Opportunities in Federal NITRD High End Computing Interagency Working Group (HEC-IWG) Agencies

Robert Bohn
HEC Technical Coordinator
NCO/NITRD

November 18, 2008

Agenda

- **Introduction to the Networking and Information Technology Research and Development (NITRD) program**
- **Today's speakers**
  - Bill Thigpen - NASA
  - Dan Hitchcock - DOE/Office of Science
  - Almadena Chtchelkanova – NSF
  - Steve Meacham - NSF
  - Judith Terrill - NIST

NITRD National Coordination Office

- **Definition:**
  - The NITRD Program is the primary mechanism by which the U.S. Government coordinates its unclassified networking and information technology (IT) research and development (R&D) investments.
- **Legislation:**
  - The High-Performance Computing Act of 1991 (Public Law 102-194) as amended by the
  - Next Generation Internet Research Act of 1998 (P.L. 105-305) and the
  - America COMPETES Act of 2007 (P.L 110-69)
- **Objectives:**
  - Serve as the Federal focal point for interagency technical planning, budget planning, and coordination for the Federal NITRD Program
  - Serve as a source of timely, high-quality, technically accurate, in-depth information on accomplishments, new directions, and critical challenges for the NITRD Program
  - Support NITRD-related policy making in the White House Office of Science and Technology Policy (OSTP)

13 Member Agencies – IT R&D Budget ~ $3B

Agency for Health Research Quality

National Institutes of Health

Defense Advanced Research Projects Agency

DOE/National Nuclear Security Administration

National Institute of Standards and Technology

DOE/Office of Science

National Oceanic and Atmospheric Administration

Environmental Protection Agency

National Security Agency

National Archives and Records Administration

Department of Defense

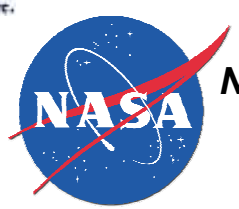

National Aeronautics and Space Administration

National Science Foundation

NITRD Program Organization

White House
Executive Office of the President

Office of Science and Technology Policy

President's Council of Advisors on Science and Technology (PCAST) serving as the **President's Information Technology Advisory Committee (PITAC)**

National Science and Technology Council

Committee on Technology

National Coordination Office (NCO) for Networking and Information Technology Research and Development

Subcommittee on Networking and Information Technology Research and Development (NITRD)

High End Computing (HEC) Interagency Working Group (HEC I&A, R&D)

Cyber Security and Information Assurance (CSIA) Interagency Working Group

Large Scale Networking (LSN) Coordinating Group

High Confidence Software and Systems (HCSS) Coordinating Group

Human Computer Interaction and Information Management (HCI&IM) Coordinating Group

Social, Economic, and Workforce Implications of IT and IT Workforce Development (SEW) Coordinating Group

Software Design and Productivity (SDP) Coordinating Group

Bill Thigpen
Deputy Manager for

NASA's High-End Computing Capability Program

HEC Partnership Opportunities with NASA

- **High-fidelity M&S is essential as NASA pursues some of greatest science and engineering challenges**

