

Getting Work out of New, High-End Systems

Michael Levine/Ralph Roskies

Scientific Directors, PSC

(Nick Nystrom, AD Apps

Ray Scott, AD Operations)

The Salishan Conference on

High-Speed Computing

April 20, 2005

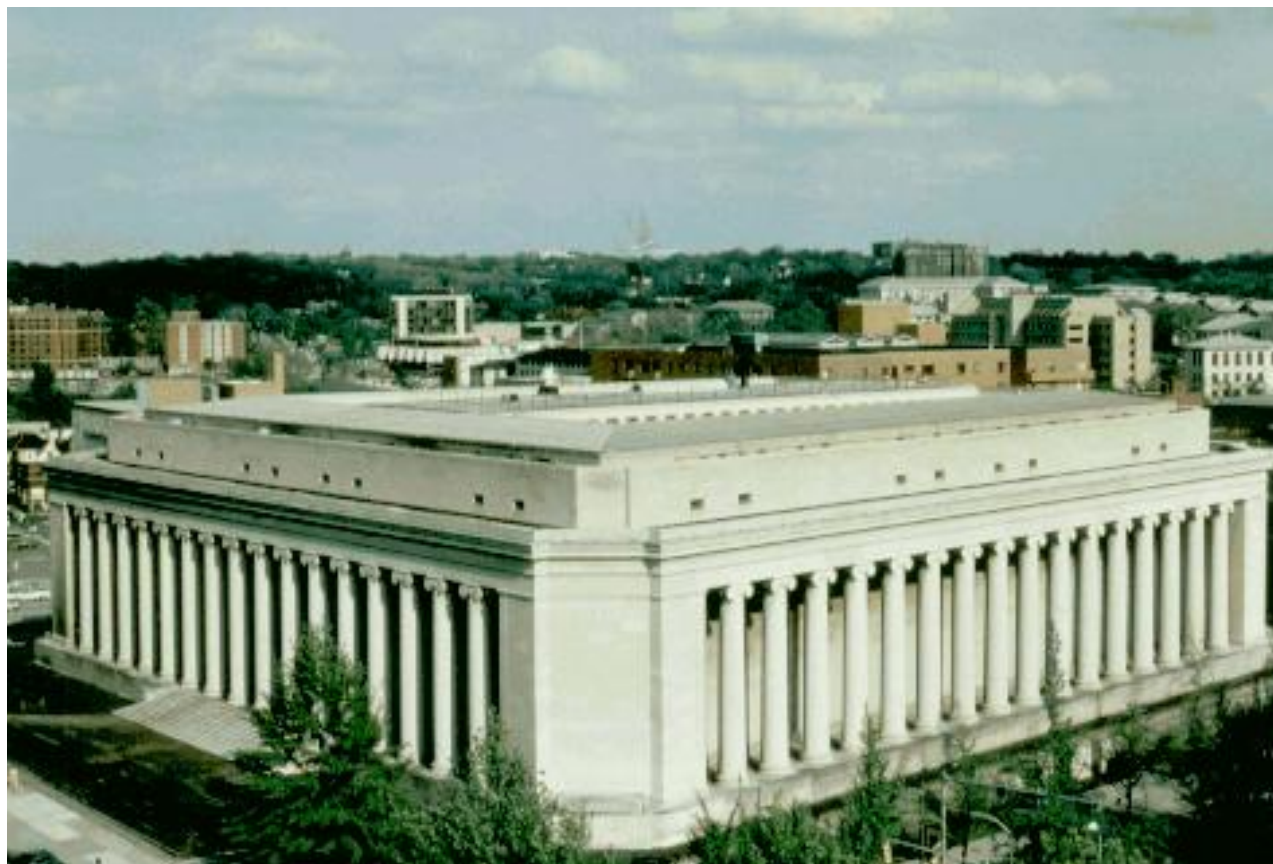
P I T T S B U R G H
S U P E R C O M P U T I N G
C E N T E R

Introduction & Background

- Thanks to Salishan Conference organizers for invitation.
- Start with a very brief intro to PSC
 - Explain basis of experience for this talk
 - Been in business since 1985
 - Predominantly, but not exclusively, an NSF sponsored resource center for HEC.
 - Notice: Science is the NSF's middle name!
 - Member of the large class envious of National Labs' budgets
 - Small size limits scope of approaches compared to Nat'l Labs
 - Ralph Roskies, my partner in crime, is also here

Pittsburgh Supercomputing Center (PSC)

- Cooperative effort
 - Carnegie Mellon University (MJL)
 - University of Pittsburgh (RZR)
 - Westinghouse (JK→LL)
- Offices in Mellon Institute
 - On CMU campus
 - Adjacent to Pitt campus.



PITTSBURGH
SUPERCOMPUTING
CENTER

Westinghouse

- *Energy Center,*
Monroeville PA
- Major computing
systems.
- High-speed
network
connections.



PITTSBURGH
SUPERCOMPUTING
CENTER

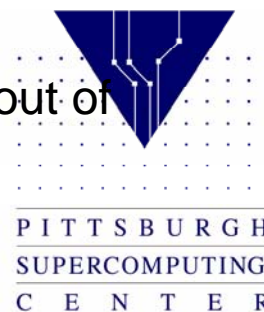
PSC's Primary Mission

- As a *Resource Provider* on the US-NSF *TeraGrid*, PSC provides HEC capability to the US open, scientific research community.
 - Other TG RP's at meeting
- A part of the NSF's *Cyberinfrastructure* effort to *revolutionize* the doing of scientific research through advanced IT.
- PSC, as for past 19 years, is charged with providing *Leadership Class Computing* resources



PSC's Major Production Platforms

- *Lemieux*: NSF's Terascale Computing System (TCS) (EV68)
- *Jonas & Rachel*: large memory SMP's (Marvel - EV7)
- Cray XT3, *being deployed*.
 - "Product" version of *Red Storm*
 - 2112 node, 3D torus, ~20% of SNL Red Storm
 - Collaborative efforts with SNL, Cray, ORNL, CSCS, ...
- History of *first and early* systems over past 19 years
 - XMP, YMP, CM2, CM5, T3D, C90, T3E
 - Lot's of experience "standing up" new systems and getting work out of them.



Framing Comments (thank you Horst)

- As with Horst's sorrow, I regret at not having *theoretical* materials to discuss
 - Those are VERY important issues but not my assignment, today.
- This talk is largely about *barnyard engineering*: making the best use of existing capabilities of new systems
 - Based on practice developed at PSC over nearly 20 yrs and with very low budgets
 - Shared experience with LANL on "Q" shows analogous problems and solutions (probably why Manuel Vigil gave me this assignment)
- Our approach to Horst's dilemma (science vs. technology) is to focus on the **Science** (it is the *National Science Foundation*, stupid)
- We can not afford to wait ~> 1 year from introduction of new technology to its productive use!
- (With XT3, as with every Cray at PSC since XMP, we push for *early* use)
 - End-of-life fixed by technology & competition
 - We can modify the start-of-life by careful preparation and ingenuity

..
..
..
..
..
..
H

SUPERCOMPUTING
C E N T E R

Outline

- Problems: getting the right configuration rapidly into productive use
- Solutions: how we make progress despite the problems.
- Example areas
 - System configuration, physical and logical
 - Platform-specific and pan-platform system management environment
 - Flexibility in all efforts and not shying away from temporary workarounds
 - Examples: file systems and additional IO capabilities
 - **Focus on Science and Users.**
- Apology: *barnyard engineering* is eclectic by nature
 - **An excise in searching each day for *low hanging fruit***
 - **Approach surely applicable to PetaScale systems!**
- Most examples from XT3 although issues are general
- Status and successes on PSC's XT3



Source of the Problems

- Switch to commodity components has
 - Reduced hardware costs
 - Reduced margins
 - Less specialized work by vendors
 - More burden on “Centers”
- No current vendor has full depth experience and expertise in HEC.
 - Some bright spots, some trolls
- Burden falls ever more to the centers or other customers
- Pick up specific problem areas in discussion of actions
 - (Again: PSC’s small size forces great selectivity in *development*)

Example: System Design (LeMieux/Q issues)

- Overly simplistic scaling up of from small systems might produce an adequate capacity system but not a good capability system (caution! re petascale)
 - Network: consultation with Quadrics → substantially shorter cable and latency; increased effective bw
 - Maintenance: quick response is no help. Locally designed use of error logs and preventative maintenance improves availability
 - I/O: artful use of Quadrics as a MAN to viz & storage
 - PSC libraries
 - PSC snapshot acceleration services
 - Coupled to scheduler producing *resubmitting* job and ameliorating modest mean time to next node failure (Similar to LANL on “Q”)
 - Application-level checkpointing. Not burdensome if it is your code.
 - (Discussed by Keshav Pingali, earlier)



Example: How Much Disk bw?

