Computational Science: America's Competitive Challenge

President's Information Technology Advisory Committee

Subcommittee on Computational Science

Daniel A. Reed, Chair

Dan_Reed@unc.edu

PITAC Meeting, Washington, D.C.

April 14, 2005

Computational Science Subcommittee

- Subcommittee
  - ***Daniel A. Reed***, Ph.D., ***Chair***, Chancellor's Eminent Professor, Vice-Chancellor for Information Technology, and Director, Renaissance Computing Institute, University of North Carolina at Chapel Hill
  - ***Ruzena Bajcsy,*** Ph.D., Director, Center for Information Technology Research in the Interest of Society (CITRIS), University of California, Berkeley
  - ***Manuel A. Fernandez***, Managing Director, SI Ventures
  - ***José-Marie Griffiths***, Ph.D., Dean, School of Information and Library Sciences, University of North Carolina at Chapel Hill
  - ***Randall D. Mott***, Senior Vice President and CIO, Dell Computer
- Consultants
  - ***Jack Dongarra***, University Distinguished Professor, Tennessee
  - ***Chris Johnson***, Distinguished Professor, Utah

Subcommittee Work Plan

- June 9, 2004: Charge from the White House
- June 17: PITAC meeting in Arlington, Virginia
- September 16: Information gathering meeting in Chicago
- October 19: Information gathering meeting in Arlington, Virginia
- November 4: PITAC meeting in Arlington, Virginia
- November 10: Birds of a Feather (BOF) at Supercomputing 2004 in Pittsburgh
- November to January 2005: Revised findings & recommendations and solicited additional input
- January 11, 2005: Call for information sent to Federal agencies involved in computational science
- January 12: PITAC meeting - Status report on findings & recommendations
- January to April: Draft report
- *April 14 PITAC meeting – PITAC deliberation on the draft report*

Computational Science Definition

Computational science is a rapidly growing multidisciplinary field that uses advanced computing capabilities to understand and solve complex problems.

Computational science fuses three distinct elements:

- *algorithms* (numerical and non-numerical) and *modeling and simulation software* developed to solve science (e.g., biological, physical, and social), engineering, and humanities problems,
- *computer and information science* that develops and optimizes the advanced system hardware, software, networking, and data management components needed to solve computationally demanding problems; and
- *the computing infrastructure* that supports both the science and engineering problem solving and the developmental computer and information science.

Computational Science: America's Competitive Challenge

This report consists of five chapters:

1. *A Wakeup Call: The Challenges to U.S. Preeminence and Competitiveness*
2. *Medieval or Modern? Research Structures for the 21st Century*
3. *Multidecadal Roadmap for Computational Science*
4. *Sustained Infrastructure for Discovery and Competitiveness*
5. *Research and Development Challenges*

A Wakeup Call:

The Challenges to U.S. Preeminence and Competitiveness

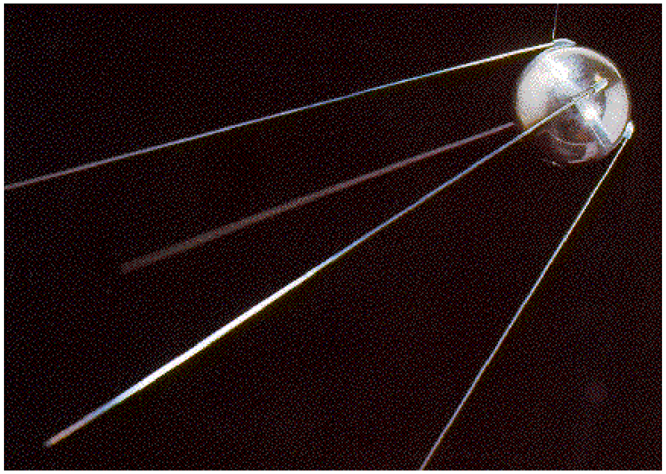

Sputnik: October 4, 1957

Challenges and Opportunities

- Computational science is central to key sectors
  - we have achieved some major successes
  - to a larger degree, we have missed opportunities
- U.S. science and engineering *leadership is in jeopardy*
  - computational science should be a major driver for
    - scientific progress
    - economic competitiveness
    - national security
- *There are obstacles to progress (lessons not learned)*
  - *investments: short-term fulfillment over long-term vision*
  - *planning: incremental and tactical rather than strategic*

The Third Pillar of 21st Century Science

- Three pillars
  - theory, experiment and *computational science*
- Computational science enables us to
  - investigate phenomena where economics or constraints preclude experimentation
  - evaluate complex models and manage massive data volumes
  - model processes across interdisciplinary boundaries
  - transform business and engineering practices

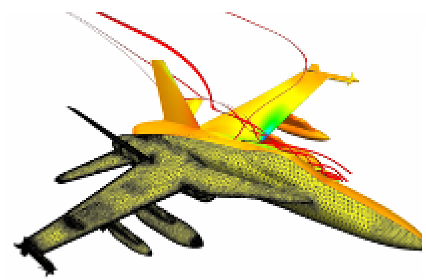
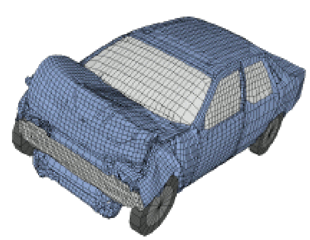
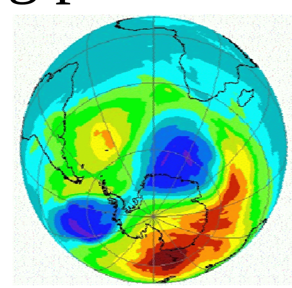