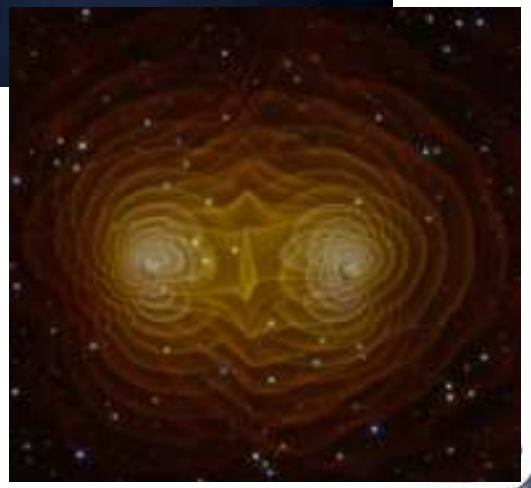

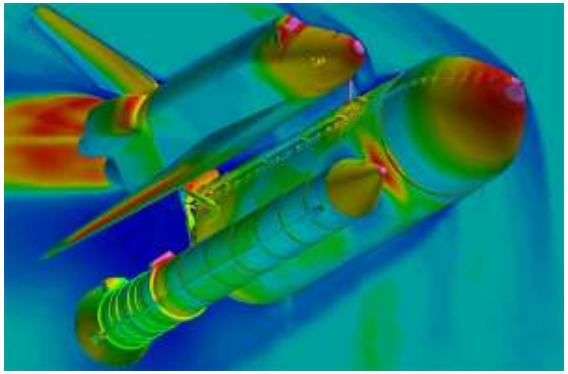

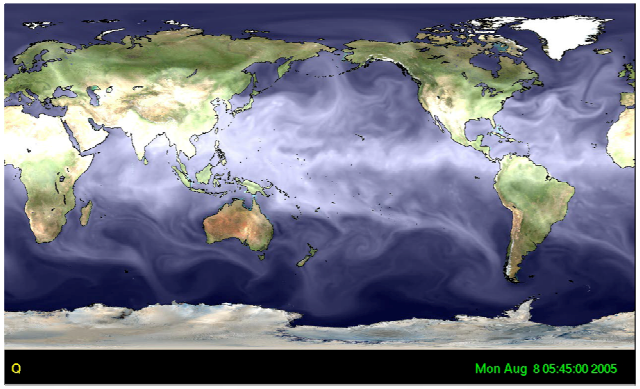

HEC Partnership Opportunities with NASA

- **High-fidelity M&S enabled by advanced HEC tech**

Pleiades: 51,200-core SGI supercomputer

hyperwall-2: 250 M-pixel data exploration environment

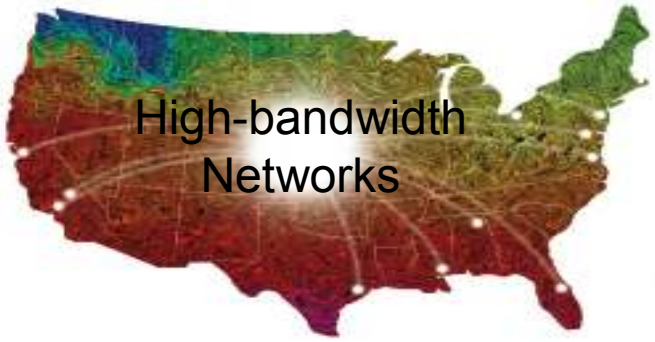

HEC Partnership Opportunities with NASA

- **Advanced HEC enabled by partnering with HEC community to enhance NASA HEC environment**
  - Industry on hardware and software innovations, new technology development, and major procurements
  - Other agencies to support architecture and technology advancements, access to complementary systems, platforms, and environments
  - Academia in novel collaborative research
  - Small businesses through the SBIR program

Dan Hitchcock
Senior Technical Advisor

NITRD

Office of Advanced Scientific Computing Research
DOE/Office of Science

U. S. Department of Energy – Office of Science

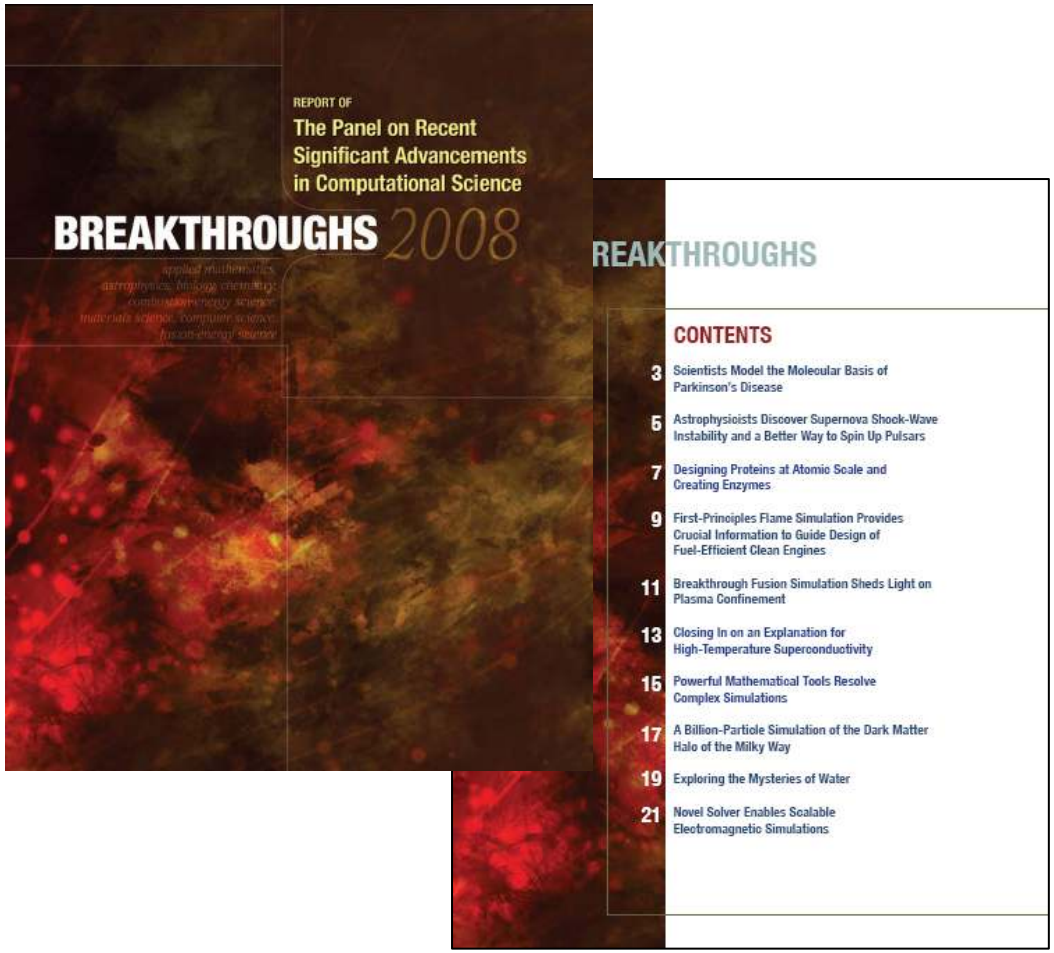

1.3 Petaflops on Scientific Application – 11/08

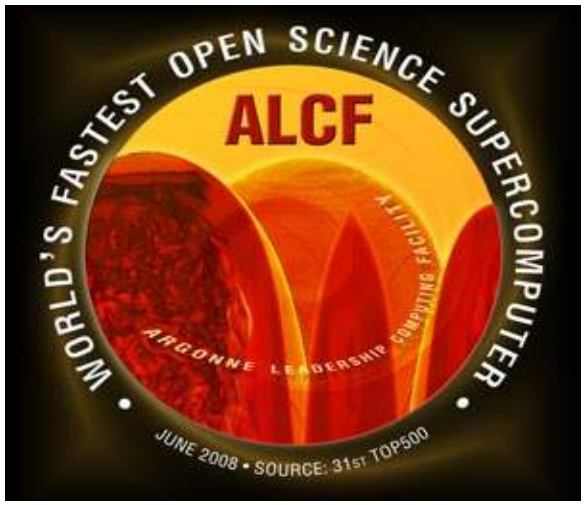

INCITE - Background

•Initiated at National Energy Research Scientific Computing Center (NERSC) at LBNL in 2004

•Provides Office of Science computing resources to a small number of computationally intensive research projects of large scale, that can make high-impact scientific advances through the use of a large allocation of computer time and data storage

