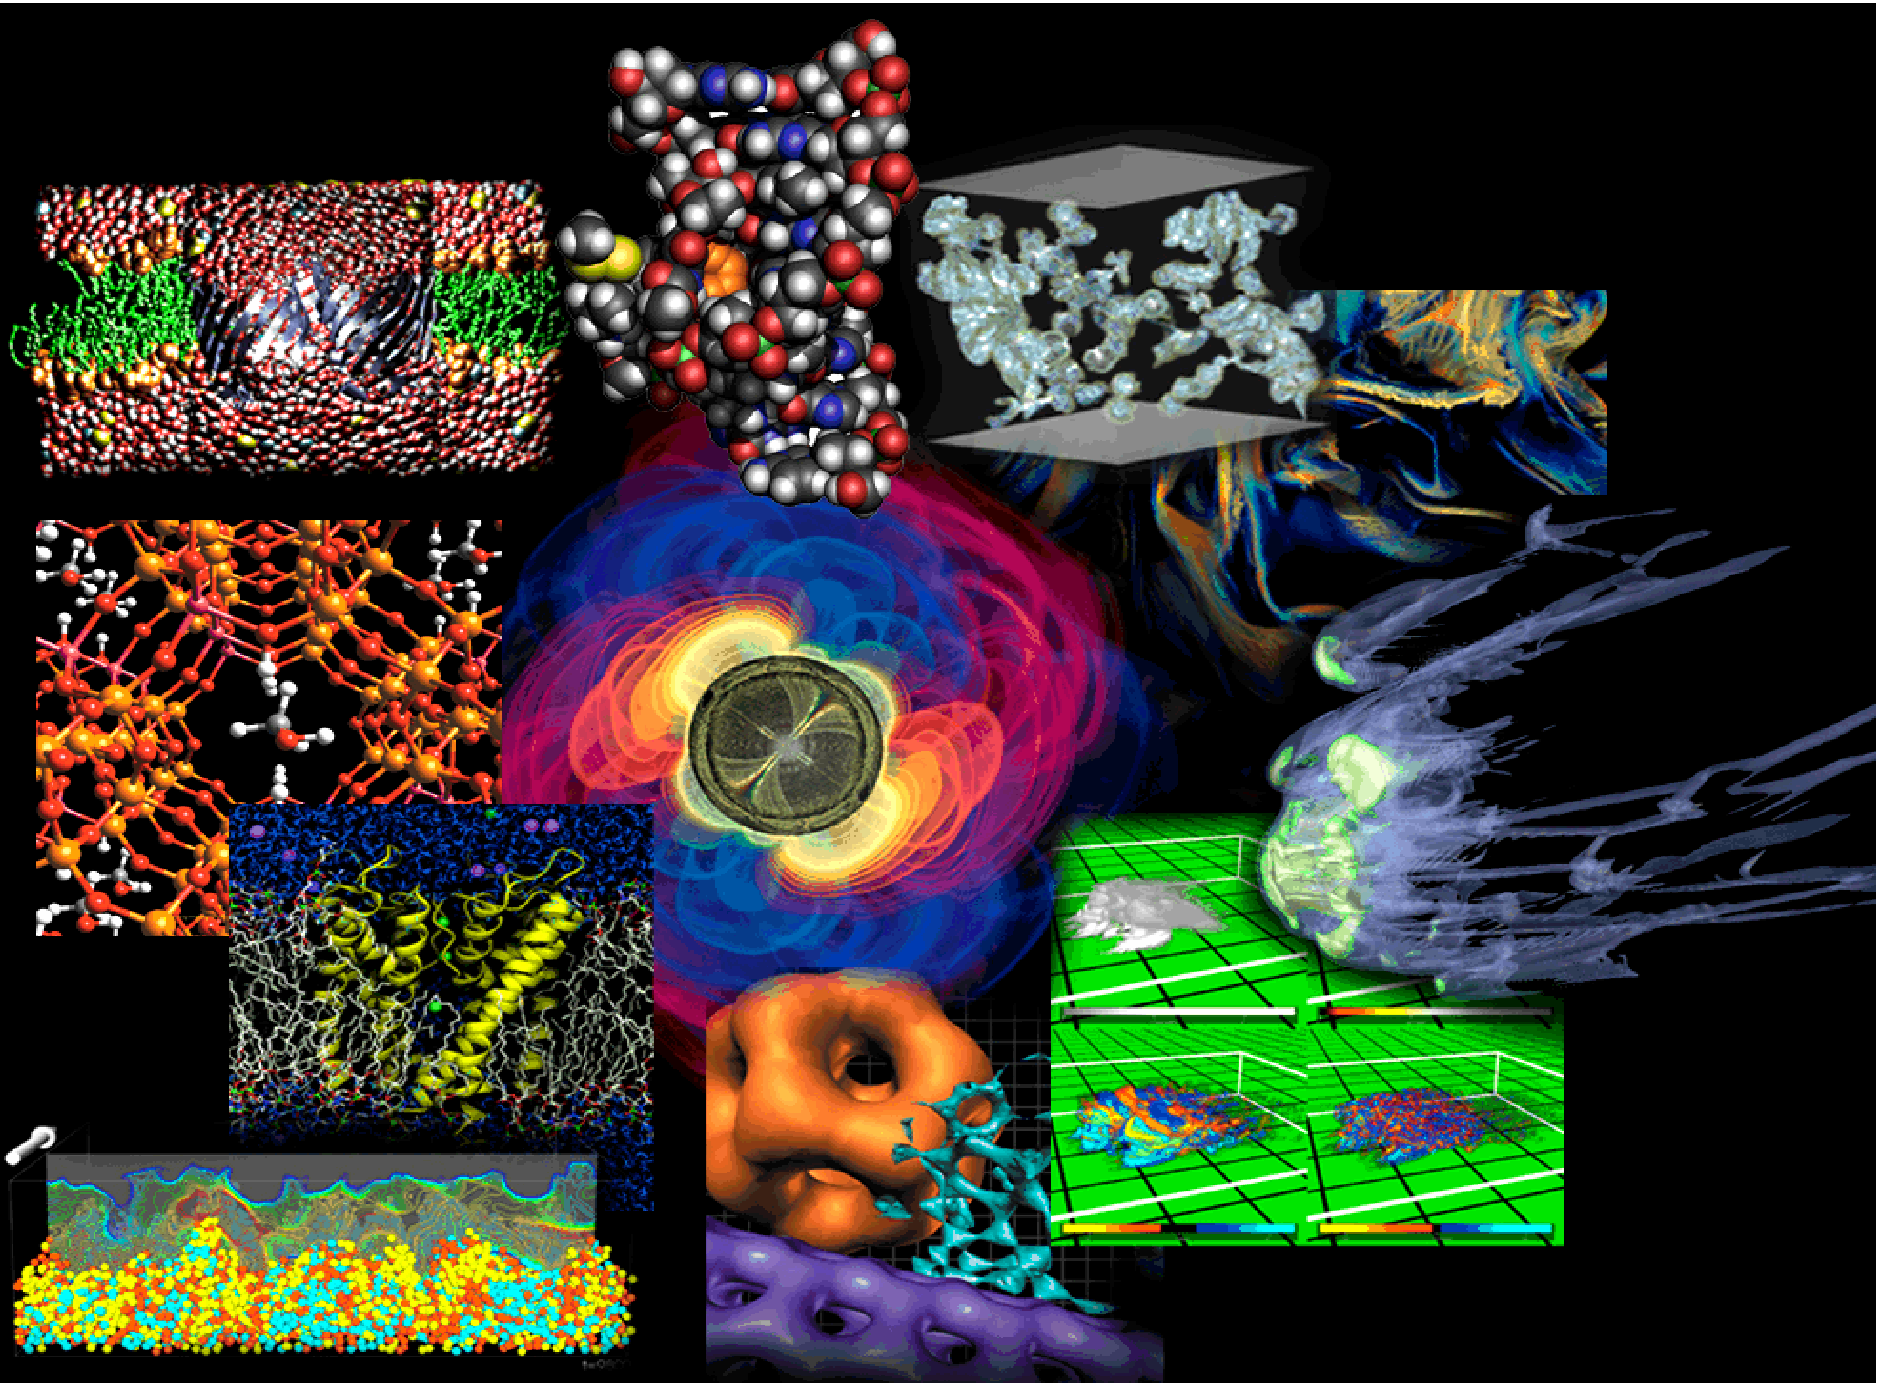

Source: NCSA

Computational Science Subcommittee: Preliminary Observations

President's Information Technology Advisory Committee

Subcommittee on Computational Science

Daniel A. Reed, Chair

Dan_Reed@unc.edu

PITAC Meeting, Arlington, Virginia

November 4, 2004

About the President's Information Technology Advisory Committee (PITAC) (1/2)

- PITAC members are appointed by the President to provide independent expert advice on maintaining America's preeminence in advanced information technology.
- The Committee's studies help guide the Administration's efforts to accelerate the development and adoption of information technologies vital for American prosperity in the 21st century.
  - Recent release: *Revolutionizing Health Care Through Information Technology.*
- PITAC members are information technology leaders in industry and academia with expertise relevant to critical elements of information technology research and development.