Analyze Complex Data

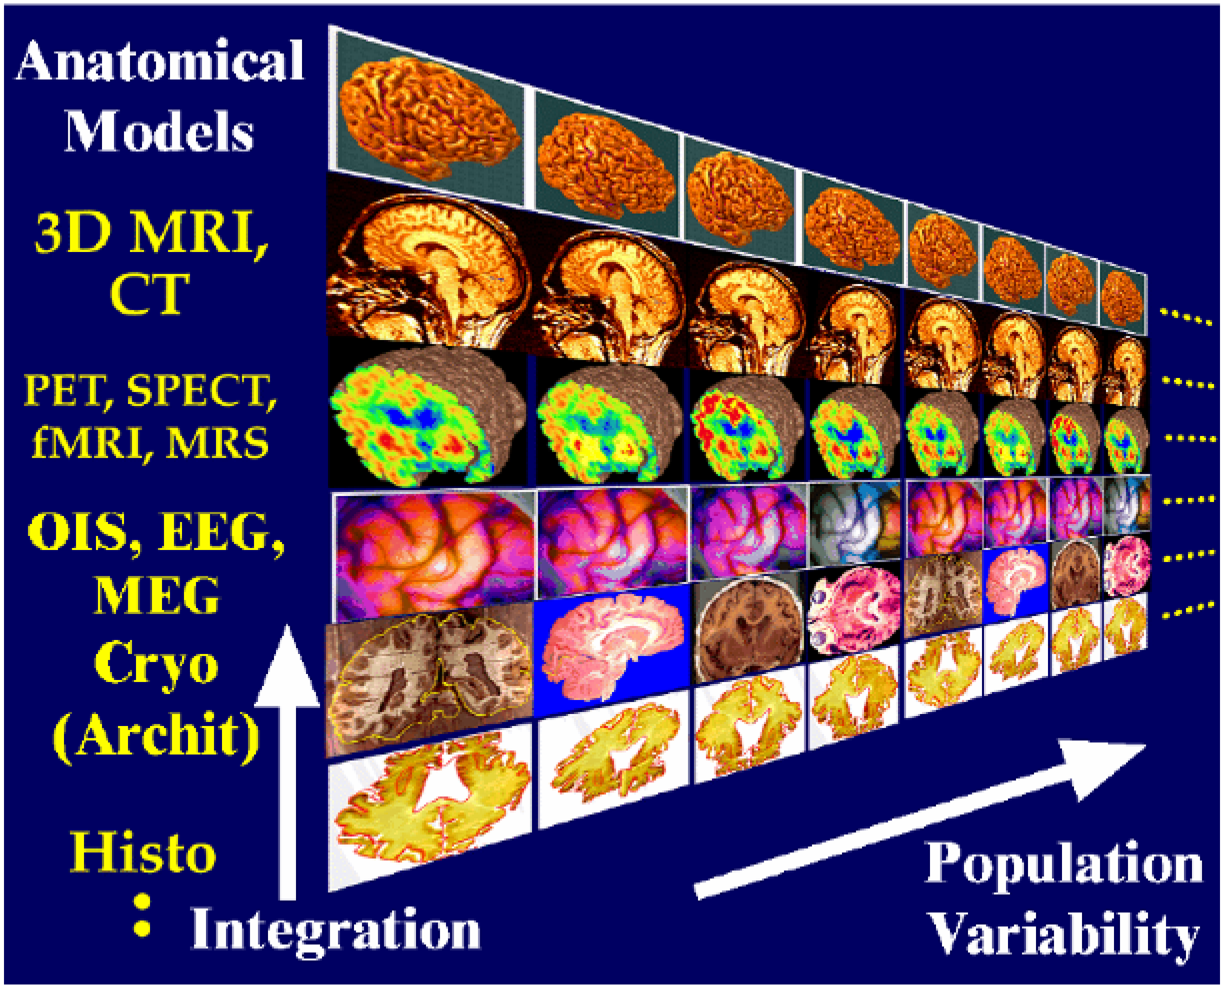

Source: Chris Johnson, Utah and Art Toga, UCLA

Conduct Interdisciplinary Analysis

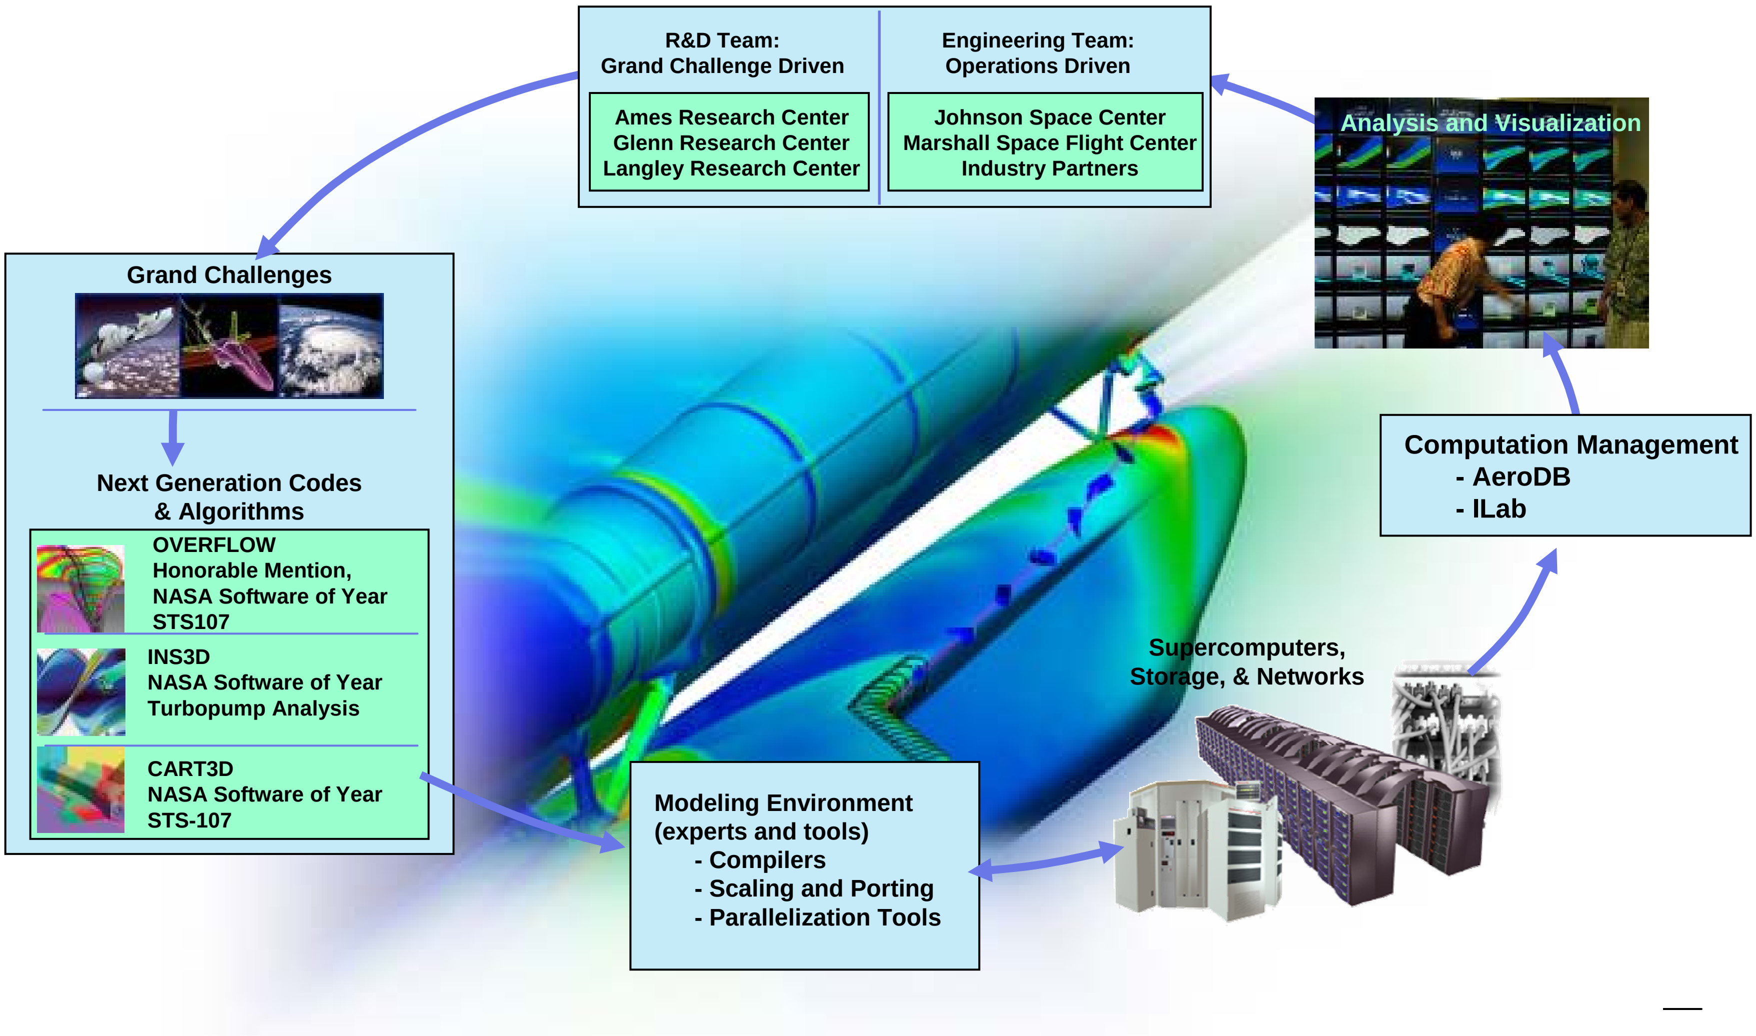

Source: Walt Brooks, NASA

Transform Business Practices

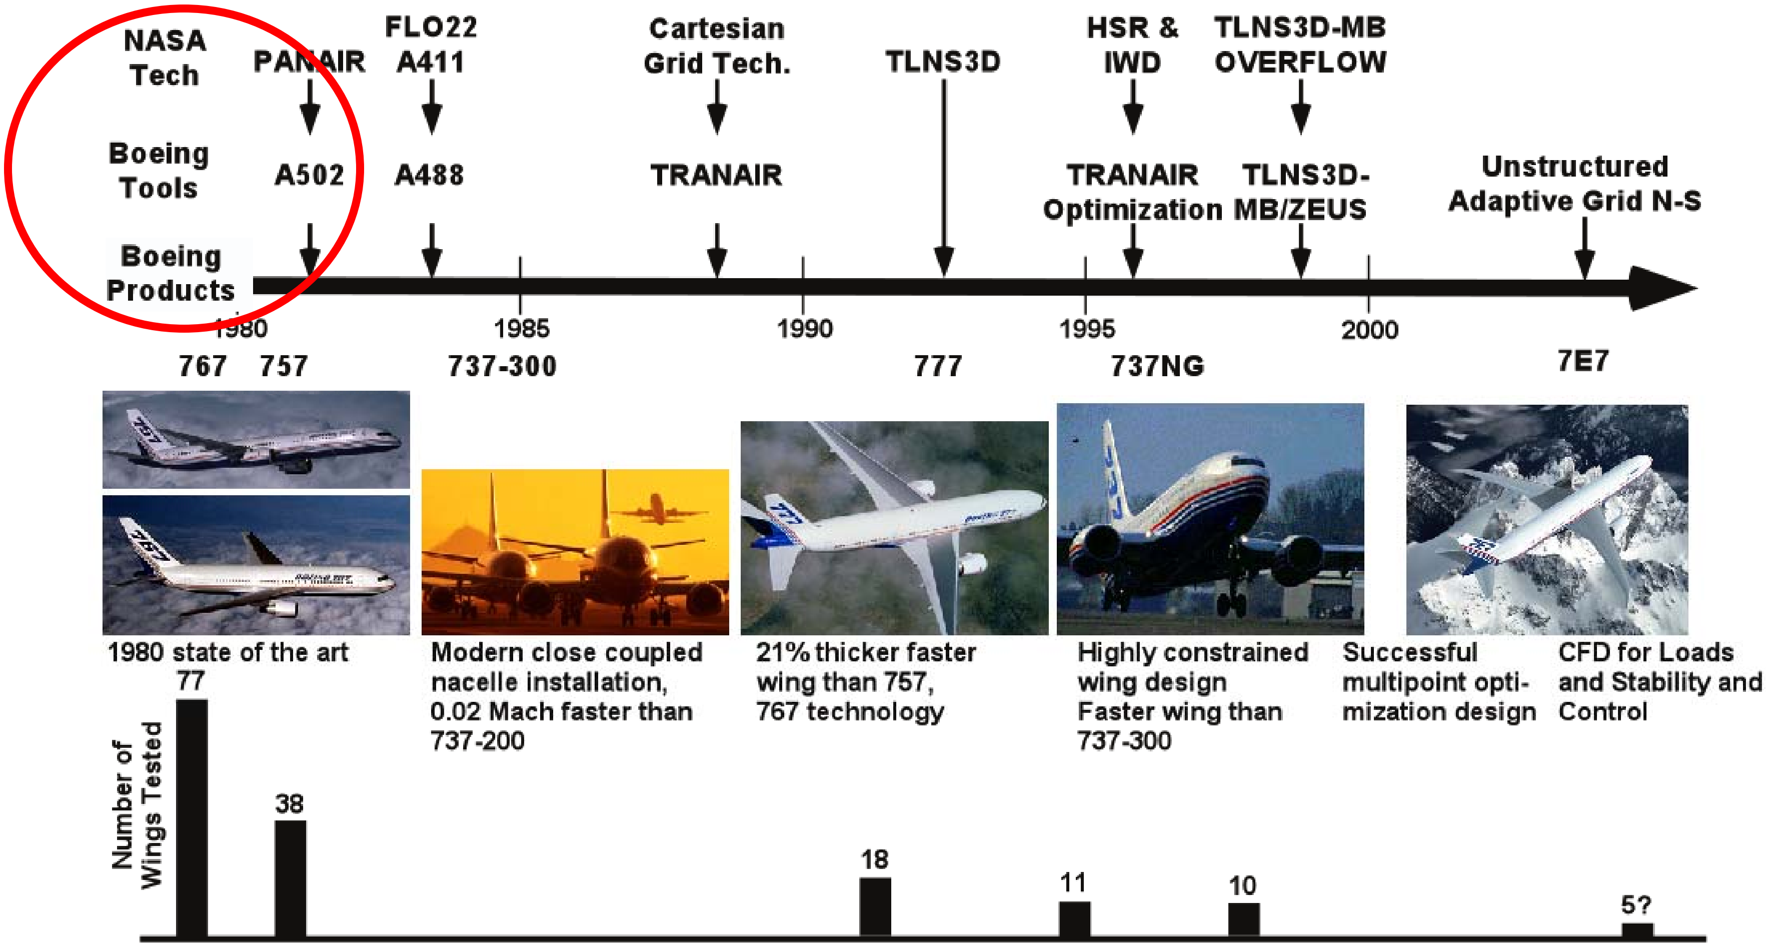

BOEING

Principal Finding

- Computational science is indispensable for solving complex problems in every sector, from traditional science and engineering domains to such key areas as national security, homeland security, and public health
- Advances in computing and connectivity and ability to capture and analyze huge amounts of data make it increasingly possible and practical to address these complex problems

- *Universities and Federal government have not effectively recognized the strategic significance of computational science*
