The Path to Petascale Science

Jay Boisseau, Director
Texas Advanced Computing Center
The University of Texas at Austin

Ter@tec 2007
June 20, 2007

Outline

- TACC Overview
- TeraGrid Overview
- Terascale/Petascale Computing
- Ranger: A Bridge to Petascale Computing
- Closing Thoughts and Summary

TACC Mission

To enable scientific discovery and enhance society through the application of advanced computing technologies.

TACC Strategic Approach

- Resource & Services
  - Evaluate, acquire & operate world-class resources
  - Provide expert support via leading technology expertise
- Research & Development
  - Produce new computational technologies and techniques
  - Collaborate with researchers to apply advanced computing technologies
- Education & Outreach
  - Inform public about impact of advanced computing
  - Educate students to increase participation in advanced computing careers

TACC Technology Focus Areas

- High Performance Computing (HPC)
  - Performance benchmarking, analysis, optimization
  - Linear algebra, solvers
  - CFD, computational chemistry, weather/ocean modeling, computational biomedicine
- Data & Information Analysis (DIA)
  - Scientific visualization
  - Data collections management
  - Data analysis & mining
- Distributed & Collaborative Computing (DCC)
  - Portals & gateways
  - Middleware for scheduling, workflow, orchestration

TACC HPC & Storage Systems

LONESTAR

Dell Linux Cluster
2900+ dual-core CPUS,
>11 TB memory, >60 Tflops peak

CHAMPION

IBM Power5 System
96 Power5 CPUs,
192 GB memory, ~1 teraflop

ARCHIVE

STK PowderHorns (2), managed by Cray DMF
2.8 PB max capacity

GLOBAL DISK

Sun SANs and
Data Direct Disk
> 50TB

TACC Advanced Visualization Systems

- *Maverick:* Sun Terascale Visualization System
  - 128 UltraSparc 4 cores, ½ TB memory
  - 16 GPUs, > 3 Gpoly/sec
- Also: SGI Prism, Dell Cluster, Workstations
- Immersive and tiled displays
  - 3x1 semi-cylinder immersive environment
    - immersive capabilities with head/motion tracking
  - 5x2 large-screen, 16:9 panel tiled display
  - 3x3 tiled 30” LCD display

TACC R&D – High Performance Computing

- Scalability, performance optimization, and performance modeling for HPC applications
- Evaluation of cluster technologies for HPC
- High performance linear algebra, solvers
- Computational fluid dynamics
- Computational chemistry
- Climate, weather, ocean modeling collaboration and support of DoD

TACC R&D – Data & Information Analysis

- Remote/collaborative interactive visualization
- Feature detection / terascale data analysis
- Hardware accelerated visualization and computation on GPUs
- Creating/hosting scientific data collections, analysis services

TACC R&D – Distributed & Collaborative Computing

- Web-based grid portals
- Grid scheduling and workflow tools
- Large-scale distributed computig
- Overall grid deployment and integration

"Scientific Computing Curriculum" Academic Classes

- Teach *applied* use of advanced computing technologies and techniques
- Comprehensive four-course curriculum:
  - Introduction to Scientific/Technical Computing
  - Parallel Computing for Science & Engineering
  - Visualization & Data Analysis for Science & Engineering
  - Distributed & Grid Computing for Science & Engineering
- Taught through UT CS department but also *cross-listed in science/engineering departments*
- Class materials available for download now
- Will record and post lectures in 2008, and teach to remote users in 2009(?)

Strategic Focus Activities in 2007+

- Petascale Computing
  - Integration, management, and operation of large-scale systems
  - Performance optimization for multi-core processors
  - Achieving extreme scalability: algorithms, libraries, community codes, frameworks, etc.
  - Fault tolerance for applications on large systems
- Petascale Visualization & Data Analysis
  - ‘In-simulation’ visualization, HPC visualization applications
  - Remote & collaborative visualization
  - Feature detection techniques
- Remote, collaborative access to petascale simulation and analysis capabilities
  - Data collections hosting with layered analysis services
  - Portals and gateways for communities, community applications

TeraGrid Overview

TeraGrid Mission

- TeraGrid provides integrated, persistent, and pioneering computational resources that will significantly improve our nation's ability and capacity to gain new insights into our most challenging research questions and societal problems.
  - Our vision requires an integrated approach to the scientific workflow including obtaining access, application development and execution, data analysis, collaboration and data management.
  - *These capabilities must be accessible broadly to the science, engineering, and education community.*

TACC

Why Science Requires Cyberinfrastructure

- Inherent **complexity** and **multi-scale** nature of todays frontier science challenges.
  - Requires **multi**-{disciplinary,investigator,institutional} approach (often international).
- High **data intensity** from simulations, digital instruments, sensor nets, observatories.
- Increased **value of data** and demand for data curation & preservation of access.
- Infrastructure **sharing** to achieve better stewardship of research funding.

Adapted from: Dan Atkins, NSF Office of Cyberinfrastructure

Who Uses TeraGrid?

