CHAPTER 1
CALL TO ACTION

I. Cyberinfrastructure Drivers and Opportunities

How does a protein fold? What happens to space-time when two black holes collide? What impact does species gene flow have on an ecological community? What are the key factors that drive climate change? Did one of the trillions of collisions at the Large Hadron Collider produce a Higgs boson, the dark matter particle, or a black hole? Can we create an individualized model of each human being for personalized health care delivery? How does major technological change affect human behavior and structure complex social relationships? What answers will we find – to questions we have yet to ask – in the very large datasets that are being produced by telescopes, sensor networks, and other experimental facilities?

These questions – and many others – are only now coming within our ability to answer because of advances in computing and related information technology. Once used by a handful of elite researchers in a few research communities on select problems, advanced computing has become essential to future progress across the frontier of science and engineering. Coupled with continuing improvements in microprocessor speeds, converging advances in networking, software, visualization, data systems and collaboration platforms are changing the way research and education are accomplished.

Today's scientists and engineers need access to new information technology capabilities, such as distributed wired and wireless observing network complexes, and sophisticated simulation tools that permit exploration of phenomena that can never be observed or replicated by experiment. Computation offers new models of behavior and modes of scientific discovery that greatly extend the limited range of models that can be produced with mathematics alone – for example, chaotic behavior. Fewer and fewer researchers working at the frontiers of knowledge can carry out their work without cyberinfrastructure of one form or another.

While hardware performance has been growing exponentially – with gate density doubling every 18 months, storage capacity every 12 months, and network capability every 9 months – it has become clear that increasingly capable hardware is not the only requirement for computation-enabled discovery. Sophisticated software, visualization tools, middleware and scientific applications created and used by interdisciplinary teams are critical to turning flops, bytes and bits into scientific breakthroughs. In addition to these technical needs, the exploration of new organizational models and the creation of enabling policies, processes, and economic frameworks are also essential. The combined power of these capabilities and approaches is necessary to advance the frontiers of science and engineering, make seemingly intractable problems solvable, and pose profound new scientific questions.

The comprehensive infrastructure needed to capitalize on dramatic advances in information technology has been termed cyberinfrastructure (CI). Cyberinfrastructure integrates hardware for computing, data and networks, digitally-enabled sensors, observatories and experimental facilities, and an interoperable suite of software and middleware services and tools. Investments in interdisciplinary teams and cyberinfrastructure professionals with expertise in algorithm development, system operations, and applications development are also essential to exploit the full power of cyberinfrastructure to create, disseminate, and preserve scientific data, information and knowledge.

For four decades, NSF has provided leadership in the scientific revolution made possible by information technology (Appendices B and C). Through investments ranging from supercomputing centers and the Internet to software and algorithm development, information tech-

The Terashake 2.1 simulation on the opposite page depicts a velocity wavefield as it propagates through the 3D velocity structure beneath Southern California. Red and yellow colors indicate regions of compression, while blue and green colors show regions of dilation. Faint yellow (faults), red (roads), and blue (coast-line) lines add geographical context.

nology has stimulated scientific breakthroughs across all science and engineering fields. Most recently, NSF's Information Technology Research (ITR) priority area sowed the seeds of broad and intensive collaboration among the computational, computer, and domain research communities that sets the stage for this "Call to Action."

NSF is the only agency within the U.S. government that funds research and education across all disciplines of science and engineering. Over the past five years, NSF has held community workshops, commissioned blue-ribbon panels, and carried out extensive internal planning (Appendix B). Thus, it is strategically placed to leverage, coordinate and transition cyberinfrastructure advances in one field to all fields of research.

Other federal agencies, the administration, Congress, the private sector, and other nations are aware of the growing importance of cyberinfrastructure to progress in science and engineering. Other federal agencies have planned improved capabilities for specific disciplines, and in some cases to address interdisciplinary challenges. Other countries have also been making significant progress in scientific cyberinfrastructure. Thus, the U.S. must engage in and actively benefit from cyberinfrastructure developments around the world.

