NAMD Algorithms and HPC Functionality

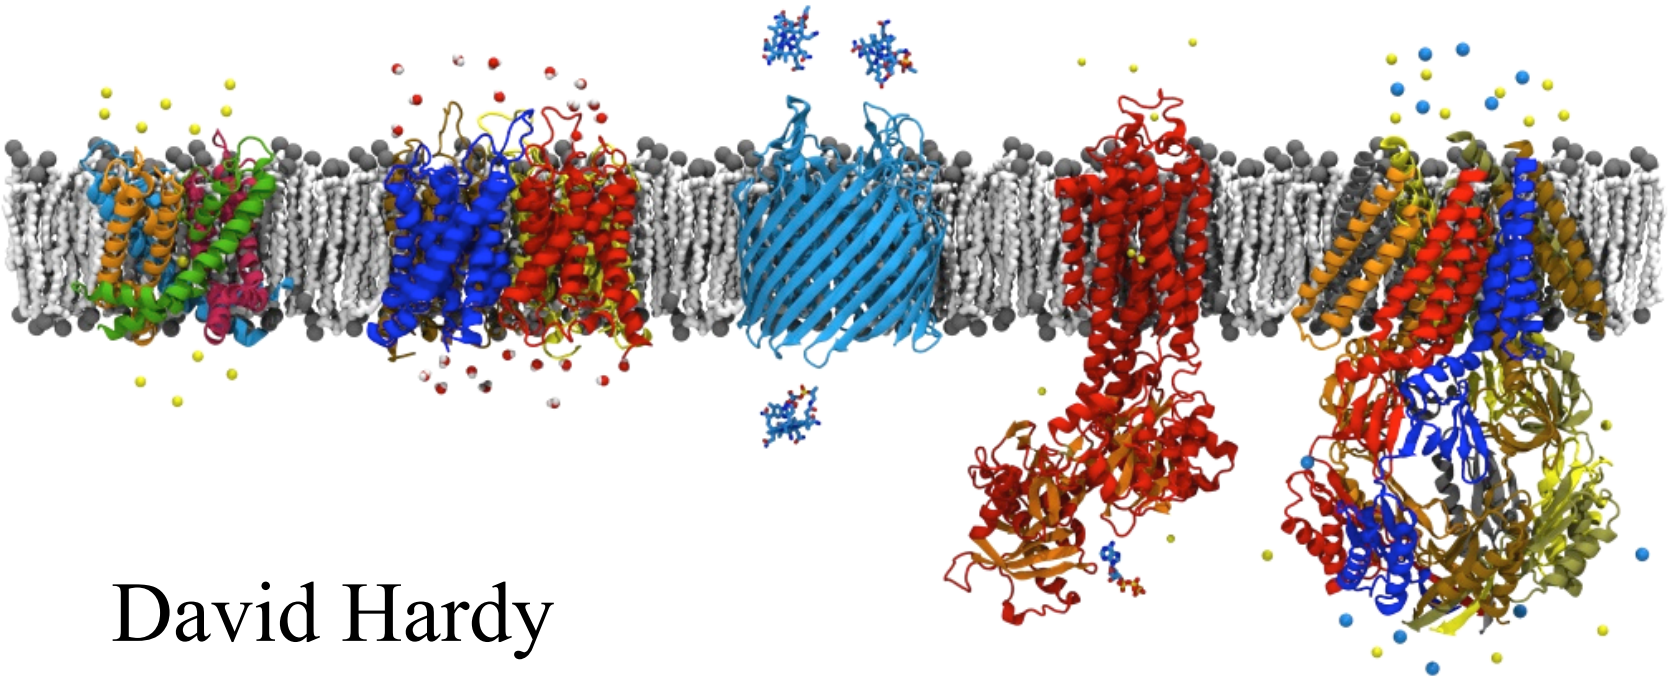

David Hardy

http://www.ks.uiuc.edu/Research/~dhardy/

NAIS: State-of-the-Art Algorithms for Molecular Dynamics

Beckman Institute
University of Illinois at Urbana-Champaign

Theoretical and Computational Biophysics Group

National Center for Research Resources

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

Beckman Institute, UIUC

Acknowledgments

Jim Phillips
Lead NAMD developer

John Stone
Lead VMD developer

David Tanner
Implemented GBIS

Klaus Schulten
Director of TCB group

NAMD and VMD: The Computational Microscope

- Study the molecular machines in living cells

Ribosome: synthesizes proteins from genetic information, target for antibiotics

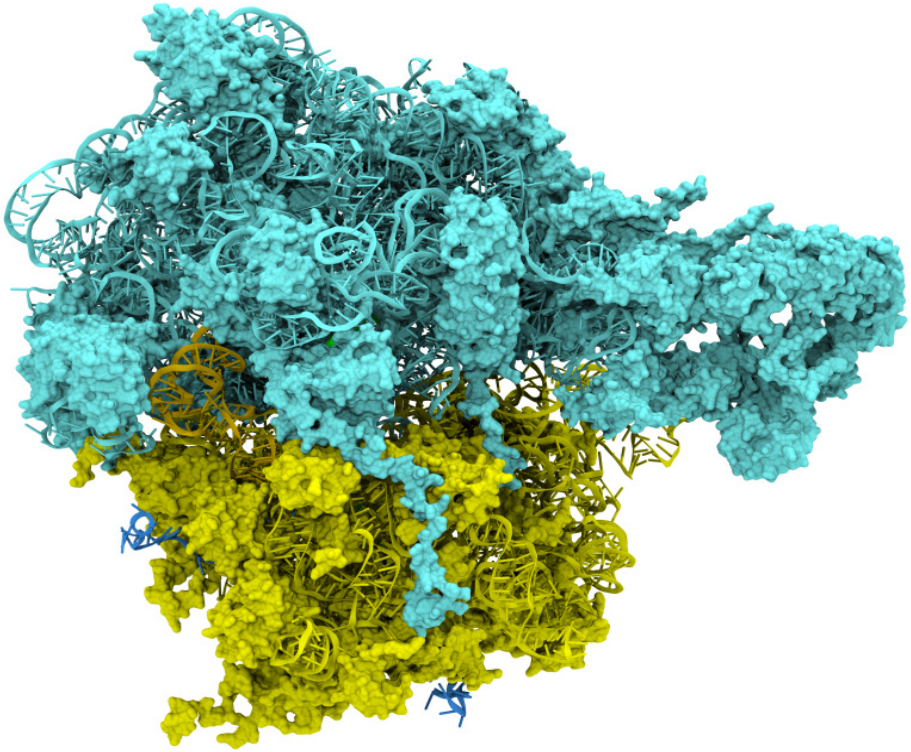

Silicon nanopore: bionanodevice for sequencing DNA efficiently

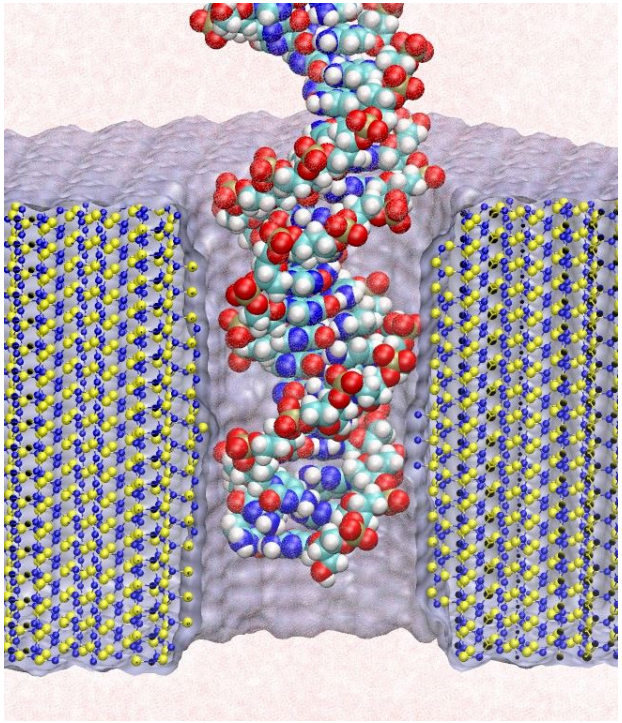

VMD – "Visual Molecular Dynamics"

- Visualization and analysis of:
  - molecular dynamics simulations
  - quantum chemistry calculations
  - particle systems and whole cells
  - sequence data
  - volumetric data
- User extensible w/ scripting and plugins