- *These inadequacies compromise U.S. scientific leadership, economic competitiveness, and national security*

Worrisome Indicators

- National security benchmark (2000 DSB study)
  - giga-updates per section (GUPS)
    - performance has not increased substantively in a decade
  - architectures not well matched to needs
- *Council on Competitiveness* survey of U.S. businesses
  - overwhelming majority: technical computing critical to survival
  - key limitations to broader use
    - lack of trained personnel and difficulty using systems
- National Academies supercomputing study (2004)
  - ~35 computational science graduates/year
  - NSF architecture grant awards down by 75%

Lessons Not Learned

- *Panel on Large Scale Computing in Science and Engineering*, interagency, 1982
- *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing*, NSF, 1993
- *Information Technology Research: Investing in Our Future*, PITAC 1999
- *The Biomedical Information Science and Technology Initiative*, NIH, 1999
- *Making IT Better*, National Academies, 2000
- *Embedded Everywhere*, National Academies, 2001
- *High-Performance Computing for the National Security Community*, DOD, 2002
- *Knowledge Lost in Information*, NSF, 2003
- *Revolutionizing Science and Engineering Through Cyberinfrastructure*, 2003
- *A Science-Based Case for Large-Scale Simulation (ScALES)*, DOE, 2003
- *Roadmap for the Revitalization of High End Computing*, Interagency, June 2003
- *Supercharging U. S. Innovation & Competitiveness*, Council on Competitiveness, 2004
- *Getting up to Speed: the Future of Supercomputing*, National Academies, 2004

Lessons Not Learned

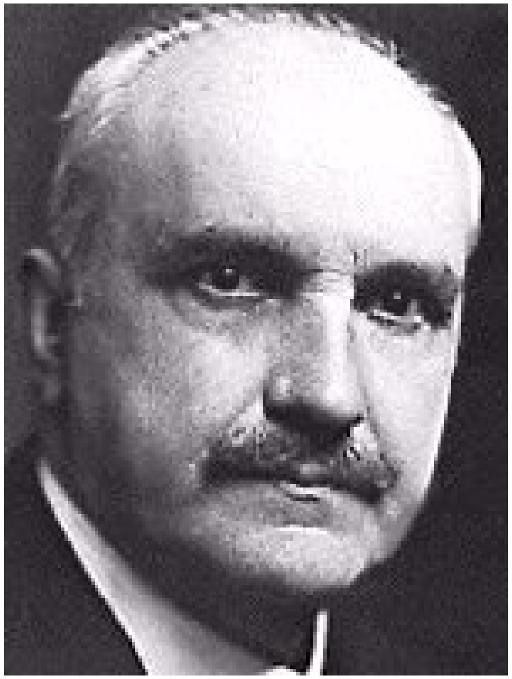

George Santayana

Those who do not learn from history are condemned to repeat it.

