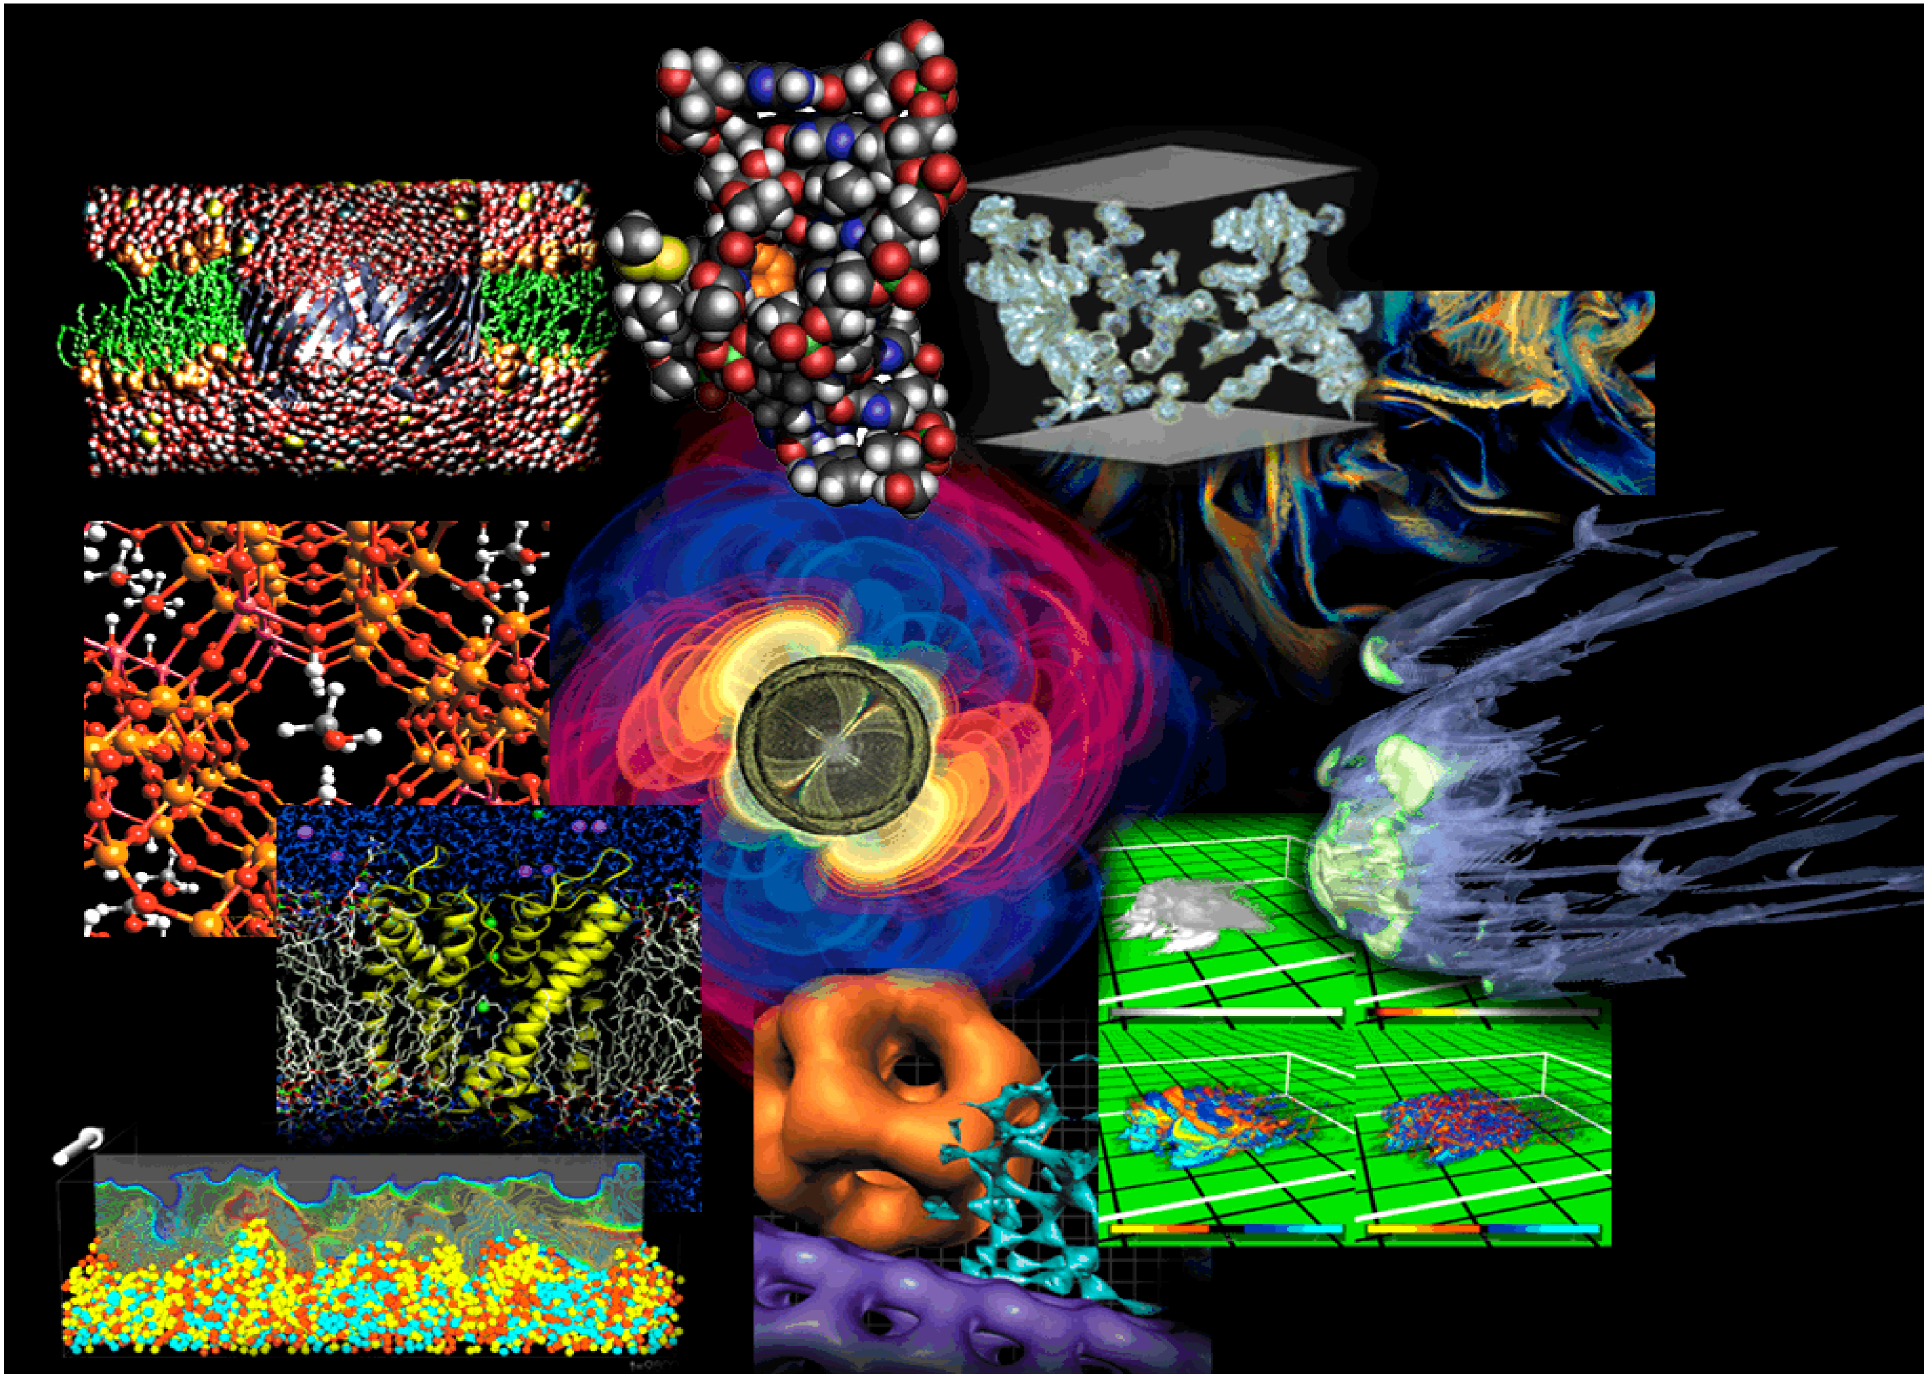

Source: NCSA

Computational Science Subcommittee: A Status Report

President's Information Technology Advisory Committee

Subcommittee on Computational Science

Daniel A. Reed, Chair

Dan_Reed@unc.edu

PITAC Meeting, Arlington, Virginia

January 12, 2005

Computational Science Subcommittee

- ***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, 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 Consultants

- ***Jack Dongarra, Ph.D.***
  - University Distinguished Professor, University of Tennessee
  - Director, Innovative Computing Laboratory
  - Distinguished Research Staff, Oak Ridge National Laboratory

- ***Chris Johnson, Ph.D.***
  - Distinguished Professor, University of Utah
  - Director, School of Computing, University of Utah
  - Director, Scientific Computing and Imaging Institute (SCI)
  - Director, Center for Bioelectric Field Modeling, Simulation, and Visualization (NIH NCRR)
  - Co-Founder, Visual Influence

