PPPL-4090

DOE Greenbook

Needs and Directions in High Performance Computing for the Office of Science

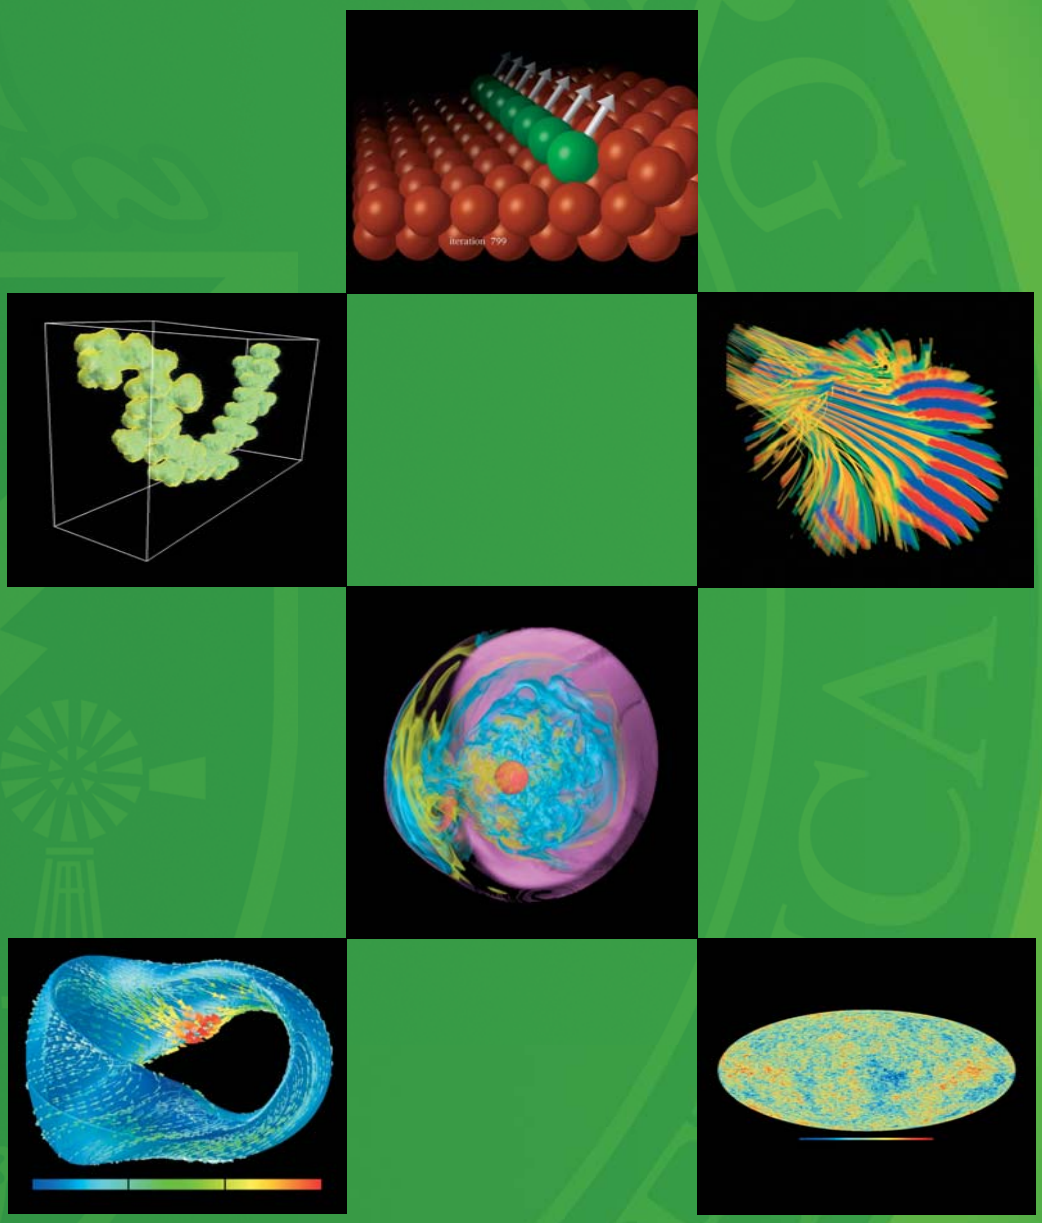

A Report from the NERSC User Group
June 2005

Prepared for the U.S. Department of Energy Office of Science

Office of Science
U.S. DEPARTMENT OF ENERGY

Princeton Plasma Physics Laboratory

PPPL-4090

PPPL-4090

DOE Greenbook:
Needs and Directions in High Performance Computing for the Office of Science

A Report from the NERSC User Group
Prepared for the U.S. Department of Energy
Office of Science

S. C. Jardin, Editor

June 2005

Prepared for the U.S. Department of Energy under Contract DE-AC02-76CH03073.

PPPL Report Disclaimers

Full Legal Disclaimer

Trademark Disclaimer

PPPL Report Availability

This report is posted on the U.S. Department of Energy's Princeton Plasma Physics Laboratory Publications and Reports web site in Fiscal Year 2005. The home page for PPPL Reports and Publications is: http://www.pppl.gov/pub_report/

Office of Scientific and Technical Information (OSTI):

Available electronically at: http://www.osti.gov/bridge.

Available for a processing fee to U.S. Department of Energy and its contractors, in paper from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062

Telephone: (865) 576-8401
Fax: (865) 576-5728
E-mail: reports@adonis.osti.gov

National Technical Information Service (NTIS):

This report is available for sale to the general public from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161

Telephone: (800) 553-6847
Fax: (703) 605-6900
Email: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/ordering.htm

Contributors:

A. Aspuru-Guzik, University of California, Berkeley
V. Batista, Yale University
E. W. Bethel, Lawrence Berkeley National Laboratory
J. D. Borrill, Lawrence Berkeley National Laboratory
J. Chen, Sandia National Laboratory
P. Colella, Lawrence Berkeley National Laboratory
D. J. Dean, Oak Ridge National Laboratory
T. DeBoni, Lawrence Berkeley National Laboratory
D. Donzis, Georgia Institute of Technology
S. Ethier, Princeton Plasma Physics Laboratory
A. Friedman, Lawrence Berkeley National Laboratory
R. Harrison, Oak Ridge National Laboratory
E. Hawkes, Sandia National Laboratory
M. Head-Gordon, University of California, Berkeley
B. E. Hingerty, Oak Ridge National Laboratory
S. C. Jardin, Princeton Plasma Physics Laboratory
Y. Jung, University of California, Berkeley
D. Keyes, Columbia University
W. A. Lester, Jr., University of California, Berkeley
F. Liu, University of Utah
M. Mavrikakis, University of Wisconsin
A. Mezzacappa, Oak Ridge National Laboratory
D. Olson, Lawrence Berkeley National Laboratory
K. Refson, Rutherford Appleton Laboratory, United Kingdon
C. Ren, University of Rochester
D. Rotman, Lawrence Livermore National Laboratory
R. Ryne, Lawrence Berkeley National Laboratory
V. Sankaran, Georgia Institute of Technology
C. Sovinec, University of Wisconsin, Madison
D. Spong, Oak Ridge National Laboratory
G. Sposito, University of California, Berkeley
G. M. Stocks, Oak Ridge National Laboratory
R. Sugar, University of California, Santa Barbara
D. Swesty, State University of New York, Stony Brook
H. Wang, New Mexico State University
V. Wayland, National Center for Atmospheric Research
P. K. Yeung, Georgia Institute of Technology

Table of Contents

1. Summary

1.1 Computational Challenges

Computational science plays an essential role within each of the research programs in the U.S. Department of Energy (DOE) Office of Science (SC). There are a diversity of traditional scientific disciplines conducting research within the SC programs; including chemistry, materials science, quantum physics, geophysics, biology, plasma physics, nuclear physics, and high energy physics, among others. Researchers from within each of these disciplines are finding that computation at the largest scale provides a capability that is now considered essential for the advancement of each of the respective disciplines. Today's most powerful computers and the associated software are being used to produce new and more precise scientific results at the cutting edge of these disciplines, and this trend is destined to continue for years to come.

Many examples of the impact of large-scale computations on the sciences are presented in this document. Our basic understanding of chemical processes has increased. New insight as to the molecular basis for vision has been obtained. A comprehensive study of the impact of pollutants on climate change has been carried out. Our understanding of the complex process of photosynthesis has increased. Significant progress has been made in our ability to simulate the folding of proteins.

The improved models within fusion energy sciences have allowed researchers to design a new class of fully 3D magnetic confinement configurations, known as quasi-symmetric stellarators, with a "hidden symmetry" that should allow greatly improved confinement and stability properties. The inertial-confinement fusion computational program has developed high-resolution beam simulations that are being used to develop and guide experiments, and fast-ignition simulations that are elucidating the physics of that potential breakthrough option.

The high energy physics program has made great progress both in accelerator modeling and in obtaining detailed computational predictions of the masses of strongly interacting particles, an important test of quantum chromodynamics (QCD) and the Standard Model. Large-scale computation has provided new insight into the properties of atomic nuclei and how they behave at extremely high energy density. Massive supernovae have been modeled with unprecedented detail, shedding light on the evolution and fate of the Universe.

These and many more scientific advances have been enabled by the extraordinary computational science resources available at the National Energy Research Scientific Computing Center (NERSC). However, a common theme in each of the disciplines is that the current computational resources available through NERSC are saturated, and this lack of additional computing resources is becoming a major bottleneck in the scientific research and discovery process. A large increase in computer power is needed in the near future to take the understanding of the science to the next level and to help secure the U.S. DOE SC leadership role in these fundamental research areas.

1.2 General Recommendations

1. Expand the high performance computing resources available at NERSC. The upcoming procurement must ensure a large increase in compute cycles available, as well as an appropriate balance of cache memory, processor memory, memory bandwidth, internode communication speed, intranode communication speed, and other computational equipment required to support the wide range of large-scale applications involving production computing and development activities in the DOE Office of Science.

2. The computing hardware and queuing systems should be configured to minimize the time-to-completion of large jobs, as well as to maximize the overall efficiency of the hardware. Here "large" can refer to jobs requiring long running times, a large number of processors, exceptionally large memory, or any combination of these. We do not recommend policies that encourage individual jobs to use more processors than is justified by their memory requirements and scaling characteristics if it degrades overall facility efficiency and overall job time-to-completion.

3. NERSC should actively support the continued improvement of algorithms, software, and database technology for improved performance on parallel platforms. This implies that NERSC should continue to retain high-quality consultants and performance support personnel to proactively assist the user community in obtaining maximum performance from the hardware.

4. Significantly strengthen the computational science "infrastructure" at NERSC that will enable the optimal use of the current and future NERSC supercomputers. Current simulations are generating data at tens of terabytes per simulation over the course of a few days. Simulations that will be performed in the not-too-distant future will generate hundreds of terabytes of data per simulation over the course of months. Moreover, these simulations will be performed, and the results analyzed, by geographically distributed teams. Consequently, local disk and archival storage, networking between NERSC's supercomputer and local storage and between NERSC and the wide-area network (WAN), and capabilities for local and remote data analysis and visualization must all be developed in order to enable the scientific workflows that will produce next-generation computational science.

5. Several of NERSC's current and potential future scientific applications are especially data or I/O intensive. These requirements should be carefully evaluated in order to support as wide a range of science as possible while also realizing significant benefits in both performance and cost in the computer configuration.

2. Introduction

The Users Group of the National Energy Research Scientific Computing Center (NERSC), known as NUG, consists of all investigators who have been allocated computational and storage resources at this Department of Energy Office of Science (SC) facility. These users elect an executive body known as NUGEX. At the June 2004 meeting of NUG, held at the flagship NERSC facility in Oakland/Berkeley, CA, the NUGEX began the process of assessing the role of computational science within DOE SC and determining the computational needs of future programs within the office. This topic was again discussed by the membership at the November 2004 NUG meeting, held in conjunction with the Supercomputing 2004 meeting in Pittsburgh, PA. At this time, a core writing team was formed which represented all offices within SC and all major user groups of the NERSC facilities. This group performed its work while communicating largely through email and conference calls, and delivered a final document in June 2005.

After serving the user community for decades as primarily a provider of sequential-vector compute cycles, NERSC has now completely transitioned to a massively parallel platform (MPP) facility. This transition has had a major impact on its user community. Researchers from every discipline needed to transform their application codes to make use of the parallel computers, and the NERSC staff played an essential role in providing education, tools for debugging, mathematical libraries, communication software, and other support needed for this to occur.

The unprecedented memory and processing power provided by these large MPP systems have allowed scientists in essentially every discipline within SC to dramatically increase the resolution, level of realism, and complexity in their simulation models. This has taken computational science to a new plateau in these application areas, and is underpinning a wide range of scientific research and discovery. As a direct result of these successes, we find that the NERSC production facilities are becoming significantly oversubscribed, in many cases leading to long delays in obtaining results of a calculation, which in turn seriously impedes the pace of the associated research. In addition, these simulations are producing enormous amounts of data, and the movement, storage, and analysis of this data presents its own challenges.

There are two broad classes of computational requirements that we refer to as *capability* and *capacity*. Capability needs refer to the overall size of the largest calculation that can be run at the center: one requiring most or all of the memory and computing power of the entire system. Capacity needs refer to the overall throughput: how many teraflop-hours are available to simultaneously meet the needs of the large, diverse user base. These are both legitimate needs, although they tend to stress the system in different ways and are sometimes at odds with one another.

There is a definite and documented need for both increased capacity and capability computer resources within the DOE Office of Science at this time. However, it is now well understood that providing teraflops alone does not meet this need. Faster processor speeds need to be matched with faster and larger memory and with faster communication

speeds throughout the system in order to be effectively utilized. Networks, storage facilities, and visualization capabilities must be enhanced in order to deal with the ever-increasing rate of data generation. Advanced algorithms need to be developed and implemented in the form of robust software packages that enable the efficient usage of the increasingly complex hardware. And the education and support role played by NERSC remains essential as the systems evolve.

In this document we illustrate some computational science accomplishments and new directions, and identify associated needs from the six Office of Science disciplines: Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Science, High Energy Physics, Nuclear Physics, and Advanced Scientific Computing Research. We also summarize the current NERSC computational and human resources. Finally, we provide a set of recommendations from the user community concerning the future directions and acquisitions at NERSC.

3. Computational Challenges in the Office of Science

3.1 Office of Advanced Scientific Computing Research

The Office of Advanced Scientific Computing Research is composed of three divisions: the Mathematical, Information, and Computational Sciences Division, the Technology Research Division, and the Office of Scientific and Technical Information.

The primary mission of the Advanced Scientific Computing Research (ASCR) program, which is carried out by the Mathematical, Information, and Computational Sciences Division, is to discover, develop, and deploy the computational and networking tools that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex phenomena important to DOE SC research. To accomplish this mission, the program fosters and supports fundamental research in advanced scientific computing, applied mathematics, computer science, and networking, and operates supercomputer, networking, and related facilities. The applied mathematics research efforts provide the fundamental mathematical methods to model complex physical and biological systems. The computer science research efforts enable scientists to efficiently run these models on the highest-performance computers available and to store, manage, analyze, and visualize the massive amounts of data that result. The networking research provides the techniques to link the data producers, e.g., supercomputers and large experimental facilities, with scientists who need access to the data.

In fulfilling this primary mission, the ASCR program supports the Office of Science Strategic Plan's goal of providing extraordinary tools for extraordinary science as well as building the foundation for the research in support of the other goals of the strategic plan. In the course of accomplishing this mission, the research programs of ASCR have played a critical role in the evolution of high performance computing and networks.

Computer Science (CS) research plays a central role in providing the underlying technologies that enable advances in computational science. CS research includes a broad set of topics ranging from computer architecture to parallel programming models to performance and data analysis. Computational science is the process of conducting scientific research using computing technology as a supplement or surrogate for experiments and is a "consumer" of the results of CS research. Elements from both programs result in scalable and distributed technologies of the type needed to address the challenges posed by modern computational science and distributed scientific research teams. The information that follows describes current Computer Science research, development, and deployment activities along with the facilities and infrastructure requirements needed to support those activities at a production supercomputing center like NERSC. The information was obtained from a sampling of both DOE- and non-DOE-sponsored CS and visualization research efforts.

CS research projects involve conceiving, developing, and deploying the technologies that underpin subsequent advances in modern computational science applications. At the beginning of the day, a scientific researcher has an idea s/he would like to develop theoretically, prototype, and test. At the end of the day, computational science is about producing

scientific results in a reliable and consistent manner in order to serve the needs of domain scientists. Between the two are a myriad of challenging technical problems that stem from the activities of creating, debugging, tuning, and maintaining complex scientific software that runs on highly parallel computer systems. The CS research community faces a gap between the research prototypes and the production results. NERSC can play a pivotal role in bridging the gap between research and production, creating a conduit for advanced CS research that will have broad impact on how scientists approach computational science in the future.

Like a finished building, the most visible aspect of computational science is on the outside—the research results. The research and development needed to create the technologies as well as the myriad details of infrastructure that make computational science possible are all too easy to overlook when planning a large computer center. These applications, which typically run on hundreds or thousands of processors, must first undergo development, testing, debugging, performance analysis, and tuning prior to be put "into production." The general theme of our findings is that computer centers like NERSC tend to craft policies and procurements for the finished, production-ready computational science applications. However, the needs of the CS research programs are much different than those required for "production computational science." This singular focus supplies little middle ground for shepherding innovative research into the production setting. Therefore, the general theme of our findings is that NERSC is in an excellent position to propel significant advances in CS research through a relatively minor shift in programmatic focus.

During the course of all the steps leading up to production, having access to the actual production machine for interactive use is a crucial part of the research, hardening, and deployment process. Debug and performance analysis tests of new ideas take place on sets of processors ranging in number from one up to thousands. Such tools form the basis of future scientific applications only after requisite hardening on target production systems.

The DOE Mathematical, Information, and Computational Sciences Division (MICS) supports a multifaceted research portfolio that impacts many branches of computational science. One area of CS research creates new languages and communication libraries that make it easier to express scientific problems and parallelism. Another branch of CS research focuses on I/O technology for management and storage of scientific data and the marshalling of distributed data storage and computational resources. Performance analysis drives new machine architectures and helps evaluate the potential of new computational architectures. The development of new numerical algorithms can improve the quality of solutions or enable performance improvements that exceed the growth in computational capability that can be derived from hardware alone. The combined benefits derived from each of these areas of CS Research enable more effective use of available supercomputing systems and greatly improve the value derived from NERSC's current and future supercomputer systems.

3.1.1 Parallel Languages and Programming Models

Computer programming languages are the vehicles for expressing algorithms. Programming models are "theories of operation" that go beyond syntax and semantics to describe high-level approaches for using a computing language on a computer platform to solve a class of problems. They typically define how data is distributed and managed, and how processing work is managed and synchronized. The characteristics of both the programming language and the programming model have a significant impact on the class of algorithms they support. For example, a programming language well suited for manipulation of strings will likely not be ideal for solving large systems of equations.

The Unified Parallel C project (UPC) [1] is an extension of the C programming language designed for high performance computing on large-scale parallel machines. The language provides a uniform programming model for both shared and distributed memory hardware. The programmer is presented with a single shared, partitioned address space, where variables may be directly read and written by any processor, but each variable is physically associated with a single processor. UPC uses a single-program, multiple-data (SPMD) model of computation in which the amount of parallelism is fixed at program startup time, typically with a single thread of execution per processor. In order to express parallelism, UPC extends ISO C 99 with the following constructs: an explicitly parallel execution model; a shared address space, synchronization primitives, and a memory consistency model; and memory management primitives. The Titanium project[2] has similar goals, but is based on Java rather than C.

Parallel languages are built atop lower-level constructs that provide interprocessor communication (messaging) and memory management infrastructure. Projects in this space include message passing interfaces (VIA, MPI) and higher-level interfaces to message passing (Global Arrays). Global Arrays[3] provide an efficient and portable "shared-memory" programming interface for distributed-memory computers. Each process in a MIMD (multiple instruction stream, multiple data stream) parallel program can asynchronously access logical blocks of physically distributed dense multi-dimensional arrays, without need for explicit cooperation by other processes. The Message Passing Interface (MPI)[4] is the de facto standard low-level messaging interface. It provides an application programming interface that allows developers to explicitly control data management, movement, and synchronization in parallel programs, as well as to explicitly manage interprocessor synchronization. While typically used in distributed memory environments, it is also used in shared memory systems. GASNet[5]provides a standard low-level interface for messaging that lies beneath high-level languages like UPC and Titanium. GASNet allows a wider range of communication semantics to be expressed with less overhead than is possible with MPI. Past experiments with low-overhead communication, such as M-VIA/MVICH,[6] have continued to have an impact on modern cluster

[1] http://upc.nersc.gov/
[2] http://titanium.cs.berkeley.edu/
[3] http://www.emsl.pnl.gov/docs/global/ga.html
[4] http://www.mpi-forum.org/docs/docs.html
[5] http://www.cs.berkeley.edu/~bonachea/gasnet/
[6] http://old-www.nersc.gov/research/FTG/via/

designs. For instance, MVICH forms the core of the MPI implementation used by the Virginia Tech Apple cluster and even NERSC's new Jacquard system.

A central, time-proven tenet of software engineering is modular design. Extending this concept to parallel, distributed, multi-language computational science applications spans a large sector of DOE's CS research portfolio, and is something not provided by industry.[7] The Common Component Architecture[8] SciDAC Integrated Software Infrastructure Center (ISIC) aims to define a minimal set of standard interfaces that a high-performance component framework has to provide to components, and can expect from them, in order to allow disparate components to be combined to build a running application.

Many of the enabling technologies from the CS research community are developed first in a workstation environment. As they evolve, they are adapted for use in other environments and purposes. These might include massively parallel environments, or use for a specific computational science application. Since there are relatively few locations where one can perform porting and testing on thousands of processors, the CS research community naturally looks to centralized computing facilities like NERSC to provide the resources needed to support research objectives. Many problems with parallel programming models become visible only at the largest scales—levels of parallelism that the language and library designers do not have access to outside of centralized computing facilities that can provide access to thousands of processors. However, the testing and performance analysis of parallel programming models at large scale can sometimes require interactive access for periods of time, because many debugging solutions cannot be executed in batch mode.

The products from the CS research community often share a common interest: garnering the attention of the wider HPC community as well as encouraging widespread adoption and use where appropriate. There is often a huge gap between the individual CS researcher and the ultimate consumer of the technology. To help bridge this gap, the central computing facility can play a crucial role. Namely, the center can make available to its users the products from the CS research community. In many cases, such deployment will require adaptation for use in a specific environment. Advanced deployment activities are beyond the reach of the individual CS researcher, but are required in order for the technology to be used successfully on production systems. In bridging the gap, an increased level of interaction between CS researcher and production facility will be beneficial to the center, to the CS research, and to the computational science users.

Central production facilities typically have uniquely strong relationships with hardware vendors. Smaller CS research groups typically are not in a position to have such relationships, since they do not have the large hardware budget that garners the attention of the vendors. Central facilities like NERSC can play an important role in helping the language and library developers to establish closer relationships with the vendors who supply the

[7] *Federal Plan for High-End Computing: Report of the High-End Computing Revitalization Task Force (HECRTF),* (Arlington, VA: National Coordination Office for Information Technology Research and Development, May 10, 2004); http://www.house.gov/science/hearings/full04/may13/hecrtf.pdf.
[8] http://www.cca-forum.org/

production systems in order to get early access to specifications for new systems, and to help drive the vendors' design process to better accommodate advanced concepts in parallel programming exposed by these languages and libraries. This idea is consistent with the recommendations of the High End Computing Revitalization Task Force (HECRTF)[9] and NERSC's vision of "science-driven computing." Without such relationships, the primary driver on vendors for supporting the needs of the computational and computer science research communities will be the limited set of benchmarks that are used for procurements. Status quo in this regard will not produce the desired long-term results: growth and innovation in all of computer and computational science, as well as the strength and vitality of the HPC market.

3.1.2 Profiling and Performance Analysis

Due to the inherent complexity of today's large parallel platforms, computational science applications benefit from careful performance and profiling analysis.[10] Such analysis helps to detect "hot spots" in code where optimization will result in a dramatic performance improvement on a particular platform or architecture. In some cases, computational science applications benefit from using specially optimized libraries that are well tuned for a particular architecture.[11] In addition to profiling and performance analysis, debugging large parallel codes is a significant challenge. All these areas—performance analysis, profiling, and debugging—are significant CS research activities. The breadth and depth of performance analysis requirements tend to not be fulfilled by products from the computer vendors due to the absence of strong economic incentives in the scientific computing marketplace.

To achieve a high level of success, which could mean broad impact on multiple computational science communities, centers like NERSC play a key role in the development and deployment of new technologies. CS researchers require access to detailed documentation for system interfaces and characterizations that are not typically readily available through public channels. Such access is required so that performance analysis, profiling, and debugging tools work effectively. NERSC should leverage its close relationship with vendors to provide such information to the CS research community.

NERSC can also engage the MICS research community to push software architectures that cover gaps in the production infrastructure. One area in particular that continues to fall well short of meeting the needs of CS and computational science research is high-quality parallel debugging systems. NERSC can take a proactive role by providing to CS researchers and developers the same special access to NERSC hardware and technical support that matches the level of "red-carpet treatment" given to the largest scientific users. NERSC has shown examples of this kind of work with the development of non-

[9] http://www.house.gov/science/hearings/full04/may13/hecrtf.pdf

[10] Products conceived, developed, and maintained by centers like ACTS, PERC, and TOPS; performance analysis tools with unique capabilities not provided by industry, e.g., PAPI.

[11] Automatically optimizing libraries originate from various CS research programs, including SciDAC ISICs and joint efforts with other research areas (math, for example). Examples include Atlas and Optimal Solver Libraries.

invasive performance data collection tools like POE+ and IPM.[12] These tools would not have come into being without the detailed knowledge of the machine acquired by the NERSC consultants who developed them. It is important that the same kind of support be extended to particular projects in the CS research community that cover other gaps in the NERSC software infrastructure. Without strong relationships between centers like NERSC and CS research programs, there is a very real risk that new technologies may never reach the level of maturity and scalability necessary for deployment at centers like NERSC.

3.1.3 Algorithm and Machine Architecture Performance Characterization

Still other areas of CS research focus on characterizing algorithm performance on current and potential future architectures.[13] These projects can similarly benefit from the facilities and expertise at centers like NERSC. NERSC has developed considerable internal expertise on benchmarking and performance evaluation for procurements, but little of this analysis or expertise is shared outside of procurement activities. NERSC has the resources and knowledge to contribute to performance evaluation, as evidenced by many years of operation and procurements. The combination of resources at NERSC and accumulated knowledge and expertise would be beneficial to CS research projects. CS research communities need access to such resources and expertise to help propel research forward.

A logical next step in providing support for CS research would be for NERSC to intersperse large-scale systems procurements with smaller procurements of experimental and evaluation systems. The smaller experimental systems would benefit NERSC users as well as the research communities. While such an approach offers a clear benefit to the general CS research community, there are longer-term benefits to NERSC as a center and DOE computational science projects as well. At any given time, the NERSC user base is to a large degree determined by the capabilities of NERSC's production facilities. NERSC's production facilities are procured on the basis of current workload characterization. Not having access to evaluation/early-release systems forces NERSC into a narrowing spiral where a new system's effectiveness is evaluated on the basis of a limited set of performance benchmarks deemed to be representative of the current user base.