George Santayana

Principal Recommendation

- Universities and Federal R&D agencies must *make coordinated, fundamental, and structural changes* that affirm the integral role of computational science
  - the most important problems are multidisciplinary, multi-agency, multi-sector, and collaborative
- Federal government, in partnership with academia and industry, must also *create and execute a multi-decade roadmap* that directs coordinated advances in computational science and its applications in science and engineering

Fundamental Change Is Needed

- Roadmap must drive
  - flexible organizational structures
    - reduce silos and barriers
  - coordinated budget plans
    - recognize that local decisions have global impact
- Coordination must increase
  - within and across Federal R&D agencies
  - within and across universities
- Planning horizons must increase
  - a succession of short term plans is *not* a long term plan
  - some elements require 10-20 year horizons
- Risk portfolio must broaden
  - short-term, derivative research is itself a risk

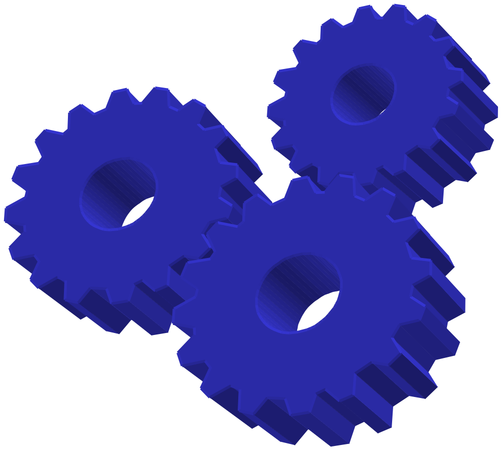

Medieval or Modern?

Research Structures for the 21st Century

Findings

- Traditional disciplinary boundaries within academia and Federal R&D agencies severely inhibit the development of effective research and education in computational science
- Paucity of incentives for longer-term multidisciplinary, multi-agency, or multi-sector efforts stifles structural innovation

Recommendations for Academia

- *Universities must significantly change their organizational structures* to promote and reward collaborative research

- *Universities must implement new multidisciplinary structures* to provide rigorous, multifaceted educational preparation for the growing ranks of computational scientists that the Nation will need to remain at the forefront of scientific discovery

Recommendations for Government

- The National Science and Technology Council must *commission a fast track study* by the National Academies to *recommend changes and innovations in Federal R&D agencies' roles and portfolios* to support revolutionary advances in computational science

- *Individual agencies must implement changes and innovations in their organizational structures* to accelerate and advancement of computational science

Organizational Structure and Practices

- Change in organizational structures is *very* slow
  - academic and Federal agencies reinforce each other
  - organizational silos limit adaptation and nimble response
- Crosscutting academic centers
  - have sunset clauses and fixed lifetimes
  - *do not address the educational issues*
- Computing's universality is a political weakness
  - everyone's "second priority" without an organizational home
- Greater coordination is required
  - most activities are short-term and low-risk
  - local priorities must be derived from global planning

Education and Leadership

- Interdisciplinary education
  - key problems are increasingly interdisciplinary
  - reward metrics and mechanisms must encourage interdisciplinary collaboration and education
    - foster experiential and collaborative learning environments and tie to ongoing R&D efforts
  - there are few computational science degree programs
- Cultivating leaders for computational science
  - limited number of senior leaders in computational science. Need sustained leadership training program
  - constraints on Federal employment policies and practices
    - significant disincentives for service

Multidecadal Roadmap for Computational Science

Finding

- Scientific needs stimulate exploration and creation of new computational techniques and, in turn, these techniques enable exploration of new scientific domains
- The continued health of this dynamic computational science "ecosystem" demands long-term planning, participation, and collaboration by Federal R&D agencies and computational scientists in academia and industry
- *Instead, today's Federal investments remain short-term in scope, with limited strategic planning and a paucity of cooperation across disciplines and agencies*

The Ecosystem Really Matters

- Local change at a single agency can
  - upset national agendas
  - have long-term repercussions
- Disjoint planning means
  - the holes may never be filled

Recommendation

- The National Science and Technology Council (NSTC) must commission the National Academies to convene, on a *fast track*, one or more task forces to *develop and maintain a multi-decade roadmap* for computational science and the fields that require it, with a goal of assuring continuing U.S. leadership in science, engineering, and the humanities
- The roadmap must be assessed and *updated every five years*, and Federal R&D agencies' progress in implementing it must be *assessed every two years by PITAC*

Roadmap Components and Needs

- At a minimum, the roadmap must address
  - computing system hardware, networking, software
  - data acquisition, storage, and visualization
  - algorithms and applications
    - science, engineering, and humanities
- Prioritize the especially problematic issues
  - inadequate software
  - lack of sustainable infrastructure
  - education and training
- Recognize ecosystem issues and interdependencies
  - effective planning must be holistic

