National Leadership Computing Facility *- Bringing Capability Computing to Science*

Jeff Nichols, Director

Computer Science and Mathematics Division,
National Center for Computational Sciences
Oak Ridge National Laboratory

Frontiers of Extreme Computing
October 24-27, 2005
Santa Cruz, CA

Outline

- Who are we: National Leadership Computing Facility (NLCF) at Oak Ridge National Laboratory (ORNL)
- Science Motivators (first results)
- Allocations and Support for Break-Through Science
- 2006-2010 Roadmap
- Wrapup

Outline

- Who are we: NLCF at ORNL
- Science Motivators (first results)
- Allocations and Support for Break-Through Science
- 2006-2010 Roadmap
- Wrapup

ORNL is DOE’s largest multipurpose science laboratory

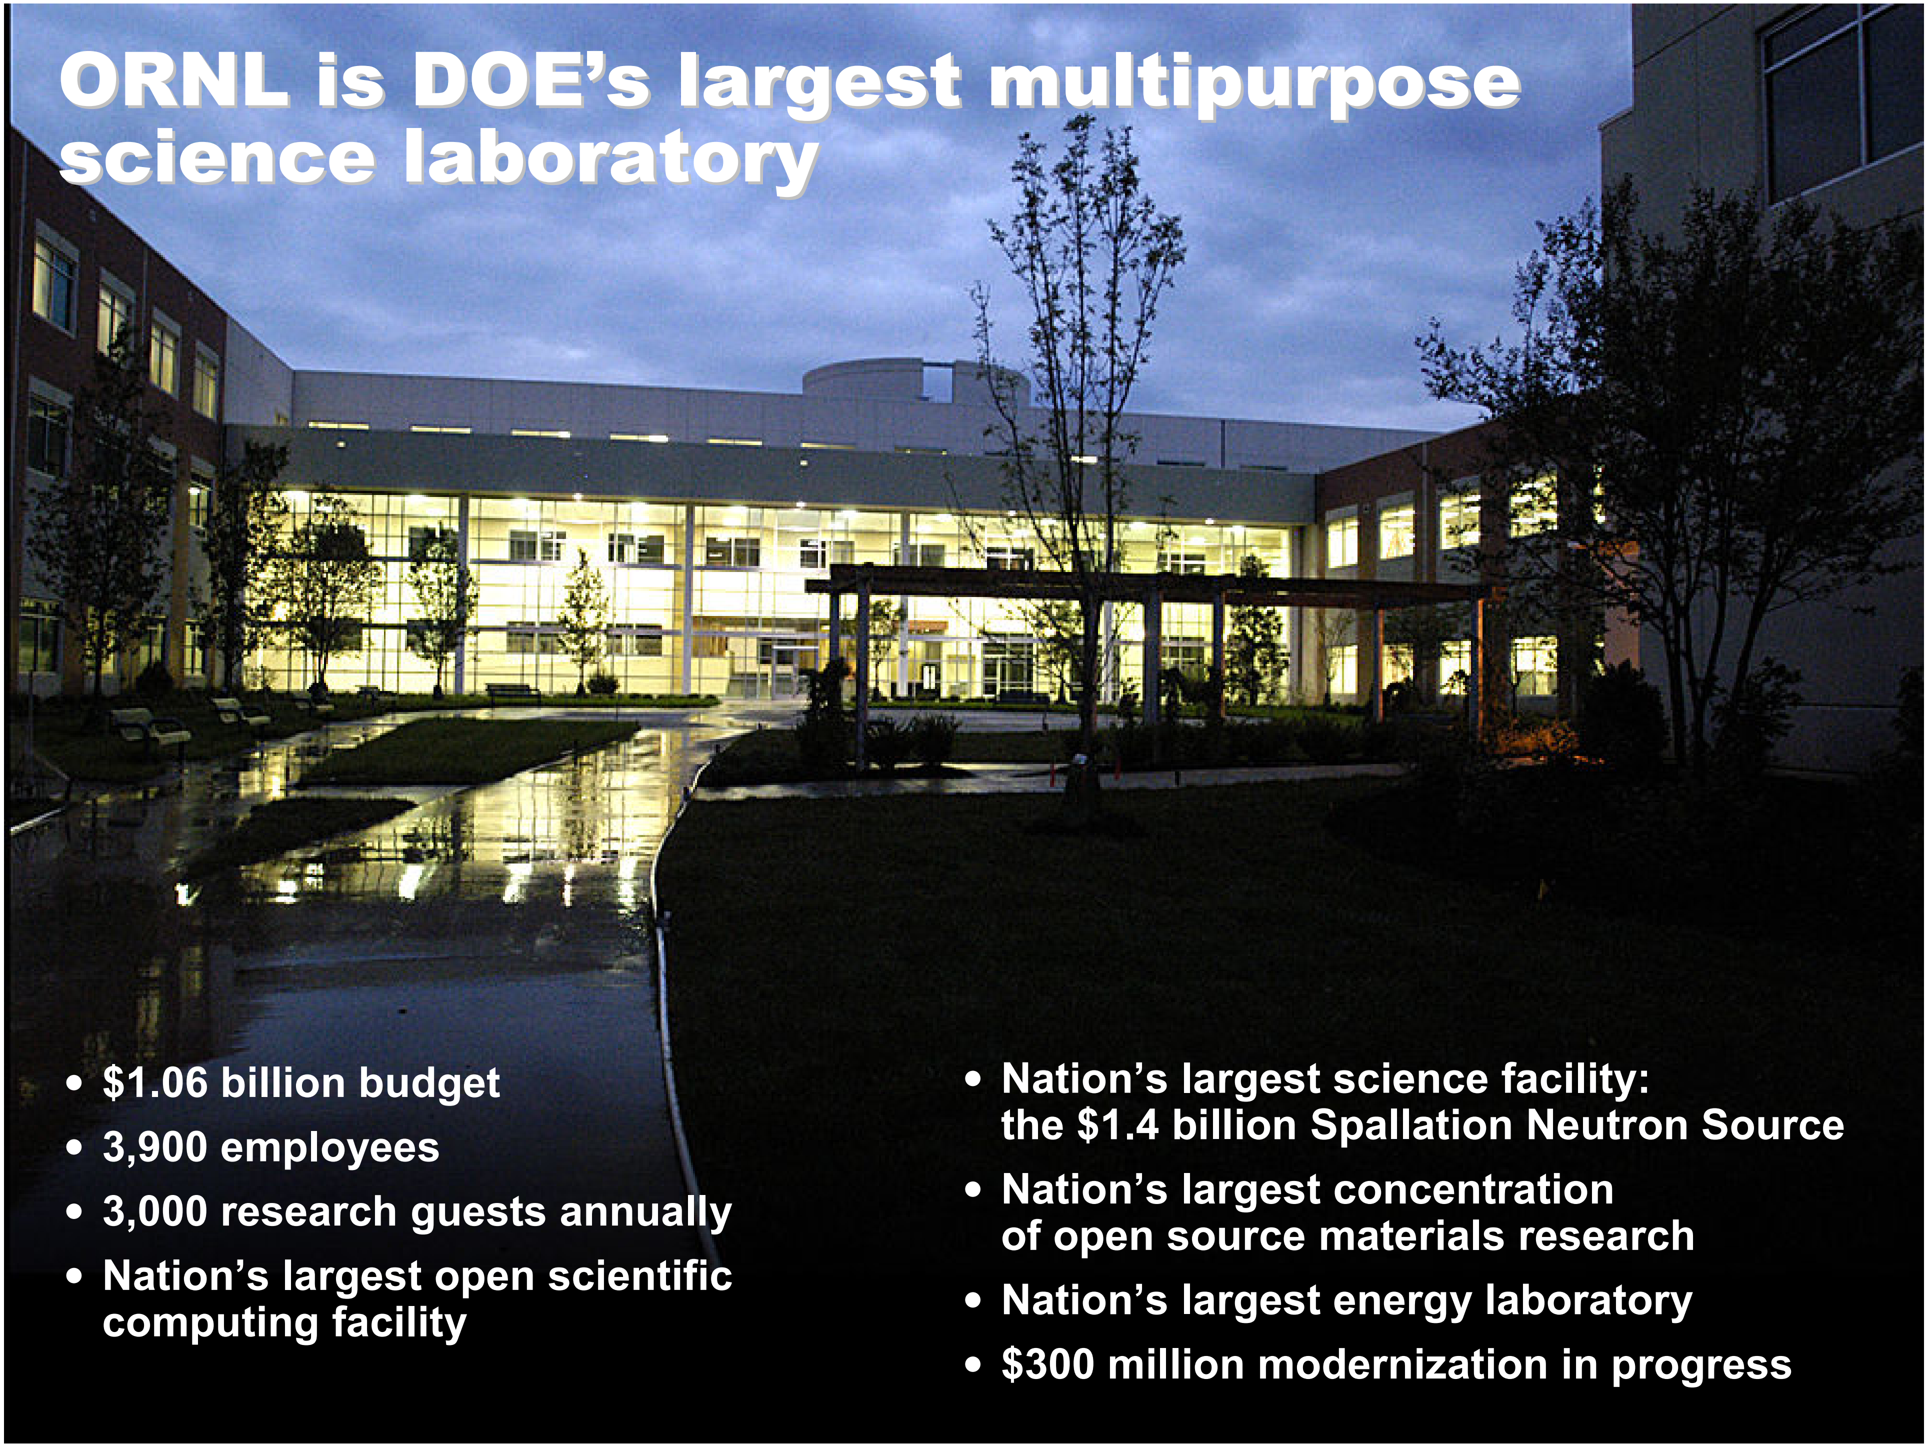

- $1.06 billion budget
- 3,900 employees
- 3,000 research guests annually
- Nation’s largest open scientific computing facility
- Nation’s largest science facility: the $1.4 billion Spallation Neutron Source
- Nation’s largest concentration of open source materials research
- Nation’s largest energy laboratory
- $300 million modernization in progress

National Center for Computational Sciences performs three inter-related activities for DOE

- Deliver National Leadership Computing Facility for science
  - **Focused on grand challenge science and engineering applications**
- Principal resource for SciDAC and (more recently) other SC programs
  - **Specialized services to the scientific community: biology, climate, nanoscale science, fusion**
- Evaluate new hardware for science
  - **Develop/evaluate emerging and unproven systems and experimental computers**

Intel Paragon: World's fastest computer

1995

IBM Power3: DOE-SC's first terascale system

2000

IBM Power4: 8th in the world (2001)

2001

Cray X1: Capability computer for science

2004

Cray XT3 and X1E: Leadership computers for science

2005

National Leadership Computing Facility

Leadership computing is a national priority

- “The goal of such systems [leadership systems] is to provide computational capability that is at least 100 times greater than what is currently available.”
- “High-end system deployments should be viewed not as an interagency competition but as a shared strategic need that requires coordinated agency responses.”
- “Traditional disciplinary boundaries within academia and Federal R&D agencies severely inhibit the development of effective research and education in computational science.”
- “The multidisciplinary teams required to address computational science challenges represent what will be the most common mode of 21st century science and engineering R&D.”