Current 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: Revised findings and recommendations and solicited additional input
- January 11, 2005: Call for information sent to Federal agencies involved in computational science
- *January 12 (today): PITAC meeting - Status report on findings and recommendations*
- January to April: Draft report
- April 14: Review and approval at PITAC meeting

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:

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

Research Areas and Priorities

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?*

Findings (Red denotes new findings)

- Today's computational science ecosystem is unbalanced, with a software base that is inadequate to support and track evolving hardware and application needs.
- The creation and long-term maintenance of software that is key to computational science requires the support of the Federal government. This software includes operating systems, libraries, compilers, software development and data analysis tools, application codes and databases.

Research Areas and Priorities

Findings (continued)

- Computational science has no clear roadmap outlining decadal priorities for investment, with a clear assessment of those priorities derived from a survey of the problems and challenges.

- 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.

Research Areas and Priorities

Recommendations

- Create a multi-decade computational science roadmap that *identifies the most important problems* (e.g., algorithms, applications, architecture, infrastructure, and software) and *prioritizes investment areas*, *funding levels*, *and recommended schedules* to guide government (both individual agency and interagency), academic and industry investment.
  - Regularly update the roadmap each two years to ensure relevance and continuity.

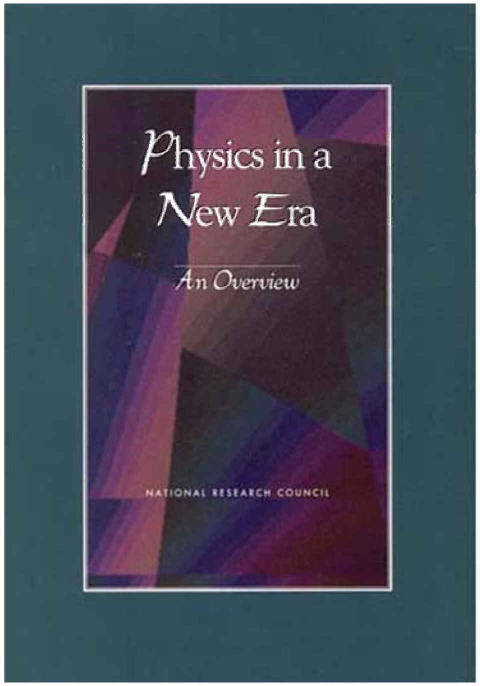

Research Areas and Priorities

Recommendations (continued)