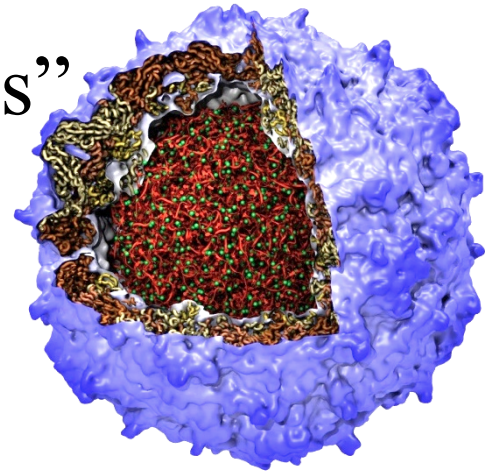

Poliovirus

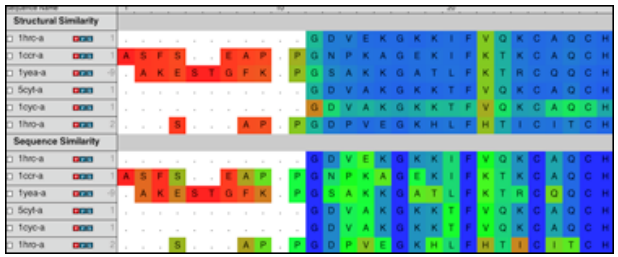

Ribosome Sequences

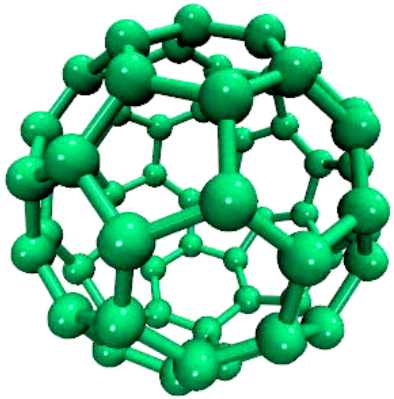

Electrons in
Vibrating Buckyball

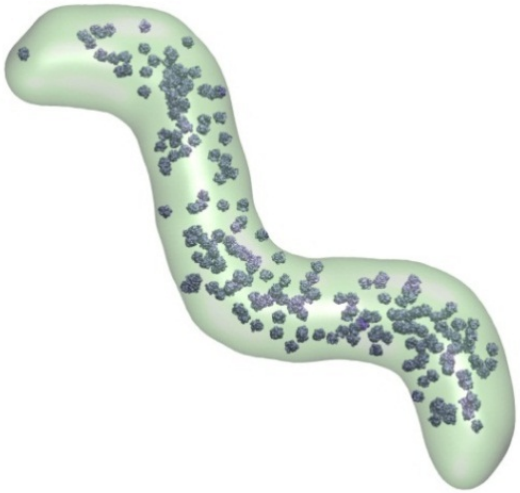

Cellular Tomography,
Cryo-electron Microscopy

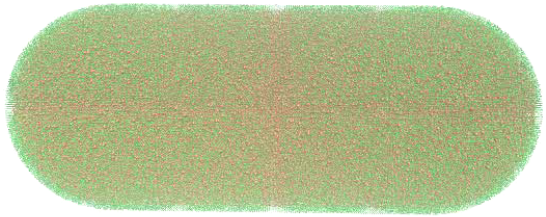

Whole Cell Simulations

VMD Interoperability – Linked to Today's Key Research Areas

- Unique in its interoperability with a broad range of modeling tools: AMBER, CHARMM, CPMD, DL_POLY, GAMESS, GROMACS, HOOMD, LAMMPS, NAMD, and many more …
- Supports key data types, file formats, and databases, e.g. electron microscopy, quantum chemistry, MD trajectories, sequence alignments, super resolution light microscopy
- Incorporates tools for simulation preparation, visualization, and analysis

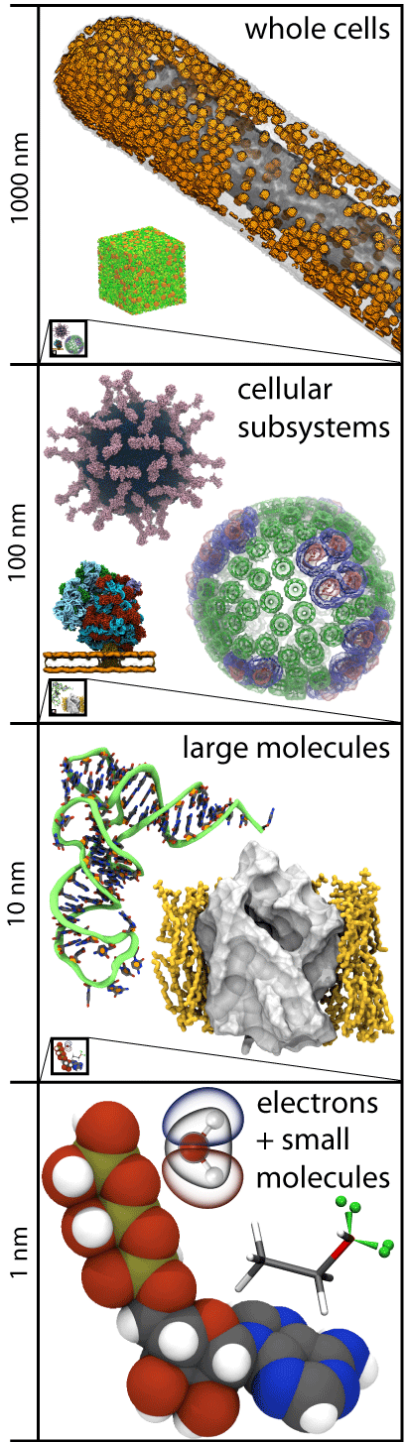

NAMD: Scalable Molecular Dynamics

2002 Gordon Bell Award

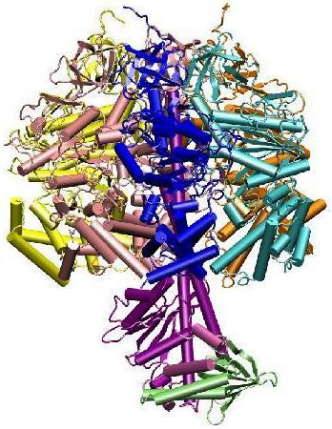

ATP synthase

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

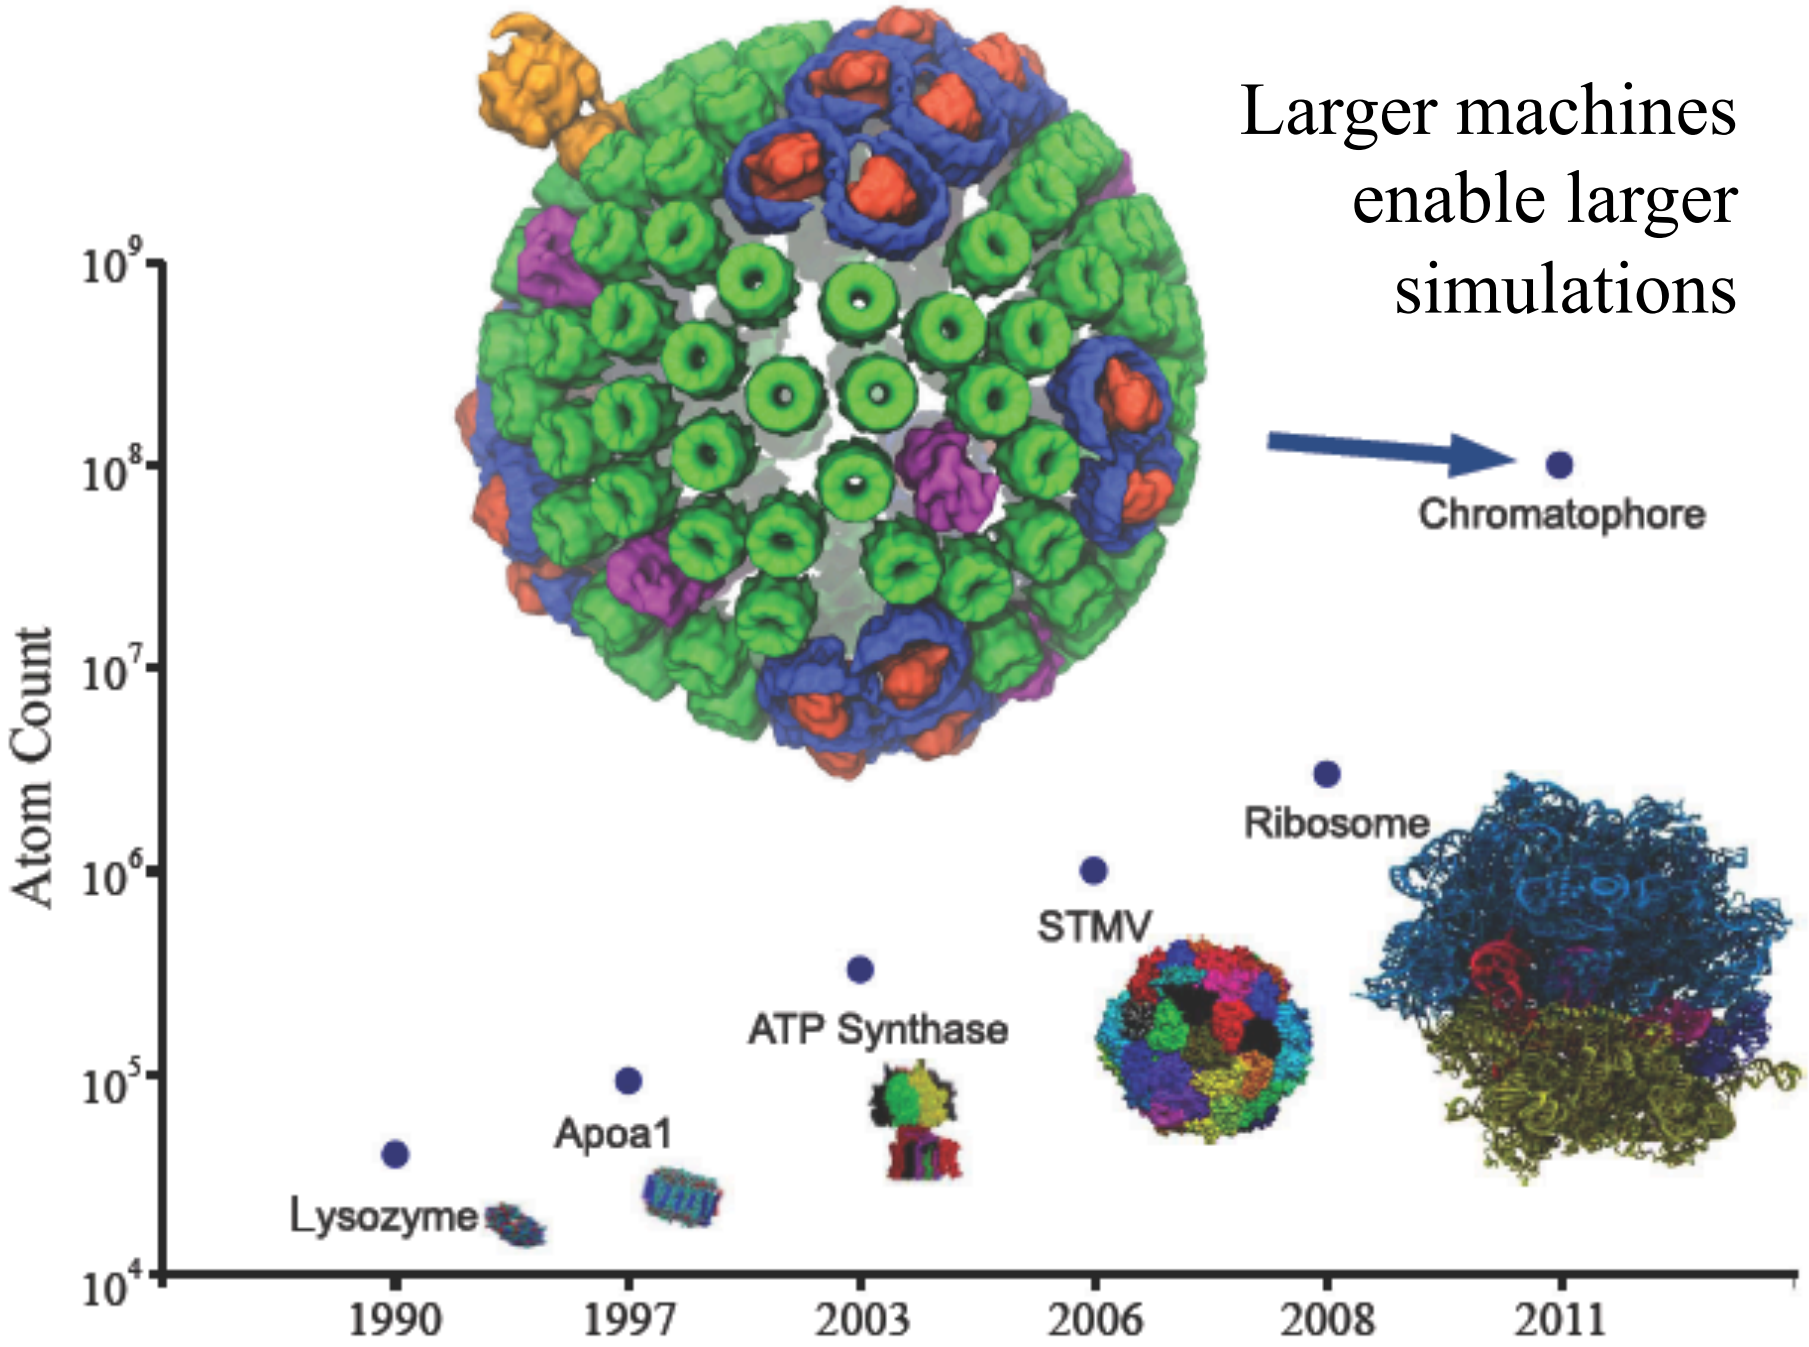
Larger machines enable larger simulations
Atom Count
10^9
10^8
10^7
10^6
10^5
10^4
1990
1997
2003
2006
2008
2011
Chromatophore
Ribosome
STMV
ATP Synthase
Apoa1
Lysozyme