About PITAC (2/2)

- Chartered by Congress under the High-Performance Computing Act of 1991 (P. L. 102-194) and the Next Generation Internet Act of 1998 (P. L. 105-305). It is formally renewed through Presidential Executive Orders.
- Federally chartered advisory committee operating under the Federal Advisory Committee Act (FACA) (Public Law 92-463) and other Federal laws governing such activities.
- Reports to the President through the Office of Science and Technology Policy, Executive Office of the President.
- Supported by the National Coordination Office for Information Technology Research and Development.

Current PITAC Membership

Subcommittee on Computational Science

- ***Daniel A. Reed***, Ph.D., ***Chair***, Chancellor's Eminent Professor, Vice-Chancellor for Information Technology and CIO, 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 with 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 Chief Information Officer, Dell Computer

Subcommittee's Charge (1/3)

June 9, 2004 letter from Dr. John H. Marburger, III, Science Adviser to the President

1. How well is the Federal Government targeting the right research areas to support and enhance the value of computational science? Are agencies' current priorities appropriate?
2. How well is current Federal funding for computational science appropriately balanced between short term, low risk research and longer term, higher risk research? Within these research arenas, which areas have the greatest promise of contributing to breakthroughs in scientific research and inquiry?

Subcommittee's Charge (2/3)

3. How well is current Federal funding balanced between fundamental advances in the underlying techniques of computational science versus the application of computational science to scientific and engineering domains? Which areas have the greatest promise of contributing to breakthroughs in scientific research and inquiry?
4. How well are computational science training and research integrated with the scientific disciplines that are heavily dependent upon them to enhance scientific discovery? How should the integration of research and training among computer science, mathematical science, and the biological and physical sciences best be achieved to assure the effective use of computational science methods and tools?

Subcommittee's Charge (3/3)

5. How effectively do Federal agencies coordinate their support for computational science and its applications in order to maintain a balanced and comprehensive research and training portfolio?
6. How well have Federal investments in computational science kept up with changes in the underlying computing environments and the ways in which research is conducted? Examples of these changes might include changes in computer architecture, the advent of distributed computing, the linking of data with simulation, and remote access to experimental facilities.
7. What barriers hinder realizing the highest potential of computational science and how might these be eliminated or mitigated?

Subcommittee Work Plan

- June 9: Charge from the White House
- June 17: PITAC meeting (Arlington, Virginia)
- September 16: Information gathering meeting (Chicago)
- October 19: Information gathering meeting (Arlington, Virginia)
- *November 4: PITAC meeting (Arlington, Virginia)*
- November 10: Birds of a Feather (BOF) at Supercomputing 2004
- November-December: Report drafting and input solicitation
- January 2005: PITAC meeting - major review of the draft report
- February - March 2005: Editing
- March 2005: Review and approval of final draft report

Outline of Disciplines Explored

- Overarching and cross-cutting issues
- Disciplinary studies
  - biomedical and biological sciences
  - engineering
  - climate, weather, and environmental science
  - more to come
- Technology and human resource assessment
  - software and algorithms
  - architecture and infrastructure

Biomedical and Biological Sciences

- NIH roadmap observations
  - Computation is integral and critical to biomedical research.
  - Computation's deficiencies and insufficient personnel and education limit progress.
- Clinical decision making requires computational science (including access to long-term archives of clinical records) to select tests, determine therapies, choose dosages and schedules, and identify combination. Examples include surgical and radiation treatment planning, which also requires computation- and data-intensive imaging.