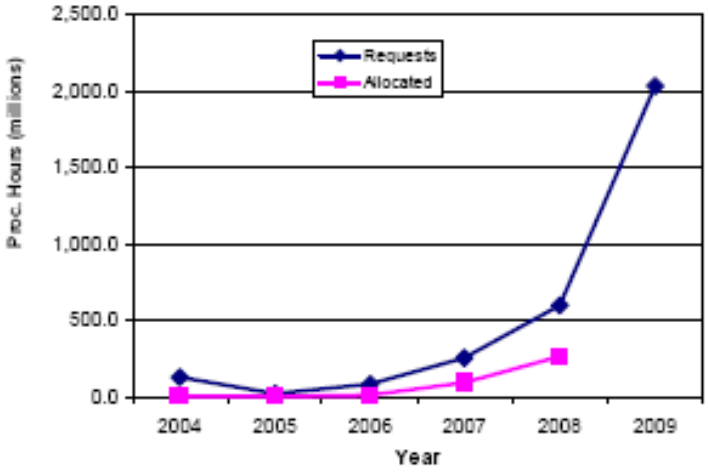

•Open to national and international researchers, including industry

•No requirement of DOE Office of Science funding

•Peer and computational reviews

INCITE - 2009

- 79 unique proposals received from scientific disciplines of accelerator physics, astrophysics, chemical sciences, climate research, computer science, engineering physics, environmental science, fusion energy, life sciences, materials science and nuclear physics
- 21 renewal proposals
- Over **2 Billion processor hours** requested for 2009 from new and renewal proposals

Over 600 Million processor hours available for 2009 awards, including renewals

Scientific Discovery Through Advanced Computing

Scientific Discovery

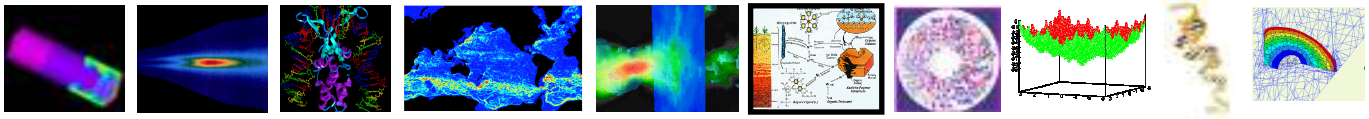

Applications

Computing/ Networking

SciDAC 2

- **Accelerator science and simulation**
- **Climate modeling and simulation**
- **Fusion science**
- **Petabyte high-energy/nuclear physics**
- **Nuclear physics**
- **Radiation transport**
- **Astrophysics**
- **Computational Biology**
- **High-energy physics**
- **Materials science and chemistry**
- **QCD**
- **Turbulence**
- **Groundwater reactive transport modeling and simulation**

- **Centers for Enabling Technology**

- **Scientific Applications Partnerships**
- **Institutes (University-lead)**

INCITE NERSC Allocation

Leadership Computing- ANL	556 TF IBM BG/P	
Leadership Computing- ORNL	263 TF Cray XT4	→ 1 PF Cray XT5
Production Computing- NERSC	104 TF Cray XT4	→ ~360 TF Cray XT4
ESnet	On path toward Dual rings 40Gbps/ 10 Gbps fault tolerant	

Applied Mathematics Research

- **FY 08 Budget:** Approximately 75% Labs, 25% Universities
- **Chart Allocations by Area**
  - **Based on $23.6M (FY08)**
  - **Does not include**
    - **Computational Science Graduate Fellowship Program ($5M)**
    - **New Multiscale Mathematics and Optimization awards**
    - **Potential new FY 09 initiatives:**
      - Mathematics for Analysis of Petascale Data
      - Joint Applied Mathematics-Computer Science Institutes
      - High Risk / High Payoff Technologies

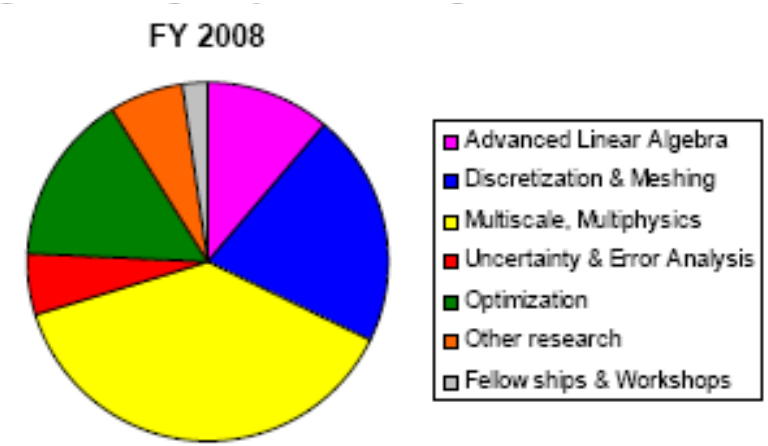

Computer Science Challenges

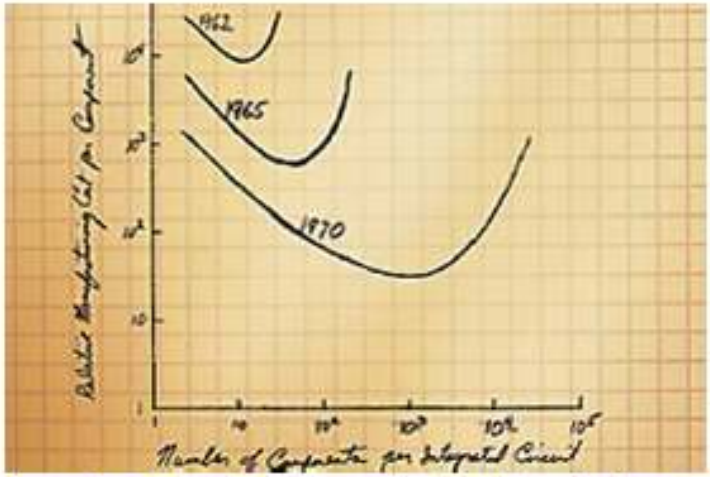

Moore's original graph predicting Moore's Law in 1965. Chip capacity will double every 2 yrs.