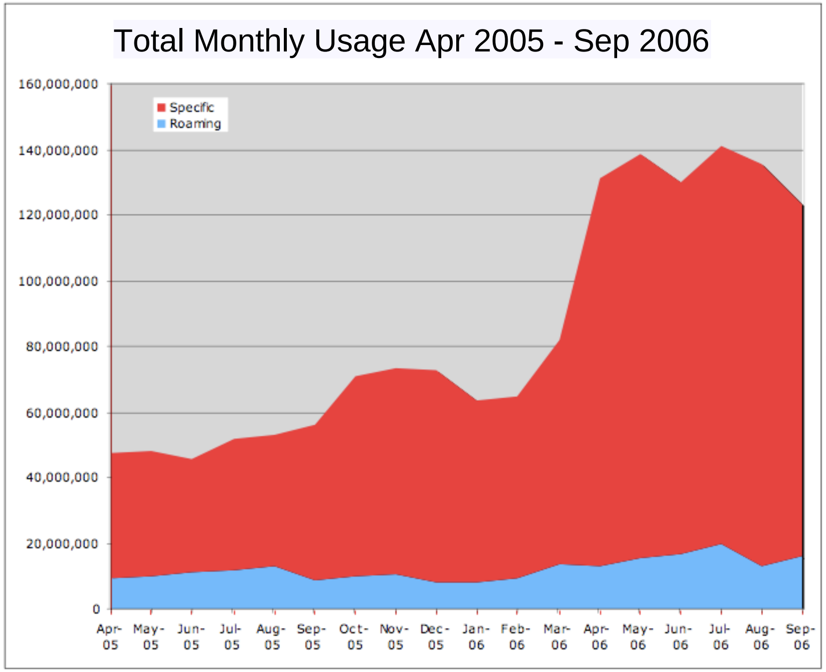

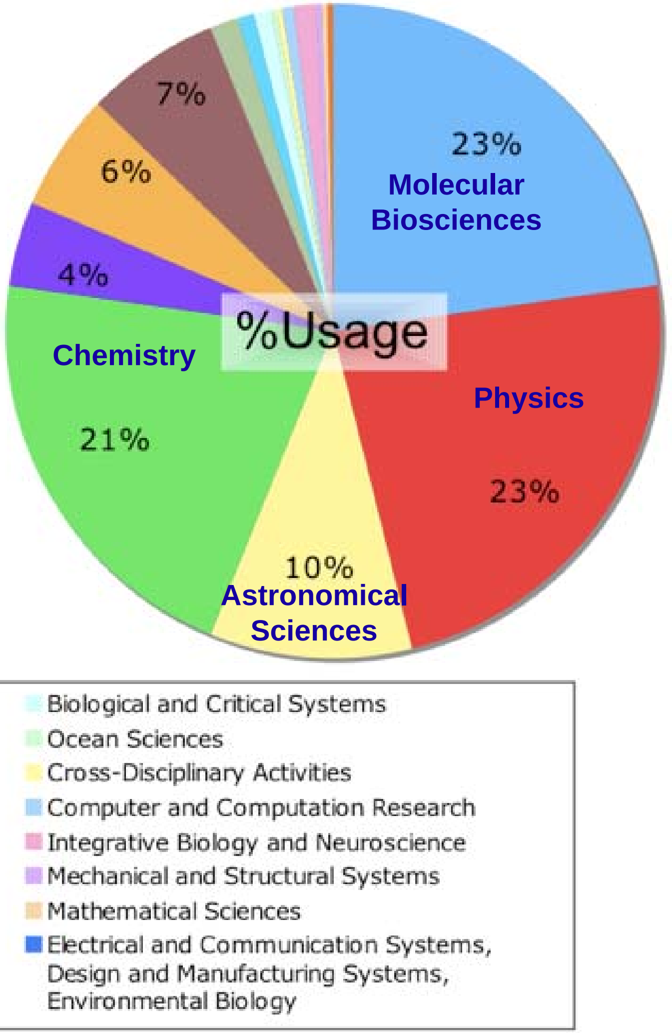

TACC

1000+ projects, 3000+ users

TeraGrid Projects by Institution

Blue: 10 or more PI's
Red: 5-9 PI's
Yellow: 2-4 PI's
Green: 1 PI

FREE Map Code by MapBuilder.net

Map data ©2006 Tele Atlas - Terms of Use

TeraGrid allocations are available to researchers at any US educational institution by peer review. Exploratory allocations can be obtained through a biweekly review process. See www.teragrid.org.

TACC

1000+ projects, 3000+ users

TeraGrid Science Examples

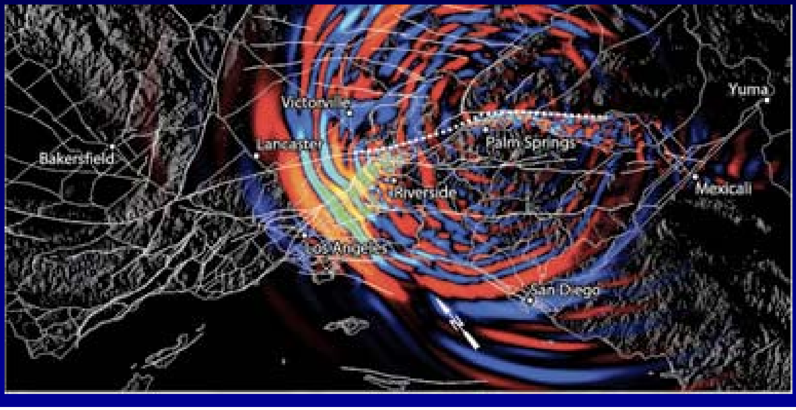

The Southern California Earthquake Center uses multiple TeraGrid resources to simulate and analyze earthquakes.

John Rust (UMd) has used TeraGrid for economic analysis and forecast.

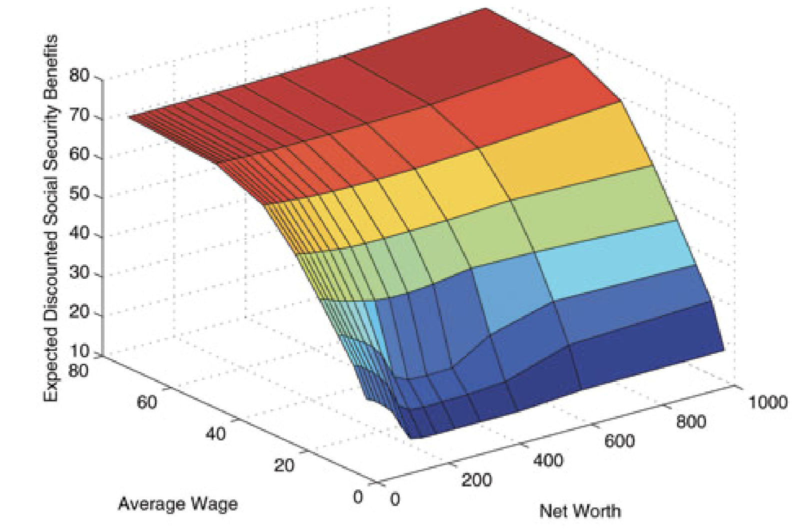

See www.teragrid.org for more science impact stories.

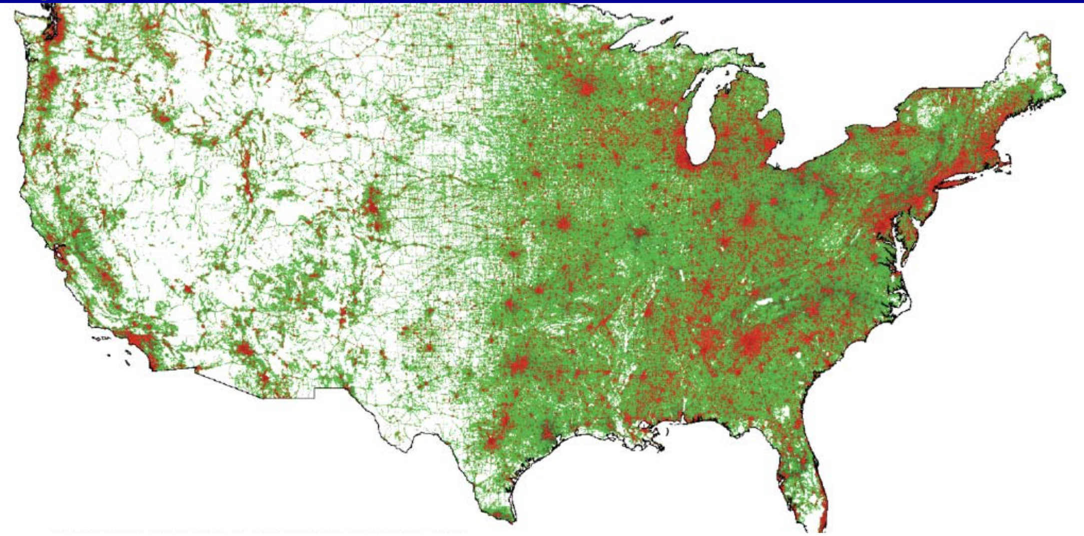

NIH MIDAS project has begun to use TeraGrid systems at NCSA to simulate pandemic mitigation strategies.

Klaus Schulten models protein binding using TeraGrid.

