Recent Activities in Networking and Information Technology R&D

Dr. David Nelson

nelson@itrd.gov

National Coordination Office for

Information Technology Research and Development

November 5, 2004

BOARD ON MATHEMATICAL SCIENCES AND THEIR APPLICATIONS

Topics

- **Initial implementation of High End Computing plan produced by High End Computing Revitalization Task Force (HECRTF)**
- **Initial findings from President's Information Technology Advisory Committee (PITAC) study of computational science**

Background: Networking and Information Technology Research and Development Program

- Helps focus interagency IT R&D:
  - Identify common research needs
  - Plan multi-agency research programs
  - Coordinate and collaborate on research announcements and funding
  - Review research results and adjust accordingly
- Includes R&D programs of twelve participating agencies totaling $2B
- Evolved from the Federal High Performance Computing and Communications Initiative, Computing Information and Communications Program , and Next Generation Internet Program
- Reports to National Science and Technology Council and OSTP in the Executive office of the President
- Assessed by the President's Information Technology Advisory Committee (PITAC)

Background: President's Information Technology Advisory Committee (PITAC)

- **Presidential Advisory IT Committee**
  - Members appointed directly by the President
- **Chartered by Congress in the High-Performance Computing Act of 1991 (P.L. 102-194) and Next Generation Internet Act of 1998 (P.L. 105-305)**
- **Recommendations have helped guide the NITRD program and its predecessors. Reports available at**
  - http://www.itrd.gov/pitac/
- **Current active reviews**
  - Computational Science
  - Cybersecurity

Revitalizing High End Computing

- *Federal Plan for High-End Computing* released May 10, 2004
- Initial implementation has begun

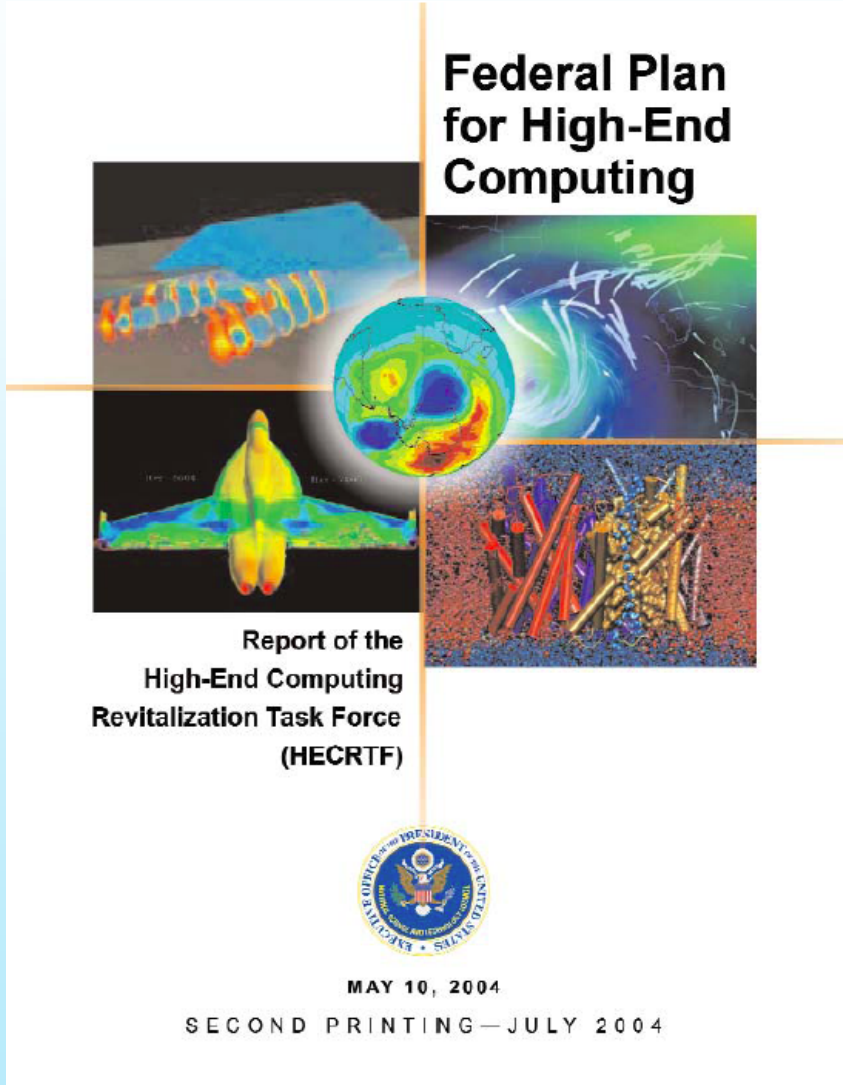

High-End Computing Revitalization Task Force (HECRTF)