(*Defensive IO II*, for Seager)

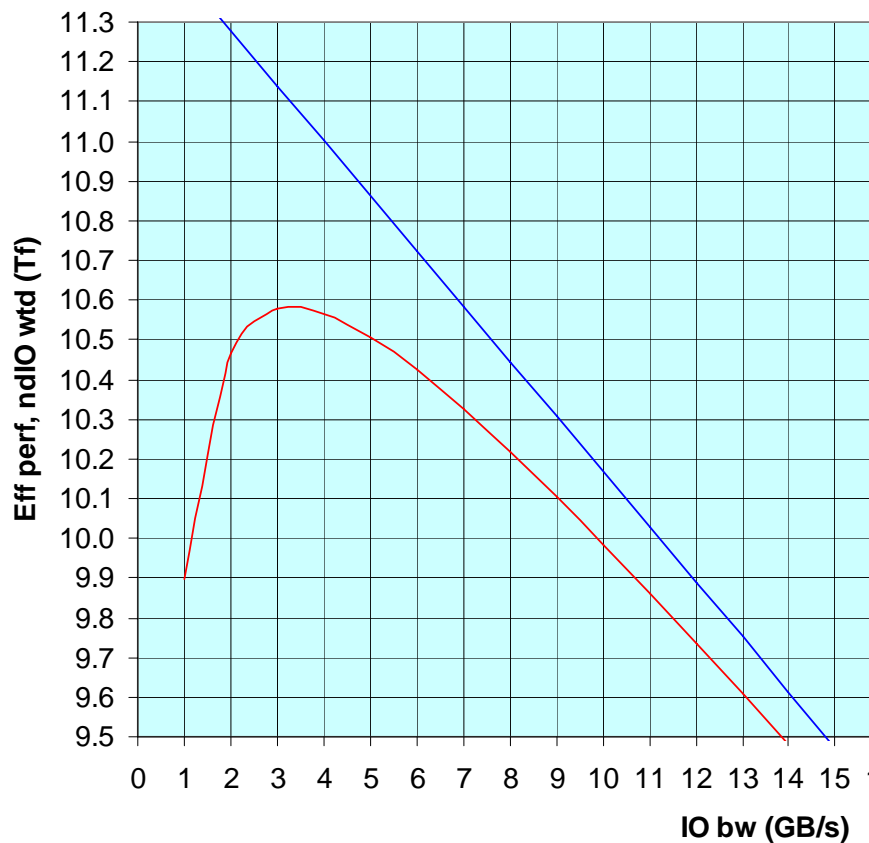
- Question: is there an optimal disk investment?
- Answer: Yes, and it depends on a lot of things.
 - Fix total investment
 - Determine
 - cost of compute,
 - cost of IO bw.
 - Compute effective computational power as function of bw for given IO load including
 - Optimal *defensive* IO (frequency of checkpoint as fn of MTTI & dump time)
 - Required IO
 - Convolute with distribution of IO as fraction of usage.
 - Get optimal IO bw and sensitivity.

How Much Disk bw? (XT3)

Bw cost: \$170/MB/s

Bw*perf: 3.4GB/s → 10.58Tf

Sensitivity: high

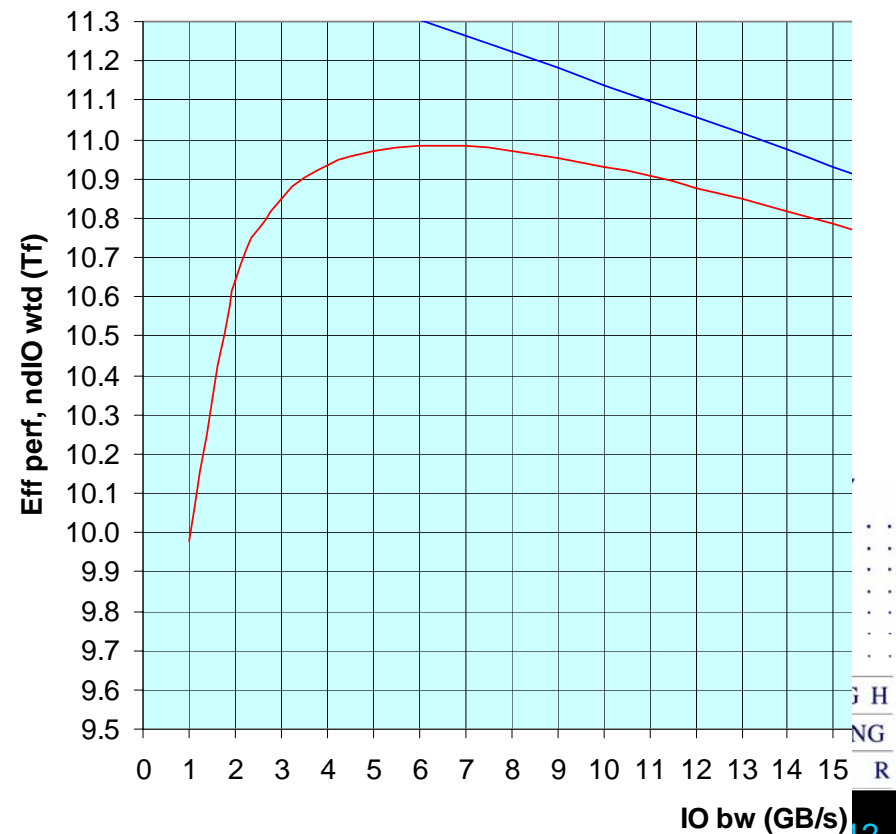


Supercor

\$22.5/MB/s

6.5GB/s → 10.95Tf

low



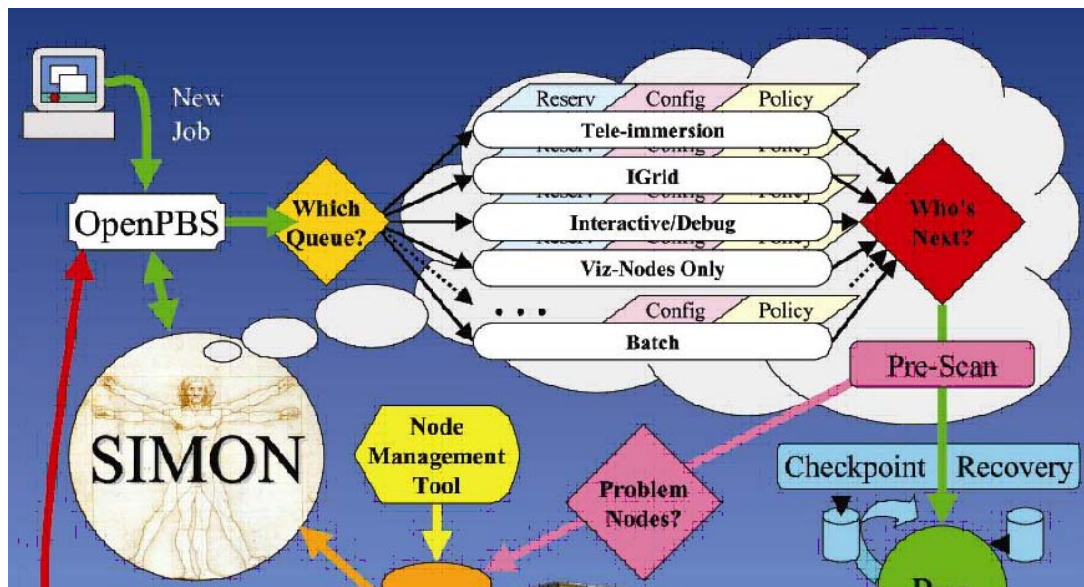
12

System Software Environment

- Platform-specific:
 - Running quasi-production at PSC for months using system *checkout* management software suite
 - Compare with *Chippewa Operating System* by Seymour Cray for the CDC-6600
 - Moral: simple but working is better!
- Pan-platform:
 - It is what users see. Separate and stabilize.
 - Batch queuing & scheduling, acctg, MAN, bulk disk included
 - This separation is why we have had queues, monitors, etc at each XT3 stage from SC04 through present