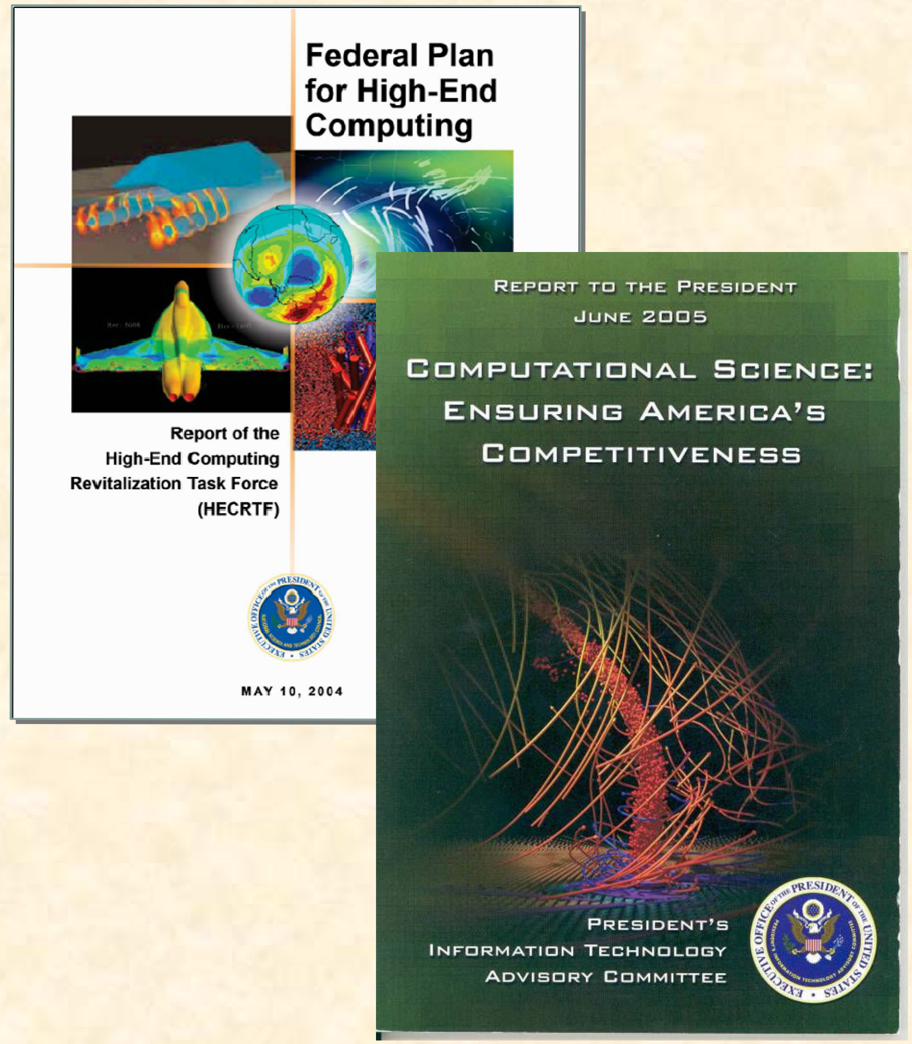

Leadership systems as enabler of science and technology

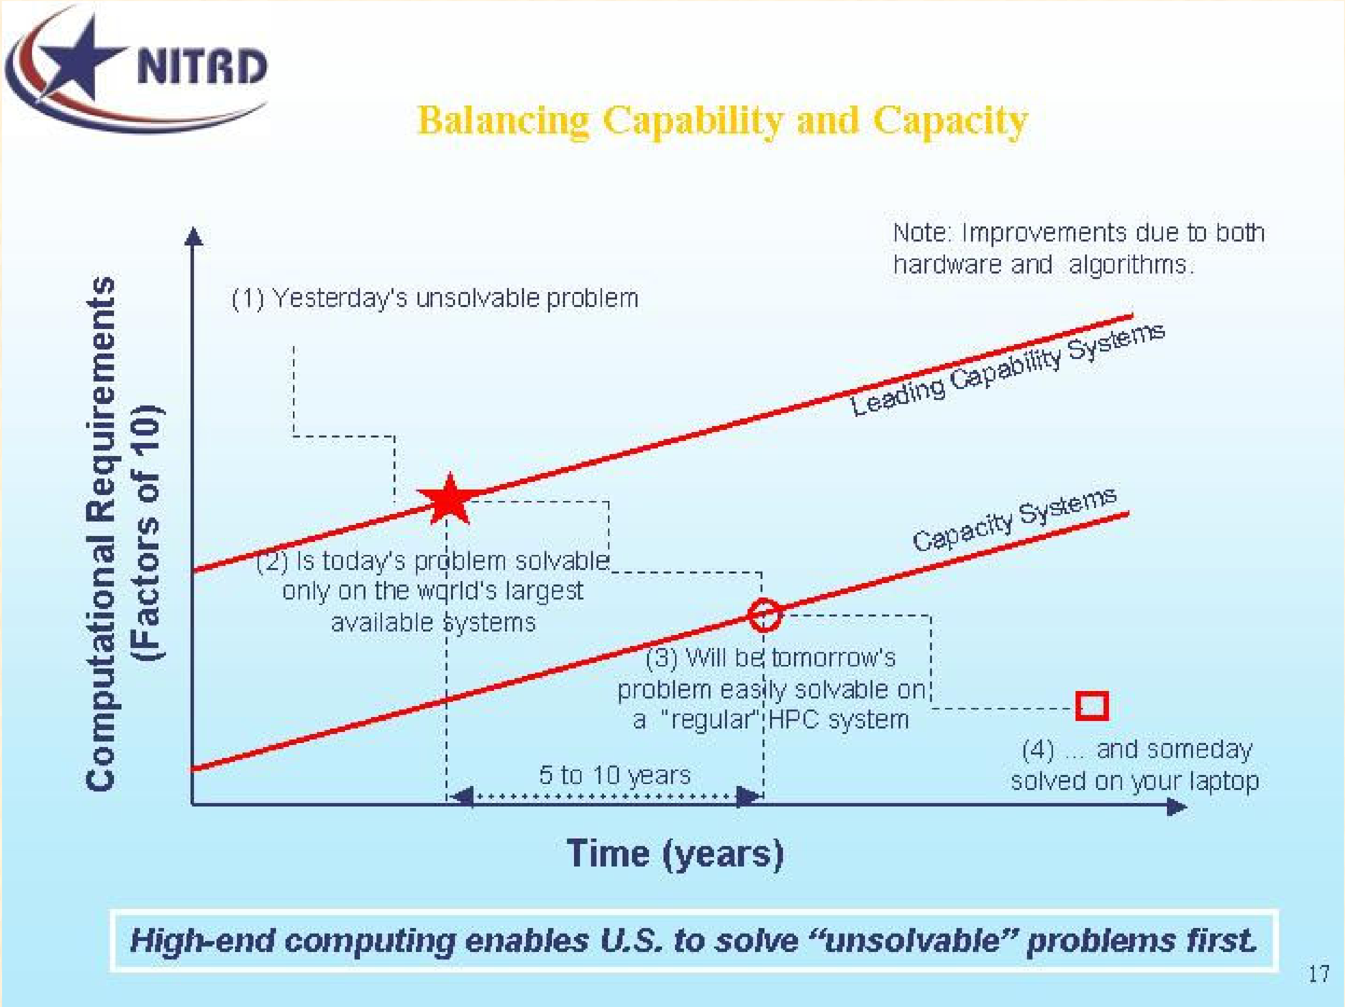

NLCF Mission

World leader in scientific computing

"User facility providing leadership-class computing capability to scientists and engineers nationwide independent of their institutional affiliation or source of funding"

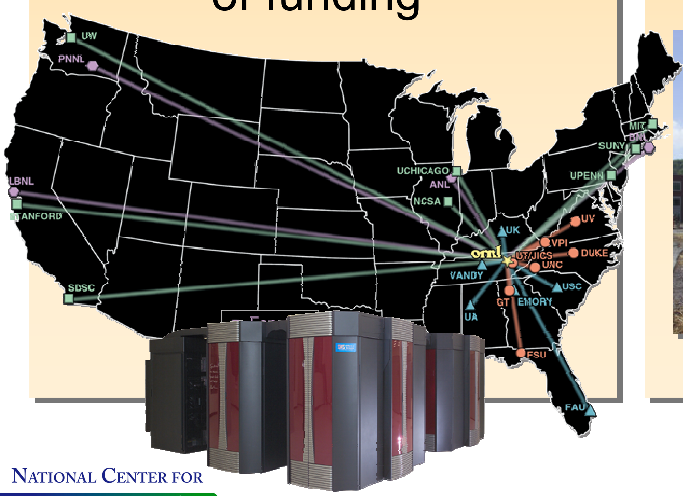

Intellectual center in computational science

Create an interdisciplinary environment where science and technology leaders converge to offer solutions to tomorrow's challenges

Transform scientific discovery through advanced computing

"Deliver major research breakthroughs, significant technological innovations, medical and health advances, enhanced economic competitiveness, and improved quality of life for the American people"

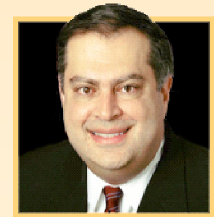

– Secretary Abraham

National Center for Computational Sciences

New world-class facility capable of housing leadership class computers

- $72M private sector investment in support of leadership computing
- Space and power:
  - **40,000 ft^2 computer center with 36-in. raised floor, 18 ft. deck-to-deck**