The narrow scope of vision in procurement activities exacerbates the narrowing spiral, which effectively inhibits the emergence of new user communities. For example, the data-intensive computing community remains under-served at NERSC because there are few systems of scale at NERSC that can meet their requirements. Because that community remains small at the present time, it does not offer a significant target for new NERSC procurements. Emerging architectures like the STI Cell processor or FPGA-based computing lack the software tools at this time necessary to make them effective for

[12] http://www.nersc.gov/projects/ipm/

[13] Examples include: the Earth Simulator/vector evaluation project; PERC activities such as APEX-Map; NERSC activities such as IPM and workload characterization project; early architecture evaluation activities like the multi-laboratory X1 evaluation.

production science. Such missing tools will likely originate from the DOE CS research community. Those tools will come to fruition only if experimental systems of considerable scale can be set up side by side with a production scientific computing environment. In many cases, significant infrastructure is needed to support these emerging architectures, and such infrastructure is well outside the reach of individual research programs. Keeping emerging systems in non-production/research facilities will discourage important new technologies and concepts from being used in production environments. The experimental machines also provide a substrate that encourages the development of new and alternative user communities who have requirements that are not governed by the strengths of NERSC's primary computing environment. An increase in available systems, and the research they support, will in turn likely influence decisions for subsequent NERSC procurements.

3.1.4 Data Management Research and Associated Infrastructure

CS research topics in the areas of distributed computing and data management infrastructure are diverse and have a broad impact on nearly all aspects of computational science applications. The topics range from high performance I/O infrastructure, to efficient storage and retrieval of data, to enabling use of multiple, distributed computing resources.

As many computational science projects collect and/or generate data that is then shared among multiple, distributed researchers or research teams, several CS research projects focus on basic or applied research in the area of distributed data management, which is a central pillar in computational and experimental sciences. During early stages of development, computational science researchers benefit from having access to centrally located data repositories, where users and developers alike have access to reference data used to validate new techniques, benchmark applications, and so forth. Such data depots benefit the communities of users that perform analysis, including visualization, on data located at NERSC that is simply too large to move. As science projects evolve, data depots serve large distributed communities who need to perform analysis on community-generated or -gathered data. MDSPlus[14] is an example of a community-centric data storage and access facility used to share data among a distributed team of researchers. Storage Resource Managers (SRM)[15] extend the idea of remotely located storage so that it is location-transparent, and may consist of many individual data stores in multiple locations. Logistical Networking (LN)[16] also includes the data store idea as well as the ability to migrate data between storage caches. Application-level projects, like the Earth System Grid[17] and Particle Physics Collaboratory Pilot,[18] build higher-level services atop the lower-level distributed computing infrastructure to deliver functionality to distributed teams of scientists.

[14] http://www.mdsplus.org/
[15] http://sdm.lbl.gov/projectindividual.php?ProjectID=SRM
[16] http://loci.cs.utk.edu/
[17] http://www.earthsystemgrid.org/
[18] http://www.ppdg.net/

While NERSC has provided some support in the past for data depot projects, such support has been minimal, and falls well short of what is needed to help transition data management infrastructure research into production quality, mature technology. For example, the MDSPlus data repository is arguably the most widely used system for data management, including distributed data management, in the Magnetic Fusion Energy program. It is currently installed at over 30 sites, including NERSC, spread over four continents. At the same time, some NERSC fusion users instead prefer to take advantage of Logistical Networking to store and move simulation data. In both cases, NERSC provided basic support services: a workstation, an operating system with current security patches, and system administration. In both cases, the MDSPlus and Logistical Networking developers could benefit from a deeper interaction with NERSC staff. During early research and development, the MDSPlus servers were set up in an ad hoc fashion in small testbeds by the developers and participating scientists. To move to production status at NERSC, MDSPlus needed a proper security model. The developers of MDSPlus were not equipped to understand the security deficiencies and how to correct them. Although NERSC's security team understood the security policies and found MDSPlus to be security-deficient, there is neither funding nor human resources available to foster the type of relationship between CS researchers and NERSC staff to remedy the situation. As a result, an opportunity to improve MDSPlus and better serve the scientific research community has been missed.

For data that can be moved, current CS research projects that focus on distributed data repositories (MDSPlus, SRM, and LN) and their respective user communities envision the need for existing 1 GB data repositories set aside for testing at NERSC to grow beyond 10 TB in the very near future. These needs will grow within the next few years to well past 100 TB of disk storage and 1 PB in tape resources. In a related matter, while initial code prototypes will be developed under funding to the CS research groups who collaborate with application scientists, in the long term, there is an unmet need for funding for additional developers that will adapt other codes to use those tools. Also, future implementations of infrastructure like Logistical Networking will require the close attention of and interaction with NERSC staff to develop and deploy appropriate policy enforcement mechanisms that are compatible with the site-specific policies of NERSC.

At the file system level, several recent CS research projects have focused on wide-area and high performance parallel file systems. The vision of Lustre[19] is a scalable file system that can serve clusters with tens of thousands of nodes and petabytes of storage, and move hundreds of gigabytes per second. SDSC, IBM, and NCSA demonstrated a wide-area GPFS deployment at SC2003,[20] while SDSC, NCSA, and LLNL demonstrated a wide-area Lustre[21]deployment, also at SC2003.

[19] http://www.lustre.org/
[20] http://www.hoise.com/primeur/03/articles/monthly/AE-PR-12-03-40.html and http://www.teragrid.org/news/apps/0312/hpcwire9.html
[21] http://www.foundrynet.com/about/newsevents/releases/pr11_18_03b.html, and http://www.clusterfs.com/pr/2003-11-12.html

As a programmatic activity, both the CS research and computational science projects hosted at NERSC will benefit greatly from a high performance, global-unified-parallel file system that is accessible by all platforms within the NERSC center. In the near future, the scientific community as a whole will benefit greatly from global file systems that are deployed across major centers in the DOE community. Projects like MADmap[22] are currently I/O bound, and envision the need to store and share upwards of 100 TB among a handful of participating sites using a wide-area file system. The software technology required for a wide-area global unified parallel file system, which is still in early research and development phases, is not likely to emerge naturally from the commercial software/hardware vendors. This means NERSC will need to play an expanded role in wide-area file system testbeds with other centers and research groups. Participation in development and deployment activities will accelerate technology hardening and also provide developers better insight into the needs of HPC centers like NERSC. Without direct NERSC involvement, solutions that emerge for wide-area file systems may prove to be unworkable, either due to NERSC policy requirements that were not well understood by the developers or to workload characteristics that were not well accommodated by the developed solution.

Complementary to file system and shared data resources is the need for effective I/O at the application level. CS research projects in this area include MPI-IO,[23] Parallel HDF5,[24] and pNetCDF.[25] The combination of file system and parallel I/O interfaces produces the bandwidth capacity required by computational science and data-intensive applications. Having "real" parallel file systems is not only helpful for research, but also offers opportunity for synergy between different groups of CS researchers: those working on parallel code interfaces and those developing and deploying parallel file systems. Early deployment, testing and feedback of evolving parallel I/O technologies at NERSC will prove beneficial to the CS research groups as well as to early adopter computational science projects, and help NERSC to better understand requirements of emerging applications.

CS research in file systems would benefit from having access to NERSC's accumulated base of file system knowledge and expertise. NERSC file system characterization activities performed over the past few years represent a wealth of information and data on the performance of real hardware that would benefit CS research. More broadly, there are many areas where the data gathered in order to qualify equipment procurements are invaluable to research efforts, so a coordinated effort to disseminate this information in a coordinated manner would prove beneficial to CS research activities.

3.1.5 Distributed Computing Research and Infrastructure

Another element of distributed computing infrastructure is middleware and tools for designing and creating reusable software for high-performance, parallel applications. The

[22] http://www.lbl.gov/Science-Articles/Archive/sabl/2005/February/planck-satellite-map.html and http://crd.lbl.gov/~cmc/MADmap/doc/
[23] http://www-unix.mcs.anl.gov/romio/
[24] http://hdf.ncsa.uiuc.edu/Parallel_HDF/PHDF5/ph5-status.html
[25] http://www-unix.mcs.anl.gov/parallel-netcdf/

Common Component Architecture (CCA) Forum[26] aims to conceive, develop, and deploy a number of related technologies aimed at providing the infrastructure for creating reusable software components for use in large-scale parallel codes, including distributed computing applications. Higher-level application frameworks, like SCIRun2,[27] support creation of distributed, parallel applications from reusable software components built atop varying types of component interfaces (CCA, CORBA). DiVA[28] focuses on methodology for achieving a high degree of interoperability at all levels of the software stack in high performance, parallel, and distributed visualization and analysis. Both SCIRun2 and DiVA must address remote component invocation, data marshaling, and task scheduling within processors on a single machine as well as between distributed, parallel resources. Distributed computing technologies are predicated on the ability for distributed software components to be launched and to communicate with one another. Furthermore, the distributed resources should be schedulable for testing and debug use as well as for full production use. Researchers and users of such technologies often must navigate through a maze of security and access policies that are different at each site, thereby compounding the complexity of research, early deployment, and production use. While NERSC and other centers can help by migrating towards unified access policies that eliminate these barriers, a program-wide mandate and support for implementing such a mandate will likely be required to achieve the desired results.

During the production phase, parallel codes need direct access from the parallel nodes directly to storage (scratch disk and/or mass storage) and in other cases, other machines on the network. Another need is the ability to access all nodes on the parallel computer from the outside world, that is, an externally visible and accessible IP address for each node. While the absence of global visibility is not a complete showstopper when connecting two distributed components, each consisting of a large set of parallel processes, at least one of the components needs to have IP address visibility. Going through a small set of head nodes is not scalable.

3.1.6 Scientific Visualization

A direct byproduct of the growth in size and speed of computational systems is a proportional growth in size and complexity of the resulting data. Visualization, which is the transformation of abstract scientific data into readily comprehensible images, is the most effective medium for presenting large amounts of complex scientific data to a human for interpretation and understanding. Generally speaking, the visualization techniques that were adequate 10 or 20 years ago, when machines were slower and data smaller and less complex, are often inadequate for use on today's data and on today's computational platforms. The gamut of visualization and computer graphics research aims to address many challenges posed by the need for understanding data produced on the world's largest machines. The topics include algorithmic research and development on parallel and distributed platforms, delivery of visual and analysis content to remote users, effective visual

[26] http://www.cca-forum.org/
[27] http://www.llnl.gov/CASC/calendar/parker.012704.html
[28] http://vis.lbl.gov/Research/DiVA/

presentation, and techniques that reduce data size and/or complexity as well as leverage human intuition to help to "find needles in haystacks."

A fact of life in modern computational science is the distributed nature of the resources and consumers. When data is small, it is simply and easily transferred to the remote location where it undergoes local analysis and visualization. For the largest problems, however, moving data is not an option. Furthermore, it is a much more efficient use of resources to move the analysis close to the data for such problems. Presenting the results of such analysis and visualization to the remote user is the domain of remote and distributed visualization research. Emerging research in remote and distributed visualization relies on the presence of fundamental infrastructure: fast and reliable networks, the ability to log in to remote machines, and so forth.

In nearly all cases, visualization and rendering algorithms substantially benefit from modern graphics hardware. Case in point—software rendering techniques reach asymptotic performance levels of hundreds of thousand triangles per second on modern microprocessors, while a $300 graphics accelerator is capable of throughput levels on the order of 50 million triangles per second—two orders of magnitude difference. Some visualization and graphics research efforts focus on techniques that leverage graphics processing unit (GPU) programs for performing vertex or fragment operations. These approaches are often used for vector and tensor field visualization, where texture/color combining operations are more effectively performed on the GPU than in host memory. Even though centers like NERSC typically do not accommodate local viewing of visualization output (e.g., at a local console or display wall), some visualization and rendering software is quite capable of leveraging graphics hardware to accelerate processing. The resulting images may then be delivered to a remote viewer using one of many different techniques. For example, hardware-based image compositing[29] or tiling systems provide the means for low-latency image combination to support either sort-first[30] or sort-last[31] parallel algorithm architectures, both of which are intended to provide increased capacity and throughput for visual data analysis. Overall, the key idea is that a highly capable visual analysis resources located near the data will serve the analysis needs of computational science projects. As with other areas of CS research, research will benefit from access to such a resource that is located close to the data source but remote from the ultimate user.

Computational platforms are now sufficiently powerful to generate data faster than users can visually inspect the results. Users are simply not capable of viewing all the data in all its detail. Instead, new avenues in visualization and analysis research leverage advances in data analysis and scientific data management to focus analysis on specific subsets

[29] Much work has been done in research (e.g., http://portal.acm.org/citation.cfm?id=827060) and product development (http://www.sgi.com/products/visualization/onyx4/modules.html#compositor) for hardware- and software-based hardware compositing and image manipulation systems.
[30] For example, see http://www.cs.princeton.edu/omnimedia/papers/piwalk.pdf, http://www.r3vis.com/Downloads/OpenRM-Chromium-WhitePaper-July2003.pdf, and http://citeseer.ist.psu.edu/kutluca97imagespace.html.
[31] For example, see http://spire.stanford.edu/raptor/, http://csdl.computer.org/comp/proceedings/pvg/2001/7223/00/72230085abs.htm, and http://www.ccs.lanl.gov/ccs1/projects/Viz/pdfs/99sgexampl.pdf.

within larger data. The result is a reduced load on the end-to-end visualization process, from less visualization processing time to a reduced visual load on the viewer who must interpret and analyze the results. Query-driven visualization,[32] data mining, and visual analytics are all labels for this kind of activity. All are amenable to deployment in either interactive or offline processing modes. Applications in this area will by definition make use of multiple resources, some of which may be geographically distributed. As with other areas of CS research, there is a very real need to be able to use such resources for algorithmic development, testing, debugging, and profiling/performance analysis. Such resources will need to be provisioned for interactive as well as ultimate production use.

In contrast to simulation codes, applications that can be classified as visualization, data analysis, visual analytics, and so forth are often characterized as *data-intensive applications*.[33] The balance of machine characteristics required to support data-intensive applications differs from that required for the leading-edge computing resources at NERSC. Whereas the primary computational platforms tend to be biased towards raw CPU performance, a resource more suitable for data-intensive applications should be biased more towards I/O throughput and large memory. As with earlier statements to the effect that CS research needs access to the large platforms at NERSC as part of the ongoing research efforts, the visualization research community needs access to such platforms that are part of larger facilities in order to conceive, develop, and deploy end-to-end solutions in the environments where they will be ultimately used.

If we assume that the largest production machines will not be used for interactive analysis —continuing the current policy and reflecting a choice to run applications on the most appropriate platform given performance requirements and system balance—the most promising approach will be to first provide a deep, high-bandwidth I/O capability for those large machines. Such an infrastructure will help to move simulation data out as efficiently as possible. On the outside, a production, data-intensive machine that is specifically designed and used for post-simulation analysis and visualization will be brought to bear. These secondary computing resources often need to be of a capacity and performance nearly on the order of the original simulation machine, though perhaps with more emphasis on memory and the bandwidth to parallel I/O systems and archival tape/disk farms. The ASCI roadmap from 1998[34] spells out a Center balance where 10% of the Center's resources (measured in processors or flops) should be dedicated to interactive analysis. The High End Computing Revitalization Task Force (HECRTF)[35] raises that ratio to 25%. In contrast to these recommendations, NERSC's current allocation of resources for data analysis, including visualization, is approximately 1.3% in FY04 and 3.0% in FY05.[36]

[32] http://crd.lbl.gov/~kewu/fastbit/
[33] W. T. Kramer et al., "Deep Scientific Computing Requires Deep Data," IBM J. Res. & Dev. **48**(2), 209 (2004); (http://www.research.ibm.com/journal/rd/482/kramer.pdf).
[34] http://www.llnl.gov/asci-pathforward/pf-desc.html
[35] http://www.house.gov/science/hearings/full04/may13/hecrtf.pdf
[36] Assuming a $10 M/yr budget for computational platforms, NERSC spent about $130 K in FY04 and will spend about $300 K in FY05 on data analysis platforms for its user community.

While size and machine balance are important considerations in data-intensive computing, it is equally important to point out that data-intensive platforms are often used interactively. Interactive use tends to be "bursty," where the platform will transition from being "idle" to "fully loaded" and back to "idle" within the span of a few seconds. Such a load pattern reflects interactive use patterns: generate some visual results, look at the results for a brief period of time, change the parameters to the visualization or analysis software, generate new results, look at new results, etc. While it is tempting to "backfill" these "idle" machines with batch jobs to use the cycles, doing so will have adverse impacts on interactive users. There will be a noticeable delay to the interactive user as batch jobs are paged out of the system. This effect will be more pronounced if batch jobs are paged back in and resumed while an interactive job is running, but temporarily idle during an examination phase. A good compromise is to temporally partition the machine so that periods of time are dedicated exclusively to interactive use, and other periods may be shared by batch and interactive jobs.

3.1.7 Summary of Recommendations and Requirements from CS Research Programs

The high performance networking community has long recognized the gap between the promising prototypes of revolutionary technologies demonstrated by the research community and production networking requirements. In the 2002 NSF/ANIR workshop on Grand Challenges in E-Science[37] and the subsequent DOE report on the future of DOE Scientific Networking,[38] leading members of the scientific networking community grappled with this issue by created distinct definitions for research, experimental, and production networks, where the experimental networks provide the missing link that supports examination of the most promising research concepts at a production scale and thereby forms a more robust conduit between innovative research and production. This conduit has fostered the rapid emergence of advanced networking services that exploit user-controlled switched light-paths—concepts that had been formerly marooned in research laboratories due to uncertainties about their feasibility and robustness for large-scale deployment. A similar situation exists in the area of high performance computing. The missing link between CS research and NERSC's production supercomputing capabilities is an expanded role to support experimental systems that foster closer ties between research and production. There is an element of risk involved in increasing support for immature CS technology—the success of any given technology is not pre-ordained. But the benefits for advancing technology far outweigh these costs.

Proactive Deployment of CS Research

Become a part of the research cycle rather than a consumer of finished, mature products. CS research is an evolutionary process that requires synergistic relationships with consumers (computational science applications) and the centers where the technology is ultimately deployed. Without direct NERSC involvement, solutions that emerge from CS research may prove to be unworkable, either due to NERSC policy requirements that

[37] http://www.calit2.net/events/2002/nsf/ExpInfostructure_Irvine_FINAL_110502.pdf
[38] W. E. Johnston et al., "A Vision for DOE Scientific Networking Driven by High Impact Science," March 15, 2002 (http://www.sc.doe.gov/ascr/high-performance_networks.pdf).

were not well understood by the developers or to workload characteristics that are not well accommodated by the solution. NERSC should be proactive about deploying, with the aid of CS researchers, new technologies before they reach maturity so that the research community has an opportunity to test and harden those systems. Waiting for CS research technology to mature before production deployment will not help it to mature, and such delays effectively act as an impediment to the evolution of computer and computational science. To achieve this objective, NERSC may consider extending the same "red carpet" treatment afforded to the largest science projects to CS research.

Increase program focus on deploying results from CS research on NERSC platforms. As CS research technologies mature, NERSC should be proactive about deploying these new technologies developed for its user community on its machines. The Center should be a delivery vehicle for CS research products to the ultimate beneficiaries—the computational science projects hosted at NERSC. Deployment of these technologies will require adaptation and tuning by NERSC staff, possibly with the involvement of the CS researchers and developers. Deployment may include providing consulting services to users to help adapt technologies into their applications and workflows. Such deployment and technology transfer from research into production is an activity that is, generally speaking, not funded under CS research projects.

Expand Program Focus to Support CS Research

Embrace a flexible and adaptive approach to system configuration. Unlike many production activities that may work well when the center uses a "set and forget" approach to system configuration, many CS research programs require variation in system configuration parameters for the purpose of research, development, analysis, testing, and related pre-production activities. While the CS researcher does not necessarily require "root access," having the ability to work closely with NERSC staff who can be perform such configurations will be immensely useful.

Embrace a more flexible and accommodating policy for investigative runs, even on large numbers of processors. Testing, debugging, and performance analysis of parallel programming models at large scale often require interactive access for periods of time because many debugging activities cannot be conducted in batch mode.

Provision systems for debug, test, and analysis runs. With debug queue wait times currently in excess of six hours at the time of this writing, the present approach to provisioning is not viable for CS research. In many cases, debug/test runs are not computationally intensive, but sometimes require the use of large numbers of nodes for relatively short periods of time. Another distinctive need of this type of work is access to platforms for routine regression testing. This is critical to insure the quality of the software being developed, and as testing suites increase in size, the computational requirements grow, and it needs done frequently. Many projects run regression tests on a daily basis. Such patterns do not fit well in the model of computation currently used as the basis for NERSC allocation requests.

Increase dialogue between NERSC experts and CS researchers. Parallel language researchers in particular can benefit from a close relationship with NERSC consultants so

that architectural details and performance analysis details of the production computing systems are available to CS research.

Break the Narrowing Spiral that Inhibits Innovation

Broaden the vision for procurements. When procurement evaluations and requirements for new platforms are based on current user requirements, the result is a platform that at best serves the need of established users bases. It takes a while for CS tools to catch up with hardware and architectural innovations, so operating innovative or new hardware/ software architectures is essential for fostering new user communities and accelerating innovation in computing for the domain sciences. These new communities have little voice in new procurements until the software methods for properly exploiting new technology enable them to reach critical mass. Reaching critical mass requires leading-edge CS research on innovative systems that are operated at a scale that is relevant for evaluating high-end science applications. However, the CS research community is not equipped to set up and front such systems. Likewise, the research community is not equipped to provide user management systems like NIM in order to manage any kind of user community.

NERSC offers the most appropriate venue for operating new technologies at a larger experimental scale because it leverages core capabilities in systems infrastructure, expertise, and user management that are required for NERSC's production operations. The CS research community would benefit greatly if NERSC could open up the procurement process to new user and research communities by using smaller procurements to provide early access for CS researchers and new user communities so they can become established. A logical next step in providing support for CS research would be for NERSC to intersperse the procurements of large-scale systems with smaller-scale procurements of experimental/evaluation systems that are made available to NERSC users as a service to both the research and the computing science communities.

Procure and provide emerging architectures and evaluation systems. The experimental machines provide a substrate that encourages the development of new/alternative user communities who have requirements that are not governed by the strengths of NERSC's primary computing environment. Early progress on emerging platforms by new user communities may have a positive, long-term impact on subsequent NERSC procurement choices.

Facilities Requirements

Provide capable systems suitable for data-intensive computing research and production. Unlike production computational science, where the emphasis is typically on flops/byte, the needs of data-intensive algorithmic research are different, placing the architectural balance point more in the direction of greater amounts of memory visible to the application and greater I/O capacity.

Strive for better architectural balance within the center. The ASCI roadmap called for a center balance where data analysis platforms accounted for about 10% of center resources; the HECRTF calls for a ratio of about 25%. The reasoning is to prevent another "Earth Simulator data bottleneck" as well as starvation of data analysis infrastructure. If

analysis resources do not meet the needs of the computational science community, they will not be utilized because they are ineffective.

Increase capacity and capability of secondary storage to better support research and production activities in both CS and computational science. CS research and computational science projects have identified near-term needs for tens of terabytes of spinning disk, with hundreds of terabytes of disk and more than a petabyte of tape storage within the next few years.

Increase I/O capability on the largest production platforms. If we assume that the largest production machines will not be used for interactive analysis—continuing the current policy and reflecting a choice to run applications on the most appropriate platform given performance requirements and system balance—the most promising approach will be to first provide a deep, high-bandwidth I/O capability for those large machines.

Local Resource Requirements

Provide high-quality parallel debugging systems. Useful software tools like POE+ and IPM were created by NERSC consultants who have an in-depth knowledge of the hardware. These kinds of activities should be encouraged and supported.

Increase the architectural diversity of resources at NERSC. A theme common to many domains of CS research is the need for access to a broad range of architectures, and at a scale larger than is possible in an individual research program. The CS research community depends on centralized computing facilities like NERSC to provide the resources needed to support research objectives.

Leverage NERSC Expertise and Stature to Strengthen CS Research Programs

Leverage NERSC relationships with vendors for CS research. NERSC can play a role by helping language and library developers to establish closer relationships with vendors who supply production systems in order to get early access to specifications for new systems, and to help drive the design process to better accommodate emerging concepts in parallel programming exposed by new languages and libraries.

Provide system details and documentation to the CS research community. To achieve a high level of success and a broad impact, NERSC should play a key role in the deployment and development of advanced technologies. This is a two-way street: NERSC needs to provide access to systems, system details, and documentation to researchers that they otherwise might not have access to via public channels—NERSC is and should be a resource center.

3.2 Office of Basic Energy Sciences

The mission of the Basic Energy Sciences (BES) program is to foster and support fundamental research in the natural sciences and engineering to provide a basis for new and improved energy technologies and for understanding and mitigating the environmental impacts of energy use. The program features many different types of research and simula-

tions, ranging from highly accurate atomistic calculations of small molecular systems, to large molecular dynamics simulations of solids and biological systems, to very large direct numerical simulations of flames. Utilization of the NERSC computational resources has a direct impact on the progress of the mission-critical objectives of four major BES subprograms: Materials Science, Chemical Sciences, Geosciences, and Energy Biosciences. As the computations in each of these areas evolve to larger systems with greater accuracy, the lines between each of these subprograms blur. For example, as material surfaces become more accurate and defects are well defined, chemical reactions on those surfaces become computationally feasible topics of research. Even with the many advances that have been made in the algorithms used in these simulations to improve the scalability and order of magnitude of the computations, each one of the computational areas will benefit from much larger computational resources.

Across BES, the balance between capacity and capability computational resources varies between disciplines from 1:1 to perhaps 5:1 (capacity:capability). The two most oft-cited issues for future resources are more capacity resources and longer-running jobs of all sizes. In the following discussion, we highlight recent BES INCITE allocations, and then present specific examples and associated computational challenges in several of the subprograms.

3.2.1 Innovative and Novel Computational Impact on Theory and Experiment (INCITE)

INCITE is a competitive program designed to support a small number of computationally intensive, large-scale research projects that can make high-impact scientific advances through the use of a substantial allocation of computer time and data storage at the NERSC Center. Three BES projects have been awarded multi-million-hour allocations of computer time under this program, two in 2004 and one in 2005. The vision and accomplishments of these projects hint at what would be possible when it becomes feasible to make far larger allocations of computer time across the BES portfolio.

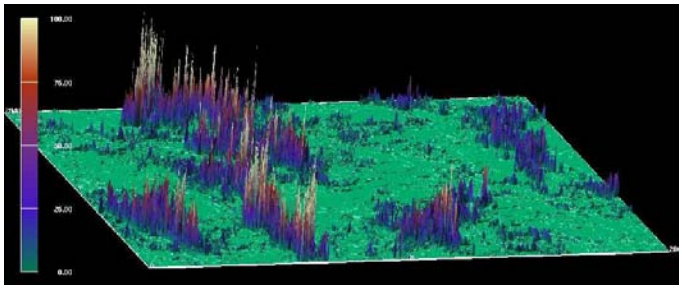

Figure 1. Localized peaks in a selected 2048^3 grid plane indicate intermittent behavior of the contaminant dissipation rate.