Batch Scheduling

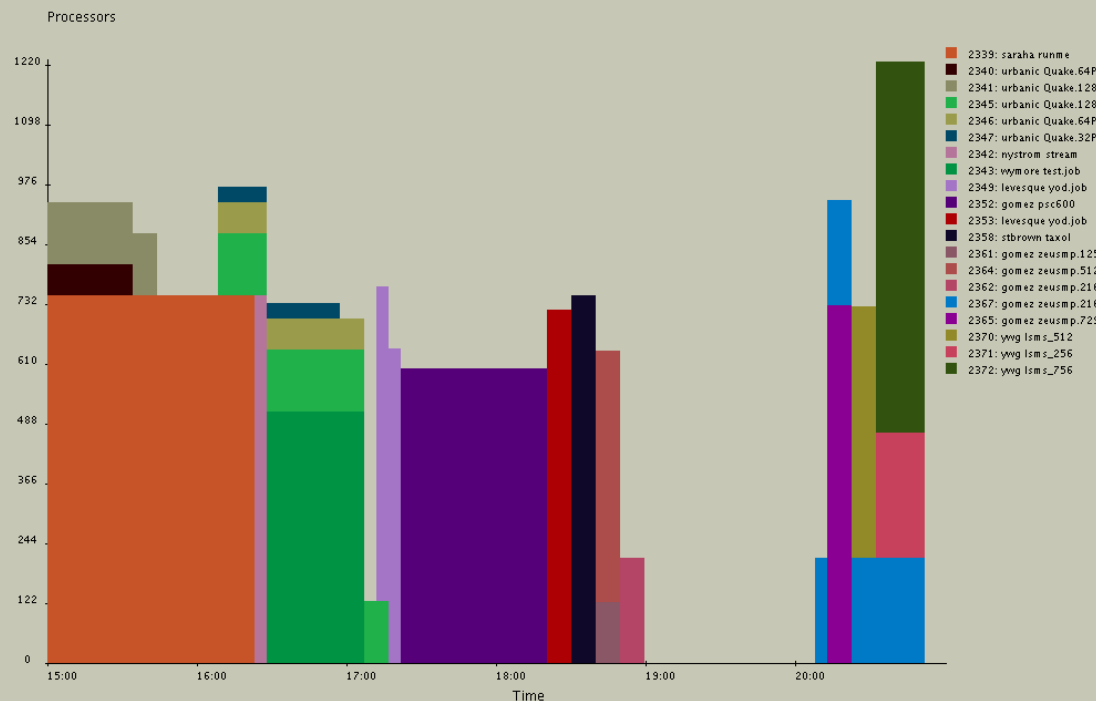
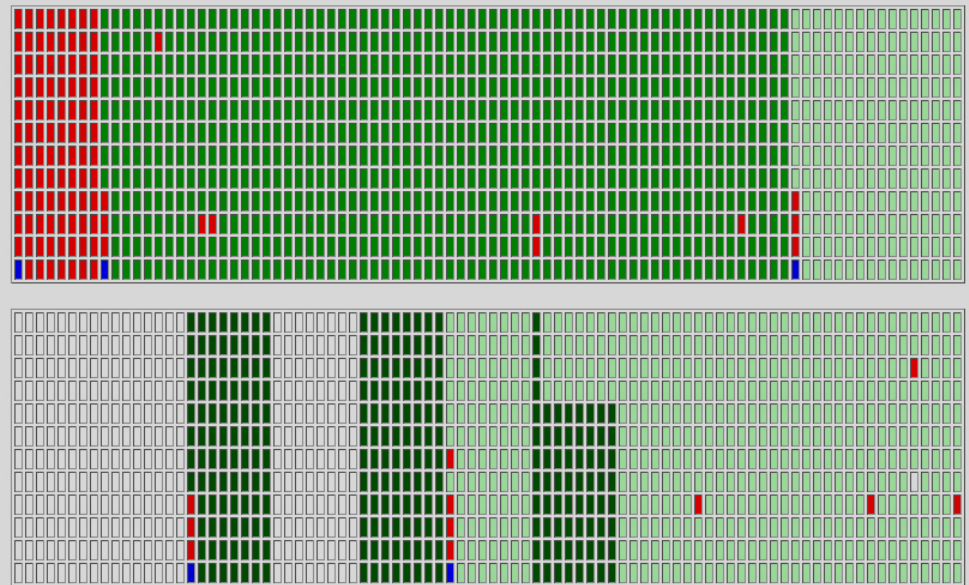


- OpenPBS port
- Scheduler based on Simon scheduler from LeMieux
 - Backfill
 - Removing troubled nodes from service
 - Integrated with XT3 job launcher (yod)
- Monitoring
 - Daily Report
 - *shownids* command with system status
 - Visual Monitors

Visual Monitors

- Job Layout
- Utilization History

XT3 Status Monitor



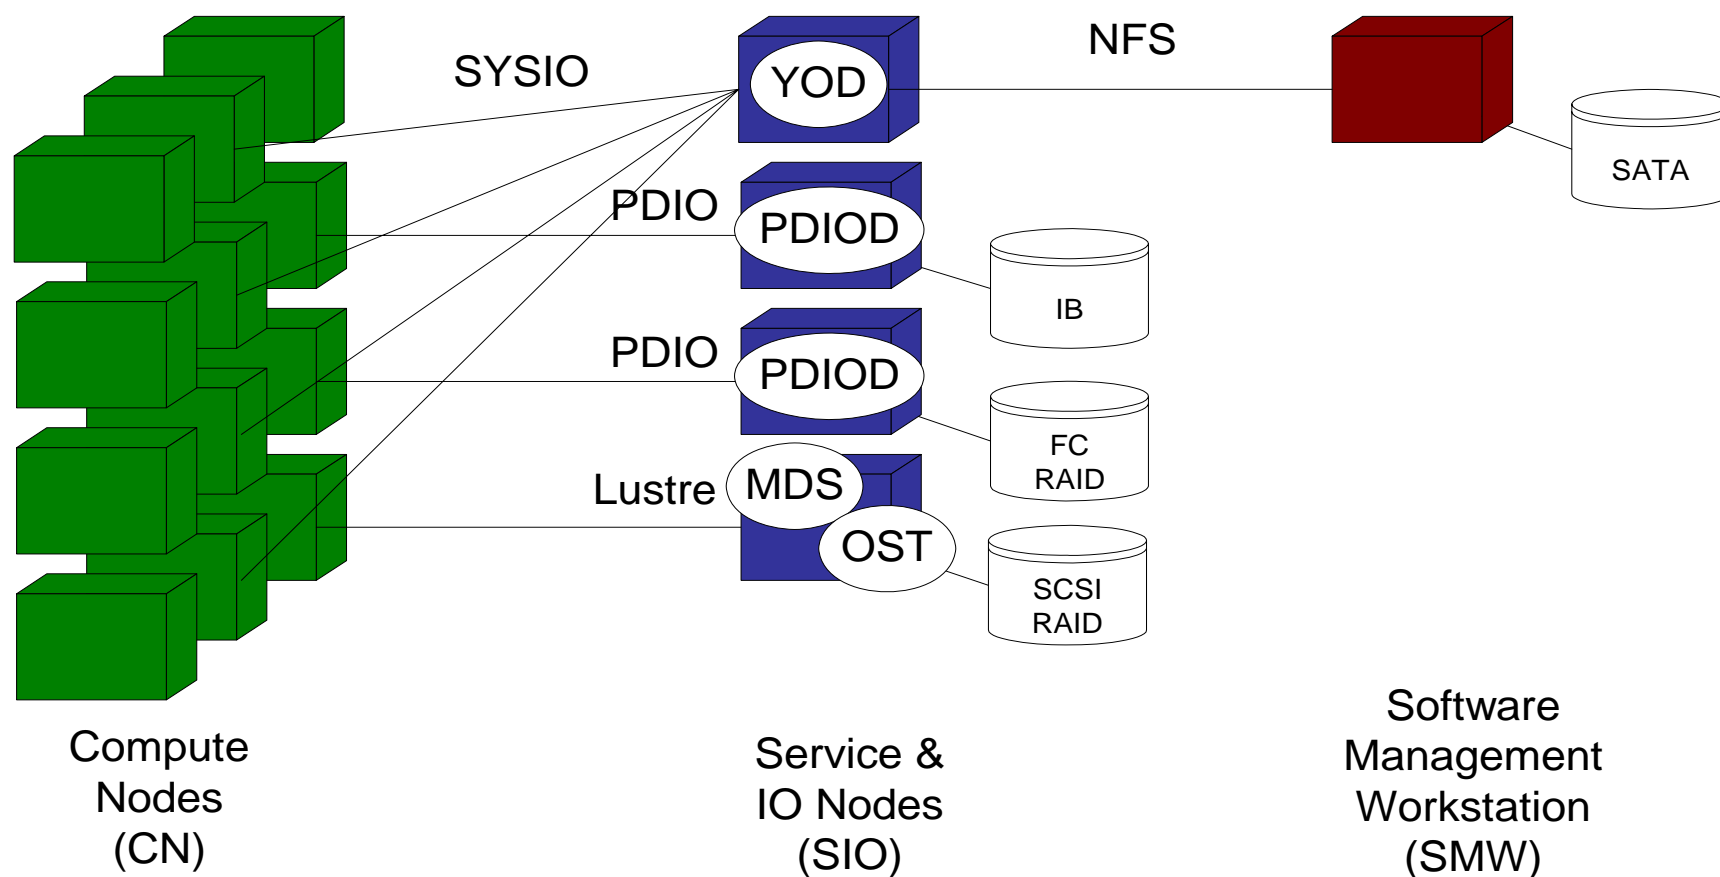
Lustre File System

- PSC experimenting with Lustre long before XT3 delivery.
- PSC's Lustre file systems in production.
 - Experimented w/Lustre long before XT3 delivery
- One of multiple file systems
 - No need for monoculture
- Lustre file system in production
- Grid FTP experiment NCSA-PSC/XT3 Lustre
 - (later, under apps)