NAMD features are chosen for scalability

- CHARMM, AMBER, OPLS force fields
- Multiple time stepping
- Hydrogen bond constraints
- Efficient PME full electrostatics
- Conjugate-gradient minimization
- Temperature and pressure controls
- Steered molecular dynamics (many methods)
- Interactive molecular dynamics (with VMD)
- Locally enhanced sampling
- Alchemical free energy perturbation
- Adaptive biasing force potential of mean force
- User-extendable in Tcl for forces and algorithms
- All features run in parallel and scale to millions of atoms!

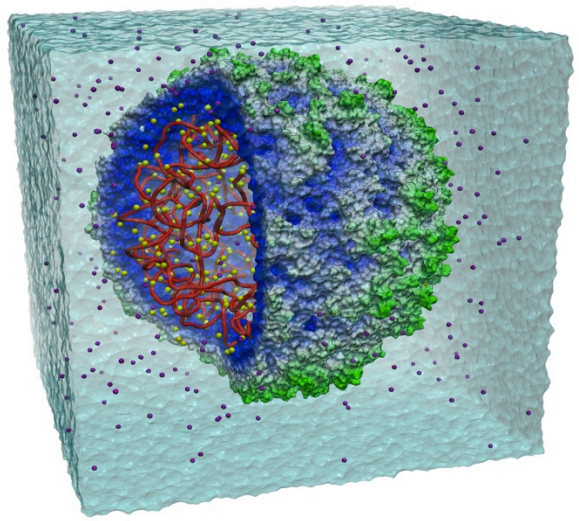

NAMD 2.9 Release

- Public beta released March 19, final version in May
- Capabilities:
  - **New scalable replica-exchange implementation**
  - **QM/MM interface to OpenAtom plane-wave QM code**
  - Knowledge-based Go potentials to drive folding and assembly
  - Multilevel Summation Method electrostatics (serial prototype)
- Performance:
  - Cray XE6/XK6 native multi-threaded network layer
  - Communication optimizations for wider multicore nodes
  - GPU acceleration of energy minimization
  - GPU-oriented shared-memory optimizations
  - GPU Generalized Born (OBC) implicit solvent
  - Faster grid force calculation for MDFF maps

} Enables **Desktop** **MDFF**

NAMD 2.9 Desktop MDFF

with GPU-Accelerated Implicit Solvent
and CPU-Optimized Cryo-EM Forces

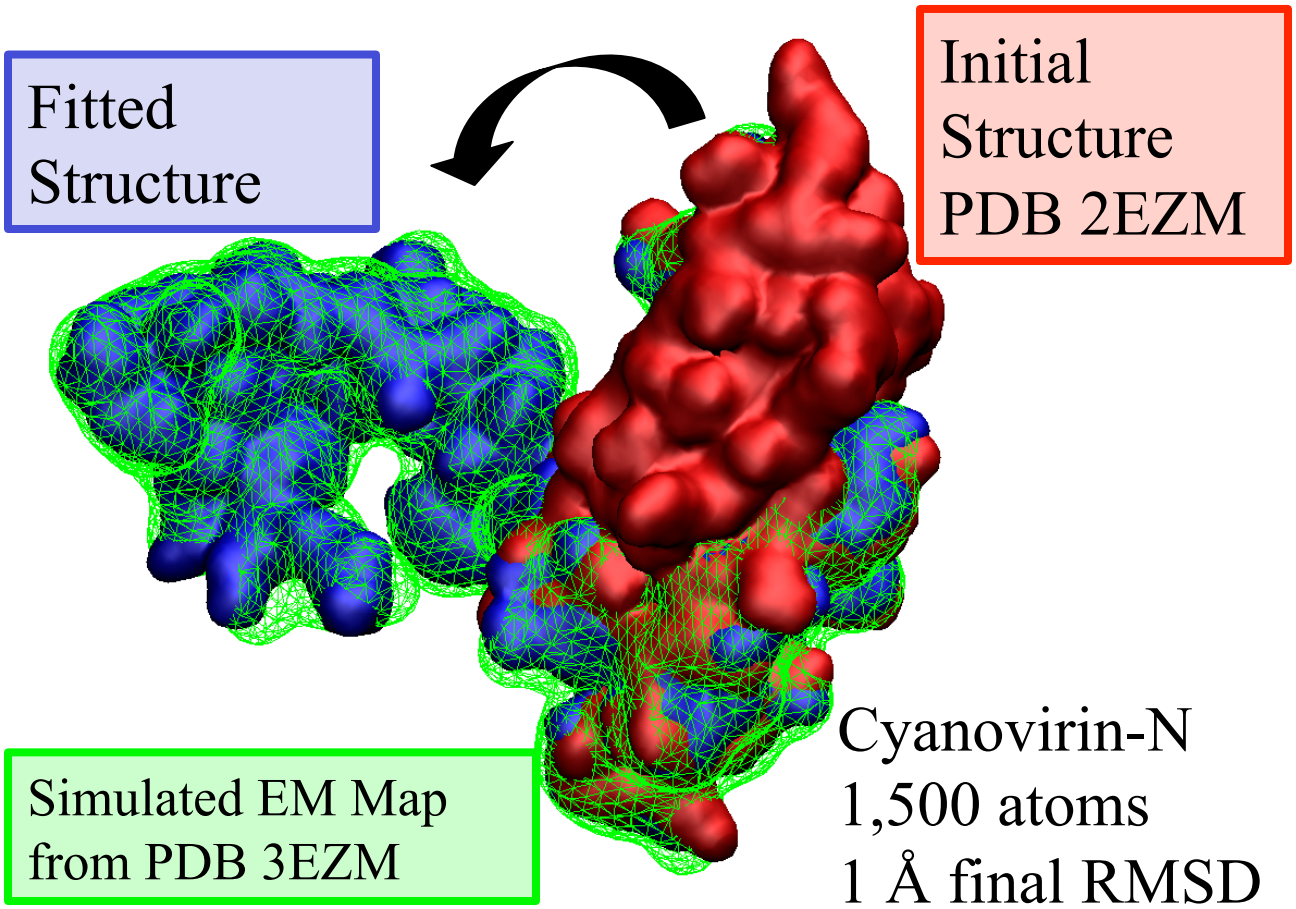

Fast: 2 ns/day

Explicit Solvent
8 cores (1X)

Faster: 12 ns/day

Implicit Solvent
8 cores (6X)

Fastest: 40 ns/day

Implicit Solvent
1 GPU (20X)

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

NAMD 2.9 QM/MM Calculations

Car-Parrinello MD (OpenAtom) and NAMD in one software

OpenAtom *(100 atoms, 70Ry, on 1K cores)*: 120 ms / step
NAMD: *(50,000 atoms on 512 cores)*: 2.5 ms / step

Permits 1000+ atom QM regions

Synchronous load-balancing of QM and MD maximizes processor utilization

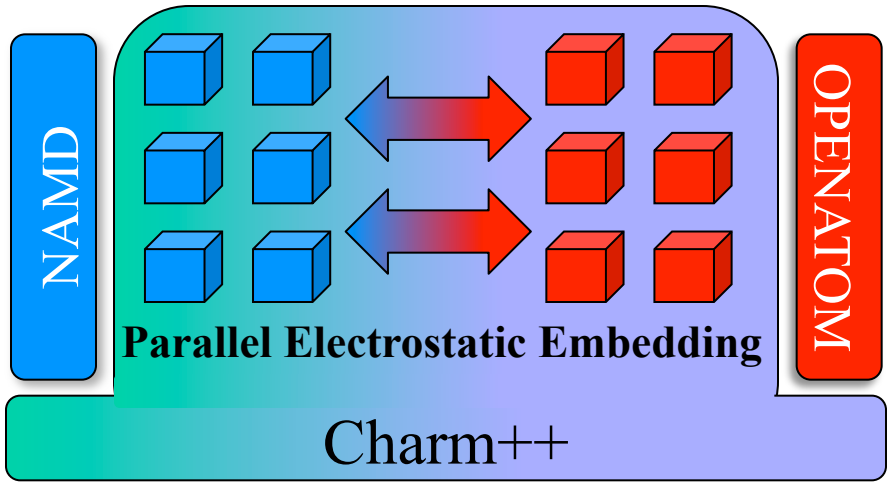

Method combining OpenAtom and NAMD

Parallel Car-Parrinello MD
DOE and NSF Funded 10 yrs
Martyna/Kale Collaboration

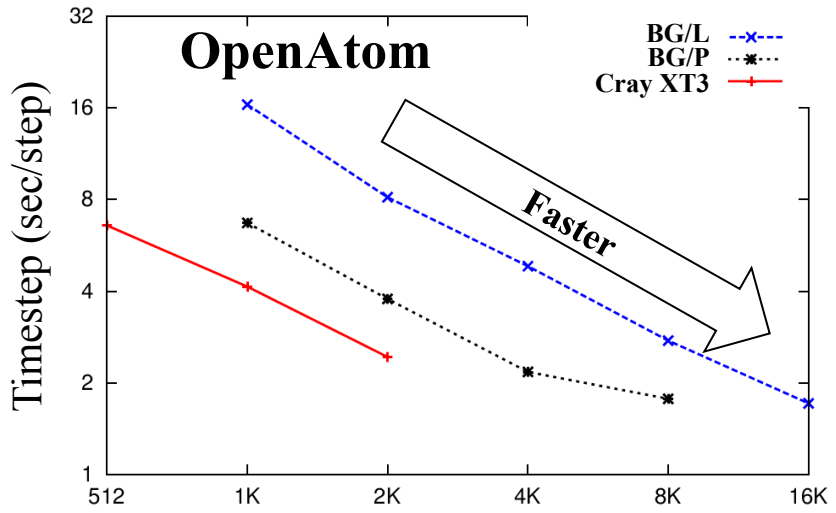

Harrison & Schulten, *Quantum and classical dynamics of ATP hydrolysis in solvent.* Submitted