- **Inter-agency planning group**
  - **Develop 5-year plan/roadmap to improve how the Federal government develops, purchases, and provisions HEC**
  - **Participants include DoD (DARPA, ODUSD (S&T), HPC Modernization Program, NSA), DOE (NNSA and Science), EPA, NASA, NIH, NIST, NOAA, NSF, OMB, OSTP, NCO (approx. 60 people)**
  - **Focus on advancing agency/end-user needs in HEC**
- **Established by OSTP, under the auspices of the National Science and Technology Council, in March 2003. Plan published May 10, 2004.**

Applications of High-End Computing:
Big Problems with Big Impacts

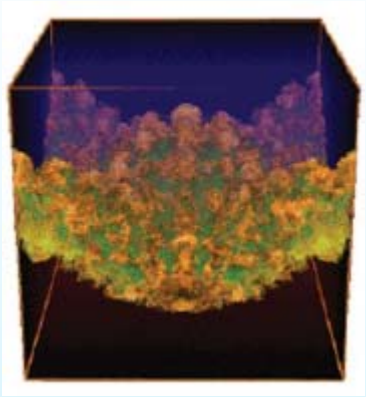

Nuclear Stockpile Stewardship

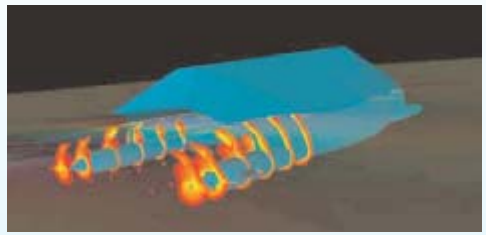

Ship Design

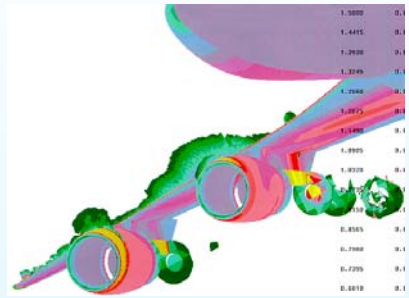

Aeronautics

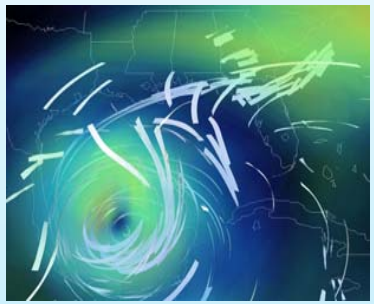

Weather Prediction

Cryptanalysis

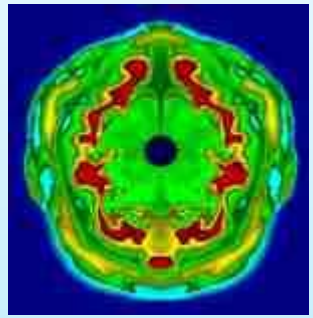

Astrophysics

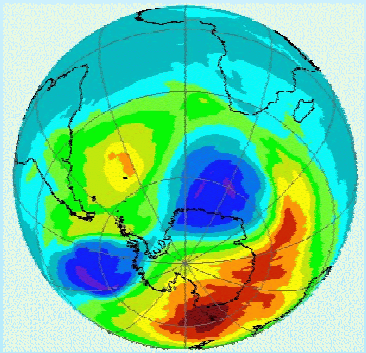

Climate Modeling

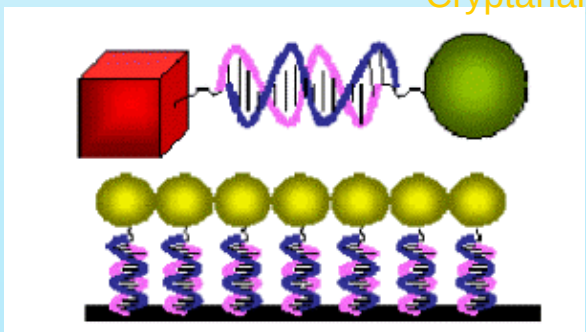

Nano-Science

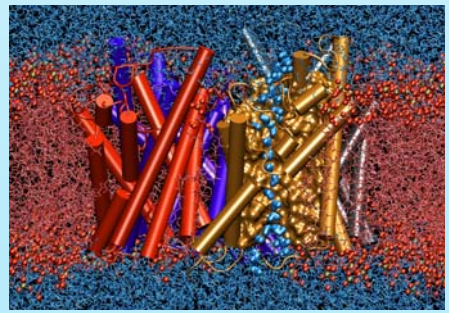

Biology

User and Agency Views on High-End Computing

- **Research pipeline dry**
- **Industrial base issues**
- **Technology improvement demanded by users**
  - **Radical improvements in time-to-solution;**
  - **Significant improvements to system bandwidth, reliability, ease of programming**
  - **Diversity of architectures**
  - **Better processes for creating and sustaining software**
- **User demand exceeds available resources for both capacity and capability**

High End Computing Revitalization Plan in a Nutshell

	Elements	Major Challenges Addressed
R&D	● Hardware, software, and systems roadmaps ● Basic and applied research, advanced development, engineering and prototypes, and test and evaluation – Research and evaluation systems – Life-cycle software strategy	● Improve performance, programmability, usability, and reliability for Agency applications ● Provide a range of robust HEC architectures and software technologies to address Agency requirements ● Re-establish research pipeline ● Ensure healthy research/tech/industry base
Resources	● Accessibility – Small / Large Agencies & Industry ● Availability – Production Computing ● Leadership – Largest systems for scientific leadership	● Lack of access to HEC resources by small agencies (e.g., NIST) ● Increasing demands for HEC exceed resources ● Large-scale systems to attack high-priority national problems
