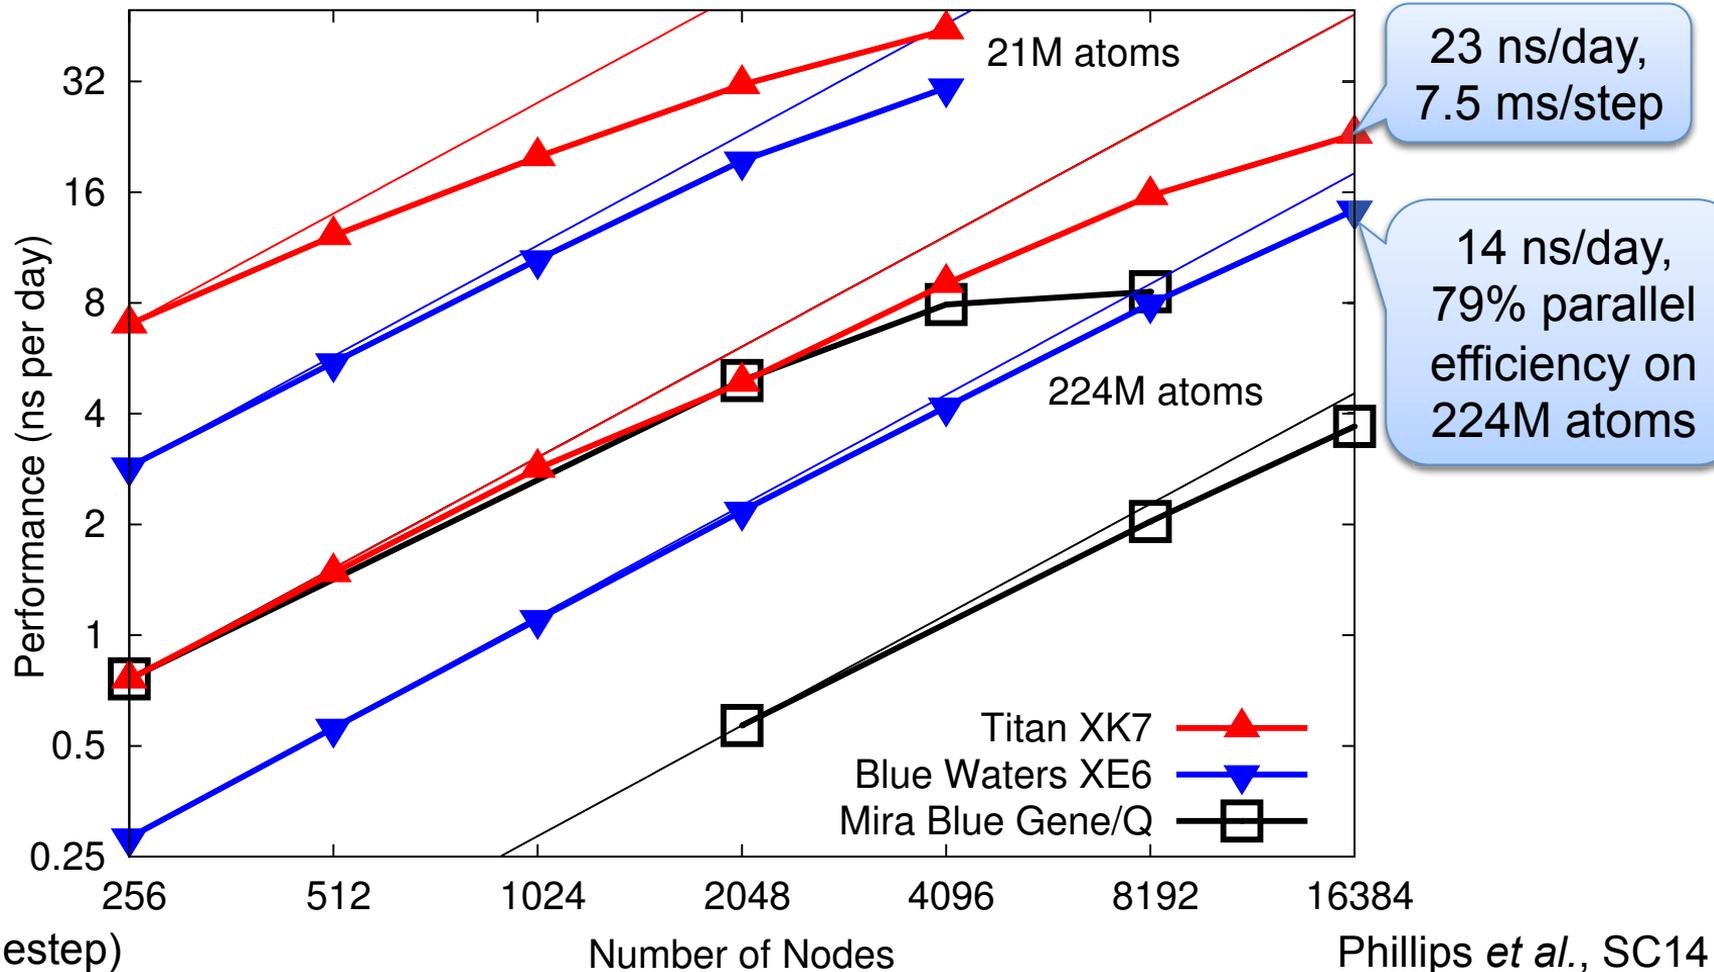


# NAMD for Large Systems

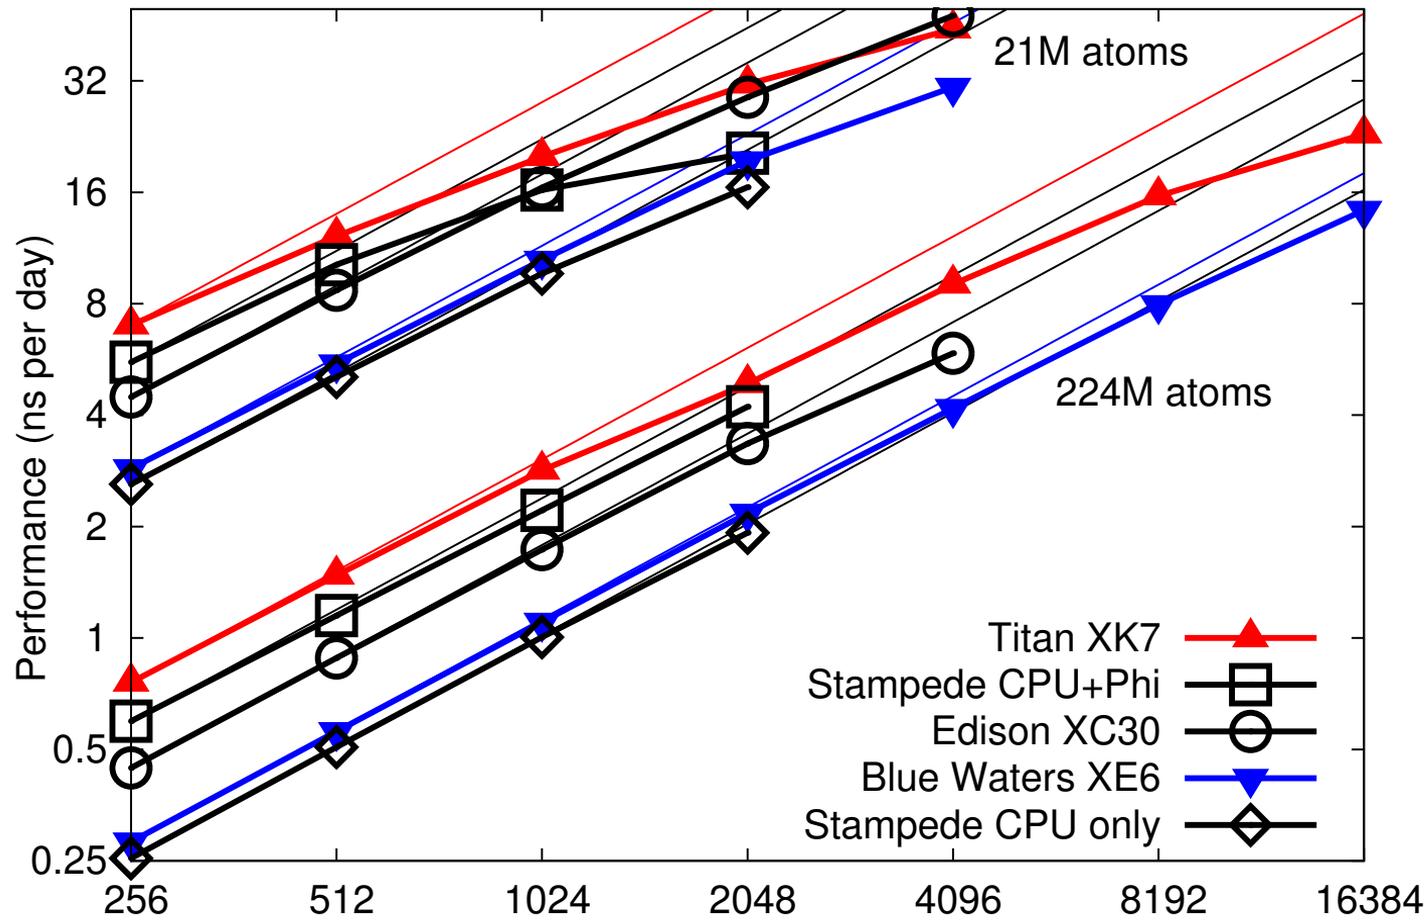
- Per-node memory usage
  - Exploit redundant structure
  - Pre-compressed static data
  - Distributed per-atom data
  - Special “memopt” build
  - Not all features supported
  - NAMD-only file formats
    - May change between versions
    - No VMD reader/writer - ever
- I/O performance
  - Data is relatively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