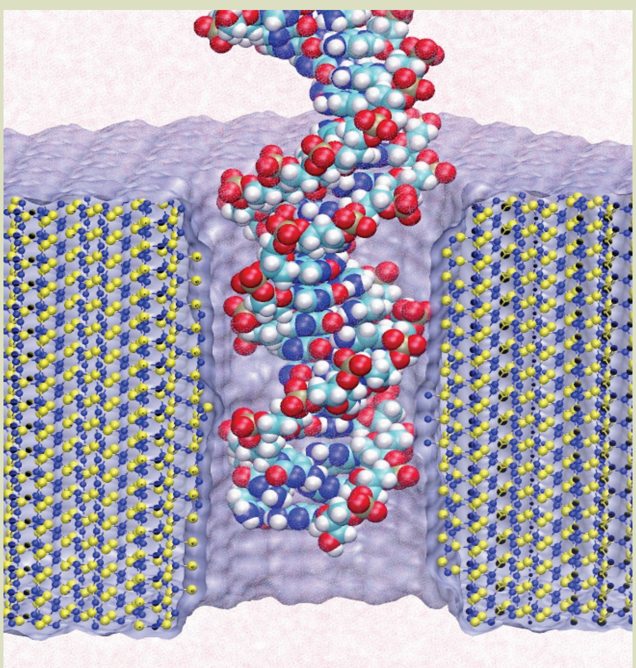

Visualization of a molecular dynamics simulation of a double stranded DNA molecule as it enters a nanopore in a silicon nitride membrane.

Not only is the time ripe for a coordinated investment in cyberinfrastructure, but progress at the science and engineering frontiers depends on it. Our communities are in place and are poised to respond to such an investment.

Working with the science and engineering research and education communities and partnering with other key stakeholders, NSF is ready to lead.

II. Vision, Mission and Principles for Cyberinfrastructure

A. Vision

NSF will play a leadership role in the development and support of a comprehensive cyberinfrastructure essential to 21st century advances in science and engineering research and education.

B. Mission

NSF's mission for cyberinfrastructure (CI) is to:

- Develop a human-centered CI that is driven by science and engineering research and education opportunities;
- Provide the science and engineering communities with access to world-class CI tools and services, including those focused on: high performance computing; data, data analysis and visualization; networked resources and virtual organizations; and learning and workforce development;
- Promote a CI that serves as an agent for broadening participation and strengthening the nation's workforce in all areas of science and engineering;
- Provide a sustainable CI that is secure, efficient, reliable, accessible, usable, and interoperable, and that evolves as an essential national infrastructure for conducting science and engineering research and education; and
- Create a stable but extensible CI environment that enables the research and education communities to contribute to the agency's statutory mission.

C. Principles

The following principles will guide the agency's FY 2006 through FY 2010 investments:

- Science and engineering research and education are foundational drivers of CI.
- NSF has a unique leadership role in formulating and implementing a national CI agenda focused on advancing science and engineering.
- Inclusive strategic planning is required to effectively address CI needs across a broad spectrum of organizations, institutions, communities and individuals, with input to the process provided through public comments, workshops, funded studies, advisory committees, merit review and open competitions.
- Strategic investments in CI resources and services coupled with enabling policy and organizational framework are essential to continued U.S. leadership in science and engineering.
- The integration and sharing of cyberinfrastructure assets deployed and supported at national, regional, local, community and campus levels represent the most effective way of constructing a comprehensive CI ecosystem suited to meeting future needs.
- Public and private national and international partnerships that integrate CI users and providers and benefit NSF's research and education communities are also essential for enabling next-generation science and engineering.
- Existing strengths, including research programs and CI facilities, serve as a foundation upon which to build a CI designed to meet the needs of the broad science and engineering community.
- Merit review is essential for ensuring that the best ideas are pursued in all areas of CI funding.
- Regular evaluation and assessment tailored to individual projects is essential for ensuring accountability to all stakeholders.
- A collaborative CI governance and coordination structure that includes representatives who contribute to basic CI research, development and deployment, as well as those who use CI, is essential to ensure that CI is responsive to community needs and empowers research at the frontier.

Robert Patterson demonstrates NCSA's 3D Visualization to Dr. Arden Bement, the Director of NSF, and others during the FY08 NSF budget roll-out.

III. Goals and Strategies

NSF's vision and mission statements on CI need well-defined goals and strategies to turn them into reality. The goals underlying these statements are provided below, with each goal followed by a brief description of the strategy to achieve the goal.

Across the CI landscape, NSF will:

- ***Provide communities addressing the most computationally challenging problems with access to a world-class, high performance computing (HPC) environment through NSF acquisition and through exchange-of-service agreements with other entities, where possible.***

NSF's investment strategy for the provision of CI resources and services will be linked to careful requirements analyses of the computational needs of research and education communities. NSF investments will be coordinated with those of other agencies in order to maximize access to these capabilities and to provide a range of representative high performance architectures.