Additional IO Capabilities

- Additional IO support
 - SCSI interface (to less expensive disk)
 - InfiniBand MAN
 - To Linux Routing Nodes (twixt TeraGrid Network and XT3)
 - To disk servers
- Portals Router
 - Computation is best use of module in XT3 messaging network.
 - Move as much Lustre disk handling as possible to external file servers
 - Interconnect using IB
 - Controlled sharing with other systems

Portals Interconnect Utilities



Focus on Science and Users

- Stay very close to users
 - PSC staff with computational research experience assigned to each major user group
- Stay in the *line-of-fire*
 - PSC staffers *take the arrows*
 - User community needs to be respected and protected
- (Lessons being transferred to NSF TeraGrid
 - Premature exposure of users to immature services only works once. If it fails, you have lost your audience.)

So, where does that put us?

- Sue Kelly was correct in her *Lessons Learned* but she underemphasized the good news.
- Take what does work!
- Remember that it is the Science and other (NSF) programmatic drivers that count (Dimitri Kusnezov's comments)
- Selective *barnyard engineering* already provides a surprisingly useful system (despite XT3's immature status)
- Quick pass through a broad-brush status report to show what has been accomplished...

NSF Program

- CyberInfrastructure (CI) organized as TeraGrid (TG)
 - Multiple resource providers
 - (ANL, SDSC, NCSA, ORNL, TACC + ?, here)
- Inclusion in a program is particularly useful in initial period
 - We can focus on the apps that do work
- All (nearly) PSC resources are “on the TG”
- XT3 is *on the TeraGrid*

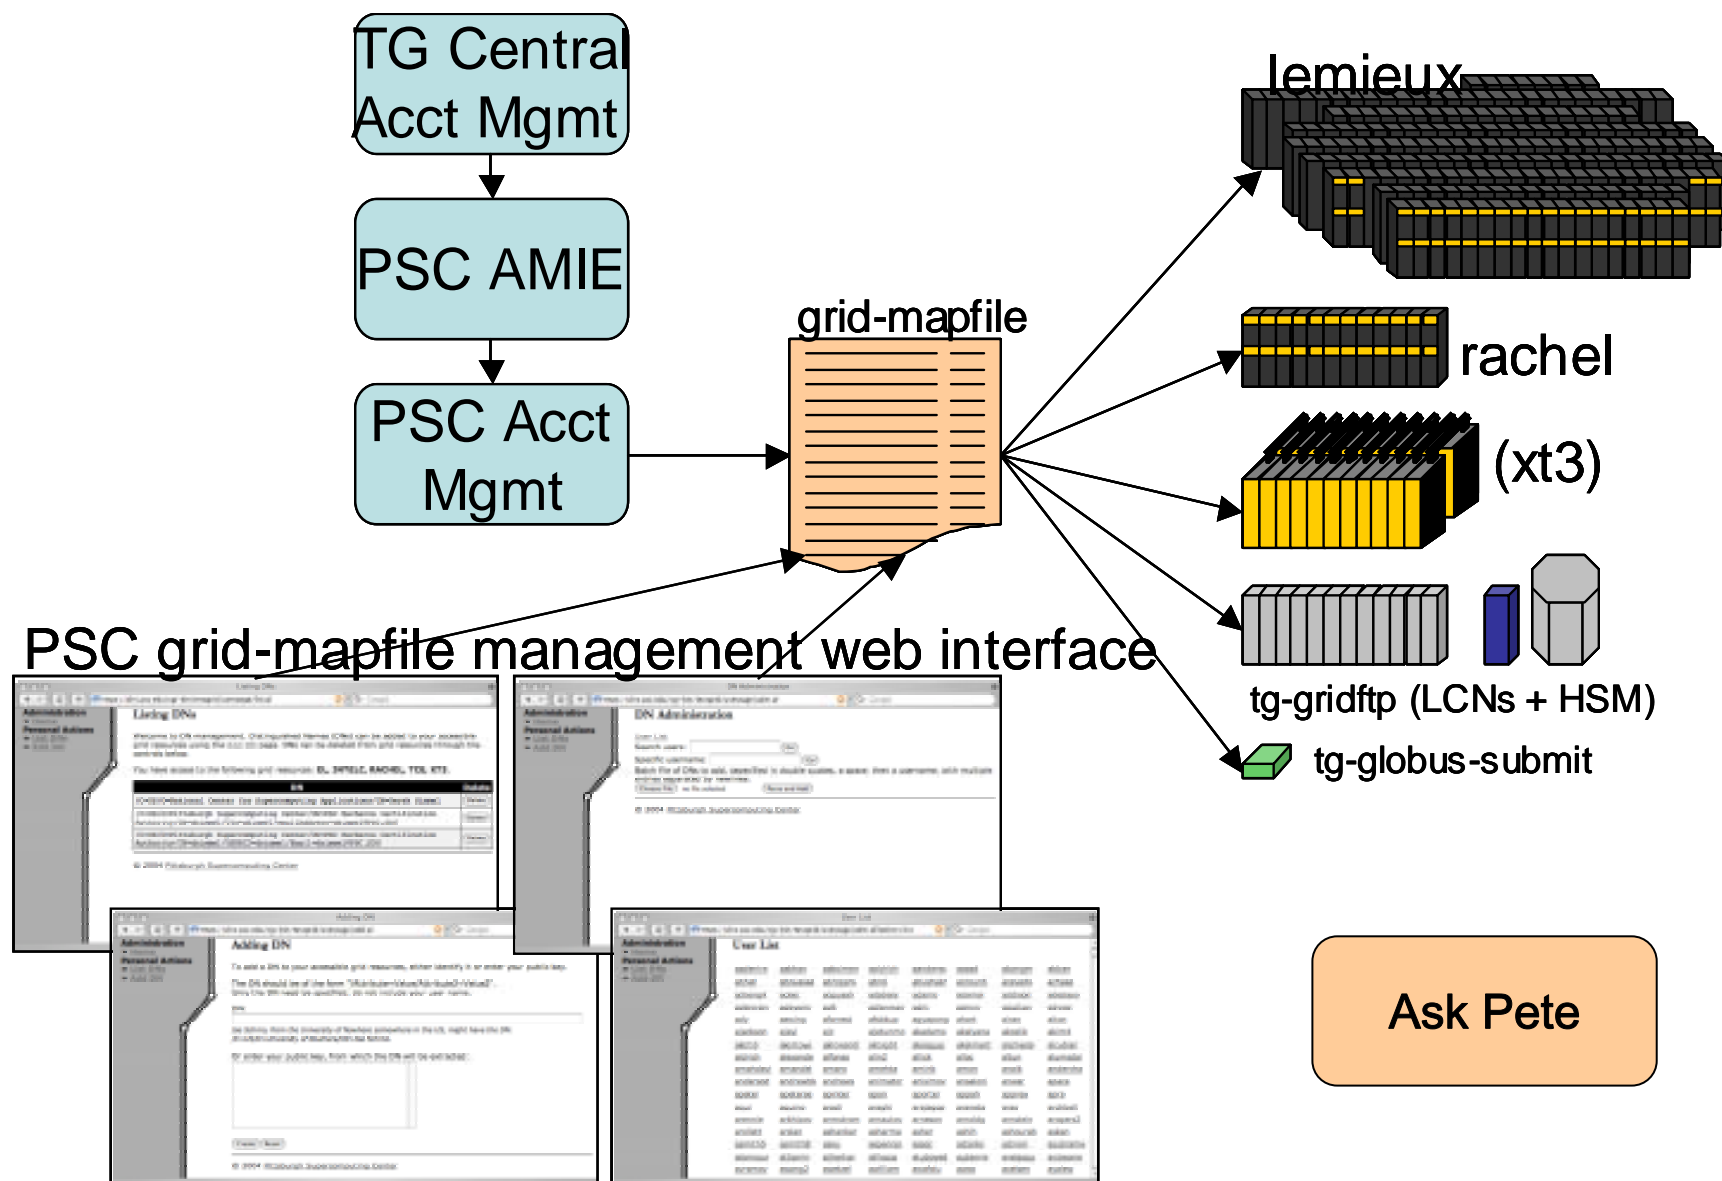


XT3 TeraGrid Presence

- Account Control
 - Account Creation using *ath*
 - *gridmap* updated automatically
- Usage Reporting
 - PSC Usage Database
 - AIME database updates
- Grid Services
 - Job submission
 - File Transfer



