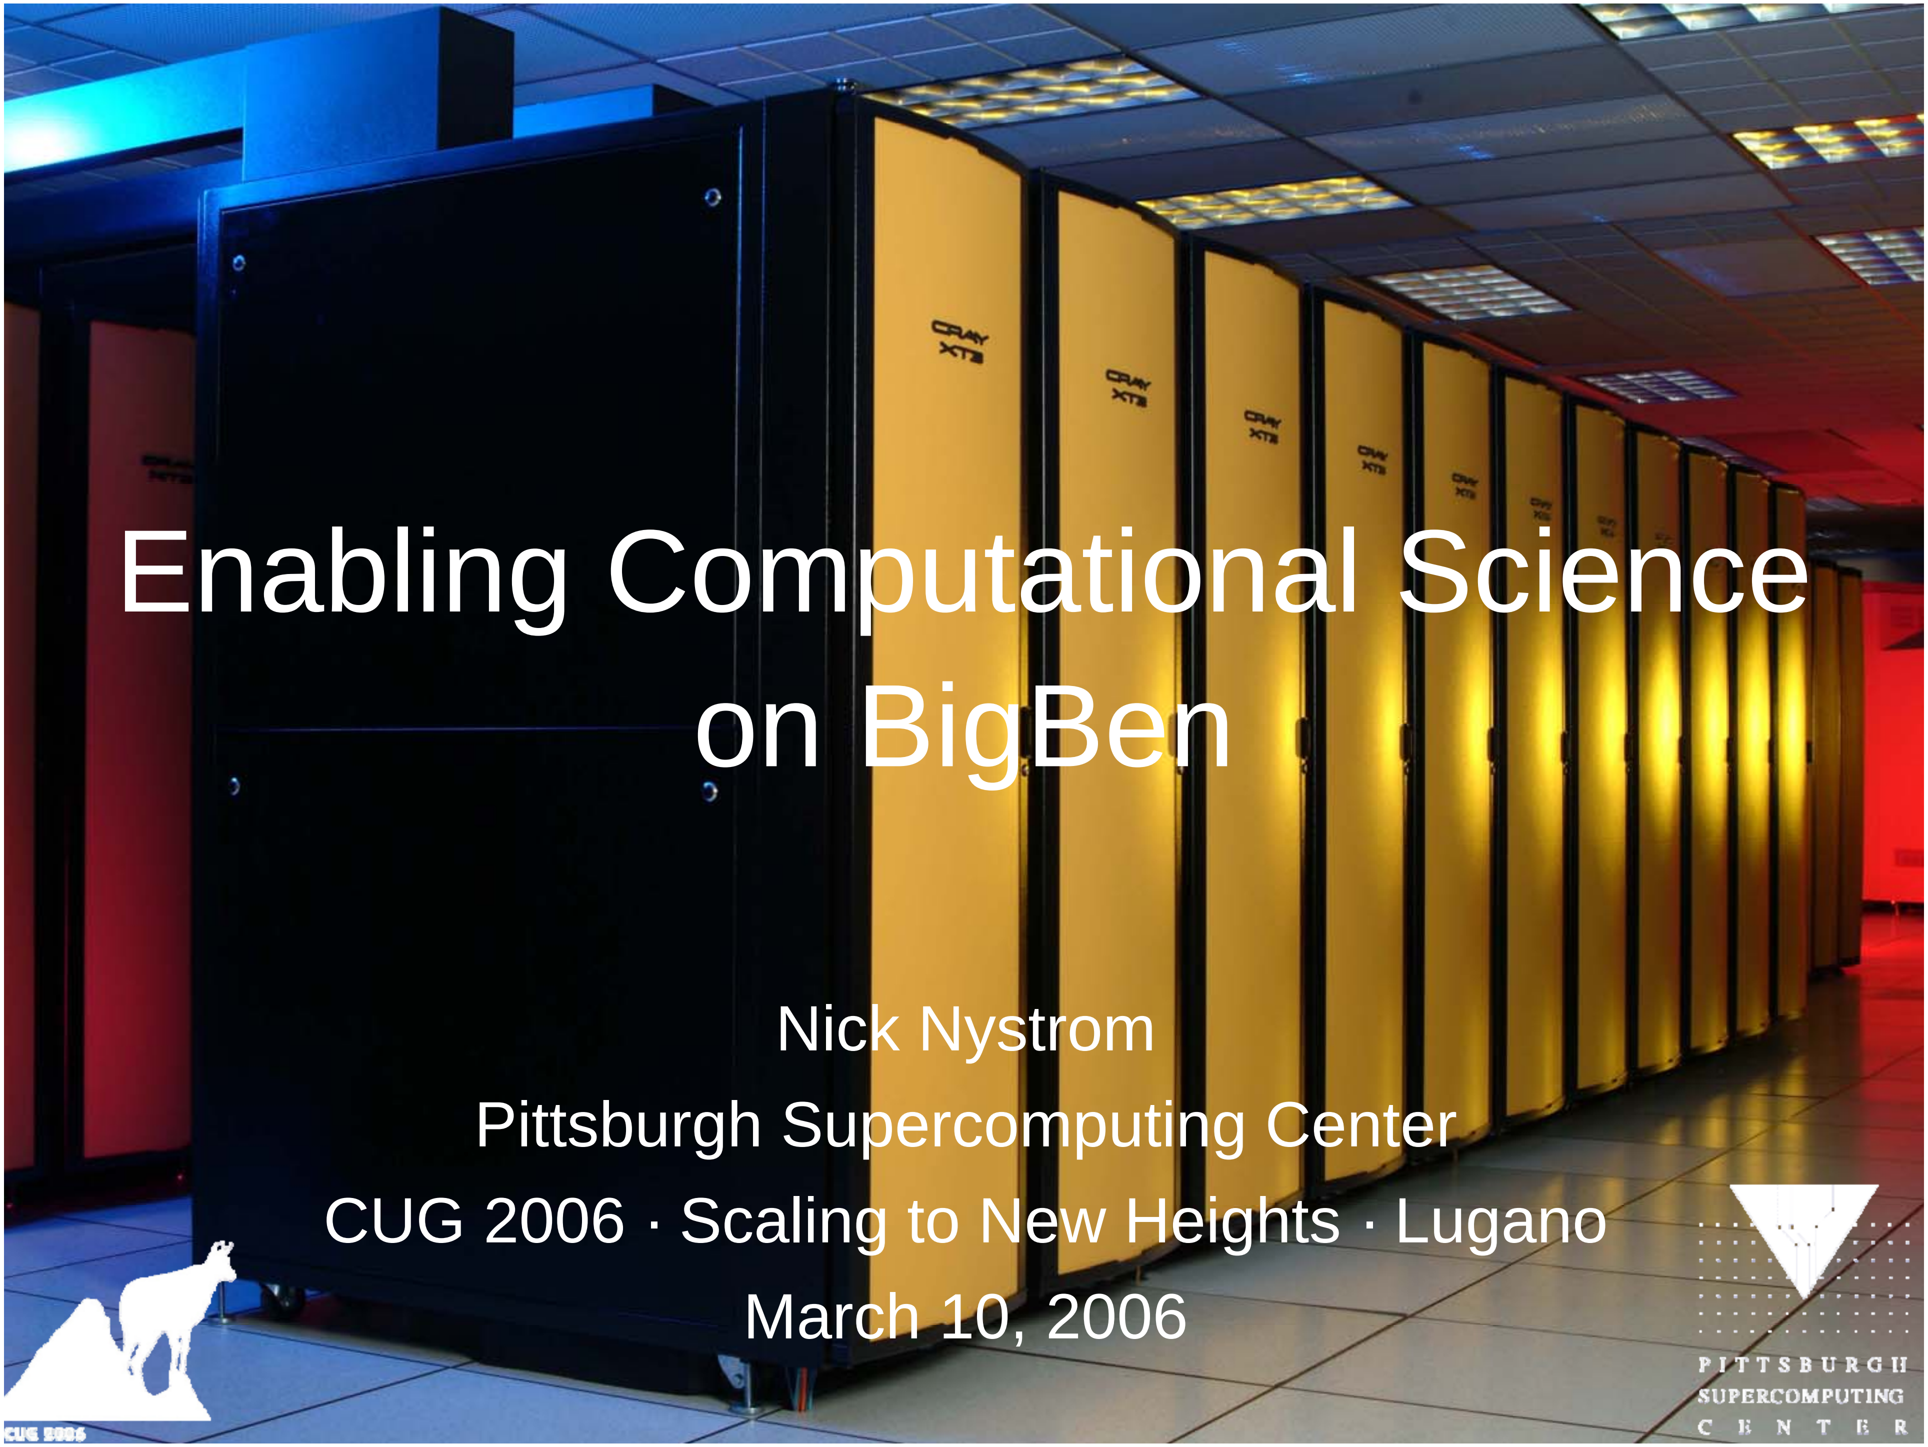

Enabling Computational Science on BigBen

Nick Nystrom

Pittsburgh Supercomputing Center

CUG 2006 · Scaling to New Heights · Lugano

March 10, 2006

Lest we run out of time…

- This presentation reflects the work of numerous PSC staff and users, including
  - Shawn Brown, David Deerfield, Jeff Gardner, Roberto Gomez, Junwoo Lim, Demian Nave, David O'Neal, Stu Pomerantz, Rich Raymond, R. Reddy, Nathan Stone, John Urbanic, Yang Wang, Deborah Weisser, Joel Welling, Art Wetzel, Troy Wymore
  - Jacobo Bielak, Tiziana Di Matteo, Kelven Droegemeier, Robert Duke, M. Eisenbach, J. S. Faulkner, Omar Ghattas, Kosuke Imai, L. V. Kale, George Lake, David O'Hallaron, D. M. C. Nicholson, Jim Phillips, David Porter, Kevin Quinn, Tom Quinn, R. Radhakrishnan , D. C. Richardson, A. Rusanu, Klaus Schulten, J. Stadel, G. M. Stocks, R. Venkatramani, J. Wadsley, Paul Woodward, Ming Xue
  - John Levesque, Luiz DeRose, Ted Packwood, Pete Johnsen, Howard Pritchard, Doug Gilmore, Norm Troullier, Jeff Brooks, Jeff Larkin, Mick Dungworth, Jim Harrell, Peter Hill, Charlie Carroll, Sarah Anderson, Cindy Nuss, and Pat Brockway (Cray Inc.) provided valuable technical and logistical assistance.
  - Any inadvertent omissions are regretted and will be corrected in a subsequent release of this presentation.

PSC: A Brief History

- Established in 1986 as a joint effort of Carnegie-Mellon University and the University of Pittsburgh, together with Westinghouse Electric Company
- Major funding provided by the National Science Foundation (NSF), supplemented by the state of Pennsylvania
- Other grants from the National Institutes of Health (major Research Resource), NSF Networking, Department of Energy, Department of Defense
- Strong focus on providing state-of-the-art facilities and support for large-scale computation
- Significant software developments: networking, filesystems, scaling, …
- Particular strengths in networking research and in biomedical applications of HPCC

PSC’s Mission

- Advance Science and Engineering Research through use of High Performance Computing and Communications (HPCC)
- Advance the state-of-the-art of HPCC
- Educate researchers in techniques of HPCC
- Help American industry exploit HPCC

PSC Facilities (subset)

Rachel (128p, 512GB)
Jonas (128p, 512GB)

XT3 (BigBen)
2090p, 10.0 TFlop/s

TCS (LeMieux)
3000p, 6.0 TFlop/s

Storage Silos
2 PB

DMF Archive Server

Storage
Cache Nodes
100 TB

Visualization
Nodes
NVidia Quadro4 980XGL

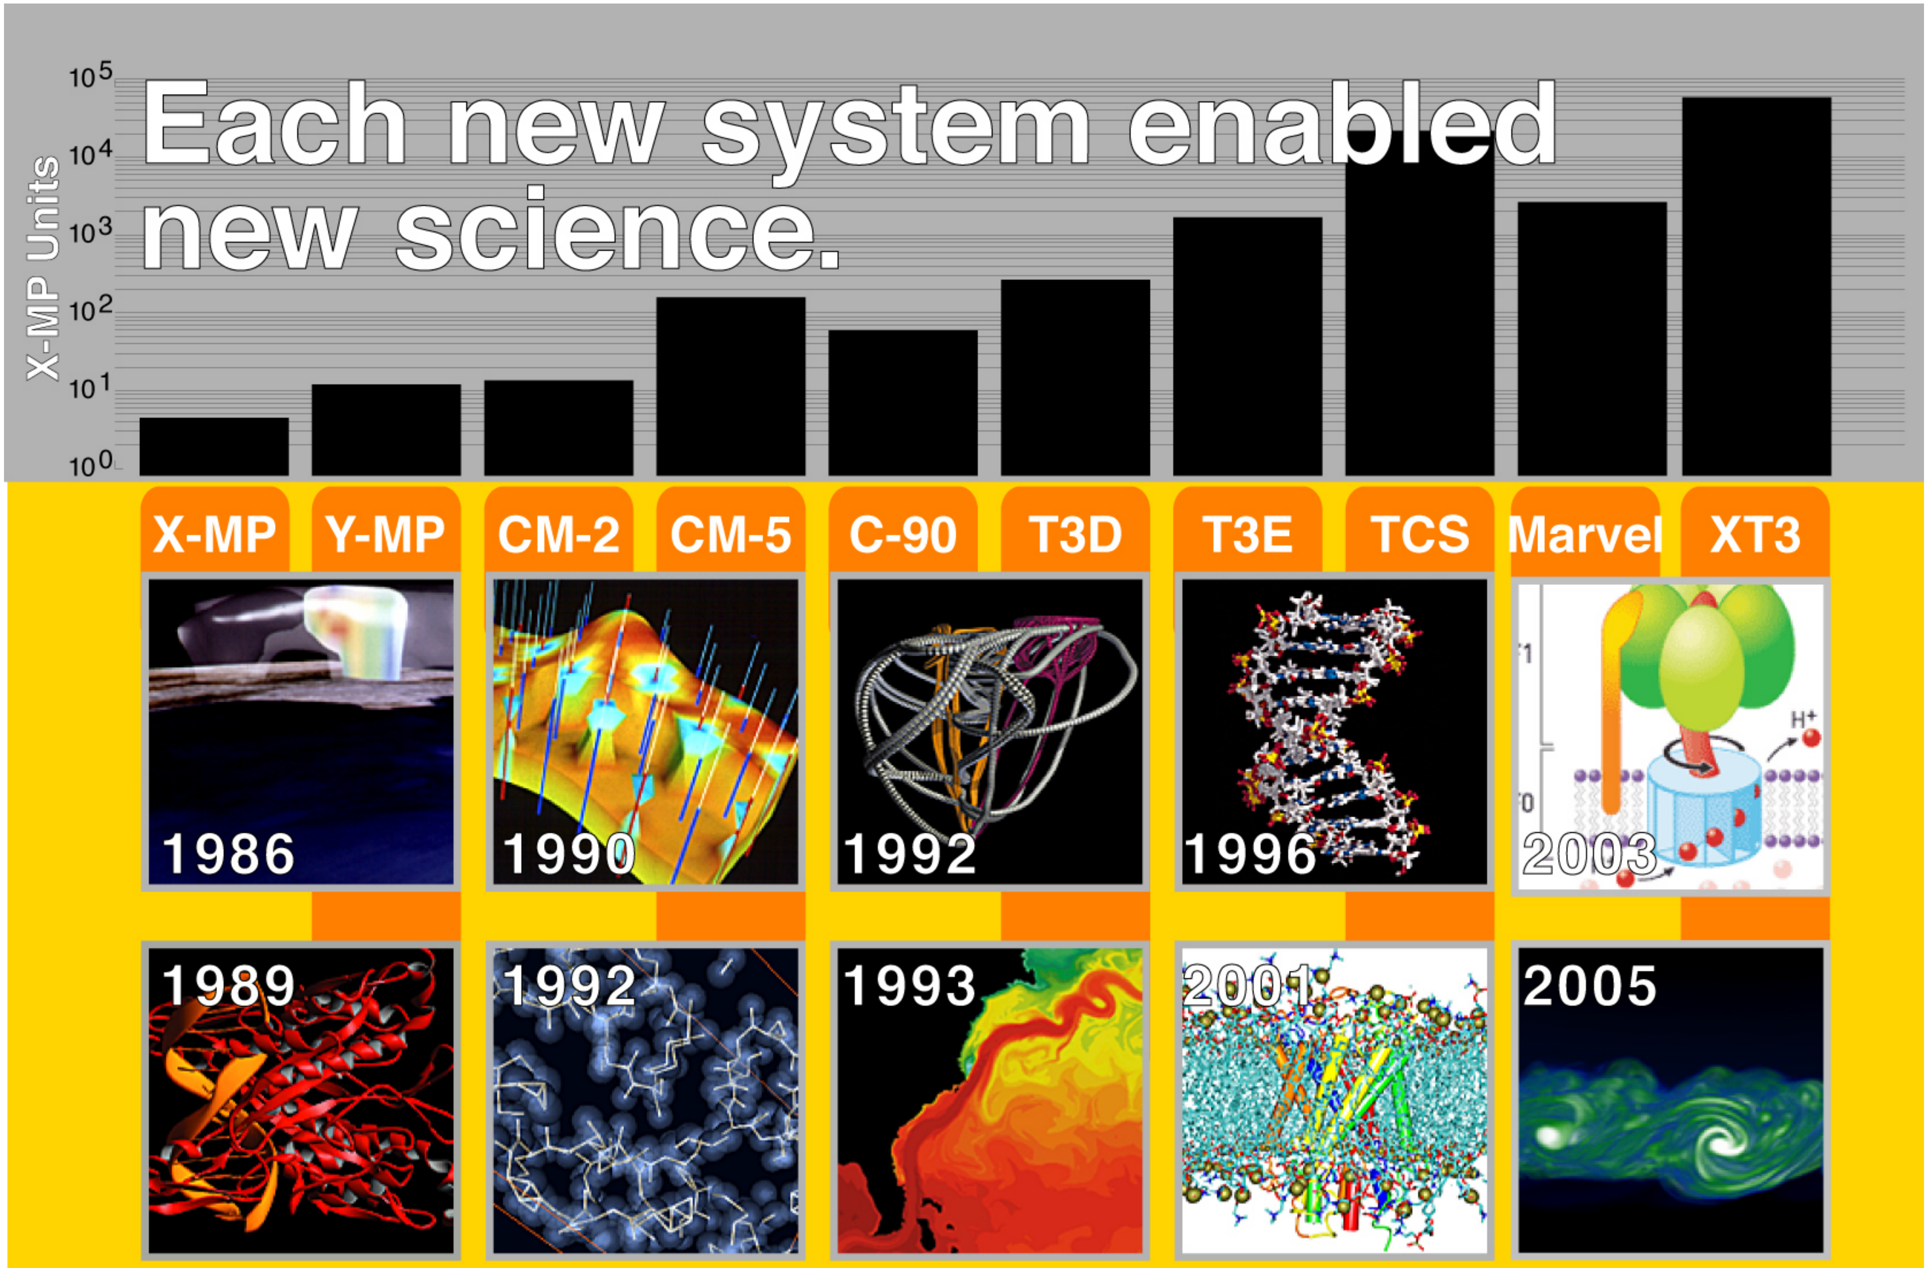

History of first or early systems

Transcending Barriers to Realism

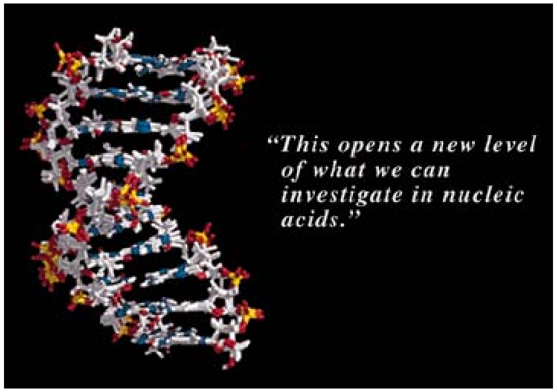

DNA structure and dynamics
P. Kollman & T. Cheatham (UCSF),
M. Crowley (PSC)