Biomedical Imaging Challenges

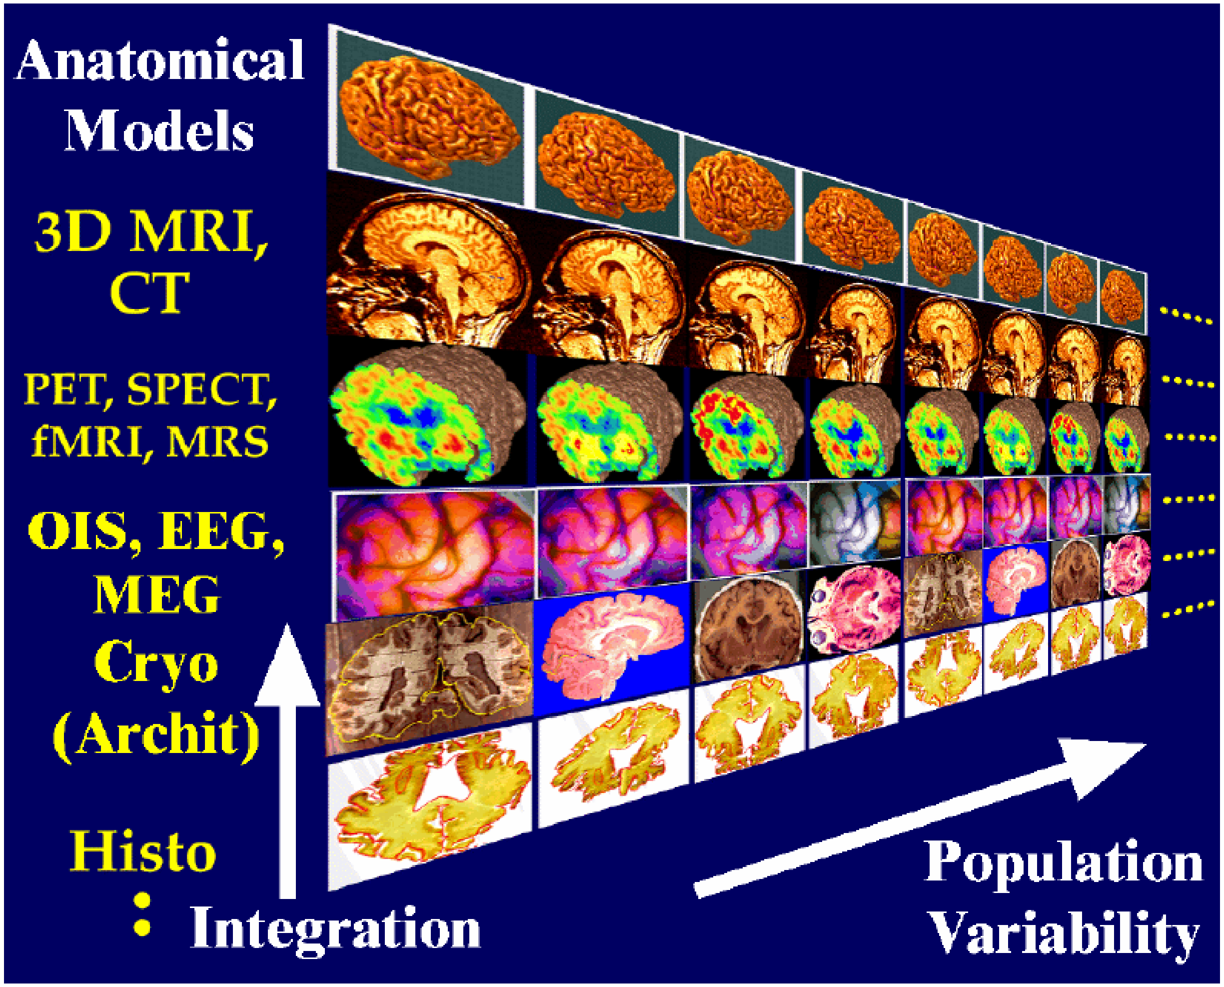

Source: Chris Johnson, Utah and Art Toga, UCLA

Biology/Bioengineering Grand Challenges

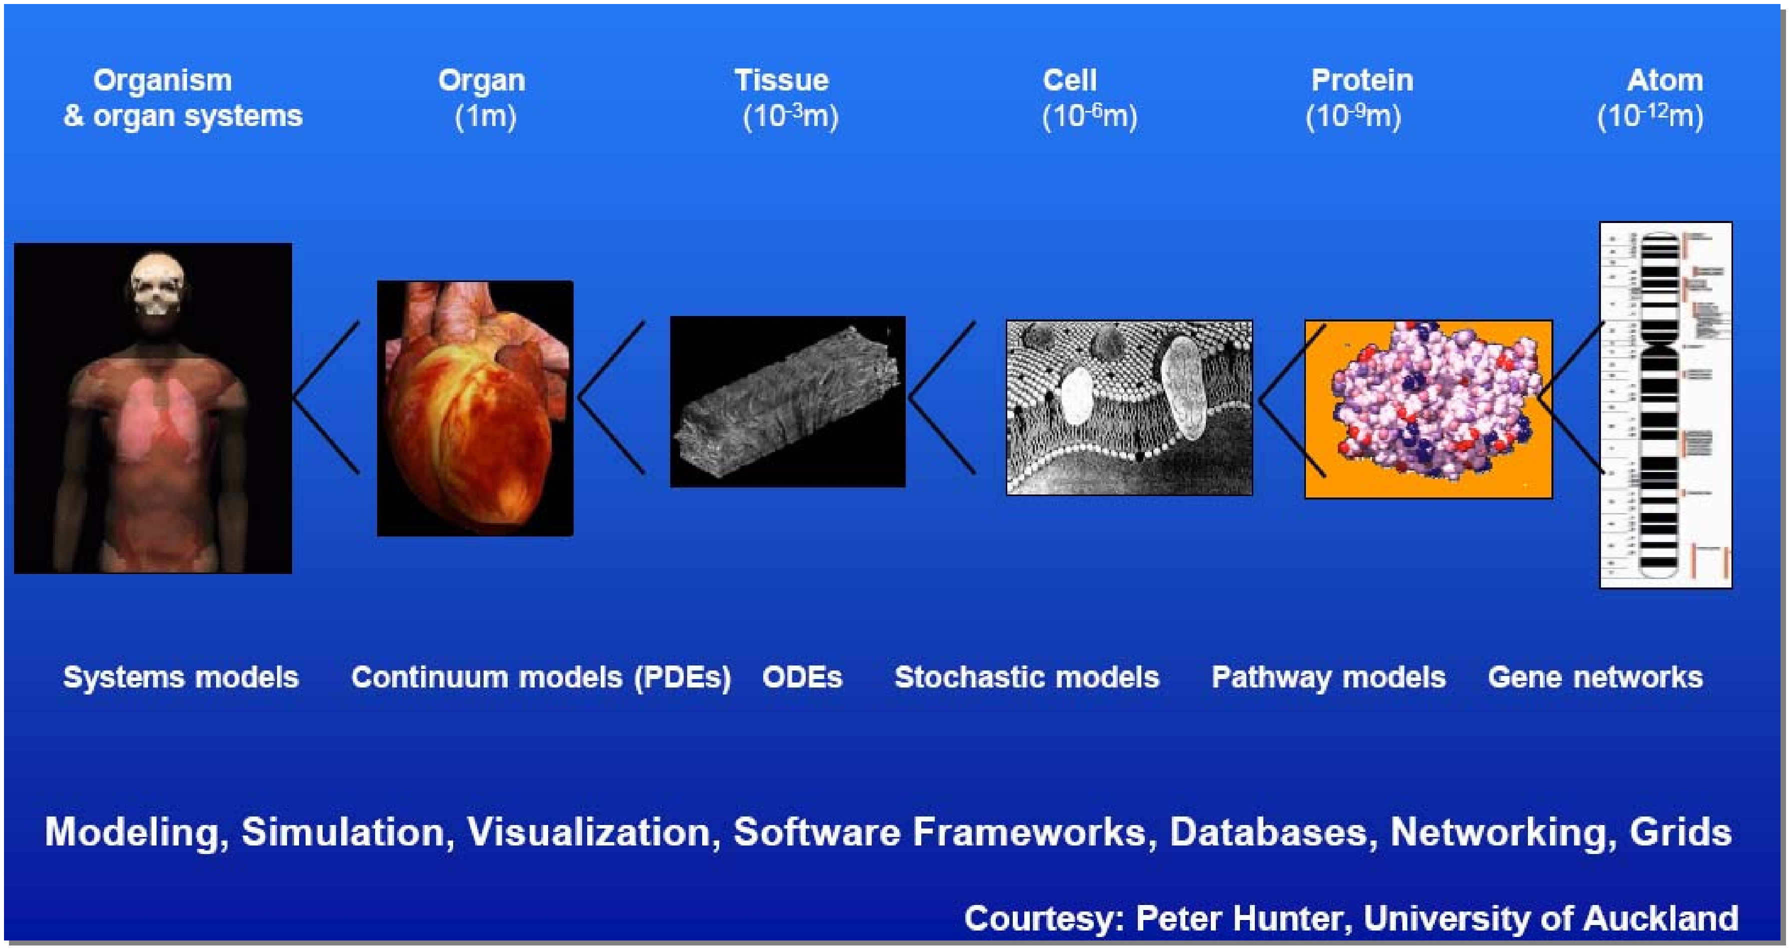

Source: Chris Johnson, Utah

Multidisciplinary Optimization

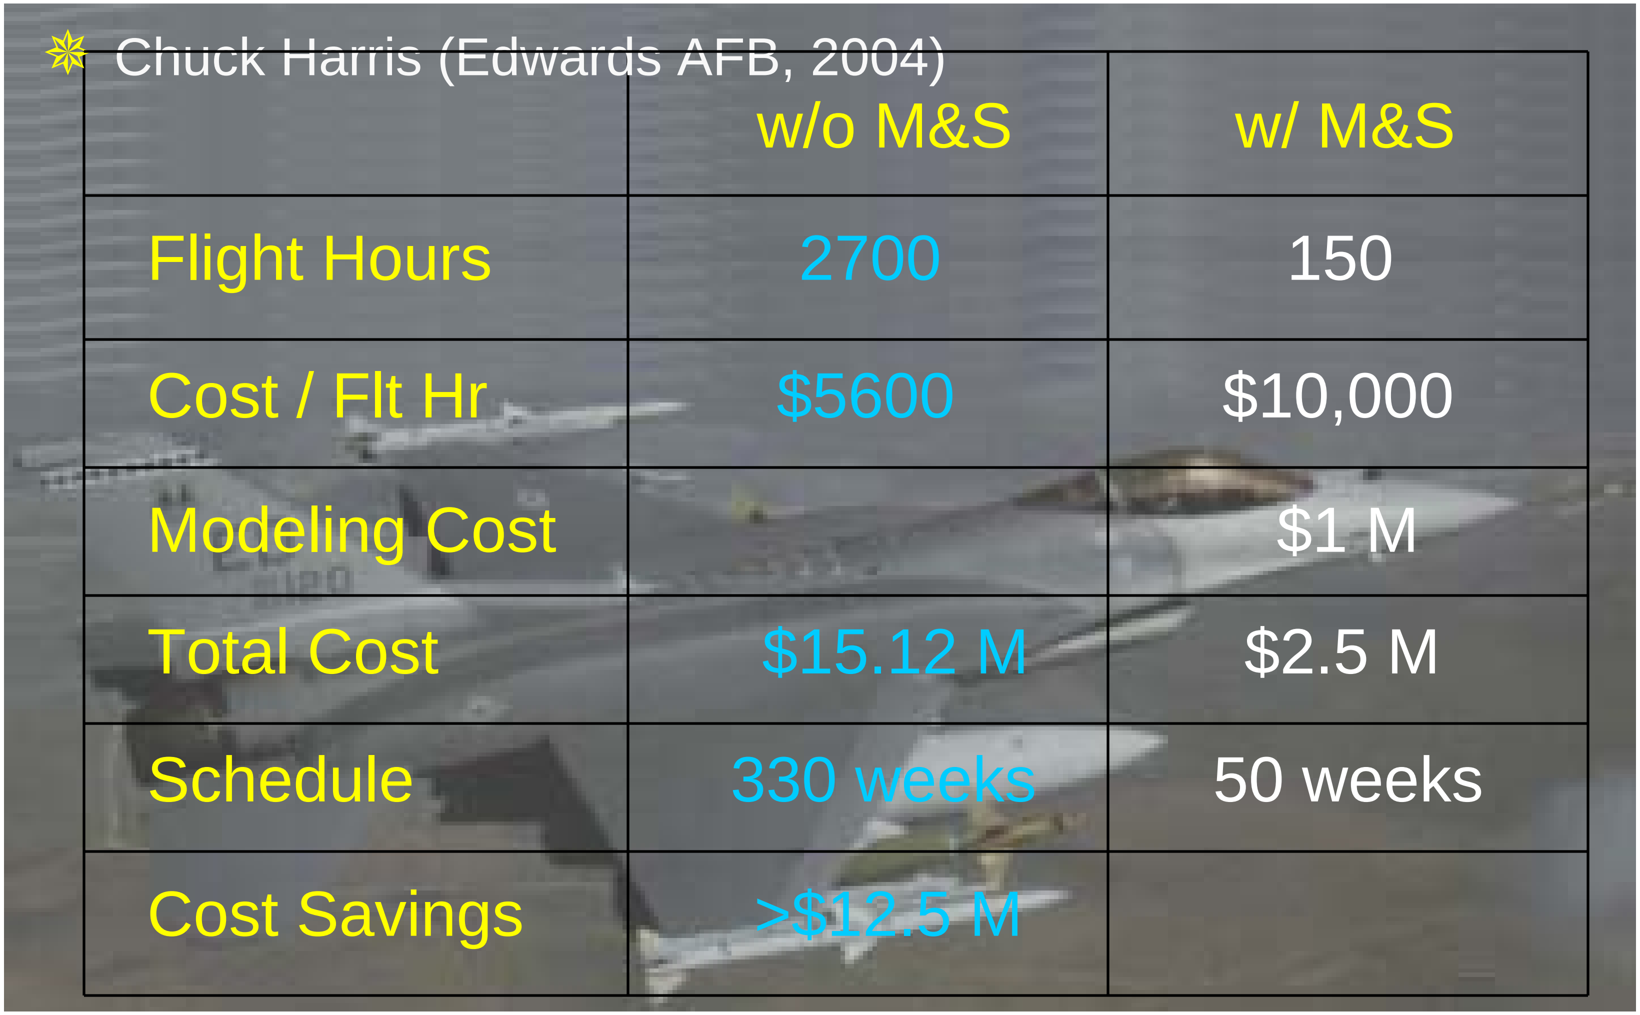

- Chuck Harris (Edwards AFB, 2004)

	w/o M&S	w/ M&S
Flight Hours	2700	150
Cost / Flt Hr	$5600	$10,000
Modeling Cost		$1 M
Total Cost	$15.12 M	$2.5 M
Schedule	330 weeks	50 weeks
Cost Savings	>$12.5 M	

Source: Charbel Farhat, Stanford

Supercomputing Support of Columbia Accident Investigation Board (CAIB)

Ames marshalled supercomputing-related R&D assets and tools to produce time-critical analysis for CAIB

R&D and Operations teams joined to form integrated, interdisciplinary team to support CAIB in evaluation of launch and re-entry events

R&D Team: Grand Challenge Driven	Engineering Team: Operations Driven
Ames Research Center Glenn Research Center Langley Research Center	Johnson Space Center Marshall Space Flight Center Industry Partners

Bringing the investigation full circle, the team utilized advanced data analysis and visualization technologies to explore and understand the vast amount of data produced by this massive simulation effort