Fluid Turbulence and Mixing at High Reynolds Number. Turbulent mixing is the process by which different substances are brought into close contact by fluid motions in important applications such as combustion and pollutant dispersion in the environment. The primary scientific problem is to quantify the statistical and scale similarity properties of contaminant fluctuations in three-dimensional space, while considering the properties of the transporting velocity field. Advances in this field require data of great detail, for which the ideal approach is direct numerical simulation, in which all scales are computed according to exact governing equations. Using three national parallel supercom-

puter sites, this team has simulated turbulent mixing at the highest resolution (8 billion grid points) and highest Reynolds number to date in the U.S. For the first time, the simulations have been performed with a sufficiently wide range of scales to show good agreement with classical scaling. The dissipation rate of the contaminant concentration shows highly intermittent and sheet-like structures (Figure 1), which is important for combustion phenomena such as flame extinction and can be described using multifractal concepts.

Quantum Monte Carlo Study of Photoprotection in Photosynthesis. Photosynthesis by plants and bacteria is the world's foremost means of carbon sequestration, which has enormous implications for climate change and global warming. Photosynthesis is also an example of fundamental electron chemistry and is an efficient energy transfer system—processes which are fundamental in many areas of scientific research. A primary topic of study is a mechanism that protects plants from absorbing more solar energy than they can immediately utilize, and, as a result, suffering from oxidation damage. The calculations performed at NERSC are the largest all-electron quantum Monte Carlo (QMC) calculations on biological systems carried out to date. A sparse linear-scaling QMC algorithm was developed and applied to these systems. Electron densities obtained from this study (Figure 2) will be compared to results from density functional theory.

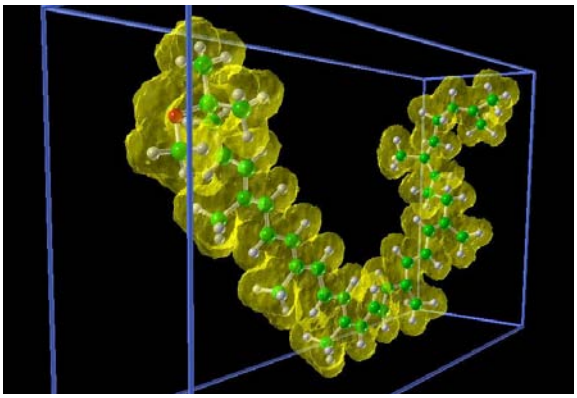

Figure 2. Electron density of spheroidene, a photoprotective carotenoid obtained using the quantum Monte Carlo method.

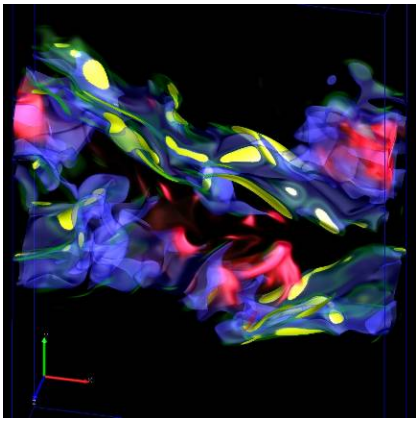

Figure 3. A snapshot of hydroxyl radical (red), local mixing rate (yellow), and stoichiometric mixture fraction (blue) in a turbulent non-premixed CO/H_2/air flame from a test run.

Direct Numerical Simulation of Turbulent Non-Premixed Combustion: Fundamental Insights towards Predictive Modeling. Detailed three-dimensional combustion simulations will be performed of flames in which fuel and oxygen are not premixed (Figure 3). The aim is to gain insight into reducing pollutants and increasing efficiency in combustion devices, for instance, in jet aircraft engines where fuel and oxidizers are not premixed for safety reasons, and in direct-injection internal combustion engines. These simulations will be the first-ever 3D direct numerical simulations with detailed chemistry of a fully developed turbulent, non-premixed flame, and will provide fundamental insight into the dynamics of extinction and reignition in turbulent flames.

3.2.2 Materials Science

The objective of the Materials Science program is to increase the understanding of phenomena and properties important to materials behavior that will contribute to meeting the needs of present and future energy technologies. It is composed of the subfields metallurgy, ceramics, solid state physics, materials chemistry, and related disciplines where the emphasis is on the science of materials. Two of the challenges facing materials science today are performing computations that cross length and time scales, including atomistic to nanoscale to macroscale computations, and creating materials by design. The latter includes the construction of multi-component molecular electronic and mechanical devices having desired properties, the control of spin as a basis of future computational devices, and the design of catalysts from first principles. These are immense and nationally important challenge areas that will require much theoretical and algorithmic development, and that also require much larger computational resources both to examine the physical systems and to describe the reactions or properties of those systems.

The NERSC supercomputers are being used to study quantum dots), enabling the first-ever quantitative analyses of the formation and stability of Ge and SiGe quantum dots (QD) on Si substrates, and, also for the first time, predicting the critical size for Ge QD nucleation or formation in very good agreement with experiments. The continued miniaturization of electronic and optoelectronic devices demands nanoscale structures whose fabrication with the desired uniformity is challenging because our fundamental understanding of the physical and dynamical processes at the nanometer scale is still limited. For self-assembly of nanostructures, forces acting at the nanometer range have to be identified. Such a force is the elastic interaction in solid materials, and devices are often made of layered structures of thin films with strain arising from lattice mismatch between layers. Using Ge dots on Si(001) as a model system, extensive first-principles calculations are being performed of surface energies, stresses, and diffusion barriers as well as their strain dependence of substrate, wetting layer, and QD surfaces (Figures 4 and 5).

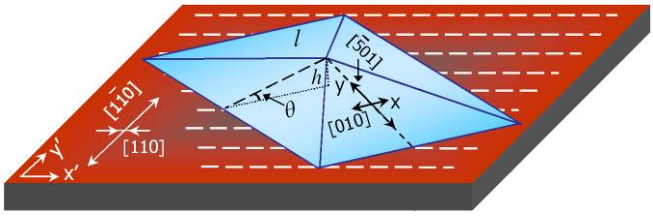

Figure 4. Schematics of a square-based Ge hut on Si(001). Arrowed crosses indicate surface stresses. Two types of Ge QDs form on the Si(001) surface. The smaller one, called "hut," forms at the early stage of growth and is bounded by {105} facets. The larger one, called "dome," forms at the later stage of growth and is bounded by {113} and other high-index facets.

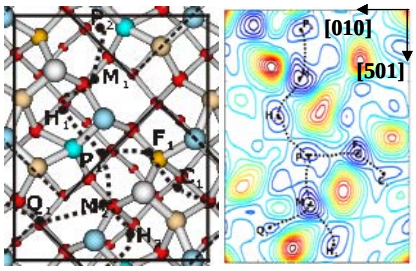

Figure 5. (a) Top view of the unit cell of Ge(105)-(2x1) surface. Higher atoms are indicated by larger spheres. (b) Contour plot of Ge adatom potential energy surface. The letters and dashed lines indicate the main binding sites and low-barrier diffusion paths.

By combining these calculations with continuum modeling, researchers are able to perform, for the first time, quantitative analyses of formation and stability of Ge and SiGe QDs on Si substrate. Related studies have also shown that on a compressive surface, Ge diffusion is 10^2–10^3 times faster than Si, while on a ten-

sile surface, they are comparable. Consequently, growth of a compressive SiGe film is rather different from that of a tensile film. The calculations employ a pseudopotential plane-wave total-energy method modeling the surfaces as a slab. Usually these calculations involve 100 to 200 atoms and ~100 plane-wave basis functions per atom. The MPI parallel computational code runs with good efficiency and scaling up to 32 nodes. Future simulations of more complex systems will demand many more calculations using larger slabs and therefore will be much more computationally demanding. Given the limited scalability of the computational scheme, an enabling machine characteristic would be more memory per processor, as well as much greater permitted runtimes.

Exploiting newly developed codes implemented on the NERSC high performance computers, the internal structure and angular tilt of the magnetic moment of a cobalt nanowire next to a platinum surface step has been predicted by first-principles theory. Due to the need for ever higher density recording media and smaller sensors, atomic-scale magnetic devices are the focus of much experimental and theoretical research. The accurate design of nanoscale magnets requires a detailed understanding of the relevant physical properties that can only be provided by first-principles, parameter-free quantum mechanical approaches. The Co-chain and the Pt-step were modeled as an extended impurity embedded in a trough in an otherwise perfect semi-infinite Pt-surface, and constrained density functional theory was used to calculate the magnetic state at each time step in the evolution towards equilibrium (Figure 6). The calculations were fully relativistic so that the large exchange interactions and much smaller magnetocrystalline anisotropy energy were treated on the same first-principles footing. In addition to this theory being in quantitative agreement with experiment, static "magnetic force theorem" calculations of the Co-chain magnetic moment were performed to demonstrate how a reduction of symmetry leads to canted magnetic state having a very non-intuitive canting angle.

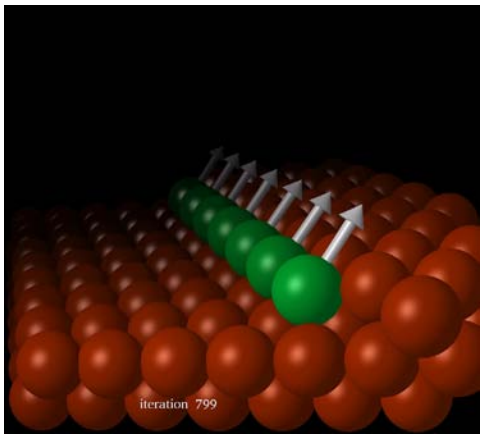

Figure 6. After 800 time steps, the system relaxes to a canting angle of 42 degrees.

This work shows it is now possible to accurately simulate many properties of experimentally realizable nanostructures. Moreover, unlike experiment, the theoretical model can reveal the underlying physical mechanisms that determine the magnetic state. Using current codes and computers, it is still only possible to treat nanostructures comprising a few hundred atoms. However, this work clearly indicates that with continued theoretical development and more powerful computers, direct first-principles simulation will play a key role in development of complex magnetic nanostructures.

First-principles quantum-mechanical calculations at NERSC of the diffusion of transition metal (TM) solutes in nickel are conceptually challenging the commonly accepted description of solute atom diffusion rates in metallic lattices. Due to consideration of

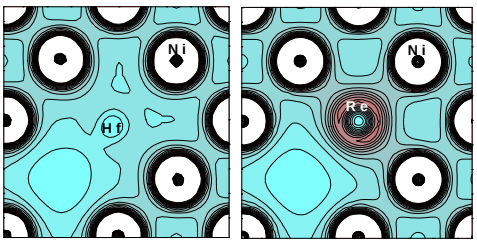

Figure 7. Directional bonds are apparent in these electron density maps of Hf and Re in nickel.

lattice strains induced by diffusing atoms, it is commonly believed (and clearly stated in the standard textbooks) that the larger the solute atom, the slower its diffusion rate. Simulations discover exactly the opposite: larger atoms, in fact, move *much faster* than smaller atoms. In other words, for the 4*d* and 5*d* transition metal rows, diffusion rates are the slowest for the mid-row solutes, e.g., Ru and Re, despite their having atomic sizes which are the smallest and closest to nickel. The calculations reveal that this is due to the existence of a significant diffusion energy barrier for smaller solutes. The origin of this counterintuitive diffusion trend lies in the electronic bonding characteristics of solute atoms: the development of bonding directionality (for elements such as Ru and Re) leads to smaller atomic radii and incompressible solute-host bonds, which hinder solute diffusion (Figure 7).

3.2.3 Chemical Sciences

The Chemical Sciences subprogram supports a major portion of the nation's fundamental research in the chemical sciences. The research covers areas of the chemical sciences that impact the DOE's energy and environmental missions. Research supported covers atomic, molecular, and optical (AMO) science; chemical physics; photo- and radiation chemistry; surface chemistry and heterogeneous catalysis; organometallic chemistry and homogeneous catalysis; analytical and separation science; heavy element chemistry; and aspects of chemical engineering sciences. This research provides a foundation for fundamental understanding of the interactions of atoms, molecules, and ions with photons and electrons; the making and breaking of chemical bonds in gas phase, in solutions, at interfaces, and on surfaces; and the energy transfer processes within and between molecules.

Chemical Sciences faces several challenges that will require orders-of-magnitude increase in computational capabilities and considerable effort in algorithm development. As in the Materials Sciences program, some challenges include crossing multiple length and timescales (e.g., taking data that is produced at the atomic or molecular level and applying it to larger domains such as nanoscale). Other challenges are predicting and controlling chemical reactivity (e.g., producing the desired product in a quantitative yield without many side products), and designing chemical systems with desired properties. Catalyst design from firs principles is just now becoming feasible and exemplifies how expanded capabilities in large-scale computing will enable significant scientific progress, relevant not only to the chemicals and pharmaceuticals industries, but also for energy-production related applications such as the President's Hydrogen Fuel Initiative. The newly established Catalysis Collaboratory includes "catalysis informatics," i.e., substructure-property relationships developed from first principles calculations supported with experimental information. Expanded computational resources will lead to chemical accuracy for metal-ligand energetics; improved modeling of heterogeneous systems; development, testing, and implementation of approaches for modeling the many different spec-

troscopies that experimentalists use to characterize catalysts; and new computational tools such as the linear scaling stochastic inversion method for *ab initio* Monte Carlo that appears promising for systems with 500+ atoms.

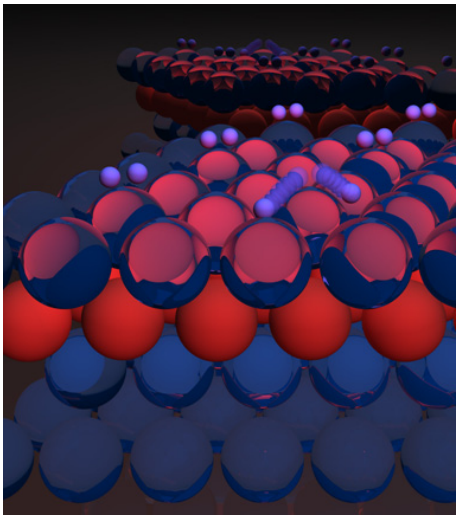

Figure 8. H_2 dissociation path on a bimetallic NSA surface.

A new class of bimetallic alloys that exhibit superior catalytic behavior for hydrogen-related reactions have been identified by simulations on the NERSC supercomputers. The rational design of heterogeneous catalysts from fundamental principles has the potential to yield new materials with improved activity and selectivity. The new class of materials has been termed near-surface alloys (NSAs) and is comprised of alloys wherein a solute metal is present near the surface of a host metal in concentrations different from the bulk. State-of-the-art, periodic, self-consistent density functional theory-generalized gradient approximation (DFT-GGA) calculations were used to predict promising NSA candidates for detailed experimental studies. Using simple but realistic nanoscale models of subsurface alloys and overlayer structures (Figure 8), the stability of a large class of NSAs were evaluated in H-rich environments. NSAs lying outside the hatched regions of Figure 9 are expected to retain their near-surface structure in H-rich environments.

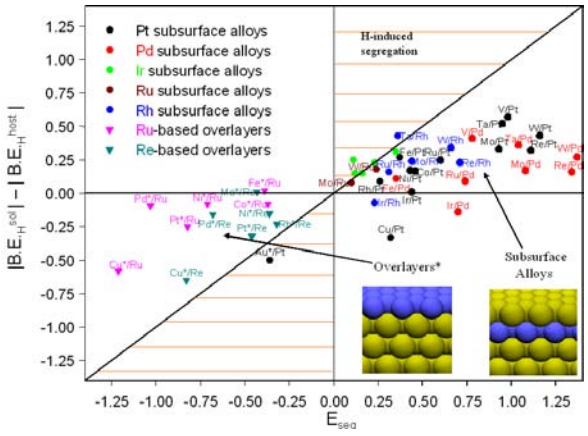

Figure 9. Identifying bimetallic NSAs stable in H-rich environments.

Certain of these stable NSAs were found to possess a remarkable attribute: they bind atomic hydrogen (H) as weakly as the noble metals (Cu, Au) while, at the same time, dissociating H_2 much more easily. This unique combination of properties is expected to permit these alloys to serve as low-temperature, highly selective catalysts for a variety of reactions involving hydrogen transfer, as robust catalysts for enhancing the uptake and release of hydrogen on light metal hydrides, and as poison-resistant electrocatalysts for low-temperature fuel cells. Extending this general catalyst identification methodology to more complicated reactions requires the availability of much more powerful computational resources.

Many-body methods are being used to study silicon surface reconstruction. When cleaved along the (100) direction, the silicon surface undergoes surface reconstruction: the silicon atoms on the surface form bonds that involve two Si atoms (dimers) that enable a favorable sigma bond. Whether these surface dimers are symmetric or buckled is

as yet unanswered due to the very flat potential energy surface of Si(100) along buckling coordinates. Previous results from DFT are inconclusive, and for a more accurate prediction of the relative energies between the two minima, inclusion of both dynamic and static (or near-degeneracy) electron correlation is critical. Multi-reference Møller-Plesset perturbation theory (MRMP2) was employed on the DFT minimum energy structures. MRMP2 predicts the symmetric structure to be more stable than the buckled structure by 4.1 kcal/mol per dimer using a threee-dimer cluster model for Si(100) (Figure 10).

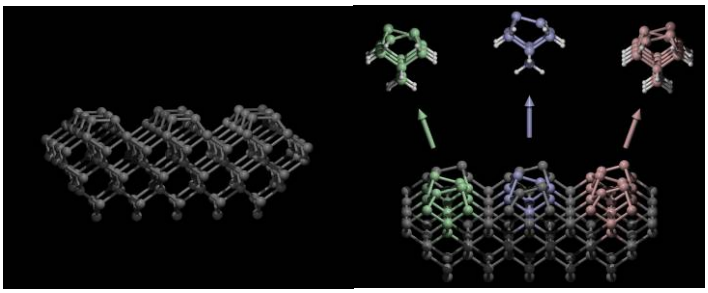

Figure 10. Silicon 1,1,0 surface conformations. The work of Jung et al. identifies the symmetric structure of the silicon (100) surface displayed on the left as the minimum in the potential energy surface, when compared to the buckled structures on the right. (Graphics by Jeffrey T. Frey)

Establishing the preferred ordering of the symmetric and buckled dimers from electronic structure calculations still requires additional work. First, it remains to be seen whether MRMP2 geometry optimization (which is infeasible today due to high computational cost) would also predict two different local minima, like DFT, or just a single minimum. If MRMP2 does indeed predict the existence of two local minima, what is the barrier height for the conversion at this level of theory? Addressing both of these questions as well as the effects of cluster size and basis set for converged results demands higher computational power.

NERSC resources are being used to develop a molecular-level understanding of vertebrate vision. Rhodopsin is a prototypical G protein coupled receptor (GPCR), responsible for turning on the signaling transmission cascade in the vertebrate vision process. A protein backbone binds covalently the retinyl chromophore where photochemistry, phototransduction energy redistribution, excited-state curve-crossing non-adiabatic dynamics, and proton transfer are facilitated and regulated by the protein environment. Computations at NERSC aim to develop a molecular-level understanding of the assembly and function of the individual residues and bound-water molecules that form the rhodopsin active site. Recently progress has been made toward these long-term objectives by developing computational studies of the primary photochemical event, based on state-of-the-art quantum-mechanics/molecular-mechanics (QM/MM) hybrid methods and quantum molecular dynamics simulations (Figure 11). The ONIOM-EE (TD-B3LYP/6-31G*//B3LYP/6-31G*:Amber) level of theory predicts S0 to S1 electronic-excitation energies for both 11-*cis* rhodopsin and all-*trans* bathorhodopsin (59.4 kcal/mol and 59.0 kcal/mol, respec-

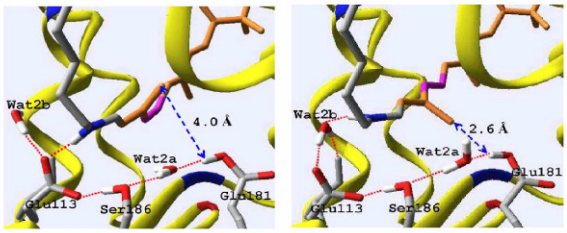

Figure 11. The reorganization of hydrogen bonds (dashed red lines) due to isomerization of 11-*cis* rhodopsin (left panel) to all-*trans* bathorhodopsin (right panel) as predicted by calculations performed at NERSC.

tively) in very good agreement with the corresponding experimental values (57.4 kcal/mol and 54.0 kcal/mol, respectively). Due to the importance of electrostatic contributions to the total chromophore-protein interaction, further progress in the description of the system requires addressing the influence of the self-consistent polarization of the protein, which significantly increases the computational expense.

Ultrafast photo-induced electron transfer processes in the condensed phase are being studied on the NERSC computers using a novel hybrid method for quantum dynamics. New rigorous theoretical methods and practical computational techniques have been developed that can provide a quantitative description of quantum chemical reaction dynamics for ultrafast photo-induced electron transfer processes in the condensed phase. Novel applications of these approaches are carried out to study charge and energy transfer processes in important chemical/biological systems, and the relation of their dynamical aspects to experimental femtosecond nonlinear spectroscopy. Due to the complexity of the problem, these computational studies can only be performed on the large parallel platforms at NERSC. The self-consistent hybrid method has been implemented efficiently on NERSC computers with excellent parallel scalability. Combined with other dynamical quantities such as the electronic population dynamics and the wave packet motion of the nuclear degrees of freedom, such simulations can provide rigorous theoretical interpretation to experimental time-resolved spectra for charge transfer reactions in complex systems.

Current calculations study photo-induced electron injection reactions from electronically excited states of dye molecules to the conduction band of semiconductor substrates (Figure 12). This process represents a key step of photonic energy conversion in nanocrystalline solar cells and is thus crucial to the practical design of molecular photovoltaics. Research aims at offering quantitative comparisons between theories and optical pump-probe spectroscopy injection processes, thus providing fundamental understanding to the quantum coherence effect that has been observed for these ultrafast processes. Due to the complexity and high accuracy required, the current computation is quite expensive, especially for the non-perturbative simulation of the time-resolved nonlinear spectra. The study of a broader class of physical problems or dynamical studies on more sophisticated models of charge and energy transfer reactions in the condensed phase demand much greater computational resources. For instance, rigorous results could be computed for the recent two-dimensional echo experiments.

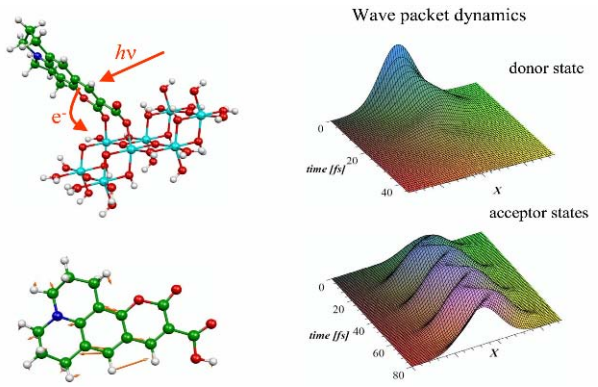

Figure 12. *Ab initio* dynamics of photoinduced electron transfer from the Coumarin 343 molecule to a TiO_2 surface.

NERSC supercomputers are being used to study combustion processes in homogeneous charge compression ignition (HCCI) engines, which are being considered as an

alternative to diesel engines. By exploiting a lean intake charge that is well mixed prior to combustion, HCCI engines may provide efficiency gains over spark-ignited engines and lower emissions compared with compression ignition engines. The influence of thermal stratification on autoignition at constant volume and high pressure was studied by direct numerical simulation (DNS) with complex hydrogen/air chemistry. In particular, the study identified the conditions for different regimes of ignition front propagation and determined the boundaries of validity for a multi-zone engine modeling approach. Figure 13 shows the nature of combustion for conditions near top dead center in an engine for different initial amplitudes of the temperature fluctuation. It is obvious that lower temperature fluctuations result in a more homogeneous, volumetric combustion process, while larger values tend to promote front-like structures. Figure 14 shows a comparison of this DNS data with corresponding multi-zone model predictions. It is noted that the comparison at low levels of fluctuation is very good, while the comparison deteriorates at higher levels. This is attributed to the increasing significance of flame propagation, and therefore diffusive transport, which the model does not take into account.

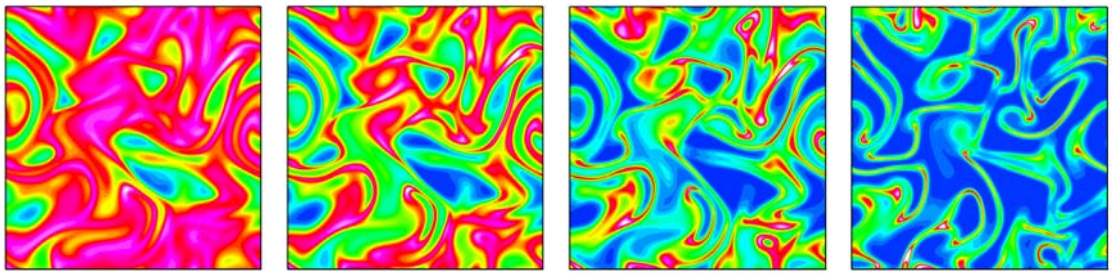

Figure 13. Normalized heat release isocontours (rainbow color scale: white denotes maximum heat release) for increasing temperature fluctuations from left to right (T = 3.75 K, 7.5 K, 15.0 K, and 30.0 K).

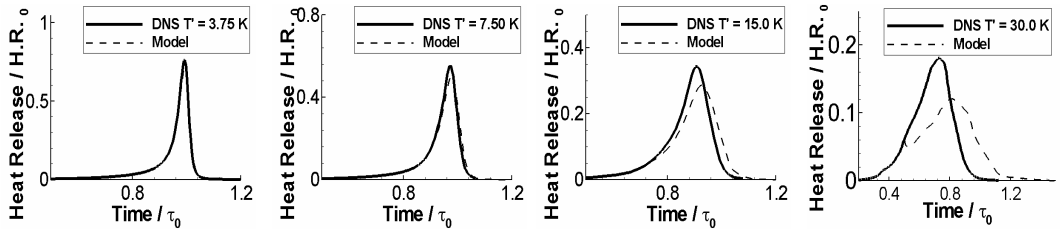

Figure 14. Normalized mean heat release versus time for the DNS and the multi-zone model for increasing temperature fluctuations from left to right (T = 3.75 K, 7.5 K, 15.0 K, and 30.0 K).

3.2.4 Geosciences

The Geosciences program supports research aimed at developing an understanding of fundamental Earth processes that can be used as a foundation for efficient, effective, and environmentally sound use of energy resources, and providing an improved scientific basis for advanced energy and environmental technologies. In the Geosciences program, the computational challenges involve many of those described for the Materials and Chemical Sciences, but also include some that are very specific to surface phenomena. A few examples of the types of challenges faced are:

- Accurate examination of the surface properties of clays and minerals with and without absorbates. This information can be used in less accurate models to improve their accuracy and to allow for larger system sizes to be examined.
- Longer simulations to model phenomena too slow to be observed with current compute resources. This includes phenomena that occur on the second timescale instead of the nanosecond timescale, but also phenomena that occur on 100-year timescales instead of 10-year timescales.
- Subsurface ground transport with the inclusion of chemical interactions, which affect the characteristics of the flow.
  - Current calculations include very simplified chemistry models that do not represent the realistic situation because the computational expense of including more accurate models is simply too great. With a couple of orders of magnitude increase in computational resources, the chemistry could be modeled much more realistically.
  - Larger sections of the Earth could be examined with a larger number of media (both porous and nonporous) in the simulation.
  - Finer grids could be used to get increased information at a smaller scale.