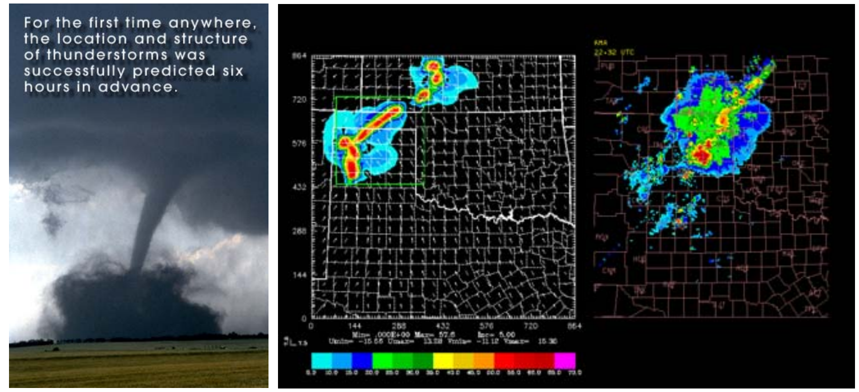

Operational storm-scale forecasting
Kelvin Droegemeier (CAPS @ U. of Oklahoma)

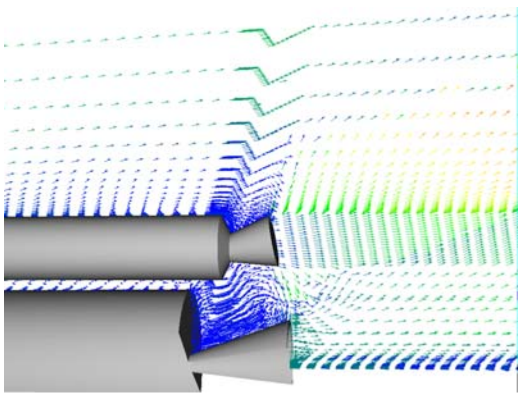

Delta II Rocket backflow
Stephen Taylor (Caltech)

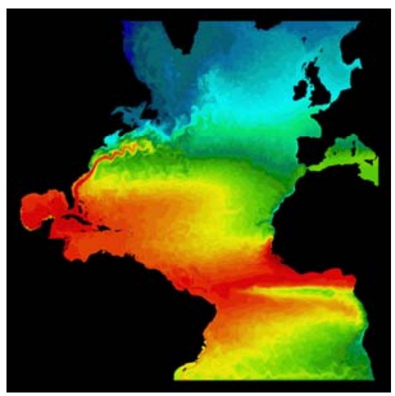

Resolution of the Gulf Stream
Rainer Bleck (U. Miami)

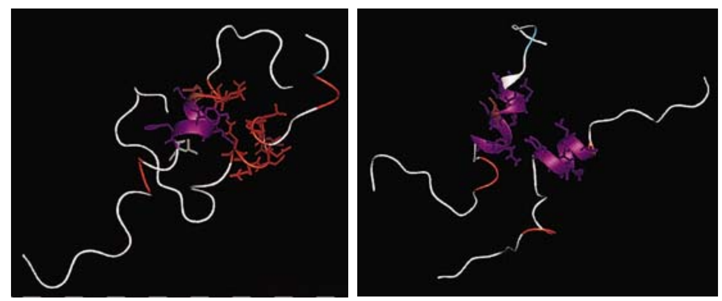

Formation of leucine zipper coiled-coil dimer
C. Brooks (CMU, Scripps), W. Young (PSC)

NSF TeraGrid Participation

- TeraGrid DEEP: Enabling the Nation's Terascale Science
  - Make science more productive through a unified set of very-high capability resources
- TeraGrid WIDE: Empowering communities to leverage TeraGrid capabilities
  - Bring TG capabilities to the broad science community
- Base TeraGrid CyberInfrastructure: integrating the nation's most powerful resources
- Provide a unified, general purpose, reliable set of services and resources
- PSC heads User Support and Security for TeraGrid

Adding Capability to the NSF TeraGrid

TERAGRID

ABOUT PRESS SITE MAP USER INFO

Home > Media News

Media News

Search TeraGrid GO

Pittsburgh Unveils Big Ben the Supercomputer

July 20, 2005

Contact:
Michael Schneider, *Pittsburgh Supercomputing Center*
schneider@psc.edu, 412.268.4960

The newest, most advanced Cray Inc. system is up and running at the Pittsburgh Supercomputing Center, where it will support research nationwide as part of the NSF TeraGrid. Researchers can learn about the system by attending a special workshop to be held at PSC, Aug. 9-12, 2005.

PITTSBURGH — The Pittsburgh Supercomputing Center (PSC) has switched sports, but its newest and most powerful system, the Cray XT3, is another black-and-gold superstar, say officials. With a nod to the Pittsburgh Steelers quarterback, Big Ben is Cray XT3 serial #1 - the newest stage in the evolution of high-performance computing technology and a major boost for computational science in the United States.

Acquired via a $9.7 million grant from the National Science Foundation (NSF) in September 2004, Big Ben - the first XT3 system to ship from Cray - comprises 2,090 processors with an overall peak performance of 10 teraflops: 10 trillion calculations per second. If every person on Earth, about 6.5 billion people, held a calculator and did one calculation per second, they would all together still be 1,500 times slower than Big Ben.

Section site map: Press

- Press Home
- User/System News
- Quote Sheet
- Media Appearances

2005 News Stories

- GADU/GNARE Uses TeraGrid For Protein Sequence Analysis
- TeraGyroid: Gyroids on the TeraGrid
- TeraShake: Simulating a Big Shake in Southern California Basins
- Seismic Modeling and Oil Reservoir Simulations with TeraGrid
- Improving Groundwater Cleanup Decision-Making with TeraGrid
- Harnessing TeraGrid to the Sky: Understanding Dark Energy
- AMANDA and TeraGrid: Exploring the Violent Universe
- TeraGrid Science Gateways: NanoHUB
- Identifying Brain Disorders with TeraGrid

Big Ben began TeraGrid production on October 1, 2005 (as scheduled)

Big Ben as a TeraGrid Resource

Sites	SDSC	NCSA	UC/ANL	PSC	IU	PURDUE	ORNL	TACC
Computational Resources	Itanium2/ IA-64 • SUSE Linux SLES8/ 2.4 SMP • 4 TF • 512 CPU Power4+ (DataStar) • IBM AIX 5L 5.2 • 0.65 TF • 96 CPU	Itanium2/ IA-64 • SUSE Linux SLES8/ 2.4 SMP • 5.2 TF Phase I: • 512 CPU Phase II: • 1262 CPU • 6 TF SGI Altix • SGI ProPack 3.4 • 6.5 TF • 1024 CPU	Itanium2/ IA-64 • SUSE Linux SLES8/ 2.4 SMP • 0.5 TF Phase I: • 32 CPU Phase II: • 92 CPU Xeon/ IA-32 • SuSE Linux SLES8/ 2.4 SMP • 0.5 TF • 192 CPU	Alpha EV68 (Lemieux) • Tru64 Unix • 6 TF • 3000 CPU Alpha EV7 (Rachel) • Tru64 Unix • 0.3 TF • 128 CPU Cray XT3 (Big Ben) • **Front End:** SuSE Linux **Compute processors:** Catamount • 10 TF • 2068 CPU	Itanium2/ IA-64 • SUSE Linux SLES8/ 2.4 SMP • 0.166 TF • 32 CPU XEON/IA-32 • Red Hat Enterprise Ed. V3/2.4.24 • 192 CPU • .55 TF	Heterogeneous IA-32 Cluster • Debian/ sarge Linux • 1.7 TF • 960 CPU IBM Power3-II • IBM AIX 5.1L • 320 CPU	Intel Xeon • SUSE Linux 9.1 • 56 CPU	Intel Pentium 4/ IA-32 (Lonestar) • Redhat Linux 7.3 • 5.2 TF • 856 CPU UltraSPARC IV (Maverick) • Solaris 9 • 128 CPU
Online Storage	540 TB	230 TB	20 TB	200 TB	6 TB			50 TB
Archival Storage	6 PB	1.5 PB		2.4 PB	150 TB			2 PB
Networking (Gbps to hub)	30 Gbps LA	30 Gbps CHI	30 Gbps CHI	30 Gbps CHI	10 Gbps CHI	10 Gbps CHI	10 Gbps ATL	10 Gbps CHI
Database & Data Collections	YES	YES			YES	YES		YES
Instruments					YES		YES	
Visualization	YES		YES	YES				YES

http://www.teragrid.org/userinfo/guide_hardware_table.html (26 Nov 05)

PSC's Cray XT3: Architecture Overview

- 2,090 AMD Opteron processors
  - 2.4 GHz clock, each 4.8 GFlop peak
  - 10 TFlop theoretical peak aggregate
- Cray SeaStar interconnect
  - extremely high bandwidth: 6.5 GB/s sustained
  - configured as a 3-D torus
- Well-designed operating systems
  - Catamount OS on compute nodes prevents jitter, allows scalability
  - SUSE Linux on SIO nodes provides full functionality and connections to TeraGrid and I/O
- *Topology proven in the Cray T3E*
- 2 TB aggregate memory (1GB/proc)
- 200 TB disk storage (DDN)

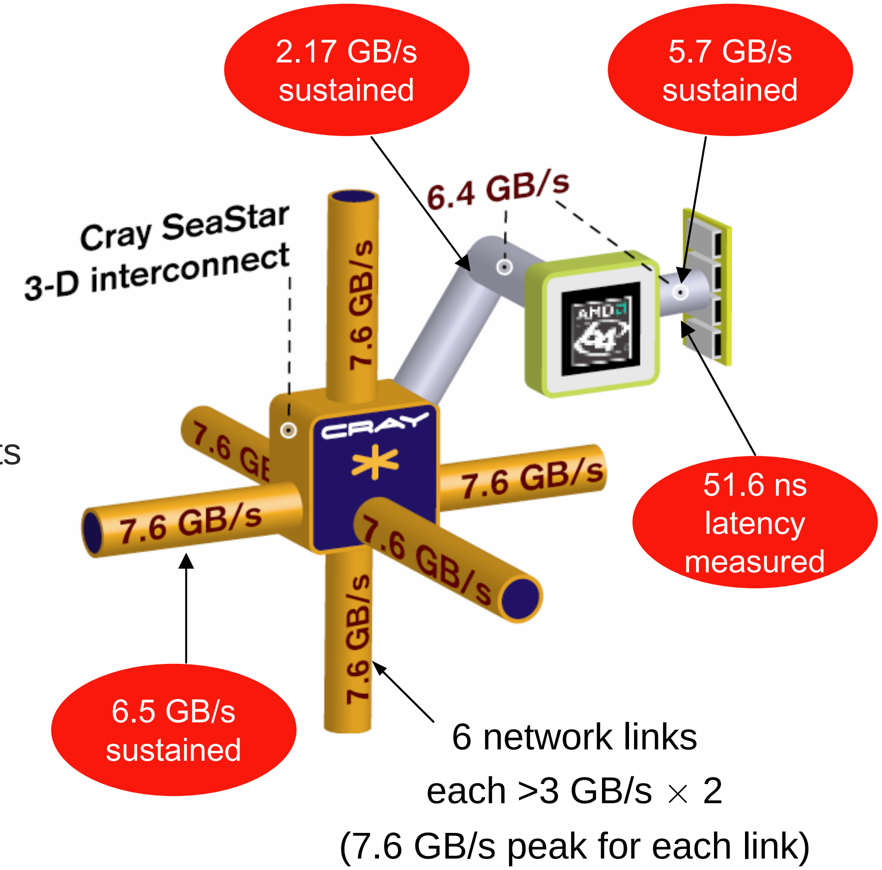

Image courtesy Jeff Brooks, Cray Inc.

Application-Enabling XT3 Architecture

- 3D torus direct-connected processor (DCP) architecture
- high bandwidth, low latency interconnect with embedded communications processing and routing
- OS system designed to support scalable applications
- integrated RAS system provides high reliability, allowing full applications to run to completion
- high-speed, global I/O
- standards-based programming environment
  - Fortran, C, C++, MPI
  - Convenient development on smaller systems

Ensuring Success for NSF TeraGrid Users

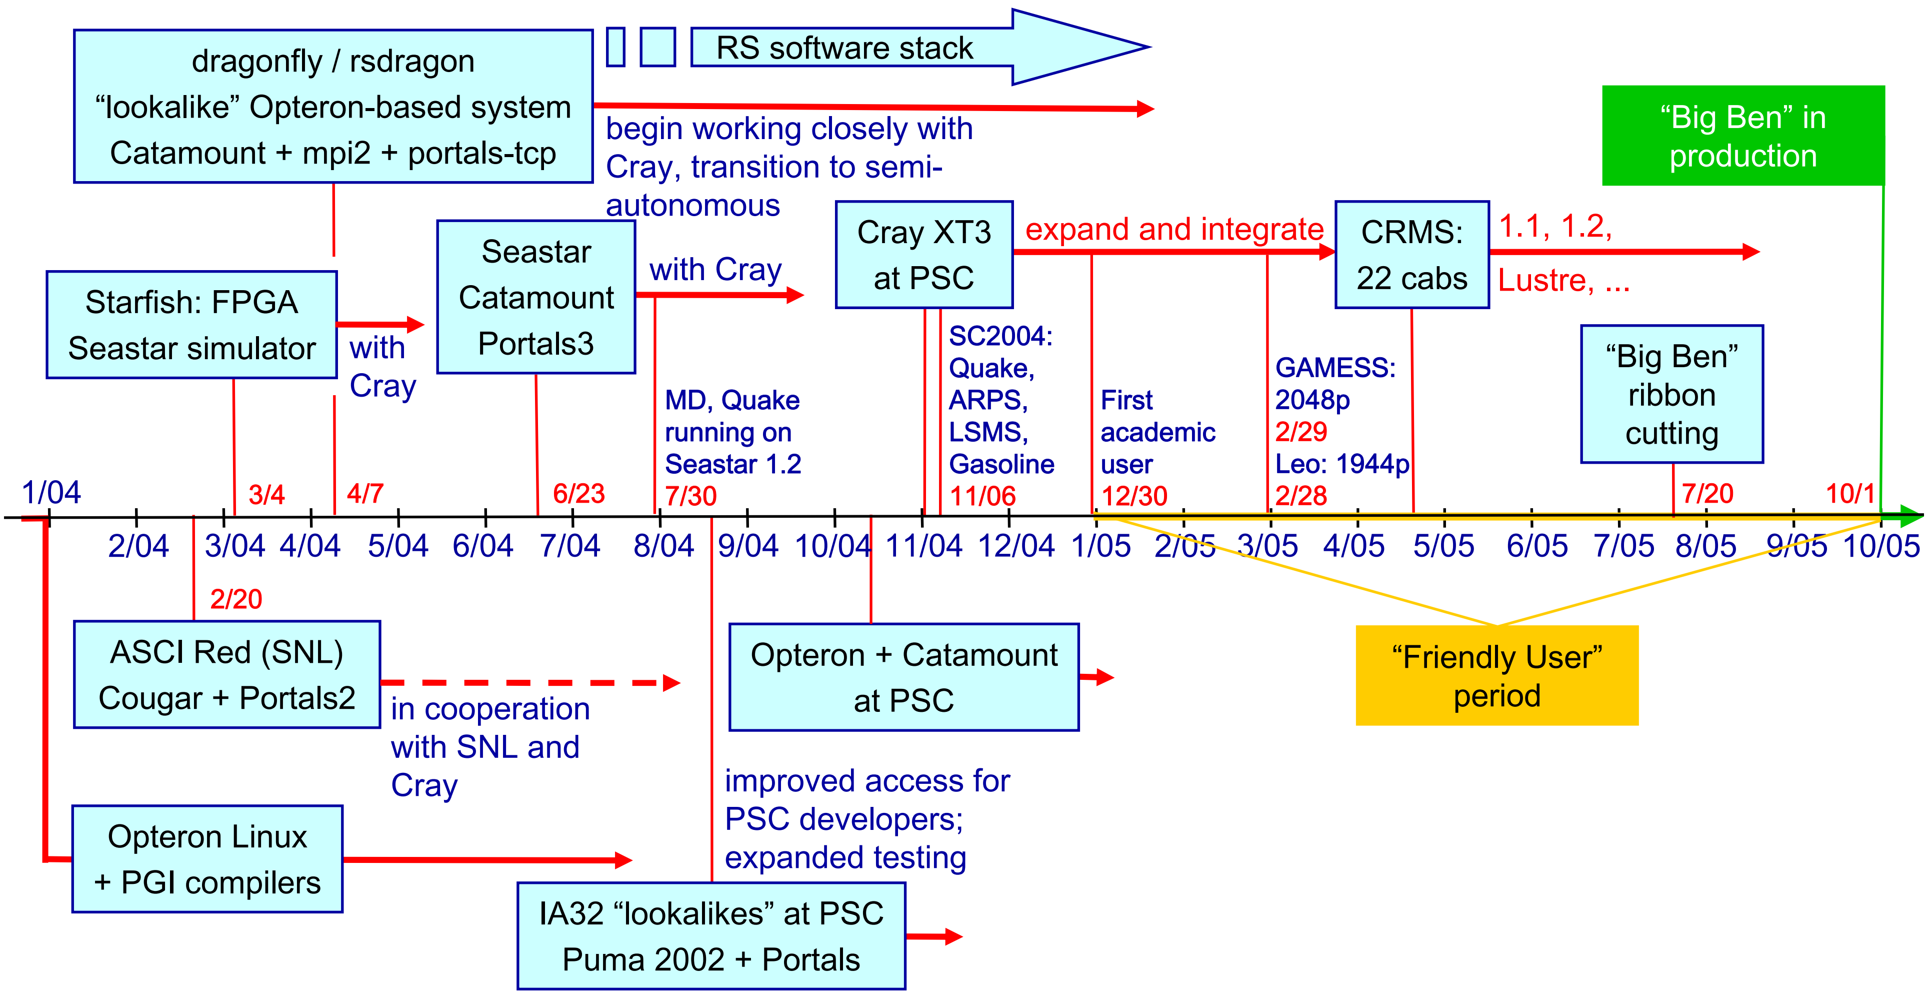

“Friendly User” Period

- Advance access to new system during build, leading to production
  - generate scientific output through the full system lifetime, maximizing return on investment
  - facilitate port and optimization of key applications
  - resolve potential issues as early as possible
- 44 groups spanning NSF directorates and divisions participated in BigBen’s “friendly user” period

“Friendly User” Highlights (1)

PI(s) / institution (users)	Involvement
Lee Pedersen, Keith Gubbins, Bob Duke / UNC Chapel Hill (2)	• AMBER/pmemd running and installed for production use • Performed long (>20ns) simulations of Factor VIIa (FVIIa) and the key coagulation complex Factor VIIa/soluble Tissue Factor (FVIIa/sTF) to permit an assessment of relative changes in the catalytic triad, the protein-protein interface region, the N-terminus and other regions thought to be critical for the enhancement of FVIIa activity by TF
Yong Duan / UC Davis (1)	• Performing biomolecular simulations to elucidate mechanisms of disease and gene expression
Steve Gottlieb / Indiana University (1)	• Ported MILC; commenced QCD calculations • In order to control sytematic errors, it is necessary to have multiple ensembles over a range of quark masses and lattice spacing. Currently running 40^3 x 96 lattices with $m_{up} = m_{down} = 0.1\ m_{strange}$ and lattice spacing = 0.9 fm
Jacobo Bielak / CMU (1)	• Ported new Quake to the XT3; performed end-to-end run, spanning mesh generation through visualization
Klaus Schulten / UIUC (7)	• Ported NAMD; performance tuning
L. V. Kale / UIUC (3)	• Ported Charm++, required for NAMD, OpenAtom

“Friendly User” Highlights (2)

PI(s) / institution (users)	Involvement
Paul Woodward and David Porter / Univ. of Minnesota (1)	• Developed new version of PPM development using one-sided MPI • Planned demos for IGRID, IEEE Vis, and SC05
Gary Glatzmeier / UC Santa Cruz (2)	• Ported Dynamo; began specification of new simulations
P. K. Yeung / Georgia Tech (3)	• Ported DNSmsp
Michael Klein / University of Pennsylvania (5)	• Ported and debugged CPMD • CPMD and QM/MM calculations with applications in biophysics, superacids, and surface science
Kathy Yelick / UC Berkeley (9)	• Ported UPC compiler
Al Malony / Univ. of Oregon (3)	• Ported the TAU performance tool
George Karniadakis / Brown Univ. (3)	• Ported NekTar
Tiziana di Matteo / CMU (1)	• Began astrophysical simulations using Gadget

“Friendly User” Highlights (3)

PI(s) / institution (users)	Involvement
Mordecai-Marc Mac Low / Am. Mus. of Nat. History (1)	• Ported ZEUS-MP (with Roberto Gomez, PSC)
John Mellor-Crummey / Rice U. (2)	• Developing Co-Array Fortran
Tom Quinn / U. of Washington (2)	• Cosmological simulations; development of Gasoline and Ntropy
Sadaf Alam, Jeff Kuehn, Patrick Worley, and Jeffrey Vetter / ORNL (4)	• ORNL collaboration: cross-system comparisons (CRMS vs. Dev Harness, Lustre, etc.), performance measurement
Martyn Guest / Daresbury (1)	• Ported and benchmarked chemistry applications
Andrew Jones / Manchester (1)	• Ported and benchmarked a data mining application for reporting at CUG and possible procurement decisions.