The team used state-of-the-art codes to meet the challenge of modeling a free object in the flow around an airframe, with automatic re-gridding at each time step

Grand Challenges

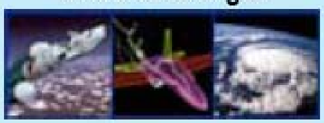

Next Generation Codes & Algorithms

OVERFLOW
Honorable Mention, NASA Software of Year
STS107

INS3D
NASA Software of Year
Turbopump Analysis

CART3D
NASA Software of Year
STS-107

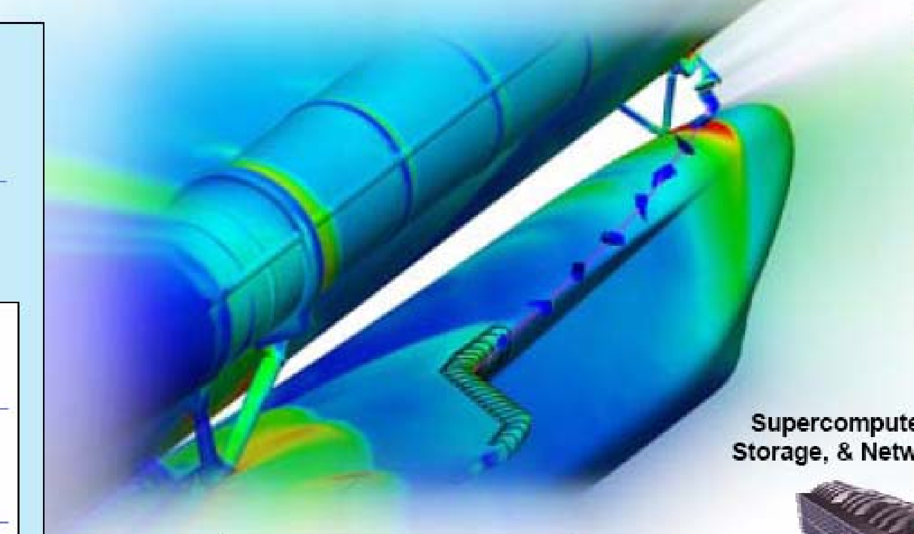

Computation Management
- AeroDB
- ILab

These codes have been honed by R&D experts in large-scale simulation, with their environment of programmer tools developed to minimize the effort of executing efficiently

Modeling Environment (experts and tools)
- Compilers
- Scaling and Porting
- Parallelization Tools

Supercomputers, Storage, & Networks

Rapid, high-fidelity computational investigations may involve hundreds or thousands of large simulations and results datasets, but automated computation and data management tools are at the ready

The codes and tools were tuned to exploit advanced hardware (supercomputers, networks, and storage) that was developed for just such large, tightly-coupled computations, enabling rapid turnaround

Source, Walt Brooks, NASA

Weather and Economic Loss

•40% of the $10T U. S. economy is impacted by weather and climate
•$1M in economic loss to evacuate each 1 mile of coastline