# NAMD on Petascale Platforms



# NAMD on Torus and Non-torus Networks

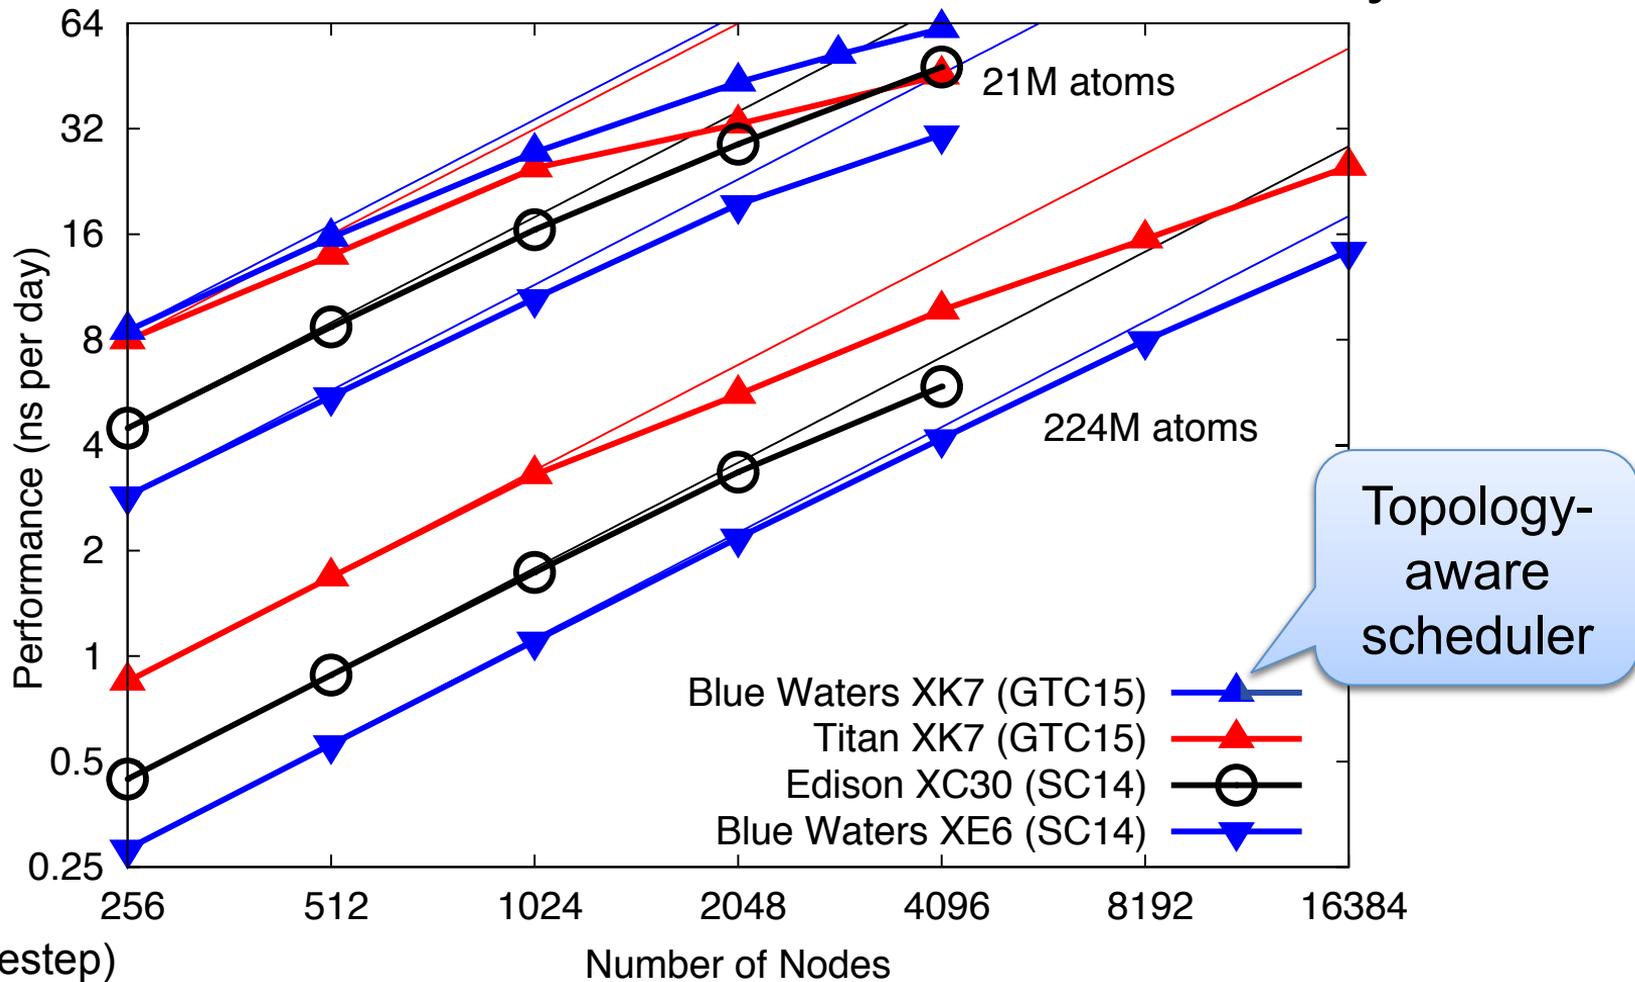


(2fs timestep)