NERSC supercomputers are being used to predict the atomic structures of environmental nanoparticles that play important roles in global geochemical cycles. Montmorillonite is a naturally occurring layered mineral with the atomic structure of mica. Substitution of Al for Si and Mg for Al in the mineral results in a negative structural charge (0.5–1.2 electron charges per formula unit) that is balanced by cations residing on its basal planes of oxygen atoms, giving the balanced chemical formula $M^{+}_{x+y}[Si_{8-x}Al_x][Al_{4-y}Mg_y]O_{20}(OH)_4$, where x is the substitution of Al for Si, as in mica, y is the substitution of Mg for Al, and M^{+} is a balancing cation on the basal plane (e.g., Na^{+}). Montmorillonite, widely used in industrial catalytic applications, water-based drilling fluids for oil production, and drug delivery, is ubiquitous in soils and sediments of temperate zones. Its high surface reactivity with metal cations and organic molecules makes it one of the key environmental nanoparticles that help govern the geochemical cycles of metals and carbon in soils, water bodies, and the atmosphere. Despite this importance, an accurate structure of montmorillonite is not yet known because it is not possible to prepare a single crystal of this mineral. Quantum mechanical structural calculations based on DFT can assist powder X-ray diffraction methods in determining the structure. Total-energy optimizations with ultrasoft pseudopoten-

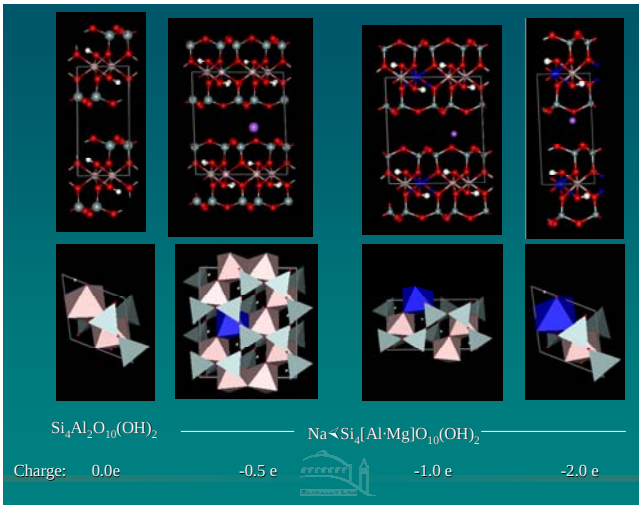

Figure 15. Total-energy optimizations were used to calculate Na-montmorillonite structures with varying charge.

tials and plane-wave basis functions were performed using 64 processors to calculate Na-montmorillonite structures with varying charge (Figure 15), thus obtaining complete crystallographic information without adjustable parameters. The DFT results were consistent with experimental X-ray diffraction data on montmorillonites with 0.0 e, –0.5 e, and –1.0 e structural charge, although predicted interatomic distances were 2–3 % too large on average, a known shortcoming of DFT simulations of minerals.

3.2.5 Energy Biosciences

The Energy Biosciences program supports fundamental research needed to develop future biotechnologies related to energy. The supported research focuses on the biological mechanisms occurring in plants and microorganisms. Plants and microbes fit readily into the energy context by virtue of serving as renewable resources for fuel and other fossil resource substitutes, as vehicles to restore previously disrupted environmental sites, and as potential components of industrial processes to produce new products and chemicals in an environmentally benign manner. Most of the computations at NERSC for this program involve large-scale molecular mechanics/dynamics to examine biomaterial involved in energy conversion (e.g., photosynthesis). The challenges here are similar to those in the other programs: larger physical systems with greater accuracy and longer timescales.

3.3 Office of Biological and Environmental Research

The mission of the OBER is to develop the knowledge needed to identify, understand, and anticipate the long-term health and environmental consequences of energy production, development, and use. This mission is carried out through research at DOE national laboratories, universities, and private institutions in the general areas of Life Sciences, Medical Sciences, and Environmental Sciences. Large-scale computing plays an important role in Life Sciences through Human Genome, Structural Biology, and Microbial Genome research. In close cooperation with improved technologies, NERSC computational capabilities are focused on genome mapping, sequencing, and information management; characterizing the molecular nature of the human genome; and determining protein structures and functions. In addition, genomes of microbes are being investigated for potential relevance to energy production/use, bioremediation, and global climate. Building on these successes, the Genomes to Life program will take the next step: to understand the composition and function of the biochemical networks and pathways that carry out the essential processes of living organisms.

DOE/OBER Global Change activities include the process research and modeling efforts needed to (1) improve understanding of factors affecting the Earth's radiant-energy balance; (2) predict accurately any global and regional climate change induced by increasing atmospheric concentrations of greenhouse gases; (3) quantify sources and sinks of energy-related greenhouse gases, especially carbon dioxide; and (4) improve the scientific basis for assessing the potential consequences of climatic changes, including the potential ecological, social, and economic implications of human-induced climatic changes caused by increases in greenhouse gases in the atmosphere, and the benefits and costs of alternative response options. Research is focused on understanding the basic chemical,

physical, and biological processes of the Earth's atmosphere, land, and oceans, and how these processes may be affected by energy production and use, primarily the emission of carbon dioxide from fossil fuel combustion. A major part of the research is designed to provide the data that will enable an objective assessment of the potential for and consequences of global warming. The program is comprehensive, with emphasis on the radiation balance from the surface of the Earth to the top of the atmosphere, including the role of clouds, and on enhancing the quantitative models necessary to predict possible climate change at the global and regional levels.

3.3.1 Climate Change Simulations

Substantial support from NERSC resources and staff enabled the Climate Change and Prediction (CCP) group at the National Center for Atmospheric Research (NCAR) to carry out extensive climate change simulations with the Community Climate System Model, Version 3 (CCSM3) and the Parallel Climate Model (PCM) in support of the U.S. submission to the Intergovernmental Panel on Climate Change Fourth Assessment Report (IPCC AR4). These simulations included developing testing runs for CCSM3; multi-member ensemble Special Report on Emissions Scenarios (SRES) runs with PCM and CCSM3 for the 21st and 22nd centuries; special simulations for the demonstration project of the Advanced Climate Prediction Initiative (ACPI); a large number of single and combined forcing runs for the 20th century; and other simulations. The data from these simulations are available from a number of sites, including NERSC, via the DOE Earth System Grid. Information on the simulations can be found at www.cgd.ucar.edu/pcm/ and www.earthsystemgrid.org/.

The IPCC production effort represents a distributed multi-institutional effort by scientists and software engineers at NCAR, Oak Ridge National Laboratory (ORNL), NERSC, and Berkeley Lab. The IPCC SRES science scenarios that were recommended to the global coupled climate modeling groups by the IPCC WG1 include SRES A2, SRES A1B, and SRES B1. The A2, A1B, and B1 scenarios consist of five-member ensembles run from the year 2000 to the year 2100. Then, commitment scenarios are run with the concentrations of all atmospheric constituents in A1B and B1 held constant at year 2100 values, and the model continues to the year 2200 with five member ensembles (Figure 16). One of the ensemble members from each is continued for an additional 100 years, to the year 2300, with constant year 2100 concentrations. There is also a 20th century commitment

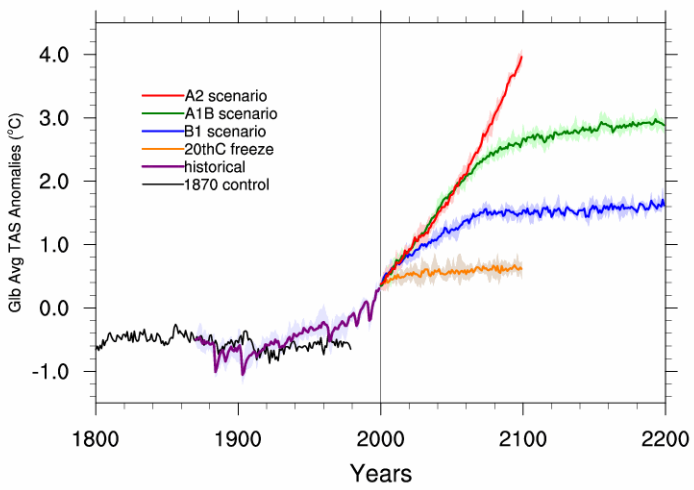

Figure 16. CCSM3 runs of various climate scenarios for the IPCC.

scenario that freezes concentrations at levels observed in the year 2000, and the model is run to the year 2100.

Due to the substantial computational requirements of the IPCC runs, the PCM and the CCSM3 IPCC runs were distributed across the NSF/NCAR, DOE/ORNL and DOE/ NERSC supercomputing sites. The complete IPCC SRES A2 scenario was run at NERSC, with the companion A1B and constant 20th century forcing scenarios at NCAR, and the B1 scenario at ORNL. The CCSM3 runs started in May 2004, ramped to full production in June, and finished during September and November 2004.

Focused model development support from DOE and NSF science and software engineering teams, significant computing resources at NERSC, ORNL, and NCAR, and 24-hour-by-7-day run monitoring and support by the CCP staff have kept the IPCC runs on track for completion in time to meet IPCC data submission and paper publication deadlines. During the run, NERSC consultants assisted the NCAR team to run the NERSC A2 five-member ensemble as a single 1020-processor massively parallel run.

The data products from the PCM and the CCSM3 IPCC simulations will be used for climate change impact studies and as boundary conditions to drive regional-scale models. To support DOE and NSF regional modeling efforts, researchers are outputting two additional sub-daily data streams for each of the CCSM IPCC scenarios. By design, the raw and post-processed data products from the different IPCC scenarios are being kept at the DOE and NSF sites where they were computed. This large distributed dataset is being freely served to the U.S. climate research and education community via the DOE Earth System Grid.

The results of this project are currently being submitted as the U.S./DOE/NSF contribution for the IPCC AR4. The integrated picture formed by these scenarios will be the basis for long-term energy and resource use policies. The higher resolution of these studies provides data to the climate impacts community that will be more useful than previous studies.

In addition to the large IPCC simulations, the NERSC computing resources enabled the CCP group to carry out important climate research science studies. The effect of future increases in pollutant haze over the Indian subcontinent on the Indian monsoon was simulated in collaboration with a group at Scripps. NERSC resources allowed quick testing of a hypothesis from the Program for Climate Model Diagnosis and Intercomparison (PCMDI) analysis group that greatly increased the fidelity of ozone simulations during the late 20th century. Finally, high-temporal-resolution output was added to the PCM A1FI IPCC SRES future scenario to enable modeling of detailed regional impacts of climate change on California water resources.

3.3.2 Global Biogeochemistry Models and Global Carbon Cycle Research

The climate modeling community has long envisioned an evolution from physical climate models to "earth system" models that include the effects of biology and chemistry, particularly those processes related to the global carbon cycle. The community generally accepts the premise that understanding and predicting global and regional climate change requires the inclusion of carbon cycle processes in models to fully simulate the feedbacks between the climate system and the carbon cycle. Moreover, models will ultimately be employed to predict atmospheric concentrations of CO_2 and other greenhouse gases as a function of anthropogenic and natural processes, such as industrial emissions, terrestrial carbon fixation, sequestration, land use patterns, etc. Additionally, these models must be run at the regional resolution needed to address processes that determine the climate state.

Predicting atmospheric CO_2 concentrations represents a substantial scientific advance, because large terrestrial vegetation and oceanic sources/sinks of carbon are key components of the present-day carbon cycle that will likely change as a result of climate change. Thus, the carbon cycle and the climate system are linked in such a way that both must be simulated simultaneously in order to produce credible predictions of future climate change.

Nevertheless, the development of coupled climate-carbon models with demonstrable quantitative skill will require a significant amount of effort and time to understand and validate their behavior at both the process level and as integrated systems. Carbon-climate models are going to be complex, with the carbon cycle strongly interacting with many other components. For example, we have very little understanding of long-term acclimation of terrestrial biological systems to changes in atmospheric CO_2 and climate, after ecosystems, nutrients, and soil biogeochemistry adjust to new conditions. Land management may play an important role in shaping ecosystems. A much more detailed and comprehensive understanding of the roles of nutrient cycles, CO_2-fertilization, water, temperature, etc. is needed to improve our predictive capabilities for multi-decadal evolution of carbon stocks on land ecosystems. Modeling the ocean component is similarly problematic. There has been no sound quantitative study of the relative importance of various process uncertainties in making carbon-cycle predictions on the multi-decadal time scale. The relationship between model ecosystem complexity and predictive skill for the problem of predicting multi-decadal carbon exchange has yet to be demonstrated. Perhaps global ocean carbon modeling needs to be put on a firmer physical basis, so that its equations describe a physically well-defined system, or perhaps very simple parameterizations have all the skill we will ever need.

At present, humans introduce about 6 gigatons of fossil fuel-derived carbon into the atmosphere each year in the form of carbon dioxide. This and previous emissions have resulted in an increase in concentration of atmospheric CO_2 from about 280 parts per million (ppmv) during the mid 19^{th} century to about 370 ppmv today. The atmospheric concentration is expected to continue to increase until it levels off at some "stabilization value" depending on governmental agreements to control emissions.

However, not all anthropogenic CO_2 stays in the atmosphere. Only about half of the emissions accumulate, the so-called "airborne fraction." The rest is taken up by the oceans or vegetation/soils as part of the carbon cycle. These carbon cycle sinks of carbon dioxide are expected to change as climate changes. The terrestrial (mostly plants) and marine (ocean circulation, chemistry, and biology) components of the global carbon cycle transfer large amounts of CO_2 into and out of the atmosphere seasonally and geographically. Thus, the net transfer of carbon that occurs, about half the man-made input, is small compared to the large gross fluxes of the system. This makes simulation a challenge, but more importantly, it helps produce a system that is delicately balanced and sensitive to climate change.

Figures 17 and 18 demonstrate current state-coupled carbon-climate modeling using the INCCA model (obtained via implementation of land and ocean biogoechemisty modules into a NCAR climate model). The magnitudes of feedbacks within the climate-carbon system are poorly constrained. Higher CO_2 concentrations increase photosynthesis and promote water-use and nitrogen-use efficiency of plants, ultimately increasing plant growth. Biomass and soil carbon, and thus terrestrial carbon uptake, may be expected to increase with higher atmospheric CO_2 levels. However, the effects of photosynthetic CO_2 "fertilization" will saturate at sufficiently high CO_2 levels, and higher global temperatures may increase the loss of soil carbon to the atmosphere.

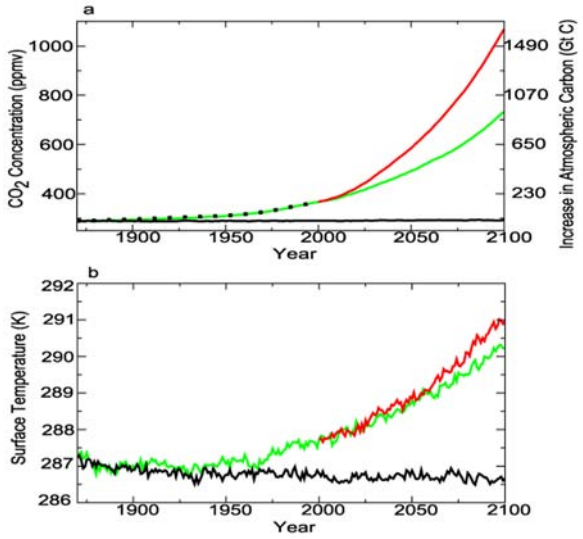

Figure 17. Atmospheric CO_2 levels (a) and globally averaged surface temperature (b) simulated by the three-dimensional coupled ocean/atmosphere/carbon cycle INCCA model. Black dots are observed amounts. The black lines show results from the model's control run, i.e., assuming no human effect on the carbon cycle. The colored lines show results with anthropogenic carbon dioxide emission specified at historical levels for 1870–2000 and prescribed for the 21st century according to the A2 (continuously increasing population) scenario of the Intergovernmental Panel on Climate Change (IPCC). Green lines show results assuming no saturation of carbon sinks; red lines show results assuming saturation occurs immediately at the present moment. These two assumptions bound the range of plausible behavior of the real Earth. The differences between the red and green curves demonstrate substantial uncertainty in future positive feedback of the carbon cycle on global warming.

Figure 17 shows that assumptions regarding CO_2-fertilization of the land biosphere greatly affect the atmospheric concentration of CO_2 (a) and temperatures(b). Year 2100 atmospheric CO_2 concentrations are 336 ppmv higher in the saturation case than in the fertilization case. Over the 21st century, 1790 gigatonnes of carbon (GtC) is emitted to the atmosphere; atmospheric CO_2 content increases by 776 GtC (366 ppmv) and 1489 GtC (702 ppmv) in the fertilization and saturation cases, respectively. The temperature difference at year 2100 between the saturation and fertilization cases is only 0.7 K (Fig.

1b). Figure 18 continues this simulation to year 2300 and shows global vegetation structure when, through continued CO_2 emissions, global temperatures have increased 8 K. Note the large changes in forest cover.

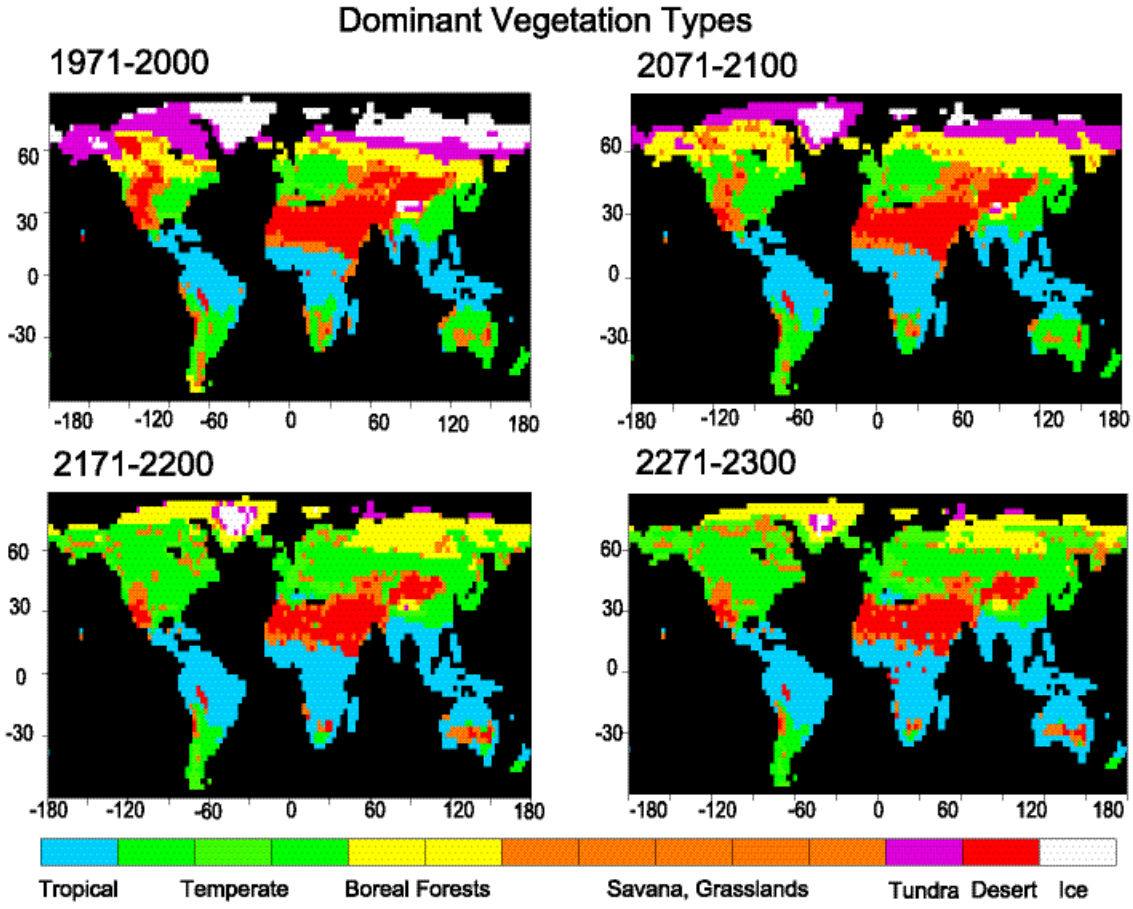

Figure 18. Vegetation simulated by the INCCA model for the present day (upper left) and for future time periods under the IPCC A2 scenario of anthropogenic CO_2 emission. The area covered by tropical and temperate forests increases dramatically by the year 2300 (lower right), when the model simulates a global surface temperature distribution averaging ~8 K warmer than the present day. Note, however, that the model does not include the direct effect of land-use change on vegetation (e.g., deforestation). The simulated vegetation is therefore what would occur "naturally" given the simulated climate change.

Future simulations must extend the length of simulated time, since the inclusion of the carbon cycle necessarily involves very long time scale processes, and increase the horizontal resolution, since small-scale heterogeneity in land types can influence the sources, sinks, and reservoirs of carbon. It will also be increasingly important to tie these carbon simulations to atmospheric chemistry and aerosol simulations. The rates of carbon exchange across atmosphere-ocean boundaries as well as those across atmosphere-land boundaries are strongly dependent on precipitation, aerosol loading, radiation/clouds, and fluxes of species such as iron. All these parameters depend on the chemical state of the atmosphere.

3.3.3 Beyond the Human Genome Project: Genomes to Life

The DNA sequences generated in hundreds of genome projects now provide scientists with the "parts lists" containing instructions for how an organism builds, operates, maintains, and reproduces itself while responding to various environmental conditions. But we still have very little knowledge of how cells use this information to "come alive." The functions of most genes remain unknown. Nor do we understand how genes and the proteins they encode interact with each other and with the environment. If we are to realize the potential of the genome projects, with far-ranging applications to such diverse fields as medicine, energy, and the environment, we must obtain this new level of knowledge.

One of the greatest impacts of having whole-genome sequences and powerful new genomic technologies may be an entirely new approach to conducting biological research. In the past, researchers studied one or a few genes or proteins at a time. Because life does not operate in such isolation, this inherently provided incomplete and often inaccurate views. Researchers now can approach questions systematically and on a much grander scale. They can study all the genes expressed in a particular environment or all the gene products in a specific tissue, organ, or tumor. Other analyses will focus on how tens of thousands of genes and proteins work together in interconnected networks to orchestrate the chemistry of life—a new field called *systems biology*.

The DOE Genomes to Life (GTL) program builds on these successes by combining DNA sequence data with advanced technologies to explore the amazingly diverse natural capabilities of microbes—the invisible organisms that thrive in every known environment on earth. The ultimate goal is to understand and use their diverse functions to meet DOE mission challenges in energy security, global climate change, and toxic waste cleanup.

Although we now have the entire genome sequences for hundreds of microbes, we still have very little understanding of how the information in DNA creates, sustains, and reproduces living systems. Obtaining this knowledge, a critical first step in harnessing microbial functions, requires a comprehensive approach extending from individual cells to many cells functioning in communities. Such studies must encompass proteins, multimolecular assemblies of components that work together (sometimes called *molecular machines*), the intricate labyrinth of pathways and networks in which they interact, and cells (Figure 19). The wealth of data to be collected must be assimilated, understood, and modeled on the scale and complexity of real living systems and processes.

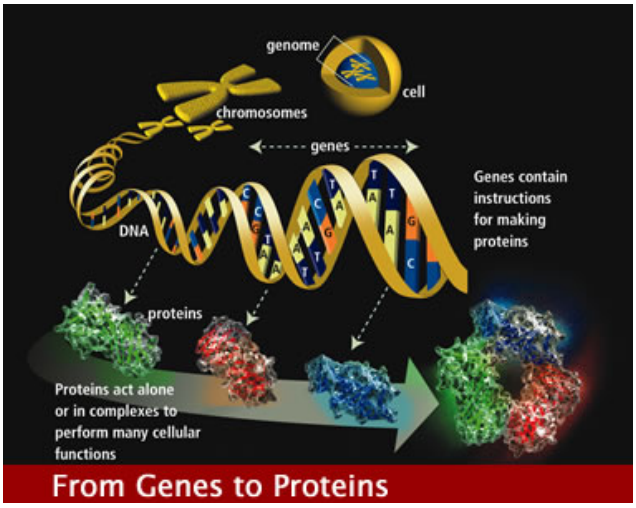

Figure 19. Understanding how DNA creates, sustains, and reproduces life requires computational modeling of complex living systems and processes.

Just as DNA sequencing capability was completely inadequate at the beginning of the Human Genome Project, the quantity and complexity of data that must be collected and analyzed for systems biology research far exceed current capabilities and capacities. Dozens of advanced large-scale technologies and approaches must be developed, with mathematics and computing guiding the research questions and interpretation at every step. Computational tools must manage and integrate the data into mechanistic models that describe how cells work. These studies eventually will enable an integrated and predictive understanding of how living cells function and respond to environmental changes, opening the door to using microbial capabilities. Several examples of such studies are described below.

Quantum Monte Carlo Study of Photoprotection via Carotenoids in Photosynthetic Centers. This project aims to increase understanding of the complex processes which occur during photosynthesis and is important on several levels. First, plants and bacteria are the world's foremost means of "carbon sequestration," or storing carbon from the atmosphere—a process which has enormous implications for climate change and global warming. Additionally, photosynthesis is an example of fundamental electron chemistry and is an efficient energy transfer system—processes which are fundamental in many areas of scientific research.

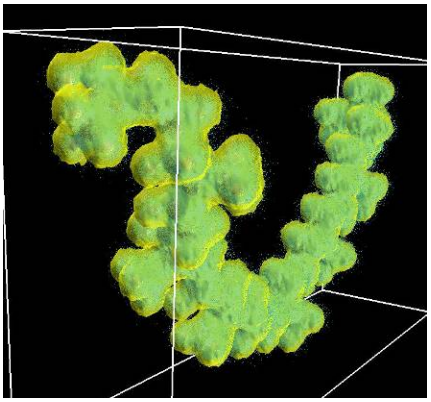

Figure 20. Quantum Monte Carlo electron density of spheroidene, a molecule responsible for photoprotection in photosynthetic reaction centers. This is the largest biological molecule ever treated using the accurate quantum Monte Carlo method.

After the project was under way, it was determined that the calculation could be made much more efficient by expressing the wave functions in a less dense representation and by using a Slater basis. In order to accomplish the latter, NERSC acquired and installed the ADF software package for first-principles electronic structure calculations on NERSC's IBM SP. As a result, the calculation is more capable of scaling up to large systems involving hundreds of electrons. A sparse representation of the wave function was accomplished via a grid acceleration technique. The improvements yielded a 16-fold reduction in wallclock time for one of the systems of interest (the spheroidene molecule shown in Figure 20). This project led to at least an order of magnitude improvement in the efficiency of the Zori quantum Monte Carlo code developed at UC Berkeley.

The goal of the project was to compute the excitation energies of the spheroidene and bacteriochlorophyll molecules and to use the obtained electron densities to estimate a rate of energy transfer between the two molecules. NERSC's storage capabilities allowed configurations from the Monte Carlo random walk to be stored for analysis, making possible an estimate of the rate of energy transfer as well as further analysis of the results using tools such as electron-pair localization function.