TeraGrid / PSC grid-mapfile deployment & management



XT3 System Status

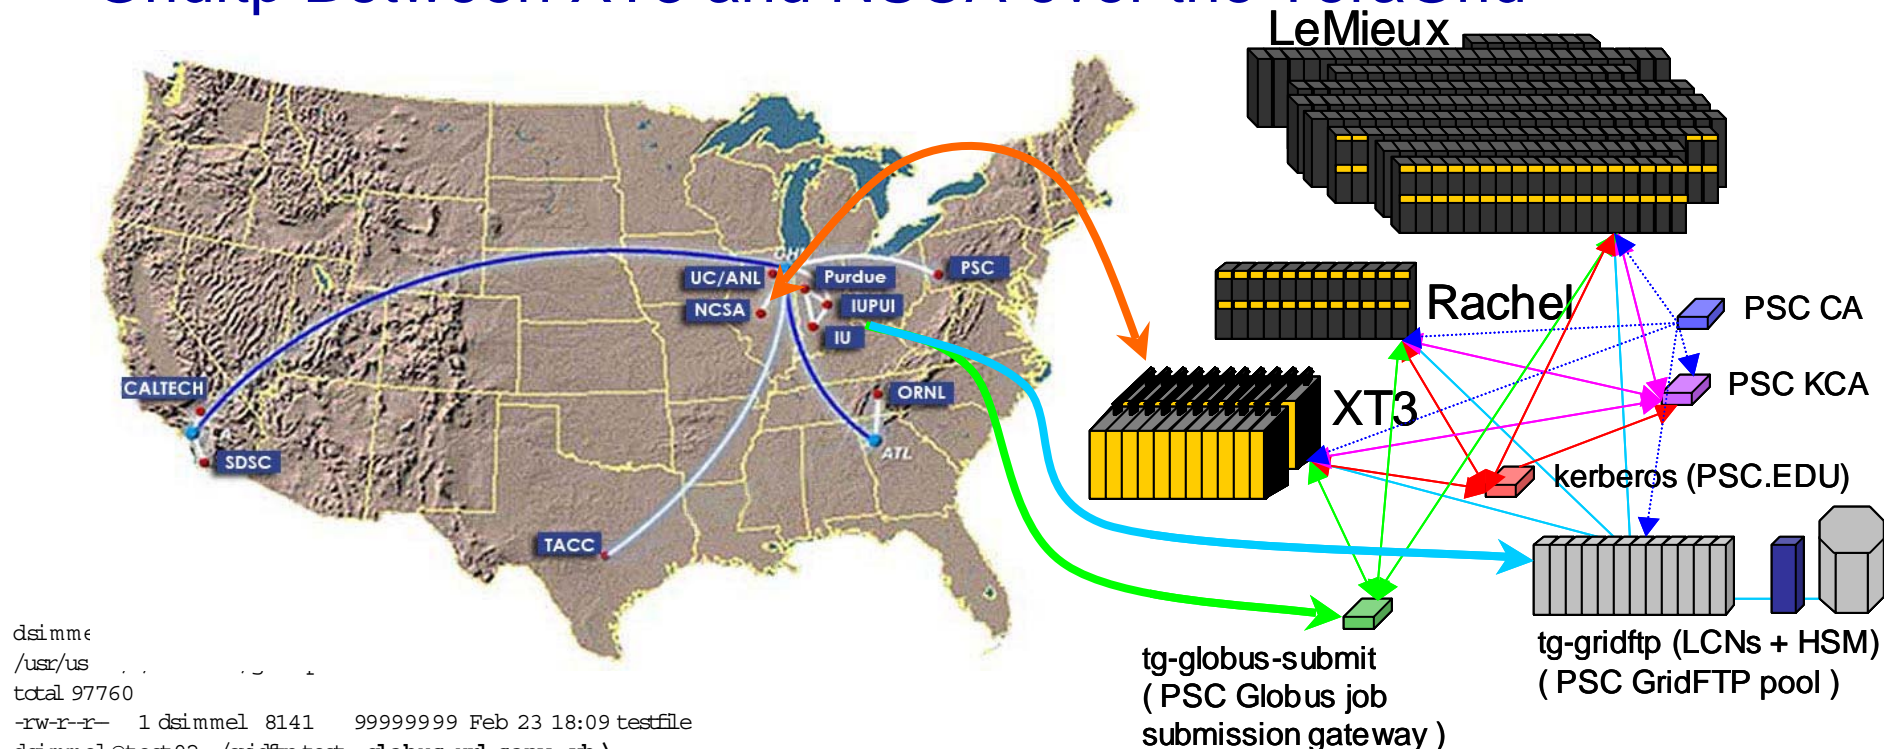
- State of installation
 - 1st chassis live, on floor, at SC04; Batch queues running (modest) apps
 - 1st row (11 cabs) installed (multiple systems) and running apps end December 2005
 - 2nd row (11 cabs) installed and running apps Feb 2005; currently in 2 parts for improved system MTBF
- Running *quasi production*
 - *6 month miasma*: initial scientific work mixed with system development
 - Started with Opteron systems, look-alike systems and network access to Cray development systems long before hardware delivered.
 - *New Science* being done
 - ~50% utilization when “up”
 - Lustre running for users

XT3 Apps Successes

- Successful GridFTP file transfers from NCSA to the XT3's Lustre file system
- Strong effort with NSF user community



Gridftp Between XT3 and NSCA over the TeraGrid



```

dsimml
/usr/us
total 97760
-rw-r--r- 1 dsimml 8141 99999999 Feb 23 18:09 testfile
dsimml@test02: ~/gridftp-test> globus-url-copy -vb \
file:///usr/users/0/dsimml/gridftp-test/testfile gsiftp://tg-login1.ncsa.teragrid.org/home/ncsa/dsimml/testfile01
Source: file:///usr/users/0/dsimml/gridftp-test/
Dest: gsiftp://tg-login1.ncsa.teragrid.org/home/ncsa/dsimml/
testfile -> testfile01
97517568 bytes 5.43 MB/sec avg 6.00 MB/sec inst
dsimml@test02: ~/gridftp-test> globus-url-copy -vb \
gsiftp://tg-login1.ncsa.teragrid.org/home/ncsa/dsimml/testfile01 file:///usr/users/0/dsimml/gridftp-test/testfile02
Source: gsiftp://tg-login1.ncsa.teragrid.org/home/ncsa/dsimml/
Dest: file:///usr/users/0/dsimml/gridftp-test/
testfile01 -> testfile02
96468992 bytes 2.71 MB/sec avg 2.51 MB/sec inst
dsimml@test02: ~/gridftp-test> ls -l
total 195520
-rw-r--r- 1 dsimml 8141 99999999 Feb 23 18:09 testfile
-rw-r--r- 1 dsimml 8141 99999999 Feb 25 12:52 testfile02

```



Pre-Production Apps Strategy

- Assign PSC consultants to users with current or potential interest in the Cray XT3 system
- Maximize return on resources (both PSC and users) by staging access according to users' needs and system availability
 - *Developers* : users who are primary developers of applications that use or can use capability-class systems
 - *End users* : users with challenging scientific problems; consumers of applications developed by developers
 - *Performance specialists* : performance modelers and systems specialists seeking to characterize relatively stable, production-ready machines
- (Stages often overlap)



Pioneers in Computational & Computer Science (1)