XT3 “Friendly Users” Represent PSC’s Production Workload

Normalized usage on PSC Platforms Feb04-Jan05

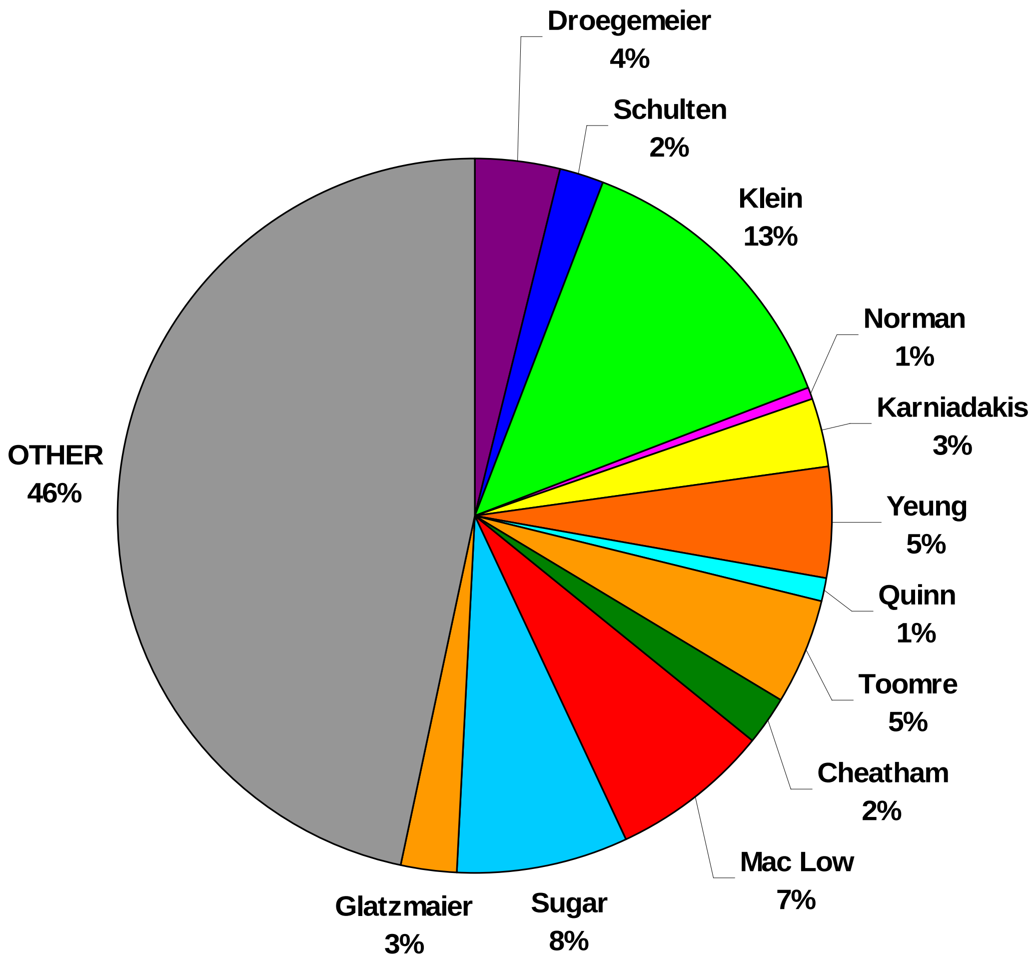

- The users with which PSC is working constitute 54% of normalized usage on PSC platforms over the last year.
- Examples:

Klein	25,563,887
Sugar	14,831,603
Mac Low	13,402,569
Yeung	9,513,116
Toomre	9,136,520
Droegemeier	7,243,520
Karniadakis	5,787,768
Glatzmaier	4,829,447
Cheatham	4,266,049
Total	191,928,345

- *Friendly users account for >50% of production utilization.*
- *Applications they develop, e.g. NAMD, impact an even broader user base.*

Cray XT3 Workshops at PSC

- ***Focus on porting and optimizing users' applications***
- August 9-12, 2005
  - 4-day workshop: ~1.5 days of instruction, ~2.5 days of time for work on applications, tour
  - 25 users (>13 institutions) + 13 PSC staff attended
  - *NAMD issue resolved; now runs stably (also ported Tcl to compute nodes)*
  - *Performance issue in DNSmsp identified and resolved, dramatically increasing performance*
  - *Cactus built; ASH debugged*
  - *shmem incorporated into Co-Array Fortran translator*

Luiz DeRose providing expert insights into performance tuning with Apprentice^2

- October 18-21, 2005
  - 16 users (5 institutions) + 2 PSC staff attended; same format as August
  - "Yesterday I improved the code 30%, and today 5-10% more." – Jinchen Liu, Pitt Chem. Eng, speaking of his Monte Carlo code to simulate water in carbon nanotubes
  - "Very great lecture series." – Pengwang Zhai, TAMU
  - CPMD debugged
  - averaged >300 jobs/day for the last 3 days of the workshop
- Next workshop: August 21-24, 2006

BigBen Allocations: Spanning Directorates and Disciplines (1)

- **Robert L. Sugar** [**1,891,500**] (UC Santa Barbara; MPS/PHY): Lattice Gauge Theory on MIMD Parallel Computers
- **Michael L. Klein** [**1,500,000**] (University of Pennsylvania; MPS/CHE): Terascale Quantum Simulations for Chemical Biology
- **Klaus J. Schulten** [**1,275,000**] (UIUC; BIO/MCB): Simulations of Supramolecular Biological Structures
- **Gregory Voth** [**1,000,000**] (University of Utah; MPS/CHE): Computational Studies of Electron Transfer at Water-Metal Interfaces and Intermolecular Hydride Transfer in Condensed Phases
- **Annick Pouquet** [**950,000**] (NCAR; MPS/PHY): Dual Imprint of Large and Small Scale Flows on Turbulence Scaling
- **Juri Toomre** [**850,000**] (University of Colorado, Boulder; MPS/PHY): Coupling of Turbulent Compressible Convection with Rotation
- **Yong Duan** [**500,000**] (UC Davis; BIO/MCB): Toward Understanding Amyloidosis: All-Atom Molecular Dynamics Simulations of Amyloid Protofibril Formation and Inhibition
- **Thomas Quinn** [**500,000**] (University of Washington; MPS/AST): Large Scale Structure and Clusters of Galaxies
- **Thomas Cheatham** [**400,000**] (University of Utah; CIE/CDA): Insight into Biomolecular Structure, Dynamics, Interactions and Energetic from Simulation
- **Pui-Kuen Yeung** [**300,000**] (Georgia Tech; ENG/CTS): High-Resolution Computations for Studies of Scaling in Turbulence and Turbulent Mixing

BigBen Allocations: Spanning Directorates and Disciplines (2)

- **Pui-Kuen Yeung** [**300,000**] (Georgia Tech; ENG/CTS):
  High-Resolution Computations for Studies of Scaling in Turbulence and Turbulent Mixing
- **George Em Karniadakis** [**295,000**] (Brown University; ENG/CTS):
  Hybrid Spectral Element Algorithms: Parallel Simulations of Turbulence in Complex Geometries
- **Okelsii Askimentiev** [**273,000**] (University of Illinois at Urbana-Champaign; BIO/MCB):
  Simulations of Nanoscale Biomolecular Systems
- **Jacobo Bielak** [**210,000**] (Carnegie Mellon University; GEO/EAR):
  High Resolution Forward and Inverse Earthquake Modeling on Distibuted Terascale Systems
- **Paul Woodward** [**200,000** + **10,000**] (University of Minnesota; MPS/AST):
  Simulation of Compressible Turbulent Astrophysical Flows
- **Hong Im** [**200,000**] (Sandia National Laboratories; ENG/CTS):
  High-Fidelity Direct Simulations of Turbulent Nonpremixed Flame Extinction by Water Droplets
- **Yang Wang** [**200,000**] (PSC; MPS/DMR):
  Ab Initio Studies of the Electronic and Magnetic Properties of Nanoparticles
- **Michael Norman** [**200,000**] (University of California, San Diego; MPS/AST):
  Testing the Condordance Model of Cosmological Structure Formation
- **Lee Pedersen** [**199,757**] (Univ. of North Carolina at Chapel Hill; BIO/MCB):
  Development of Solvated 3D Structures of VKD Proteins
- **Ravi Radhakrishnan** [**150,000**] (Cornell University; BIO/MCB):
  Theoretical Investigations of DNA Polymerases

BigBen Allocations: Spanning Directorates and Disciplines (3)

- **Jeremiah Ostriker** [**135,000**] (Princeton; MPS/AST):
  The Structure of Dark Matter Halos and the LCDM Model
- **Keh-Fei Liu** [**120,000**] (U. of Kentucky; MPS/PHY):
  Lattice Gauge Calculation of Hadronic Physics
- **Maria G. Kurnikova** [**110,000**] (CMU; MPS/CHE): Molecular Dynamic Simulation of the Interaction of the Adapter with the Genetically Modified Channel
- **Tiziana Di Matteo** [**80,000**] (Carnegie Mellon University; MPS/AST):
  Simulating the Joint Evolution of Black Holes and Galaxies)
- **Mordecai-Mark Mac Low [51,600**] (American Museum of Natural History; MPS/AST):
  Formation of Stars and Stellar Clusters in the Turbulent Interstellar Medium
- **Gary A Glatzmaier [50,000**] (UC Santa Cruz; MPS/AST):
  Convection and Magnetic Field Generation in Planets
- **Charles Musgrave [50,000]** (Stanford University; ENG/CTS):
  Quantum Chemical Investigation of Engineering Applications of Transition Metal Complexes
- **Bernard Montgomery Pettitt [50,000**] (Univ. of Houston; MPS/CHE):
  Salt Effects in Solutions of Peptides and Nucleic Acids
- **Ali Uzun** [**40,000** + **10,000**] (Florida State University; ENG/CTS):
  Aeroacoustics of Wind Turbine Blades
- **Pablo Laguna** [**30,000** + **10,000**] (Penn State Univ; MPS/PHY):
  Black Hole Evolutions

BigBen Allocations: Spanning Directorates and Disciplines (4)

10,000 SU starter allocations (DACs) typically precede LRAC or MRAC requests

- **Kelvin K. Droegemeier** [**25,000**] (U. of Oklahoma; GEO/ATM): Center for Analysis and Prediction of Storms
- **Allan Snavely** [**21,376**] (SDSC; CIE/ASC): A Framework for Performance Modeling and Prediction
- **Alexander MacKerrell** [**20,000**] (University of Maryland; BIO/MCB): Atomic Detail Investigations of the Structural and Dynamic Properties of Biological Systems
- **Yiannis Kaznessis** [**15,000**] (University of Minnesota; BIO/MCB): Biomolecular Interaction and Gene Network Simulations
- **Dave Yuen** [**15,000**] (University of Minnesota; BIO/MCB): Modeling Blood Clotting with the Fluid Particle Model
- **Kosuke Imai [10,000]** (Princeton University; SBE/SES): Bayesian Dynamic Estimation of Legislators' Ideal Points in the U.S. House and Senate
- **Eric P. Chassignet** [**10,000**] (University of Miami; GEO/EAR): Interamural Modeling of the North Atlantic Ocean with HYCOM
- **Yong Duan** [**10,000**] (UC Davis; BIO/MCB): Deciphering Histone Code by Computer Simulations
- **Katherine Yelick** [**10,000**] (UC Berkeley; CIE/CCR): Portability and Scaling of Global Address Space Languages
- **Kenneth Jordan [10,000]** (Univ of Pittsburgh: MPS/CHE): Theoretical Studies of Water Clusters with Excess Protons and Electrons and of Biomolecule Energy Landscapes

BigBen Allocations: Spanning Directorates and Disciplines (5)

- **Kathleen Carley [10,000]** (Carnegie Mellon University; SBE/SES):
  Large-scale Agent-based Modeling of Weaponized Biological Attacks on a Realistic Population within the Background of Naturally-Occurring Diseases and Its Validation
- **Michael Lipson** [**10,000**] (Cornell University of Michigan; MPS/PHY): Evolutionary Photonics
- **Maria G. Kurnikova** [**10,000**] (CMU; MPS/CHE): Molecular Dynamic Simulation of the Interaction of the Adapter with the Genetically Modified Channel
- **Brian D. Athey** [**10,000**] (University of Michigan; BIO/MCB):
  Benchmarking of Computational Medicine and Biology Research Applications
- **Donald Wuebbles** [**10,000**] (UIUC; GEO/ATM): Application of Multiprocessing NCSA Platforms to Study Global Change Issues Relating to Climate and Ozone
- **Junwoo Lim** [**10,000**] (PSC; ENG/CTS):
  Investigation on the Near-wall Behavior of Stratified Turbulence
- **Henry Schaefer [10,000]** (University of Georgia: MPS/CHE):
  The Next Generation: A Rigorous Quantum Mechanical Study of a Nucleotide-Nucleotide Pair and its Anion, with Complete Structural Optimizations
- **Sameer Shende [100]** (University of Oregon; CIE/ASC): Performance Tool Scalability
- **Juri Toomre** [**10,000**] (University of Michigan: MPS/AST): Coupling of Turbulent Compressible Solar Convection with Rotation, Shear, and Magnetic Fields
- **Jesus Izaguirre** [**10,000**] (University of Notre Dame; CIE/ASC):
  Parallel Multigrid Summation

BigBen Allocations: Spanning Directorates and Disciplines (6)

- **Michael Norman** [**10,000**] (UC San Diego; MPS/AST):
  Benchmarking Enzo and ZEUS-MP
- **Junwoo Lim** [**10,000**] (PSC; ENG/CTS): Terascale Simulations of Boltzman Equation for Turbulence: An Initial Performance Test
- **Vijay Pande [10,000]** (Stanford University: BIO/MCB):
  Molecular Dynamics Simulation of Membrane Fusion Mechanism
- **Andrew Rappe [10,000]** (University of Pennsylvania: MPS/CHE):
  Classical Potential Model for Perovskite Complex Oxides
- **Thomas Cheatham** [**10,000**] (University of Utah; CIE/ASC):
  Evaluation of AMBER 9 for Nucleic Acid Simulation
- **Adrian Roitberg** [**10,000**] (University of Florida: BIO/MCB):
  Biomolecular Simulations
- **David Lambeth [10,000]** (Carnegie Mellon University; MPS/DMR):
  Electronic Structure and Charge Transport in Polythiophene
- **Carlos Lousto [10,000]** (University of Texas, Brownsville; MPS/PHY):
  Waveforms from Orbiting Black Holes
- **Gabrielle Allen [10,000]** (Louisiana State University; CIE/ASC):
  Performance Profiling and Optimization of Community Frameworks
- **Linda Jen-Jacobson** [**10,000**] (University of Pittsburgh; BIO/MCB):
  Molecular Dynamics Simulations of Site-Specific Protein-DNA Interactions

BigBen Allocations: Spanning Directorates and Disciplines (7)

- **Allen Taflove** [**10,000**] (Northwestern University; GEO/EAR):
  Global ELF Propagation about the Earth
- **William Lytton** [**10,000**] (SUNY Health Science Center at Brooklyn; BIO/IBN):
  Epileptiform Patterns Vary with Network Size
- **Douglas James** [**10,000**] (Carnegie Mellon University; CIE/IRI):
  Supercomputing Kernels for Interactive (Surgical) Simulation using Reduced Deformable Models
- **Daniel Wang** [**10,000**] (University of Massachusetts at Amherst; MPS/AST):
  Hot Gas in Galactic Bulges
- **Jayathi Murthy** [**10,000**] (Purdue University; MPS/DMR):
  A Highly Scalable Simulation Model for Atomistic Calculation of Thermal Properties of Silicon
- **Glenn David Starkman and David Spergel** [**10,000**] (Case Western Reserve University and Princeton; MPS/AST): Finding the Topology of the Universe
- **Harold Scheraga** [**10,000**] (Cornell University; MPS/CHE):
  Development and Application of a Hierarchical Protocol for *Ab Initio* Prediction of Protein Structure
- **Michael Hines** [**10,000**] (Yale University; BIO/MCB):
  Parallel Network Simulation with NEURON
- **Meng Cui** [**10,000**] (Mount Sinai School of Medicine; BIO/MCB):
  Mechanism of Agonism-Antagonism of Sweet Taste Receptors
- **Eui Joong Kim** [**10,000**] (Duke University; ENG/MSS):
  Development of Parallel Mortar Contact Algorithm

BigBen Allocations: Spanning Directorates and Disciplines (8)

- **Zdzislaw Meglicki** [**5,000**] (Indiana University; ENG/ECS):
  Computational Nanophotonics with Chombo
- **Roberto Gomez** [**10,000**] (Pittsburgh Supercomputing Center; MPS/PHY):
  Code development for large scale characteristic simulations
- **Robert Harkness** [**10,000**] (San Diego Supercomputer Center; CIE/ASC):
  XT3 Benchmarking
- **Shanhui Fan** [**3,840**] (Stanford University; MPS/DMR):
  Computational Micro and Nano-Photonics
- **Aytekin Gel** [**10,000**] (Aeolus Research; ENG/CTS):
  MFIX Research