Molecular Dynameomics. This project combines molecular dynamics and proteomics to create an extensive repository of the molecular dynamics structures for protein folds, including the unfolding pathways. There are approximately 1130 known, non-redundant protein folds, of which about 30 have been simulated. Researchers plan to use the information from these simulations to improve algorithms for predicting protein structure. Structure prediction remains one of the elusive goals of protein chemistry. It is necessary to successfully predict native states of proteins, in order to translate the current deluge of genomic information into a form appropriate for better functional identification of proteins and drug design.

New types of databases (both hardware and operating system) must accommodate large data volumes, high schema complexity, and rapid query retrieval. Along with this, research should be done on new scalable storage hardware and software systems that can accommodate petabyte-scale data volumes and provide rapid analysis, data query, and retrieval. Rapid retrieval will require environments for large-scale data analysis on clusters and massively parallel programming technology for tools, libraries, and repositories.

Support will be vital for the development of reusable component and middleware analysis codes. One computational challenge of reverse engineering is to rigorously solve a network model that best matches known data and knowledge of the biology modeled. Data mining is an essential first step in solving the reverse-engineering problem. Much existing information is hidden in the often noisy, incomplete, and sometimes conflicting data.

Computational prediction and modeling and data collection through experiments should be one integrated process; computation should be a key driver for designing experiments. Networking and computing hardware also are required across the community, along with robust network technologies to support GTL facility-oriented community access, analysis, and archival activities. Stable computing power (i.e., in a production-oriented environment) is needed to run long time-scale biological simulations as well as real-time experiment drivers for the GTL facilities.

In summary, the future needs of GTL will require ever greater facilities for the analysis of complex biological systems and the integration of these technologies with advanced computational resources.

3.4 Office of Fusion Energy Sciences

Fusion Energy Science research is concerned with creating the conditions whereby light atomic nuclei can fuse together to form heavier nuclei and release energy. The dominant research focus is that of the physics of hot, fully ionized gas (plasma) in a magnetic field, and with the design and optimization of magnetic plasma confinement systems for future fusion experiments. Also included in the Fusion Energy Science portfolio is the physics of intense ion beams and high-powered lasers for inertial fusion energy, and the emerging field of high energy density physics.

A high-temperature magnetically confined plasma is an extraordinarily complicated medium, and its simulation represents one of the biggest challenges in computational physics. It is extremely anisotropic: a typical particle will travel a distance of ten kilometers along the direction of the magnetic field before it suffers an effective collision, whereas it is confined to scales on the order of millimeters in the directions perpendicular to the field. The collision time is long compared to the timescale of many dynamical phenomena, so that ordinary fluid theory is in general not applicable and it is often necessary to employ kinetic modeling or kinetic theory corrections to fluid equations. Timescales of interest form a continuum from the sub-microsecond timescale for plasma turbulence, to fractions of milliseconds for typical magnetohydrodymamic (MHD) phenomena, to many seconds for slow plasma evolutions. The magnetic field is itself a dynamical quantity of interest, and the electric field and magnetic field perturbations created by charges and currents in the plasma provide global coupling terms describing collective phenomena. The confined plasma supports a large variety of electromagnetic waves of interest for plasma heating and current sustainment and for plasma diagnostics. Present interest is in fully three-dimensional magnetic configurations, called stellarators, which may have advantages over the axisymmetric tokamak for fusion reactor applications because they do not require large plasma currents to be supplied and maintained.

Because of the many timescales present in a fusion plasma, and because of the different approaches, there are several different sub-disciplines within fusion science for which there exist mature code lines that are actively being developed and used. We describe highlights from five of these:

- ***Turbulence and transport:*** the science of the microscopic processes that occur in a magnetized plasma that lead to the transport of particles and energy across the magnetic field.
- ***Macroscopic stability:*** the study of the gross stability of fusion devices to large-scale instabilities, and of magnetic reconnection.
- ***Stellarator physics:*** the design and analysis of fully 3D plasma confinement configurations with hidden symmetries that lead to favorable transport and stability properties.
- ***Ion-beam dynamics:*** the study of the physics of intense ion beams for fusion, from their sources, through the accelerator driver and the fusion chamber environment, and on to the target.
- ***Fast ignition:*** A variant on laser fusion where the compressed capsule is ignited by a separate very intense, highly focused, very short laser.

3.4.1 Plasma Turbulence and Transport

One of the fundamental grand challenge problems in magnetic fusion energy research is the understanding and control of the turbulent transport of energy observed in the core of many fusion experiments. Drift-wave turbulence has been identified experimentally as a primary mechanism in degrading energy confinement in tokamak core plasmas. For some years there has been a large effort in the fusion community to simulate drift-wave turbulence. This simulation activity has led to a suite of 3D toroidal simulation codes that have

been used by a national collaboration (first under the auspices of the Numerical Tokamak Turbulence Project and subsequently as the SciDAC Plasma Microturbulence Project). These models have been extensively benchmarked against independent linear calculations of the basic underlying micro-instabilities and nonlinearly against one another to obtain results for the nonlinear saturation of drift-type instabilities in operating experiments, e.g., Princeton's TFTR, General Atomics' DIII-D, JET at Culham, Alcator C-Mod, and NSTX. The simulation results have been used to calibrate reduced models of the turbulent transport and to derive scaling relations for use in comparing and predicting experimental results with increasing success. Specific features associated with the moderation of the turbulent transport by means of externally imposed and self-generated velocity shear have been illustrated and demonstrated with the simulations. The elucidation of the physics of shear-flow inhibition by simulation and theory, and its confirmation in experiments exhibiting internal transport barriers, have led to major advances in producing tokamak plasmas with improved energy confinement in a more predictable and repeatable fashion.

Three related models have been used, all of which solve for the self-consistent electric or electromagnetic fields and the associated nonlinear plasma response. The three models solve the coupled Maxwell and Vlasov equations for plasmas supporting drift-type microinstabilities in three spatial dimensions and two velocity-space variables (the third velocity-space variable, the gyrophase angle in the applied magnetic field, has been analytically removed by gyroaveraging the equations). The three models are gyrokinetic particle codes (Lagrangian description), gyrofluid codes with Landau closure (mostly retired at this time), and gyrokinetic continuum codes (Eulerian description). All three models run efficiently on the NERSC IBM SP and other massively parallel platforms. By developing and applying multiple approaches to microturbulence simulation, the magnetic fusion community has been able to carefully explore the comparative computational efficiencies of the approaches and perform important code crosschecks on the nonlinear simulations. Because the approaches differ significantly in their algorithms, their diversity also has been useful for understanding how to optimize code efficiency for the specific architecture of the host supercomputer.

The fusion research community has been very successful in developing codes for modeling plasma turbulence for which computer run-time and problem size scale well with the number of processors on massively parallel machines. However, the gyrokinetic codes have evolved substantially during the past few years. Most of them now include full kinetic electron dynamics with electromagnetic coupling of the drift waves to kinetic shear Alfvén waves that modify the microinstabilities in the plasma at finite plasma pressures (Figure 21). The single ion species simulations of ion temperature gradient with adiabatic electrons, which have been the focus of plasma microturbulence research until recently, were already very computationally demanding. The new simulations with kinetic electrons and improved physics considerably increase the computational requirements in terms of number of operations, memory size, and time to solution. To date, these simulations have been limited by computing resources that constrain us to simulate experiments with either smaller plasmas or plasmas with less than the optimal spatial resolution, or to undertake fewer simulations and limit parameter studies, or to target an annu-

lar region of a tokamak experiment, albeit with realistic parameters, whose computational requirements are generally less stringent than those for a full global simulation. Another important limitation on research progress due to limited computer resources has been the turnaround time for a researcher to be able to undertake a series of simulations addressing a parameter scan, which profoundly impacts the pace of physics progress. Thus, upgrades at NERSC addressing both capability and capacity simultaneously are vital to taking the next steps in increasingly realistic physical simulations in magnetic fusion research on microturbulence.

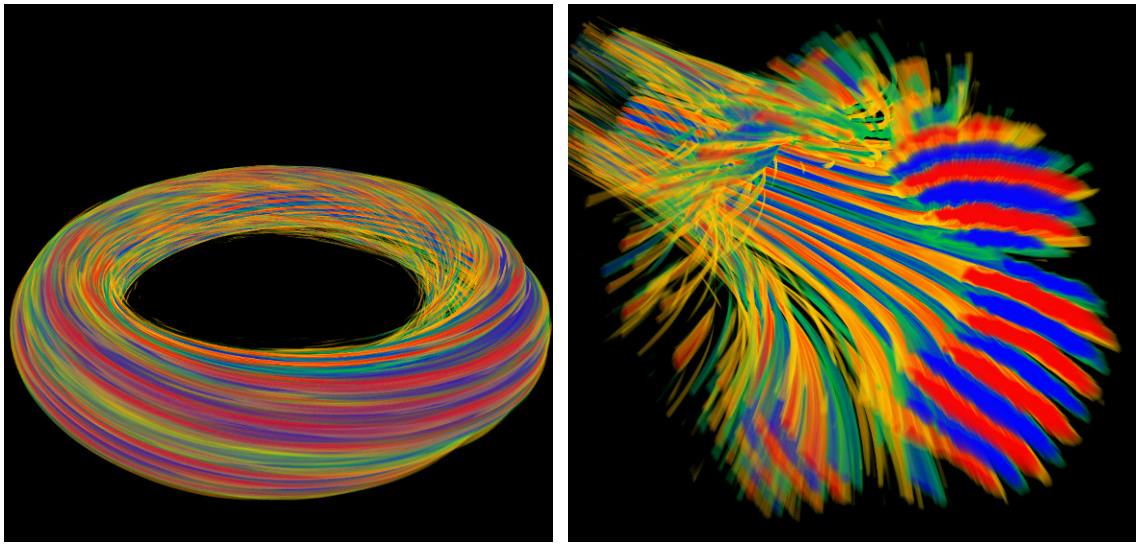

Figure 21. Advanced volume visualization of microturbulence data generated by PPPL's Gyrokinetic Toroidal Code (GTC). The field-line following perturbed electrostatic potential inside the tokamak is shown using a rainbow color map, with the lowest negative potential values mapped to red and the highest positive values mapped to blue. Higher transparency depicts values of potential near zero. The image on the left shows the full torus, while the one on the right shows a cross-section where the elongated eddies of the linear phase are dominant. (Images courtesy of Kwan-Liu Ma's research group at UC Davis, part of the SCIDAC Gyrokinetic Particle Simulation Project.)

To illustrate the importance of upgrading computing power, we consider a full-device simulation of a next-step tokamak capable of achieving ignition or high fusion gain, e.g., ITER. We have performed detailed scaling studies with two of the leading global gyrokinetic codes, GYRO and GTC, both extensively utilized for simulating core tokamak turbulence. GYRO is a Eulerian continuum code, while GTC is a particle-in-cell code. We project that 300 hours on 4096 processors on Seaborg would be required to find a steady-state transport profile in ITER. This assumes a peak performance of 6 teraflops. Current superscalar processors are about 4 times faster than Seaborg for most codes, so we expect NERSC-5 to deliver at least 20 Tflops of peak performance for 4096 processors. Thus, approximately 90 hours would be required on such a computer with that many processors to complete one ITER simulation, assuming that we get the same percentage of peak performance.

We conclude that full-device simulations of ITG turbulence in an ignition-scale magnetic fusion device is ambitious, but is achievable on a 20 Tflops computer. Such simulations

would allow detailed investigation of scientific questions regarding the role of meso-scales in ITG turbulence, the dynamics of spectral transport, and the formation and evolution of transport barriers. To explore physics parameters adequately and do justice to the science will require 50–100 such simulations per calendar year. Given that NERSC is a shared facility, an appropriately capable computer to host ITER simulations would need to deliver 100 Tflops or more. In order to accommodate all of the needs of the fusion community in simulating various aspects of ITER, a capability machine approaching 1 petaflop is justified. The expense of such computing power is significant, but is small compared to the $\$5–10 \times 10^9$ investment that will be made in ITER.

3.4.2 Macroscopic Stability

There are many critical scientific problems in fusion science that can be addressed with the macroscopic simulation model known as the extended MHD model. Most of these share the common features of extreme temporal and spatial stiffness, severe spatial anisotropy, and complex boundary conditions.

The emphasis in large-scale MHD modeling is on simulating global nonlinear dynamic phenomena essential to the operation and design of tokamaks and other fusions devices. The next generation of burning plasma fusion devices will need to operate near their stability limits. These limits define an operational window that is bounded by several nonlinear processes, including internal reconnection (sawtooth), the neoclassical tearing mode, and the resistive wall mode. Each of these phenomena depends on the detailed pressure and current distribution in the device, as well as plasma properties associated with collisionality and ion orbit (Larmor radius) size. The next device will be in a new parameter regime, and its stability will be affected by the presence of a large energetic fusion-product (alpha particle) population so that simple extrapolations from existing devices are not possible. An assessment of the nonlinear stability of these devices requires a computational approach. The computational models being developed for this purpose are based on sets of magneto-fluid equations for magnetized plasma that include the effects of realistic geometry and boundary conditions. The nearly collisionless nature of high-temperature plasmas is also taken into account by supplementing the fluid equations with analytic and particle-based closures. When implemented on computer hardware capable of delivering teraflops of computing power, these models can be applied to understand nonlinear magnetohydrodynamics in parameter regimes that characterize both present and proposed fusion experimental devices (Figure 22).

The programmatic importance of this work is high. Applying the extended MHD model with increased numerical resolution will greatly improve our ability to evaluate and understand the essential nonlinear mechanisms that set pressure limits in tokamaks and other plasma confinement devices. In addition to its programmatic benefits, macroscopic simulation offers an exciting opportunity to advance the scientific understanding of several fundamental plasma processes of wider scientific interest. These include plasma relaxation and self-organization, magnetic reconnection in low-collisionality plasmas, and the effect of wave-particle resonance on macroscopic plasma behavior.

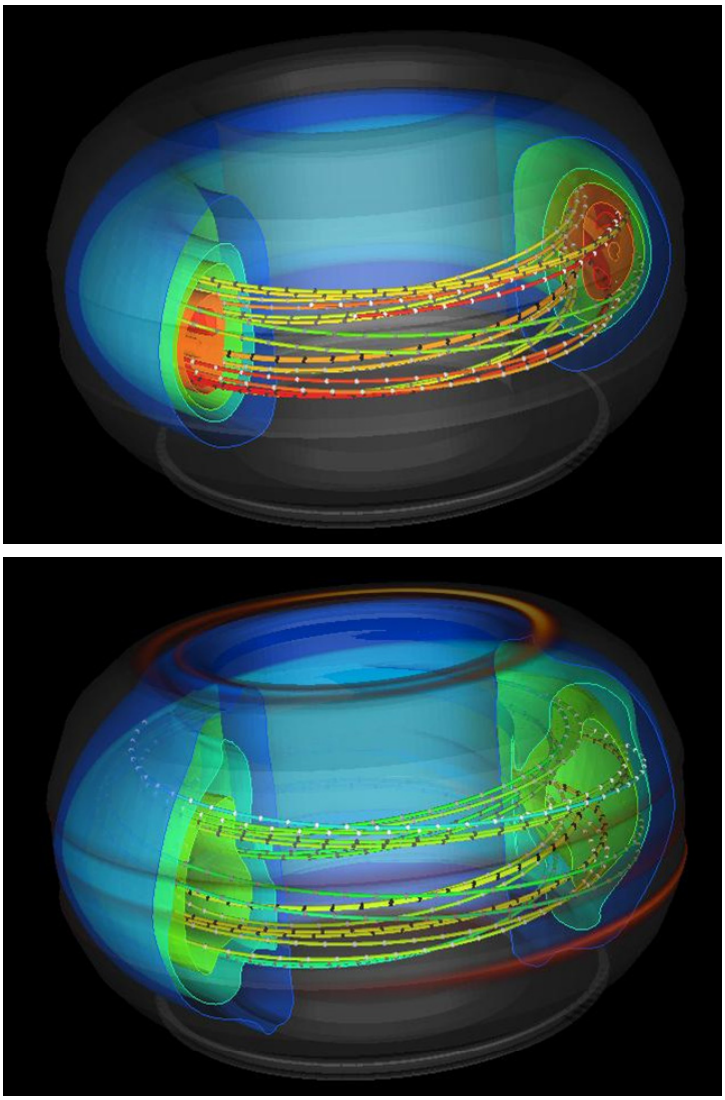

Figure 22. A recent application of macroscopic nonlinear simulation studies how heat flow to the wall becomes localized and asymmetric during a disruptive instability in the DIII-D experiment. An internal instability shown developing in the top frame rearranges magnetic field-lines. The color scheme of the constant-pressure surfaces indicates a loss of high pressure later in time (bottom frame), while localized heat flow brightens elongated spots on the shadowy wall. Understanding and control of disruptive processes are essential for burning plasma experiments. (Courtesy of S. Kruger, Tech-X Corp., and A. Sanderson, University of Utah)

The computational challenges of extended MHD simulation are also high, particularly with respect to the range of temporal scales. The fastest physical behavior results from waves and anisotropic diffusion, which propagate information across the domain. The slowest behavior is associated with profile adjustments over global transport timescales. Successful simulations must evolve the nonlinear macroscopic dynamics, which arise at intermediate timescales, long enough to understand their effects on profiles, while maintaining the subtle balances imposed by the fastest effects. The range of temporal scales for resistive MHD alone is indicated by the *Lundquist number*, *S,* which is the ratio of typical resistive diffusion times to Alfvén-wave transient times. Present computing power limits nonlinear computations to *S*-values of 10^5–10^7, at least an order of magnitude less than what will be realized in burning plasma experiments. Moreover, extending the MHD model to include electron fluid dynamics introduces faster timescales together with an increased spatial resolution requirement. Both increased computing power and algorithmic improvements are required, and the two must be compatible. For example, application of a direct parallel linear system solver provided a performance breakthrough for the implicit time-stepping used in multi-scale MHD simulations.[39] Parallel direct

[39] "AMR Methods Accelerate MHD Simulations," NERSC 2003 Annual Report, http://www.nersc.gov/news/annual_reports/annrep03/advances/2.2.superflusolver.html.

methods are also expected to play a critical role for simulating the faster electron fluid dynamics. However, parallel direct methods are communication-intensive, and gains in CPU processing speed will not be effective without more impressive gains in the ability to communicate data among parallel processors.

3.4.3 Stellarator Physics

Stellarators are three-dimensional toroidal plasma confinement devices that rely on a numerically determined plasma surface shape in order to achieve optimized plasma confinement, stability, and steady-state operation. These devices have the potential of providing a low-cost path to the development of fusion power with a much lower risk of the current-driven disruptive instabilities that are present in tokamaks. Stellarator physics and simulation are of increasing interest in the U.S. fusion program due to the fact that new experimental facilities will be built at several U.S. research centers. The stellarator devices that are either operating or in the planning/construction phase in the U.S. program target the three possible forms of quasi-symmetry (quasi-helical, quasi-toroidal, and quasi-poloidal). Quasi-symmetry implies that there is an ignorable coordinate present for the magnetic field strength (i.e., a direction in which the magnetic field does not vary significantly) when viewed within a particular set of coordinates. Figure 23 shows devices (HSX, NCSX, and QPS) that have been designed based on these forms of quasi-symmetry. Due to their inherently three-dimensional nature, stellarators rely heavily on numerical simulation and computation. The optimized design and physics analysis of these devices would not be possible without access to high-performance parallel computers.

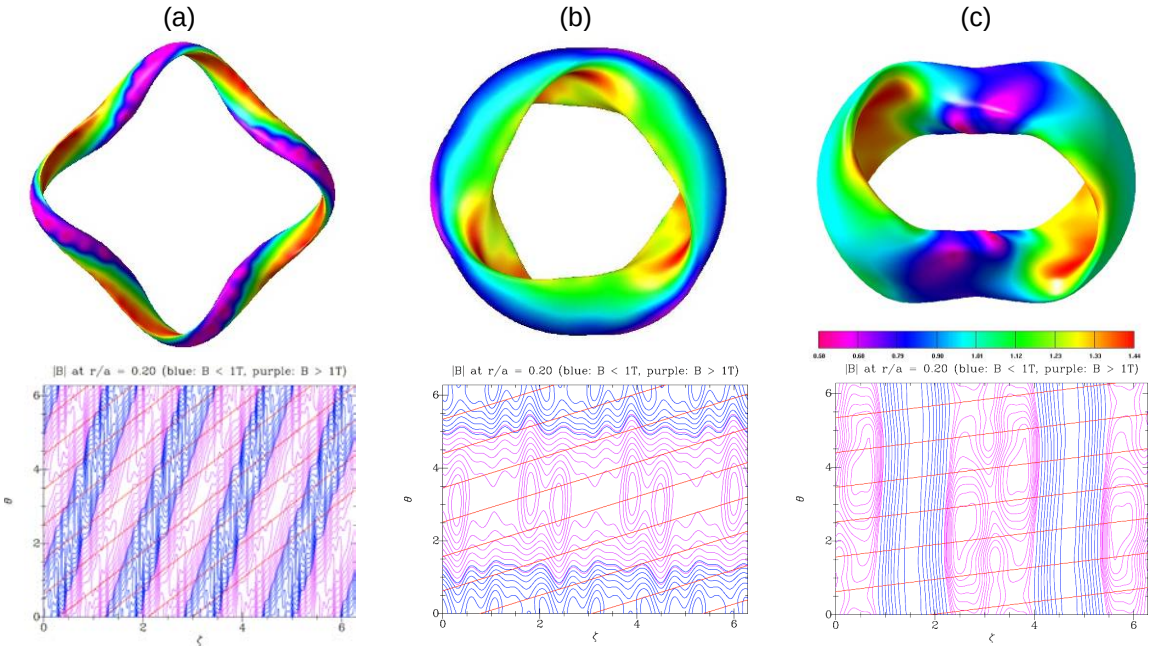

Figure 23. Outer flux surfaces (top) with magnetic field strength color contours of devices with (a) quasi-helical, (b) quasi-toroidal, and (c) quasi-poloidal symmetry. The lower figures show contours of magnetic field strength for an interior flux surface in the transformed coordinates appropriate to quasi-symmetry for each device.

Three-dimensional configurational flexibility is one of the significant new features that stellarator experiments offer. Recent designs allow variation of currents in up to ten different major coil groups; various smaller trim coils may also be present. However, the value of all this flexibility would be largely lost if it were not for powerful optimization codes that can find coil current combinations in this multi-dimensional parameter space that lead to improved plasma configurations. The STELLOPT code has been the primary tool used for both the design of recent stellarators and in flexibility studies. One example of the use of flexibility is in the reduction of magnetic islands. These are defects in the magnetic configuration that can be induced by a variety of sources, such as small manufacturing or alignment errors in the magnetic coils, changes in coil geometry due to thermal or magnetic stresses, or internally by plasma instabilities. In the vicinity of magnetic islands, the magnetic field lines no longer are confined to nested, closed toroidal flux surfaces. They move radially away from the initial equilibrium surfaces, leading to much more rapid plasma energy losses. Figure 24 shows an example of island suppression.

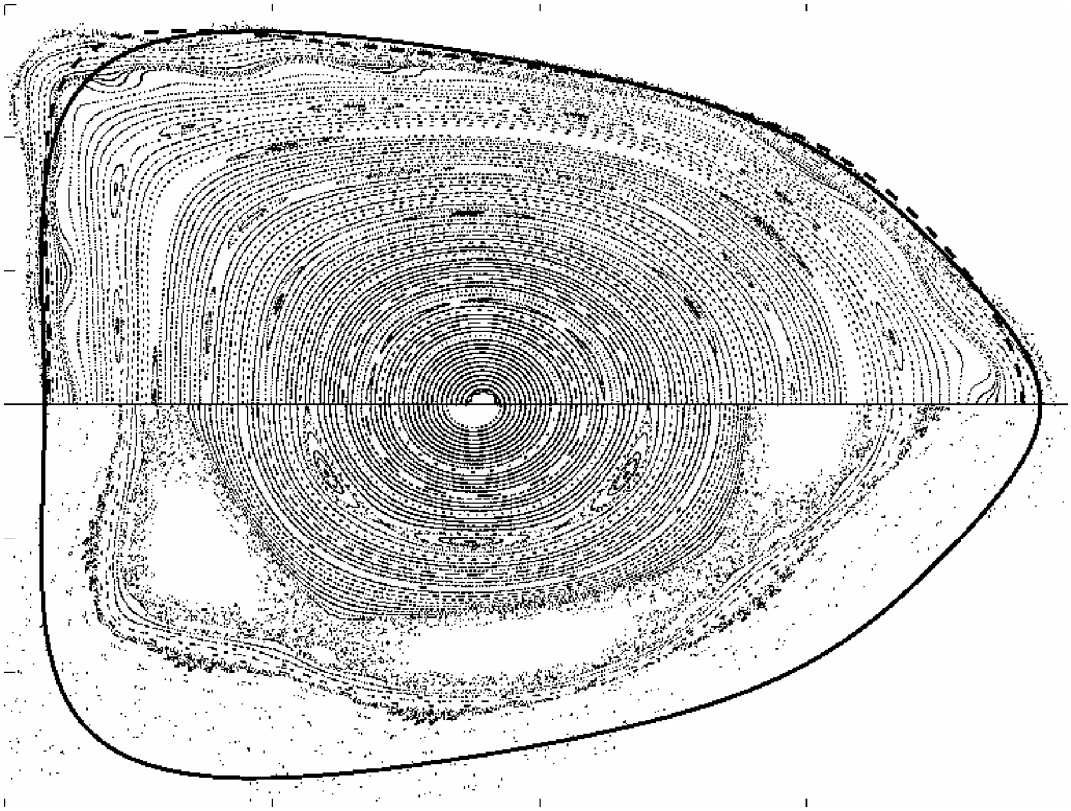

Figure 24. PIES magnetic field puncture plot *t* at a fixed toroidal plane for the NCSX stellarator showing the initial state with sizeable magnetic islands (lower half) and the final state after optimization (upper half) with relatively small magnetic islands and nested, closed flux surfaces.

This example is based on an early coil design algorithm for the National Compact Stellarator Experiment (NCSX) that did not explicitly target resonant field reduction, and resulted in configurations with large magnetic islands. Subsequently an optimizer was built around the PIES equilibrium code which modified the design of candidate coil sets to give good finite current, finite pressure flux surfaces, while preserving desired physics and engineering properties of the configuration. The optimizer was applied to obtain the

reference coil set for NCSX. Figure 24 shows the results of following the field lines for the NCSX coil set before (bottom half) and after (top half) the resonant field reduction process.

Computation of plasma equilibria is the first step for both the optimization and the physics analysis of stellarators. The PIES equilibrium code can calculate full 3D equilibria, including island regions. A second equilibrium code, VMEC, which assumes nested closed flux surfaces, is also extensively used in the optimization and physics analysis of stellarators due to its ability to rapidly calculate somewhat more approximate equilibria. The PIES code will be upgraded to include neoclassical transport effects on magnetic islands. Its efficiency will also be improved for 3D equilibria with islands and stochastic regions so it can be used more routinely for analysis of experimental data.