  - **8 MW of power (expandable) @ 5c/kWhr**
- High-ceiling area for visualization lab (Cave, Powerwall, Access Grid, etc.)
- Separate lab areas for computer science and network research

NATIONAL CENTER FOR COMPUTATIONAL SCIENCES

Where we are today

September 2005 Summary

Network Routers

Control Network
1 GigE
10 GigE
UltraScience

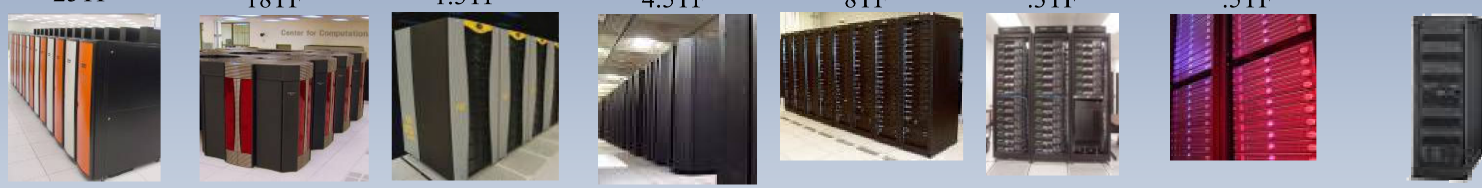

	Cray XT3 Jaguar	Cray X1E Phoenix	SGI Altix Ram	IBM SP4 Cheetah	SGI Linux OIC	IBM Linux NSTG	Visualization Cluster	IBM HPSS
Performance	25TF	18TF	1.5TF	4.5TF	8TF	.3TF	.5TF	
Processors / Memory	5294 2.4GHz 11TB Memory	1024 .5GHz 2 TB Memory	(256) 1.5GHz 2TB Memory	(864) 1.3GHz 1.1TB Memory	(1376) 3.4GHz 2.6TB Memory	(56) 3GHz 76GB Memory	(128) 2.2GHz 128GB Memory	Many Storage Devices Supported
Shared Disk	120TB	32TB	36TB	32TB	80 TB	4.5TB	9TB	5TB

Scientific Visualization Lab
- 27 projector PowerWall

Test Systems
- 96 processor Cray XT3
- 32 processor Cray X1E
- 16 Processor SGI Altix

Evaluation Platforms
- 144 processor Cray XD1 with FPGAs
- SRC Mapstation
- Clearspeed
- BlueGene (at ANL)

Backup Storage

5PB

Summary
- 8 Systems
- Supercomputers: 8,998 CPUs, 19TB Memory, 58 TFlops
- 318.5 TB
- 5 PB