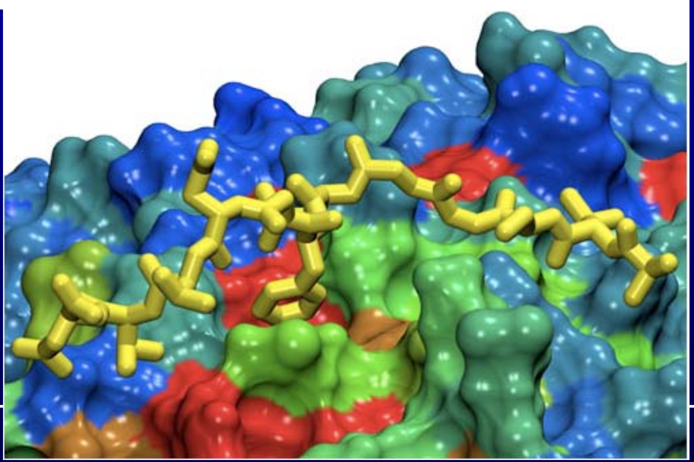

TeraGrid User Portal

System Accounts

Resource	Username
IU	
tg-login-ia32.iu.teragrid.org	tg-catle
tg-login1.iu.teragrid.org	tg-catle
NCSA	
login-co.ncsa.teragrid.org	catlett
login-cu.ncsa.teragrid.org	catlett
tg-login.ncsa.teragrid.org	catlett
login-w.ncsa.teragrid.org	catlett
ORNL	
tg-login.ornl.teragrid.org	catlett
PSC	
tg-login1.lemieux.psc.teragrid.org	ccatlett
tg-login.rachel.psc.teragrid.org	ccatlett
tg-login.bigben.psc.teragrid.org	*no account*
Purdue	
tg-login.purdue.teragrid.org	catlett
SDSC	
bglogin.sdsc.edu	catlett
dslogin.sdsc.edu	catlett
tg-login.sdsc.teragrid.org	catlett
TACC	
tg-login.tacc.teragrid.org	catlett
maverick.tacc.utexas.edu	catlett
UC/ANL	
tg-login.uc.teragrid.org	catlett

Allocation Usage

Start Date	End Date	Resource	Project Allocation (SU) Remaining / Awarded	My Usage (SU)	Alloc. Type
Project Title: TG Staff Project: Project Management **Charge No.:** TG-STA040014N **Grant No.:** STA040014N **Project PI?** No					
2004-10-19	2013-12-31	teragrid_roaming	99999 /99999	0.0	new
2003-12-19	2013-12-31	teragrid	99999 /99999	0.0	new
Project Title: TG RP UC/ANL **Charge No.:** TG-STA060016N **Grant No.:** STA060016N **Project PI?** No					
2006-05-19					

Tracking usage for my allocations

TeraGrid™ User Portal

Logout
Welcome, Charles E Catlett

Home | My TeraGrid | Resources | Documentation | Consulting | Allocations

Systems Monitor | Science Gateways | Data Collections

TeraGrid Systems Monitor

Refresh

High Performance Computing Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Load	Jobs* R	Jobs* Q	Jobs* O
Lonestar	TACC	Dell PowerEdge Linux Cluster	5200	55.00	10.40	94.90		186	64	0
Big Red	IU	IBM e1350	2048	20.40	8.20	266.00		1	126	1892
Tungsten	NCSA	Dell Xeon IA-32 Linux Cluster	2560	16.38	3.75	109.00		76	1149	34
DataStar p655	SDSC	IBM Power4+ p655	2176	14.30	5.75	115.00		34	48	33
TeraGrid Cluster	NCSA	IBM Itanium2 Cluster	1744	10.23	4.47	60.00		57	430	70
Bigben	PSC	Cray XT3	2090	10.00	2.02	48.80		1	32	59
Lear	Purdue	Dell EM64T Linux Cluster	1024	6.60	2.00	28.00		218	385	4
Cobalt	NCSA	SGI Altix	1024	6.55	3.00	100.00		33	262	12
Lemieux	PSC	HP Alpha Cluster	3000	6.00	2.93	78.13		46	2	38
Blue Gene	SDSC	IBM Blue Gene	2048	5.70	0.50	19.50		3	5	7
TeraGrid Cluster	SDSC	IBM Itanium2 Cluster	524	3.10	1.02	48.80		6	246	68
Copper	NCSA	IBM Power4 p690	384	2.00	1.44	30.00		88	108	2
DataStar p690	SDSC	IBM Power4+ p690	192	1.30	0.88	115.00		18	3	12
TeraGrid Cluster	UC/ANL	IBM Itanium2 Cluster	124	0.61	0.24	4.00		7	10	0
NSTG	ORNL	IBM IA-32 Cluster	56	0.34	0.07	2.14		0	0	0
Rachel	PSC	HP Alpha SMP	128	0.31	0.50	6.00		20	75	1
		Total:	**24322**	**158.82**	**47.17**	**1125.27**		**794**	**2945**	**2232**

Advanced Visualization Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Graphics HW
TeraGrid Cluster	UC/ANL	Intel Xeon Cluster	192	0.61	0.38	4.00	nVIDIA GeFORCE 6600GT AGP graphics cards
Maverick	TACC	Sun E25K	128	0.27	0.50	0.56	16 nVIDIA QuadroFX 3000G graphics cards
		Total:	**320**	**0.88**	**0.88**	**4.56**	

Managing Credentials

Current State of all Resources

TACC

TeraGrid Science Gateways Initiative: Service-Oriented Approach

The science and engineering community has been building discipline-specific cyberinfrastructure in the form of portals, applications, and grids. Our objective is to enable these to use TeraGrid resources transparently as “back-ends” to their infrastructure.

The TeraGrid Science Gateways program has developed, in partnership with 20+ communities and multiple major Grid projects, an initial set of processes, policies, and services that enable these gateways to access TeraGrid (or other facilities) resources via web services.

TACC

TeraGrdid Science Gateway Projects

Science Gateways

Below is a complete list of current science gateways, to see a detailed project description please click on the name science gateway.

Title	Field of Science	Portal Homepage
Biology and Biomedicine Science Gateway	Molecular Biosciences	Visit Portal
Computational Chemistry Grid	Chemistry	Visit Portal
Computational Science and Engineering Online	Chemistry	Visit Portal
GEON(GEOsciences Network)	Earth Sciences	Visit Portal
GIScience Gateway	Geography and Regional Science	Visit Portal
Grid Analysis Environment	Physics	N/A
Linked Environments for Atmospheric Discovery	Atmospheric Sciences	Visit Portal
National Biomedical Computation Resource	Integrative Biology and Neuroscience	Visit Portal
National Virtual Observatory	Astronomical Sciences	Visit Portal
Network for Computational Nanotechnology and nanoHUB	Emerging Technologies Initiation	Visit Portal
Network for Earthquake Engineering Simulation	Earthquake Hazard Mitigation	Visit Portal
Neutron Science Instrument Gateway	Physics	Visit Portal
Open Life Sciences Gateway	Molecular Biosciences	Visit Portal
Open Science Grid	Advanced Scientific Computing	N/A
SCEC Earthworks Project	Earthquake Hazard Mitigation	Visit Portal
Special PRiority and Urgent Computing Environment	Advanced Scientific Computing	Visit Portal
TeraGrid Visualization Gateway	Visualization, Graphics, and Image Processing	Visit Portal
The Earth System Grid	Global Atmospheric Research	Visit Portal
The Telescience Project	Neuroscience Biology	Visit Portal
Virtual Laboratory for Earth and Planetary Materials	Materials Research	Visit Portal

For more information on the science gateways effort please visit the Science Gateways program page.