Procurement	● Pilot studies in benchmarking, total cost of ownership, and procurements	● Improve efficiencies in procurement for government and industry ● Improve evaluation methodologies of HEC systems for procurements and systems designs

High End Computing University Research Activity (HECURA)

- **Coordinated Solicitations**
  - **NSF, DOE/Office of Science, DARPA**
  - **Working group chair: Candy Culhane (NSA)**
- **Public Portal**
  - **http://www.itrd.gov/hecrtf-outreach/hec-ura/index.html**

HEC R&D Area	Lead Agency
Operating systems	DOE/Office of Science
Languages, compilers, and libraries	National Science Foundation

DOE/SC HEC-URA

- **Six awards totaling $4.5M in FY04 (including $1M from DARPA)**
- **Awards scope:**
  - Reliable, scalable, adaptive Linux
  - Dynamic OS/Application optimization
  - MicroKernels
  - Plan 9
  - Framework for OS/Runtime adaptability
- **FY2005 plans**
  - Grow by $2M
  - Additional awards
  - Continued emphasis on research community building

NSF HEC-URA

- **Nine projects were awarded totaling $7M per year (including DARPA funding)**
  - Tools and libraries for high-end computing
  - Resource management
  - Reliability of high-end systems
- **NSF hopes to continue this activity in FY 2005**

(Proto) Leadership-class Systems

- **NASA Project Columbia**
  - Partnership with SGI and Intel
  - Cluster of 20 interconnected SGI® Altix® 512-processor systems; 10,240 Itanium 2 processors
- **DOE Leadership-class system at ORNL**
  - 20 TF Cray X1 + 20 TF Cray Red Storm growing to 100TF in 2006
- **DOE and NASA discussing possibility for coordinated availability of resources to community**

Augmentation of HPCS Program

- **Several agencies working with DARPA**
  - NASA, NSF, NSA, NRO, DOE/NNSA, DOE/SC, DOD/HPCMP
- **Emphasizing performance characterization and software productivity**
  - HPC Challenge
- **Likely to be further expanded in scope and participants**

Agency-Industrial Partnerships

- **Purpose: to bring together agencies and companies interested in application area**
  - Possible examples: combustion, multi-scale modeling, CFD
- **Organized jointly by Council on Competitiveness (CoC) and NITRD program**
- **Outgrowth of CoC workshop on HPC last summer**
- **Process:**
  - NITRD agencies selecting areas of interest
  - CoC will seek industry interest
  - Future joint meetings to define projects
- **Could involve substantial computer resources**