Applications Run at Scale on the XT3

application	domain	nodes
HPCC	HPC Challenge benchmarks	5208
NAMD 2.5	Molecular dynamics	5000
Charm++	Parallel C++ (used by NAMD)	5000
WRF	Weather Research & Forecasting	4096
MILC	QCD	4096
HOMME	Atmospheric dynamical core	4096
MPQC	Massively Parallel Quantum Chemistry	2067
LSMS 2.0 LSMS 1.6	Materials science / electronic structure	2048 2048
HYCOM	Ocean modeling	2048
GAMESS	Quantum chemistry	2048
PSCC	Parallel Spectral Channel Code	2048
sPPM	PPM benchmark	2048
Quake	Earthquake modeling	2048
Gasoline	N-body astrophysics	2048
CPMD 3.9.1	Car-Parrinello Mol. Dynamics	2048

application	domain	nodes
DNSmsp	Direct Numerical Simulation	2048
Dynamo	QM/MM biochemistry	2048
ABINIT	Electronic struture	2048
CHARMM	Molecular dynamics	2048
PARATEC	Parallel Total Energy Code	2048
S3D	Turbulent combustion	2048
MHD	Magnetohydrodynamics	2048
PPM	Piecewise Parabolic Method	2048
OOCORE	Out-of-Core Solver	2048
Leo	Numerical relativity	1944
Gadget	Smoothed Particle Hydrodynamics	1800
AMBER	Molecular dynamics	1024
QChem	Quantum chemistry	1024
ZEUS-MP	Astrophysics	1000

Elucidating the Molecular Basis of Disease

April 2005

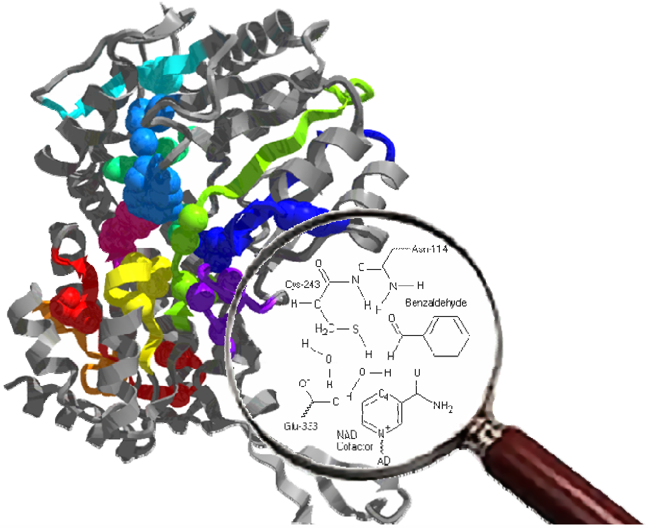

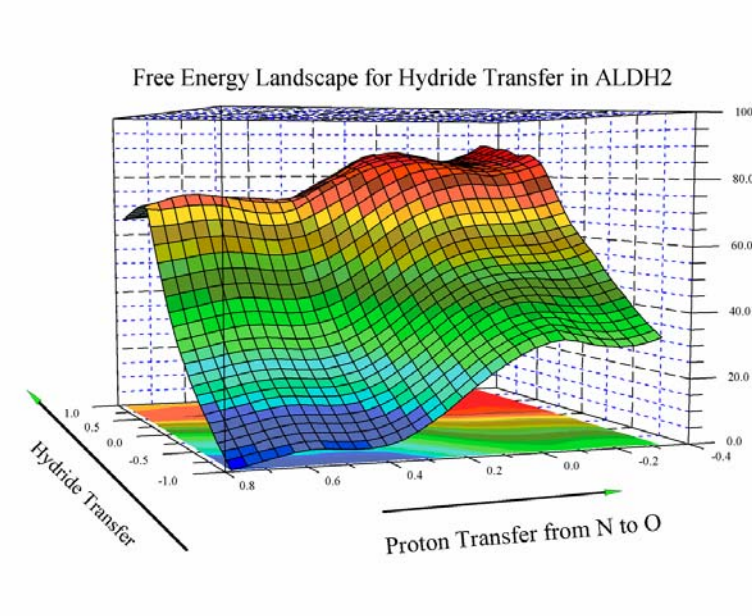

- Path-integral QM/MM simulations reveal the significance of quantum dynamical effects (proton tunneling) in Glutathione S-Transferase (GST) and Aldehyde Dehydrogenase (ALDH)
- Preliminary results support hypothesis that lack of intermediate stabilization is the molecular basis behind two different metabolic diseases, Hyperprolinemia Type II and Sjogren-Larsson Syndrome
- QM/MM simulations on coupled proton transfer-hydride transfer events in ALDH:
  - Dynamo 2.2
  - 36,738 atoms, 10ps, 10^4 timesteps
  - 900 processors: linear scaling
  - 9,450 CPU hours

Understanding the Formation of the Universe

- Tiziana Di Matteo (Carnegie Mellon University)
- Full hydrodymical (stars and gas and black holes as well as dark matter) cosmological simulation; 2×486^3 particles
- 5,000 kpc slices from the full 35,000 kpc simulation

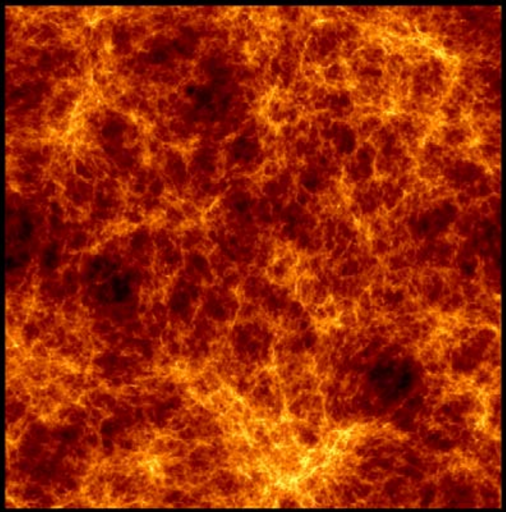

Redshift z=14
Very early universe, about 200 million years after the big bang

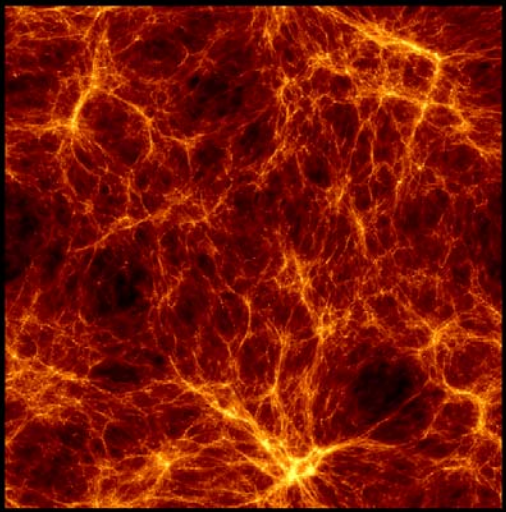

Redshift z=4
About 1.5 billion years old; about 1/10 present age

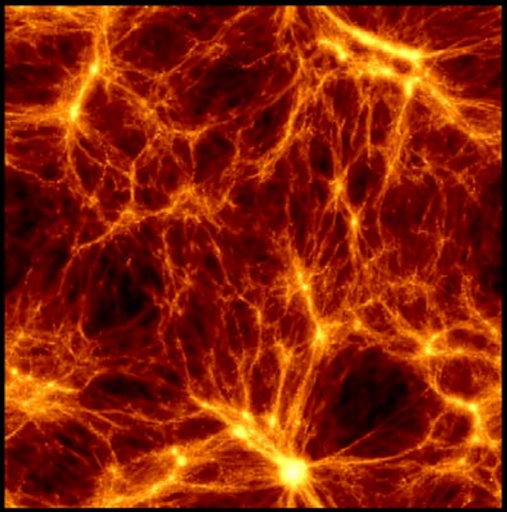

Redshift z=1.6
About 5 billion years old; what the universe looks like ~today

- Full hydro at this resolution reveals unprecedented structure in filaments, galaxies, and black hole formation

Understanding the Formation of the Universe (2)

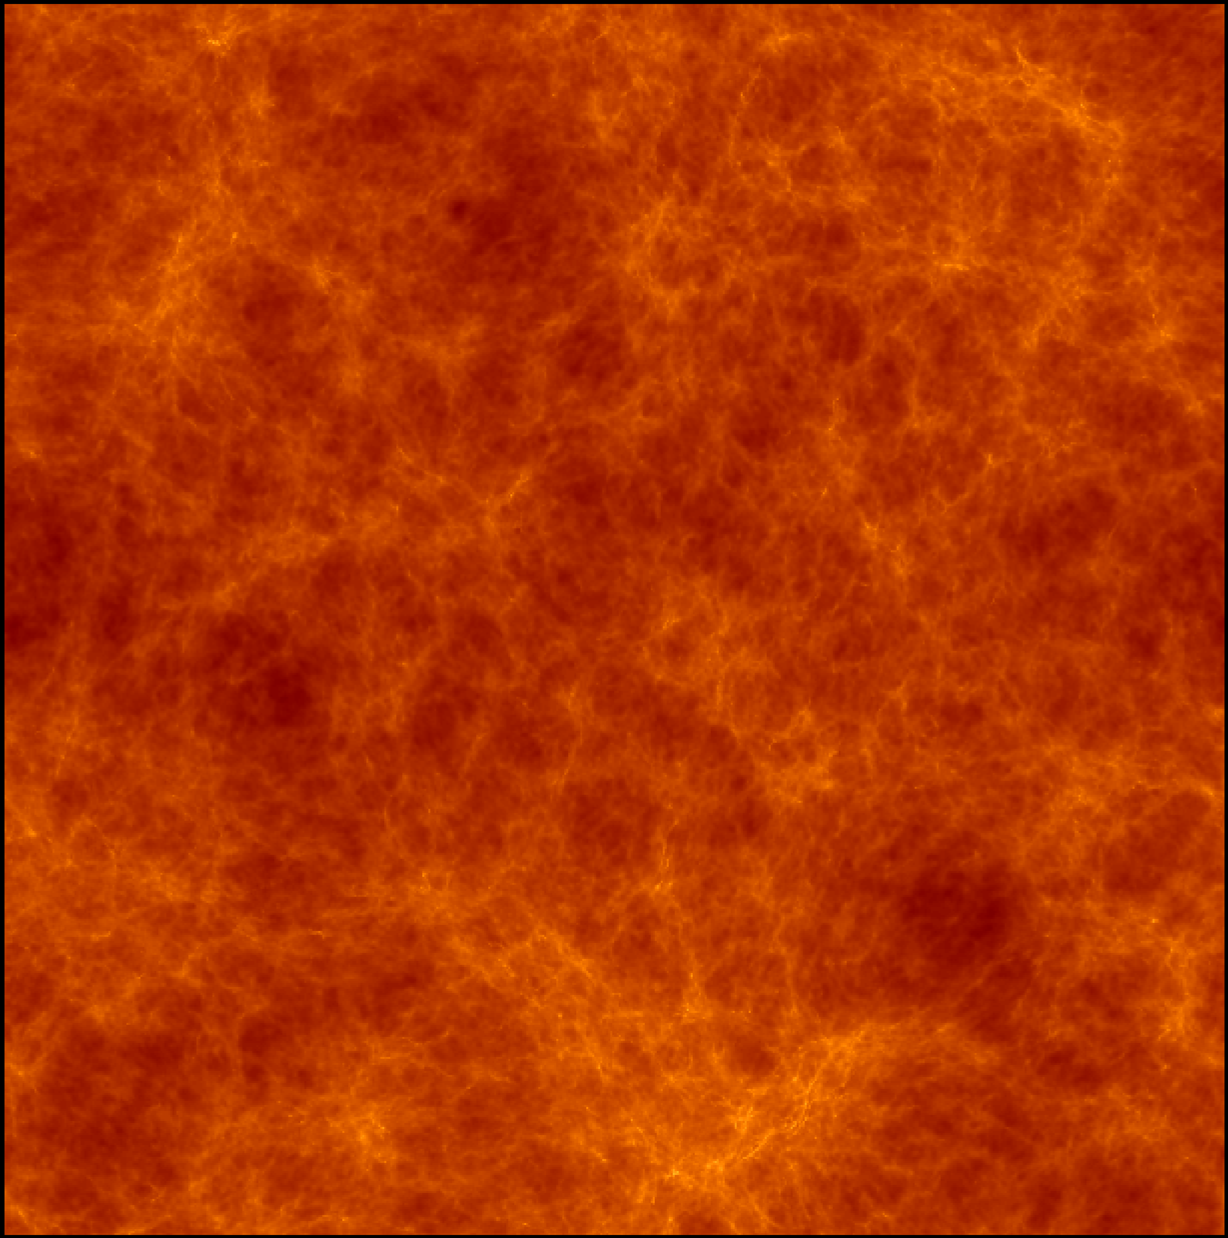

Ideal Point Estimation of Legislative Trends

- Kosuke Imai (Princeton) and Kevin Quinn (Harvard)
- Markov Chain Monte Carlo to infer probability distributions for legislators' voting patterns over time
- Prior work addresses Supreme Court decisions since 1953
  - 9 justices, 1 month on desktop
- Goal: model Congress, ~550 U.S. Senators and Representatives and 1600 decisions/year, over 40 years
- Parallelizing MCMC application for the XT3 (Deborah Weisser, PSC)
  - first use of massively parallel computing by statistical political science
  - invited paper to Journal of Computational and Graphical Statistics

Conservative vs. Liberal Trends of U.S. Supreme Court Justices: 1986-2002

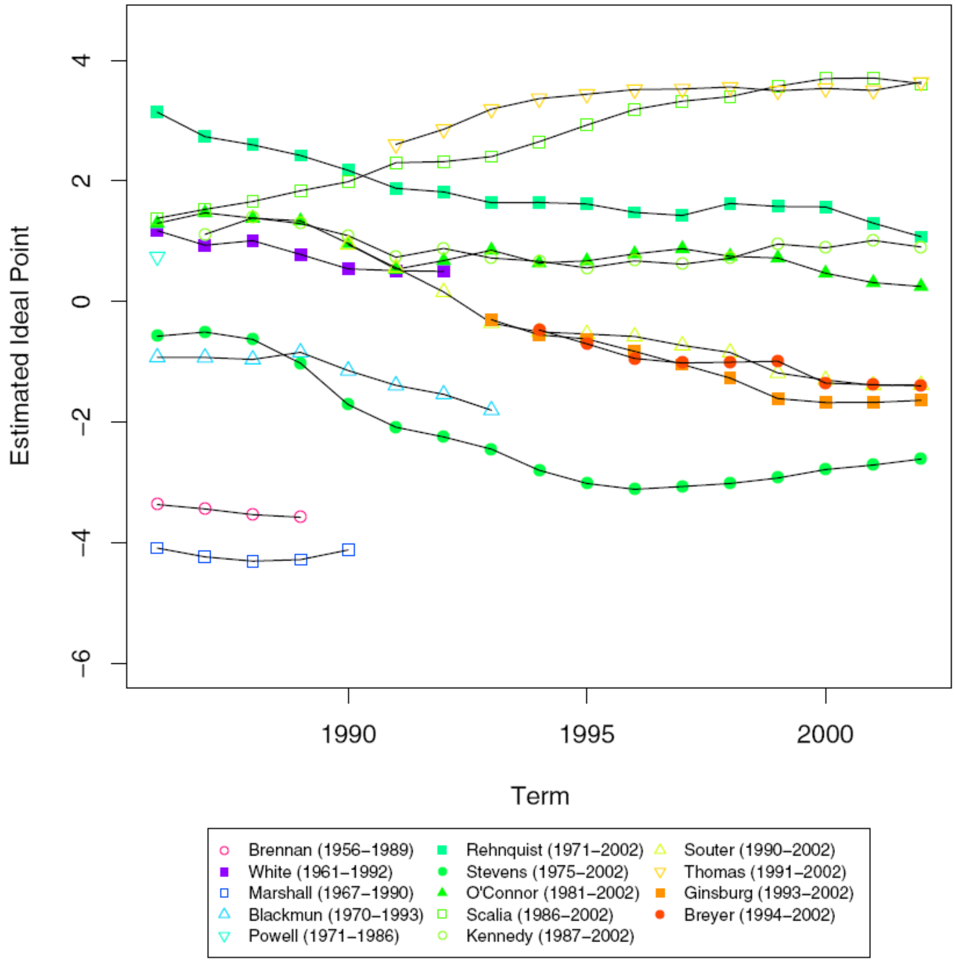

Steering CFD: Compressible Turbulent Astrophysical Flows

- Interactive steering and visualization of parallel PPM runs: Paul Woodward (U. Minnesota)
- PDIO: Portals-to-TCP routing. PDIO daemons aggregate hundreds of portals data streams into a configurable number of outgoing TCP streams. Heterogeneous Portals: both QK + Linux nodes.
- Utilizing the ETF network: 10GigE end-to-end. Benchmarked >1Gbps.