Bandwidth explosion will impact how we do science

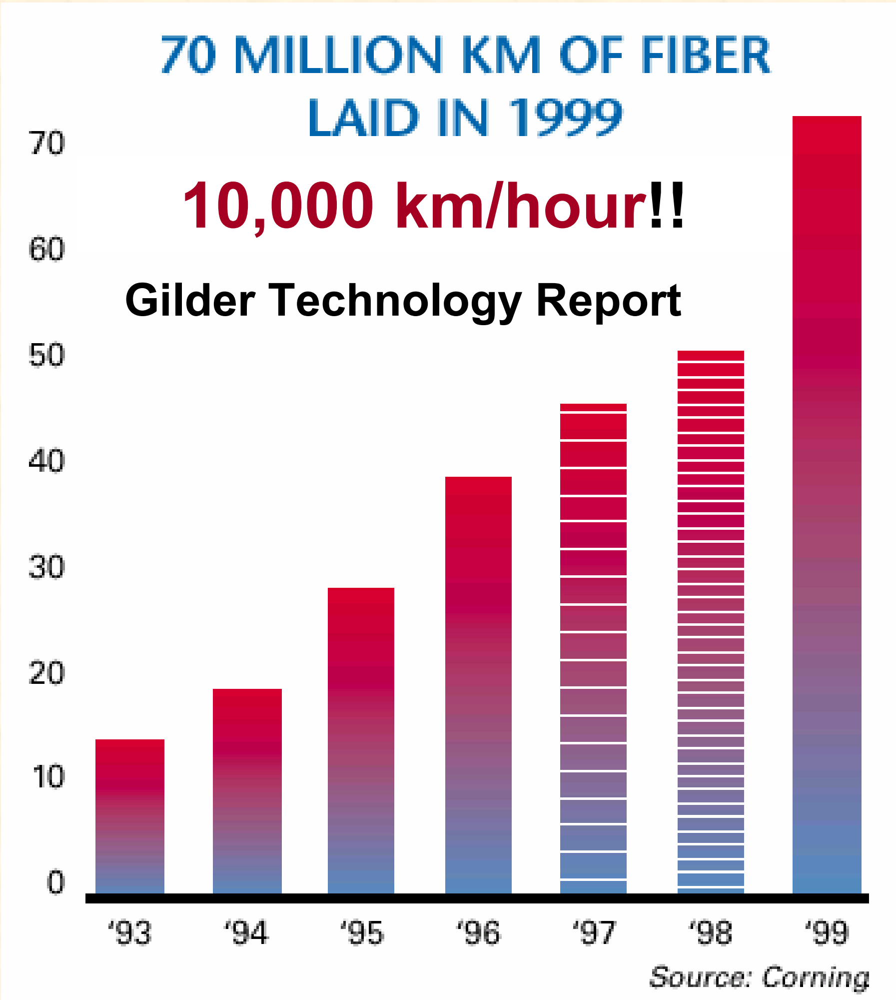

"…. a huge overinvestment in fiber-optic cable companies, which then laid massive amount of fiber-optic cable on land and under a oceans, which dramatically drove down the cost of making a phone call or transmitting data anywhere in the world."

--Thomas Friedman, *The World is Flat* (2005)

Already a resource provider on TERAGRID

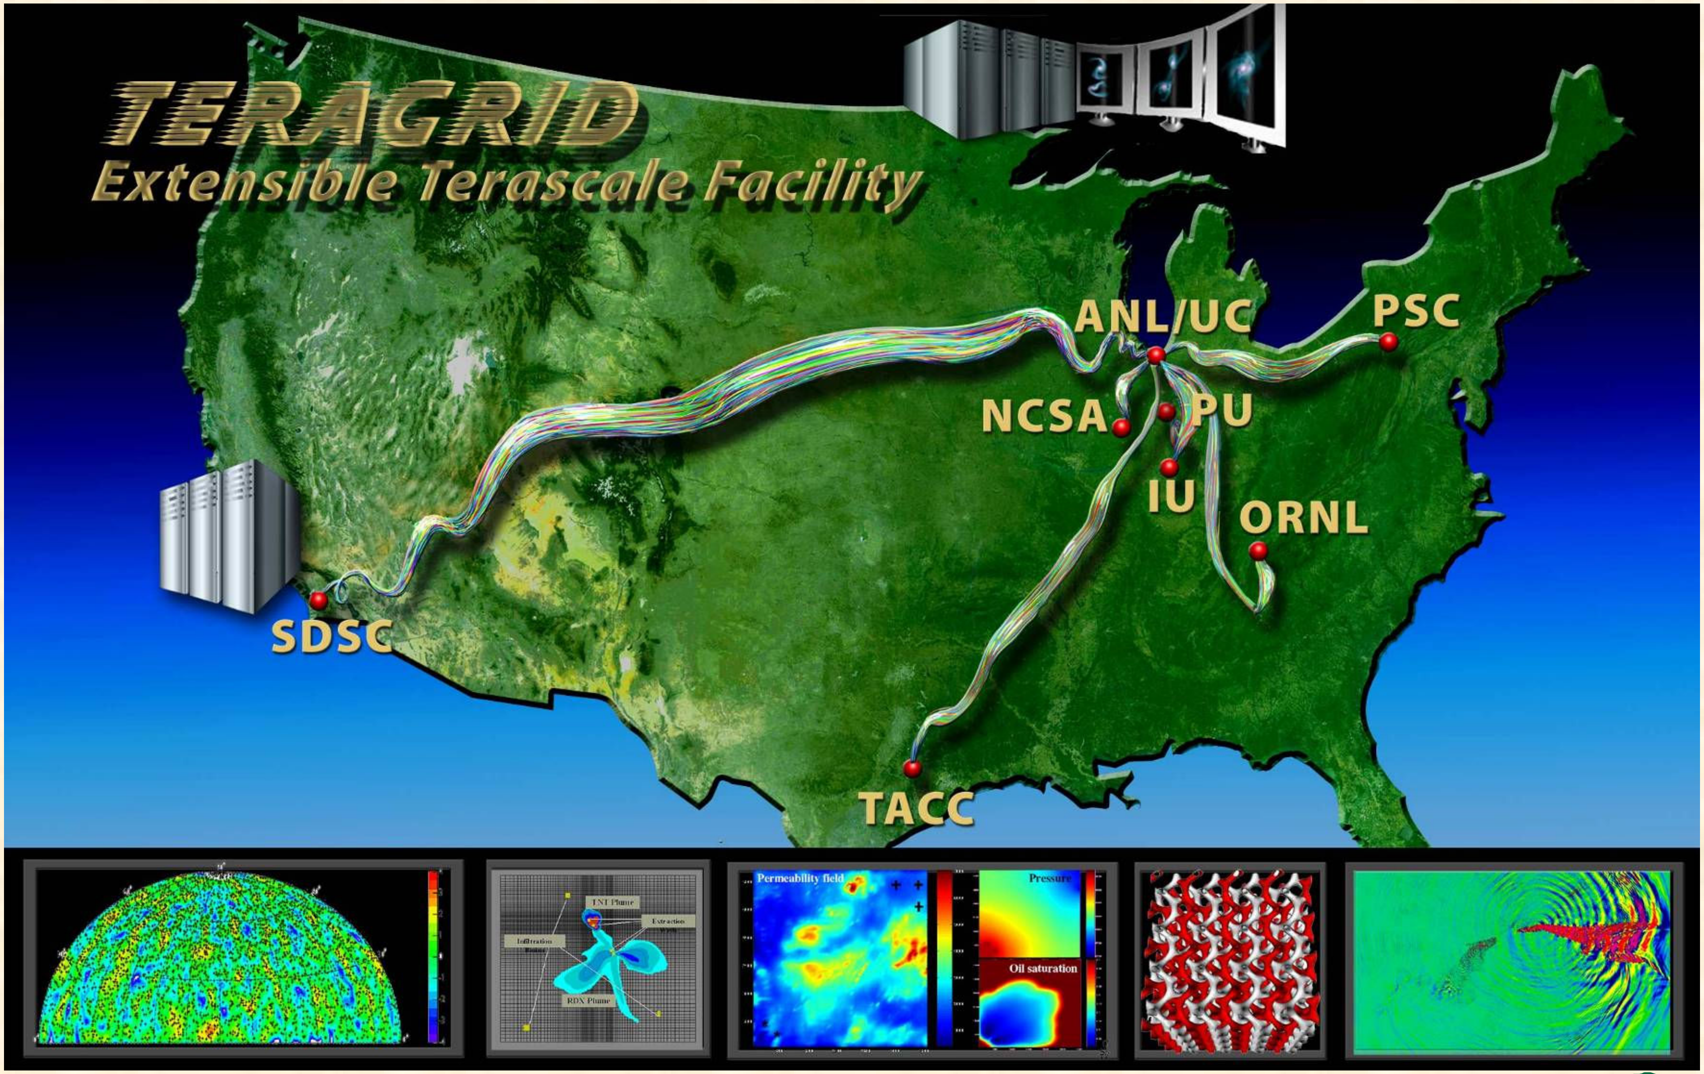

In 5 years ORNL went from backwaters to forefront in networking

Many countries are interconnecting optical research networks to form global super network