Plasma transport is another computationally intensive physics area of importance for stellarators. Due to the three-dimensional nature of stellarator magnetic fields, analytic expressions for transport coefficients cannot be derived in general to the extent that has been possible for axially symmetric (tokamak) devices. Computational solutions have been developed based on Eulerian localized transport codes, such as DKES, as well as Lagrangian orbit-following Monte Carlo codes such as DELTA5D, GTC, and ORBIT. A shell code has recently been developed to run DKES in parallel over multiple flux surfaces with subsidiary loops over collisionality and electric field so that the extensive transport coefficient databases required for each configuration can be more rapidly generated. In addition to predictions of cross-field transport of particles and energy, this type of calculation has also allowed prediction of plasma flow velocity moments within each magnetic flux surface (Figure 25). Such plasma flow predictions are important for understanding enhanced transport barriers (related to sheared flows) as well as allowing improved calculations of collisional bootstrap current and the ambipolar electric field. Future extensions of this moments method include applications to multi-species impurity transport and improvements in the calculation of the viscosity coefficients. The Monte Carlo codes have been developed to address non-local transport problems, such as energetic particle transport (neutral beam slowing-down, RF tail populations, fusion alpha particles) and transport in the presence of turbulent fluctuating fields. They have also been benchmarked against thermal neoclassical transport. Since they follow non-interacting particles (a test-particle/static field model is typically used), such codes are readily parallelized and can efficiently use large numbers of processors.

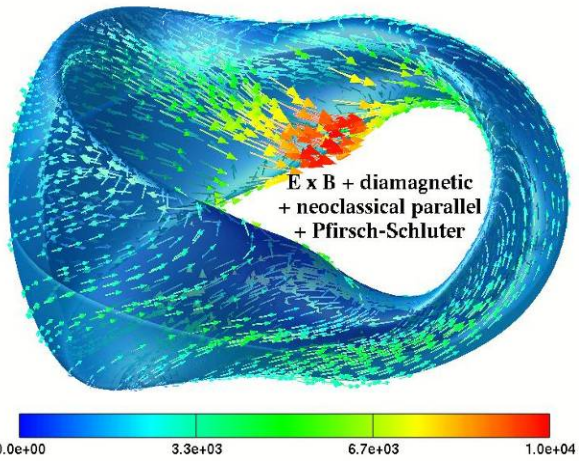

Figure 25. Plasma velocity flow vectors on the outer flux surface of the QPS stellarator.

Plasma heating RF (radio frequency) heating is used extensively in stellarator experiments due to its low cost and high flexibility. New computationally intensive tools are under development to calculate the RF wave field and its plasma absorption for three-dimensional systems. Although the wave fields enter the plasma initially at moderate wavelengths, they can become mode-converted and absorbed in local regions at much shorter wavelengths, resulting in strong couplings over a range of length scales.

Plasma stability is a final area of stellarator physics where access to high performance parallel computers is essential. Stability issues for compact stellarators include current-driven modes, short-scale-length pressure-driven ballooning modes, energetic-particle-driven instabilities, and plasma microturbulence. Several linear variational stability codes (TERPSICHORE, CAS3D, COBRA) for addressing current-driven and ballooning instabilities have been developed and are coupled into the optimization efforts. Resistive MHD and hybrid fluid-particle models are addressed by the M3D code, which can follow the nonlinear evolution of large-scale instabilities. M3D can also model linear and nonlinear energetic particle-driven modes. The STELLGAP code has been developed to calculate high-resolution Alfvén gap structures for three-dimensional systems and to identify modes that can be driven unstable either by in situ energetic tails or by external antenna excitation.

The future computational needs of stellarator physics modeling activities are driven by the need for faster processors, increased memory per processor (2 GB and above), improved turnaround time for parallel calculations, and continued use of systems with fast general parallel file systems (GPFS). The first two needs are driven by the highly computationally intensive nature of three-dimensional plasma configurations. Our ability to model the various physical processes in these systems is limited by processor speed and, in some cases, memory due to the large amount of data that must be available to resolve even equilibrium (with multiple island chains) physics in stellarators. The need for rapid turnaround is driven by the fact that, although these codes have been adapted for parallel architectures, they have not yet benefited from the many person-years of development that characterize more mature codes in other areas. Thus, there is more time spent in the debug/testing part of the coding cycle in stellarator computational research, and more rapid turnaround of moderate-sized parallel jobs (up to 128 processors) is essential for progress in this area. Many of the parallel codes in this area (e.g., the STELLOPT code) use disk writes to preserve and pass large amounts of data around among a variety of heterogeneous subsidiary codes. Problems have been encountered when this code has been ported to clusters that do not support general parallel file systems. Therefore, it is critical that future NERSC systems continue to support GPFS and use techniques such as disk striping to maintain high performance disk access. A final administrative issue for future NERSC systems is that stellarator research is a very international activity, and it is essential that good access be available to these systems from our international collaborators.

3.4.4 Intense Ion Beams for High Energy Density Physics and Inertial Fusion Energy

The U.S. Heavy Ion Fusion (HIF) program is developing intense beams of heavy ions as drivers that can heat matter to the regime of high energy density physics (HEDP), that is, 10^{11} J/m^3 or greater. Near-term beam applications will center on fundamental physics studies of the warm dense matter (WDM) regime, enabled by the unique energy deposition properties of heavy ions in matter whereby the heating can be both volumetric and shockless. Indeed, ion drivers promise a means for the thorough exploration of the entire HEDP region of the temperature-density parameter space. In the longer term, because ion beams can be created efficiently and focused by electromagnetic lenses that can be robust to the effects of target explosions, a number of studies and reviews have concluded that intense beams of heavy ions represent an attractive approach to inertial fusion energy: heavy ion fusion.

The WDM regime corresponds to densities of ~0.01 to 10 times solid density and temperatures ~0.1 to 10 eV. Here the space charge potential energy of neighboring ions is of the same order as, or somewhat greater than, the ion thermal energy; the traditional "weak-coupling" assumption of ordinary plasmas breaks down, and theoretical descriptions become more complex. Exploring this regime experimentally is thus critical to assessment of the various theoretical models. Establishing the equation of state and transport properties in this regime is of considerable importance not only for its intrinsic scientific value and its relevance to inertial fusion target physics, but also because understanding these quantities is crucial to the quantitative understanding of the structure of planets and stars. The HIF program in the U.S. plans to use ion beams to heat thin foils uniformly, with beam energy deposition occurring as the ions slow down through their Bragg peak energy. Complementary programs in Germany, Japan, and Russia are centered on the use of higher-energy ions depositing a small fraction of their energy in long cylindrical targets.

Both the beams themselves and their interactions with the targets they heat require extensive simulations for a full understanding. The target interaction studies are similar in many regards to those carried out in the larger Inertial Confinement Fusion program, and indeed rely to a large degree on the same computational tools. The key question in Heavy Ion Fusion beam science is: "*How do intense ion beams behave as they are accelerated and compressed into a small volume in space and time?*" The challenges and opportunities associated with simulating these beams are described briefly below.

Intense ion beams are non-neutral plasmas and exhibit collective, nonlinear dynamics that must be understood using the kinetic models of plasma physics. This physics is rich and subtle: a wide range in spatial and temporal scales is involved, and effects associated with instabilities and non-ideal processes must be understood. Ion beams have a long memory, and initialization of a simulation at mid-system with an idealized particle distribution is often unsatisfactory; thus, a key goal is to develop and extensively exploit an integrated and detailed source-to-target beam simulation capability. Major issues include:

Long-term evolution of space-charge-dominated beams. The beams are accelerated by inductive electric fields and confined by applied "focusing" fields. The dynamics are mainly governed by a balance between applied fields and space charge forces, and are Liouvillean (collisionless): the phase space density remains constant along particle orbits. As a result, emittance growth (dilution of the phase space) takes place through complicated distortions driven by collective processes, imperfect applied fields, image fields from nearby conductors, and inter-beam forces. Such dilution must be minimized because of the need to focus the beams onto small (1 to a few mm) focal spots on the targets. This area is challenging because of the need for an efficient but detailed description of the applied fields, and the needs for good statistics and mesh resolution.

Beam halo generation and multispecies effects in driver. Oscillations of the beam core can parametrically pump particles into an outlying, or halo, population. To avoid the adverse effects of ions impinging on walls, beam halo must be kept minimal. Here, particle-in-cell (PIC) methods have been used, but emerging nonlinear-perturbative and continuum-Vlasov methods may offer advantages. A new and comprehensive set of models of the physics of stray electron "clouds" and gas in positive-particle accelerators is well along in development. These models account for the complex set of interactions among the various species, with the walls, and with the applied and self-fields. This area is computationally challenging because of the ratio between the time scales for electron motion and those for electron build-up. To this effect researchers have developed a new large-timestep electron "mover" that smoothly interpolates between direct-orbit and drift formulations as the magnetization of the electrons varies.

Beam interactions with the target chamber environment. Three-dimensional simulations of the propagation of beams through the final focusing optics and onward through the target chamber's environment of gas and plasma are required in order to provide a realistically complete model of the target illumination. The beam and background plasma dynamics include multibeam effects; return current formation and dynamics (streaming instabilities); imperfect neutralization; beam stripping; emittance growth; and photoionization of the beam ions and background gas. Multiple-beam interactions near the target will be one focus of research efforts; collective instabilities, such as resistive hose, filamentation, and two-stream modes, will be another. Here the challenges include the needs for complex physics models, outgoing-wave boundary conditions, and an implicit hybrid model (i.e., a blend of fluid and discrete-particle electrons). An implicit electromagnetic (EM) field model (which can stably under-resolve fast time scales not essential to the physics), or a magneto-inductive (Darwin) model that eliminates light waves, is also needed.

Representative images from simulations of a variety of problems are shown in Figure 26. The principal computer codes have been adapted to NERSC's parallel computers and achieve good scaling. However, since simulation needs are progressing toward self-consistent full-system studies, and multi-species effects must be understood, the research effort stands to benefit greatly from a qualitative increase in available computational resources. A detailed end-to-end particle-in-cell simulation of a single full-scale beam in a full-scale general HEDP facility or inertial fusion driver will require 10^8–10^9 particles,

10^5–10^6 time steps, and a (moving and adaptive) 3D computational grid of order 10^9 mesh points or more. This estimate assumes success in the development of an improved large-timestep electron mover and the use of optimized parallel adaptive mesh refinement with particle-in-cell dynamics (AMR-PIC, a technique pioneered by the HIF program). Thus, simulations to date have covered only sections of a system, and in most cases have used approximate models, e.g., neglecting stray electrons.

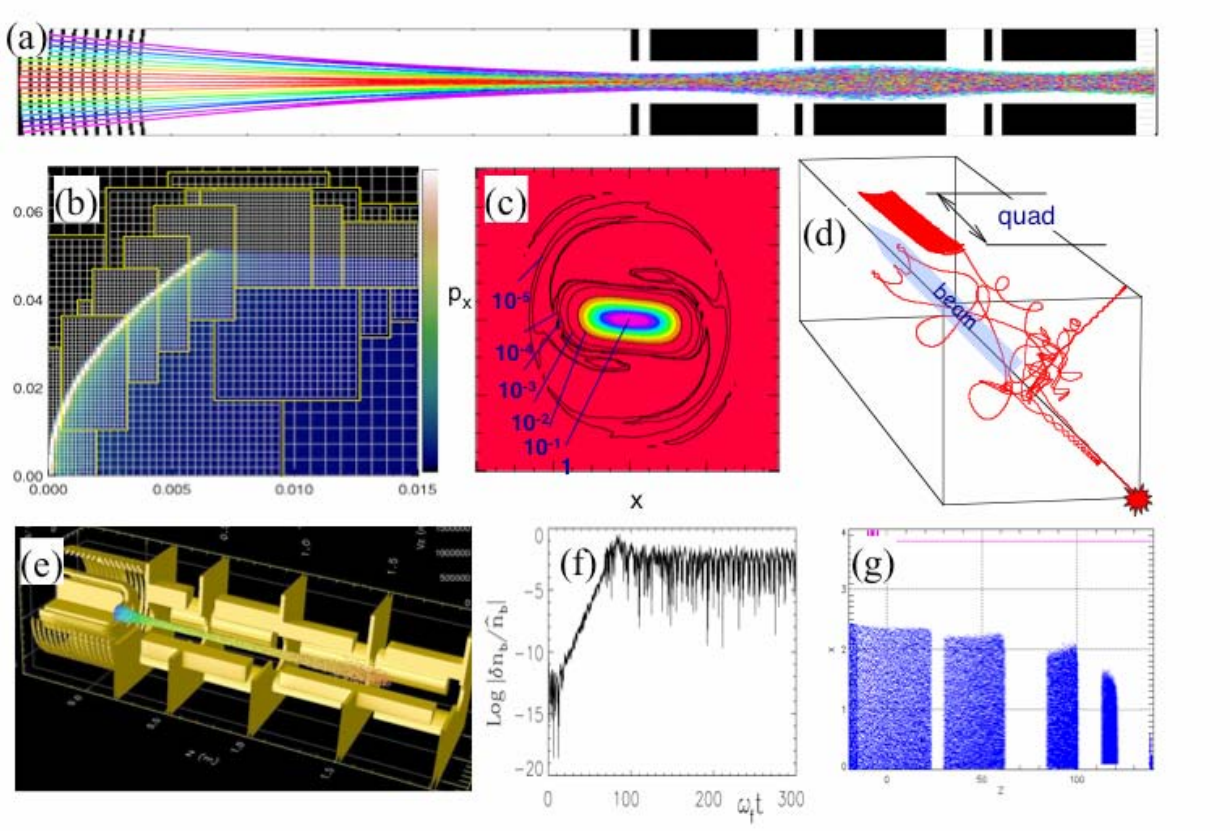

Figure 26. Representative output from beam simulations carried out in the Heavy Ion Fusion program: (a) Simulation of merging-beamlet intense-beam ion injector experiment at LLNL; (b) expanded view of a WARP3D simulation of the injector diode in the HCX experiment at LBNL, showing mesh-refinement patches; (c) semi-Lagrangian Vlasov simulation of beam halo generation due to anharmonic applied fields, using a prototype model in WARP-SLV; (d) WARP3D simulation of a stray-electron orbit in a quadrupole magnet and adjacent regions; (e) WARP3D PIC simulation of space-charge-limited emission off a curved surface, and acceleration in the 3D structure of the HCX injector, including sub-grid-scale placement of conductor boundaries (cut-cell method); (f) BEST nonlinear-perturbative simulation of electrostatic anisotropy-driven beam mode, showing linear growth and saturation; and (g) LSP simulation of neutralized pulse-compression and focusing in support of upcoming NDCX experiments at LBNL. (Image courtesy of A. Friedman, Heavy Ion Fusion Virtual National Laboratory, a partnership among LBNL, LLNL, and PPPL)

3.4.5 Fast Ignition

In conventional inertial confinement fusion (ICF), the hot spot is a low-density, high-temperature, thermonuclear plasma heated to ignition temperatures by compressional

work of the surrounding cold, overdense shell driven by the laser energy. The hot spot mass is typically much smaller than the shell mass, and the energy required to heat the hot spot to 10 keV is comparable to the energy required to compress the cold shell to several hundred g/cc (~500 g/cc for a NIF-like target). In conventional ICF, the driver must provide the compressional work to simultaneously heat the hot spot and compress the deuterium-tritium (DT) shell.

The assembly of both the hot spot and the surrounding shell is difficult. The hot spot can only be heated to thermonuclear temperatures if the surrounding shell remains reasonably uniform during the hot-spot assembly. However, since the hot spot-shell interface is hydrodynamically unstable during the assembly phase, small nonuniformities of the shell inner surface will grow exponentially, causing "fingers" of the cold shell plasma to penetrate and cool the hot spot. Two- and three-dimensional simulations indicate that this "deceleration phase instability" often prevents the achievement of ignition conditions inside the hot spot. All ICF implosions must be designed to minimize the impact of such hydrodynamic instabilities.

The overdense cold fuel assembly can also be problematic in conventional ICF. The imploding cold shell is unstable to the "acceleration phase" Rayleigh-Taylor instability when a fraction of the absorbed laser energy is converted into the shell kinetic energy. This instability causes the nonuniformities on the shell outer surface to grow exponentially, resulting in a large-scale distortion of the accelerating shell. A distorted shell supplies large initial seeds for the deceleration phase instability and does not provide efficient compressional heating of the hot spot. Both the acceleration and deceleration instabilities impose severe constraints on the target design and driver requirements for conventional ICF.

The fast ignition scheme has the great advantage, with respect to conventional ICF, of removing the constraints associated with the hot-spot formation while at the same time alleviating some of the requirements imposed by the cold fuel assembly. In fast ignition, the target is first compressed to the required density (~300 g/cc) by either direct drive from nanosecond lasers or indirect drive from x-rays using a hohlraum driven by nanosecond lasers, ion beams, or a z-pinch. The ignition is initiated by a fast laser pulse (the so-called "ignitor pulse"), which "bores" through the outer, lower-density edges of the target, then deposits its energy into fast electrons (~MeV energy per particle) near the relativistic critical-density surface. Since the igniting spark is externally supplied, the hydrodynamic constraints associated with the hot spot assembly are relaxed. The cold fuel still needs to be assembled fairly carefully in order to achieve the required high densities, which are typically less than the densities required in conventional ICF. The optimum compressed-fuel configuration for fast ignition is an approximately uniform density spherical assembly of high-density DT fuel without a central hot spot. Gain in fast ignition can be considerably larger than in conventional ICF because lower fuel densities can be used with an acceptable ignition threshold. Furthermore, the central hot spot is not needed, and the driver energy can be entirely used to compress the cold fuel.

Key questions in fast ignition research include: (1) How can we assemble an optimum fast ignition fuel pellet with well-connected uniform fuel mass density (no hot spot)? (2) What are the characteristics of the energetic electrons produced through laser-plasma interactions? and (3) How do the huge currents of the energetic electrons transport through the dense plasma, and how do the energetic electrons deposit their energy to create a hot spot? All research in these areas relies heavily on computer simulations. The fuel assembly simulations use hydrodynamic codes. The laser-plasma interactions are simulated by PIC codes. For the electron transport and energy deposition in the dense fuel region, even though the physics there is most likely highly kinetic, due to computational limitations, a hybrid approach is envisioned with a fluid description of the background plasma and a particle description of the energetic electrons.

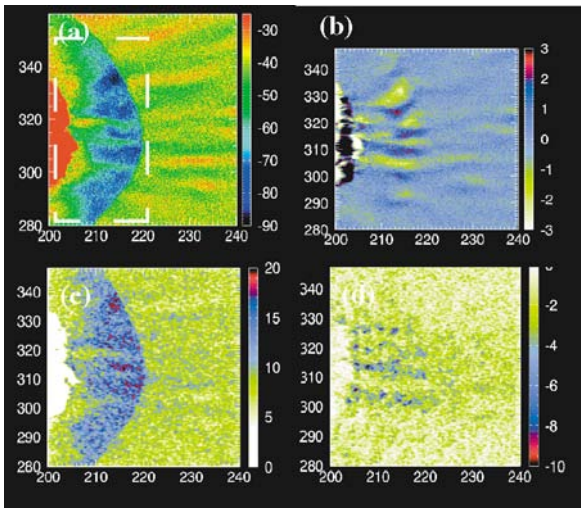

Figure 27. For a circular target with *p* polarization at t = 964 fs, (a) electron density (in units of n_c), (b) magnetic field b_3 (in units of $mc\omega/e$), (c) return current (in units of $n_c ec$), and (d) fast electron current. Distance in units of $\lambda_0/2\pi$. (Ren et al., Phys. Rev. Lett. **93**, 185004 [2004])

The computational challenges faced in fast ignition research can be illustrated by considering a full-scale PIC simulation of the laser-plasma interaction region. The relevant volume is at least 200 × $(100)^2$ $(\mu m)^3$ with an average plasma density of 10 times the critical density. A fully explicit PIC simulation would need >250 billion cells and >10^{12} particles. About 1.5 million timesteps are needed for simulating a 10-ps-long ignition pulse. Therefore, researchers are restricted to scaled-down simulations at present. Figure 27 shows a plot from a recent 2D PIC simulation, which employs 280 million particles on a 12,000 × 12,000 grid (corresponding to a region of 100 μm × 100 μm) for 60,000 steps (corresponding to 1 ps).

3.5 Office of High Energy Physics

The mission of the High Energy Physics (HEP) program is to understand the universe at a fundamental level by investigating the elementary particles that are the basic constituents of matter and the forces between them. Our understanding of how matter interacts requires a significant interplay among theory, computation, and experimental efforts. In many instances, the significant discoveries focus our computational efforts. In the following sections we discuss the primary areas of HEP research taking place at NERSC. These are computational studies of quantum chromodynamics (QCD), simulations of accelerator design, and examples from experimental physics.

3.5.1 Accelerator Physics

Background

Particle accelerators have enabled remarkable scientific discoveries and important technological advances that span all programs within the DOE Office of Science. In the HEP and Nuclear Physics (NP) programs, experiments associated with high-energy accelerators have led to important discoveries about elementary particles and the fundamental forces of nature, quark dynamics, and nuclear structure. In the Basic Energy Sciences (BES) program, experiments with synchrotron light sources and spallation neutron sources have been crucial to advances in the materials and biological sciences. In the Fusion Energy Sciences (FES) program, great strides have been made in developing heavy-ion particle accelerators as drivers for inertial fusion energy. The importance of accelerators to the DOE SC mission is evident from an examination of the DOE planning report *Facilities for the Future of Science: A Twenty-Year Outlook.*[40] Of the 28 facilities listed, 14 involve accelerators.

Particle accelerators and their associated detectors are among the largest, most complex scientific instruments in the world. The successful development of large accelerator facilities involves enormous investments in the three paradigms of scientific research: theory, experiment, and simulation. Neglecting any one of these can lead to an inability to meet performance requirements, to cost overruns, and ultimately to project failure. Simulation is already playing an increasingly prominent and valuable role in the theory, design, and development of accelerators and their associated technologies. Simulation is being used to understand and improve the performance of existing accelerators such as the PEP-II B-Factory at the Stanford Linear Accelerator Center (SLAC), the Fermilab Tevatron, the Fermilab Booster, and the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. In the future, a new level of simulation will be needed as researchers push the frontiers of beam intensity, beam energy, and system complexity. The three-dimensional, nonlinear, multi-scale, many-body aspects characteristic of future accelerator design problems, and the complexity and immensity of the associated computations, account for their extreme technical difficulty. But despite the difficulty, the corresponding return to science and to the DOE SC mission is equally significant. High-end simulations, used in concert with theory and experiment, will help to ensure the success of next-generation accelerators, providing optimal design while reducing design time, cost, and risk. Used as a tool of discovery, high-end simulations are helping researchers to explore systems under extreme conditions of high energy density. Simulations will furthermore help advance the frontiers of accelerator science and technology, leading to ultra-high-gradient laser- and plasma-based accelerators that will have huge consequences in scientific research, industry, and medicine.

Great progress in accelerator modeling has been made under the SciDAC Accelerator Science and Technology (AST) project. Thanks to SciDAC (and its predecessor, the DOE Grand Challenge program), a new suite of parallel simulation tools has been developed and applied to important problems in accelerator design and analysis with great success, laying the groundwork for the next advancement in simulation capability. NERSC has

[40] *Facilities for the Future of Science: A Twenty-Year Outlook,* U.S. Department of Energy Office of Science, November 2003.

been a key computational resource for the AST project. NERSC resources have been used to model existing machines such as the Fermilab Tevatron and Booster, the SLAC PEP-II B factory, and the RHIC. NERSC resources have also been used to simulate future accelerators such as the Large Hadron Collider (now under construction) and the International Linear Collider. They have also been used to study advanced accelerator concepts through simulations of experiments at SLAC (E-157, E-162, E-164) and at Lawrence Berkeley National Laboratory (L'OASIS laboratory). Though there are many examples of NERSC resources being brought to bear on important problems in accelerator science and technology, below we present four examples.

Accelerator Modeling at NERSC: Four Examples

Fermilab Booster simulation. Under the SciDAC AST project, codes such as BeamBeam3D, IMPACT, Synergia, and MaryLie/IMPACT have been developed and used to model beam dynamics in a number of important DOE SC accelerators and to study the physics of intense charged particle beams. In one such example, NERSC resources were used to perform simulations of the first 1000 turns (~2 ms) after injection of the Fermilab Booster, using a fully 3D space charge model. The simulation results were then compared with results from Booster experiments (Figure 28). Members of the SciDAC AST project participated in these experiments and made significant contributions in understanding and calibrating Booster diagnostic devices. Simulation studies are now being used to help optimize the injected beam parameters to achieve a matched beam in the presence of space charge and large momentum spread.

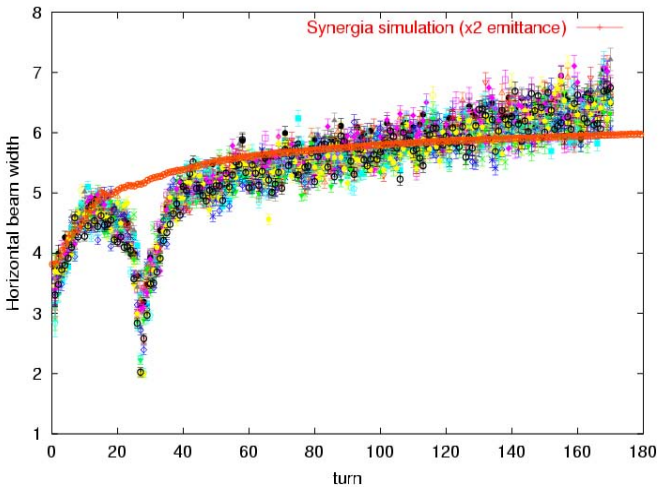

Figure 28. Comparison of Fermilab Booster experimental data with results from a Synergia application. The simulation was performed for 11 turns at 440 mA of current. The simulated beam size is in good agreement with the measured size. This simulation takes into account the calibration of the ionization monitors. (The notch in the beam is due to the decay of the current of the injection orbit bump magnet, which is not included in the simulation.)

Wakefield accelerator simulation. There is presently intense activity worldwide aimed at accelerating particles using the extremely large fields that can be produced in laser- and plasma-based systems. Previously such schemes produced beams with unacceptably large energy spread. However, recent breakthrough experiments have demonstrated acceleration of beams with just a few percent energy spread to energies of more that 100 MeV. Using NERSC resources, members of the SciDAC AST project have performed large-scale simulations using the codes OSIRIS and VORPAL to help interpret laser- and plasma-wakefield accelerator experiments and to gain insight in the acceleration mechanism (Figure 29). The highly nonlinear behavior of such systems in the "blowout regime" makes theoretical analysis especially difficult. Fully 3D parallel PIC simulations are

essential to help understand the physical processes involved and will play a key role in the design of future wakefield accelerators to produce high-quality beams accelerated to a GeV per acceleration stage. Under the SciDAC AST project, great progress has also been made in the development of reduced description models for laser/plasma simulation. One such example, QuickPIC, has been shown for some problems to provide answers as accurate as OSIRIS but with 2 to 3 orders of magnitude less computation time.