- Investment priorities should include
  - Software, including operating systems, libraries, compilers, software development and data analysis tools, application codes and databases.
  - Data analysis tools for heterogeneous, multimodal data, including business intelligence, data and information visualization, mining & processing capabilities.
  - Numerical algorithms and tools for solving complex problems.
  - Next-generation architectures better matched to the characteristics of complex computational science applications.
- Agencies strategies for computational science should be shaped in response to the decadal roadmap. The result should be strategic plans that recognize and address roadmap priorities and funding requirements.

Short Term vs. Long Term

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?

Findings

- Short-term investment and limited strategic planning have led to excessive focus on incremental research rather than on long-term sustained research with lasting impact that can solve important problems.
- Developing leading-edge computational science applications is a complex process involving teams of people that often must be sustained for a decade or more to yield the full fruits of investment.

Short Term vs. Long Term

Findings (continued)

- A sustained infrastructure is needed that provides access to leading-edge capabilities for computational science. This will require long-term investments and strategic procurements coupled with evolving scientific roadmaps.

- Data intensive 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.

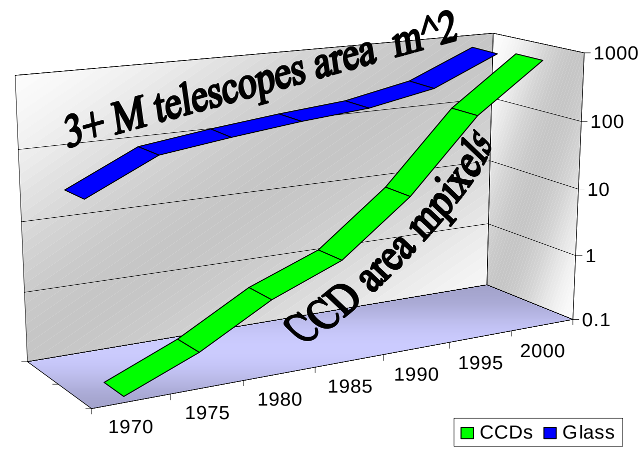

Source: Alex Szalay/Jim Gray

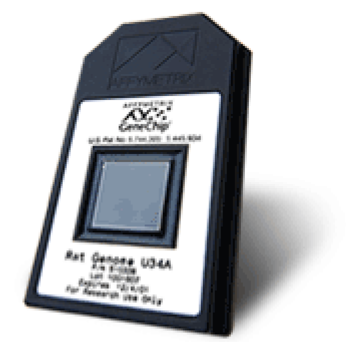

AFFYMETRIX GENECHIP®

Short Term vs. Long Term

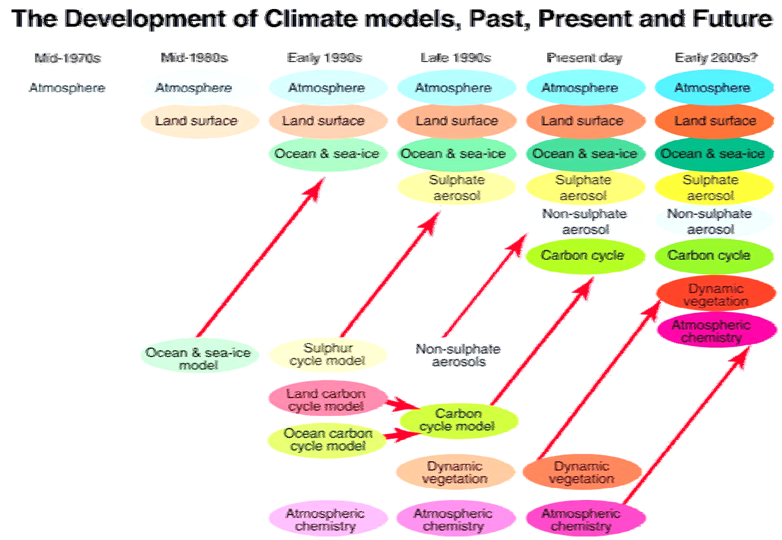

Recommendations

- Sustain investment in long-term computational science research and infrastructure development.
  - The Internet emerged as an international phenomenon and economic driver only after over twenty years of federally-funded research and development.
  - Developing and validating climate models that incorporate ocean, atmosphere, sea ice, and human interactions have required multiple cycles of development, computational experimentation, and analysis spanning multiple decades.