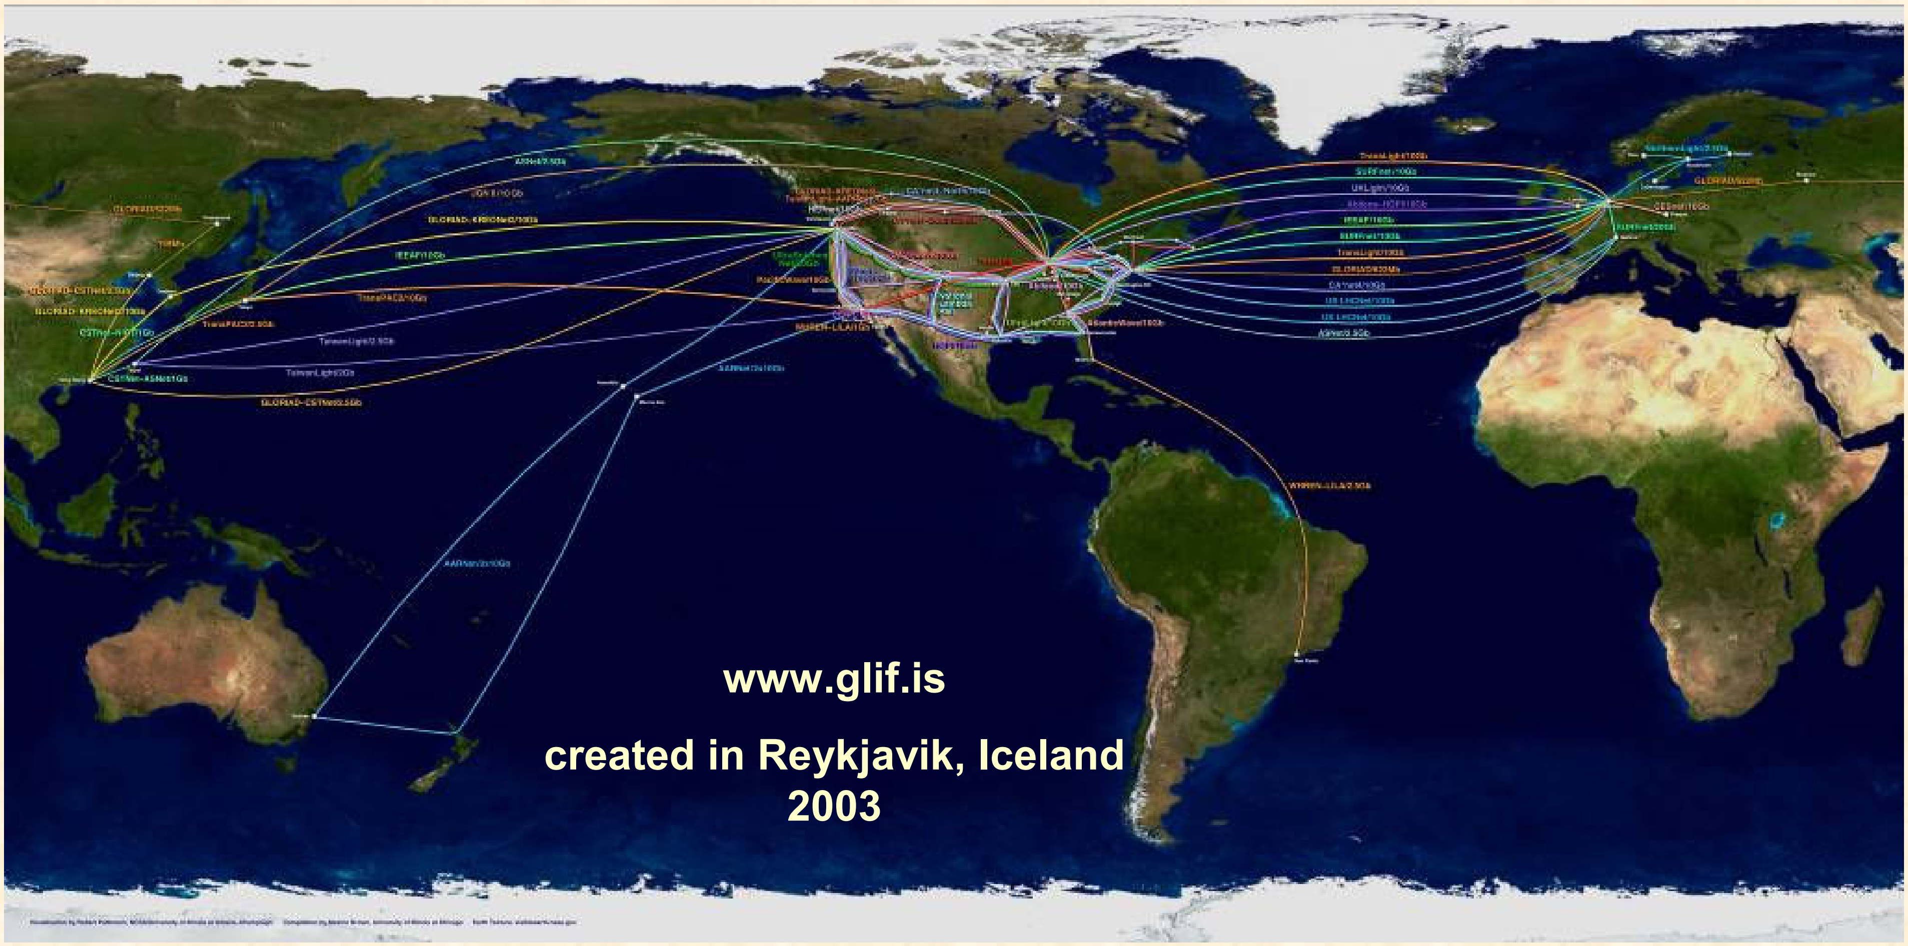

Integrate core capabilities to deliver computing for frontiers of science

Develop and evaluate next-generation architectures with industry

IBM

CRAY

sgi

Provide leadership-class computing resources for the Nation

Create math and CS methods to enable use of resources

SciDAC ISICs

Scientific Applications Partnerships

Modeling and simulation expertise

Transform scientific discovery through advanced computing

Computational End Stations

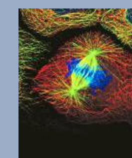

Biology

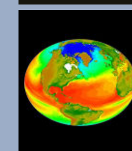

Climate

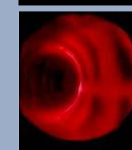

Fusion

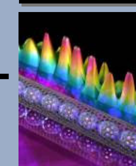

Materials

Industry/ other agencies

Outline

- Who are we: NLCF at ORNL
- **Science Motivators (first results)**
- Allocations and Support for Break-Through Science
- 2006-2010 Roadmap
- Wrapup