TeraGrid User Community in 2006

Use Modality	Community Size (est. number of projects)
Batch Computing on Individual Resources	850
Exploratory and Application Porting	650
Workflow, Ensemble, and Parameter Sweep	160
Science Gateway Access	100
Remote Interactive Steering and Visualization	35
Tightly-Coupled Distributed Computation	10

Terascale/Petascale Computing

The Terascale Era

- June 1997:
  - "ASCI Red" (Sandia) entered TOP500 list of most powerful supercomputers at #1 with 1.07 TF/s
  - Held position until Nov. 2000
- As of Nov. 2006:
  - #1 is LLNL (NNSA/DOE) IBM Blue Gene: 367 Tflops peak, 280 Tflops HPC
  - #500 machine nearly 5 Tflops peak (nearly 3 Tflops HPL)
  - 11 countries represented in top 50
- IBM Thinkpad T43p would have been in the top 100 on first TOP500 list (1993)

The Case for More Powerful Computational Science Capabilities

- Many recent federally-commissioned reports have urged sustained, long-term U.S. investments in HPC to realize full benefits of computational science:
  - NSF: Cyberinfrastructure (2003)
  - DOE: Facilities for the Future of Science (2003)
  - NIH: Biomedical Information Science and Technology Initiative
  - Council on Competitiveness: Supercharging U.S. Innovation and Competitiveness (2004)
  - Interagency: High End Computing Revitalization Task Force (2004)
  - DOE: Capability Computing Needs (2004)
  - NAS: Getting up to Speed: The Future of Supercomputing (2005)
  - PITAC: Report on Computational Science (2005)
  - NSF: Simulation-Based Engineering Science (2005)

The Case for More Powerful Computational Science Capabilities

- National Academies report “Rising Above the Gathering Storm” urges reinvestment in Science/Technology/Engineering/Math
- American Competitiveness Initiative calls for doubling of NSF, DOE/SC, NIST budgets over 10 years; largest federal response since Sputnik
  - identifies petascale computing for modeling and simulation as one of 12 goals
- NSF 5-year Strategic Plan fosters research to further U.S. economic competitiveness by focusing on fundamental science & engineering
  - Advance fundamental research in computational science and engineering, and in fundamental, applied and interdisciplinary mathematics and statistics
- NSF Cyberinfrastructure Vision for 21st Century Discovery lays blueprint for investments in CI, including HPC, data, collaboration, workforce development

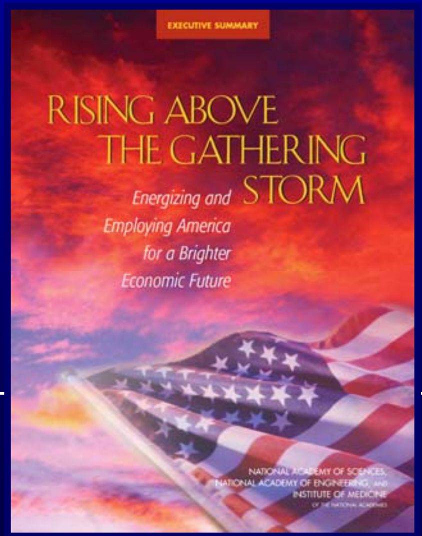

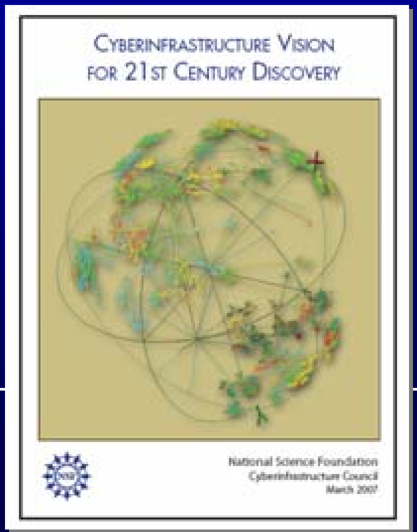

The Petascale Era

- DOE, NSF, and other US agencies now aggressively pursuing programs to deploy
  - peak petaflops systems *now*
  - *sustained* petaflops systems in the next 4 years
- A few US petascale projects
  - NSF Track2 systems deployed annually (2007-11)
  - DOE NNSA Roadrunner system @ LANL (1 PF+, 2008/09)
  - DOE Office of Science systems @ ORNL, ANL (1 PF, 2008/09)
  - NSF Track 1 Petascale Acquisition (10-20 PF, 2011)
- Cost of hardware/ & operations for NSF and DOE *sustained* petaflops systems alone: >$1B

Petascale Computing Opportunities

- Petascale will be here next year: up to science & engineering communities to make effective use
- Modeling and simulation can contribute significantly to making headway on many of the ‘grand challenge’ problems facing ***society*** as well as science:
  - future energy, climate change, environmental sustainability, clean water, natural disasters, neuroscience, drug design, predictive medicine, intelligent manufacturing, supply chain management, first principles materials design, etc.
- Petascale systems preset unprecedented capabilities, opportunities to make headway on many of the societal grand challenges

Petascale Computing Opportunities

- Raw throughput/memory will permit many enhancements to current “terascale” simulations:
  - Increased resolution
  - Greater fidelity of physics models
  - Inverse problem (a.k.a. model calibration, parameter estimation, data assimilation)
  - Uncertainty quantification
  - Optimization (design and control)
- Ultimately: simulation-based decision-making under uncertainty
  - Likely an exascale (zetascale, yottascale) computing problem for terascale deterministic forward problems

The Billion Dollar Question:

Will we be able to make effective use of PF systems?

- Enormous challenges for petascale computational science:
  - Mathematical models
  - Numerical approximations
  - Scalable numerical algorithms
  - Scalable geometric algorithms
  - Scientific visualization and data management
- Petascale computing challenges been underappreciated at agency levels for the past 15 years, still remain to be solved
  - Major troubles ahead unless sufficient resources are aimed at creating "scalable computational science"
- Indications of change - Example: NSF is planning a 5-year, $0.75B program: Cyber-enabled Discovery and Innovation (CDI)