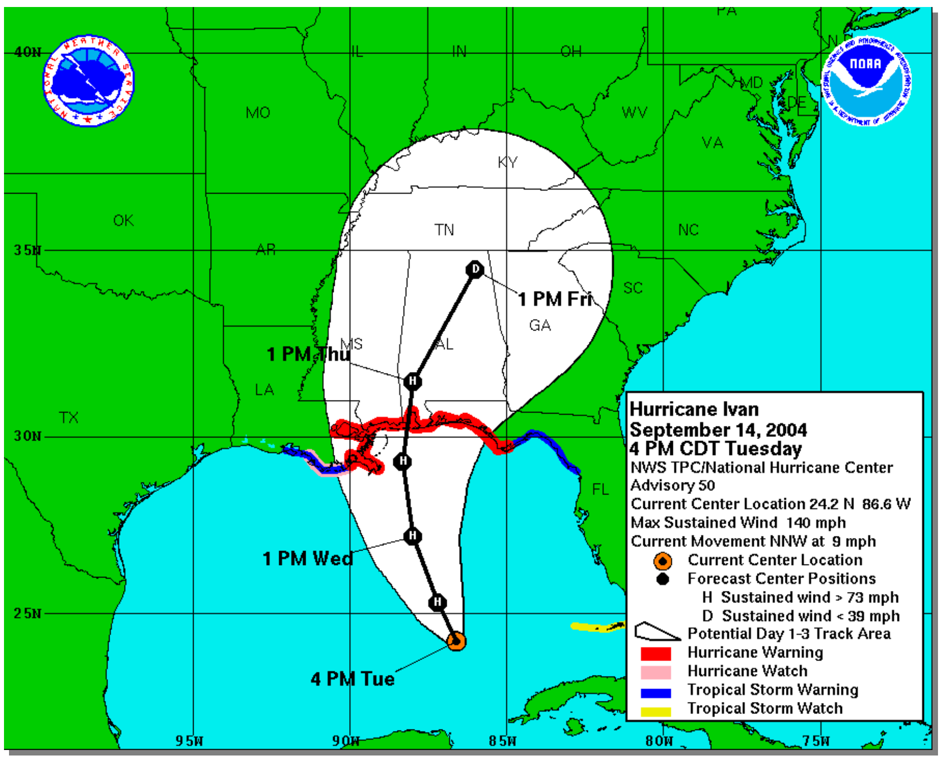

- We now over-warn by a factor of 3
- Average over-warning is 200 miles, or $200M per event
- Improved forecasts
  - saving lives and resources

Source: Kelvin Droegemeier, Oklahoma

Multidisciplinary Models

The Development of Climate models, Past, Present and Future

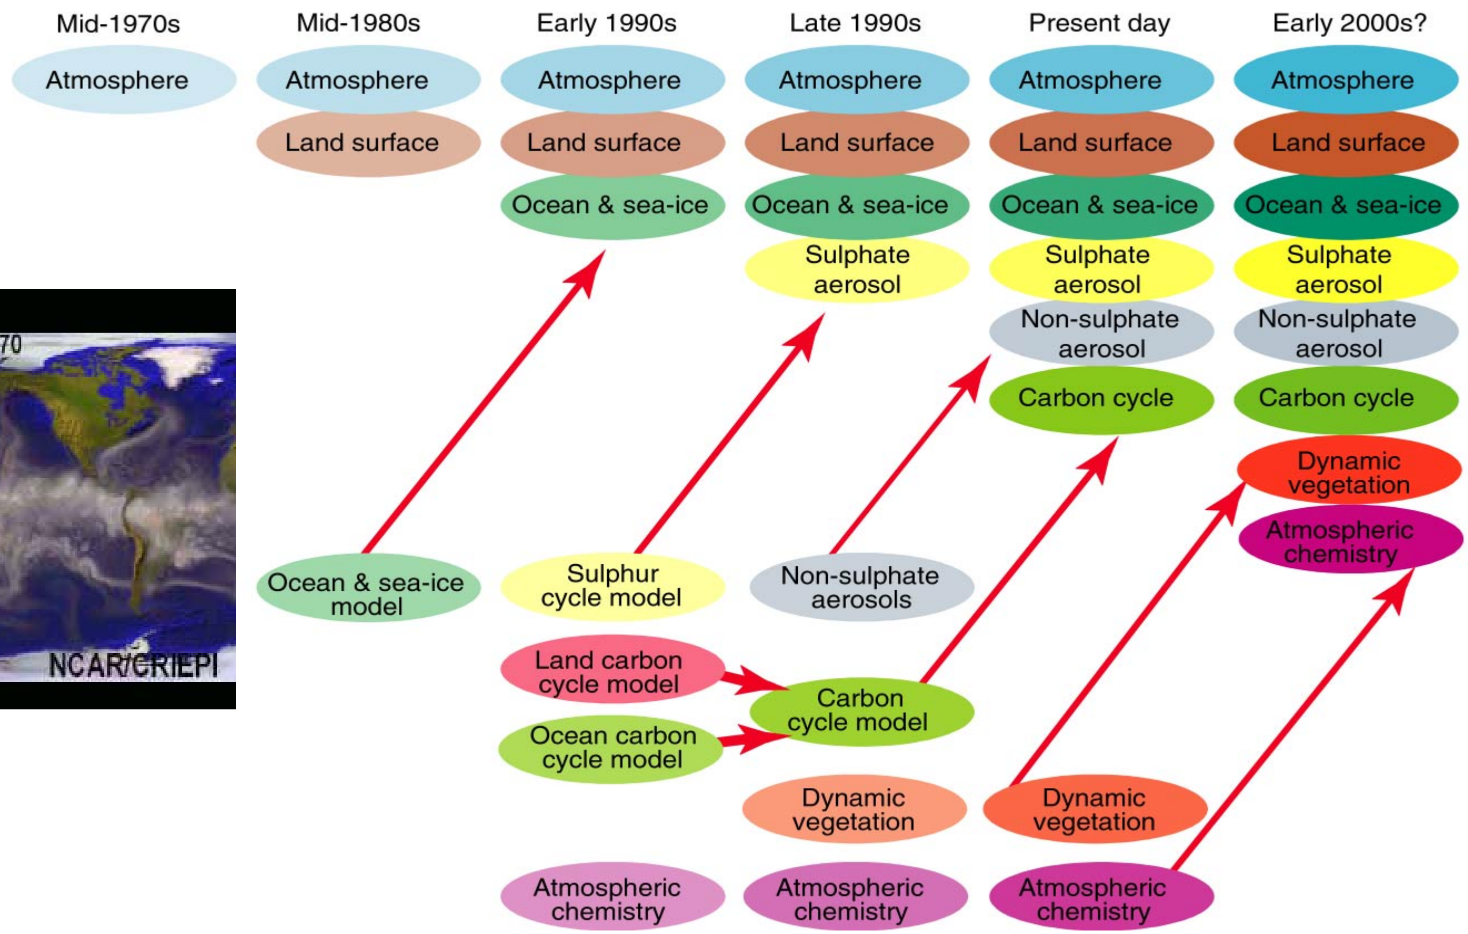

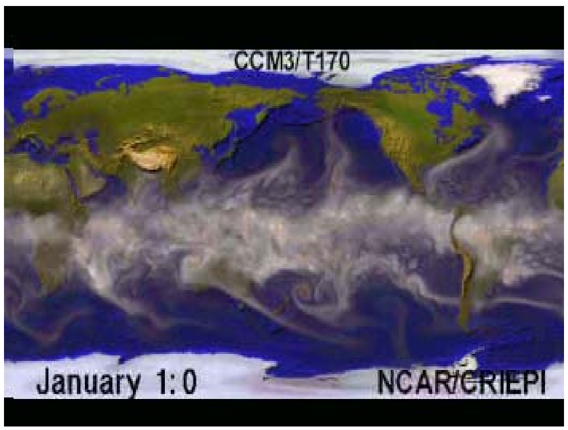

Source: Tim Killeen, NCAR

ESMF

The Earth System Modeling Framework: A High-Performance Framework for Earth Science Modeling and Data Assimilation