NCCS Cray X1E – Phoenix

- Largest Cray X1E in the world – 18.5TF
- 1024 processors – 400 MHz, 800 MHz vector units
- 2 TB globally addressable memory
- 32 TB of disk
- Most powerful processing node
  - 12.8 GF CPU, 2-5x commodity processors
- Highest bandwidth communication with main memory
  - 34.1 GB/sec

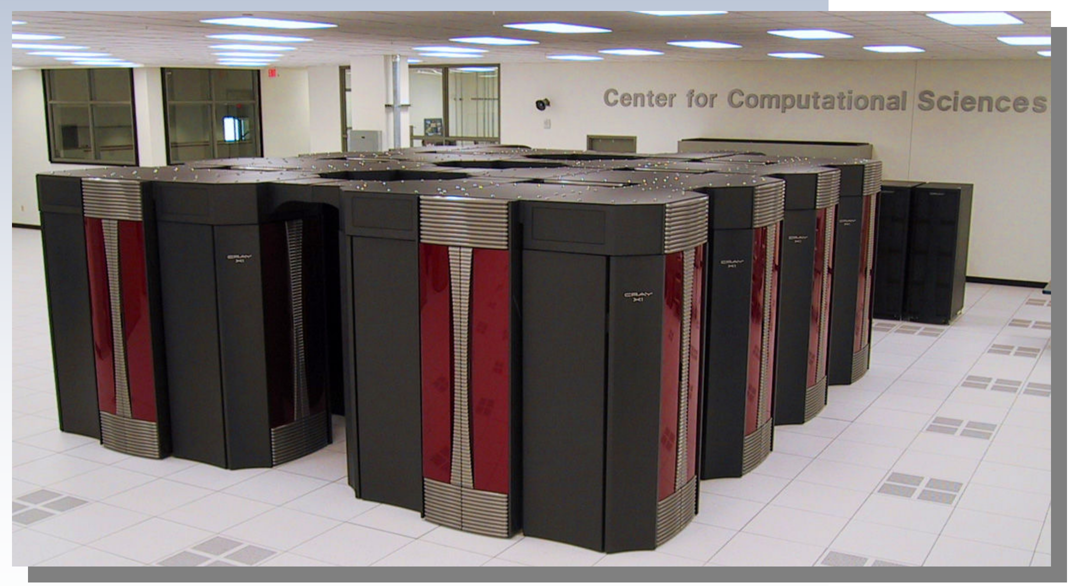

Highly scalable hardware and software
High sustained performance on real applications

FY05 X1E Allocations

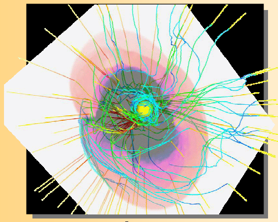

3D Studies of Stationary Accretion Shock Instabilities in Core Collapse Supernovae

A. Mezzacappa (Oak Ridge National Laboratory) and J. Blondin (North Carolina State University)
415,000 processor-hrs

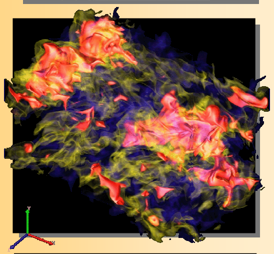

Turbulent Premix Combustion In Thin Reaction Zones

J.H. Chen (Sandia National Laboratories)
360,000 processor-hrs

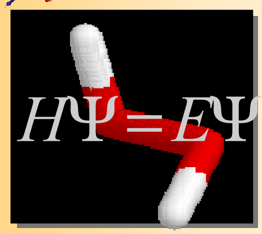

Full Configuration Interaction Benchmarks for Open Shell Systems

R. Harrison (Oak Ridge National Laboratory) and M. Gordon (Ames Laboratory)
220,000 processor-hrs

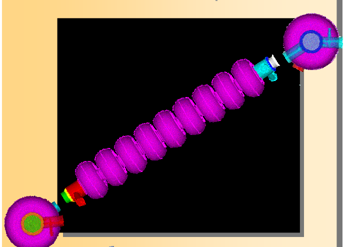

Computational Design of the Low-Loss Accelerating Cavity for the ILC

Kwok Ko (Stanford Linear Accelerator Center)
200,000 processor-hrs

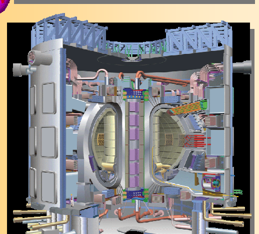

Advanced Simulations of Plasma Microturbulence

W. M. Tang (Princeton University, Plasma Physics Laboratory)
50,000 processor-hrs

Astrophysics Simulation:
Stationary accretion shock instability

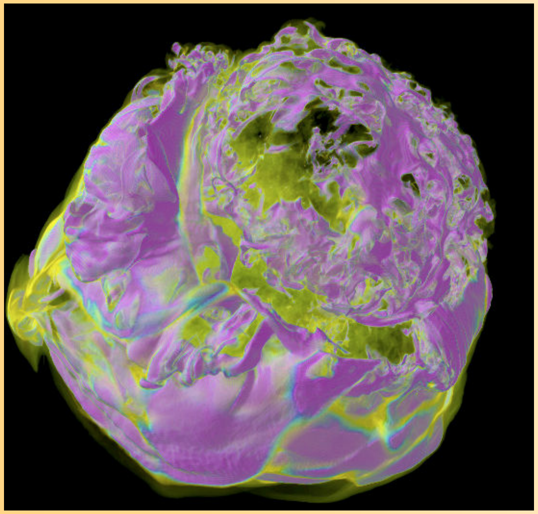

Volume renderings of TSI data showing entropy around a proto-neutron star.

Principal Investigators
John Blondin
North Carolina State University

Anthony Mezzacappa
Oak Ridge National Laboratory

- **The Problem**
  The core collapse of a massive star at the end of its life generates a shock wave that disrupts the star. TeraScale Supernova Initiative simulations have shown that this shock is dynamically unstable. This stationary accretion shock instability (SASI) will break the spherical symmetry of the parent star and perhaps aid in driving the supernova and causing it to explode.
- **The Research**
  Three-dimensional simulations on the Cray X1E are expected to yield results different than those found in previous two-dimensional simulations.
- **The Goal**
  The 3-D simulations will facilitate the exploration of the SASI and its implications for the supernova mechanism and spin up of the proto-neutron star.
- **Impact of Achievement**
  Further exploration of the SASI will enable researchers to better understand its role in the supernova mechanism and in producing key observables such as neutron star kicks and the polarization of supernova light.
- **Why NLCF**
  The throughput per run will be converted from one month to 1-2 days and will greatly facilitate the three-dimensional exploration of the SASI.

Combustion Simulation:
Non-premixed flame

- **The Problem**
  Detailed computer models are needed for design of cleaner, more efficient and environmentally friendly combustors.
- **The Research**
  This is the first 3-dimensional direct numerical simulation of a non-premixed flame with detailed chemistry.
- **The Goal**
  Simulations will provide essential data for understanding effects of turbulence and fuel-air mixing rate, flame extinction and re-ignition.
- **Impact of Achievement**
  Advancing basic understanding of turbulent combustion and developing predictive combustion models are essential to deliver reliable data for manufacturer design of combustors and to limit hardware testing costs.
- **Why NLCF**
  Code runs significantly faster on Cray X1E vector processors than on scalar processors; on NLCF computers, process runs in weeks rather than months or years.