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

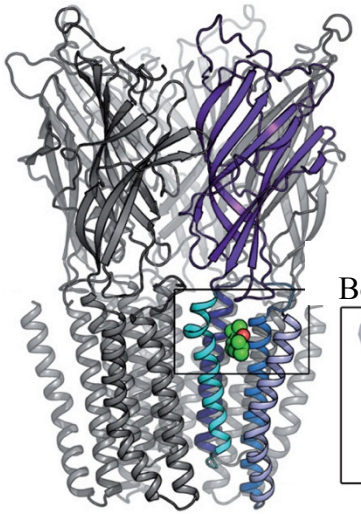

180K-atom 30 ns study of anesthetic binding to bacterial ligand-gated ion channel provided "complementary interpretations…that could not have been deduced from the static structure alone."

Bound Propofol Anesthetic

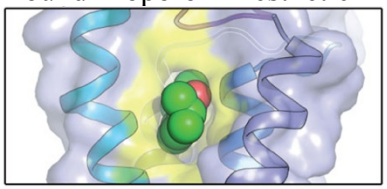

Voth, et al., *PNAS*, 2010

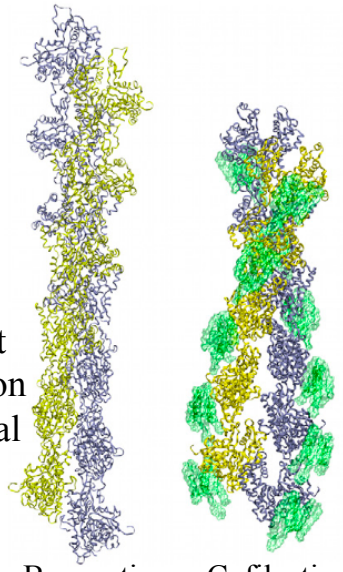

500K-atom 500 ns investigation of effect of actin depolymerization factor/cofilin on mechanical properties and conformational dynamics of actin filament.

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)

Challenges of New Hardware

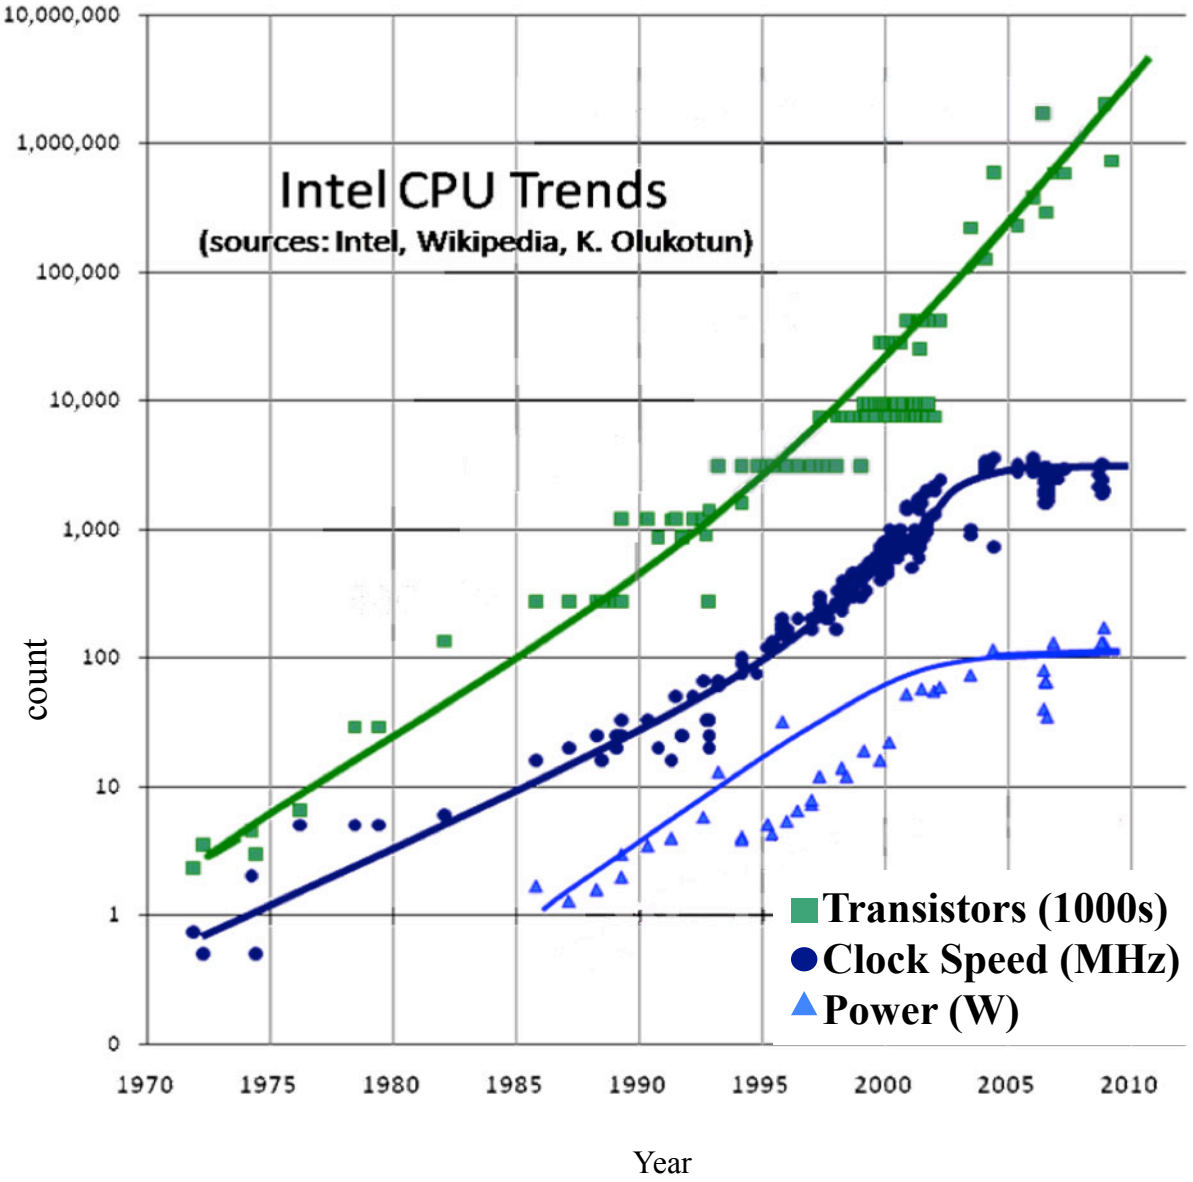

The number of transistors on a chip keeps increasing (and will, for 10 years)

BUT the frequency has stopped increasing (since 2003 or so)

Due to power limits

Harnessing Future Hardware

- Challenge: a panoply of complex and powerful hardware
  - Complex multicore chips, accelerators

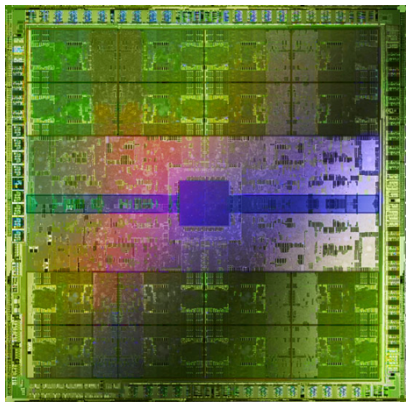

Kepler GPU
(Blue Waters)

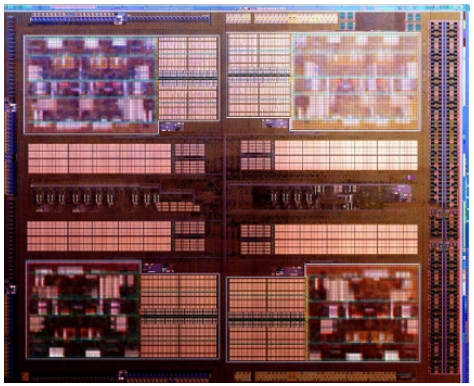

AMD Interlagos
(Blue Waters)

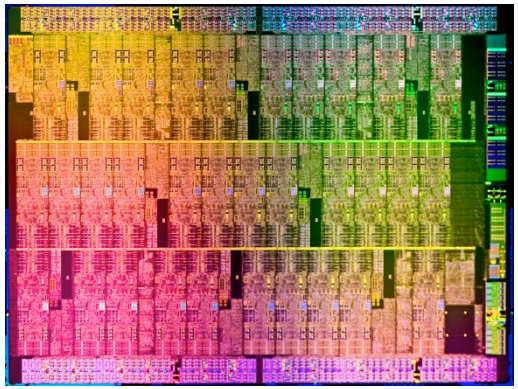

Intel MIC
(TACC Stampede)

- Solution: BTRC computer science expertise
  - Parallel Programming Lab: leading research group in scalable parallel computing