PI	Application Area	NSF Directorate
Michael Klein/U Penn	Materials Science (CPMD)	MPS
Kelvin Droegemeier/Oklahoma	Storm Modeling (ARPS)	GEO
Klaus Schulten/Illinois	Complex Biological Systems (NAMD)	BIO
Charles Brooks/Scripps	Molecular Dynamics (AMBER, CHARMM)	BIO
Joel Stiles/CMU	Physiological Modeling (MCell)	BIO
Michael Norman/UCSD	Astrophysics (ENZO)	MPS
Steve Gottlieb/Indiana	QCD (MILC)	MPS
Tom Quinn/Washington	Cosmology (Gasoline)	MPS
George Karniadakis/Brown	Fluid Dynamics and Turbulence (NekTar 2.5d)	ENG
Paul Woodward/U Minnesota	Fluid Dynamics & Astrophysics (PPM)	MPS
Omar Ghattas/CMU	Blood Flow	ENG
Jacobo Bielak/CMU	Earthquake Modeling (Quake)	ENG
John Michalakes/NCAR	Climate Modeling (WRF 2.0)	GEO
PK Yeung/Georgia Tech	CFD	MPS

Pioneers in Computational & Computer Science (2)

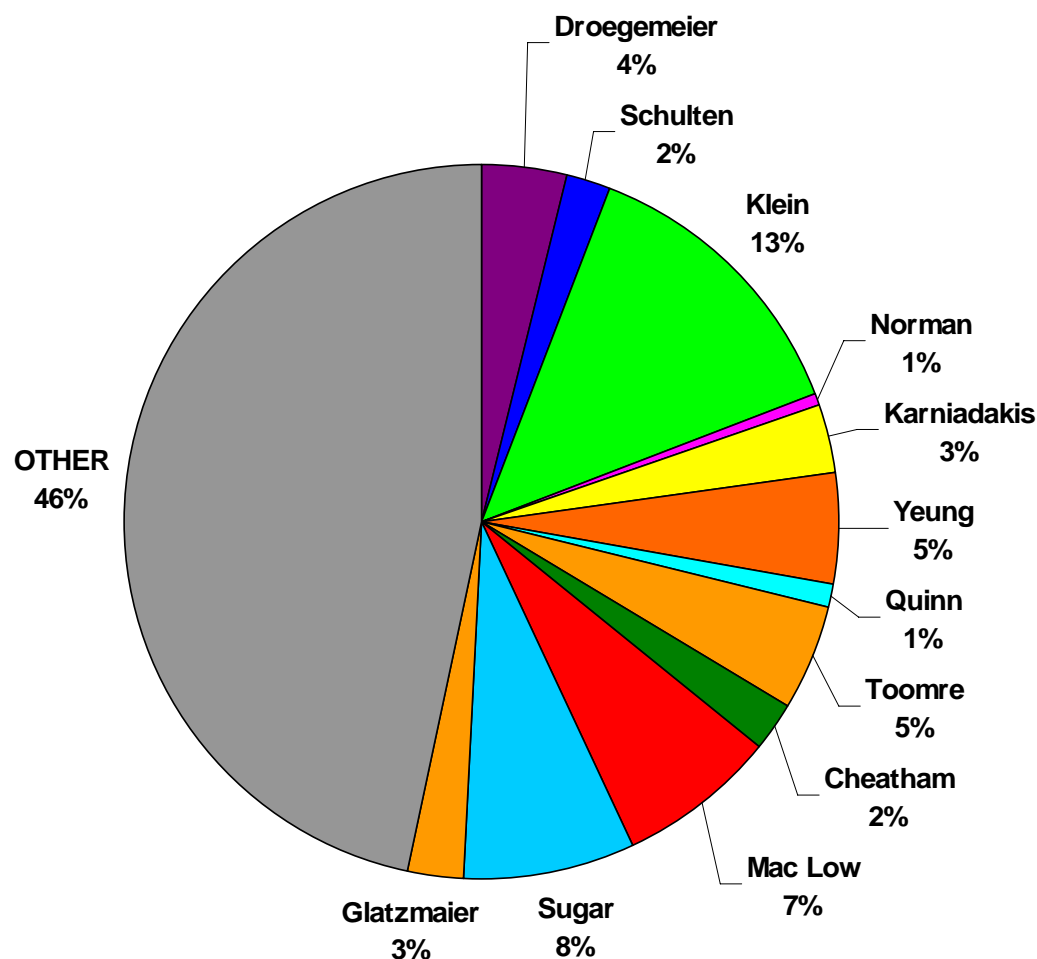
PI	Application Area	NSF Directorate
Richard Fujimoto/Georgia Tech	Magnetosphere modeling	MPS
Gary Glatzmeier/UC Santa Cruz	Structure & dynamics of planet and star interiors	GEO
Greg Voth/Utah	Chemistry	MPS
Susan Atlas/UNM	Chemistry (ATLAS)	MPS
Mark Ingber/UNM	Mechanical Engineering, classical MD (pdQ)	MPS
Juri Toomre/Colorado	Astrophysical Fluid Dynamics	GEO
Yang Wang/PSC	Materials Science (LSMS)	MPS
Roberto Gomez/PSC	General Relativity (Leonardo)	MPS
Thomas Cheatham/Utah	Medicinal Chemistry (MD)	MPS
Carlos Simmerling/SUNY Stony Brook	Computational Structural Biology (MD)	MPS
Adrian Roitberg/Florida	Protein and Peptide Folding (MD)	MPS
Mordecai-Marc Mac Low/American Museum of Natural History, Columbia	Numerical gas dynamics and MHD (ZEUS)	MPS
Troy Wymore/PSC	Enzyme dynamics (Dynamo)	MPS

Pioneers in Computational & Computer Science (3)

PI	Application Area	NSF Directorate
Lee Pedersen, Robert Duke/UNC-Chapel Hill	Molecular Dynamics (AMBER)	MPS
Tiziana di Matteo/CMU	Cosmology	MPS
David Yuen/Minnesota	Large-scale Geophysics	GEO
Craig Douglas/Yale and U. Kentucky	Data Driven Numerical Simulation of Wildfires	GEO
George Riley/Georgia Tech	Network simulation	CISE
Sanjay Kale/Illinois	CHARMM++ and Performance Modeling	CISE
John Mellor-Crummey/Rice Jarek Nieplocha/PNNL	Co-Array Fortran, HPCToolkit ARMCI	CISE
Kathy Yelick/Berkeley	UPC	CISE
Pat Worley/ORNL	Performance Analysis	CISE
Jack Dongarra/UTenn	Performance Analysis	CISE
Barney McCabe/UNM	Lightweight Kernels	CISE
Fabrizio Petrini/LANL	System modeling and performance	CISE

Early XT3 Users

Normalized usage on PSC Platforms Feb04-Jan05



- 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
Glatzmeier	4,829,447
Cheatham	4,266,049

...

Total 191,928,345

Representative Applications: PSC

<u>applications</u>	<u>domain</u>	<u>nodes</u>
STREAM	Memory bandwidth benchmark	1035
Quake	Earthquake modeling	1024
Gasoline	N-body astrophysics	1024
MILC	QCD	1024
Leonardo	Numerical relativity	1014
ZEUS-MP	Astrophysics	1000
MPQC	Massively Parallel Quantum Chemistry	941
Dynamo	QM/MM biochemistry	900
HPCC	HPC Challenge benchmarks	900
sPPM	Piecewise Parabolic Method bmk	768
PD	Portals Direct	760
cpu_burn	System stability	760
GAMESS	Quantum chemistry	708

<u>application</u>	<u>domain</u>	<u>nodes</u>
LSMS	Materials science / electronic structure	700
NBP EP	NAS Parallel Benchmarks: EP	600
HPCC	HPC Challenge benchmark	552
FFTW	Adaptive FFT library	256
NPB: LU, CG, MG	NAS LU Decomposition, Conjugate Gradient, Multigrid	256
NPB: FT	NAS Fourier Transform	256
AMBER	Molecular dynamics	128
ARPS	Storm modeling	128
CHARM++	Parallel language & runtime sys.	128
PMB	Pallas MPI benchmarks	64
svr	Volume Rendering	64
TBLC	Turbulent Boundary Layer Code	32
CHARMM	Molecular dynamics	16
NAMD	Molecular dynamics	16
VASP		128

Applications and libraries
 Benchmarks and kernels
 Systems and stress tests

Representative Applications: Cray

application	domain	nodes
LSMS 1.9	Materials science	1024
POP 2.1	Ocean model, 0.1 degree	960
HPCC HPL	LINPACK benchmark	625*
HPCC PTRANS	Transpose benchmark (interconnect bandwidth	625*
Patisan	SNL application and benchmark	564
ITS	SNL application and benchmark	564
sPPM	3D gas dynamics benchmark (piecewise parabolic method)	560
POP 1.4.3	Ocean model, 1 degree	480
CTH	Shock wave physics	256
Sweep3D	3D Discrete Ordinates Neutron Transport benchmark	192
Sage	SNL application and benchmark	188