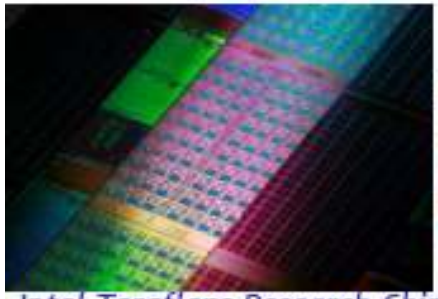

Intel Teraflops Research Chip

IBM Stacked Chip

- Unpredictable evolution of hardware
- Multilevel and heterogeneous parallelism; memory hierarchies
- Programming models must work at scale (numbers of cores, lines of code, numbers of components)
- Managing data, simulation, experimental and observed
- Communications:
  synchronous → asynchronous
- Reliability

It's not just extreme scale, it's also extreme complexity

- **Staffing Changes**
  - Susan Turnbull - detailee from GSA, Team Lead
- **Program elements:**
  - **Network research** - core network research
  - **Middleware research** – Grid technologies

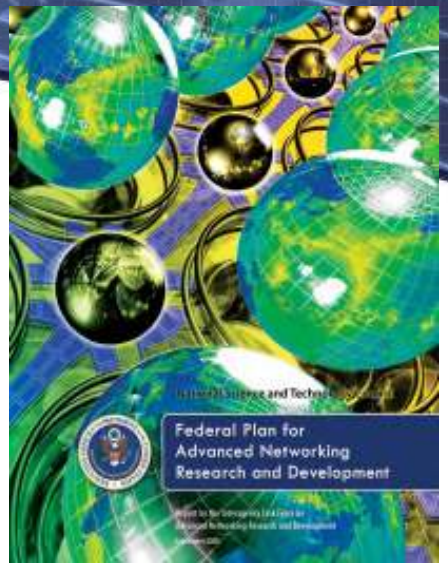

- **The next-generation program**
  - Supports R&D activities to develop advanced networks to enable distributed high-end science
  - Coordinates with ESnet to develop and deploy networks that enable scientists to push the limits of today's networks
- Next-generation network technologies have enabled the efficient and rapid distribution of massive data generated by the LHC experiment and climate modeling
- Major activities in FY09: Network research program announcement

High-End Computing University Research Activity HECURA National Science Foundation

Almadena Chtchelkanova
Program Director
Directorate for Computer and Information Science and Engineering (NSF/CISE)
Computing and Communication Foundations

High-End Computing University Research Activity (HECURA)

- FY 2006 Budget $14.5M
- NSF/DARPA/DOE/EPSCoR activity focused on
  - Input/Output capabilities
  - File Systems
  - Storage Systems

- 19 projects awarded
  - I/O, file and storage systems design for efficient, high throughput data storage, retrieval and management in the HEC environment.
  - hardware and software tools for design, simulation, benchmarking, performance measurement and tuning of file and storage systems.

High-End Computing University Research Activity (HECURA)

- FY 2008 Budget $10.5M
- NSF/EPSCoR activity focused on
  HEC Programming Models, Languages and Compilers
  18 projects awarded

HECURA BOF Kick-off meeting
Tuesday 5:30 – 7 pm 12A/12B

High-End Computing University Research Activity (HECURA)

- FY 2009 Budget $10M (proposed, subject to availability of funds)
- NSF activity focused on HEC
  - Input/Output capabilities
  - File Systems
  - Storage Systems

  Anticipated publication date – December 2008

  Proposals due – April 2009

More Funding Opportunities at NSF

Expeditions in Computing

http://www.nsf.gov/pubs/2008/nsf08568/nsf08568.htm

- **Created to inspire bold, transformational research that explores new scientific frontiers that promise disruptive innovations in computing**

- **Program Goals:**
  - Catalyze far-reaching research in the computing and information fields motivated by hard, emerging problems and/or compelling applications
  - Inspire current & future generations to pursue CISE careers
  - Stimulate significant research and education outcomes that benefit society through effective knowledge transfer
- **Scope: Research that cuts vertically or horizontally across CISE**
- **Awards: 3 anticipated, w/ budgets of $2,000,000/yr for five years**
- **Preliminary Proposal Submission Deadline September 10, 2009**
- **FY08 4 awards include**

Computational Sustainability: Computational Methods for a Sustainable Environment, Economy, and Society

More Funding Opportunities at NSF

Cyber-Enabled Discovery and Innovation (CDI)

http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503163

Computational thinking for science and engineering

-computational concepts, methods, models, algorithms, and tools

Type I Preliminary Proposal Submission Deadline December 08, 2008

Type II Preliminary Proposal Submission Deadline December 09, 2008

FY09 Budget $26M

- **CDI is a five-year initiative.**
- **All NSF directorates & programmatic offices involved in CDI.**
- **CDI goal is to make revolutionary, paradigm-changing research outcomes possible.**
- **Discovery and innovation are enabled by computational thinking.**
- **Requires bold, multidisciplinary activities.**

NSF's High-End Computing Activities

Steve Meacham
Sr. Science and Technology Advisor
Office of Cyberinfrastructure
National Science Foundation

SC08
November 2008

Strategic Plan (FY 2006 – 2010)