- ***Broaden access to state-of-the-art computing resources, focusing especially on institutions with less capability and communities where computational science is an emerging activity.***

Cyberinfrastructure will broaden access to state-of-the art resources for learning and discovery, creating new opportunities for participation by emerging and underserved communities.

Building on the achievements of current CI service providers and other NSF investments, the agency will work to make necessary computing resources more broadly available, paying particular attention to emerging and underserved communities.

- ***Support the development and maintenance of robust systems software, programming tools, and applications needed to close the growing gap between peak performance and sustained performance on actual research codes, and to make the use of HPC systems, as well as novel architectures, easier and more accessible.***

NSF will build on research in computer science and other research areas to provide science and engineering applications and problem-solving environments that more effectively exploit innovative architectures and large-scale computing systems. NSF will continue and build on its existing collaborations with other agencies in support of the development of HPC software and tools.

- ***Support the continued development, expansion, hardening and maintenance of end-to-end software systems – user interfaces, workflow engines, science and engineering applications, data management, analysis and visualization tools, collaborative tools, and other software integrated into complete science and engineering systems via middleware – in order to bring the full power of a national cyberinfrastructure to communities of scientists and engineers.***

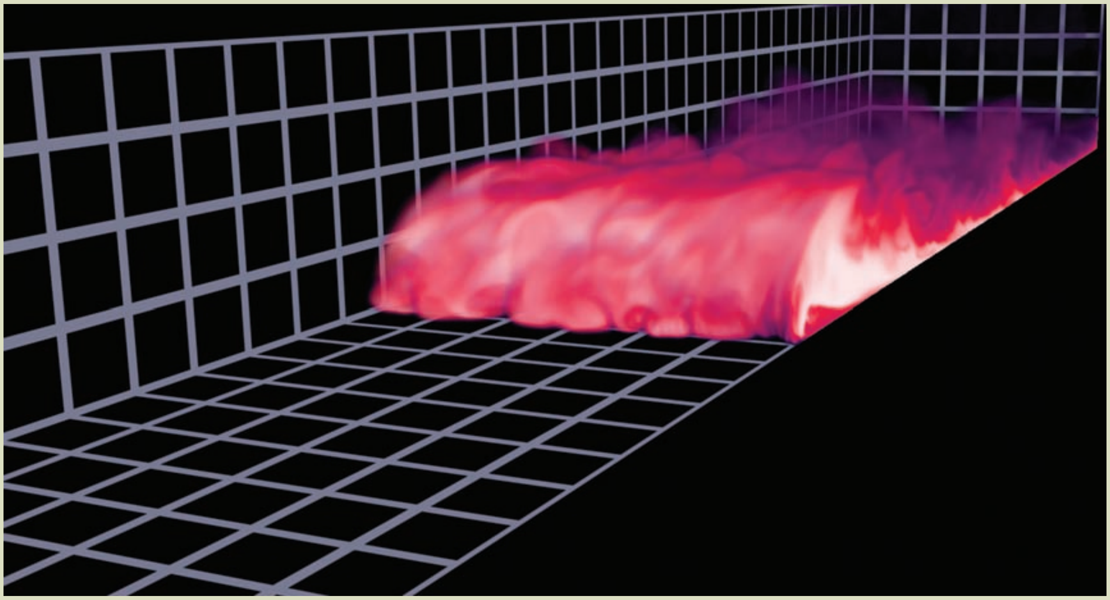

NCSA's Cobalt computing system uses a 3D cylindrical configuration to model the sediment discharge of a river into the ocean and the initial stages of alluvial fan formation at the river's mouth.

These investments will build on the software products of current and former programs, and will leverage work in core computer science research and development efforts supported by NSF and other federal agencies.

- ***Support the development of the computing professionals, interdisciplinary teams, enabling policies and procedures, and new organizational structures such as virtual organizations, that are needed to achieve the scientific breakthroughs made possible by advanced CI, paying particular attention to opportunities to broaden the participation of underrepresented groups.***

NSF will continue to improve its understanding of how participants in its research and education communities, as well as the scientific workforce, can use CI. For example, virtual organizations empower communities of users to interact, exchange information, and access and share resources through tailored interfaces. Some of NSF's investments will focus on appropriate mechanisms or structures for use, while others will focus on how best to train future users of CI. NSF will take advantage of the emerging communities associated with CI that provide unique and special opportunities for broadening participation in the science and engineering enterprise.