Applications and libraries
 Benchmarks and kernels
 Systems and stress tests

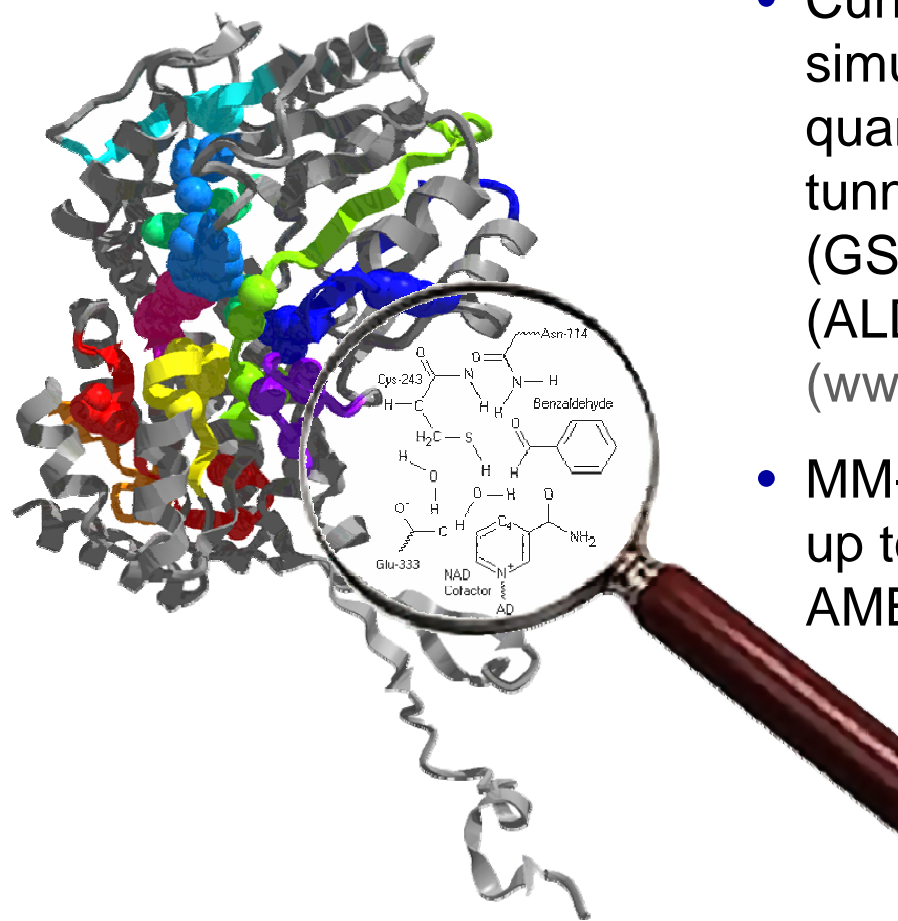
Scientific Targets

- As applications are developed for the Cray XT3, scientific runs begin. Goals include:
 - Generate scientific results
 - Serve as a workload, facilitating deployment of a productive, stable system when production use commences
 - As the XT3 grows, scale applications to larger processor counts, enabling increasingly ambitious simulations



**RUNNING:
900 NODES**

Hybrid Quantum Mechanical/Molecular Mechanical simulations on the Cray XT3 will enable new insight into the mechanisms of enzymes



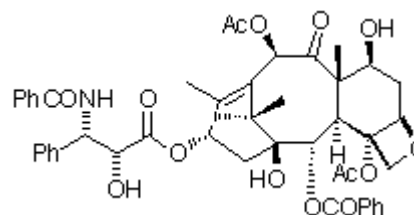
- Currently running path-integral QM/MM simulations to reveal the significance of quantum dynamical effects (proton tunneling) in Glutathione S-Transferase (GST) and Aldehyde Dehydrogenase (ALDH) using Dynamo (www.psc.edu/biomed/dynamo/www_v2)
- MM-MD simulations of enzyme systems up to 100k atoms using CHARMM and AMBER (pmemd) are being tested



**RUNNING:
708 NODES**

Developing Faster Integration Schemes for Kohn-Sham Density Functional Theory

- Integrals that arise from the formulation of KS-DFT are of such a complicated nature that it necessary to use numerical quadrature in their solution.
- One method to increase efficiency of the numerical algorithm is to use interpolation from a Cartesian based grid to the angular grids commonly used.
 - Interpolation algorithm implemented in Q-Chem 2.1
- It is critical to analyze the convergence quality of current numerical schemes in order to assess the reliability of the interpolation algorithm
- GAMESS is being used on the Cray XT3 to validate numerical quality of the solution obtained from interpolation
 - 6-31G*/BLYP Taxol; 959 basis functions
 - 708 Cray XT3 nodes
 - Vary angular and radial grid size until convergent energy reached



**RUNNING:
700 NODES**

LSMS: Towards Petacomputing in Nanotechnology

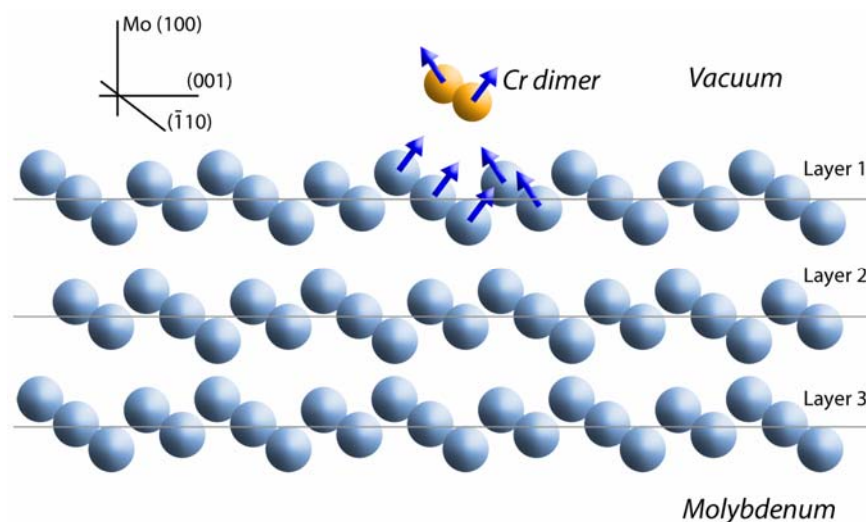
- Locally self-consistent multiple scattering (LSMS) method
 - a first-principles $O(N)$ scaling technique
 - LSMS achieves 4.65 TFlops on TCS; 1998 Gordon Bell award
- The Cray XT3 and future computing systems will enable realistic quantum mechanical simulation, e.g. study of the dynamics of magnetic switching processes, of real nanostructures.
- Planned calculations will investigate electronic and magnetic structure of a 5nm cube of Fe, which contains approximately 12,000 atoms.
- Yang Wang, PSC
Malcolm Stocks, D.M.C. Nicholson, and Markus Eisenbach, ORNL
Aurelian Rusanu and J.S. Faulkner, Florida Atlantic University



**RUNNING:
700 NODES**

First principles approach to non-collinear magnetic structure of Cr dimers on Mo(110) Surface

- Small clusters may display innovative properties in ultra small scales, suggesting promising applications.
- Cr has antiferromagnetic ordering in its bulk, described by a spin density wave with a wavelength incommensurate with the lattice constant.
- The locally self-consistent multiple scattering (LSMS) method for spin-dynamics and the full potential linearized augmented plane wave (FLAPW) method for magnetic anisotropy are adopted in the calculations.
- Yang Wang, PSC; Ruqian Wu, University of California, Irvine

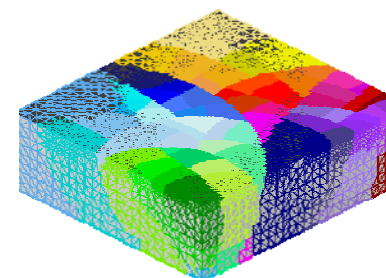
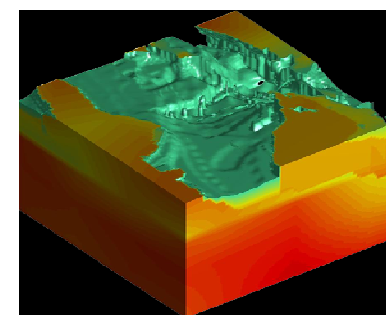
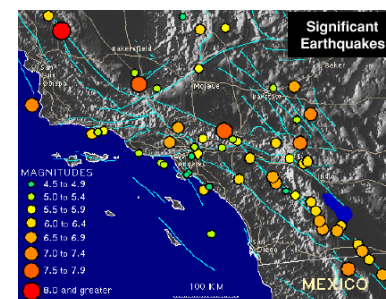


Using electronic structure calculations done with LSMS, Yang and Wu seek to understand interactions between magnetic moments of molybdenum atoms on the (110) surface and the chromium dimer.