- Ch. 2: High Performance Computing

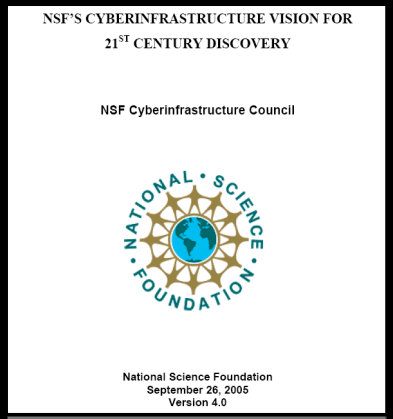

5yr Goal:

- To enable petascale science & engineering thro' deployment & support of a world-class HPC environment comprising the most capable combination of HPC assets available to the academic community

Components:

- Acquisition, deployment, operation of science-driven HEC systems
- Development & maintenance of supporting software
  - New design tools, performance modeling tools, system software & fundamental algorithms
- Development and maintenance of portable, scalable applications software

HEC Program Elements

- **Acquisitions**
  - Track 1 - Petascale
  - Track 2 - Mid-range supercomputers
- **Operations**
- **HEC System Software**
  - Compilers, fault-tolerant OS, fault-survivability tools, system status monitoring, file-systems, PSEs, …
- **HEC Petascale Application Development**
  - Scalable math libraries, scalable algorithms, data exploration tools, performance profiling and prediction, large application development
- **Coordinated with other agencies**

Recent High-end HPC Investments

Petascale Software Development (PetaApps, SDCI and STCI) 2006...

2007- 2011- 2016 Track 1 Petascale Acquisition (UIUC/NCSA)

TeraGrid Operations
TACC, NCSA, SDSC, PSC, ORNL, ANL, Indiana, Purdue, NCAR, LSU

TeraGrid Phase III (XD)

Track 2

Track 2A: 2006-2011 TACC

Track 2B: 2007-2012 UTK

Track 2C: 2008-2013 PSC

Track 2D: 2009-2014 Multiple smaller resources

TACC - Texas Advanced Computing Center, Austin.
UTK - Univ. of Tennessee, Knoxville, Joint Institute of Computational Science.
NCSA - National Center for Supercomputer Applications, Univ. of Illinois, Urbana-Champaign.
SDSC - San Diego Supercomputer Center, Univ. of California, San Diego.
PSC - Pittsburgh Supercomputing Center

Impact

Greatly expanding capacity of the TeraGrid for digital exploration with reduced oversubscription and queue wait times.