Number of Nodes

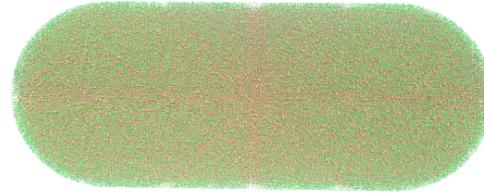
Phillips *et al.*, SC14

# NAMD on Petascale Platforms Today

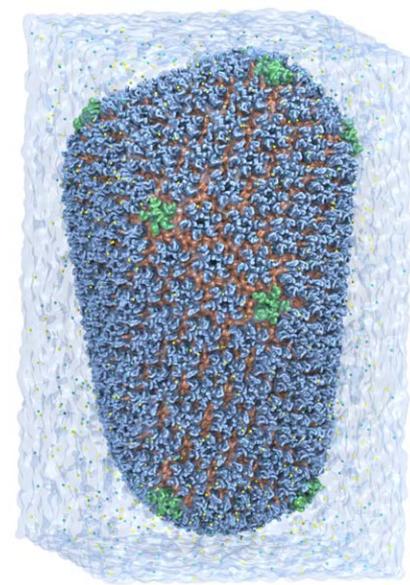


# VMD – “Visual Molecular Dynamics”

- Visualization and analysis of:
  - molecular dynamics simulations
  - particle systems and whole cells
  - cryoEM densities, volumetric data
  - quantum chemistry calculations
  - sequence information
- User extensible w/ scripting, plugins
- <http://www.ks.uiuc.edu/Research/vmd/>



Whole Cell Simulation



MD Simulations

Structural Similarity	
1hrc-a	GDVVERGKIKIFVQKCAQGC
1ccr-a	GNPKAGEKIKYKTKCAQGC
1yaa-a	GSAAKKAATLTKTRCQGC
5cya-a	GDVAKGKKTIVQKCAQGC
1cyc-a	GDVAKGKKTIVQKCAQGC
1hrb-a	GDVVEGKHLFTICITC

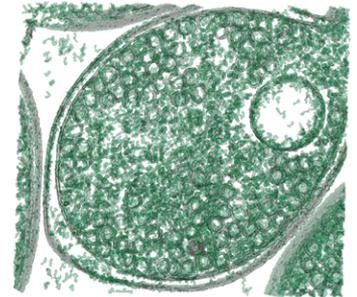
  

Sequence Similarity	
1hrc-a	GDVVERKIKIKIFVQKCAQGC
1ccr-a	GNPKAETIKIKTKCAQGC
1yaa-a	GSAAKKAATLTKTRCQGC
5cya-a	GDVAKGKKTIVQKCAQGC
1cyc-a	GDVAKGKKTIVQKCAQGC

Sequence Data



Quantum Chemistry



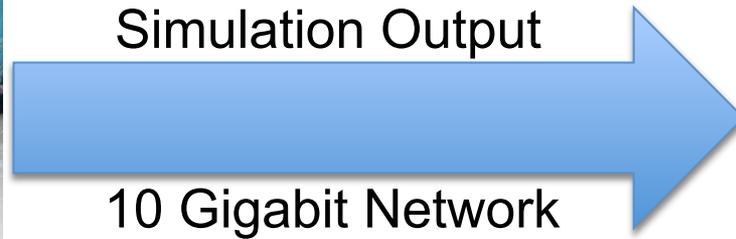
CryoEM, Cellular Tomography

# NIH Center Facilities Enable Petascale Biology

Over the past several years the Center has assembled all necessary hardware and infrastructure to prepare and analyze petascale molecular dynamics simulations, and ***makes these facilities available to visiting researchers.***



External Resources,  
90% of our  
Computer Power



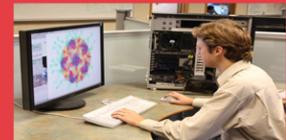
Petascale Gateway Facility

Storage



Compute

Visualization



High-End Workstations  
Accessible to Visitors

# Virtual Facilities Enable Petascale Anywhere



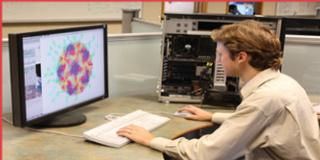
High-end visualization and analysis workstations currently available only in person at the Beckman Institute must be *virtualized and embedded at supercomputer centers*.

Storage



Compute

Visualization



Compressed Video



1 Gigabit Network



# NAMD Serves NIH Users and Goals

## *Practical Supercomputing for Biomedical Research*

- 72,000 users can't all be computer experts.
  - 18% are NIH-funded; many in other countries.
  - 20,000 have downloaded more than one version.
  - 4000 citations of NAMD reference papers.
- One program available on all platforms.
  - Desktops and laptops – setup and testing
  - Linux clusters – affordable local workhorses
  - Supercomputers – free allocations on XSEDE
  - Blue Waters – sustained petaflop/s performance
  - GPUs - next-generation supercomputing
- User knowledge is preserved across platforms.
  - No change in input or output files.
  - Run any simulation on **any number of cores.**
- Available free of charge to all.