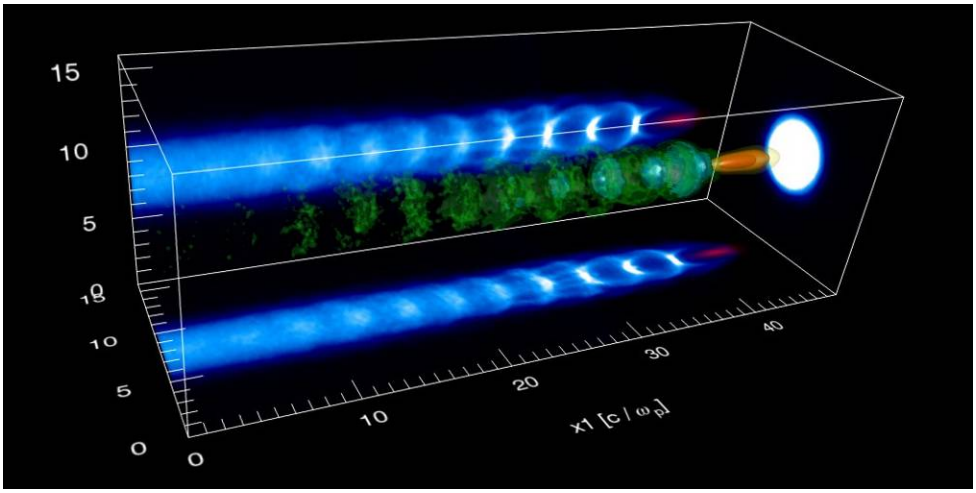

Figure 29. Results from a 3D OSIRIS simulation, performed at NERSC, of a beam with 2×10^{10} particles, $\sigma_r = 20$ μm and $\sigma_z = 63$ μm, propagating through a lithium gas cell of density $n_0 = 5.143 \times 10^{16}$ cm^{-3}. The simulation box was $512 \times 128 \times 128$ grid cells wide, and the simulation was run for 3100 time steps. This 3D plot of the electron density shows the ionization electrons and wakefield as well as the drive electron bunch. Both isosurfaces and projections of this data are shown: the drive beam in red-orange-yellow colors, and background electrons in blue-cyan-green colors. (Image: R. A. Fonseca and L. O. Silva, Instituto Superior Técnico, Portugal; S. Deng and T. Katsouleas, USC; and F. S. Tsung and W. B. Mori, UCLA)

Analysis of the PEP-II interaction region. Under the SciDAC AST project, a comprehensive suite of parallel, 3D electromagnetic modeling codes has been developed to model large, geometrically complex accelerator structures with unprecedented accuracy. These codes include Omega3P, S3P, Tau3P, T3P, and Track3P. The codes have been used to model electromagnetic structures for a number of DOE SC projects, including PEP-II, RIA, and a future linear collider. In regard to PEP-II, simulations were performed at NERSC using the codes Omega3P and Tau3P to analyze beam heating in the interaction region (IR) (Figures 30 and 31). Initially, PEP-II was prevented from operating at higher current due to the presence of trapped modes in the IR chamber that led to excessive heating. Omega3P was used to calculate the trapped modes in the IR from crotch to crotch. Figure 30 (left) shows the trapped mode with the highest power loss, whose frequency was close to the one measured in the experiment in the region of excessive heating. Figure 30 (right) is the heating power distribution centered about the interaction point, which shows the heating is significantly higher on one side, in agreement with what was observed. It was obtained by summing 330 modes found with Omega3P that have the highest wall loss. These results provided input on what modifications were necessary in the redesign of the IR to remedy this situation. Since the upgrade, PEP-II

has been able to raise the current by 15%, which led to higher luminosity and physics discovery.

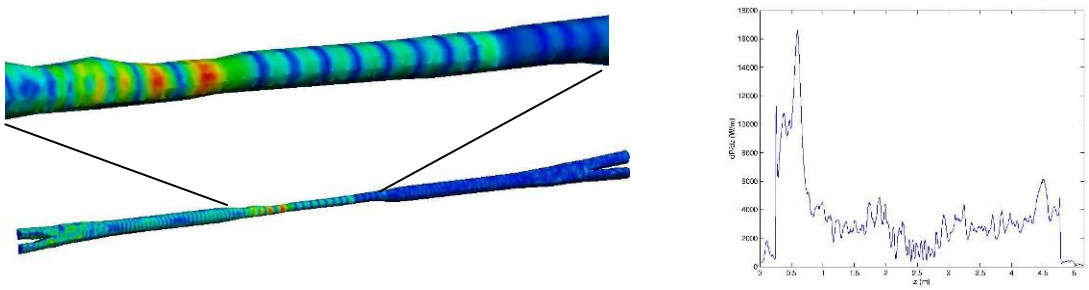

Figure 30. (Left) Trapped mode with highest power loss calculated by Omega3P (5.28 GHz, 230 W). (Right) Power loss distribution about interaction point (17.2 kW total from 330 modes).

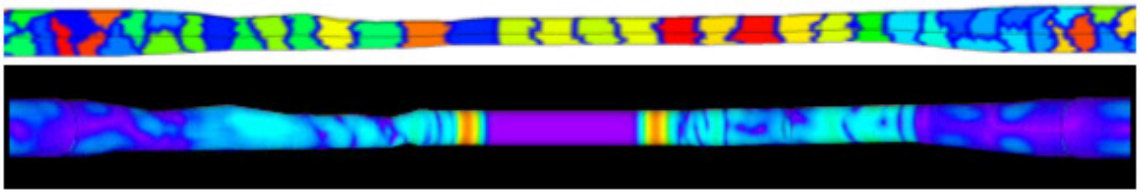

Figure 31. (Top) Distributed mesh of the IR between the crotches only. (Bottom) Snapshot in time of electric field due to two colliding beams from Tau3P time-domain simulation.

Phase space reconstruction. The applications scientists in the SciDAC AST project have close collaborations with researchers in the SciDAC Integrated Software Infrastructure Centers (ISICs) and with researchers supported by the DOE Office of Advanced Scientific Computing Research's Scientific Application Pilot Programs (SAPP) in areas such as applied mathematics, numerical analysis, computer science, visualization, and statistics. For example, statistical methods are being developed to aid in calibration, prediction, and uncertainty quantification in beam dynamics simulations. In particular, statistical methods are being developed to reconstruct a phase space distribution from one-dimensional wire scans. This involves solving an inverse problem in the presence of a small number of data samples which themselves have intrinsic errors. One approach to this problem involves performing a large number of forward simulations that, combined with beam profile monitor data, yield an estimate of the phase space distribution function. Figure 32 shows an example of the procedure, applied to the transport of a beam with space charge in a quadrupole channel. Successful application of this procedure will make it possible to determine the phase space initial distribution in experiments, including quantitative uncertainty bounds, and help make large-scale beam dynamics simulation a tool for prediction that can be used to help improve accelerator performance.

Future Resource Needs of the Accelerator Community at NERSC

Currently the accelerator community receives a combined allocation of approximately 2 million processor-hours on the IBM SP3 (Seaborg) at NERSC. But important problems have been identified, and CPU requirements estimated, where the simulations required to solve a single problem require the equivalent of nearly 10 million SP3 hours. These problems include, for example, the self-consistent simulation of beam-beam effects in colliders, and a design study for a 1 GeV plasma wakefield accelerator. Given the demand in

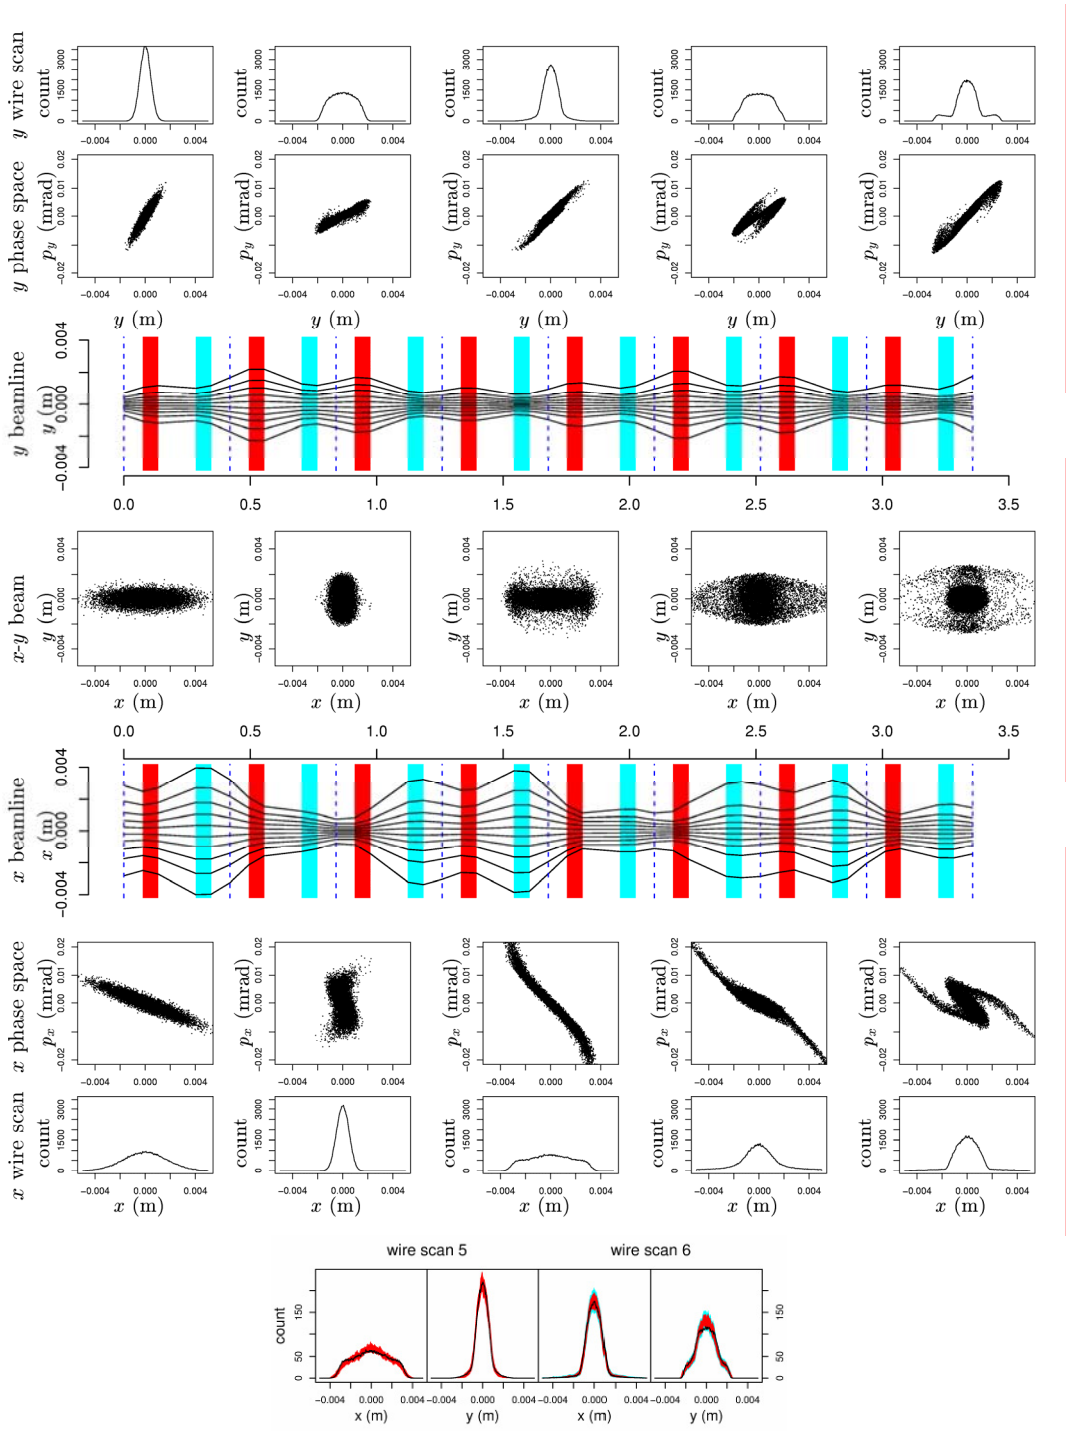

Figure 32. Simulation of a high-intensity proton beam through a series of quadrupole magnets. Working with the SciDAC AST Project, statistical scientists are developing new techniques for forecasting, calibration, and uncertainty quantification. Shown here, Bayesian techniques were used to combine 1D profile monitor data with simulations to infer the 4D beam distribution. The bottom figure shows the 90% intervals for the predicted profile at scanner #6 (shaded regions), and, for comparison, the observed data (black line). Only data from the odd-numbered scanners were used to make the prediction.

the accelerator community, the community would like to see as large an increase in CPU hours as is practical in regard to the next acquisition at NERSC. As to the type of resource, the community's codes have wide-ranging needs. Problems involving unstructured grids, parallel PIC space-charge simulations, and strong-strong beam-beam simulations involve significant interprocessor communication of varying sizes of data sets. As such, high bandwidth and low latency are both important, depending on the application. Other problems, such as beam-beam simulations in the weak-strong limit, involve little communication but require significant capacity. Besides raw CPU hours, some problems, such as direct Vlasov codes, require large memory, since the memory requirements scale as N^6 (for a 3D code). Given the massive number of grid points and/or particles, and the need to do visualizations and analyses, parallel I/O is also a high priority; this is especially true in regard to analyses involving the time evolution of complex phenomena such as beam-plasma interactions. Finally, the importance of algorithms cannot be overemphasized. Already under SciDAC, the AST project has seen performance increases of over a factor of 100 through the use of optimal algorithms and software provided by the ISICs. Continued progress in the development of efficient algorithms, performed at NERSC and elsewhere, will be crucial to getting the most science out of the next phase of NERSC resources.

3.5.2 Quantum Chromodynamics (QCD)

The long-term goals of high energy and nuclear physicists are to identify the fundamental building blocks of matter and to determine the interactions among them that lead to the physical world we observe. The DOE supports major experimental, theoretical, and computational programs aimed at reaching these goals. Remarkable progress has been made through the development of the Standard Model of high energy and nuclear physics, which provides fundamental theories of the strong, electromagnetic, and weak interactions. This progress has been recognized through the award of Nobel Prizes in Physics for the development of each of the components of the Standard Model: the unified theory of weak and electromagnetic interactions in 1979, and quantum chromodynamics (QCD), the theory of the strong interactions, in 1999 and 2004. However, our understanding of the Standard Model is incomplete because it has proven extremely difficult to determine many of the predictions of QCD. To do so requires terascale numerical simulations on four-dimensional space-time lattices.

The study of the Standard Model is at the core of the DOE's experimental programs in High Energy and Nuclear Physics. Major goals are to verify the Standard Model or discover its limits; determine the properties of strongly interacting matter under extreme conditions; and understand the internal structure of nucleons and other strongly interacting particles.

Lattice QCD calculations are essential to research in all of these areas, and NERSC resources have enabled significant progress in these calculations. Recent advances in algorithms and calculational methods, coupled with rapid increases in the capabilities of massively parallel computers, have created opportunities for major advances in the next few years.

A very large fraction of the computational resources consumed in lattice QCD calculations goes into Monte Carlo simulations that generate representative samples of the QCD ground state. These representative gauge configurations are saved and used to calculate a wide variety of physical quantities. In the last several years, Seaborg has been used to carry out some of the largest and most realistic simulations yet undertaken.

Because the generation of gauge configurations is so computationally expensive, it is important to maximize the physics that is obtained from them. To this end, a number of U.S. groups have made their configurations available for use by others, and the U.S. lattice QCD community as a whole is playing a leadership role in efforts to promote the international sharing of configurations. NERSC is playing an important part in this effort by providing hardware and system administration support for the first open lattice QCD archive, the Gauge Connection (qcd.nersc.gov).

NERSC resources are being used to study a broad range of topics with gauge configurations generated on Seaborg or at other supercomputer centers. Some highlights are described below.

Masses of strongly interacting particles. The determination of the mass spectrum of strongly interacting particles is one of the major goal of lattice QCD. The calculation of the masses of the lightest particles provides an important check on computational methods. Furthermore, quark model assignments of many excited states are not well established, and lattice calculations can help to pin them down. Lattice calculations of sufficient accuracy would also help experimental searches for particles that lie outside the naive quark model, such as glueballs, exotic mesons, and pentaquarks. Another important objective of spectrum calculations is to determine the masses of the quarks, the fundamental constituents of strongly interacting matter.

Weak decays of strongly interacting particles. Many strongly interacting particles decay via the weak interactions. A major component of the DOE's large experimental program in high energy physics is devoted to the study of such decays. One goal is to determine the elements of the CKM matrix, which describes how quarks are coupled to the weak interactions. The CKM matrix elements are fundamental parameters of the Standard Model, but some of them remain poorly known. In most cases the determination of a CKM matrix element requires both an experiment and an accompanying lattice QCD calculation. By determining a CKM matrix element through different experiments and lattice calculations, one can explore the range of validity of the Standard Model and possibly discover new physics beyond it. At the moment, it is uncertainty in the lattice calculations rather than in the experiments that limits the precision with which one can determine the least well known CKM matrix elements and test the Standard Model. So, improving these calculations is one of the major goals of lattice QCD.

Internal structure of strongly interacting particles. Strongly interacting particles are bound states of quarks and gluons, the fundamental entities of QCD. Understanding the internal structure of these bound states is an important objective of experimental programs at high energy and nuclear physics accelerator laboratories, and of lattice QCD

calculations. Lattice calculations are needed to determine the distribution of quarks and gluons within the nucleon, and to understand how these distributions give rise to the basic properties of nucleons and to the nuclear forces.

Strongly interacting matter at extreme conditions. The properties of nuclear matter at high temperatures and finite density are important for interpreting heavy ion collisions in progress at Brookhaven, and for understanding the development of the Universe immediately after the Big Bang. Under ordinary laboratory conditions, one does not observe quarks and gluons directly, only their bound states: protons, neutrons, and a host of other strongly interacting particles produced in accelerator experiments. However, at sufficiently high temperatures and/or densities, one expects to find a new form of matter consisting of a plasma of quarks and gluons. Lattice QCD calculations seek to understand the phase diagram of strongly interacting matter, to determine the equation of state of the quark-gluon plasma, and to calculate quark number susceptibilities which are related to event-by-event fluctuations in heavy ion collisions.

In addition to generating gauge configurations and to studying specific problems, such as those mentioned above, physicists have used NERSC resources to develop new algorithms and computational techniques. The development and testing of new methods for formulating QCD on the lattice has been of particular importance. These "improved actions" have already led to major advances and are likely to transform the field in the future. An indication of the current state of the art is given in Figure 33, where the ratio of a number of quantities calculated on the lattice to their experimental values is shown. The quantities shown in this figure are among the most straightforward to calculate on the lattice, but many other quantities of considerable importance to our understanding of fundamental physics are expected to be determined to similar accuracy in the near future.

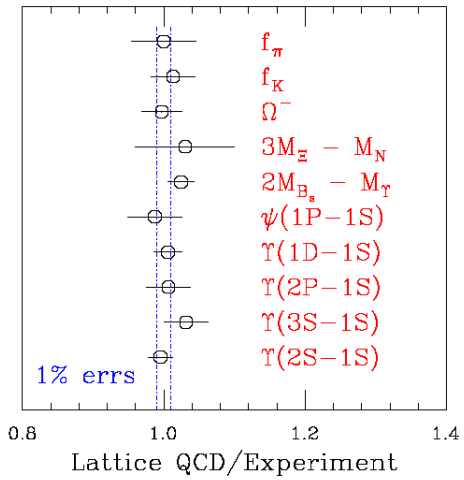

Figure 33. The ratio of several quantities calculated on the lattice to their experimental values. All results agree with experiment to within statistical and systematic errors which are estimated to be 3% or less. Some of these quantities involve only light quarks, some heavy quarks, and still others a mixture of the two. The quantities shown in this figure are among the most straightforward to calculate on the lattice, but many others are expected to be determined to similar accuracy in the near future.

The Department of Energy's High Energy Physics, Nuclear Physics, and Advanced Scientific Computing research programs have recently begun to fund dedicated computing hardware for lattice QCD. This hardware includes the QCDOC, a computer specially designed for lattice QCD, and clusters specifically optimized for it. We expect computationally intensive configuration generation to be moved to these machines once they are available for production work. However, the very flexible computers that NERSC has traditionally

provided will continue to be of great importance for measurements of complex quantities on these configurations and for the development and testing of new algorithms. The goal is to maximize the output of each type of computer by running on it the jobs for which it is most suitable.

3.5.3 Cosmic Microwave Background Data Analysis

Scientific Background

As the Universe expanded after its creation in a hot Big Bang, it cooled down, eventually dropping below the ionization temperature of hydrogen. At this point the previously free electrons were captured by protons, and the primordial photons, which until then had been scattering off them, began to stream freely. The continued expansion of the Universe then stretched the wavelengths of these photons, shifting the peak of their blackbody distribution from 3000 K at last-scattering to 3 K at detection. This radiation is the cosmic microwave background (CMB), and it provides the earliest possible image of the Universe.

A snapshot of the Universe only 400,000 years after the Big Bang, the CMB is a uniquely powerful probe of cosmology. Serendipitously first detected in 1965, the puzzling extraordinary uniformity of the CMB's temperature across the sky—when then-standard theory predicted that points more than a degree or so apart on the sky today would have been out of causal contact at last-scattering—ultimately inspired the current cosmological paradigm of inflation. However, the primary scientific value of the CMB lies in the tiny variations in its temperature and polarization, and in particular on the statistics of their power on different angular scales. After decades of searching, the temperature (T-mode) anisotropies were finally detected at the microK level on the largest angular scales by the COBE satellite in 1992, consistent with a cosmology in which tiny initial density perturbations are generated in the early Universe and act as gravitational seeds for the formation of the large-scale structure—from individual galaxies to clusters and superclusters of galaxies—that we observe today. Subsequent balloon-borne and ground-based observations in the late 1990s pushed this detection to sub-degree scales, revealing an angular power spectrum of fluctuations consisting of the harmonic series of peaks and troughs predicted by inflation, whose location in angular scale measured the geometry of the Universe to be flat, and whose relative amplitudes provided an independent (and happily consistent) measurement of the baryon fraction of the Universe. These observations also complemented other astrophysical data, particularly from Type 1a supernovae, which taken together led to the entirely unexpected conclusion that the Universe was dominated by some unknown dark energy, and that its matter component was dominated by some unknown dark matter. In 2003 the WMAP satellite provided a much higher-resolution spectrum of the large- and mid-scale fluctuations that sharpened all of the previous constraints on the dozen or so fundamental parameters of cosmology. The CMB E- and B-mode polarization anisotropies—orders of magnitude fainter than their temperature counterparts—carry even more information, but are only now beginning to be observed. Already, though, WMAP's measurement of the TE cross-spectrum suggests a surprisingly early reionization of the Universe.

Since the early Universe is the ultimate particle accelerator, in addition to constraining the cosmological models and their parameters, the CMB is a unique source of information about ultra-high-energy physics. Processes associated with physics beyond the current Standard Model must have occurred in the first moments of the Universe—in particular from the hypothesized fundamental-force fracturing by symmetry-breaking phase transitions—and will have left their imprint on the CMB. Inflation, dark energy, and dark matter must all be incorporated in any model of physics at such energies; for example, the very faintest CMB signal—the as-yet undetected B-mode polarization—has the extraordinary potential to tell us about gravity wave production during inflation, providing both the "smoking gun" for this whole paradigm and a powerful discriminant between models.

Realizing the full scientific potential of the CMB will require very high-resolution, very high-sensitivity observations of the temperature and polarization of the entire microwave sky. This challenge will be taken up first by the Planck Surveyor, a joint European Space Agency/NASA satellite mission scheduled for launch in 2007, and subsequently by NASA's CMBpol satellite, one of the five missions in its "Beyond Einstein" program. At the same time, a new generation of ground- and balloon-based experiments will survey smaller patches of the sky extremely deeply with the more restricted goal of detecting the mid- and small-scale polarization signals. The overall U.S. CMB research program, involving detector development, experimental design, mission operations, and data analysis, has been recognized as meriting a coordinated multi-agency collaboration, with DOE, NASA, and NSF all contributing their key skills and resources; and a roadmap for U.S. CMB research is currently being prepared by the Weiss Interagency Taskforce.

Computational Challenges

The fundamental challenge in making observations of the CMB is obtaining sufficient signal-to-noise to support well-resolved, high-significance estimates of the angular power spectra. Given the extraordinarily low level of the signal—$O(10^{-5})$ in temperature and $O(10^{-7})$ and lower in polarization—this necessarily requires large data volumes.

The analysis of CMB data typically proceeds in three steps: first, the time-ordered data is processed to generate pixelized maps of the sky components; then the sky maps are analyzed to estimate the underlying angular power spectra; and finally the power spectra are used to constrain cosmological parameters. At each step in the analysis, researchers must also carry forward an estimate of the statistical and systematic uncertainties in the data as the data are compressed. As researchers push toward ever fainter signals and higher resolution, the resulting datasets inevitably increase in size: Table 1 shows how the number of time samples, sky pixels, and spectral multipoles increases over time for past, present, and future CMB satellite missions.

Note that from COBE to WMAP and Planck, simultaneous increases in the sensitivity and resolution of the detectors allowed a huge increase in the number of sky pixels and spectral multipoles with a relatively modest increase in the number of time samples. However, we are now approaching the fundamental limits of detector sensitivity, so that CMBpol will have to gather many orders of magnitude more time samples (by increasing the number of detectors by many orders of magnitude) to support the increase in signal-to-noise per pixel needed to resolve the B-mode polarization.

Table 1. Numbers of samples, pixels, and multipoles for CMB satellite missions by launch date (all CMBpol entries estimated).

Mission	Launch Date	Time Samples N_t	Sky Pixels N_P	Spectral Multipoles N_l
COBE	1989	2×10^9	6×10^3	3×10^1
WMAP	2001	7×10^{10}	4×10^7	1×10^3
Planck	2007	5×10^{11}	6×10^8	6×10^3
CMBPol	est. 2015	$0(10^{14})$	$0(10^9)$	$0(10^4)$

CMB data analysis algorithms and their implementations fall into two broad classes:

- exact maximum likelihood methods scaling as $O(N_l N_p^3)$ with a small prefactor and reaching around 50% or more of peak performance
- approximate iterative/Monte Carlo methods scaling as $O(N_t)$ with a very large prefactor and reaching less than 5% of peak performance.

Moreover, the analysis invariably proceeds iteratively, with each iteration improving our understanding of the dataset. Given the dimensions of these datasets (and of their ground- and balloon-based counterparts) and the scalings, efficiencies, and iterations of the analysis algorithms, CMB data analysis has necessarily become a massively parallel computational endeavor.

Over the last eight years, NERSC has become the primary high performance computing center for CMB data analysis, currently supporting some 80 CMB data analysts from a dozen different experiments sharing an annual allocation that has risen to almost 1 million processor-hours. Scientific highlights of the work supported by NERSC to date include the first unambiguous determination of the large-scale geometry of the Universe from a long-duration Antarctic balloon flight by the BOOMERanG team (Figure 34), and the first analysis of simulated data from all of the detectors at Planck's key CMB frequency (Figure 35), which was performed on 6000 processors of Seaborg.

High Performance Computing Infrastructure for Planck

The near- and medium-term high performance computing infrastructure needs of the CMB community in general are dominated by the particular needs of the Planck mission for three reasons:

- Planck will generate the largest and hence most computationally challenging dataset (at least until CMBpol).
- Planck has the largest and most diverse team of collaborators and the widest distribution of computing resources to manage.
- Planck is a satellite mission and therefore has to satisfy the most rigorous operational protocols.