Notional Computational Science Roadmap

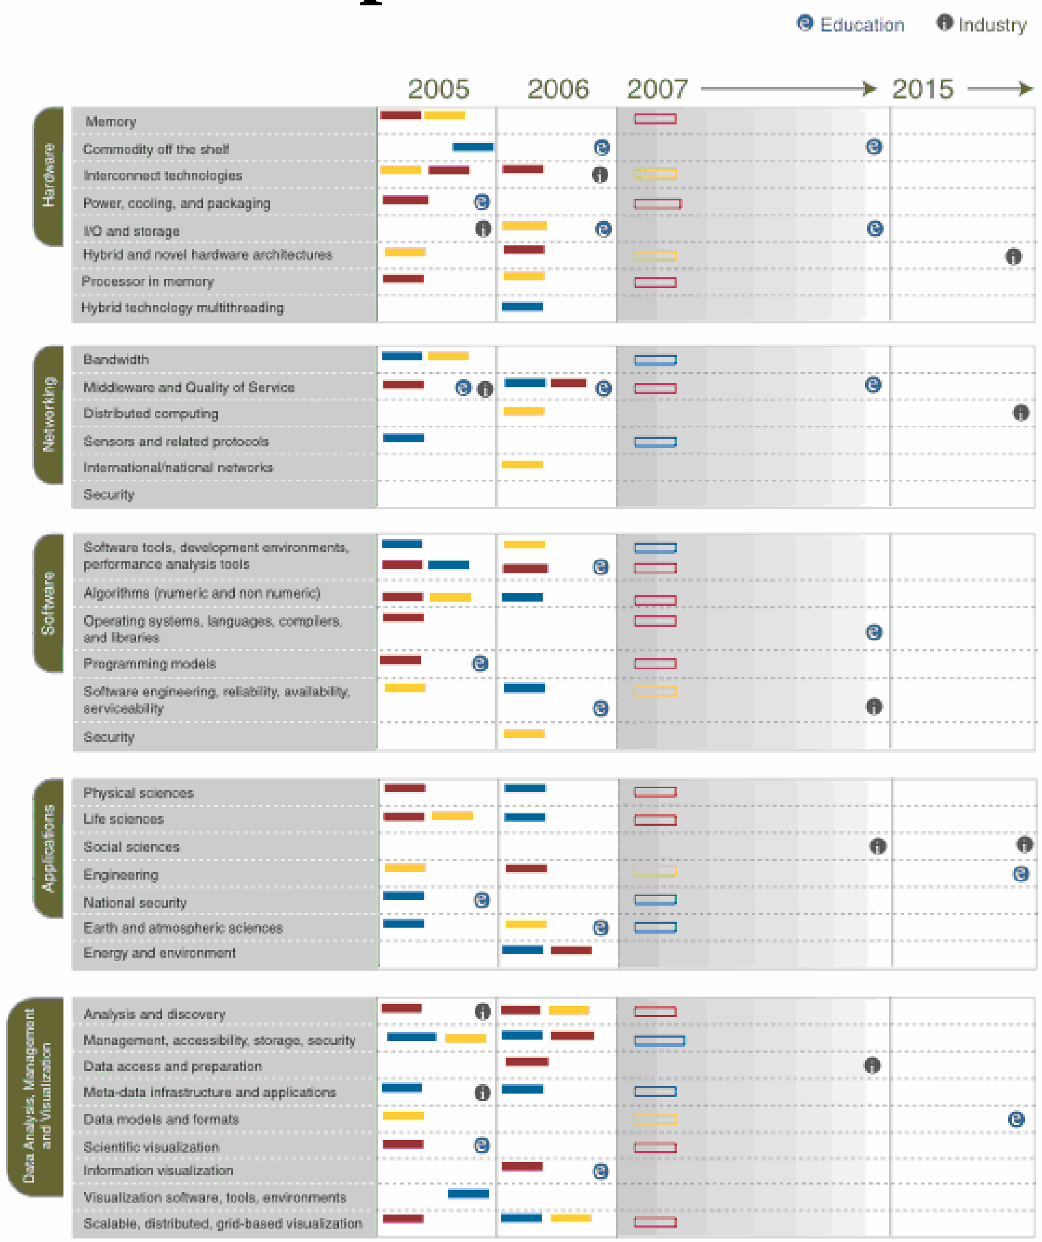

Roadmap Process

- Quantitative and measurable milestones and timelines should be an integral part of the roadmap process
- Roadmap creation and maintenance should occur via an open process that involves broad input from industry, academia and government
- Roadmap creation and evaluation should include the status of computational science activities across industry, government and academia
- Roadmap should specify opportunities for cross-fertilization of agency activities, successes and challenges
- Agency strategies for computational science should be shaped in response to the roadmap
  - strategic plans should recognize and address roadmap priorities and funding requirements

Sustained Infrastructure for Discovery and Competitiveness

Nothing tends so much to the advancement of knowledge as the application of a new instrument. The native intellectual powers of men in different times are not so much the causes of the different success of their labors, as the peculiar nature of the means and artificial resources in their possession.

Sir Humphrey Davy

The Need for Sustained Infrastructure

- At least four national elements
  - software sustainability centers
  - data and software repositories
  - high-end computing leadership centers
  - community integration and sustenance

- The National Science Board (2003) noted that academic research infrastructure "… has not kept pace with rapidly changing technology, expanding research opportunities, and an increasing number of (facility) users"

Software Sustainability Centers Finding

- The computational science ecosystem is unbalanced
  - software base is woefully inadequate
- Imbalance forces researchers to build atop crumbling and inadequate foundations rather than on a modern, high-quality software base
- The result is greatly diminished productivity for both researchers and computing systems

Software Sustainability Centers Recommendation

- *Federal government must establish national software sustainability centers* whose charge is to harden, document, support, and maintain vital computational science software whose useful lifetime may be measured in decades
- Software areas and specific software artifacts must be chosen in consultation with academia and industry
- *Software vendors must be included* in collaborative partnerships to develop and sustain the software infrastructure needed for research

National Data and Software Repositories Finding

- Explosive growth in the number and resolution of sensors and scientific instruments has engendered unprecedented volumes of data, presenting historic opportunities for major scientific breakthroughs in the 21st century