RUNNING: 1024 NODES

Quake

- 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 64x 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 4x 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*



**RUNNING:
128 NODES**

AMBER / PMEMD

- Enabling larger molecular dynamics calculations will allow significant advances in biochemistry and structural biology.
- Simulations of 1-20ns, throttled by computational resources, are typical now. Being able to simulate 6ns/day on 128 Cray XT3 processors will allow simulations of 50-100ns to become commonplace.
- Simulation accuracy can be increased through use of polarizable force fields.
- Protein-protein interactions, critical for many biological phenomena, require simulations of sizes rarely attempted today.
- PMEMD XT3 work being done by Robert Duke (UNC-Chapel Hill), John Urbanic (PSC), Jim Maltby (Cray Inc.), and Troy Wymore (PSC).
- PSC users: Brooks, Simmerling, Cheatham, Roitberg

PMEMD: Factor IX
(90,906 atoms, constant pressure, 10,000 timesteps)

<u>Cray XT3 nodes</u>	<u>ns/day simulated</u>
16	1.29
32	2.42
64	4.26
96	5.40
128	6.00

**RUNNING:
1024 NODES**

Gasoline

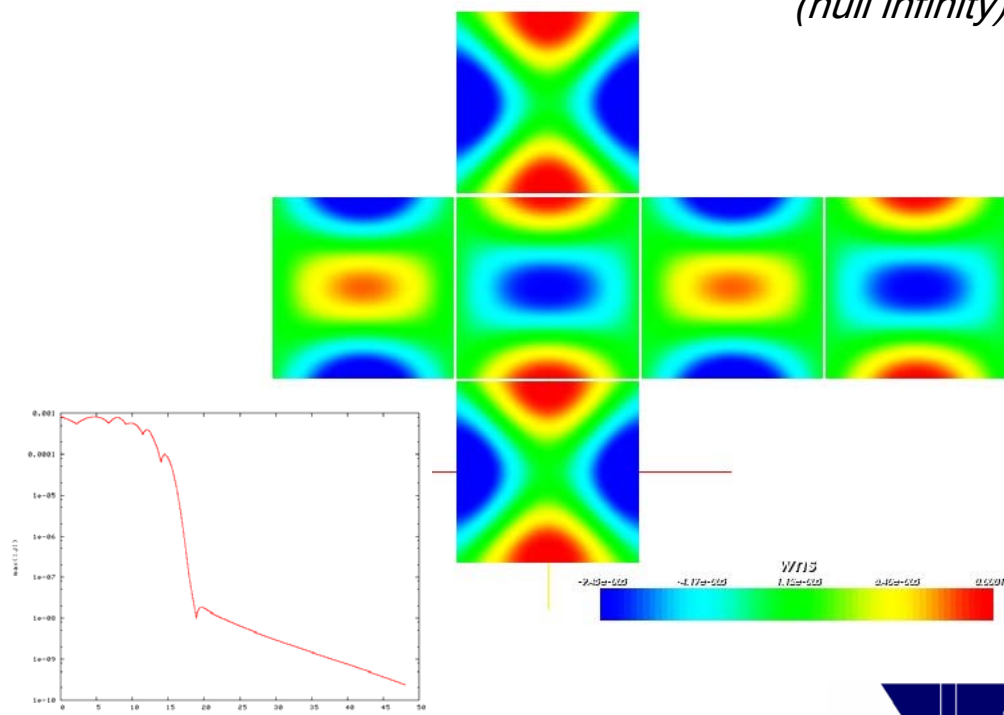
- Multi-platform, massively parallel N -body tree code used to simulate a variety of astrophysical processes
 - formation of disk (spiral) galaxies
 - tidal stirring of dwarf galaxies
 - cosmic microwave background & the Sunyaev-Zel'dovich effect
 - gas giant planet formation
 - intracluster light emission
 - galaxy cluster X-ray emissions
- Collaborators
 - George Lake, Tom Quinn, Joachim Stadel, *University of Washington*
 - James Wadsley, *McMaster University*
 - Jeffrey P. Gardner, *Pittsburgh Supercomputing Center*
 - Derek C. Richardson, *University of Maryland*

**RUNNING:
1014 NODES**

General Relativity

- Solves Einstein equations in vacuum, modeling single black hole spacetimes
- Code runs on 486 ($6 \times 9 \times 9$) nodes for 12,288 timesteps (execution time: 3h37m)
- Grid size:
 6×144^2 (angular) \times 512 (radial)
- Norm of metric fields indicates correct execution
- PI: Roberto Gomez, PSC

*Metric field W on the outermost sphere
(null infinity)*



**RUNNING:
128 NODES**

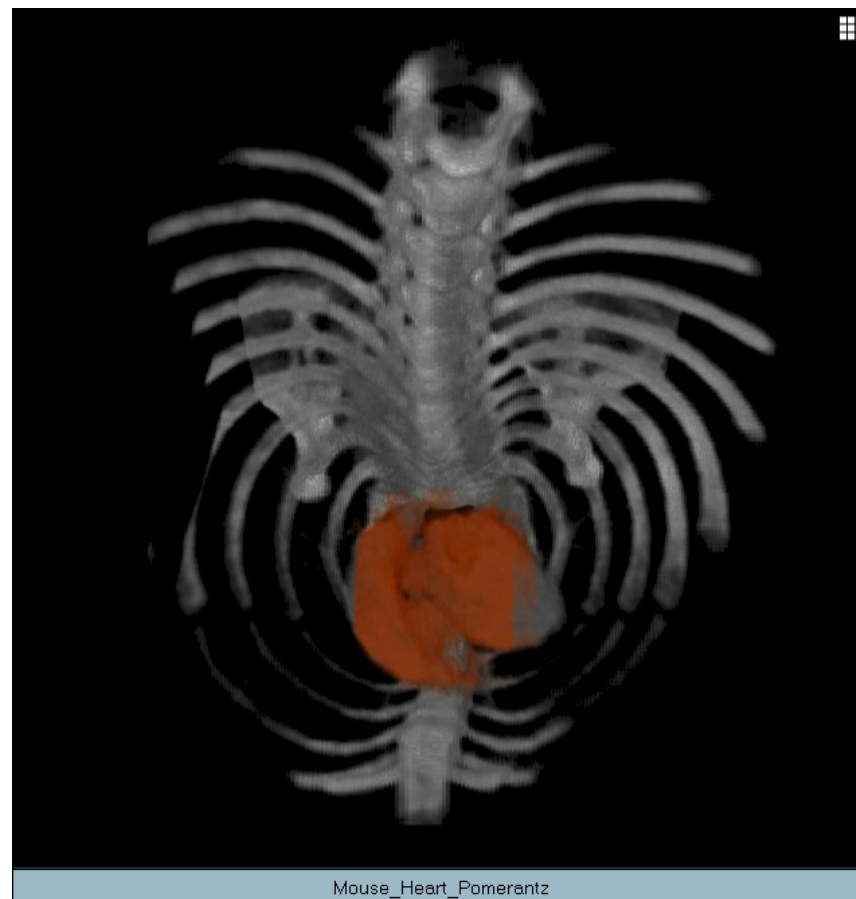
ARPS

- Advanced Regional Prediction System
 - comprehensive regional to storm-scale atmospheric modeling/prediction
 - includes real-time data analysis and assimilation, forward prediction, and post-analysis
 - observations confirm validity of ARPS simulations of intensive convective systems
- Kelvin Droegemeier and Ming Xue, *Center for Analysis and Prediction of Storms, U. of Oklahoma*

**RUNNING:
64 NODES**

Volume Rendering

- 4D volume rendering of beating mouse heart
- CT mouse data (200³) courtesy of the Duke Center for In-Vivo Microscopy
- Rendered on the Cray XT3
- Rendering and animation: Art Wetzel, Stu Pomerantz, Demian Nave (PSC)



Next Steps

- Progress developing applications and preliminary performance have been highly encouraging
 - *Developers* : External developers are already on the XT3 (Gottlieb: Dec 04; Duke: Feb 05). PSC and Cray developers are continuously at work. Other external developers will be gated in.
 - *End users* : “Friendly” users will soon be able to explore the potential of the XT3 for their science, e.g. molecular dynamics, materials science, and QCD.
 - *Performance specialists* : As Portals moves into firmware and Lustre is deployed, focused performance testing will become appropriate.
- Building on recent advances in integration, scheduling, and systems software, we expect all three phases to proceed to fruition over coming months.