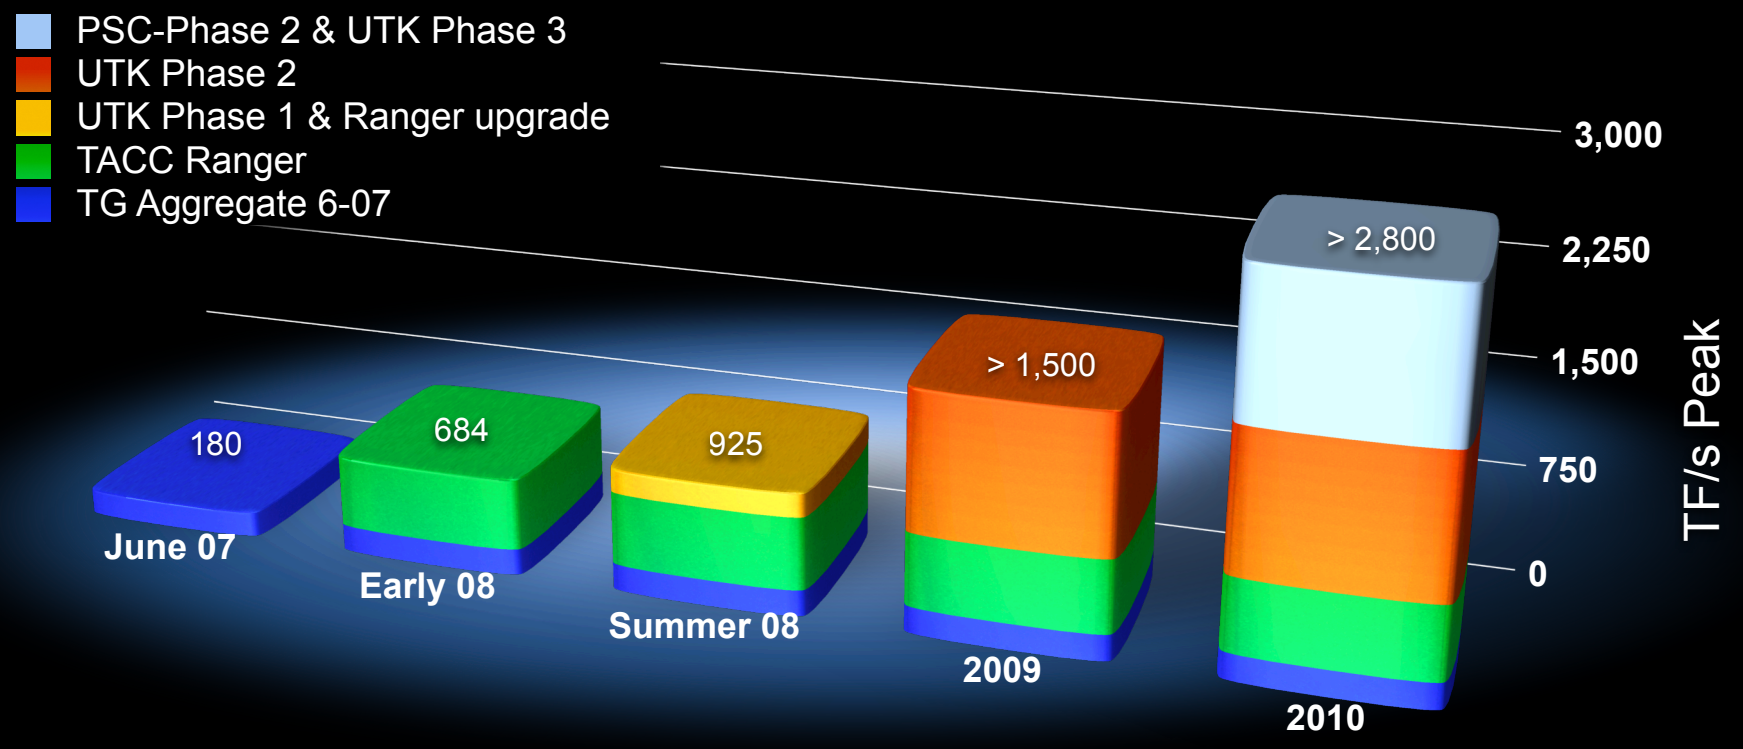

National Science Foundation
Where Discoveries Begin

Office of Cyberinfrastructure

Other Trends

- Growing power dissipation
- Complexity of system
- Complexity of programming
- Increase in uptake
- Declining investment

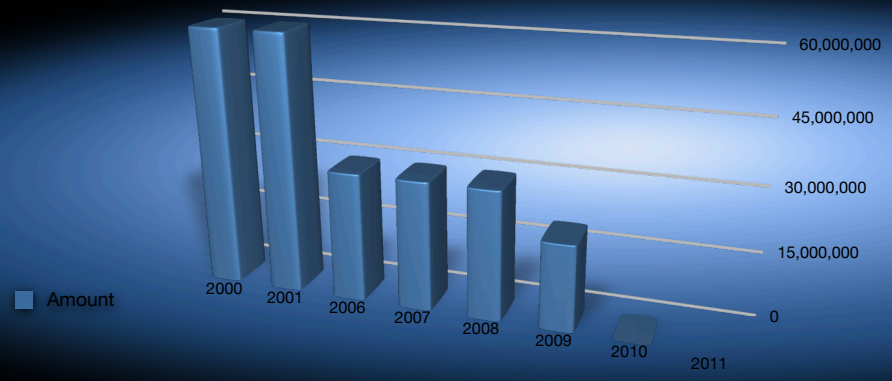

Blue Waters: the next-generation supercomputer

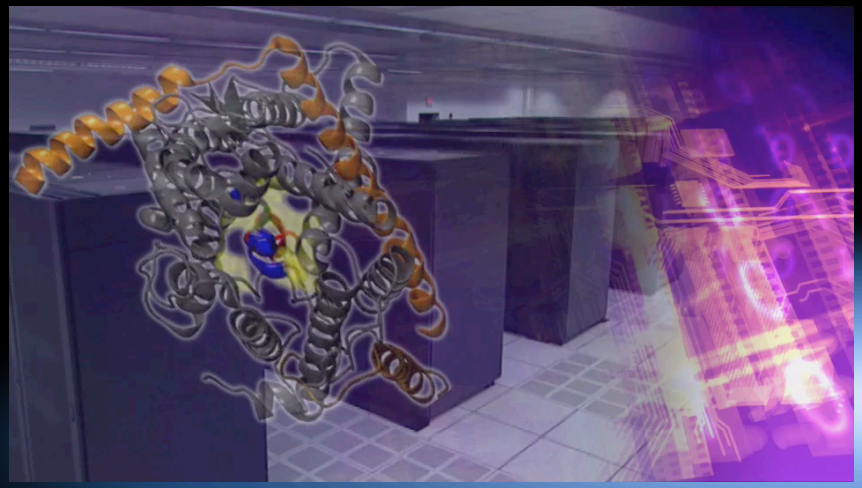

- climate science, particle physics, materials science, stellar evolution
- fluid dynamics, astrophysics and cosmology, condensed matter physics
- cell biology, nano-engineering, chemistry, plasma physics
- the influence of social networks on the spread of contagion

National Science Foundation
Where Discoveries Begin

Office of Cyberinfrastructure

HPC is an increasingly important tool for understanding:

Life

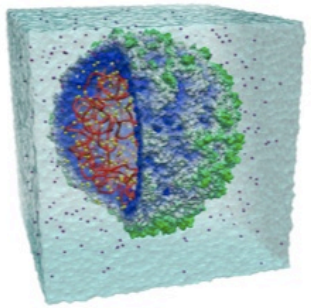

Satellite tobacco mosaic virus, P. Freddolino et al.

Aldehyde dehydrogenase, T. Wymore and S. Brown

Matter

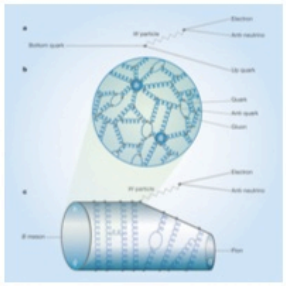

I. Shipsey

The Environment

K. Droegemeier et al.

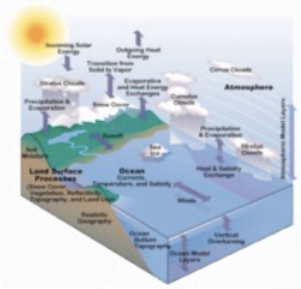

Community Climate System Model

Society

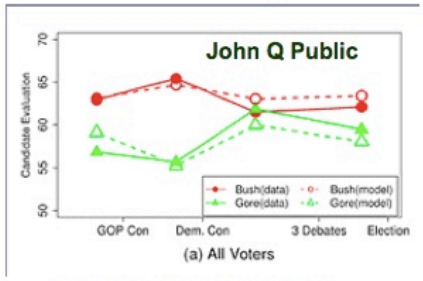

S.-Y. Kim, M. Lodge, C. Taber.

D. E. Atkins

Some Science We'll Be Able To Do

Pandemic Network Simulation and Analysis: Models for simulating the spread of infectious diseases using graph-theoretic methods. The influence of social networks. In silico testing of mitigation strategies.