Parallel Programming Lab
University of Illinois at Urbana-Champaign

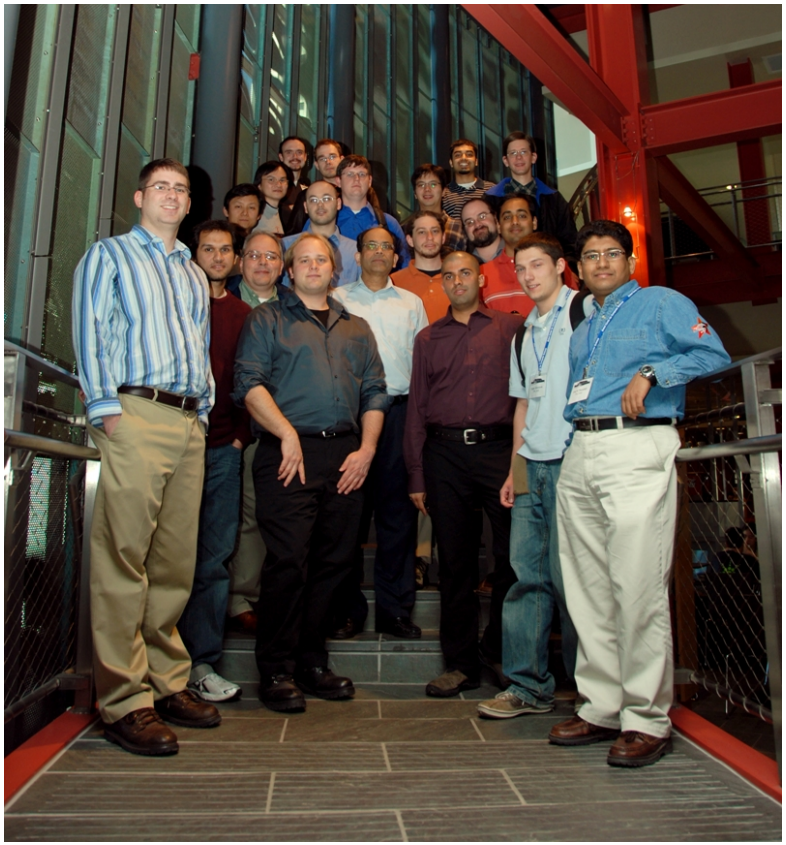

Siebel Center for Computer Science

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

Develop abstractions in context of full-scale applications

Quantum Chemistry (QM/ MM)

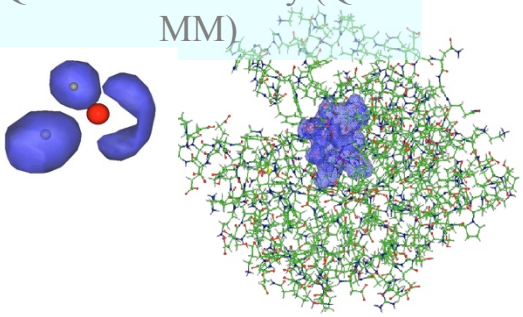

NAMD: Molecular Dynamics

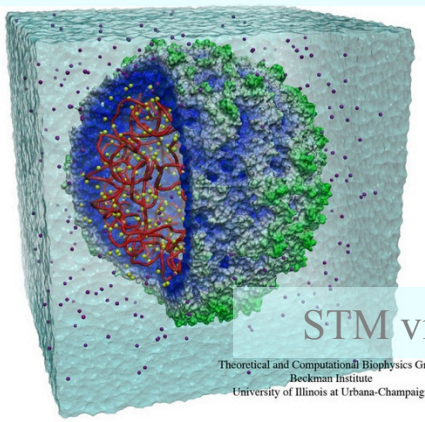

Protein Folding

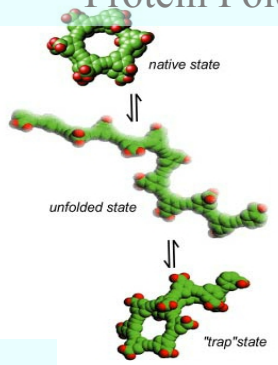

STM virus simulation

Computational Cosmology

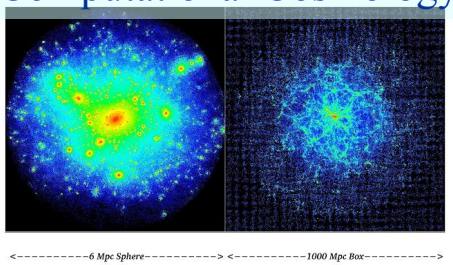

Parallel Objects,

Adaptive Runtime System

Libraries and Tools

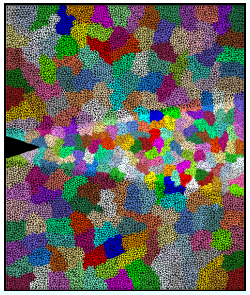

Crack Propagation

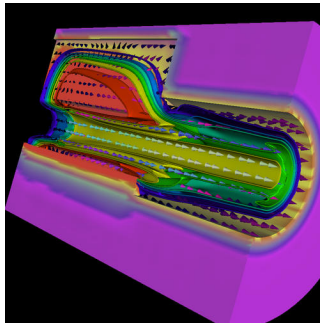

Rocket Simulation

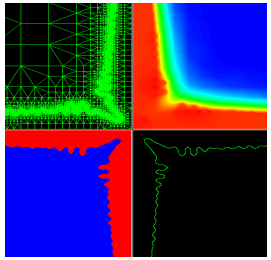

Dendritic Growth

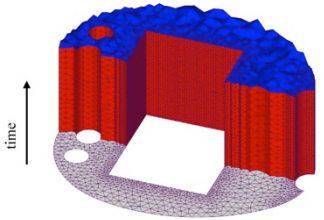

Space-time meshes

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 – directed by Prof 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.