Hands-On Workshops



Oak Ridge TITAN

# 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



# Software Engineering

- Charm++:
  - Git revision control
  - Redmine issue tracking
  - Gerrit code review
  - Nightly build & test
- NAMD:
  - CVS revision control
  - Manual tracking & review
  - Nightly build & release
  - Automated build & install
- Configuration Management
  - “Latest” builds for users
  - Older builds preserved
  - Modules in use logged
  - Charm++ ships with NAMD
- Verification and Validation
  - New tests for new features
  - Library of historical tests
  - User reports of crashes
  - Internal checksums



# Every Run is a Test

- Internal consistency checks are critical
  - Relatively few users, fewer developers
  - Variety of use cases and feature sets – runs are unique
  - Configuration changes are rampant on supercomputers
  - Testing at scale is expensive – leverage production runs
  - Crashes are annoying but harmless
  - Goal is to avoid generating bad science
  - Users do not read warning messages!
    - They barely read “FATAL ERROR” messages on stderr.
    - Exit on cases that risk incorrect results, link to web page.



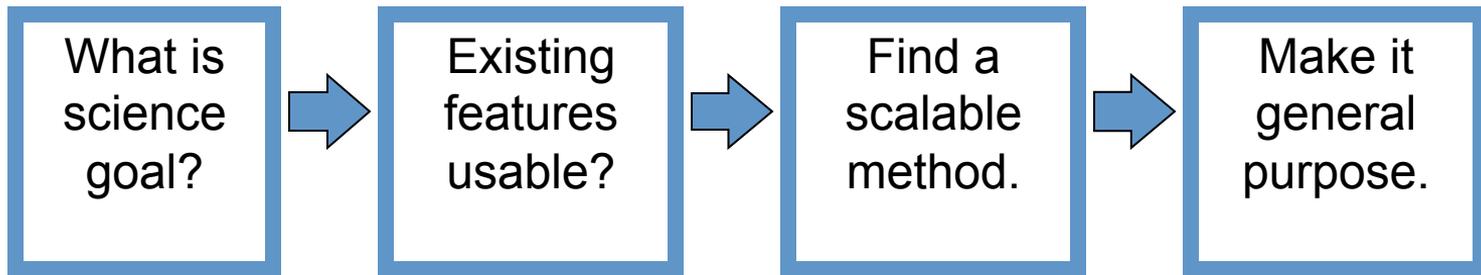
# Development Process/Philosophy

- 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



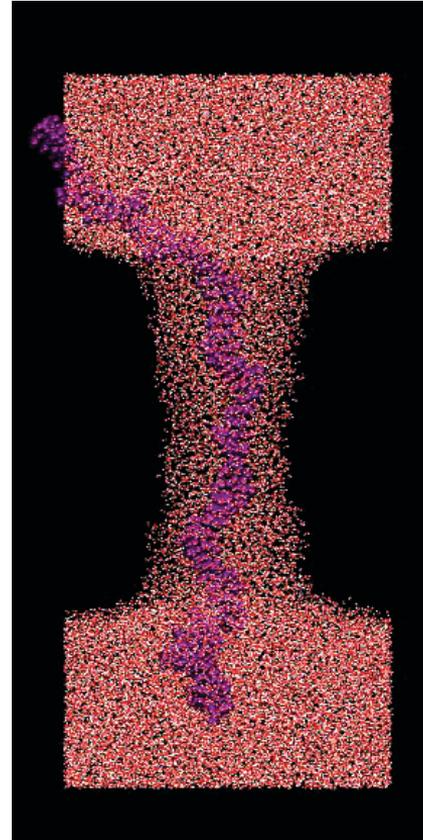
# Collaborative Driving Projects

- 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

- 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 optimizations for efficiency
  - Shrinking phantom pore for DNA



# Why NAMD and VMD Use Tcl

- History: Programs are ~20 years old.
- Maturity: Package management, portable.
- Stability: Interfaces haven't changed.
- Flexibility: Encapsulates mini-languages.
- Approachability: Looks like a simple scripting language, doesn't scare non-programmers.