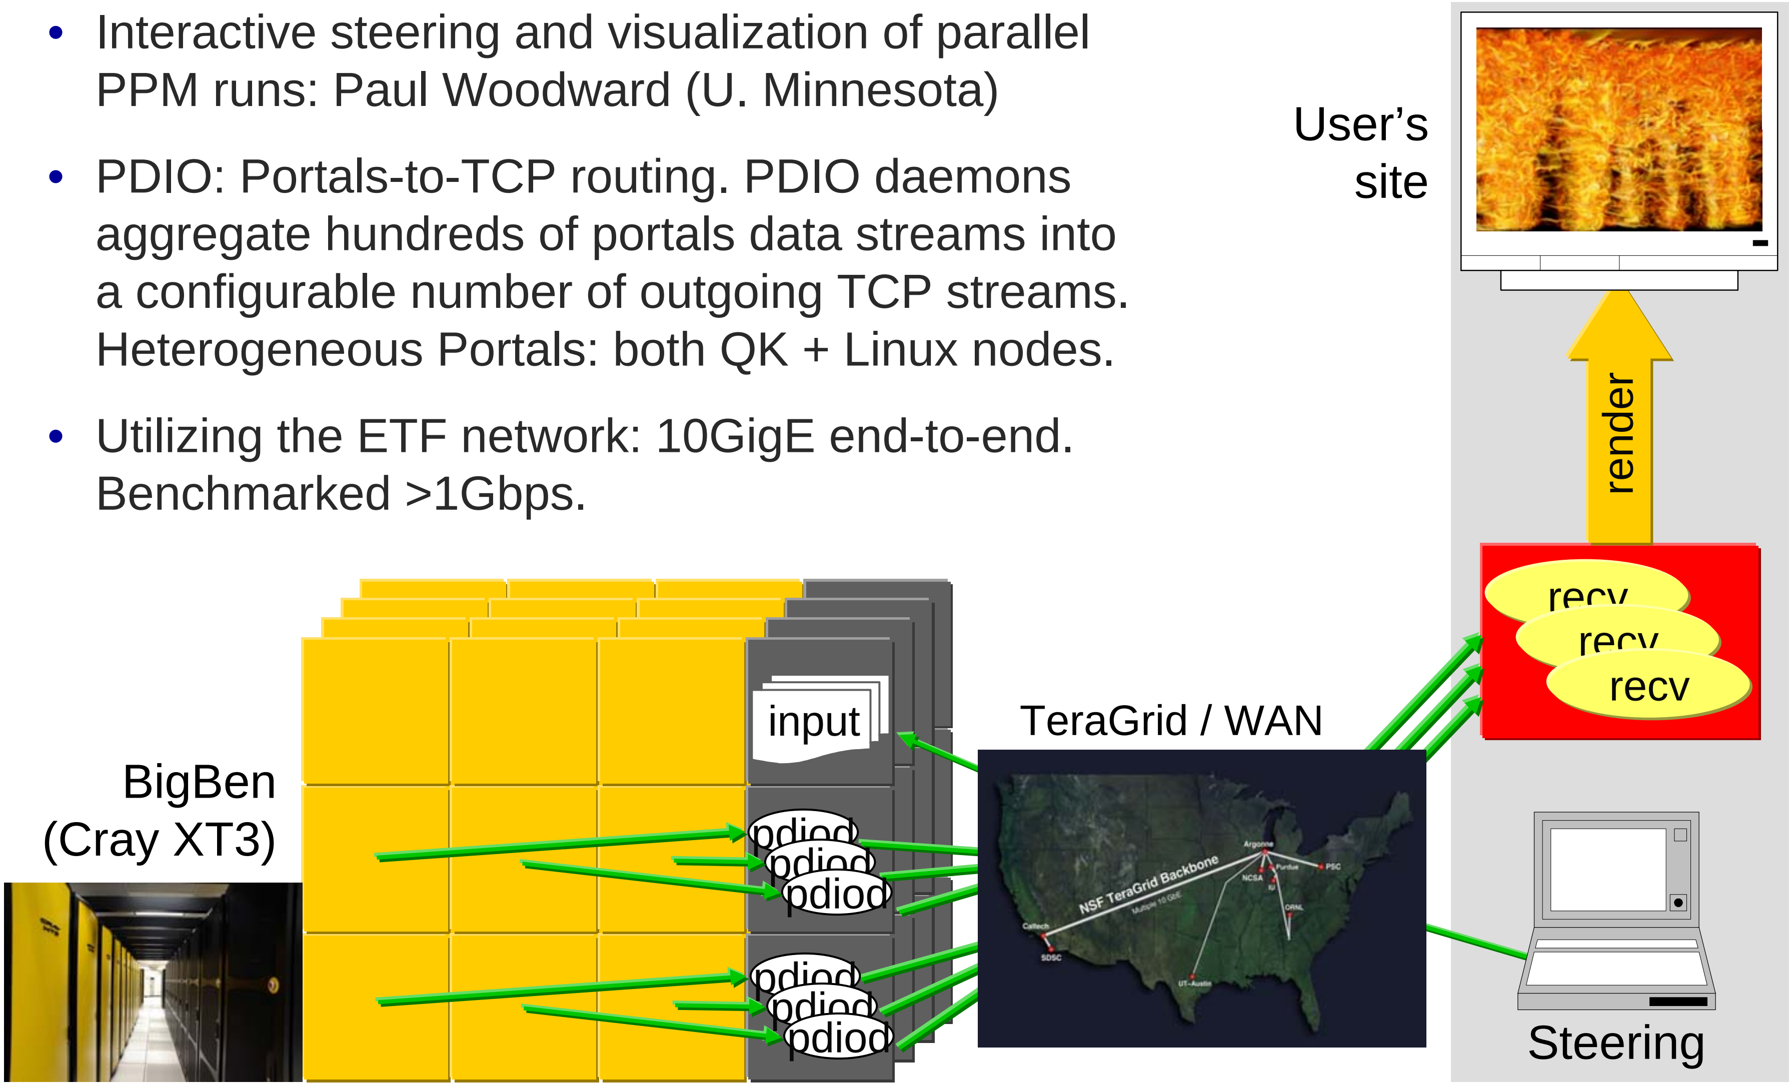

Probing Fundamental Physics: Quantum Chromodynamics

- Gauge configuration generation on $40^3 \times 96$ lattices with a lattice spacing of 0.09 fm and a light quark mass of one-tenth the strange quark mass.
- These configurations will significantly increase the accuracy of a wide range of calculations
- Very difficult run
  - 2003 – 8 July 8, 2005: first 708 time units: LeMieux, 2.5y
  - July 2005 – 5 May 2006: next 1296 time units: BigBen, 10m
  - expect to finish 3000 time units (500 stored configurations) in the next few weeks
- Increased scientific throughput due to high performance and excellent access
- This run will be used in many physics projects including light quark masses, pion and kaon decay constants, B and D meson decay constants, B and D meson semi-leptonic form factors, bottomonium and charmonium spectrum and determination of the strong coupling.

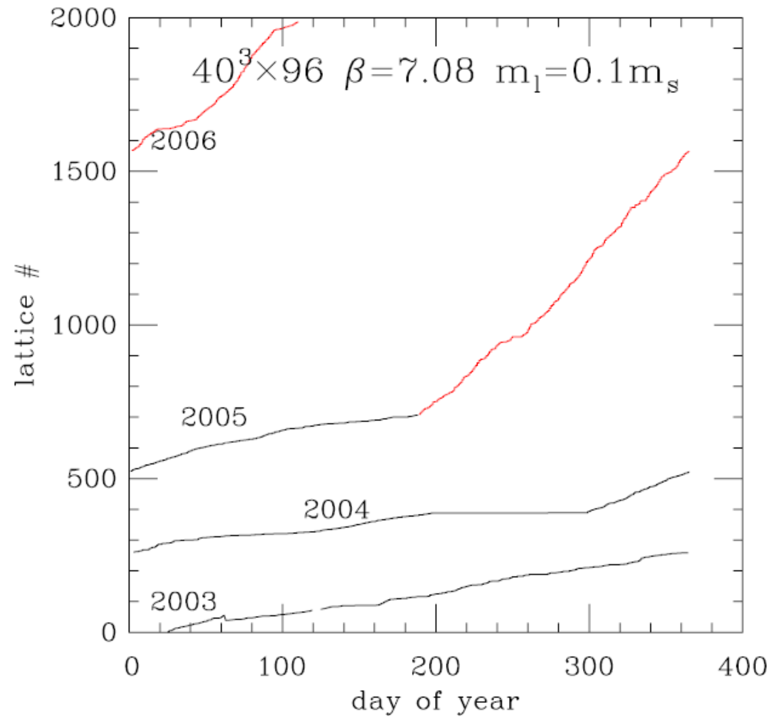

- Configurations to be made available through the International Lattice Data Group

Dynamics of Pollutants: Particle Motions in Stratified Turbulent Boundary Layer Flow

- PSCC (Parallel Spectral Channel Code), developed by Junwoo Lim at PSC, is a highly scalable parallel flow solver to enable more realistic simulations of turbulent boundary flow, including tracking the traces of massless particles under stably stratified conditions.
- Application: understanding the dynamics of pollutant particles in atmospheric boundary layers.

The Criticality of Bandwidth

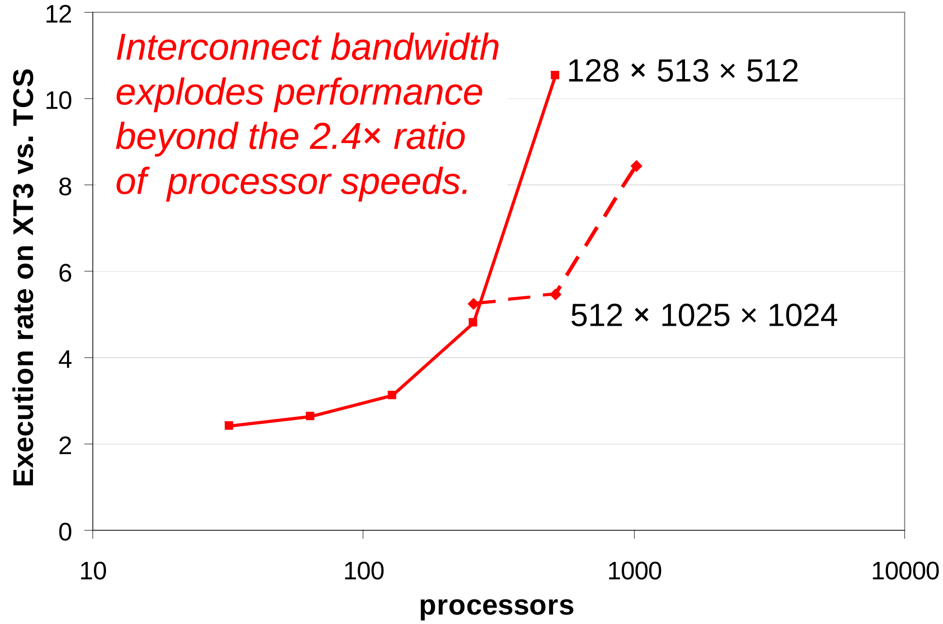

- This project is sponsored by NSF and is a part of the international collaboration research program between PSC and KISTI (Korea Institute of Science and Technology Information) Supercomputing center.

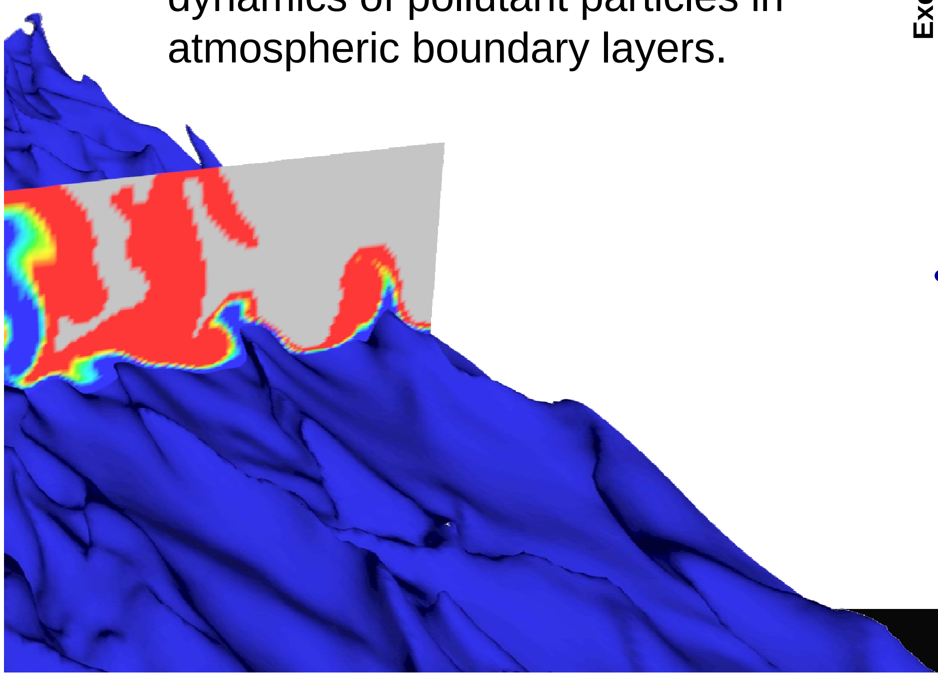

BigBen vs. LeMieux: SeaStar Bandwidth Dwarfs 2-Rail Quadrics

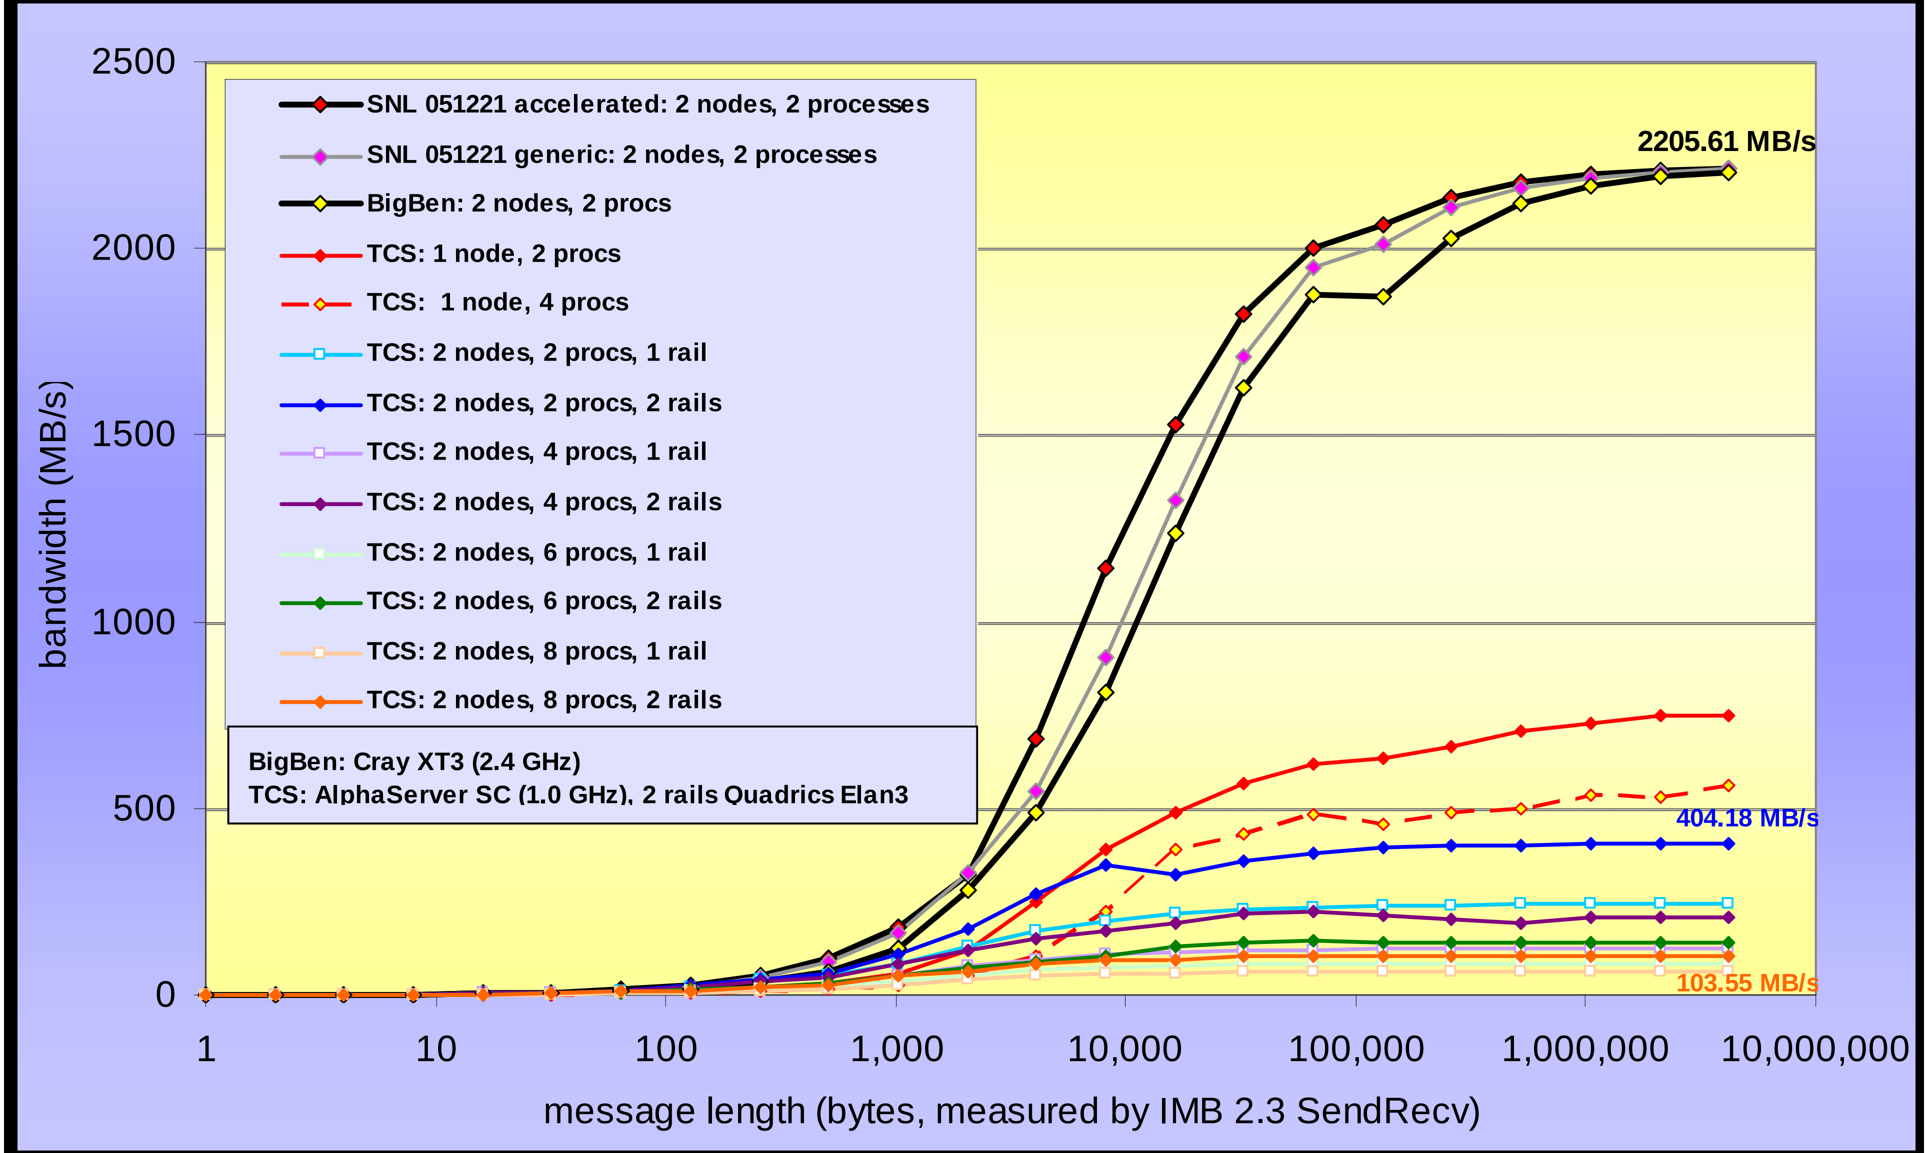

Cray XT3 Outperforms on HPCC

	procs	G-HPL	G-PTRANS	G-FFTE	G-Random Access	G-STREAM Triad	EP-STREAM Triad	EP-DGEMM	Random Ring BW	Random Ring Latency	HPL percent of peak
	count	TFlop/s	GB/s	GFlop/s	GUP/s	GB/s	GB/s	GFlop/s	GB/s	usec	percent
TCS baseline	3000	4.215	72.500	24.202	0.140	660.576	0.220	1.762	0.041	19.162	70.2%
Columbia baseline	2024	9.320	18.190	45.776	0.049	3998.493	1.976	6.236	0.123	6.98	71.9%
IBM BG/L baseline	2048	1.408	34.251	96.192	0.454	1484.583	0.725	0.905	0.021	4.98	49.1%
XT3 baseline	2048	7.713	256.383	425.660	0.348	9756.078	4.764	4.387	0.251	10.314	78.5%
XT3 optimized	2048	8.113 (1GB/p) 8.279 (2GB/p)	261.472	575.799	0.355	11856.712	5.789	4.387	0.251	9.743	82.5% 84.2%
XT3 optimized	5208	20.409	944.227	761.729	0.672	24707.000 (5000p)	4.660	4.412	0.206	9.200	81.6%

Source for Columbia and BG/L values: http://icl.cs.utk.edu/hpcc/hpcc_results.cgi, 28 Nov 05