Computational Proteomics: Modeling proteins with ~1000 amino acids to examine protein-protein interactions and trans-membrane signaling. Inverse methods for macromolecular design.

Environmental Decomposition of Complex Molecules: Develop highly scalable, locally correlated, numerical quantum chemistry methods for molecules containing hundreds of heavy atoms. Study decomposition pathways of complex molecules in the environment.

Ocean Modeling: Study ice-ocean interactions using whole-globe coupled ice-ocean models.

Cosmology: Multi-scale cosmological modeling in support of research efforts in precision cosmology, including galaxy surveys, weak lensing, and the analysis of the cosmic microwave background.

Computational Fluid Dynamics: Study the evolution of transient, localized structures in inhomogeneous turbulence using daptive mesh refinement (AMR).

Mining the Internet Archive: Studying the propagation of ideas and interactions.

Analysis of Socio-Economic data: Extraction of patterns and pattern evolution in multi-year socio-economic data collections.

Track 2d

- Track 2 = TeraGrid hardware funding line
- Track 2d = up to $20M in year 1 & constrained O&M
- Looking for innovation
- Up to four awards - two “production” and two “experimental” categories
  - A data-intensive, high-performance computing system
  - An experimental, high-performance computing system of innovative design
  - An experimental, high-performance grid test-bed for grid research (CS & applied)
  - A pool of loosely coupled grid-computing resources
- All “systems” to be made available to national community of users via the TeraGrid

Some examples of existing programs

NSF's investments in HEC R&D are spread over several offices and divisions

- HECURA
- Science and Engineering Beyond Moore's Law - proposed new effort, led by MPS

  Focuses on new paradigms for HW architecture, algorithms, & application software

  E.G. basic research on new types of logic devices: nanophotonics, spintronics, molecular computing, quantum computing
- STCI & SDCI
- Multicore Chip Design and Architecture (NSF and SRC)

  Innovative research on design, fabrication, architecture and programmability of homogeneous and heterogeneous multicore systems. - Includes hardware and software aspects
- PetaApps

  Funding for development of future simulation and analysis tools that can use petascale computing to advance the frontiers of scientific and engineering research.
- Application development in domain science and engineering programs. e.g. climate, nanoscale engineering, biosciences, chemistry, SBE
- National Center for Atmospheric Research

The next five years 2011 - 2016

- Over next year - develop a new five-year plan
- Looking for input on opportunities and priorities
  - From academia and industry
  - From consumers of leading-edge HPC and from researchers and developers of leading-edge HPC
- Anticipate greater emphasis on software development
- Over the current decade, NSF investments in HPC acquisition have been declining; future funding levels are uncertain
- NSF seeks to maintain a balanced cyberinfrastructure portfolio. Other CI areas include: data curation and data management, networking, middleware development, grid/cloud computing
- Exascale and beyond in computing and data are hard challenges
  - hope to expand partnerships with other agencies and foster international collaborations
- No investment similar to Track 2 is planned for 2010

Questions

- What are some of the research areas - hardware & software - which academia & industry would be interested in tackling together?
- What are the science grand challenges that would require the combination of international HPC resources to tackle?
- How does the cloud paradigm fit into the provisioning of high-performance computing resources?

Questions or Comments?