PITAC Computational Science Subcommittee

- Requested by Director of OSTP on June 9, 2004
- Charge letter available at http://www.itrd.gov/pitac/20040609_compsci_charge.pdf
- Charge includes study of
  - Appropriateness of agency research priorities
  - Balance between short and long term research
  - Development vs. application of Computational Science
  - Integration of Computational science training and research into scientific disciplines that apply CS
  - Coordination among agencies
  - Response of agencies to change
  - Elimination or mitigation of barriers to potential of CS
- Subcommittee led by Prof. Dan Reed
- Public input very welcome. Email to nco@itrd.gov

Context for Request that PITAC Examine Computational Science

- **Government has played important role in fostering CS**
- **CS is very important to agency missions**
- **CS Contributes to economic prosperity**
- **Yet, a sense that "all is not well"**
  - Ageing population of computational scientists
  - Difficulty of using current computational environments
  - Limited progress on recognized CS problems (e.g. Lax Report)
  - Concern that CS will not achieve potential
  - Debates about Federal investment strategy (e.g. software maintenance)

PITAC Preliminary Observations*

Computational Science: Essential to Scientific Discovery (1/2):

- **Computing has become the third component of scientific discovery, complementing theory and experiment.**
- **Computing is so integral to the scientific process that its limitations now constrain scientific discovery.**
- **The explosive growth in the resolution of sensors and scientific instruments has led to unprecedented volumes of experimental data. Computational science now broadly includes modeling, simulation, and scenario assessment using sensor data from diverse sources.**

* http://www.itrd.gov/pitac/meetings/2004/20041104/compsci.pdf

PITAC Preliminary Observations

Computational Science: Essential to Scientific Discovery (2/2):

- Complex multidisciplinary problems, from public policy through national security to scientific discovery and economic competitiveness, have emerged as new drivers of computational science, complementing the historical focus on single disciplines.
- Developing leading edge computational science applications is a complex process involving teams of people that must be sustained for a decade or more to yield the full fruits of investment.

PITAC Preliminary Observations

Current State of Computational Science:

- **There is a disconnect between commercial practice and the computing infrastructure needs of government and academia. Commercial needs are (in several cases) no longer driving technology acceleration.**

- **Short-term investment and limited strategic planning have led to excessive focus on incremental research rather than on the long-term research with lasting impact that can solve critical problems.**

Preliminary Observations

Paths to Progress (1/3):

- **Computational science would benefit from a roadmap outlining decadal priorities for investment, with a clear assessment of those priorities derived from a survey of the problems and challenges. Agencies could then respond to these with a strategic plan in recognition of those priorities and funding requirements.**

Preliminary Observations

Paths to Progress (2/3):

- **Solutions must be "right sized" for the problems**
  - temporally, recognizing the time to solution
  - socially and fiscally, recognizing complexity and sustainability

- **Diverse solutions are needed for different structural issues**
  - community organization/coordination for increased leverage
  - structural infrastructure investment
    - creating baselines for community research and development
  - coordination across agencies and missions for R&D transfer

Preliminary Observations

Paths to Progress (3/3):

- **Strategic execution, based on systemic assessment of programs and components within a long-term, strategic plan, is needed within and across agencies to create a vibrant, holistic research environment and infrastructure. Individual programs and solicitations must be viewed and managed within the context of strategic and tactical goals.**

- **Sustained investment in computational science infrastructure, defined broadly to include people, software, data, and systems, is needed to fully realize the promise of computational science.**

Preliminary Observations

Paths to Progress - People (1/2):

- **The limited number of senior leaders in computational science has constrained community advocacy and agency leadership.**

- **Interdisciplinary education in computational science and computing technologies is inadequate, reflecting the traditional disciplinary boundaries in higher education. Only systemic change to university organizational structures will yield the needed outcomes.**

Preliminary Observations

Paths to Progress - People (2/2):

- **There are few, if any, rewards for interagency coordination and collaboration on science, technology, and infrastructure development pipelines. The result has been loss of critical opportunities to sustain and develop critical capabilities, and transfer them to the commercial sector.**

Preliminary Observations

Paths to Progress - Software:

- **Easy to use, accessible, scalable software that interoperates with existing user environments is not adequately available.**

- **Community verification and validation of computational science results, via access to the software and data, are needed to accelerate scientific discovery.**

Preliminary Observations

Paths to Progress - Hardware:

- **National computing resources, high end computers, are not readily accessible/available to both small and large agencies and industry. Even when such systems are available, they are not sufficiently easy to use.**

- **A sustainable infrastructure is needed that provides access to leading edge capabilities for computational science. This requires long term investments.**