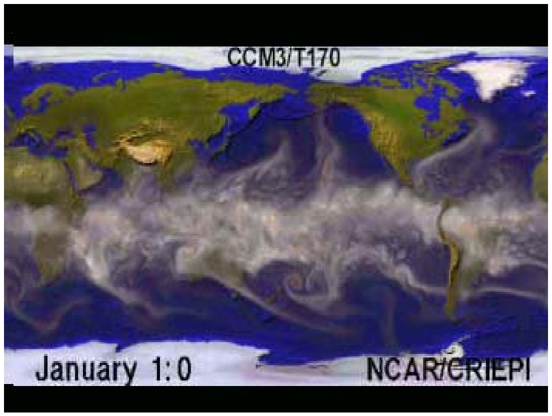

Short Term vs. Long Term

Recommendations (continued)

- Encourage diversification of funding agency research portfolios to create a more balanced mixture of long-term and high-risk projects on the one hand and shorter-term, lower-risk activities on the other.
  - Implementation mechanisms might include allocation of a portion of each agency's research budget to programs whose *raison d'etre* is fostering high-risk exploration, with concomitant changes to the peer-review and funding-decision mechanisms to ensure that risk diversification actually occurs.
  - The National Institutes of Health (NIH) roadmap includes such high risk investment mechanisms.
  - see http://nihroadmap.nih.gov/

Short Term vs. Long Term

Recommendations (continued)

- 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.
  - Ubiquitous sensors and high-resolution sensors will soon routinely produce petabytes of data. We must act now to develop the requisite data-mining, visualization and information extraction tools.

Short Term vs. Long Term

Recommendations (continued)

- Create a next-generation software, architecture, and algorithms program whose goal is to build advanced prototypes of novel computing systems.
  - Much as DARPA funded creation of ARPANet, ILLIAC IV and other systems in the 1970s and other systems in the 1980s and 1990s, these prototyping projects would have lifetimes of sufficient length and budgets of sufficient scope to develop, test and assess the capabilities of alternative designs.
  - These "expeditions to the 21st century" were also outlined in the 1999 PITAC report as a means to create systems better matched to the needs of computational science applications.

Short Term vs. Long Term

Recommendations (continued)

- In recognition of the rapidly rising complexity of today's computational science applications and the human resources and time needed to create them, establish an interagency computational science program that funds integrated, multidisciplinary teams of discipline researchers and professional software developers.
  - Like physical infrastructure, computational science applications are highly complex structures whose development requires skilled professionals.