Big Ben Performance Often Exceeds that of TCS by >2.4×

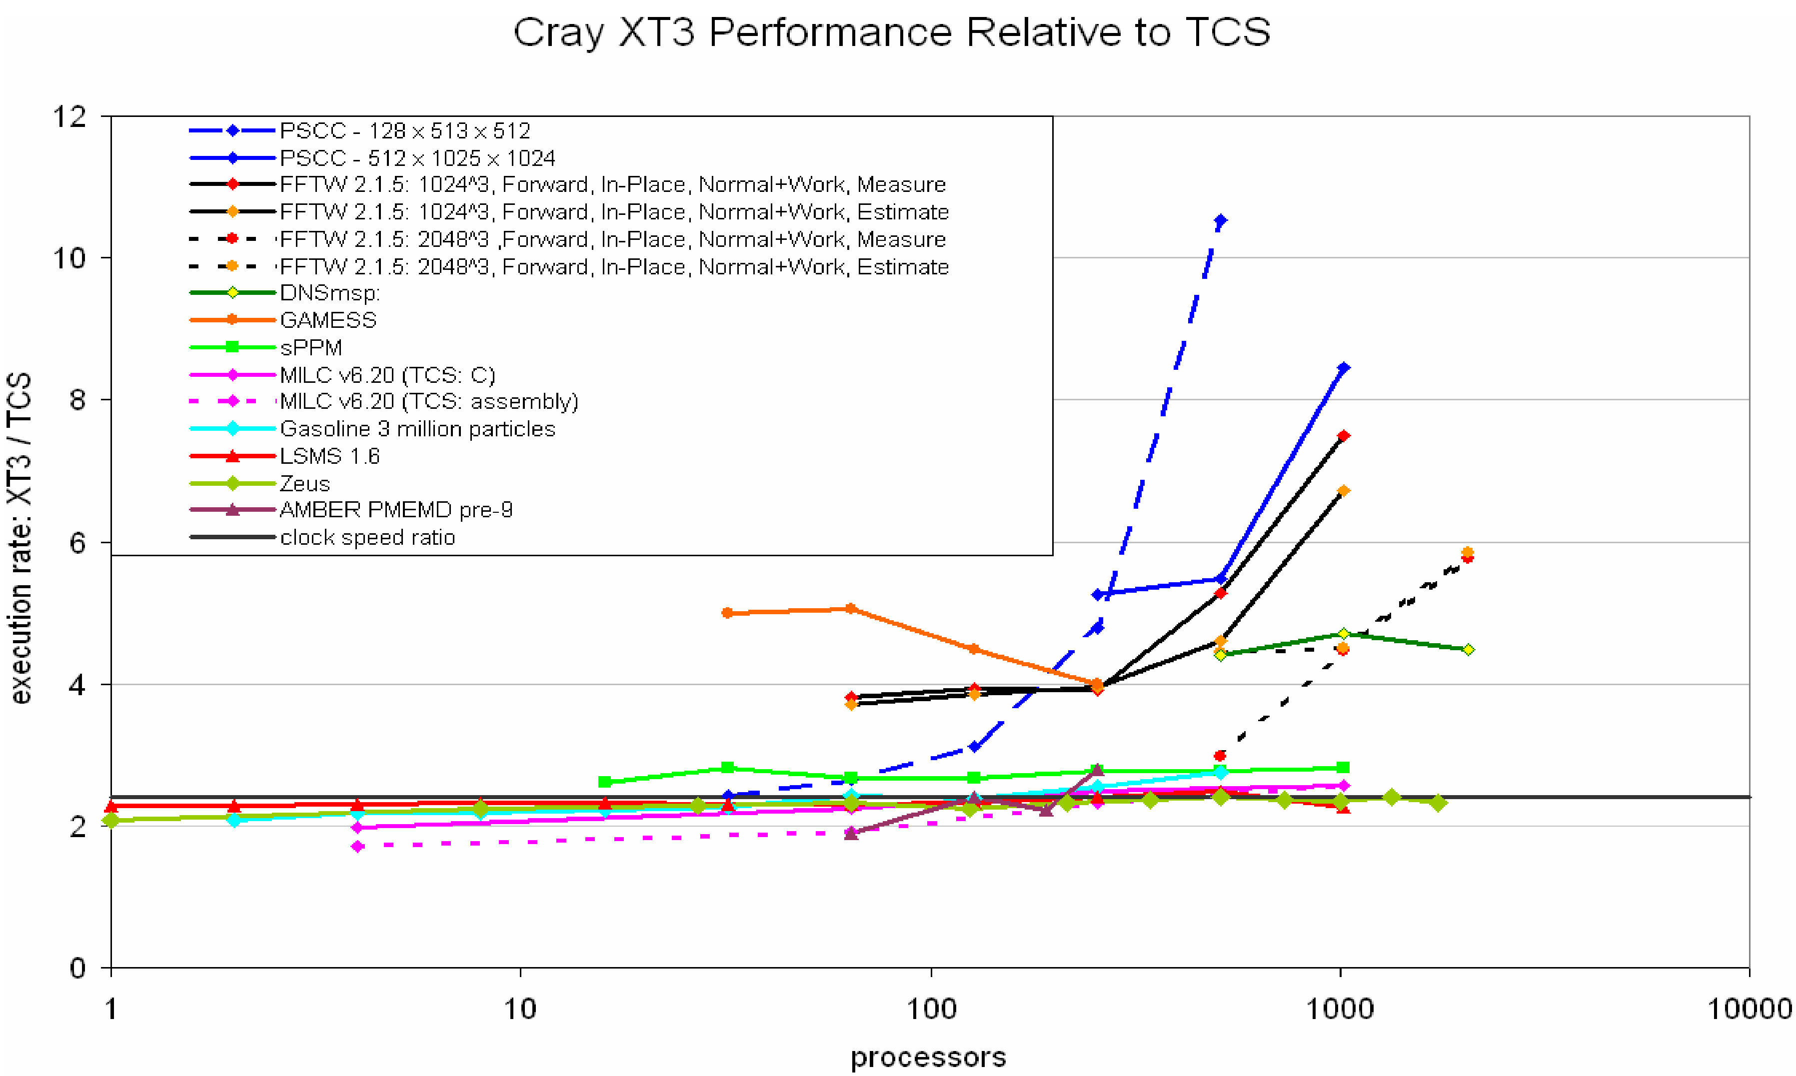

Scalable Million-Atom Molecular Dynamics with NAMD

- NAMD: Scalable simulation of large biomolecular systems; 2002 Gordon Bell performance prize
- 1,066,628 atoms (167,063 protein atoms + 299,855 waters)
- Each timestep at 5000p takes only 0.27s.
  Requires high-bandwidth interconnect and jitter-free OS (Catamount)

NAMD 2.6b1 ApoA1 (PME): 92,224 atoms

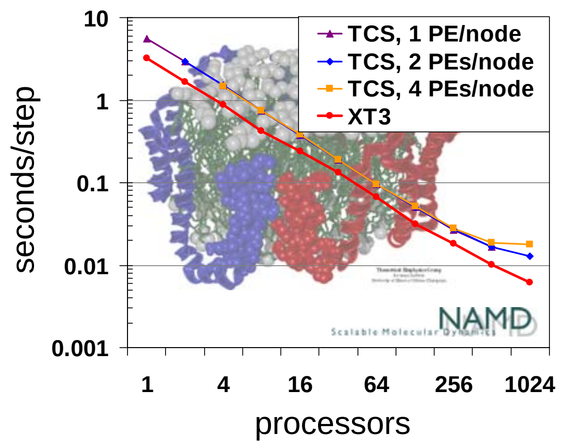

NAMD scales to at least 5000 processors on the Cray XT3

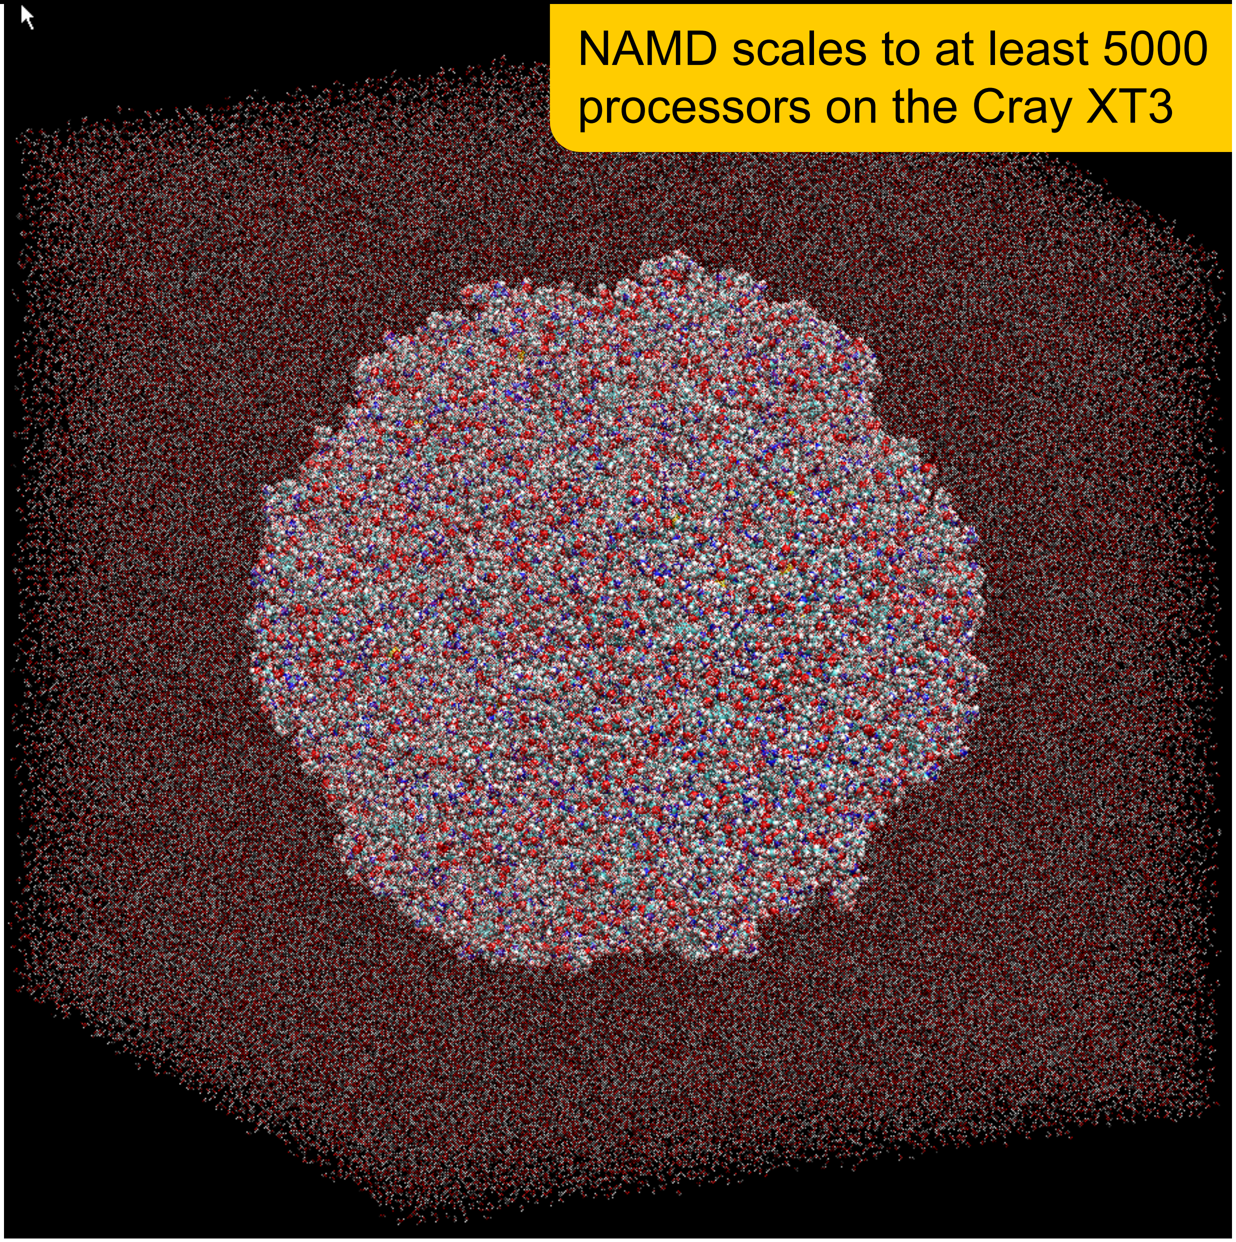

Theoretical Investigations of DNA Polymerase Mechanisms

- R. Venkatramani and R. Radhakrishnan, University of Pennsylvania
- Goal: provide a molecular mechanism for how high fidelity polymerases achieve their specificity in selecting nucleotides
- Solvated and neutralized BF-DNA-dNTP ternary complex. BF is shaped like a right hand (typical of high fidelity polymerses) with thumb, fingers and palm domains
- The enzyme is shown in a closed state with the DNA fragment in blue and the incoming nucleotide in red

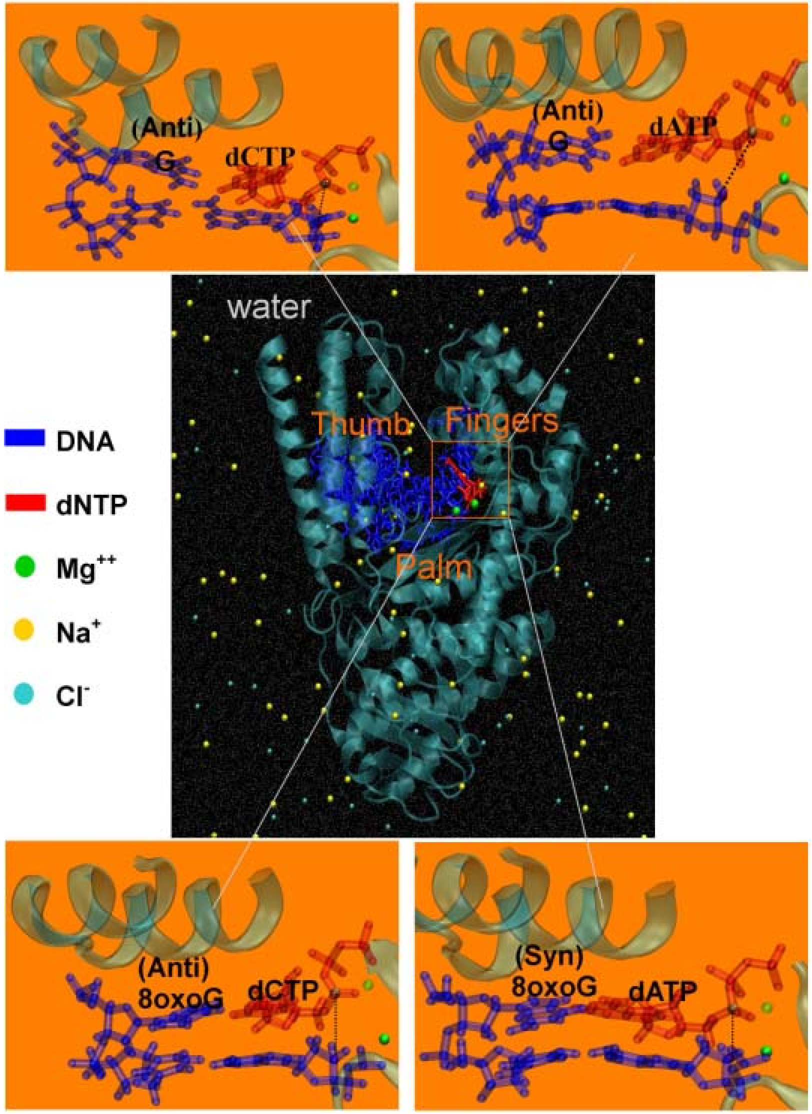

The Road to NAMD on the Cray XT3

NAMD

A highly scalable molecular dynamics code built upon the CHARM++ programming language.

Porting Issues

- PGI compiler issues
- MPI only interface
- Unsupported calls
- FFTW 2.1.5 port
- TCL 8.x porting
- Race conditions and other bugs

Performance Enhancements

- Using gnu malloc (-lgmalloc) – 20%
- Using –small_pages – 20%

NAMD Accounts for a Significant Portion of BigBen Usage

Prof. Mike Klein from Univ. of Penn.

Prof. Gregory Voth from the Univ. of Utah

Prof. Klaus Shulten from the Univ. of Illinois, Urbana

…

The Journey Continues

- Improving efficiency over large processor counts
- Process mapping
- Improving Charm++ on the XT3

Current Performance

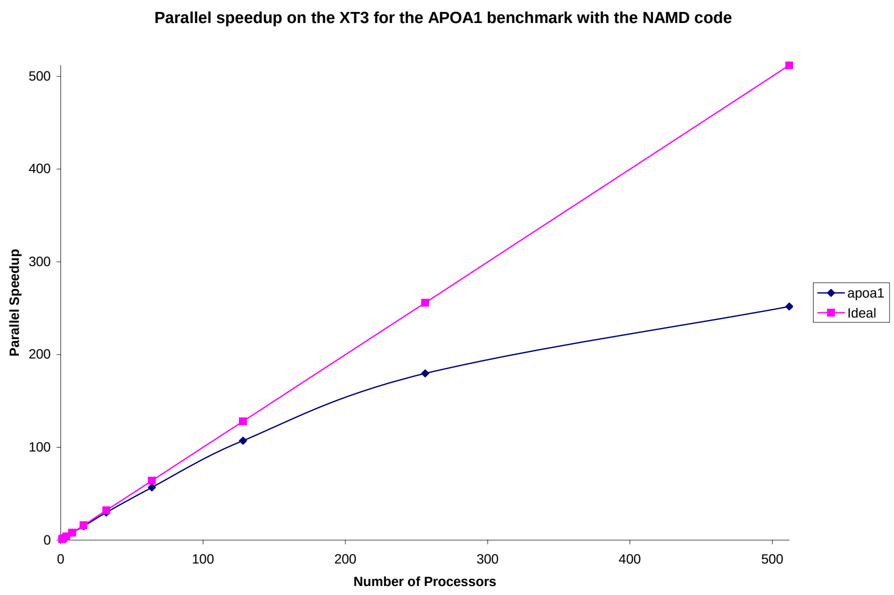

Apoa1 Benchmark: 29k atoms

STMV virus benchmark: 1M atoms

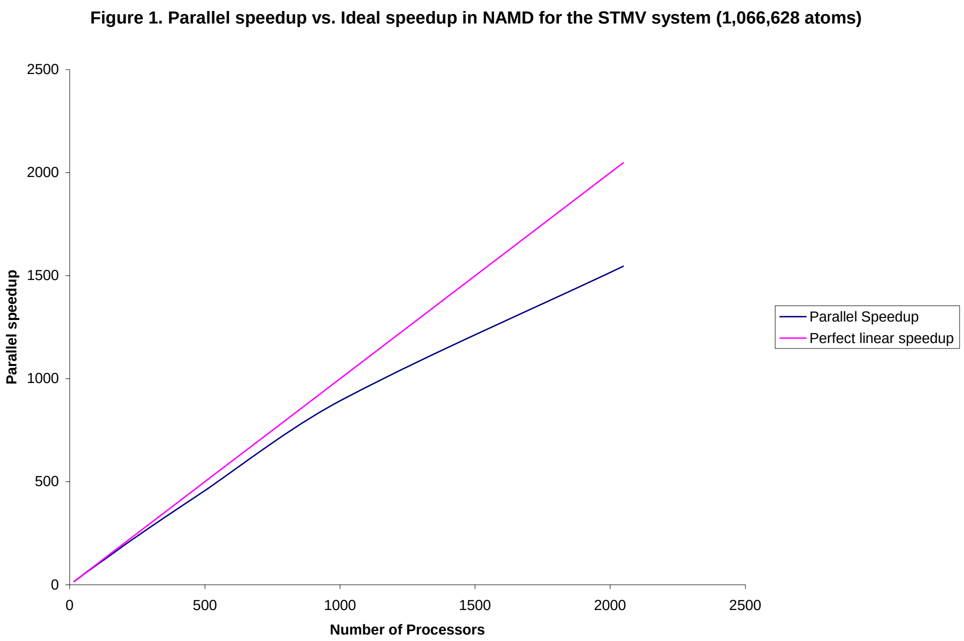

Optimizing NAMD for the XT3

- Close collaboration with Prof. L. V. Kalé and his group at UIUC to improve the efficiency and scaling of NAMD on the XT3
  - Developers of the Charm++ programming language
  - Profiling NAMD with Projections (why does ApoA1 scaling flatten?)