- ***Support state-of-the-art innovation in data management and distribution systems, including digital libraries and educational environments that are expected to contribute to many of the scientific breakthroughs of the 21st century.***

NSF will foster communication among forefront data management and distribution systems, digital libraries, and other education environments sponsored in its various directorates. NSF will ensure that its efforts take advantage of innovation in large data management and distribution activities sponsored by other agencies and through international efforts. These developments will play a critical role in decisions that NSF makes about stewardship of long-lived data.

- ***Support the design and development of the CI needed to realize the full scientific poten-***

The DANSE project at CalTech integrates new materials theory with high-performance computing, using data from facilities such as DOE's new Spallation Neutron Source in Oak Ridge, TN.

tial of NSF's investments in tools and large facilities, from observatories and accelerators to sensor networks and remote observing systems.

NSF's investments in large facilities and other tools require new types of CI – such as wireless control of networks of sensors in hostile environments, rapid distribution and analysis of petascale data sets around the world, adaptive knowledge-based control and sampling systems, and innovative visualization systems for collaboration. NSF will ensure that these projects invest appropriately in CI capabilities, promoting the integrated and widespread use of the unique services provided by these and other facilities. In addition, NSF's CI programs will be designed to serve the needs of these projects.

- ***Support the development and maintenance of the increasingly sophisticated applications needed to achieve the scientific goals of research and education communities.***

The applications needed to produce cutting-edge science and engineering have become increasingly complex. They require teams, even communities, to develop and sustain wide and long-term applicability, and they leverage underlying software tools and increasingly common, persistent CI resources such as data repositories and authentication and authorization services. NSF's investments in applications will involve its directorates and offices that support domain-specific science and engineering. Special attention will be paid to the cross-disciplinary nature of much of the work.

- ***Invest in the high-risk/high-gain basic research in computer science, computing and storage devices, mathematical algorithms, and the human/CI interfaces that are critical to powering the future exponential growth in all aspects of computing, including hardware speed, storage, connectivity and scientific productivity.***

NSF's investments in operational CI must be coupled with vigorous research programs in the directorates to ensure that operational capabilities continue to expand and extend in the future. Important among these programs are activities to understand how humans adopt and use CI. NSF is especially well-placed to foster collaborations among computer scientists; social, behavioral and economic scientists; and other domain scientists and engineers to understand how humans can best use CI, in both research and education environments.

- ***Provide a framework that will sustain reliable, stable resources and services while enabling the integration of new technologies and research developments with a minimum of disruption to users.***

NSF will minimize disruption to users by realizing a comprehensive CI with an architecture and framework that emphasizes interoperability and open standards, thus providing flexibility for upgrades, enhancements and evolutionary changes. Pre-planned arrangements for alternative CI availabilities during competitions, changeovers and upgrades to production operations and services will be made, including cooperative arrangements with other agencies.

A strategy common to achieving all of these goals is partnering nationally and internationally, with other agencies, the private sector, and with universities to achieve a worldwide CI that is interoperable, flexible, efficient, evolving and broadly accessible. In particular, NSF will take a lead role in formulating and implementing a national CI strategy.

IV. Planning for Cyberinfrastructure

To implement its cyberinfrastructure vision, NSF will develop interdependent plans for each of the following aspects of CI, with emphasis on their integration to create a balanced science- and engineering-driven national CI:

- High Performance Computing
- Data, Data Analysis, and Visualization
- Virtual Organizations for Distributed Communities, and
- Learning and Workforce Development.

Others may be added at a later date.

While these aspects are addressed separately as a means for organizing this document, the central goal is the development of a fully-integrated CI framework comprised of the balanced, seamless blending of these components. This will require integrative management structures (such as the newly formed Office of Cyberinfrastructure, the NSF-wide Cyberinfrastructure Council, and the Cyberinfrastructure Coordinators' Committee), as well as science-driven, community-based planning and implementation processes that span all the elements of a truly comprehensive CI framework.

These plans will be reviewed annually and will evolve over time, paced by the considerable rate of innovation in computing and communication, and by the growing needs of the science and engineering community for state-of-the-art CI capabilities. Through cycles of use-driven innovation, NSF's vision will become reality.

Researchers upgrade the software of an automated weather station that transmits data to help track the iceberg's position in the Antarctic and reports on the microclimate of the ice surface.


IBM
System p5
IBM
System p5