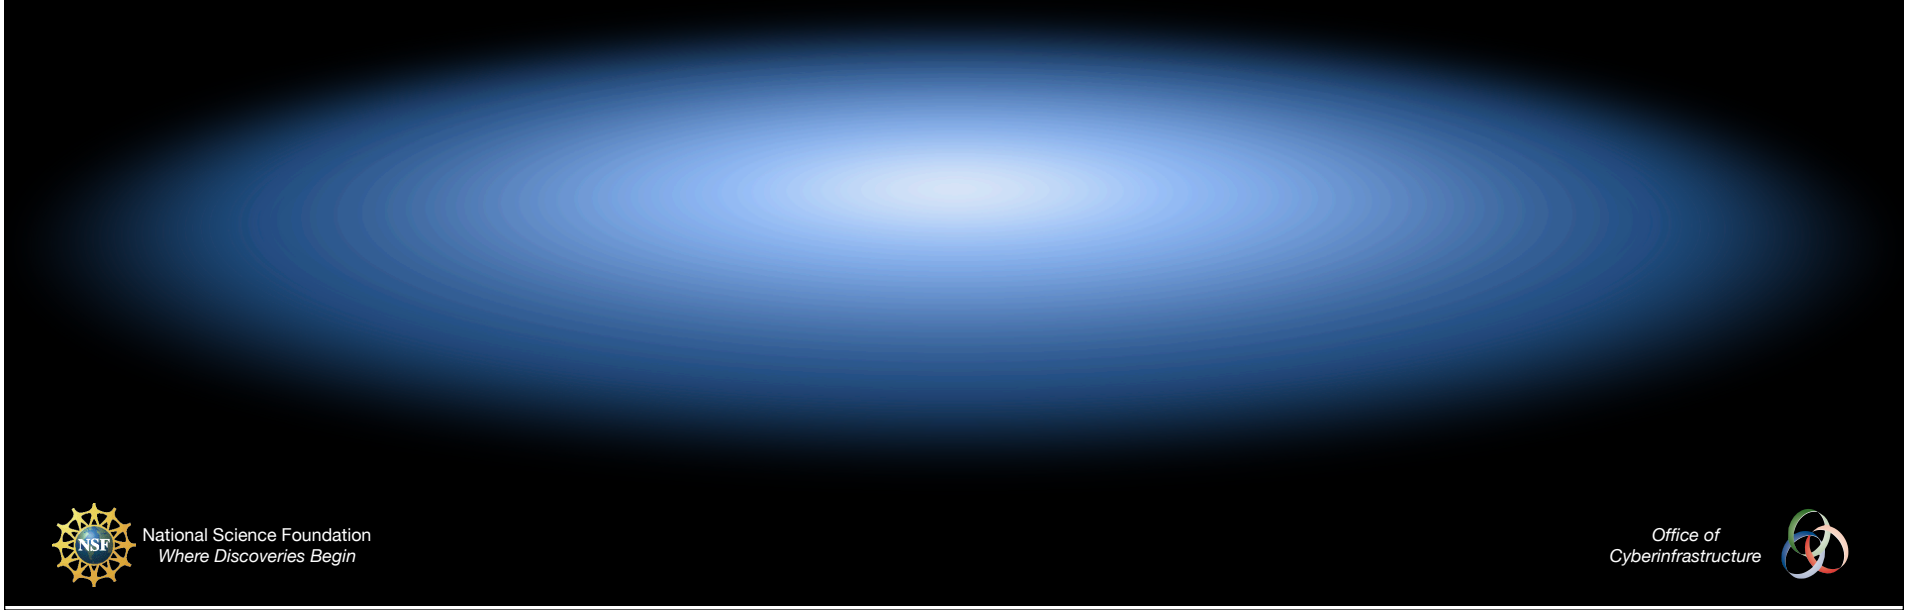

Judith E. Terrill
Leader, Scientific Applications and Visualization Group

Mathematical and Computational Sciences Division
Information Technology Laboratory
National Institutes of Standards & Technology

Some HEC work going on at NIST in the Mathematical and Computational Sciences Division

- Parallel Adaptive Multigrid: http://math.nist.gov/phaml/
- Modeling High Performance Concrete with Parallel Dissipative Dynamics: http://math.nist.gov/mcsd/savg/parallel/dpd/ (INCITE)
- Modeling Cement Paste Hydration and Microstructure Development: http://math.nist.gov/mcsd/savg/parallel/hydration/index.html
- Interoperable MPI: http://impi.nist.gov/
- Computation of Atomic Properties with the Parallel Hy-CI Method: http://math.nist.gov/mcsd/savg/parallel/atomic/
- Nanostructure Modeling: http://math.nist.gov/mcsd/savg/parallel/nano/
- OOMMF: Micromagnetic Modeling System: http://math.nist.gov/oommf/
- Screen Saver Science: http://math.nist.gov/mcsd/savg/parallel/screen/
- Modeling Fluid Flow in Complex Geometries with Parallel Lattice Boltzmann:
- http://math.nist.gov/mcsd/savg/parallel/lb/
- Immersive Scientific Visualization: http://math.nist.gov/mcsd/savg/vis/
- Others…

NIST funding through the Technology Innovation Program (TIP)

TIP Purpose

"Assisting United States businesses and institutions of higher education or other organizations, such as national laboratories and nonprofit research institutions, to support, promote, and accelerate innovation in the United States through high-risk, high-reward research in areas of *critical national need*."

America COMPETES Act
(PL 110-69)
August 9, 2007

Key Features of TIP

- **Novel Purpose: address societal challenges not being addressed in areas of Critical National Need with benefits that extend significantly beyond proposers**
- **Rich Teaming: businesses, academia, national labs, nonprofit research institutions and other organizations**
- **Scientific & Technical Merit: high-risk, high-reward research**
- **Transformational Results: strong potential for advancing state-of-the-art and contributing to U.S. science and technology base**
- **Societal Challenges: justifies government attention**

Key Features of TIP (continued)

- Clear Government Need: no other funding sources are reasonably available
- Allows institutions of higher education to lead a joint venture R&D project
- Intellectual property: Resides with U.S. company or any joint venture member (including a university joint venture member)
- Opportunities for state involvement with R&D planning
- Program assessment and annual reports from Program and Advisory Board to Congress are required

TIP Project Funding

- Funding
  - Single company projects up to $3M over a maximum of three years
  - Joint Venture (JV) projects may be funded up to $9M over a maximum of five years
  - Note: TIP funds direct project costs only
- Cost share
  - At least 50% of the yearly total project costs – direct plus indirect
  - Composed of both cash and in-kind

TIP Evaluation Criteria

- Two Evaluation Criteria
  - Scientific and Technical (S&T) Merit (50%)
  - Potential for S&T and National Impacts (50%)
- Subject to multidisciplinary peer review

A Pipeline of Critical National Need Topics

- Leverage nationally recognized science and technology reports and know-how
- Evaluate a large field of areas where transformative research could be expected to have large societal impact
- Use a TIP evaluation framework to assess a diversity of areas and challenges
  - Maps to Administration Guidance
  - Justifies Government Attention
  - Essentials for TIP Funding
- Identify interest areas that fit TIP

A Pipeline of Critical National Need Topics (cont'd)

- **TIP seeks input from a host of external stakeholders and organizations**
  - **Government agencies and advisory bodies (such as the National Research Council, the National Academy of Sciences, the National Academy of Engineering and the Institute of Medicine)**
  - **Science and Technology Policy Institute**
  - **Industry organizations, leading researchers from academic institutions, and others**

A Pipeline of Critical National Need Topics (cont'd)

- **TIP solicits white papers http://www.nist.gov/tip/call_for_white_papers.pdf**

2008: Civil Infrastructure

Advanced Sensing Technologies For Infrastructure: Roads, Highways, Bridges and Water Systems

$9M in grants expect to be awarded

- **Now accepting white papers to complement future critical national need identification: Deadline Dec 1, 2008**
- **TIP website available**
  - **http://www.nist.gov/tip/**

Technology Innovation Program (TIP)

For Info on TIP and to

Join the TIPMailing List . . .

- **Phone: 1-888-TIP-NIST (1-888-847-6478)**
- **Fax: 301-926-9524 or 301-590-3053**
- **E-mail: *tip@nist.gov***
- **Visit TIP's website:** *www.nist.gov/tip/*