ESMF Architecture

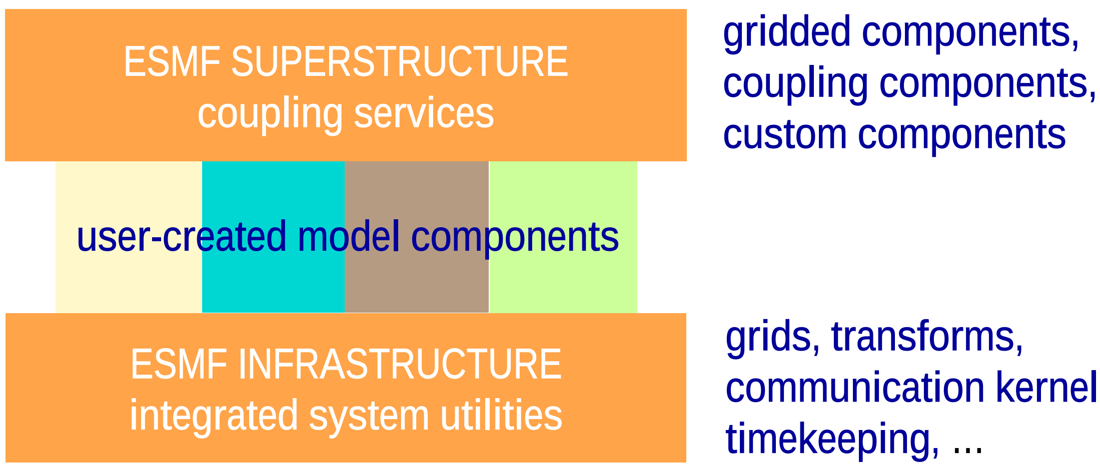

Source: Tim Killeen, NCAR

Earth System Modeling Framework Rollout

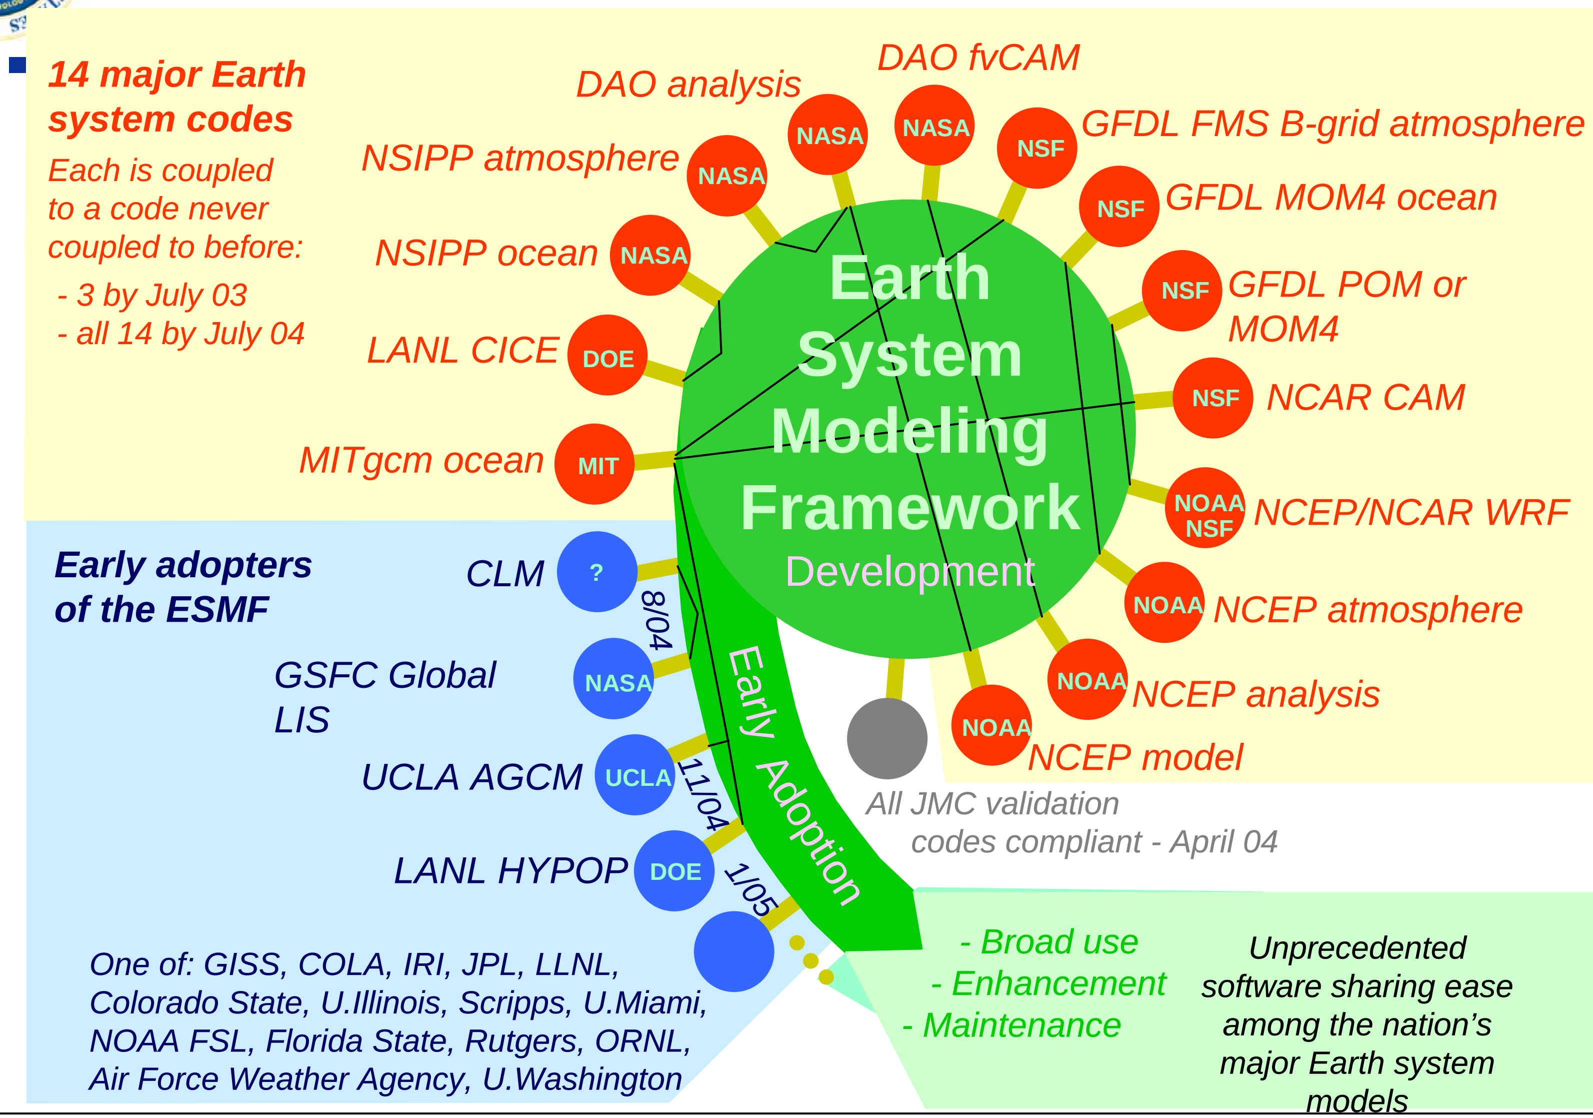

Source, Jim Fischer, NASA

Climate, Weather, and Environment

- More and better observations
- Finer scales in models and simulations
- Better data assimilation
- Improved model physics
- Ensemble forecasting
  - uncertainty analysis
- Adaptive, on-demand forecasting