TACC

NSF Vision, Strategic Plan Recognize Importance of HPC

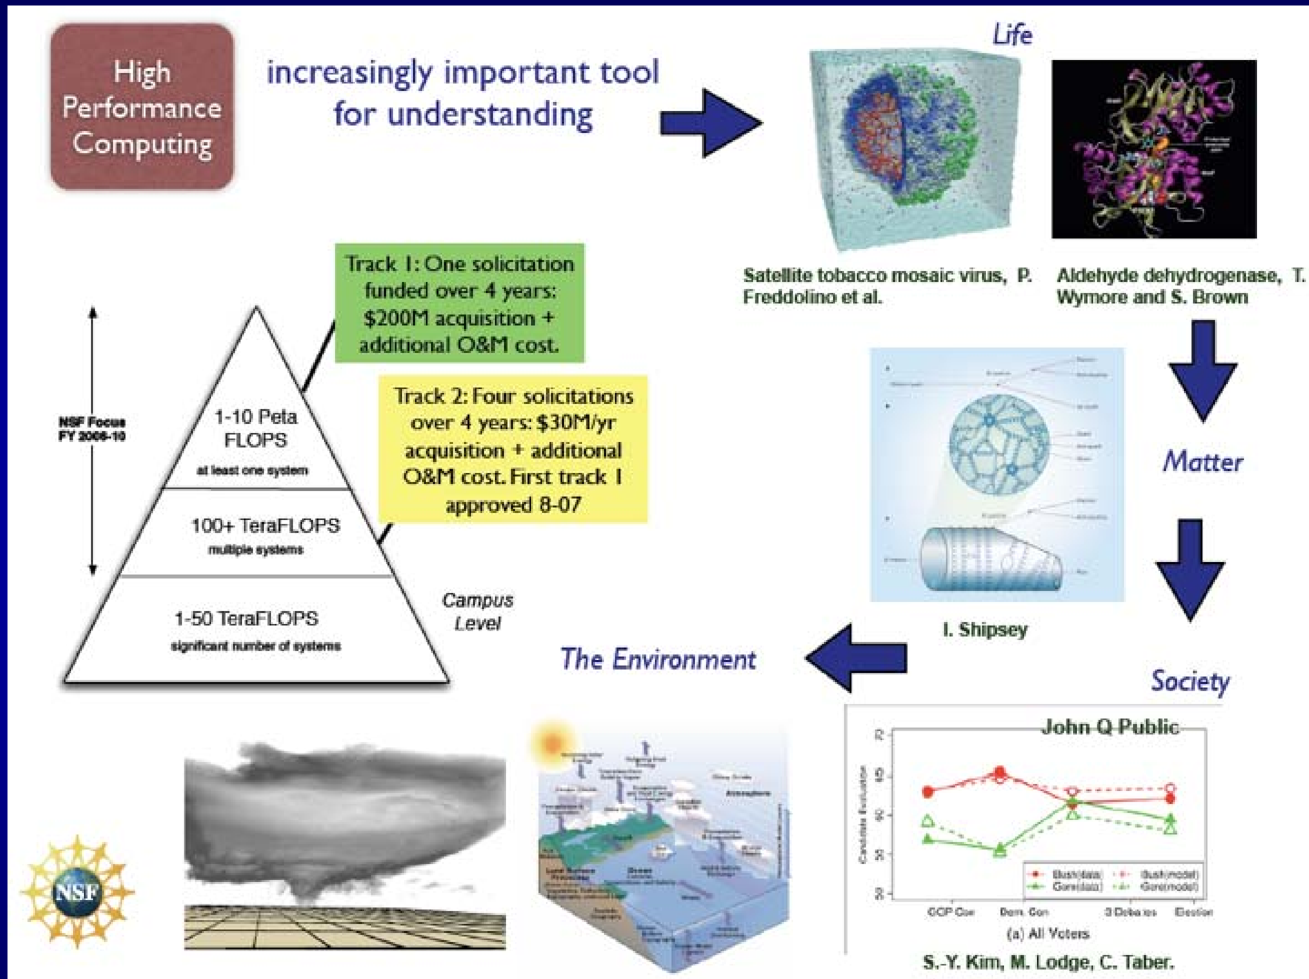

Adapted from: Dan Atkins, NSF Office of Cyberinfrastructure

NSF Cyberinfrastructure Strategic Plan

- NSF Cyberinfrastructure Strategic Plan just released!
  - Articulates importance of CI overall
  - Chapters: HPC, data, collaboration, workforce development
- NSF investing in world-class HPC
  - Annual “Track2” solicitations ($30M)
  - Five-year “Track1” solicitation ($200M)
- Complementary solicitations out or forthcoming
  - Software Development for CI (SDCI)
  - Strategic Technologies for CI (STCI)
  - Petascale applications development late 2007
  - Cyber-enabled Discovery & Innovation (CD) 2008

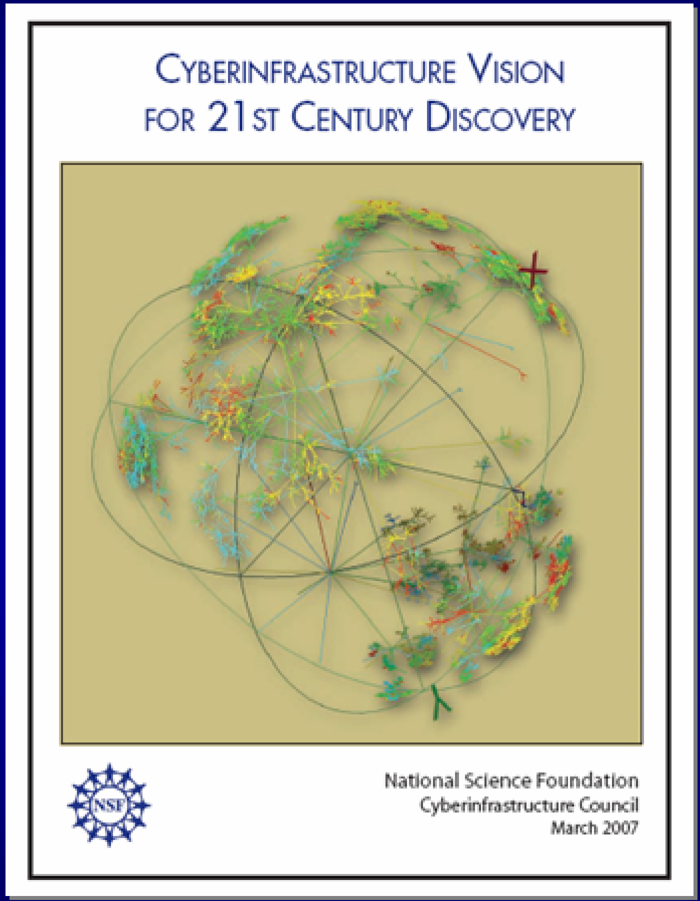

http://www.nsf.gov/od/oci/CI_Vision_March07.pdf

Ranger: A Bridge to Petascale Computing

First NSF Track2 System: 1/2 Petaflop!