- Computational science now encompasses modeling and simulation using data from these and other sources, requiring data management, mining, and interrogation

National Data and Software Repositories Recommendation

- Federal government must *provide long-term support for computational science community data repositories*
  - defined frameworks, metadata structures
  - algorithms, data sets, applications
  - review and validation infrastructure
- Government must *require funded researchers to deposit their data and research software in these repositories* or with access providers that respect any necessary or appropriate security and/or privacy requirements

National High-End Computing Leadership Centers: Finding

- High-end computing resources are not readily accessible and available to researchers with the most demanding computing requirements
- High capital costs and the lack of computational science expertise preclude access to these resources
- Moreover, available high-end computing resources are heavily oversubscribed

National High-End Computing Leadership Centers: Recommendation

- Government must *provide long-term funding for national high-end computing centers* at levels sufficient to ensure the regularly scheduled deployment and operation of the fastest and most capable high-end computing systems that address the most demanding computational problems
- In addition, capacity centers are required to address the broader base of users.
- Federal government must *coordinate high-end computing infrastructure across R&D agencies* in concert with the roadmapping activity

Infrastructure, Community and Sustainability: Finding

- Computational science ecosystem is a national imperative for research and education in the 21st century. Like any complex ecosystem, the whole flourishes only when all its components thrive – the computational science applications, the human resources and time needed to create them, and the physical infrastructure on which they depend.
- Only sustained, coordinated investment in people, software, hardware, and data, based on strategic planning, will enable the U.S. to realize the promise of computational science to revolutionize scientific discovery, increase economic competitiveness, and enhance national security.

Infrastructure, Community and Sustainability: Recommendation

- The Federal government must *implement coordinated, long-term computational science programs* that include funding for interconnecting the software sustainability centers, national data and software repositories, and national high-end leadership centers with the researchers who use those resources, *forming a balanced, coherent system* that also includes regional and local resources.
- Such funding methods are customary practice in research communities that use scientific instruments such as light sources and telescopes, increasingly in data-centered communities such as those that use the genome database, and in the national defense sector.

Research and Development Challenges

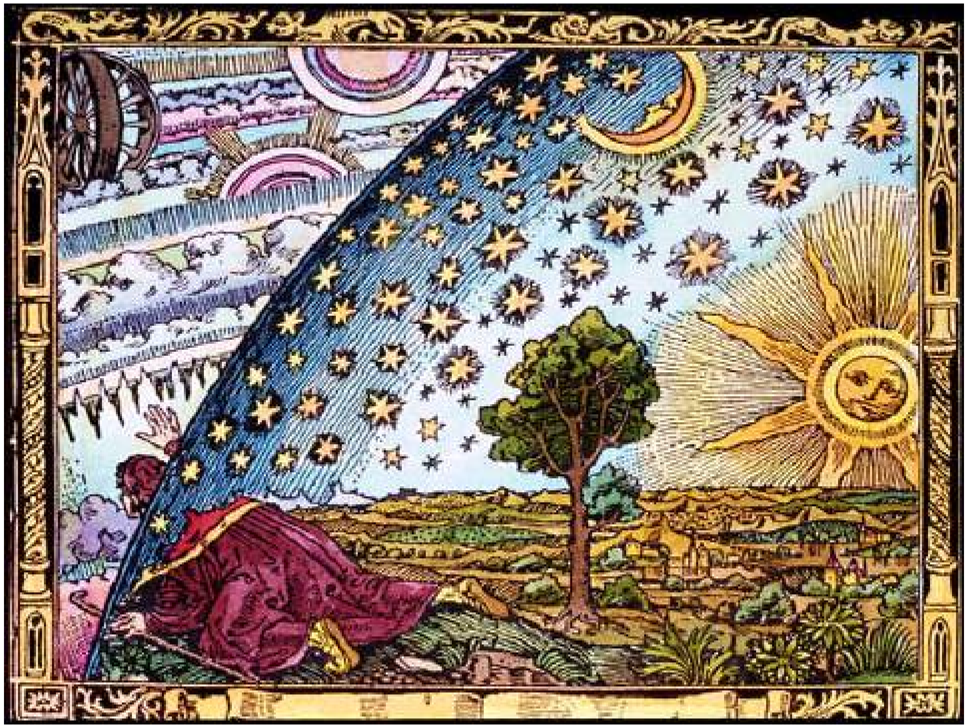

Finding

- Leading-edge computational science is possible only when supported by *long-term, balanced research and development investments* in software, hardware, data, networking, and human resources.

- Inadequate investments in robust, easy-to-use software, an excessive focus on peak hardware performance, limited investments in architectures well matched to computational science needs, and inadequate support for data infrastructure and tools have endangered U.S. scientific leadership, economic competitiveness, and national security.

Recommendation

The Federal government must rebalance research and development investments to:

a) *create a new generation of well-engineered, scalable, easy-to-use software* suitable for computational science that can reduce the complexity and time to solution for today's challenging scientific applications and can create accurate simulations that answer new questions;

b) *design, prototype, and evaluate new hardware architectures* that can deliver larger fractions of peak hardware performance on scientific applications; and

c) *focus on sensor- and data-intensive computational science applications* in light of the explosive growth of data.