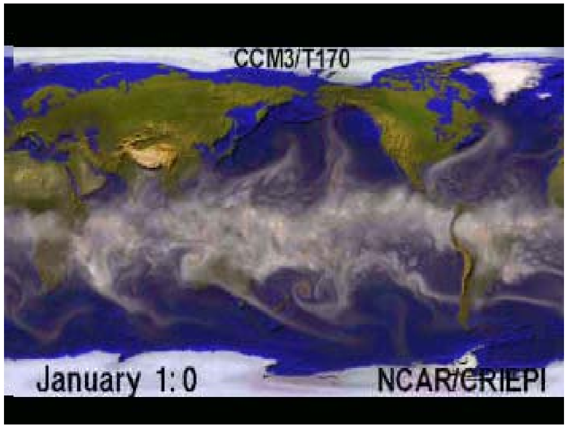

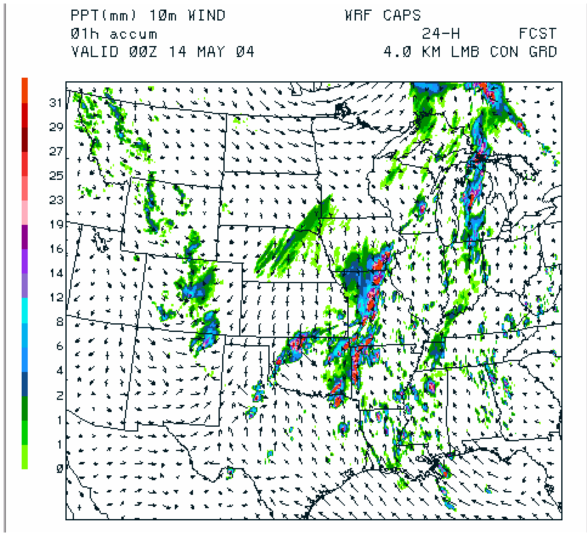

Outline of Preliminary Observations

- Importance of computational science
  - essential to scientific discovery
- Current state of computational science
- Paths to progress
  - leadership, education, and people
  - software, algorithms, and applications
  - architecture, hardware, and infrastructure
- Computational science opportunities
  - national security
  - economic return on investment
  - scientific discovery

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.

Preliminary Observations

Computational Science is 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.

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.

Preliminary Observations

Opportunities:

- Broadly speaking, computational science opportunities and benefits exist at three levels:
  - increasing national security
  - providing return on investment relative to other opportunities
  - catalyzing scientific discovery and novel ideas

Community Messages

The following slides represent information presented during the Subcommittee's information gathering meetings.

1. How well is the Federal Government targeting the right research areas to support and enhance the value of computational science? Are agencies' current priorities appropriate?

- Well designed programs do advance computational science
  - the multi-agency Grand Challenge program
  - NSF Information Technology Research (ITR)
  - DOE/SC Scientific Discovery through Advanced Computing (SciDAC)
- Past investment (1980s-early 1990s)
  - fostered architectural innovation with new system designs and products
    - today, clustered SMP systems are the least common denominator
  - fostered exploration of alternate, high-level programming models
    - today, MPI, the parallel equivalent of assembly language, is the standard
- There are opportunities for collaboration
  - across the Federal Government, academia, and industry
  - targeting common R&D simulation and data management challenges
- A computational science roadmap is needed
  - to set government-wide priorities and to guide decision making

2. How well is current Federal funding for computational science appropriately balanced between short term, low risk research and longer term, higher risk research? Within these research arenas, which areas have the greatest promise of contributing to breakthroughs in …

- Current R&D programs are too risk averse
  - there is very little long-term, high-risk research being conducted
  - funding agencies have a low tolerance for failure
  - this forces researchers into short-term, low-risk projects

- There are breakthrough opportunities in many areas, e.g.,
  - models and systems that mix simulation, data, and sensors
    - climate, weather, air traffic control, …
    - oil exploration, radio astronomy, …
  - biomedical systems

3. How well is current Federal funding balanced between fundamental advances in the underlying techniques of computational science versus the application of computational science to scientific and engineering domains? Which areas have the greatest promise of contributing …

- Over the last decade, the Federal investment strategy has been excessively biased in favor of hardware over software.
- Applications and algorithm research need more emphasis.
- The shift to distributed memory systems had many effects
  - major changes in software and algorithms were required
  - many older applications were lost
  - a new legacy of software was created by this decade of investment
- Computer scientists were marginalized to focus on incremental issues
  - fine tuning 20-year old programming abstraction and related tools
  - limited work even on evolutionary programming models (e.g., UPC)
- Most funding appears coupled to large scale application endeavors
  - the application goals set the timetable and scope of algorithmic development
  - this limits risk taking and innovation in computational science
- One of the greatest promises in life sciences
  - algorithms that navigate, filter, and mine large stores of heterogeneous data

4. How well are computational science training and research integrated with the scientific disciplines that are heavily dependent upon them to enhance scientific discovery? How should the integration of research and training among computer science, mathematical science, and the …

- Academia has been slow to respond to changing needs
  - courses and research programs are not well integrated
  - student computational science skills remain poor
  - skills acquired via research, not education, which inherently is limiting
- The key to an integrated approach is a peer environment in which the computational scientist, computer scientist, and biologist serve as colleagues with equal responsibilities. This theme applies equally to the educational, government, and industrial environment. This could involve cross-disciplinary project teams, matrix management, or new institutions.
- Better project management is required
  - stronger, knowledgeable PIs with fewer projects!!
  - clearer deliverables and milestones
  - explanations of technology use and integration

5. How effectively do Federal agencies coordinate their support for computational science and its applications in order to maintain a balanced and comprehensive research and training portfolio? (1/2)

- This is a major problem, despite many inter-agency initiatives.
- Each agency appears to fund its own computational science program, aligning it with the agency's mission or research objectives.
- There are institutions with multiple funding streams that leverage capabilities and cross agency efforts to share ideas, but there is no cross agency approach to balancing efforts.
- Coordination mechanisms have limited effectiveness
  - there are few linked agency requests for proposals (RFPs)
  - resource sharing is rare
  - communities are split, unaware of each other's work and/or tools

5. How effectively do Federal agencies coordinate their support for computational science and its applications in order to maintain a balanced and comprehensive research and training portfolio? (2/2)

- What incentive is provided for agencies to collaborate?
- There are exemplars of success
  - HPCS, with multiple agencies making coordinated investments
  - system software collaboration by DARPA, DOE/SC, and NSF
- Success require common vision and commitment
  - they occur today because of a handful of committed people
  - they also require long-term commitment to specific research agendas
- The process for success should be institutionalized and suggests need for technology "roadmap"

6. How well have Federal investments in computational science kept up with changes in the underlying computing environments and the ways in which research is conducted? Examples of these changes might … (1/2)

- Technological change in computing outpaces the ability of researchers to use technology effectively
- Computer scientists are developing new tools and methods while application and domain scientists are still learning to use those from the previous three generations!
- We have no shared national vision for the provision of stable cyberinfrastructure – yet we build and support highways just fine!
- We have many challenges – fault tolerance, performance efficiency and scaling, linking data with simulation, Grid use, networking tools, …

6. How well have Federal investments in computational science kept up with changes in the underlying computing environments and the ways in which research is conducted? Examples of these changes might … (2/2)

- There is a large focus on platforms and applications, with software infrastructure for capability computing lagging.
- There is limited focus on capability vs. capacity computing, for technical and political reasons.
- Infrastructure costs are significant and not typically considered in funding national programs.
- A long-term, stable commitment to the solution of scientific and engineering problems is needed.
- Curiosity-driven research is not antithetical to long-term government leadership.

7. What barriers hinder realizing the highest potential of computational science and how might these be eliminated or mitigated?

- There are many social barriers, including the lack of
  - coordinated national leadership and project management
  - a critical mass of peers in academia, government, & industry
  - focus on capability computing and prioritizing key grand challenges
    - adequate resources to have an impact
  - multidisciplinary education and training
- There are also technical barriers
  - the programming complexity of HPC platforms
  - inadequate or mismatched algorithms
  - latency and bandwidth (physics constraints)
  - looming threats to Moore’s law and continued scaling

Report Schedule

- November 10
  - Birds of a Feather (BOF) at Supercomputing 2004
- November-December
  - report drafting and input solicitation
- January 2005
  - PITAC meeting with major review of draft report
- February - March 2005
  - editing
- March 2005
  - approval of final draft report at PITAC meeting