- Working with Howard Pritchard and Doug Gilmore (Cray) to ensure that the XT3's messaging system is used to maximum effect

Performance Overview with *Projections*

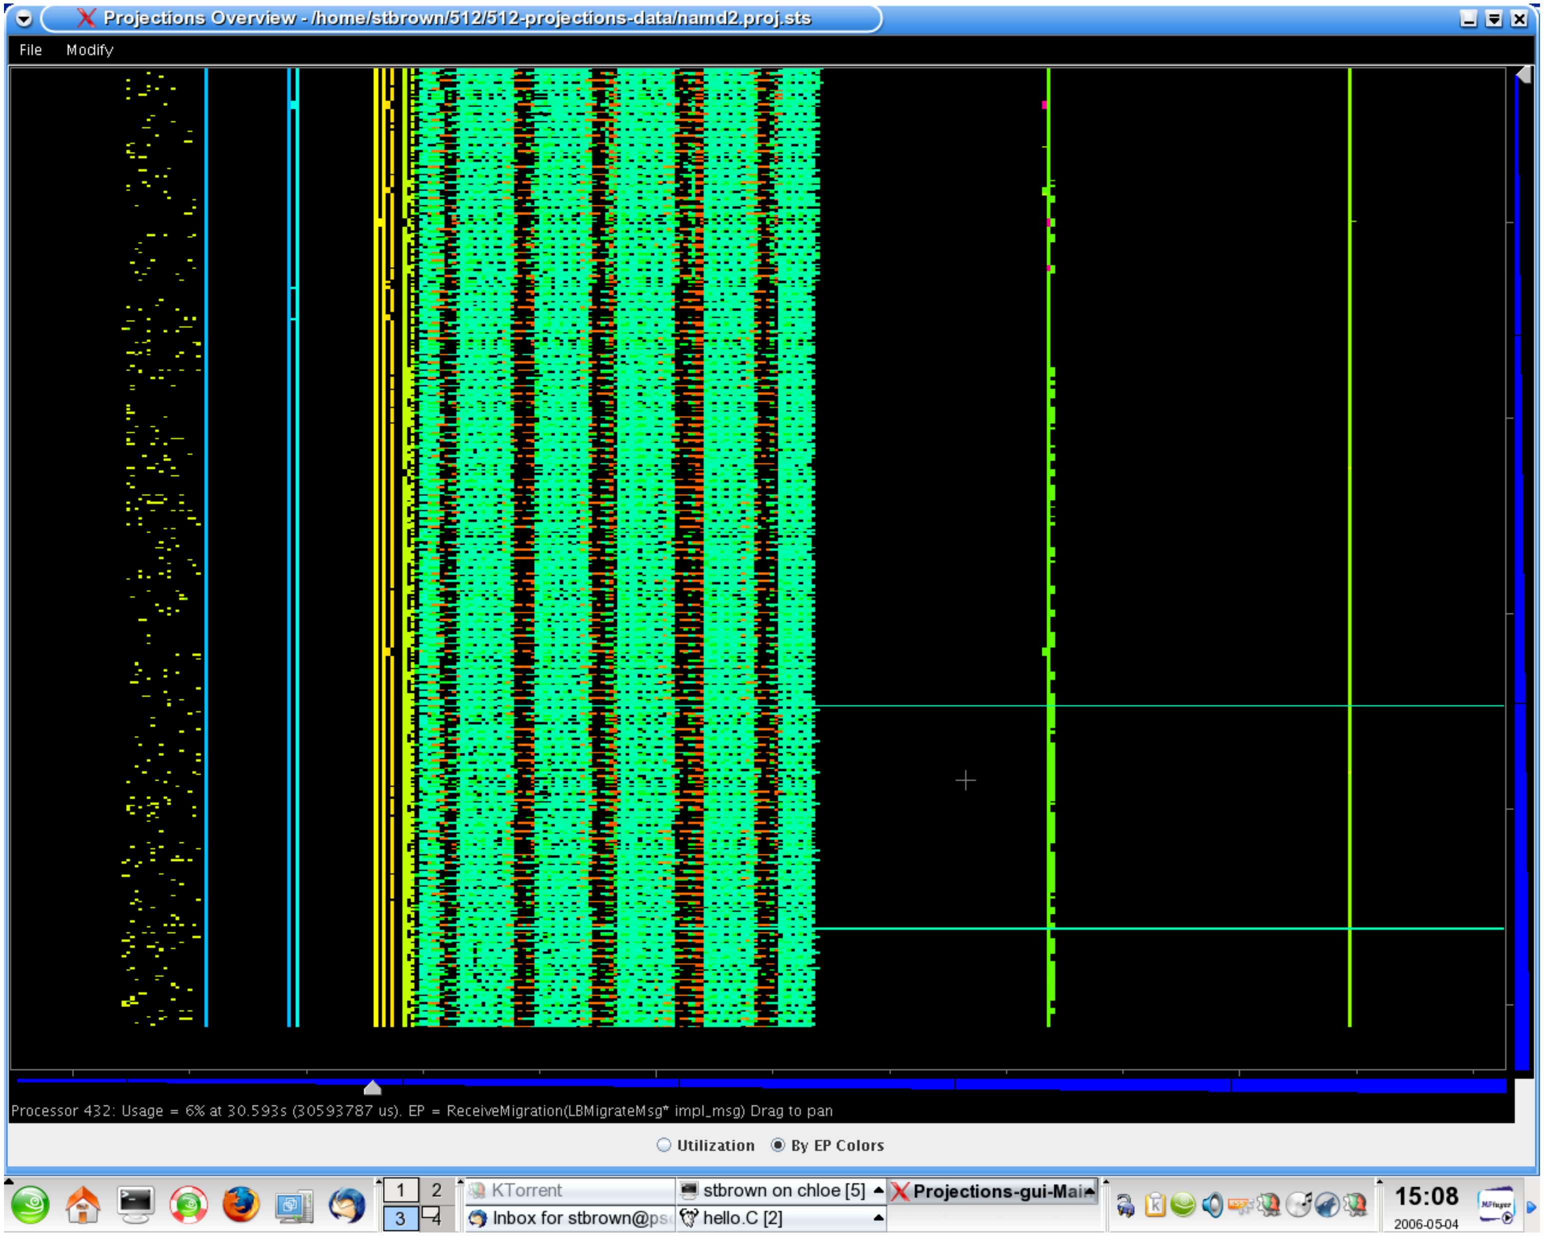

Load Imbalance

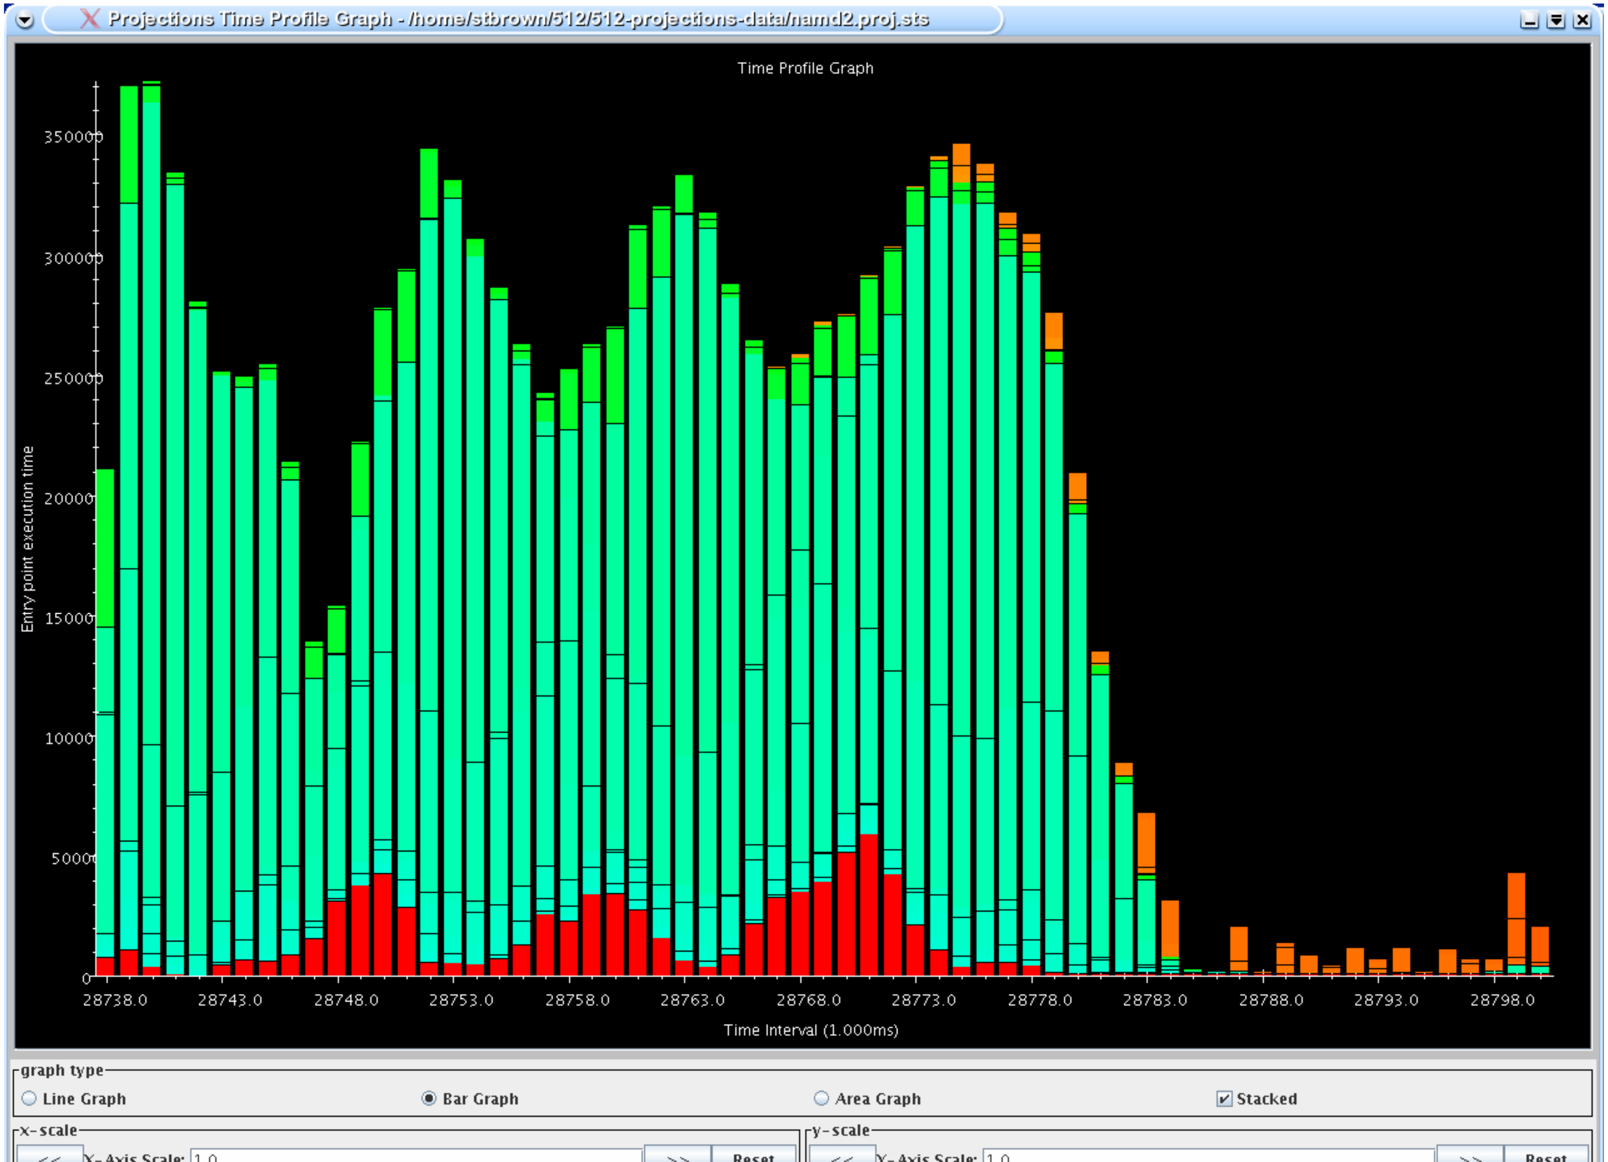


Projections Timelines - /home/stbrown/512/512-projections-data/namd2.proj.sts
File
Tools
Colors
Help
28,818,000
28,828,000
28,838,000
28,848,000
28,858,000
28,868,000
28,878,000
28,888,000
PE 0 (94.5%)
PE 23 (26.0%)
PE 46 (54.2%)
PE 69 (58.2%)
PE 92 (15.7%)
PE 115 (27.3%)
PE 138 (55.4%)
PE 161 (48.4%)
PE 184 (22.5%)
PE 207 (26.0%)
PE 230 (54.7%)
PE 253 (58.4%)
PE 276 (22.7%)
PE 299 (30.4%)
PE 322 (24.3%)
PE 345
Display Pack Times
Display Message Sends
Display Idle Time
View User Events (10)
Select Ranges
Change Entry Point Colors
<<
SCALE:
4.0
>>
Reset
Highlight Time
Selection Begin Time
Selection End Time
Selection Length
Zoom Selected
Load Selected

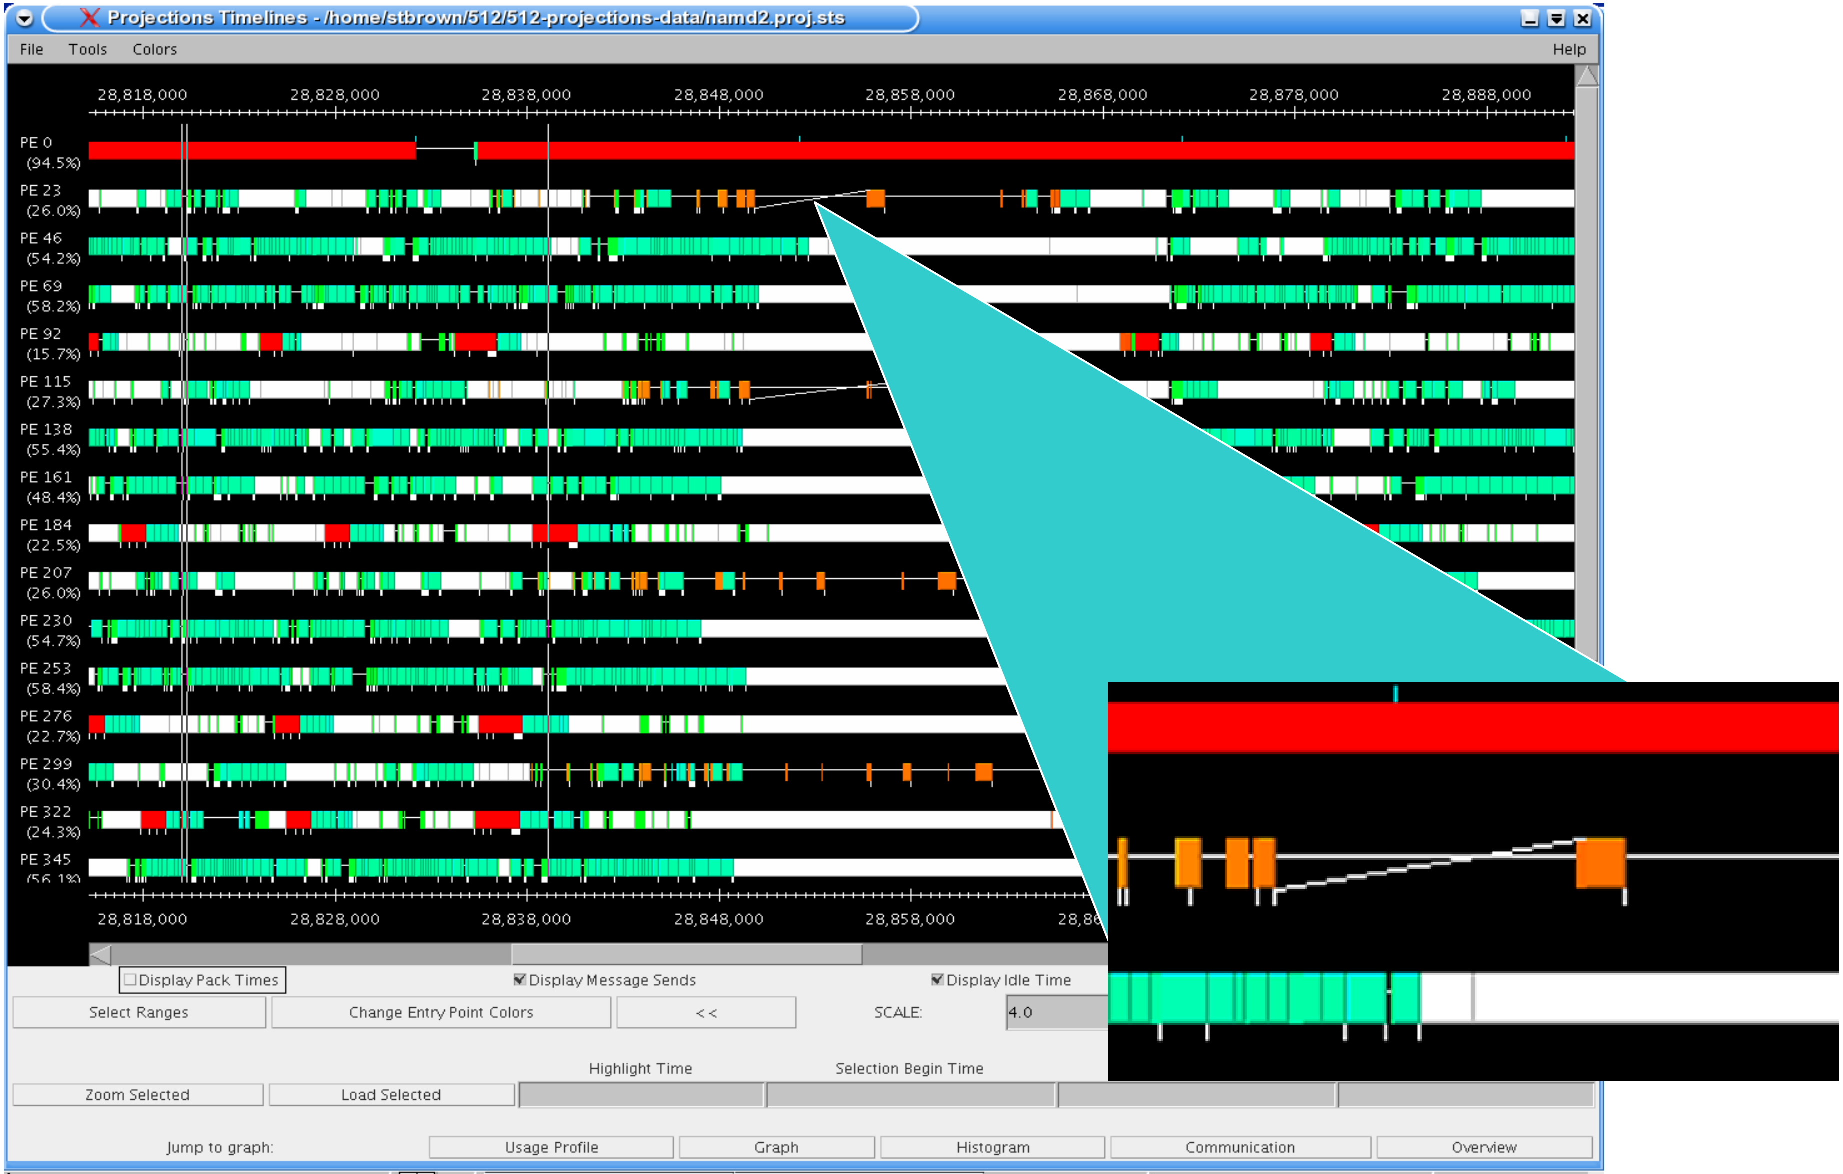
Projections Timelines - /home/stbrown/512/512-projections-data/namd2.proj.sts
File Tools Colors Help
28,818,000 28,828,000 28,838,000 28,848,000 28,858,000 28,868,000 28,878,000 28,888,000
PE 0 (94.5%)
PE 23 (26.0%)
PE 46 (54.2%)
PE 69 (58.2%)
PE 92 (15.7%)
PE 115 (27.3%)
PE 138 (55.4%)
PE 161 (48.4%)
PE 184 (22.5%)
PE 207 (26.0%)
PE 230 (54.7%)
PE 253 (58.4%)
PE 276 (22.7%)
PE 299 (30.4%)
PE 322 (24.3%)
PE 345
Display Pack Times
Display Message Sends
Display Idle Time
Select Ranges
Change Entry Point Colors
<<
SCALE:
4.0
Highlight Time
Selection Begin Time
Zoom Selected
Load Selected
Jump to graph:
Usage Profile
Graph
Histogram
Communication
Overview