Software Complexity and Growth

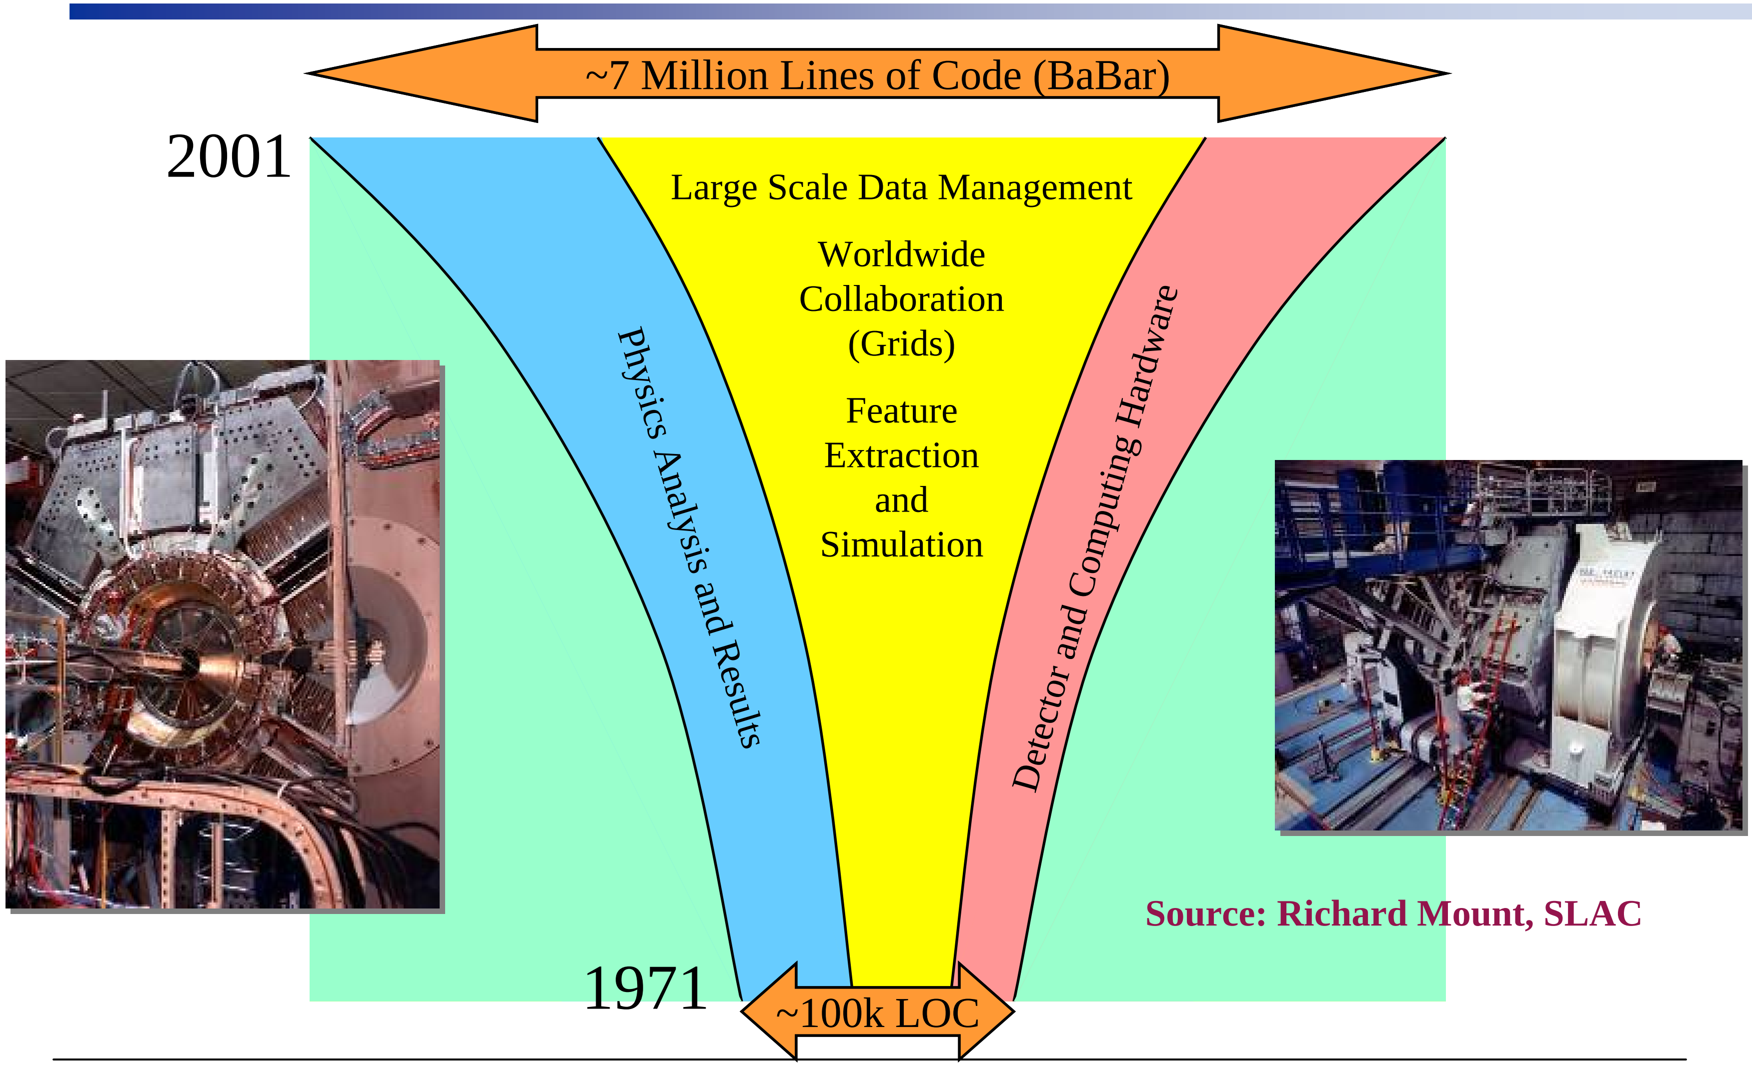

Source: Richard Mount, SLAC

System Balance

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?

Findings

- Computational science is an ecosystem, where scientific needs stimulate exploration and creation of new techniques, and new techniques enable exploration of new scientific domains.

Wetland Restoration

System Balance

Findings (continued)

- Balanced investment in both computational science techniques and their application are required for continued progress.
  - Research in high-end architecture, systems software, programming models, algorithms, tools, and mathematical methods is not the same as research in using high-end computing to address challenging applications.
  - *Both kinds of research are important, but they require different expertise. In general, they are conducted by different people, and it is a mistake to confound them.*

System Balance

Findings (continued)

- Fragile software and the lack of sustainable infrastructure often limit the ability of disciplinary and interdisciplinary teams to integrate and support complex computational science infrastructure. They also consume the intellectual energies of students and research staff, to the detriment of research.

- Academic software tools often fail to leverage commercial software capabilities and best practices, placing unduly large intellectual burdens on researchers and research teams.

System Balance

Recommendations

- Fund software development as a peer to research instrument infrastructure, with explicit software infrastructure line items in agency budgets that more accurately reflect the true costs of robust, reliable and well-documented software.

- Recognize the need for greater investment in commercial-quality software as a direct cost component of research projects, leveraging either professional software developers or extant commercial software.

System Balance

Recommendations (continued)

- Engage commercial software vendors in collaborative partnerships to develop and build elements of a national cyberinfrastructure.
- Create national software sustainability centers whose charge is to harden, document and support vital, non-commercial computational science software (both enabling technologies and domain-specific software) whose useful lifetime may be measured in decades.
  - An independent body, whose members are drawn from government, academia, and industry, would choose the software targets for these centers.

Computational Science Training

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?

Findings

- Interdisciplinary education in computational science and computing technologies is inadequate, reflecting the traditional disciplinary boundaries in higher education.
  - Interdisciplinary computational science research and education would also benefit from inclusion of the social sciences and humanities, particularly as complex scientific and engineering problems touch public policy.
  - Only systemic change to university organizational structures will yield the needed outcomes.

Computational Science Training

Findings (continued)

- The limited number of senior leaders in computational science willing to assume national roles has constrained community advocacy and agency leadership.