Partial List of Presenters to PITAC

- Jim Crowley, SIAM
- Bob Lucas, UCLA
- Phil Colella, LBNL
- Ed Seidel, LSU
- Charbel Farhat, Stanford
- Kelvin Droegemeier, U. Oklahama
- Michael Vannier, U. Chicago
- Jonathan Silverstein, Lily Research
- Vernon Burton, NCSA
- Jack Dongarra, U. Tenn.
- Al Trivelpiece, retired
- Andre van Tilborg, DoD
- Walt Brooks, NASA
- Tim Killeen, NCAR
- Chris Johnson, U. Utah
- Mike Holland, OSTP

Recent Workshops and Reports

- **Blueprint for Future Science Middleware and Grid Research and Infrastructure, August 2002**
  - http://www.nsf-middleware.org/MAGIC/default.htm
- **NSF Cyberinfrastructure Report, January 2003**
  - http://www.cise.nsf.gov/evnt/reports/toc.htm
- **DOE Science Network Meeting, June 2003**
  - http://gate.hep.anl.gov/may/ScienceNetworkingWorkshop/
- **DOE Science Computing Conference, June 2003**
  - http://www.doe-sci-comp.info
- **DOE Science Case for Large Scale Simulation, June 2003**
  - www.pnl.gov/scales/
- **DOE ASCR Strategic Planning Workshop, July 2003**
  - http://www.fp-mcs.anl.gov/ascr-july03spw
- **High End Computing Revitalization Task Force, 2003-2004**
  - http://www.itrd.gov/hecrtf-outreach/

Upcoming Public Input Opportunities

- **SC04 Birds of a Feather (BOF)**
  - November 10, 5:30-7:00pm, Room 303/304/305
  - National Priorities for Computational Science: A PITAC Town Hall Meeting
  - *3 minute position statements*
  - *Written statements also welcome*

- **Written input to nco@nitrd.gov, copy to Dan_Reed@unc.edu**
- **Public PITAC meeting on Computational Science in January**

Summary

- **High End Revitalization Task Force was major planning activity for Federal agencies**
  - Focus on needs of computational science and engineering
  - Implementation begun already in FY04
- **PITAC review of computational science underway under Dan Reed**
  - complements and extends work of Task Force
  - Aim is to realize potential for computational science
  - Public input welcome

For Further Information

Please contact us at:

nco@nitrd.gov

Or visit us on the Web:

www.nitrd.gov

Backup

HECRTF Goals

- Make high-end computing easier and more productive to use.
- Foster the development and innovation of new generations of high-end computing systems and technologies.
- Effectively manage and coordinate Federal high-end computing.
- Make high-end computing readily available to Federally Agencies that need it to fulfill their missions.

OSTP-OMB Memo on FY06 Research Priorities

"Networking and Information Technology R&D

The Networking and Information Technology **R&D (NITRD)** program is a high Administration priority. While the importance of each of the NITRD program areas continues, high-end computing (supercomputing) and cyberinfrastructure R&D should be given higher relative priority due to the potential of each in furthering progress across a broad range of scientific and technological application areas. The recent report of the High-End Computing Revitalization Task Force (HECRTF) describes a coordinated R&D plan for core high-end computing technology, as well as multi-agency approaches for addressing high-end computing capability, capacity, and accessibility issues. Agency plans in high-end computing should be consistent with the HECRTF plan, emphasize coordination, leverage the efforts of all agencies and, where appropriate, provide explicit benefit to multiple agencies through coordinated multi-agency investments."

http://www.ostp.gov/html/m04-23.pdf

Agencies and Departments Participating in NITRD Program

- Department of Defense
  - Defense Advanced Research Projects Agency (DARPA)
  - National Security Agency (NSA)
  - Office of the Director of Defense Research and Engineering (ODDR&E)
  - Defense Information Systems Agency (DISA)
- Department of Energy
  - Office of Science (DOE/SC)
  - National Nuclear Security Administration (DOE/NNSA)
- Department of Health and Human Services
  - National Institutes of Health (NIH)
  - Agency for Health Research and Quality (AHRQ)
- Department of Commerce
  - National Institute of Standards and Technology (NIST)
  - National Oceanic and Atmospheric Administration (NOAA)
- National Science Foundation (NSF)
- National Aeronautics and Space Administration (NASA)
- Environmental Protection Agency (EPA)
- Observers: Federal Aviation Administration (FAA), Food and Drug Administration (FDA)