Ongoing Optimization of NAMD and Charm++

- *Projections* revealed that significant time is being lost in MPI_Iprobe
  - this has been seen in several other MPI implementations as well

- Ongoing optimizations include:
  - Native Portals implementation of Charm++: Kalé et al.
    - will benefit other Charm++ applications as well
  - Improved MPI performance: Pritchard and Gilmore
    - will benefit other MPI applications which do extensive asynchronous messaging to overlap computation and communication

AMBER Breaks 10ns/day Barrier on the XT3

- "*For a constant volume factor ix run of my latest (pre-9) code, we now can exceed 10 nsec/day on the XT3.*"
  – Bob Duke, NIEHS
- In a related calculation, simulations of 134,427 atoms (plus water molecules) were performed on 128 processors of the Cray XT3 at PSC using PMEMD8, with long-range electrostatics interactions calculated using the Particle Mesh Ewald method.
- Coray M. Colina[1], Divi Venkateswarlu[2], Robert E. Duke[1,3], Lalith Perera[3], Tom A. Darden[3], and Lee G. Pedersen[1,3].
  [1]University of North Carolina at Chapel Hill
  [2]North Carolina A&T University
  [3]National Institute of Environmental Health Science

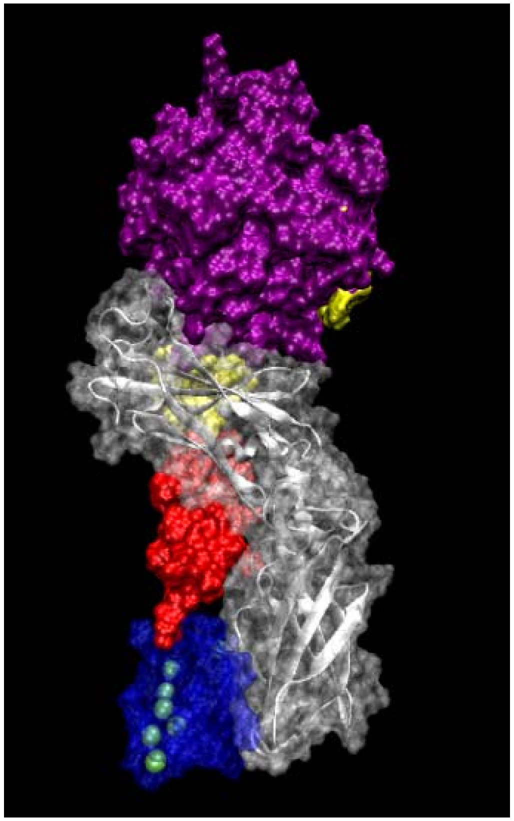

Snapshop of a 22 ns molecular dynamics simulations of the triggering complex of blood coagulation: soluble Tissue Factor-Factor VIIa. Tissue factor is represented in white, and the four domains of factor VIIa in: blue (GLA-domain), red (EGF1 domain), yellow (EGF2 domain) and purple (SP domain). Green spheres are calcium ions.

LSMS: Towards Petacomputing in Nanotechnology

- Goal: Understanding the magnetic domain walls from first principles. Historically, this is one of the greatest challenges to *ab initio* electronic structure calculations.
- Locally self-consistent multiple scattering (LSMS) method: a first-principles *O*(*N*) scaling technique
  - First application to sustain 1 TFlop/s; 4.65 TFlops on TCS; 1998 Gordon Bell award
  - ***52h57m continuous run using 2000 Cray XT3 processors***
- The Cray XT3 and future computing systems will enable realistic quantum mechanical simulation, e.g. study of the dynamics of magnetic switching processes, of real nanostructures.
- Y. Wang (PSC), G. M. Stocks, D.M.C. Nicholson and M. Eisenbach (ORNL), A. Rusanu and J. S. Faulkner (Florida Atlantic University).

LSMS sustains 8.03 TFlop/s (82% of theoretical peak) on 2048 Cray XT3 processors

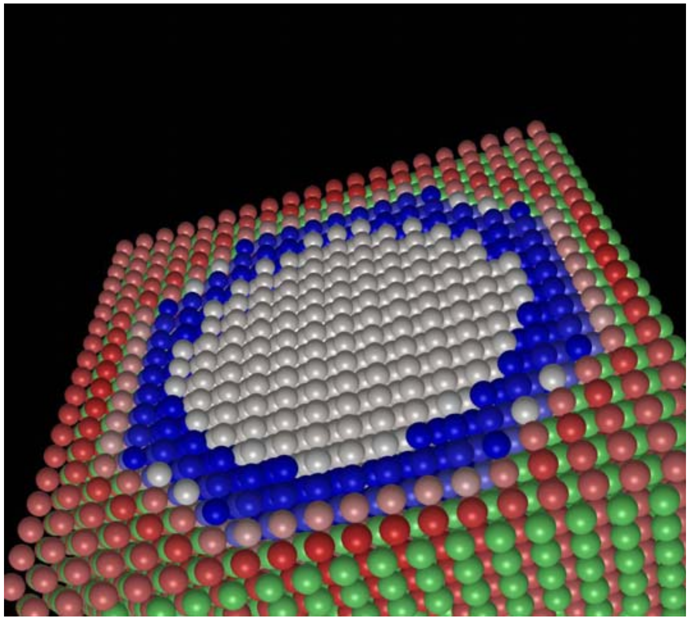

A sliced view of the magnetic Fe nanoparticle (with 4409 atoms on a BCC lattice) together with the surrounding FeAl matrix (with 11591 atoms on a B2 lattice). The calculated charge distribution within the nanoparticle and its surrounding atoms is indicated by the color change from the center to the edge.

Designing Magnetic Materials for High-Density Storage

- 3.86 nm spherical FePt nanoparticle embedded in FePt random alloy
- Y. Wang (PSC), G. M. Stocks, D. M. C. Nicholson, A. Rusanu, and M. Eisenbach (ORNL)

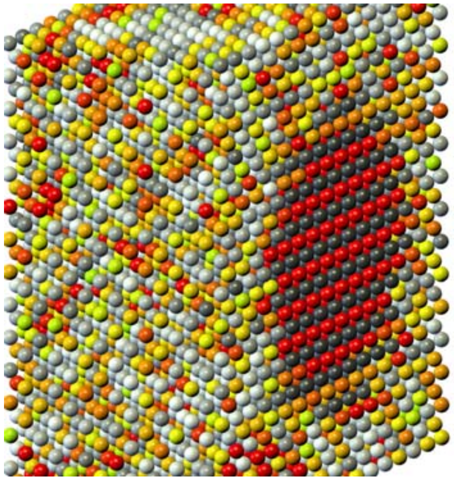

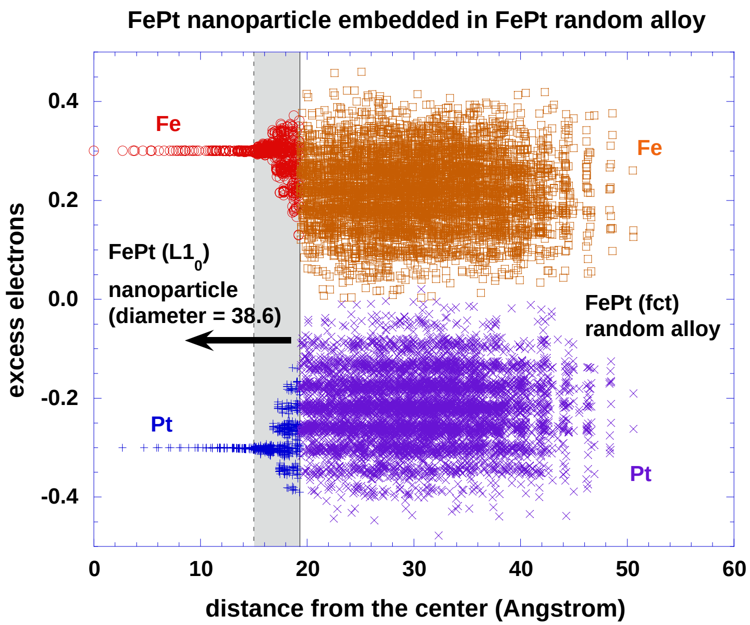

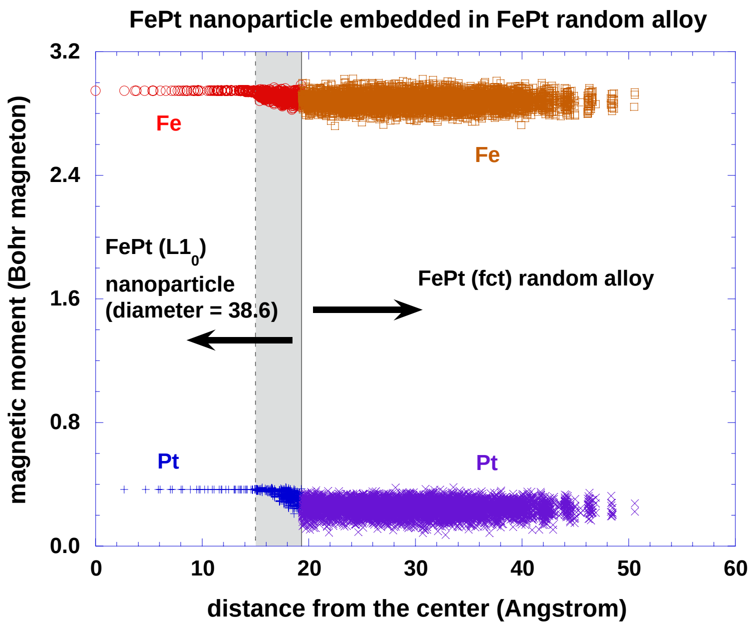

Gasoline: Astrophysical N-body Simulations

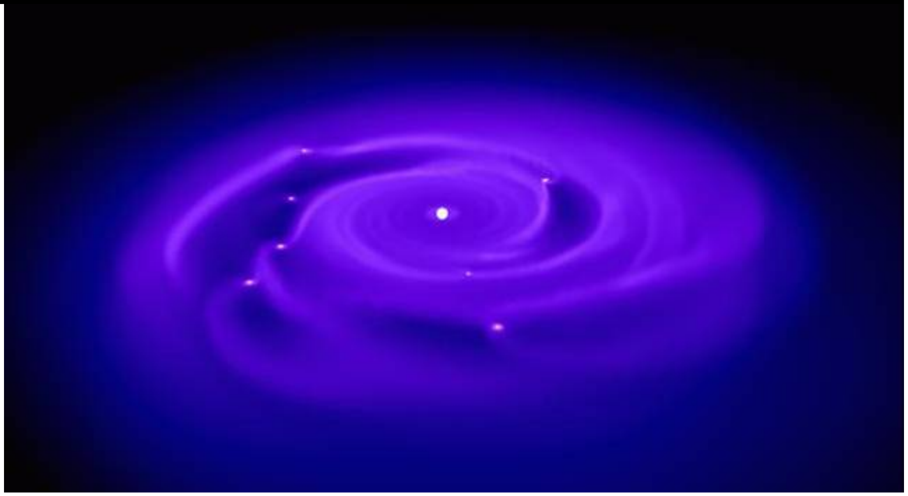

- The scaling of GASOLINE on the XT3 is very impressive: we have achieved up to 50% better scaling on the XT3 than the TCS for certain problems. Since we can run on more processors and the individual processors are much faster, the speed with which we can calculate a given scientific problem (our "science rate") will increase by nearly an order of magnitude. — *Tom Quinn, Astronomy, University of Washington*
- Multi-platform, massively parallel *N*-body tree code used to simulate a variety of astrophysical processes from “small” (on astronomical scales) to large:
  - Asteroid collisions
  - Planet formation
  - Tidal stirring of dwarf galaxies
  - Galaxy formation and evolution
  - Galaxy cluster intracluster light emissions
  - Sunyaev-Zel’dovich effect in the Cosmic Microwave Background

Gasoline: Astrophysical N-body Simulations (2)

- One particular strength of GASOLINE is that it can model how galaxies form and evolve at very high resolution
- Prior to the XT3, GASOLINE simulations could *either*:
  1. Simulate a very small number of galaxies in a "hand picked" part of the universe (i.e. a region not statistically representative of the universe as a whole) *or*
  2. Simulate a statistically representative piece of the universe, but not resolve the galaxies inside of it
- With the XT3, astrophysicists can now do *both*. Current calculations are simulating the formation and evolution of galaxies at high resolution, and in a volume that is large enough that it models a representative piece of the universe.

Hercules: End-to-End Earthquake Simulation

- ***Complete simulation: parallel meshing, solution, rendering***
- Complex ground motion simulation
  - multiple spatial scales: of O(10m-100km); multiple temporal scales: O(0.01-100s); highly irregular basin geometry; highly heterogeneous soil material properties; geology and source parameters observable only indirectly
- Scientific Goals
  - Simulation of a magnitude 7.7 earthquake centered over a 230km portion of the San Andreas fault in southern California
  - 2Hz simulation, using a new adaptive mesh (~10B elements), will afford a 64× larger grid than the SCEC “Terashake” simulation (0.5Hz, 1.8B grid points)
- Scientific motivation: the 2Hz simulation will provide much better resolution and incorporate 4× high frequencies than the SCEC calculation, quantifying the effect of higher frequencies.
- Collaborators: Volkan Akcelik, Jacobo Bielak, Ioannis Epanomeritakis, Antonio Fernandez, Omar Ghattas, Eui Joong Kim, Julio Lopez, David O'Hallaron, Tiankai Tu, *Carnegie Mellon University;* George Biros, *University of Pennsylvania;* John Urbanic, *PSC*

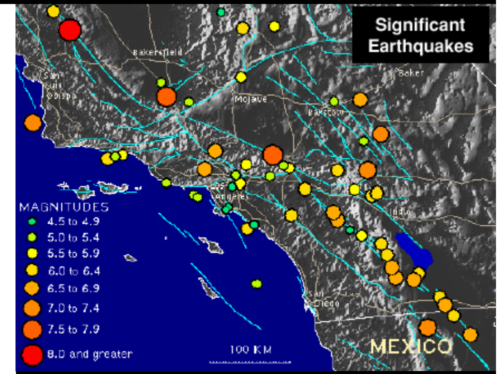

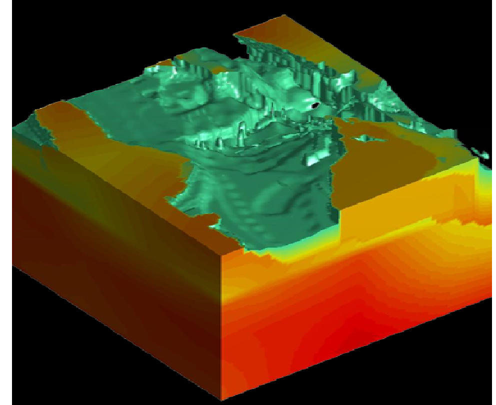

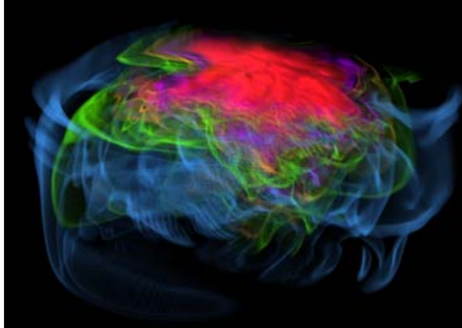

Volume Rendering

- 4D volume rendering of beating mouse heart
- CT mouse data (200^3) courtesy of the Duke Center for In-Vivo Microscopy
- Rendered on the Cray XT3 using svr
- Rendering and animation: Art Wetzel, Stu Pomerantz, Demian Nave (PSC)

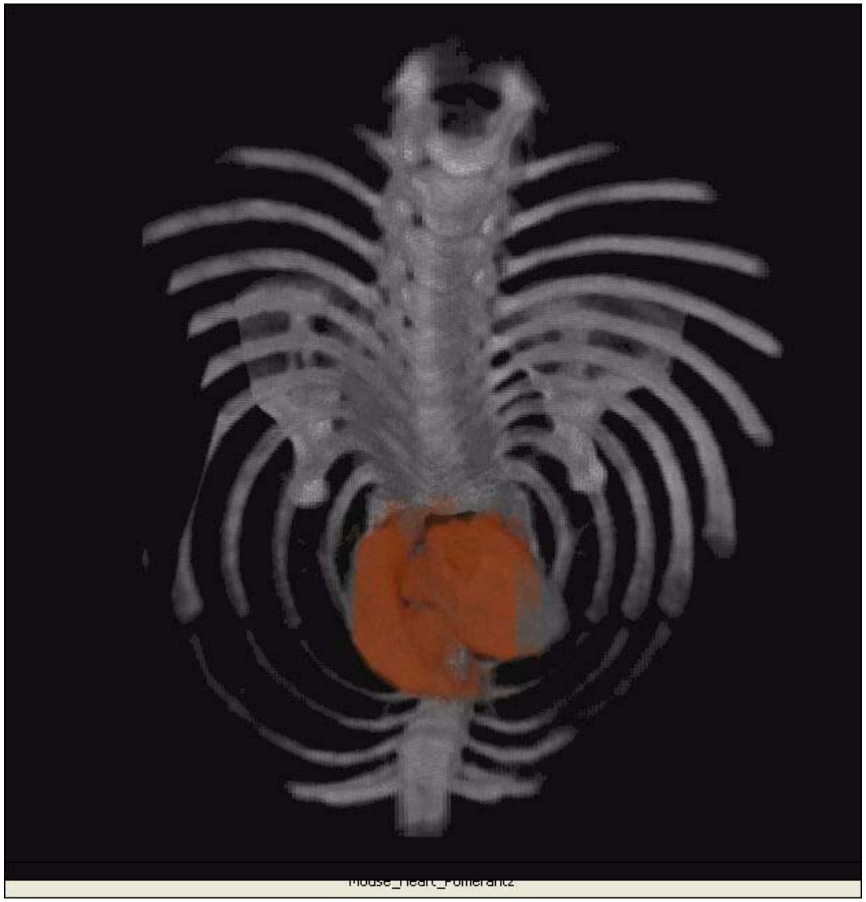

Conclusions

- BigBen, the PSC's Cray XT3, is heavily used for a very broad range of scientific inquiry, regularly obtaining important scientific results
- Cray XT3 performance is excellent, up to 10.5× that seen on TCS, due largely to the high-bandwidth SeaStar interconnect
- PSC's PDIO library extends communications from compute nodes to the TeraGrid, allowing researchers at the University of Minnesota to interactively steer and visualize large-scale CFD calculations
- Applications are routinely breaking performance records and scaling to new heights on the XT3
- PSC's XT3 workshops have been very productive in enabling new applications quickly
- PSC is doing significant performance optimization including single-processor optimization, job layout, infrastructure, and networking