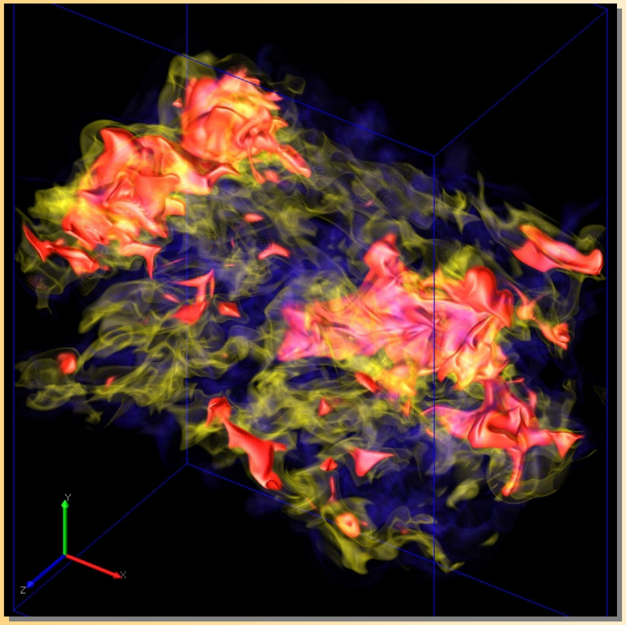

Hydroxyl radical in a turbulent jet flame.

Principal Investigator
Jackie Chen
Sandia National Laboratories

Small-Molecule Benchmark Calculations:
Calibrating quantum chemistry

- **The Problem**
  Understanding and controlling the structure, interactions and reactions of molecules are of critical importance to a wide range of phenomena, from the fate of contaminants in the environment to the treatment of genetic diseases.
- **The Research**
  This research will make use of a new parallel-vector algorithm for full-configuration interaction (FCI) calculations of molecular structures.
- **The Goal**
  Essentially exact benchmark calculations on small molecules will enable researchers to calibrate various approximate models that can then be used in calculations for much larger molecules.
- **Impact of Achievement**
  Large, fast computational power will enable advancement from approximate to exact models of molecules, especially for complex open-shell systems and excited states.
- **Why NLCF**
  The capabilities of the NLCF Cray X1E and the efficiency of the new algorithm will enable FCI calculations many times larger than were possible on other systems.

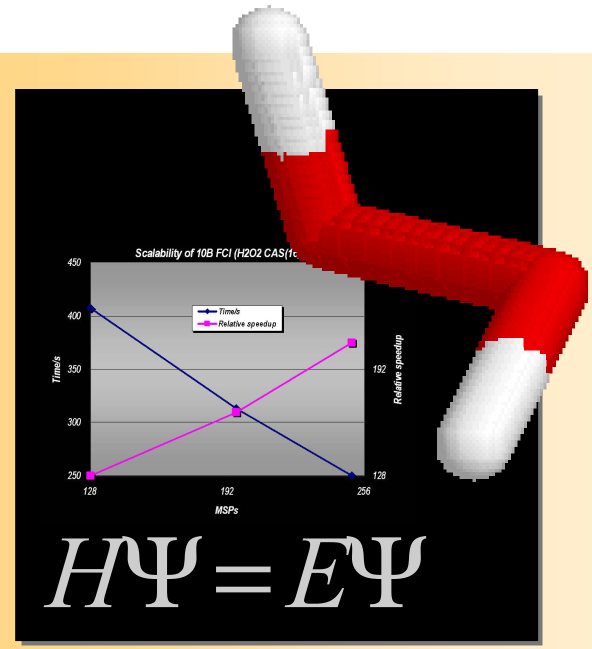

Characterizing matter at detailed atomic and molecular levels is enabled by large-scale calculations.

Principal Investigators
Robert Harrison
Oak Ridge National Laboratory

Mark Gordon
Ames Laboratory

Accelerator Design:
Low-loss accelerating cavity

- **The Problem**
  High-order-modes (HOMs) in the accelerating cavity of the International Linear Collider (ILC) can dilute the emittance of the beam and disrupt the transport of bunch trains down the accelerator. It is essential that the HOMs be sufficiently damped for stable operation of the ILC.
- **The Research**
  Researchers will use the parallel, three-dimensional electromagnetic eigensolver Omega3P – developed under SciDAC – to design a new low-loss (LL) accelerating cavity for the ILC that can meet the HOM damping criteria.
- **The Goal**
  The goal is to find the optimal geometry for the cavity and the HOM couplers to obtain the most effective damping of the HOMs.
- **Impact of Achievement**
  The ILC is the highest-priority future accelerator project in high-energy physics. Computer simulation will provide the input for determining the baseline cavity design for the ILC's main linac, which is the heart of the accelerator.
- **Why NLCF**
  Researchers are using the electromagnetic modeling capability of the Cray X1E to calculate the higher-order-mode damping needed to maintain beam stability.

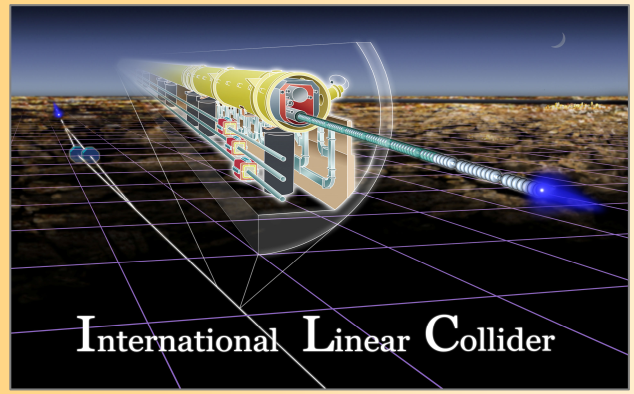

The ILC will provide a tool for scientists to address compelling questions about dark matter, dark energy, extra dimensions, and the fundamental nature of matter, energy, space and time.

Principal Investigator
Kwok Ko
Stanford Linear Accelerator Center

Fusion Simulation:
Particles in turbulent plasma

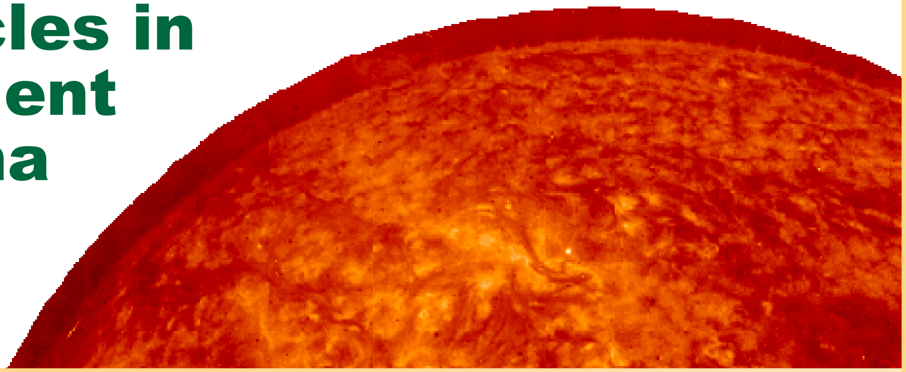

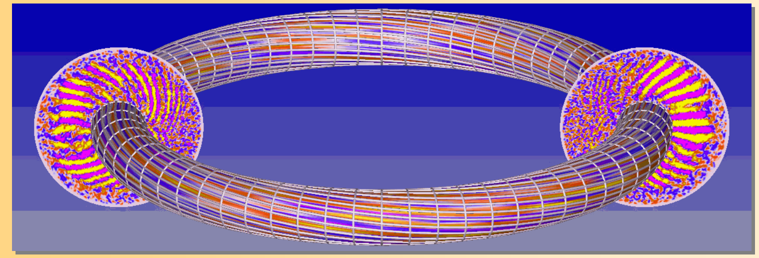

A twisted mesh structure is used in the GTC simulation.

Principal Investigators
William Tang and Stephane Ethier
Princeton Plasma Physics Laboratory

- **The Problem**
  Ultimately, fusion power plants will harness the same process that fuels the sun. Understanding the physics of plasma behavior is essential to designing reactors to harness clean, secure, sustainable fusion energy.

- **The Research**
  These simulations will determine how plasma turbulence develops. Controlling turbulence is essential because it causes plasma to lose the heat that drives fusion. Realistic simulations determine which reactor scenarios promote stable plasma flow.

- **The Goal**
  The NLCF simulations will be the highest-resolution Gyrokinetic Toroidal Code (GTC) models ever attempted of the flow of charged particles in fusion plasmas to show how turbulence evolves.