# Looking Forward

- NERSC Cori / Argonne Theta (2016)
  - Knight's Landing (KNL) Xeon Phi
  - Single-socket nodes, Cray Aries network
  - Theta Early Science Project:
    - “Free Energy Landscapes of Membrane Transport Proteins”
- Oak Ridge Summit (2018)
  - IBM Power 9 CPUs + NVIDIA Volta GPUs
  - 3,400 fat nodes, dual-rail InfiniBand network
  - CAAR Project “Molecular Machinery of the Brain”
- Argonne Aurora (2018)
  - Knight's Hill (KNH) Xeon Phi



# Portability Requires Trade-Offs



# Aspects of Portability

- Operating systems
  - Linux, Mac, Windows
  - Lustre filesystem errors
  - System library and OS bugs
- CPU architecture
  - Compiler directives (e.g., ivdep)
  - OpenMP 4.0 “#pragma omp simd”
  - Occasional vector intrinsics
- Networks
  - Infiniband, Gemini, BG/Q
  - Offload to MPI/Charm++
  - Do not use Charm++ on MPI
  - Charm++ relatively fast to port
- Coprocessors
  - CUDA is mature and best in class
  - OpenCL isn't performance-portable
  - OpenACC supported by Intel?
  - Intel offload directives only for MIC



# Offload Challenges

- Writing CUDA kernels is acceptable
  - Vendor is fully engaged
  - Knowledge is widespread
  - Tools are mature *and* up-to-date
- Offload aggregation is the challenge
  - Charm++ has multi-threaded control on CPU
    - Independent threads that share a memory space
  - GPU works best with unified stream of work
- Work remaining on CPU will become bottleneck
  - Can stream results off GPU to enable overlap
  - Need to optimize, use all CPU cores/threads available



# Charm++ and MIC Options

- Today: Aggregated Offload
  - Keep NAMD work decomposition
  - Collect and bulk-copy data
  - Bulk-launch tasks in single offload
  - Method initially developed for CUDA
  - NAMD MIC offload is clone of CUDA offload



# Charm++ and MIC Options

- KNL Option: OpenMP Thread Teams
  - Grainsize too large for single MIC thread
  - Grainsize too small for entire MIC
  - Let Charm++ control OpenMP thread teams
    - E.g, MIC = 15 Charm++ PEs
    - Each PE = 4 cores and 16 threads
    - Parallelize loops using OpenMP directives



# Code Modernization

- Re-vectorization? Post-modernization?
- NAMD is currently array-of-structures
  - Aligned to cache lines for random access
- Emerging idiom for MD kernels:
  - Load atoms in neighbor list from AoS
  - Transpose in register for vectorization
  - Gather-scatter may provide similar performance



# MIC Vectorization Options

- Intrinsic – what we have
  - Written by David Kunzman of Intel
  - Currently in production
- Compiler – what we want
  - `#pragma [omp] simd assert`
  - Currently ~20% slower with refactored kernel
    - “One missed optimization” but improving fast
- ISPC – our backup plan
  - Similar to CUDA, consider if compiler fails