- TACC selected for first NSF ‘Track2’ HPC system
  - $30M system
  - Sun is integrator
  - 15,744 quad-core AMD Opterons
  - ~ 1/2 petaflops peak performance
  - 1.7 petabytes disk
  - 125 TB memory
  - ~2 μsec MPI latency

- TACC, ICES, CS (all at UT) plus Cornell, ASU supporting system, users four 4 years ($29M)

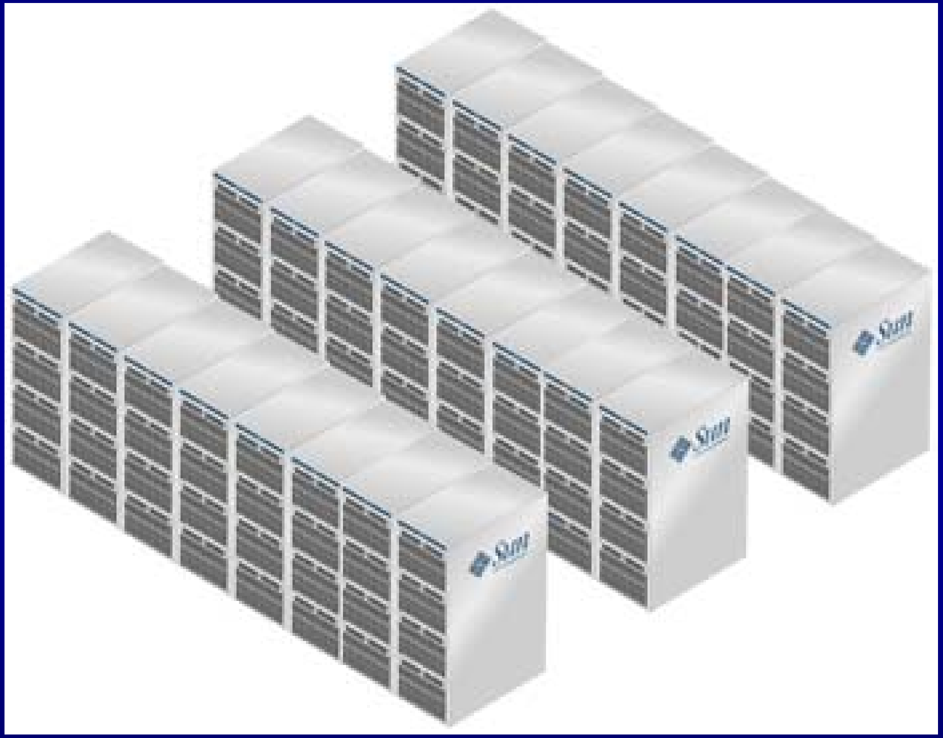

Sun System Configuration

- Compute power
  - 15744 Opteron “Barcelona” processors
    - Quad-core: 62,976 cores!
    - Four flops/cycle (dual pipelines) per core
  - 504 teraflops aggregate peak performance (currently…)
- Memory
  - 2GB/core
  - 125 TB total memory
- Expandable
  - May add more compute nodes (may vary memory)
  - May add different compute nodes (GPUs?)

TACC

Ranger Configuration

[Most switch data non-disclosure]

- Interconnect
  - Sun proprietary switch based on IB
    - Minimum cabling: robustness and simplicity!
    - SDR implementation initially
  - MPI latency:
    - 1.6-1.8 μsec per switch
    - max 2.3 μsec across system
  - Peak bi-directional b/w: ~ 1 GB/sec
  - Peak bisection b/w: 7.9 TB/sec

Ranger Configuration

- File system
  - 72 Sun X4500s (“Thumper”)
    - 48 disks per 4U
  - 1.7 PB total disk
    - 3456 drives total
    - 1 PB in largest /work file system
  - Lustre file system
  - Aggregate b/w: 40 GB/s

Ranger Configuration

[Some data non-disclosure]

- System Management
  - ROCKS (customized) Cluster Kit
    - *perfctr* patch, etc.
  - Sun N1SM for lights-out management
  - Sun N1GE for job submission
    - Backfill, fairshare, reservations, etc.

Ranger Space & Power

- System power: 2.4 MW
- System space
  - ~80 racks
  - ~2000 sqft for system racks and in-row cooling equipment
  - ~4500 sqft total
- Cooling:
  - In-row units and chillers
  - ~0.6 MW
- Observations:
  - space less an issue than power (almost 3 MW)!
  - power generation distribution less an issue than distribution!

Applications Performance Notes

- Obviously, no data for final system
  - Switch doesn’t exist yet
  - Processors don’t exist yet
- Performance predictions can be made from previous & pre-production versions, prototypes, etc.
- Applications performance projections for NSF benchmarks look **very** promising
- Expect some applications to sustain 50-100+ Tflops
  - On very large problem sizes: up to 100 TB!

User Support Challenges

- NO systems like this exist yet!
  - Will be the first general-purpose system at ½ Pflop
  - Quad-core, massive memory/disk, etc.
- NEW apps challenges, opportunities
  - Multi-core optimization
  - Extreme scalability
  - Fault tolerance in apps
  - Petascale data analysis
- System cost $50K/day--must enable user to conduct world-class science every day!

User Support Plans

- User support: the “usual” +
  - User Committee dedicated to this system
  - Applications Engineering
    - algorithmic consulting
    - technology selection
    - performance/scalability optimization
    - data analysis
  - Applications Collaborations
    - Partnership with petascale apps developers and software developers

User Support Plans

- Also
  - Strong support of ‘professionally optimized’ software
    - Community apps
    - Frameworks
    - Libraries
  - *Extensive* Training
    - On-site at TACC, partners, and major user sites, and at workshops/conferences
    - Advanced topics in multi-core, scalability, etc
    - Virtual workshops
  - Increased contact with users in TACC User Group

Technology Insertion Plans

- Technology Identification, Tracking, Evaluation, Recommendation are crucial
  - Cutting edge system: software won't be mature
  - Four year lifetime: new R&D will produce better technologies
- Chief Technologist for project, plus other staff
  - Must build communications, partnerships with leading software developers worldwide
  - Grant doesn't fund R&D, but system provides unique opportunity for determining, conducting R&D!

Technology Insertion Plans

- Aggressively monitor, and pursue:
  - NSF Software Development for Cyberinfrastructure (SDCI) proposals
  - NSF Strategic Technologies for Cyberinfrastructure (STCI) proposals
  - NSF Cyber-enabled for Discovery and Innovation (CDI) proposals (forthcoming)
  - Relevant NSF CISE proposals
  - Corresponding awards in DOE, DOD, NASA, etc.
- Some targets: fault tolerance, algorithms, next-generation programming tools/languages, etc.

Impact in TeraGrid, US

- 500M+ CPU hours to TeraGrid: more than double current total of all TG HPC systems
- 500+ Tflops : almost 10x current top system
- Enable unprecedented research
- Re-establish NSF as a leader in HPC
- *Jumpstarts progress to petascale for entire US academic research community*