- **Impact of Achievement**
  High-resolution computer simulations are needed for preliminary data to set up experiments that make good use of limited and expensive reactor time. Engineers will use the resulting data to design equipment that creates scenarios favorable to efficient reactor operation.

- **Why NLCF**
  The fusion simulations involve four billion particles. The Cray X1E's vector processors can process these data 10 times faster than non-vector machines, achieving the high resolution needed within weeks rather than years.

NCCS Cray XT3 – Jaguar

Cray XT3 Scalable Architecture

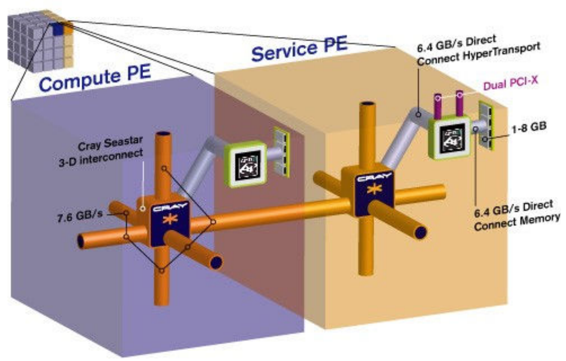

Accepted in 2005 and routinely running applications requiring 4,000 to 5,000 processors.

System Statistics	
Cabinets	56
Compute Processors	5,212 2.4 GHz Opteron
Lustre Object Storage Servers	58
10 Gigabit Ethernet nodes	2
System Services Nodes	8
Disk space	120 TB
Power	900 Kilowatts
Peak Performance	25.1 TeraFLOP/s

National Center for Computational Sciences

Cray XT3 Applications

Aero	GTC	Partisn	Benchmarks
Alegra	GYRO	POP	HALO
Amber/PMEMD	HYCOM	Presto	Hello World
AORSA	ITS	QCD-MILC	HPCC
ARPS	LAMMPS	Quake	HPL
AVUS (Cobalt-60)	Leonardo – Relativity Code	Quantum-ESPRESSO Suite	LINPACK
Calore	LM	S3D	NPB
CAM	LS-DYNA	Sage	OSU
CCSM	LSMS 1.6, 1.9, 2.0	Salinas	Pallas MPI
CHARM++	MAD9P	Siesta	PSTSWM
CHARMM	MILC	SPF	SMG2000
CPMD	moldyPSI	syr	sPPM
CTH	MPQC	TBLC	STREAM/triad
Dynamo	NAMD	Trilinos	Sweep3D
ECHAM5	NWChem	UMT2000	
FLASH	Overflow	VASP	
GAMESS	Paratec	WRF	
Gasoline – N-body astro.	Parmetis	ZEUS-MP	
Gromacs			9/1/05

55 applications + 13 Benchmarks

Largest ever AORSA Simulation
3,072 processors of NCCS Cray XT3

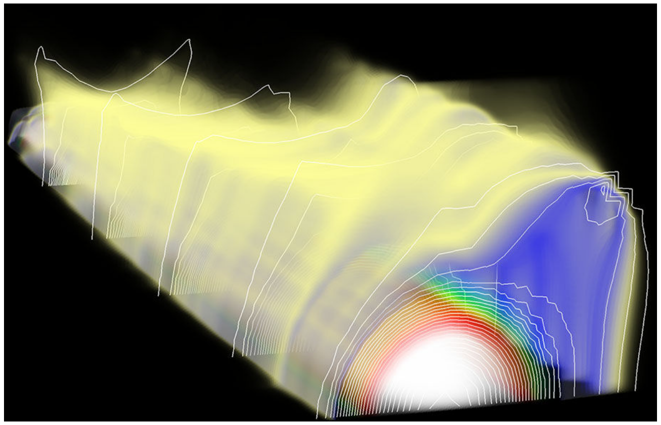

In August 2005, just weeks after the delivery of the final cabinets of the Cray XT3, researchers at the National Center for Computational Sciences ran the largest ever simulation of plasma behavior in a tokamak, the core of the multinational fusion reactor, ITER.

Velocity distribution function for ions heated by radio frequency (RF) waves in a tokamak plasma.

The code, AORSA, solves Maxwell's equations – describing behavior of electric and magnetic fields and interaction with matter – for hot plasma in tokamak geometry. The largest run by Oak Ridge National Laboratory researcher Fred Jaeger utilized 3,072 processors: roughly 60% of the entire Cray XT3.

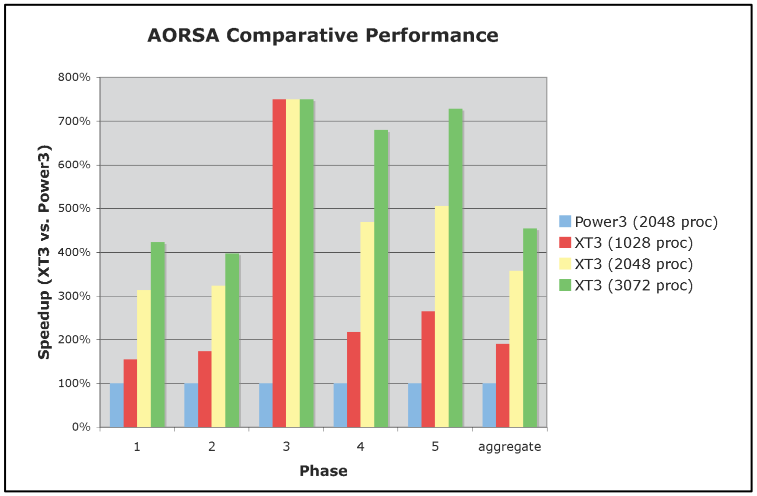

AORSA on the Cray XT3 "Jaguar" system compared with an IBM Power3. The columns represent execution phases of the code: Aggregate is the total wall time, with Jaguar showing more than a factor of 3 improvement over Seaborg.

Our Motivation:
Opportunities for Breakthrough Science

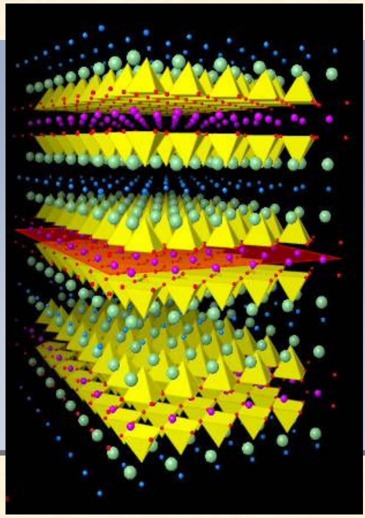

Two recent examples:

High-T_C superconducting materials:

– First solution of 2D Hubbard Model (T. Maier, *PRL*, accepted 10/2005)

Fusion plasma simulation:

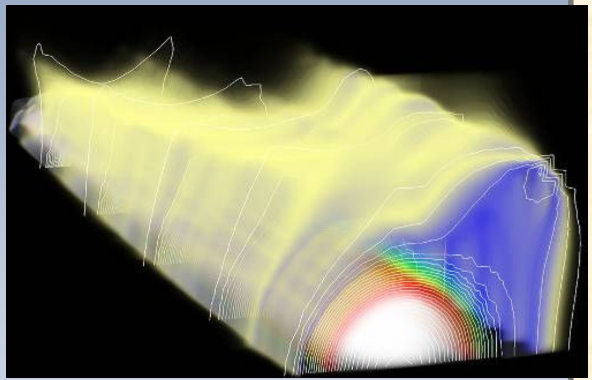

– Largest simulation of plasma behavior in a tokamak

(F. Jaeger, APS-DPP invited presentation, 10/2005)

FLASH Benchmarks

A weak scaling (problem size grows with the number of processors) plot for a standard FLASH test problem which compares the total time for solution on the Cray XT3, IBM Power3, and BG/L.