Computational Science Training

Recommendations

- Real world, complex problems require collaborative research.
  - Challenge universities to substantially change their structures to value and reward interdisciplinary and collaborative research and education.
  - Set aside funds to support the demonstration of new educational models that are multidisciplinary.
  - Offer experiential and collaborative learning environments at the graduate and undergraduate level and tie these environments to ongoing research and development efforts (which may be supported through centers and institutes).
  - Learning experiences should place students in real-world situations, including internships and field experiences.
- Fund curriculum development in computational science, targeting best practices, models and structures.

Computational Science Training

Recommendations

- Establish a leadership development program for computational sciences that targets younger researchers and exposes them to the processes and challenges associated with project management, community planning and government relations. Conduct case studies and interviews to determine those factors that most often influence researchers to become national leaders.
- Examine the current rules for government service, which often impose substantial financial and personal hardships on those who serve. Seek interpretations that are more flexible and engaging for dual career families. For example, consider strategies that would allow participants to remain at their home institutions, relying on electronic collaboration tools for distributed meetings.

Interagency Coordination

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?

Findings

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

Interagency Coordination

Findings (continued)

- Only sustained investment in computational science infrastructure, defined broadly to include people, software, data, and systems, will enable us to realize fully the promise of computational science.

- The National Coordination Office (NCO) for Information Technology Research and Development (ITR&D) has been an effective vehicle for information sharing and communication across the Federal government and the research community.

Interagency Coordination

Recommendations

- Derive agency and interagency computational science investment priorities from a regularly updated roadmap of research opportunities; see the roadmap recommendation under question 1. The recent High-End Computing Revitalization Task Force (HECRTF) was one such interagency response to a perceived need, albeit without a formally defined community roadmap.

- Create rewards for interagency coordination. In particular, ensure that the interagency working group (IWG) in ITR&D within the National Science and Technology Council (NSTC) is composed of individuals with the authority to make agency budgetary commitments. This would simplify creation and coordination of interagency programs.

Keeping Pace With Change

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.

Findings

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

Keeping Pace With Change

Findings (continued)

- 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.
- 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 and simulation using sensor data from diverse sources.

Keeping Pace With Change

Recommendations

- Lead investment in alternative architectures that are better matched to the needs of computational science, notably systems with better memory architectures (higher bandwidth, latency hiding), lower performance variability and improved I/O systems.

- Require deposit of scientific data and research software as a component of funding review. Concomitantly, develop one or more national repositories of data and research software for community use.

Keeping Pace With Change

Recommendations (continued)

- Science and computational science are in social flux as they undergo major changes in the tools, techniques and team structures used to conduct research.
  - Through case studies and the computational science roadmap of major research opportunities, determine the appropriate distribution of project team sizes to maximize scientific return. Use this research to shape agency structures, programs and funding profiles.
- Large-scale scientific instruments (e.g., accelerators, telescopes and environmental observatories) have operational lifetimes measured in decades and are expensive to relocate. It is similarly costly to replicate the physical plant and ancillary support environment for computational science infrastructure.
  - In recognition of these costs, separate periodic review of infrastructure management and processes from an assessment the infrastructure's utility and continued support.

Barriers

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

Findings

- There is a shortage of easy to use, accessible, scalable software that interoperates with existing user environments.
- Community verification and validation of computational science results, via access to the computational model software and associated data, could accelerate scientific discovery.

Barriers

Findings (continued)

- National computing resources, namely 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.

- Current agency progress and reporting structures may create an undue administrative burden on many researchers.

Barriers

Recommendations

- Fund the definition, design and implementation of computational science community repositories including defined frameworks, metadata structures, datasets, algorithms/applications and review/validation processes, deposit and access policies, etc.
- Increase the investment in software. Adequate software will emerge only if there is an aggressive, stable, sustained and long-term research program focused on software.
- Investigate mechanisms to reduce the administrative reporting burden for research projects, allowing researchers to focus on scientific discovery.

The Ongoing Search

Our immediate neighborhood we know intimately. But with increasing distance our knowledge fades. ...The search will continue. The urge is older than history. It is not satisfied, and it will not be denied. *Edwin Hubble*