TeraGrid HPC Systems plus Ranger

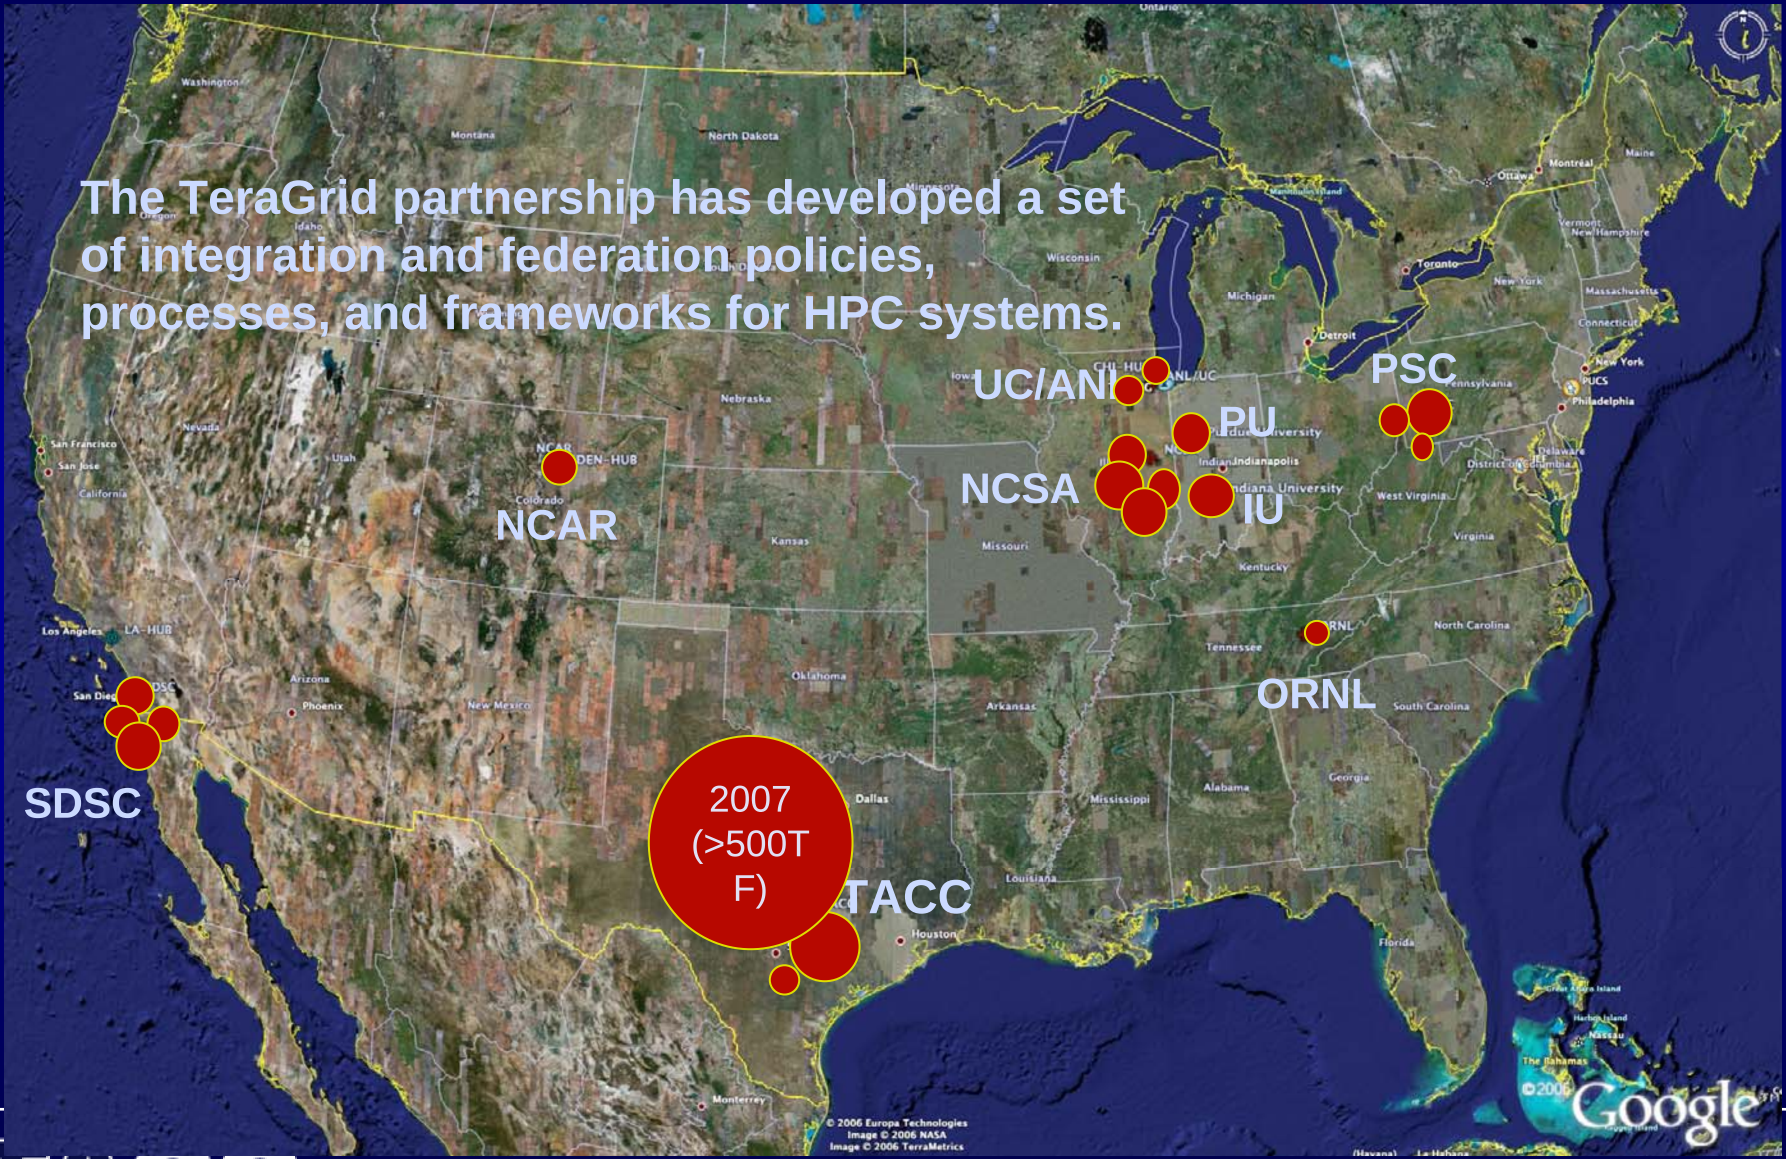

Computational Resources (size approximate - not to scale)