Computational Science Software

- There is a *crisis* in computational science software
  – many years of inadequate investments
  – lack of useful tools
  – dearth of widely accepted standards and best practices
  – paucity of third party software vendors
  – simple lack of perseverance by the community

- Improvement in computational science software is needed urgently along multiple dimensions

Low Level Software Tools

- Message Passing Interface (MPI)
  - "assembly language" of parallel computing
  - lowest common denominator
    - portable across architectures and systems
- High-Performance Fortran (HPF)
  - higher level data parallel specification
    - limited to regular data structures
  - we expected too much too soon
    - see Earth System Simulator
- Costs and implications
  - human productivity
    - low-level programming model
  - software innovation
    - limited development of alternatives

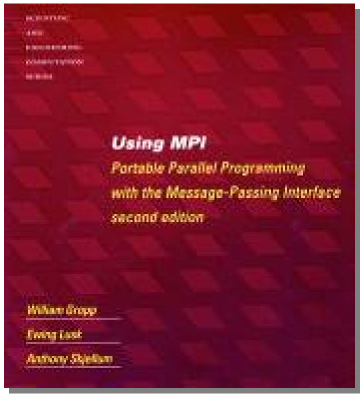

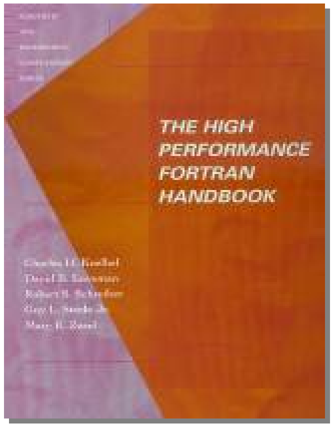

Architecture and Hardware

- COTS products are useful and cost-efficient for some applications
- However, some important complex problems can only be addressed through "purpose built" computing systems
- Demand for high-end systems does not sustain a market for such products
- Federal government must take primary responsibility for accelerating advances in computer architectures and ensuring that there are multiple strong domestic suppliers

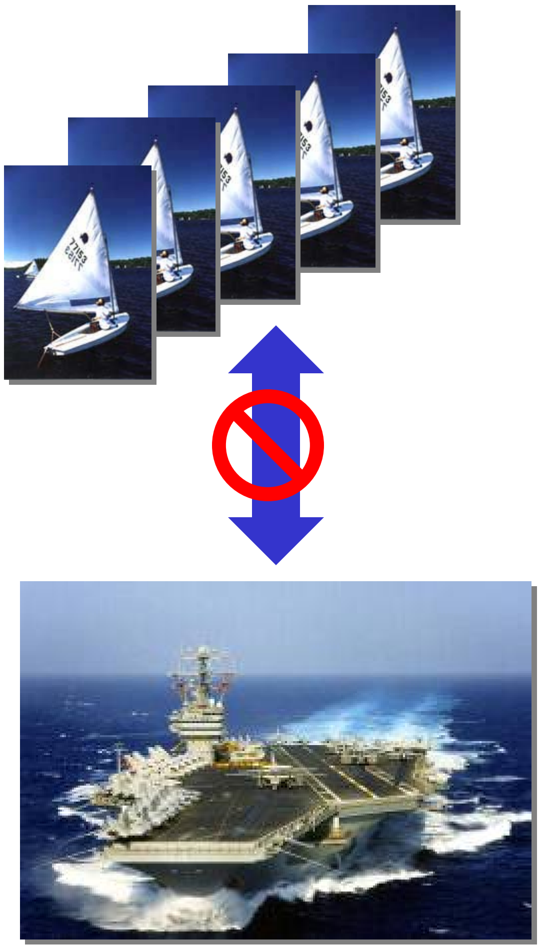

GUPS: Mismatched Needs

- GUPS benchmark
  - Giga-updates per second
  - generalized memory references
- Different optimization points
  - commercial workloads
    - small memory footprints
  - technical workloads
    - more irregular memory patterns
- Source
  - Defense Science Board, *Task Force on DoD Supercomputing Needs*, 2000

Architecture (Year)	GUPS (4 GB Memory)
Cray Y-MP (1988)	0.16
Cray C90 (1991)	0.96
Cray T90 (1995)	3.2
Cray SV1 (1999)	0.7
Cray T3E (1996)	2.2
SMPs (2000)	0.35-1.0
Clusters (2000)	0.35-1.0

Algorithms and Applications

- Multidisciplinary teams of specialists are needed, each with complementary expertise and an appreciation of the interdisciplinary aspects of the system, and each supported by a software infrastructure that can leverage specific expertise from multiple domains and integrate the results into a complete application software system
- We must continue to develop and improve the mathematical, non-numeric, and computer science algorithms that are essential to the success of future computational science applications

Data Management and Sensors

- Computational science, based on ubiquitous sensors and high-resolution detectors, is an emerging opportunity to couple observation-driven computation and analysis, particularly in response to transient phenomena.
- Explosive growth in the resolution of sensors and scientific instruments is creating unprecedented volumes of experimental data.
- We must increase investment and focus on sensor- and data-intensive computational science in recognition of the explosive growth of experimental data, itself a consequence of increased computing capability.

Kindling the Fire

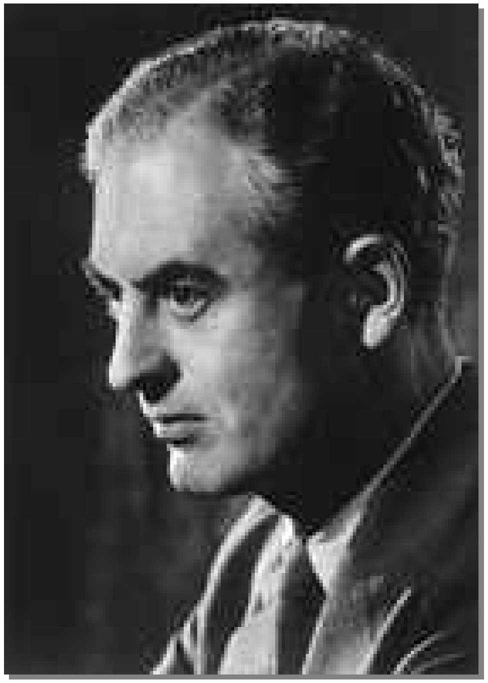

I am often asked, “What made you become scientist?” But I can't stand far enough away from myself to give a really satisfactory answer, for I cannot distinctly remember a time when I did not think that a scientist was the most exciting possible thing to be.

Sir Peter Medawar