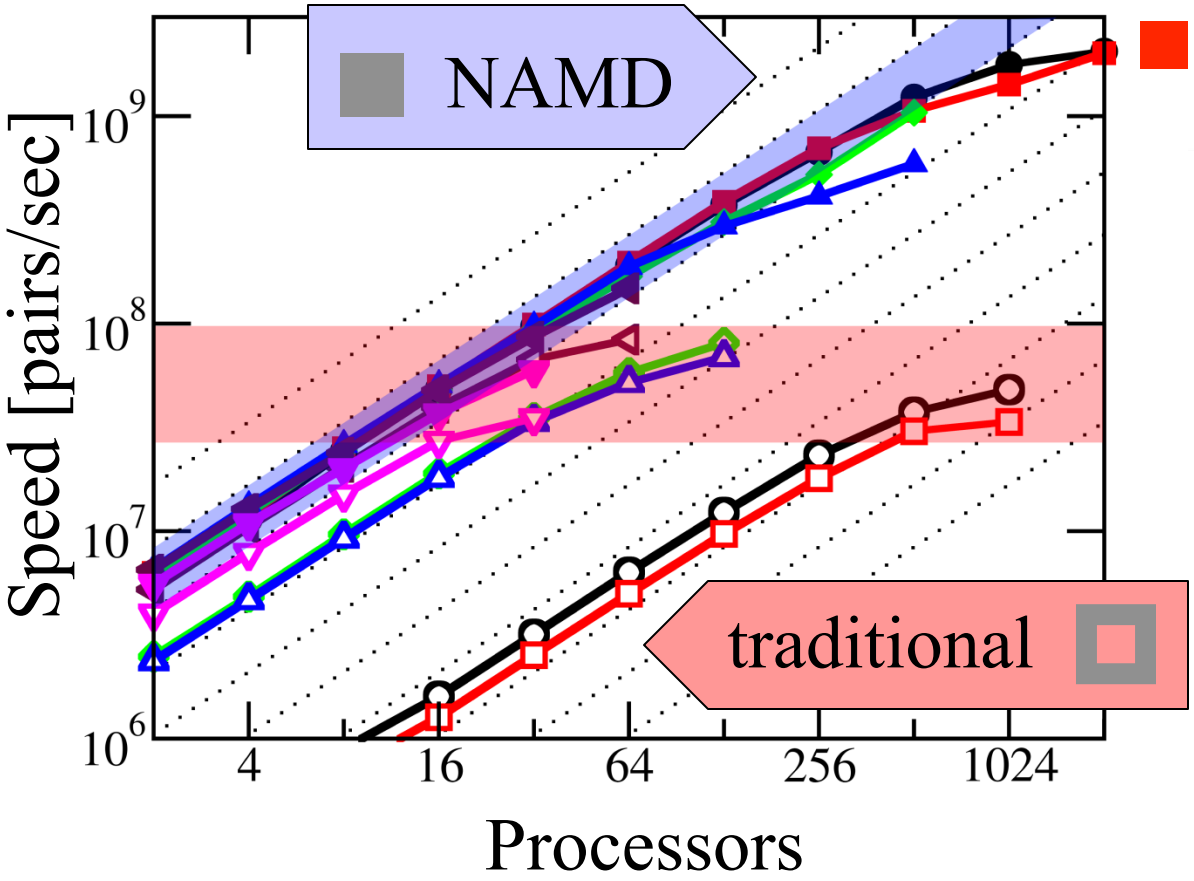

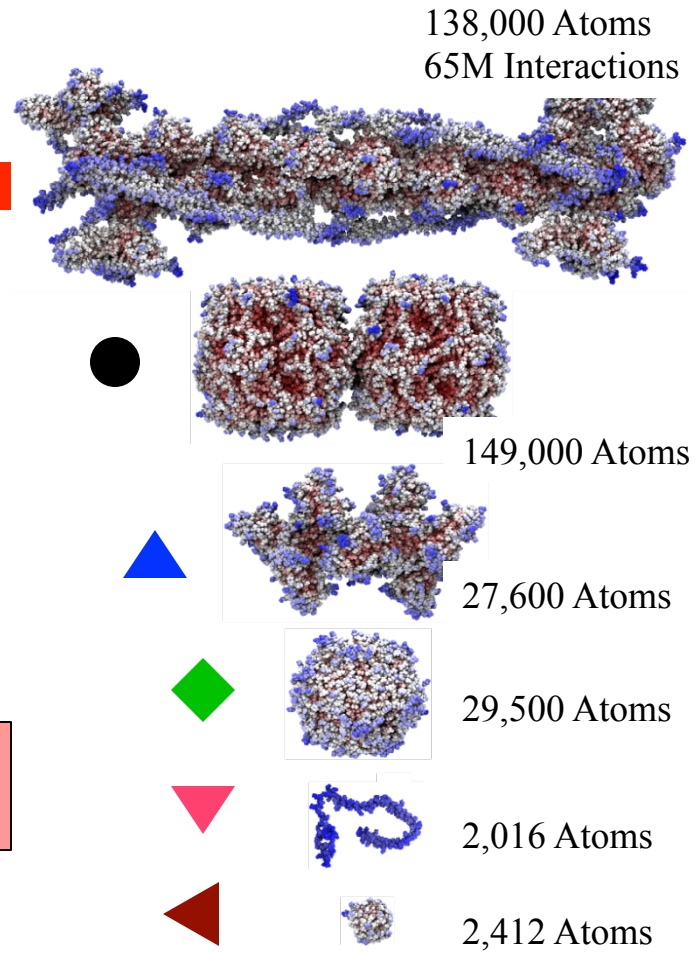

Tanner et al., J. Chem. Theory and Comp., 7:3635-3642, 2011

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

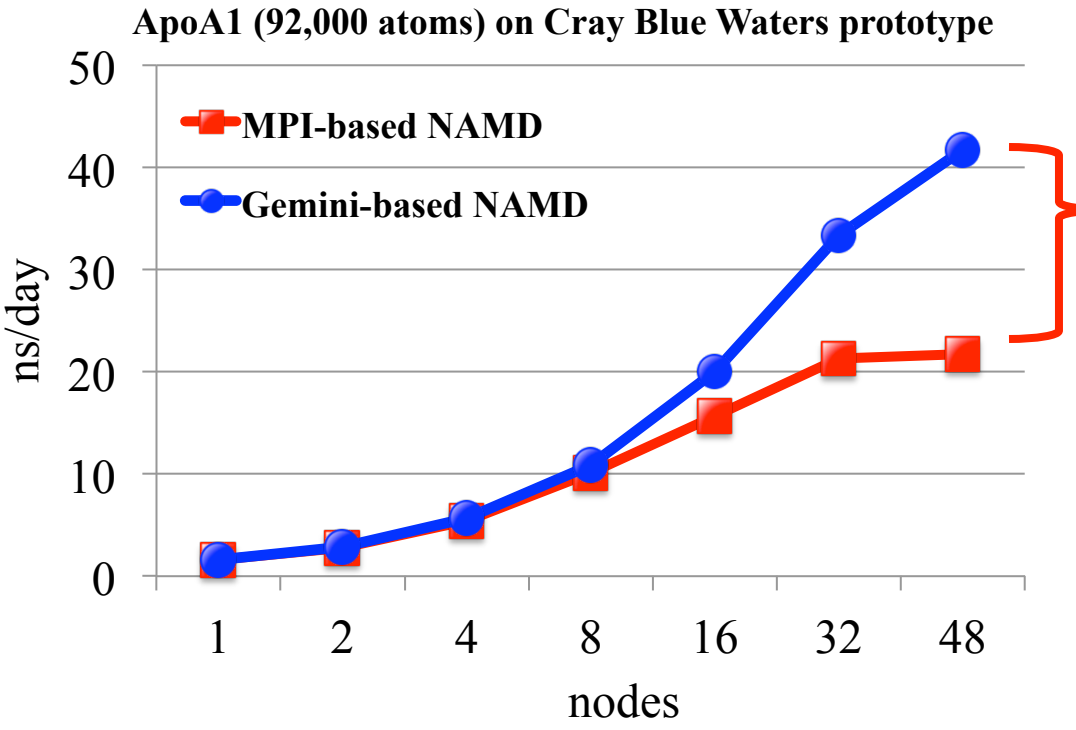

Gemini provides at least 2x increase in usable nodes for strong scaling

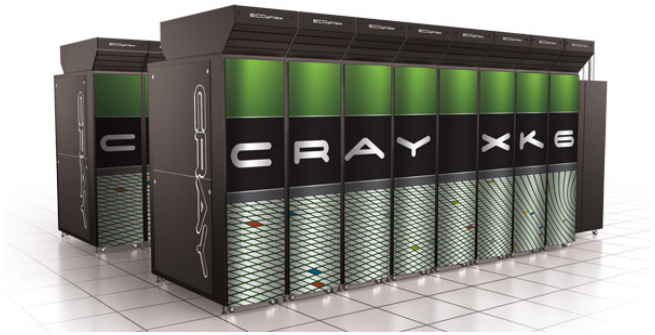

100M Atoms on Titan vs Jaguar

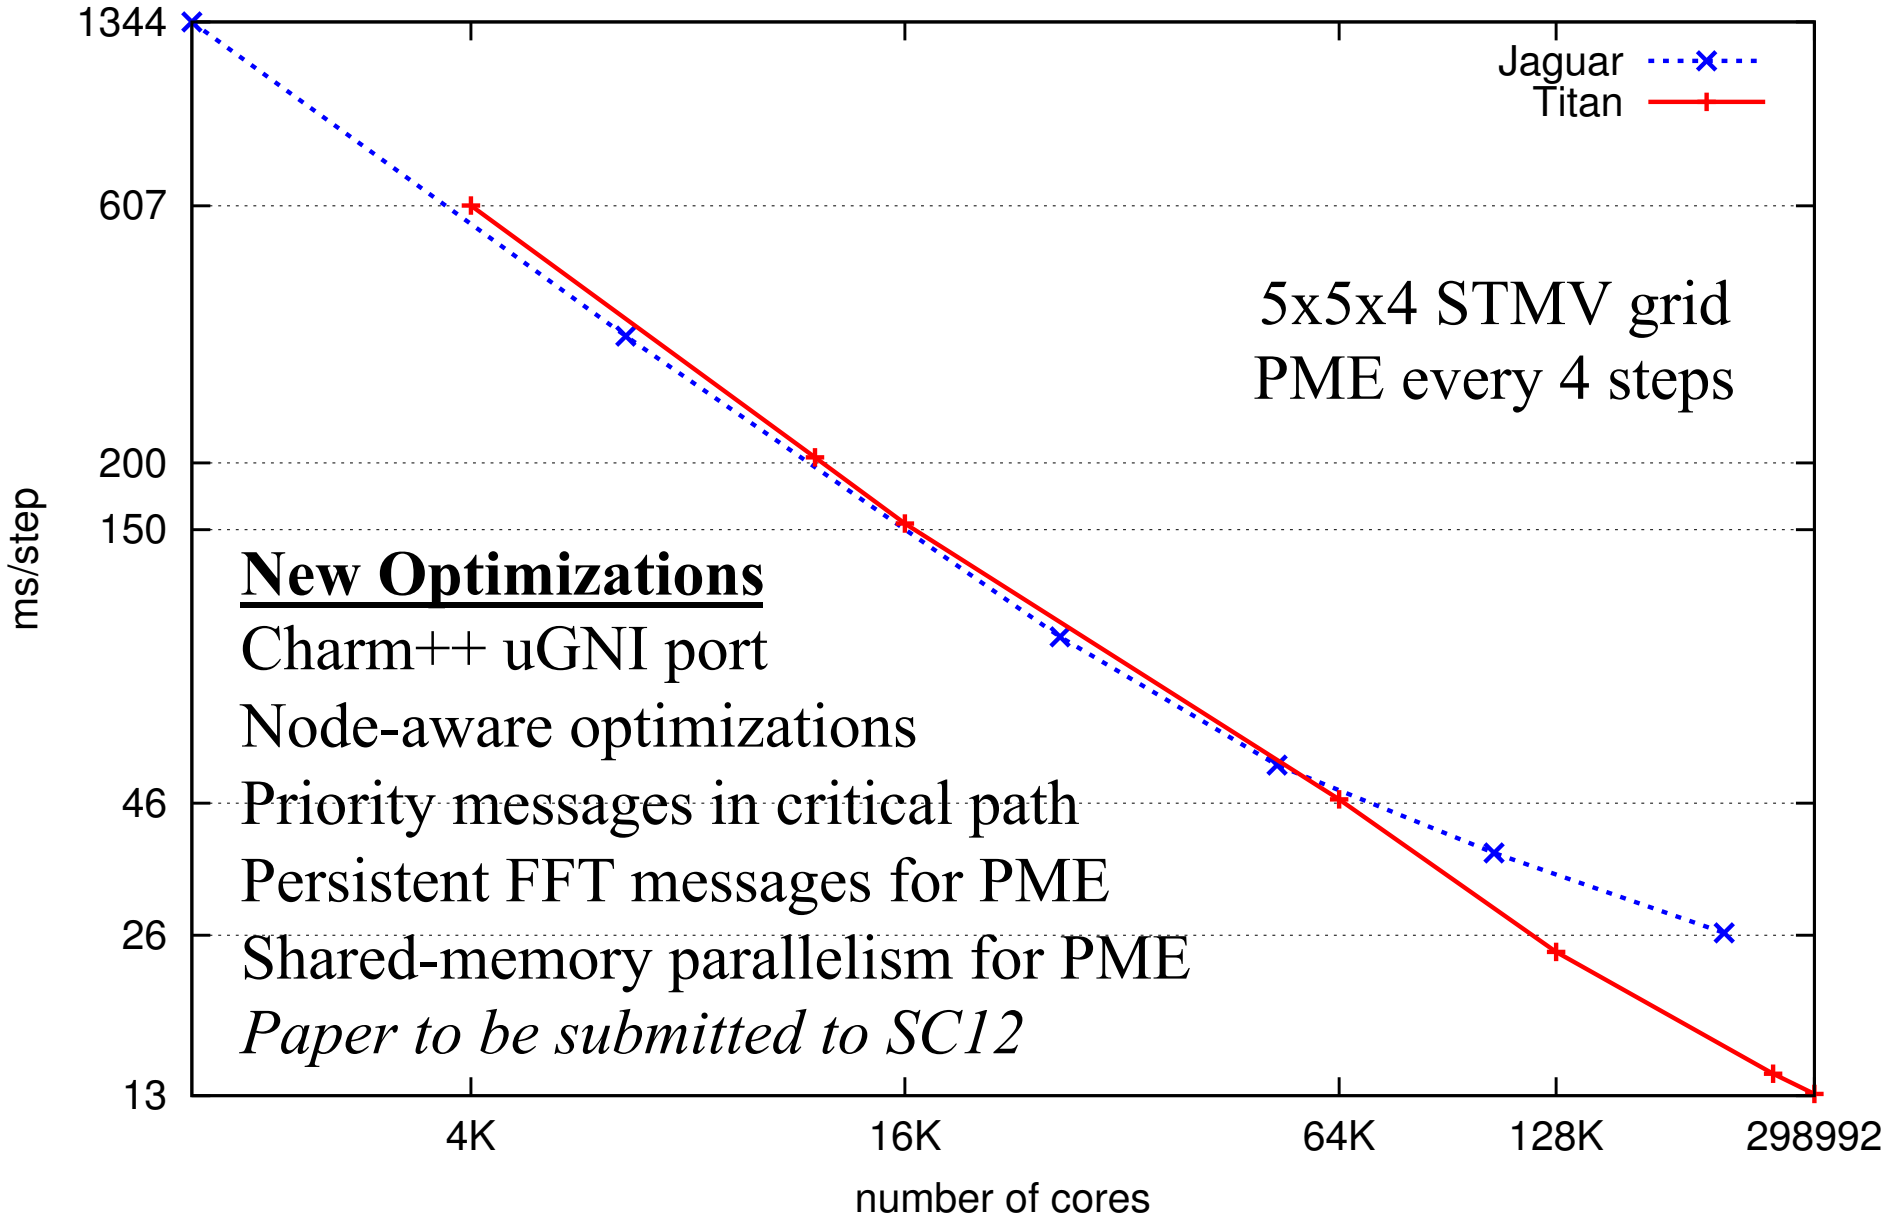