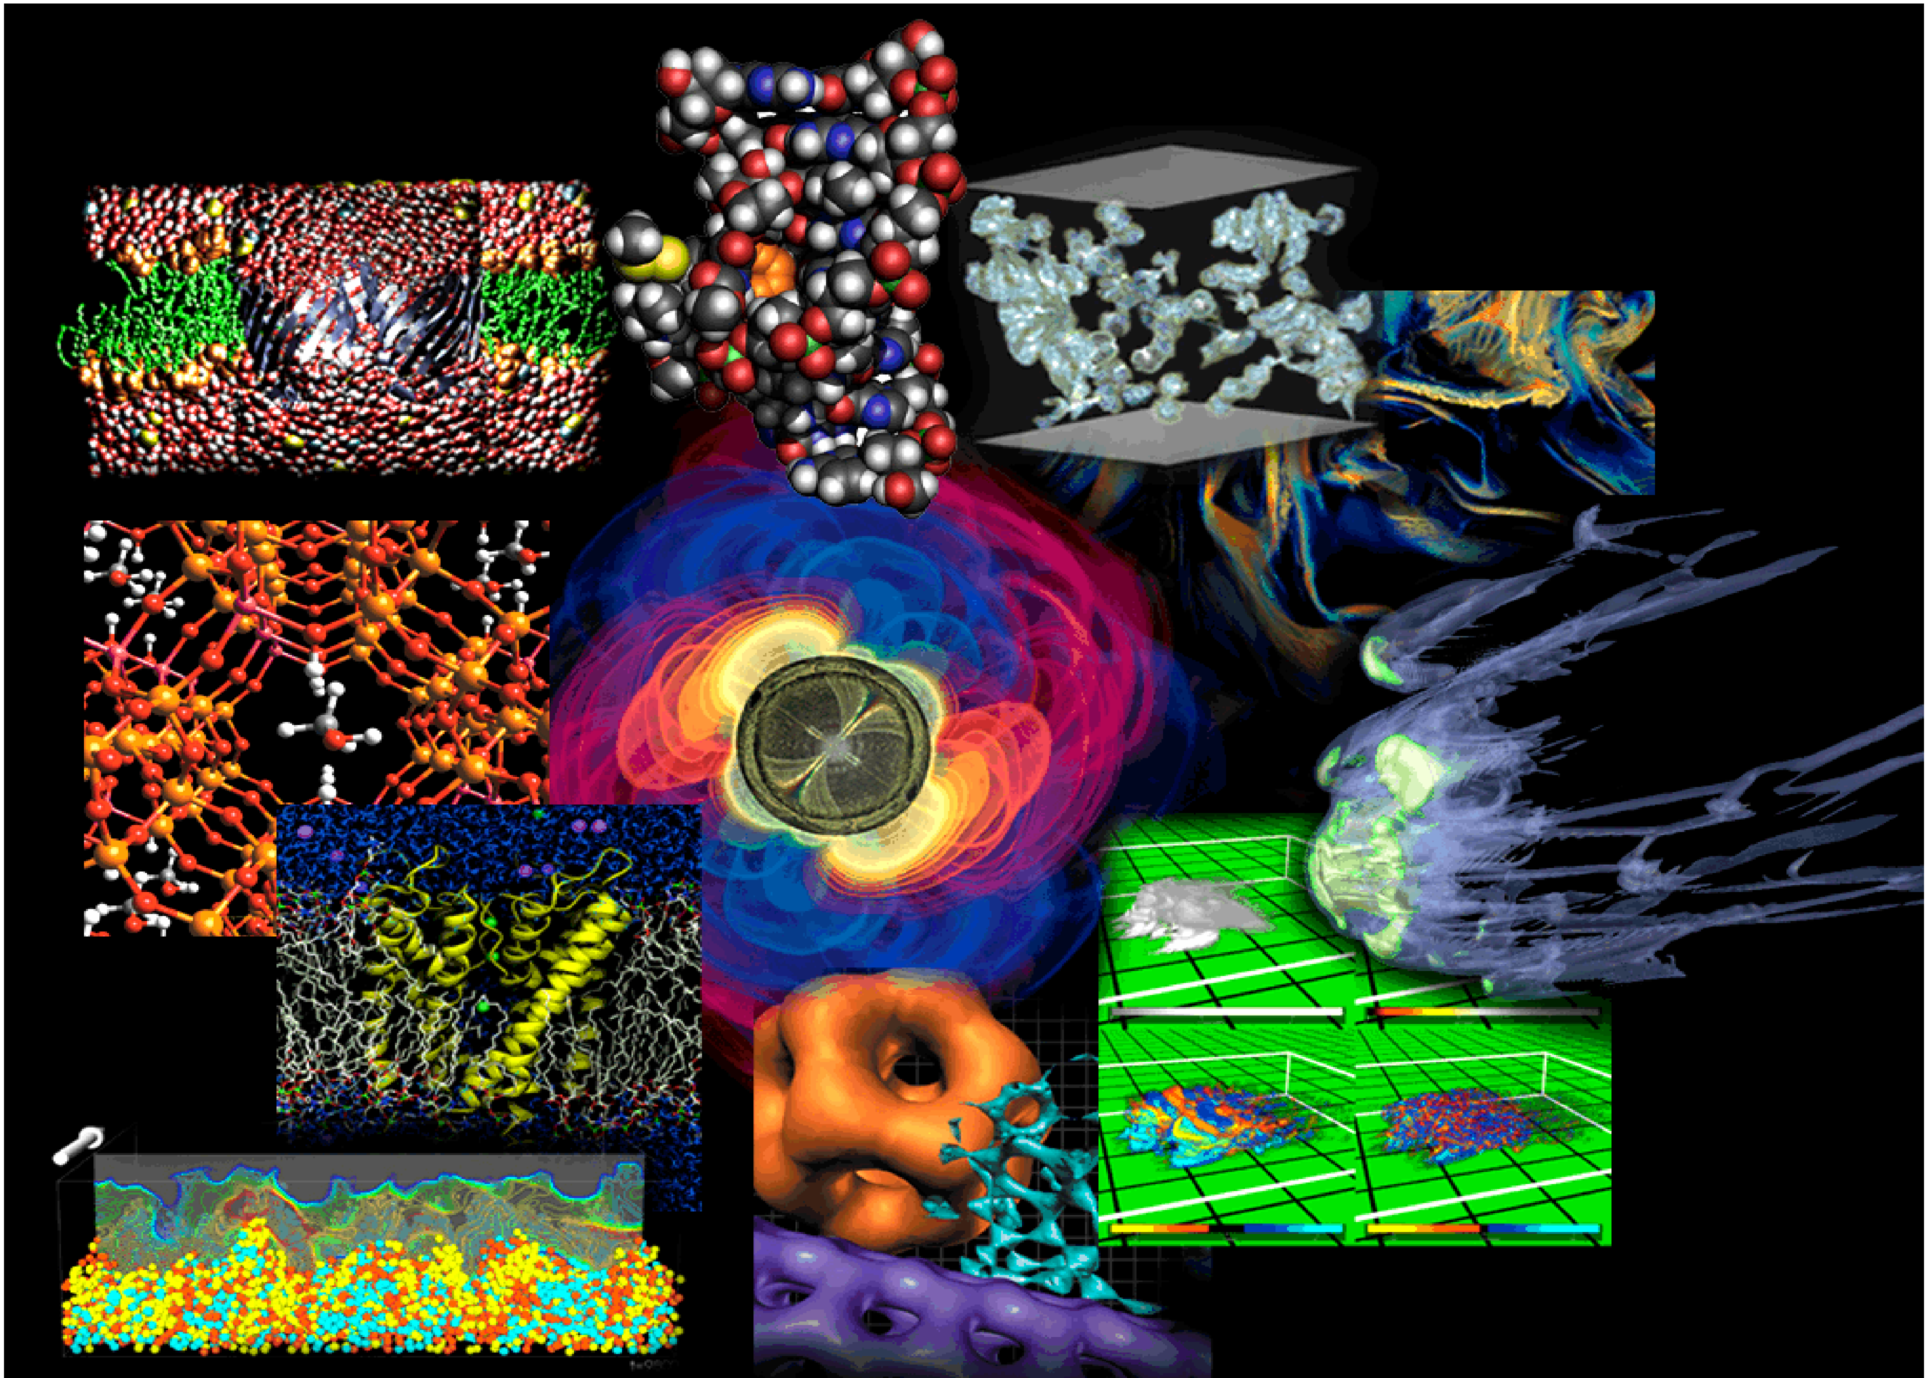

Source: NCSA