Who Might Use Ranger?
Past TeraGrid HPC Usage by Domain

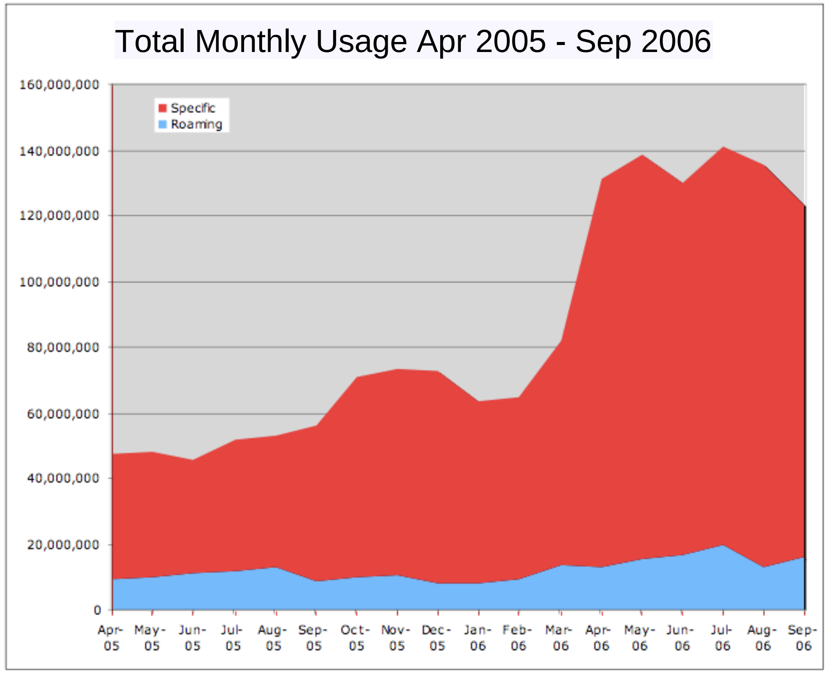

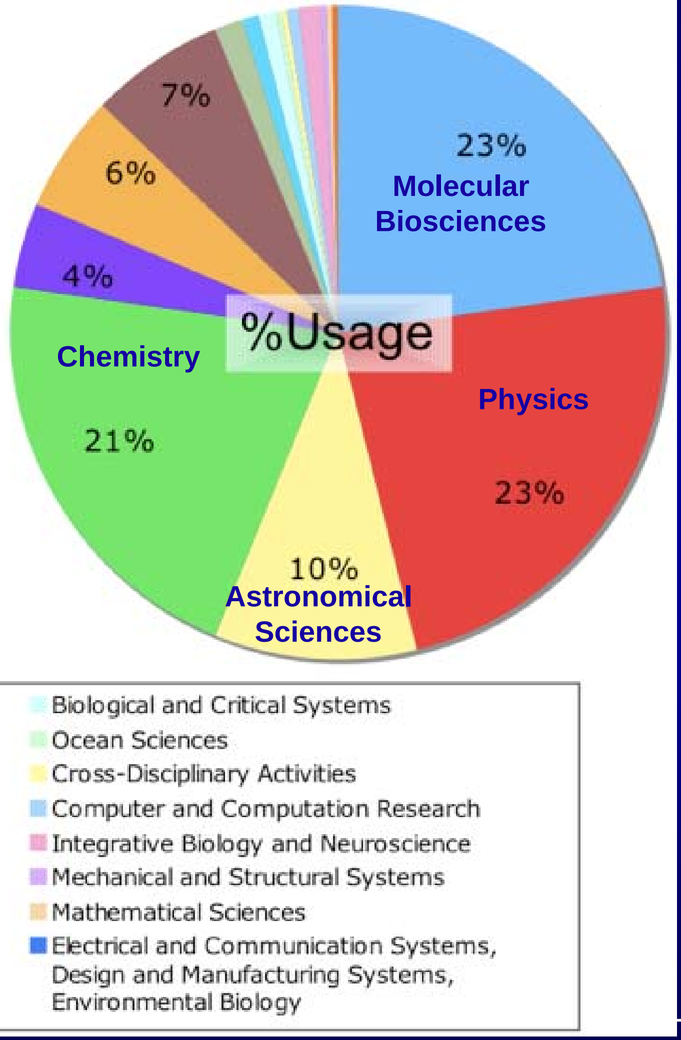

1000 projects, 3200 users

Ranger Project Timeline

Sep06	award, press, relief, beers
1Q07	equipment begins arriving
2Q07	facilities upgrades complete
2Q-3Q07	construction of system
3Q07	very friendly users
4Q07	more early users
Dec07	production, many beers
Jan08	allocations begin

What’s Next for NSF HPC?

- NSF:
  - Another Track2 system every year
    - TACC won first, 2nd is being decided right now
    - Rumors of three semi-finalists
    - Rumors of 1 PF peak on at least one option!
  - Track1 award made this year for 2011-2016
    - To ***sustain*** 1 PF
    - Four competitors
    - Each option almost certainly > 10 PF peak
    - Each competitor has extensive team or partners

TACC

Some “Peta-Challenges”

- Achieving performance on many-core
  - Processor/memory bandwidth gap increasing
- Scalable algorithms
  - To 10K-100K+ cores
  - also, must be effectively implemented
- Scalable programming tools
  - debuggers, optimization tools, libraries, etc.
- Fault tolerance
  - Increased dependence on commodity (MTBF/node not changing) and increased number of nodes -> uh oh!

Petascale Data Analysis Challenge

- Data analysis ‘in the box’
  - Data will be too big to move (network, file system bandwidths not keeping pace)
  - Analyze in simulation if able
  - Or at least analyze while data still in HPC parallel file system
  - Must develop CPU-based scalable techniques
  - Or must develop better packaging for GPUs, include on more nodes

Petascale Power Challenge

- Power constraints--generation and distribution--limit number and location of petascale computing centers
  - Remember: flops/watt I getting better, but we're building much larger systems!
  - Track1 system power budget will be more than staffing (2x!)
  - But HPC expertise becomes even more important than hosting expertise due to other challenges

Some Predictions

- Next NSF Track2 will be also homogeneous, but 3rd or 4th will not (some Cell, GPGPU, or…)
  - But not solely Cell or GPGPU at petascale!
  - Los Alamos building hybrid petascale Opteron-Cell system in 2008!
- Commodity switches will increase in port count greatly (thousand-way+) very soon (2008?)
- Serious *community* efforts on optimizing Linux for many-core compute nodes (not just vendor-specific)
- Lightweight checkpoint restart for Linux clusters
- Leading centers limited by location, infrastructure, but become islands: host compute, data, vis, etc.

Summary

- Push to petascale is driving HPC vendors like the push to 1 GHz drove AMD, Intel
- NSF determined to be a leader in petascale computing as component of world-class CI
  - solicitations for hardware, software, support, applications
- Ranger and other forthcoming NSF petascale systems (and software, and apps) will enable unprecedented high-resolution, high-fidelity, multi-scale, multi-physics applications
- It is an incredibly exciting time to be involved supercomputing again!   (And our jobs are safe)

Thanks To…

- The National Science Foundation (NSF) for giving TACC the opportunity to deploy Ranger and help the science community move to petascale computing
- Omar Ghattas, Charlie Catlett, Karl Schulz and Tommy Minyard for many contributions to this presentation
- Christian Saguez, Jacques Duysens, and many others for being such excellent hosts!

The University of Texas at Austin
Distinguished Lecture Series in Petascale Computation

- Web accessible: http://www.tacc.utexas.edu/petascale/
- Past Lectures
  - “Petaflops, Seriously,” Dr. David Keyes, Columbia University
  - “Discovery through Simulation: The Expectations of Frontier Computational Science,” Dr. Dimitri Kusnezov, National Nuclear Security Administration
  - “Modeling Coastal Hydrodynamics and Hurricanes Katrina and Rita,” Dr. Clint Dawson, The University of Texas at Austin
  - “Towards Forward and Inverse Earthquake Modeling on Petascale Computers,” Dr. Omar Ghattas, The University of Texas at Austin
  - "Computational Drug Diagnostics and Discovery: The Need for Petascale Computing in the Bio-Sciences,” Dr. Chandrajit Bajaj, The University of Texas at Austin
  - "High Performance Computing and Modeling in Climate Change Science,” Dr. John Drake, Oak Ridge National Laboratory
  - "Petascale Computing in the Biosciences - Simulating Entire Life Forms," Dr. Klaus Schulten, University of Illinois at Urbana-Champaign
- Suggestions for future speakers/topics welcome