1M Atom Virus on TitanDev GPU

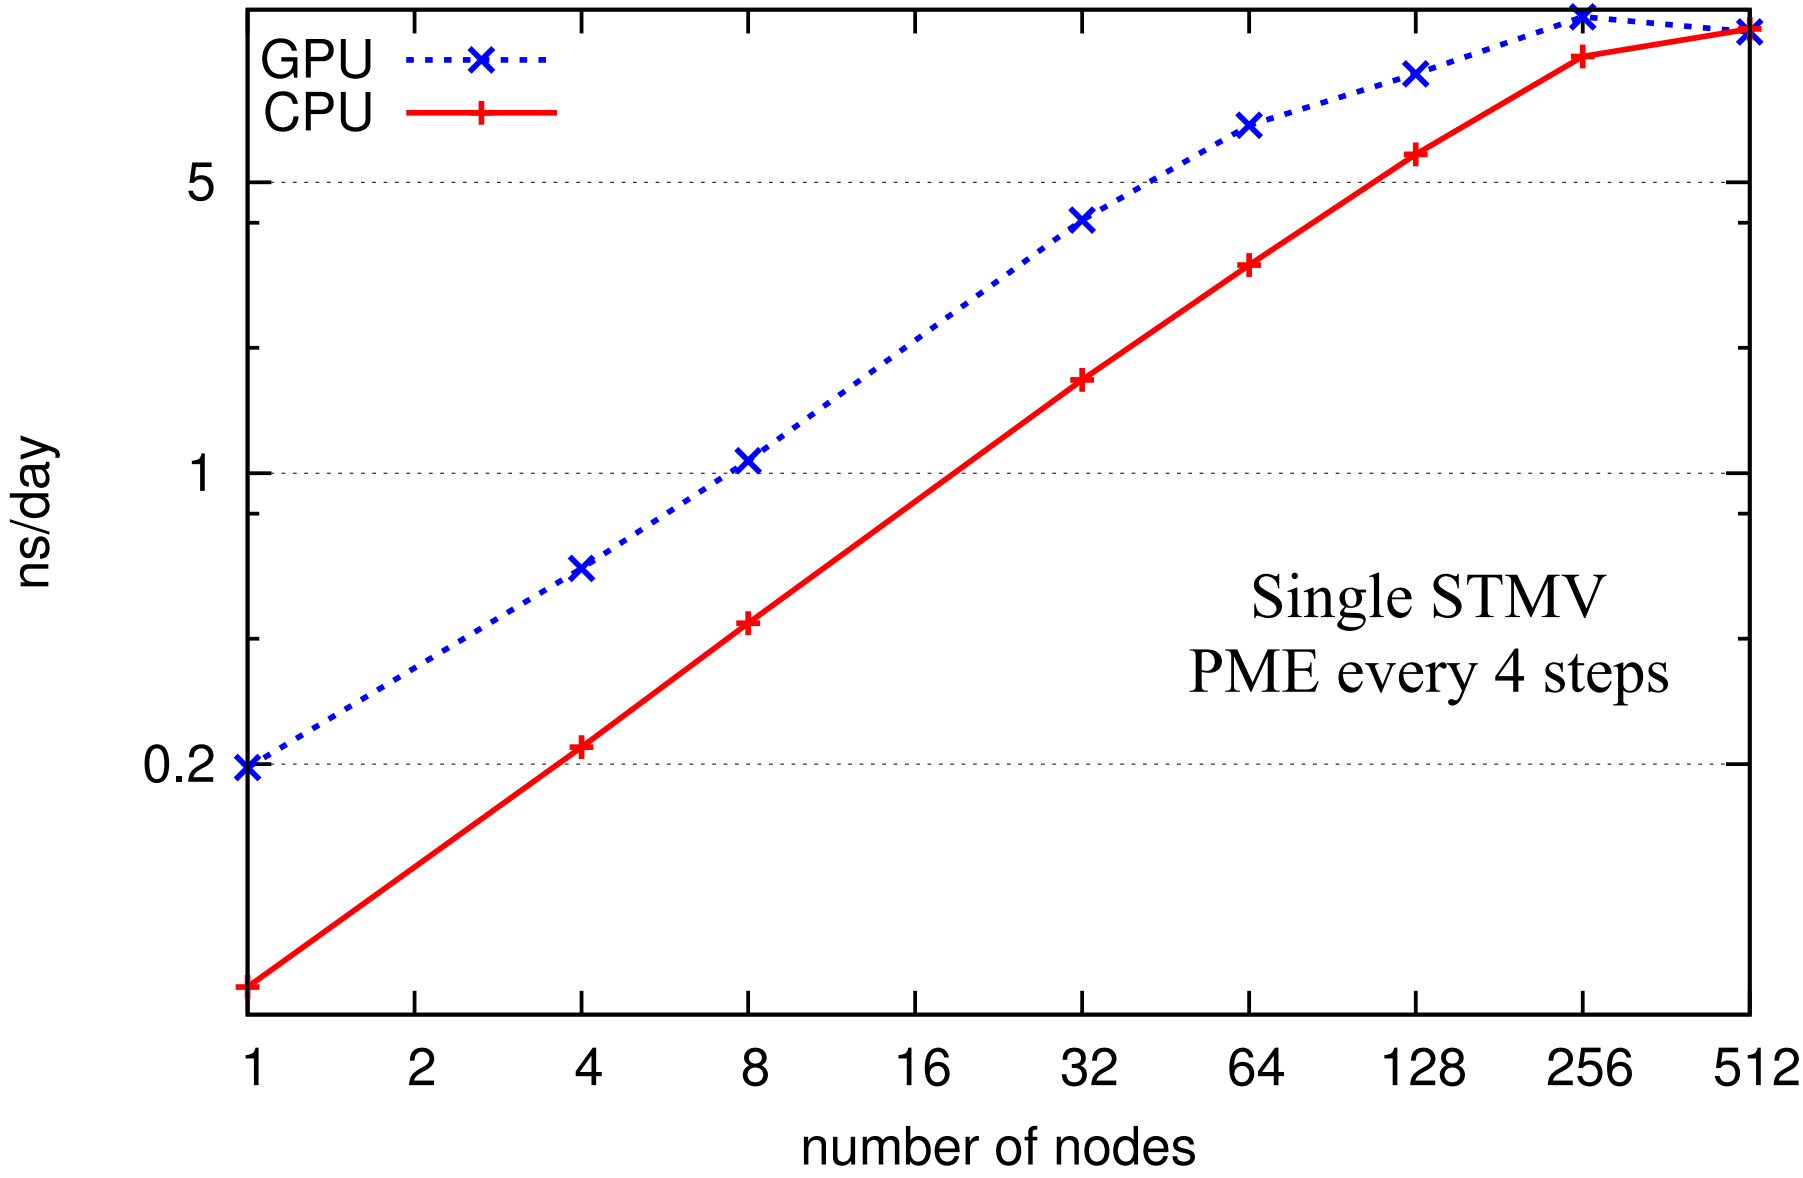

100M Atoms on TitanDev

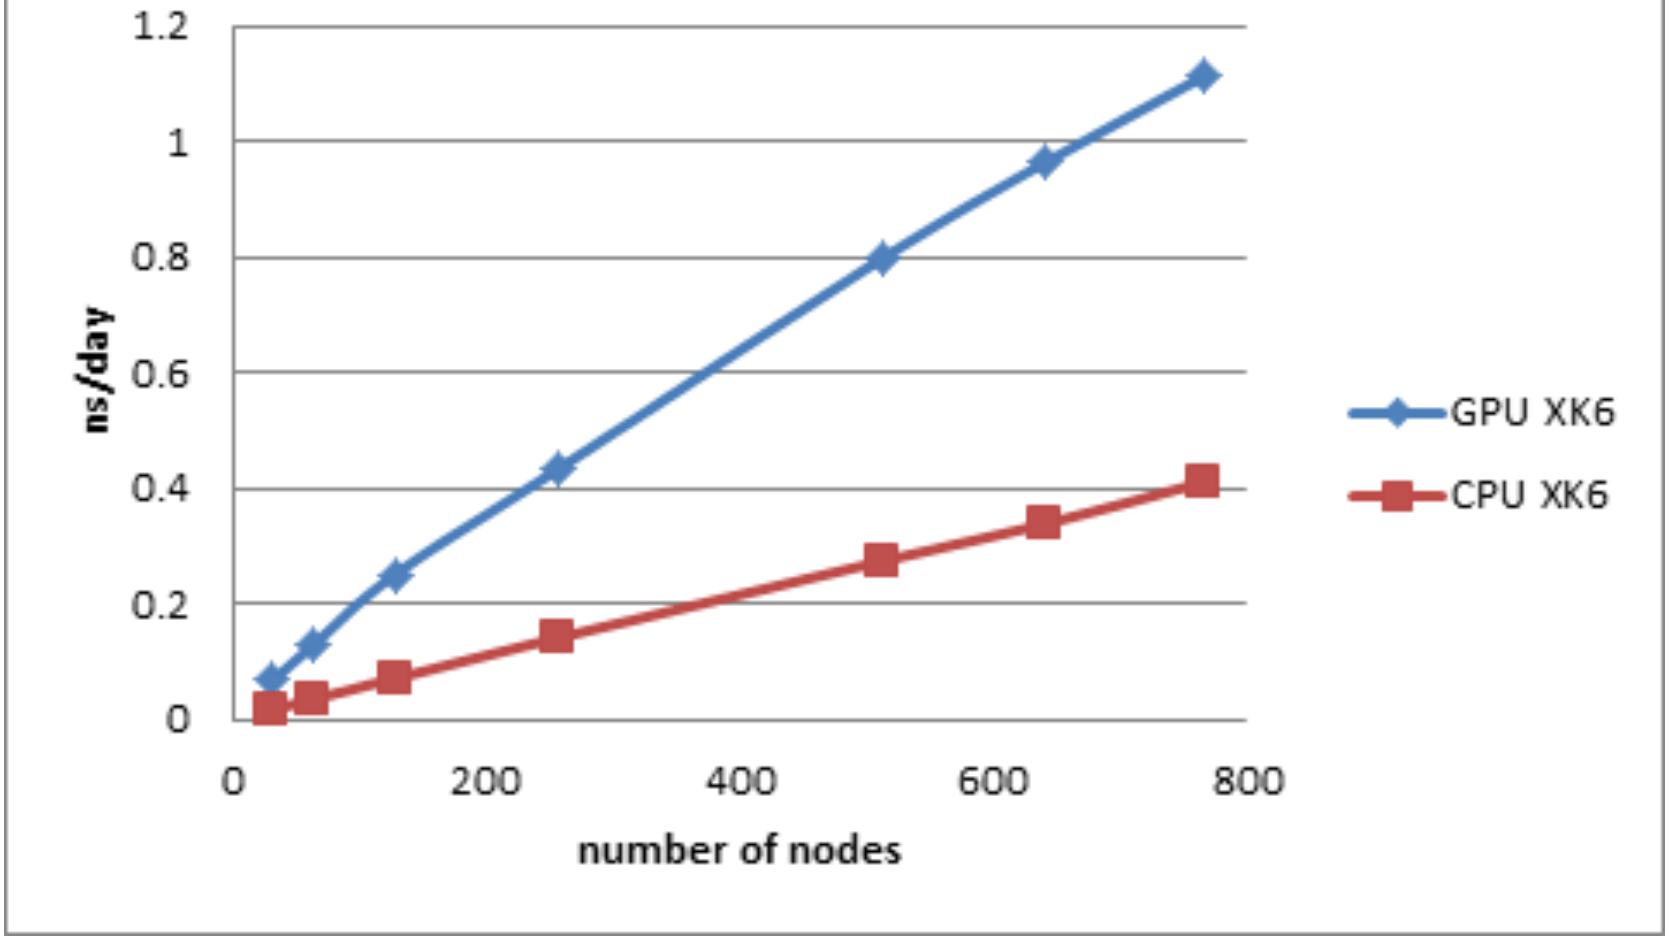

Tsubame (Tokyo) Application of GPU Accelerated NAMD

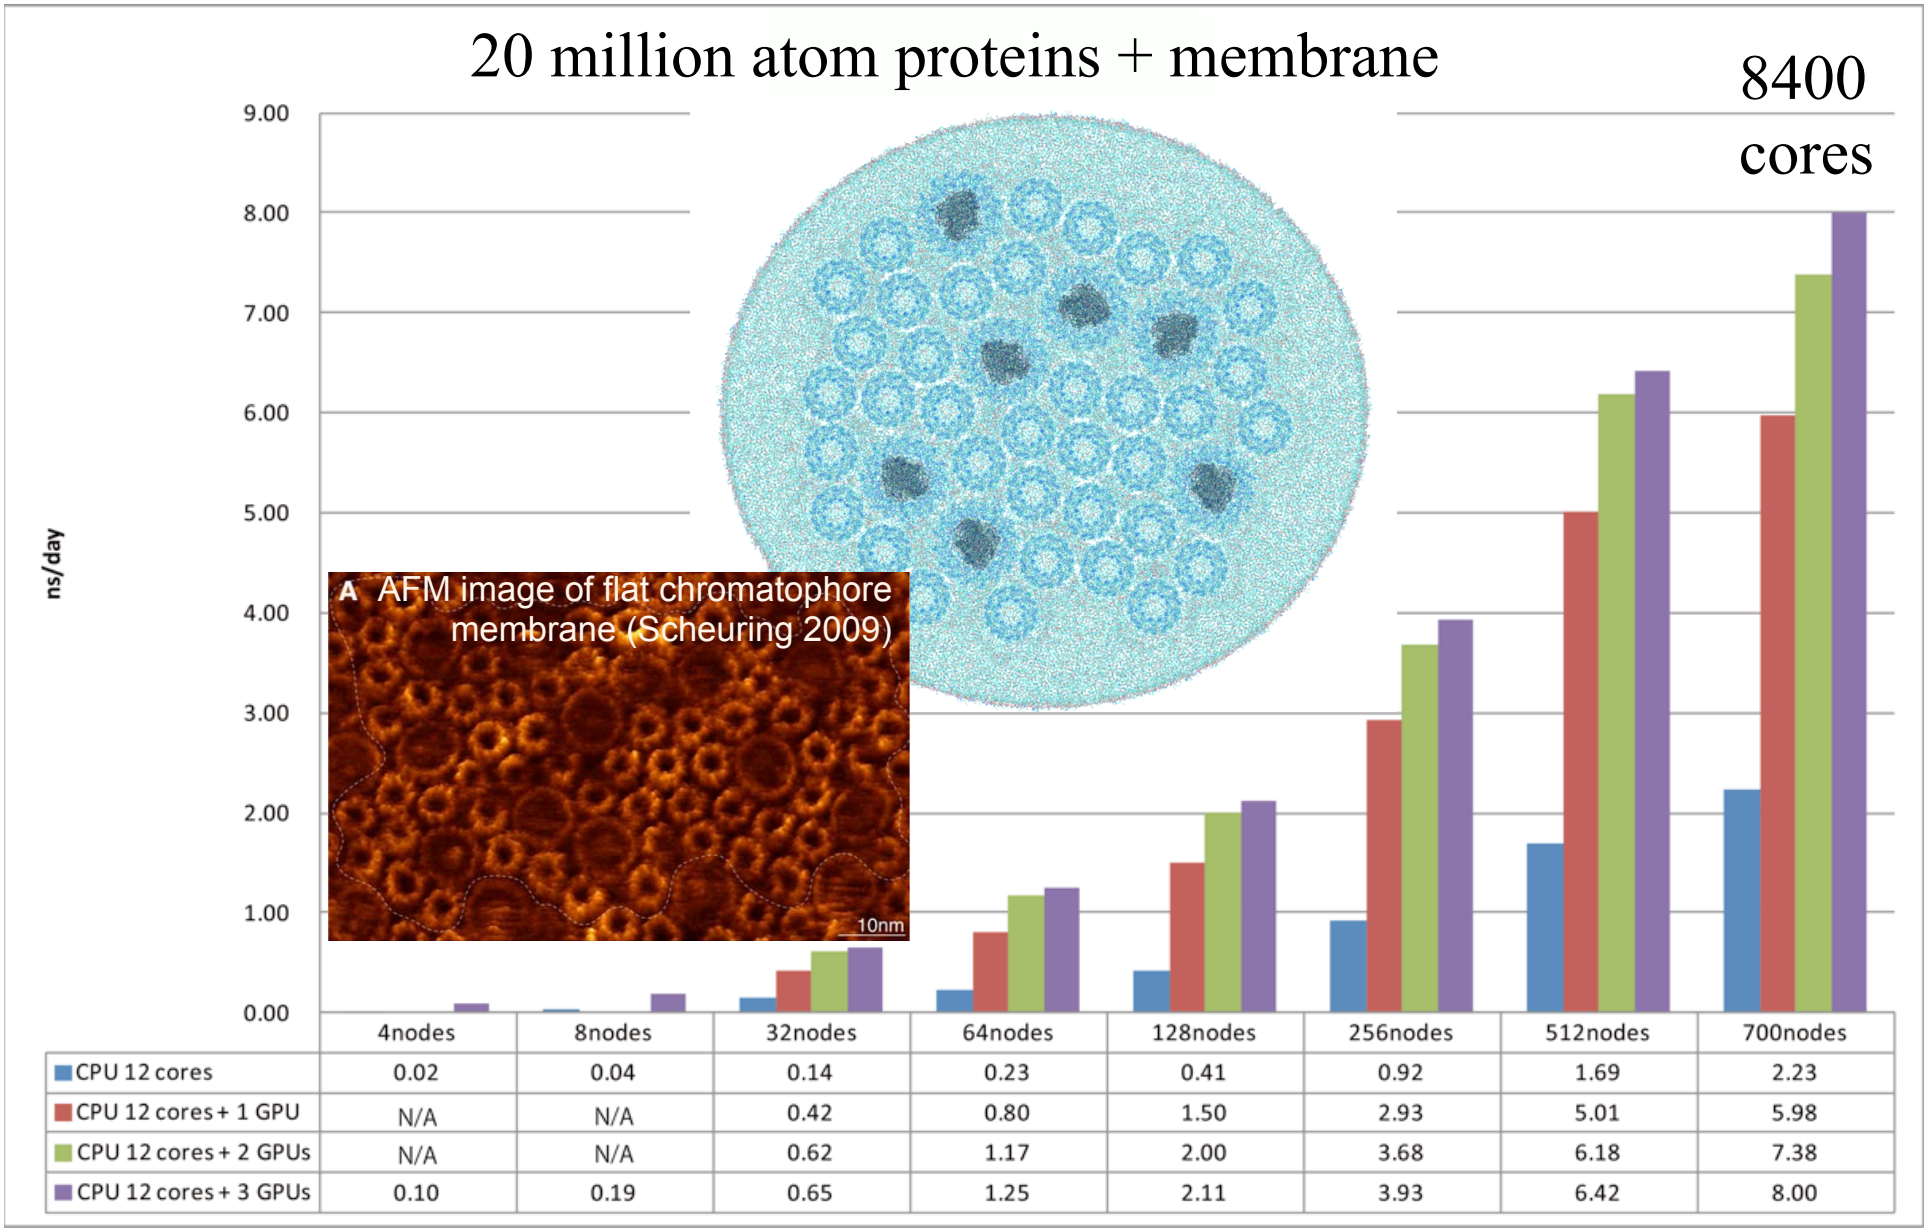

	4nodes	8nodes	32nodes	64nodes	128nodes	256nodes	512nodes	700nodes
CPU 12 cores	0.02	0.04	0.14	0.23	0.41	0.92	1.69	2.23
CPU 12 cores + 1 GPU	N/A	N/A	0.42	0.80	1.50	2.93	5.01	5.98
CPU 12 cores + 2 GPUs	N/A	N/A	0.62	1.17	2.00	3.68	6.18	7.38
CPU 12 cores + 3 GPUs	0.10	0.19	0.65	1.25	2.11	3.93	6.42	8.00

GPU Computing in NAMD and VMD

- NAMD algorithms to be discussed:
  - Short-range non-bonded interactions
  - Generalized Born Implicit Solvent
  - Multilevel Summation Method
- VMD algorithms to be discussed:
  - Electrostatic potential maps
  - Visualizing molecular orbitals
  - Radial distribution functions
  - “QuickSurf” representation