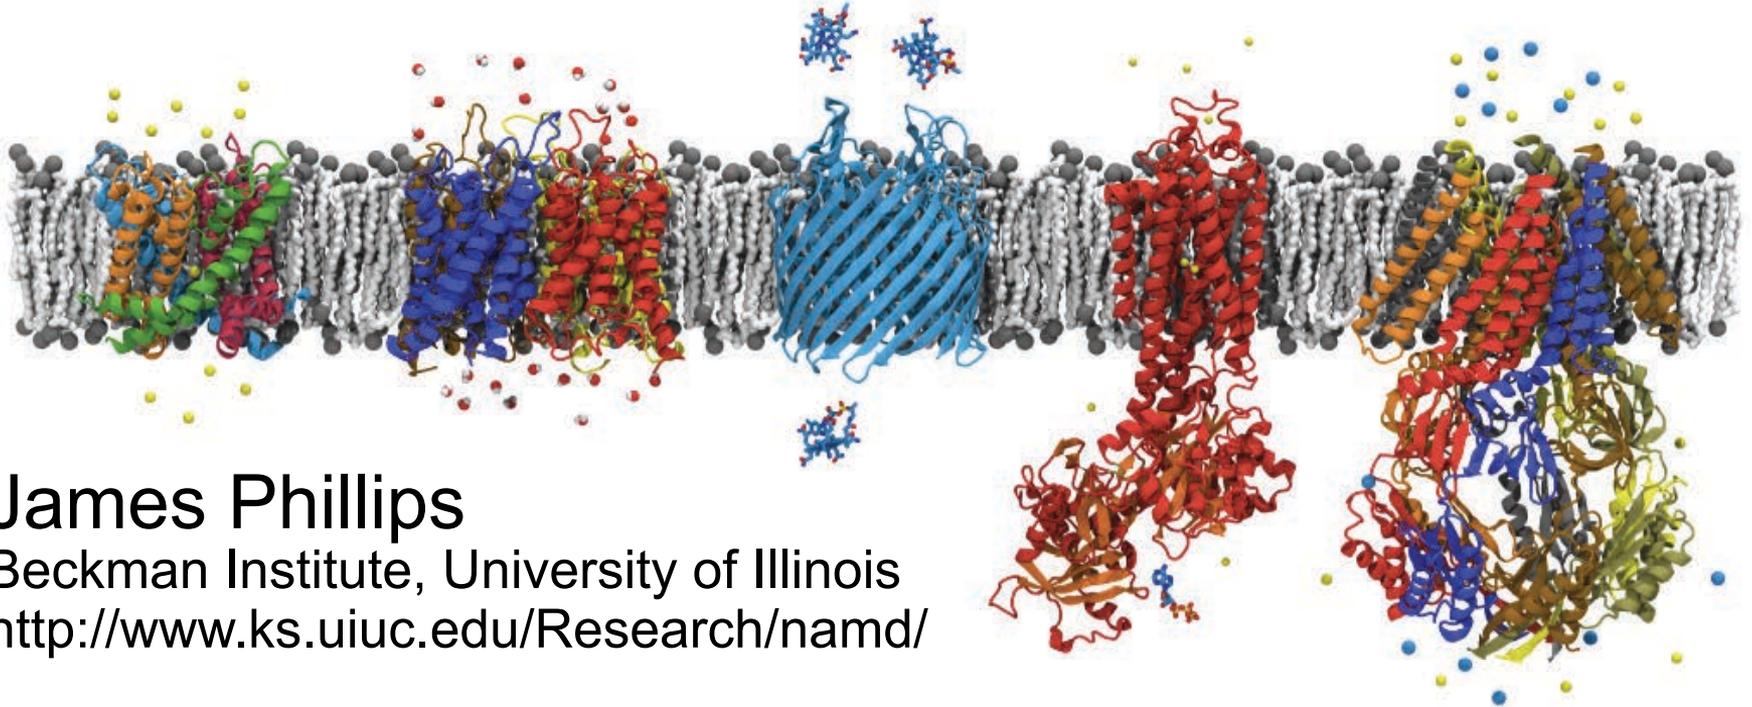


# Conclusions and Ramblings

- Science, software, and supercomputing are all hard.
  - If you get good science from any supercomputer, you are winning.
- Solve problems you have before problems you might have.
  - Performance and correctness on one platform, then portability.
  - Complexity is forever - try the simplest thing that might work.
- Do look ahead - don't paint yourself into a corner.
  - But don't worry about things you don't yet understand well.
  - If you do get stuck, don't be afraid to refactor.
- If a problem has many solutions, it is probably unsolved.
  - But even a limited tool may work for your case.



Thanks to NIH, NSF, DOE, and 20 years of  
NAMD and Charm++ developers and users.



**James Phillips**

Beckman Institute, University of Illinois

<http://www.ks.uiuc.edu/Research/namd/>