A detailed analysis of its data analysis pipeline has concluded that Planck's annual HPC needs will be:

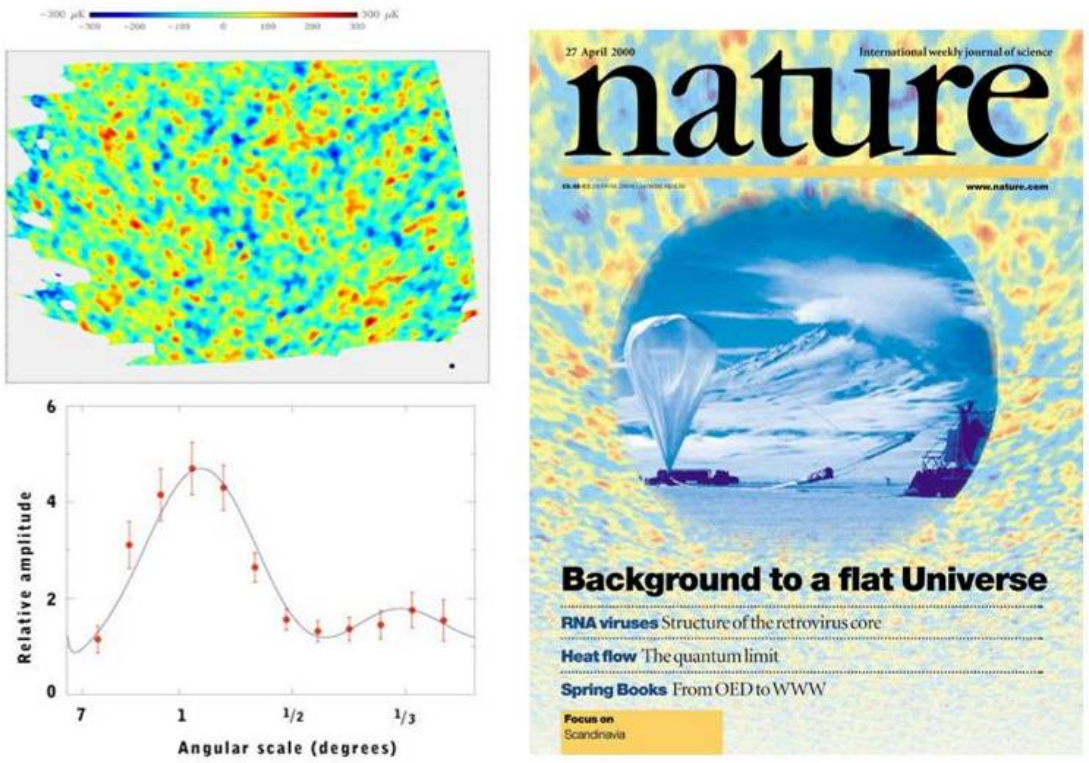

Figure 34. A 200,000-pixel map of the CMB temperature fluctuations over the patch of sky observed by BOOMERanG and its 12-bin recovered and full best-fit theoretical TT angular power spectra, together with the announcement of these results in Nature showing the balloon being launched in front of Mount Erebus.

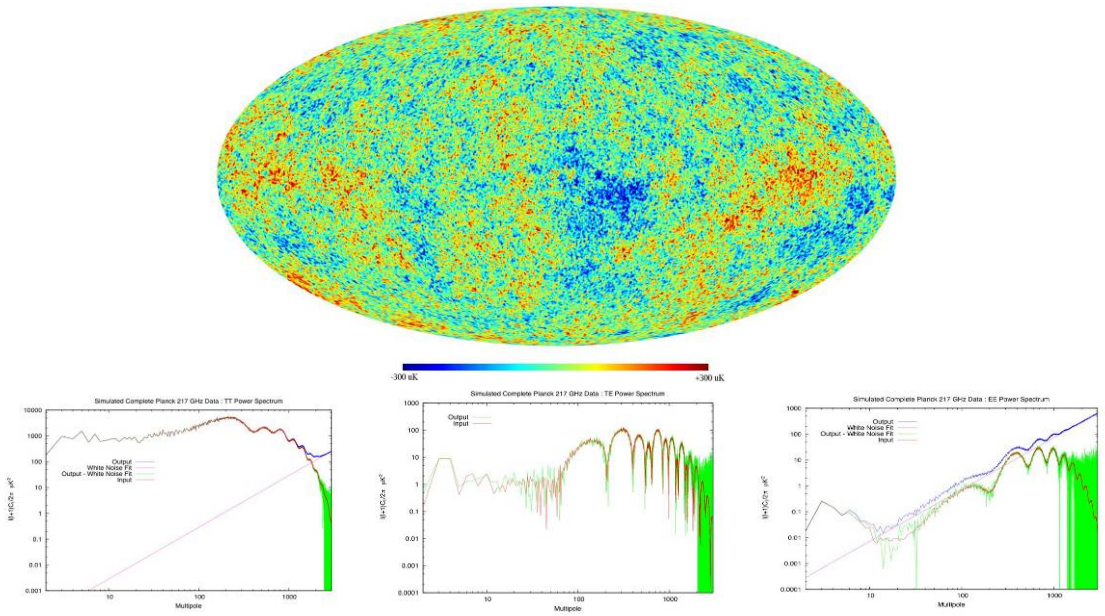

Figure 35. The 200,000,000-pixel full-sky temperature map, and the 3000-multipole input and recovered TT, TE, and EE angular power spectra, obtained from the first analysis of one year of simulated data from all 12 of Planck's 217 GHz detectors.

- O(10–100) exaflops of total processing capacity, primarily provided by NERSC's Seaborg and its successors and NASA Ames' Project Columbia.
- O(100) TB of archival file storage for primary data and derived data products.
- O(10) TB of scratch file storage at any one time to support a particular analysis, ideally simultaneously accessible from all of NERSC's machines (as in the General Unified Parallel File System model) and NASA's Project Columbia (as in the Wide-Area Network extension to GUPFS).
- O(1–10) GB of local tmp file storage on each processor or node to stage intermediate data products and enable out-of-core computations without having to repeat costly I/O subsystem calls.
- Fast (Gigabit or better) network connections between NERSC, NASA Ames, and NASA's Infrared Processing and Analysis Center (IPAC) in Pasadena, where the Planck data will first arrive from Europe, both for archival transfers and, between NERSC and Ames, for runtime I/O across a distributed file system.
- Scalable, massively parallel I/O supporting the simultaneous transfer of very large volumes of data across the entire processor set being used; currently much of the Planck-scale CMB data analysis is I/O bound. Stability with respect to the volume of all user traffic over the I/O subsystem is also highly desirable.
- An inter-processor communication system supporting the fast global reductions of gigabytes of distributed data.
- Since the majority of the floating-point operations in a CMB data analysis are within Fourier and spherical harmonic transforms, some mechanism for accelerating these would be beneficial too, provided its implementation did not inhibit the portability of top-level analysis codes.

Finally, as a satellite mission, Planck needs to have some certainty that it will have long-term (i.e., mission lifetime) access to any high performance computing resources on which it is basing its data analysis plans. To this end, scope for commitments to project allocations longer than a single year is highly desirable.

3.6 Office of Nuclear Physics

The mission of the Nuclear Physics (NP) program is to advance our knowledge of the properties and interactions of atomic nuclei and nuclear matter in terms of the fundamental forces and particles of nature. Broadly speaking, research efforts focus on understanding the ways in which matter interacts.

Our understanding of how matter interacts requires a significant interplay among theory, computation, and experimental efforts. The next few paragraphs outline the significant role of experimental physics in focusing computational problems. Several of the interesting computational problems being pursued today by NP researchers at NERSC will then be highlighted.

Nuclear Physics is a program of fundamental scientific research that provides new insight into the observed behavior of atomic nuclei. With the establishment of quantum chromodynamics as the fundamental theory of the strong nuclear interaction, the ultimate goal of nuclear theory is to understand nuclei and the nucleon in terms of their constituent quarks and gluons. It is at the highest energy scales that the nuclear quark and gluon aspects manifest themselves, and it is precisely these high-energy nuclear scales that are probed by the two nuclear physics flagship facilities, the Thomas Jefferson National Accelerator Facility (TJNAF) and the Relativistic Heavy Ion Collider (RHIC). More traditionally, nuclear theorists have understood the structure of the atomic nucleus most fundamentally in terms of interacting protons and neutrons. In certain regimes of energy and momentum transfer, this approach is not only valid but has been very successful in understanding nuclei of ordinary matter when advancing computational power has allowed much more detailed descriptions of nuclei on this microscopic level.

Computational research in NP aims to simulate the physics of the research described in the previous paragraphs. In the following sections, the primary areas of computational research utilizing NERSC resources will be discussed: nuclear structure, supernova physics, and nuclear matter at extreme energy density.

3.6.1 Nuclear Structure

The nucleus is a fascinating quantum mechanical system exhibiting diverse and rare phenomena. Governed by the strong interactions between nucleons, nuclei exhibit strong correlations resulting in both single-particle and collective modes of excitation; examples of the latter include Goldstone modes such as rotation and tunneling between spherical and deformed intrinsic nuclear configurations. Nuclear theory attempts to understand these excitations and the response of nuclei to diverse external probes within a coherent framework. This framework must encompass a wide range of energy and momentum scales for nuclei ranging from the deuteron to the superheavy elements. Nuclear theorists strive to describe the structure and dynamics of these often-disparate systems and to apply our knowledge of these systems to help unravel mysteries of our Universe.

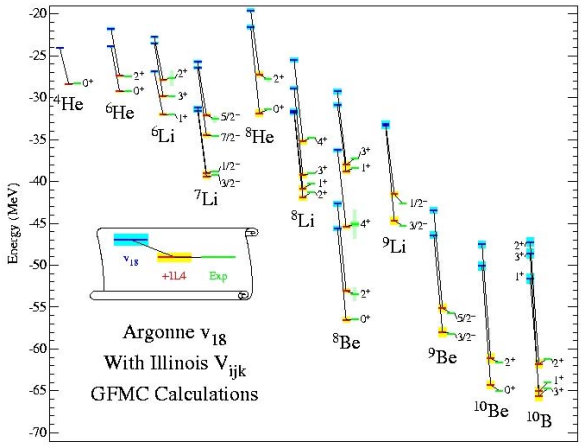

Figure 36. GFMC calculations of mass 4–10 systems.

One of the primary goals of nuclear physics is to explain the properties and reactions of nuclei in terms of interacting nucleons (protons and neutrons). There are two fundamental aspects to this problem: (1) determining the interactions between nucleons, and (2) given the interactions (i.e., the Hamiltonian), making accurate calculations of many-nucleon systems. The Argonne/Urbana/Los Alamos group uses NERSC computing resources to evaluate six- through ten-nucleon systems with realistic interactions also developed by that

group (Figure 36).[41] The accuracy of these calculations is at the 1% level for the binding energies. The resulting wave functions can be used to compute properties measured at electron and hadron scattering facilities (in particular TJNAF), and to compute astrophysical reaction rates, many of which cannot be measured in the laboratory.

Last year the Green's Function Monte Carlo (GFMC) group studied the effects of new three-nucleon potential terms on multiple states of the same angular momentum and parity.[42] Previously the GFMC group had demonstrated that the Hamiltonian that has been used successfully for more than a decade in studies of *s*-shell nucleons is inadequate in the *p* shell. Some of the possible new potential terms, whose forms are derived from meson-exchange arguments, result in considerable additional complications in the Green's function propagator. New potential models were developed that reproduce all of the known stable or narrow-width levels of up to ten-body nuclei with an average error of only 300 keV.

This project uses GFMC and variational Monte Carlo (VMC) methods to compute ground state and low-lying excited state expectation values of energies, densities, structure functions, astrophysical reaction rates, etc., for light nuclei and neutron drops. Realistic two- and three-nucleon potentials are used. The group is developing new computational techniques, optimizing them for different computer architectures, and improving the nuclear Hamiltonian used in the calculations. An area of increasing interest is the use of GFMC- or VMC-generated wave functions to compute reaction rates of astrophysical interest.

Other approaches to the nuclear many-body problem involve expansion in a basis of single-particle states. Much in the same way that chemists expand single-particle wave functions in the atomic-orbital basis, nuclear physicists typically invoke an oscillator basis in which to perform either Hamiltonian diagonalization or coupled-cluster calculations, as discussed below.

The no-core-shell model (NCSM) approach involves diagonalization of a very large dimensional matrix (typically with a rank of 1 billion) using Krylov space techniques (Lanczos or Davidson-Liu algorithms). These demanding calculations begin from a bare Hamiltonian consisting of two-nucleon and three-nucleon interactions. The implementation of a many-body model space requires that the interactions be renormalized to the model space in which one performs the diagonalization. This process of "decimation" removes the hard-core potentials normally associated with the bare nucleon-nucleon interaction and yields an effective interaction appropriate for the model space in which calculations are to be performed. The diagonalization procedure yields the energy spectrum of the nucleus, and the resulting wave functions may be used to obtain information on transition properties of the Hamiltonian. This method is currently being applied (using only two-body interactions) to ^{16}O (Figure 37). Recently this method was applied to neu-

[41] See, for example, Wiringa et al., Phys. Rev. C **62**, 014001 (2000); Pieper et al., Phys. Rev. C **64**, 014001 (2001); Pieper et al., Phys. Rev. C **66**, 044310 (2002).
[42] S. C. Pieper, R. B. Wiringa, and J. Carlson, Phys. Rev. C **70**, 054325 (2004).

trino-nucleus scattering.[43] The neutrino scattering cross section on ^{12}C remains one of the only measured neutrino-nucleus events. This weak interaction requires very accurate descriptions of nuclear ground- and excited-state wave functions in order to be accurately obtained from theory. The NCSM calculations included both two- and three-nucleon interactions and, with no adjustable parameters, recovered more than 60% of the experimental cross section from a theory that included *no* adjustable parameters. This represents a significant advance in our understanding of neutrino-nucleus scattering cross sections on nuclei.

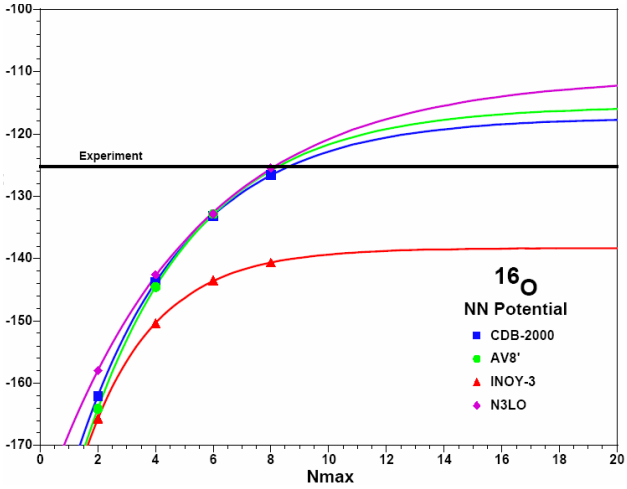

Figure 37. No core shell model results the ^{16}O ground state using two-body interactions.

While the NCSM diagonalizes a Hamiltonian in a given many-body basis, coupled-cluster theory obtains information on nuclear correlations through the solution of coupled-nonlinear algebraic equations for particle-hole excitation amplitudes generated by a given effective Hamiltonian appropriate for the chosen single-particle model space.[44] Recent large-scale applications of the method including non-iterative triples corrections to ground- and excited-state energies yielded converged results for the ground-state of ^{16}O and for the first excited negative-parity J = 3 state (the positive parity J = 0 state is not well described at this level of the theory). Figure 38 plots both ground-state and excited-state energies as a function of increasing model space. The nuclear two-body interaction was taken from chiral effective field theory. We know from GFMC and NCSM calculations in mass 10 systems that three-body forces significantly affect spin-orbit splitting of the single-particle states. The negative-parity J = 3 state is essentially a one-particle-one-hole state and should be well described by coupled-cluster techniques. We can therefore surmise that three-body forces will also play a significant role in determining the nuclear structure of these systems. Therefore in the future, we will necessarily need to develop coupled-cluster algorithms that incorporate three-body forces.

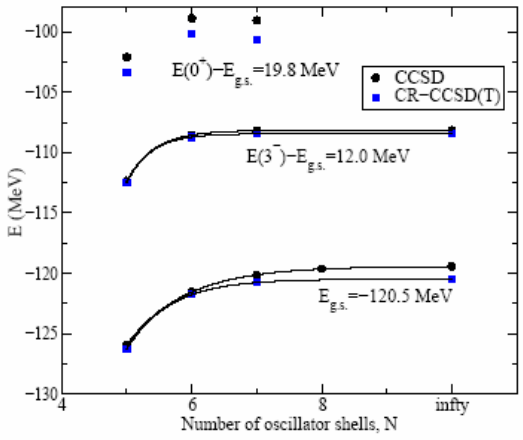

Figure 38. Nuclear coupled-cluster results for ground-and excited states of ^{16}O.

[43] A.C. Hayes et al., Phys. Rev. Lett. **91**, 012502 (2003).
[44] The application of this technique to closed-shell nuclei was described in K. Kowalski et al., Phys. Rev. Lett. **92**, 132501 (2004).

Nuclear structure studies in medium mass nuclei require an implementation of auxiliary field Monte Carlo (AFMC) techniques to calculate ground state and thermal properties of these systems. The calculation of the heat capacity of a finite interacting nuclear system including correlations is a difficult problem. The residual interaction among nucleons generates correlations that must be taken into account. The interacting nuclear shell model is an appropriate framework for such calculations, but very large model spaces are necessary to obtain reliable results. AFMC methods developed for the nuclear shell model enable zero- and finite-temperature calculations in large model spaces. One interesting effect involves breaking of pairs and the melting deformation as temperature is increased in a nucleus. Recent calculations suggest that peaks will appear in the nuclear specific heat (around T = 0.5–1.0 MeV) when pairing correlations are destroyed by temperature. Deformations wipe out this effect.[45] Figure 39 shows deformation plots for several nuclei in the mass 70–80 region with 40 neutrons. In the figure, β describes how deformed the nucleus is, while γ indicates whether the nucleus is prolate (γ = 0), oblate (γ = 60), or soft (uniformly filling in γ). ^{68}Ni and ^{70}Zn are fairly spherical nuclei, while ^{72}Ge exhibits shape-coexistence. ^{80}Zr is very well deformed (and prolate), as seen in the figure. At intermediate temperatures (middle column), all nuclei retain some of their ground-state character, but the shape distributions become softer in the γ-direction. At the higher temperature of T = 2.0 MeV, the four N = 40 systems considered become soft in both β and γ, filling uniformly the phase space in the β–γ plot.

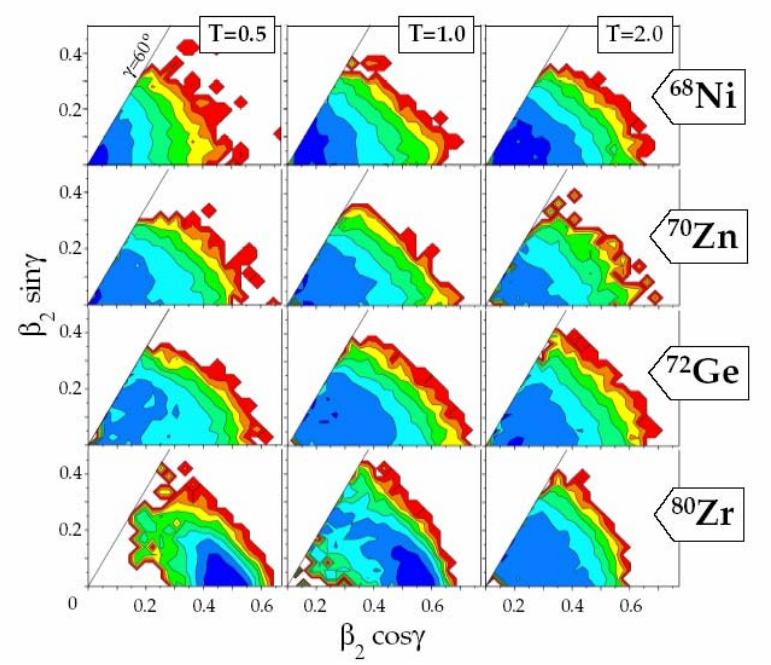

Figure 39. Nuclear deformation as a function of temperature in the mass 80 region.

Electron capture on proton nuclei is an important process that occurs during the early stages of the evolution of core-collapse supernovae. These weak captures serve to deleptonize the core of the massive star and determine the final electron fraction within the core. Therefore they set the size of the homologous core that must then collapse. The size of the core influences the amount of energy required to explode the star. Accurate nuclear structure calculations are required to understand these processes, and such calculations also require significant computational resources. While the iron region has been suffi-

[45] K. Langanke, D. J. Dean, and W. Nazarewicz, "Thermal properties of *N* = 40 isotones," Nucl. Phys. A (in press).

ciently investigated, the distinct possibility exists for electron capture to occur in nuclei beyond this region. Historically, this possibility has not been included in core-collapse simulations, where most capture took place on protons. Inclusion of newly calculated electron-capture rates (calculated with a combination of AFMC and random phase approximation theories) on nuclei indicate that nuclei up to mass 120 dominate electron capture on protons, thus leading to significant modifications of the evolution of core collapse models.[46]

Computational Needs

Calculations of high priority during the next five years include the following: (1) GFMC needs 150×10^{15} operations for one state in ^{12}C when the nuclear three-body interaction is included in the calculation. For something of both astrophysical and nuclear interest, calculation of the first excited 0^+ will require 1–2 million Seaborg processor hours. Upgrades to the processor speed, while keeping memory per processor at current or enhanced levels, is important for this effort. (2) The no-core shell model approach would run optimally with 4 GB of memory per processor. The NCSM effort will require approximately 5 million Seaborg-equivalent processor hours at NERSC to continue production runs through mass 16 nuclei. These runs will also include realistic three-nucleon forces. While some of these cycles will be obtained from other resources, NERSC remains an important part of the NCSM effort. (3) The nuclear coupled-cluster calculations would benefit from an increased memory per processor (to 8 GB) and an increased processor speed. Current plans include calculations of mass 40 nuclei and the inclusion of three-body terms into the Hamiltonian. Compared with today's efforts (mass 16 nuclei in eight oscillator shells), mass 40 nuclei will require at least nine oscillator shells for convergence. While the largest ^{16}O runs (with two-body interactions) cost 2 Tflop-hours, the largest ^{40}Ca runs will require 36 Tflop-hours. For the same number of basis states, we estimate that an ^{16}O calculation including a three-body interaction will require roughly 500 Tflop-hours per calculation. Memory needs for a three-body interaction are also significant. For an eight-oscillator shell run, the memory required to store the three-body interaction matrix elements is well over 10 TB.

3.6.2 Supernova Physics

NERSC resources have been brought to bear on two of the most significant problems in computational astrophysics: (1) understanding core collapse supernova explosions and the production of many of the elements in the Universe, and (2) understanding thermonuclear supernovae and their use as standard candles in illuminating the evolution of the Universe and its ultimate fate.

Core Collapse Supernovae

Explosions of massive stars (stars more massive than ten of our suns), known as core collapse supernovae, are arguably the most important link in our chain of origins from the Big Bang to the formation and evolution of life on Earth. They are the dominant source of most elements in the Periodic Table between oxygen and iron, and are believed to be responsible for producing half of all elements heavier than iron. These explosions also

[46] Langanke et al., Phys. Rev. Lett. **90**, 241102 (2003); Hix et al., Phys. Rev. Lett. **91**, 201102 (2003).

serve as “cosmic laboratories” for nuclear and particle physics at extremes that may be inaccessible in terrestrial experiments. Observations of supernova neutrinos (nearly massless, radiation-like particles), gravitational waves (ripples in space predicted by Einstein’s theory of gravity), and photons across the electromagnetic spectrum, in conjunction with realistic three-dimensional models, will provide windows into these cosmic laboratories.

The principal goals of core collapse supernova research are (1) to understand the mechanism(s) responsible for the explosions of massive stars; (2) to understand all of the phenomena associated with these stellar explosions, such as their contribution to the synthesis of the elements, their emission of neutrinos, gravitational waves, and, in some cases, intense bursts of gamma radiation; (3) to provide the theoretical foundations supporting the scientific mission of the Office of Science’s existing and proposed premier experimental facilities, such as the Relativistic Heavy Ion Collider (RHIC), the Rare Isotope Accelerator (RIA), and the Deep Underground Science and Engineering Laboratory (DUSEL), whose scientific missions are in part defined by supernova science; (4) to develop the methods to simulate three-dimensional, multiangle, multifrequency, precision radiation transport on terascale to petascale computers; (5) to develop the theory and methods to predict, using terascale computers, the physical states of the complex nuclei found in stars involved in supernova explosions; and (6) to serve as a testbed for the development of enabling technologies such as data management and analysis, networking, and visualization, of relevance to many applications.

As their name suggests, core collapse supernovae result from stellar core collapse and the formation of a shock wave that is ultimately responsible for the explosion. They are radiation- (neutrino-) driven, and perhaps also magnetically driven, turbulent events. One-, two-, and three-dimensional simulations that include different physics components —neutrino transport, turbulent fluid flow, rotation, magnetic fields, gravity, weak (neutrino) interactions, equations of state (thermodynamic state of the stellar core)—must be performed systematically in order to understand the role of each physics component in the explosion. NERSC resources have therefore been used to perform a variety of simulations. These include three-dimensional simulations focusing on the hydrodynamics of the stellar core during the explosion and the development of a newly discovered shock wave instability (by John Blondin, NCSU, and Tony Mezzacappa, ORNL) (Figure 40), and the first realistic two-dimensional supernova simulations, which include realistic two-dimensional multi-frequency flux-limited diffusion neutrino transport (by Doug Swesty and Eric Myra, SUNYSB) (Figure 41). The latter are currently exploring the interaction of the neutrino radiation and stellar core fluid flow with an eye toward the development of fluid instabilities and turbulence and their impact on the explosion, and will continue in order to explore whether or not a supernova explosion can be powered by neutrinos and turbulent fluid flow alone (i.e., *sans* rotation and magnetic fields). NERSC resources will also be used to perform the first contemporary simulations to include realistic neutrino transport, rotation, and magnetic fields, with an eye toward understanding the role of magnetic fields in the explosion mechanism (by John Hayes, UCSD, Steve Bruenn, FAU, and Tony Mezzacappa, ORNL). At the moment, the supernova community has very little guidance from detailed computational models on this issue.

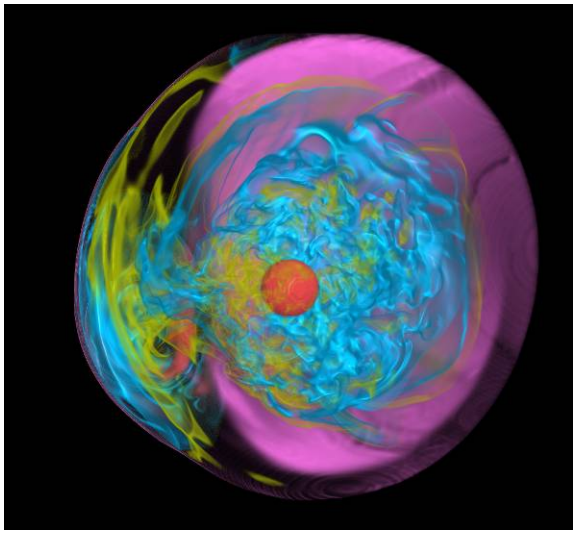

Figure 40. Visualization of the supernova shock wave instability by K.-L. Ma (UCD) based on a three-dimensional simulation performed by J. M. Blondin (NCSU) under the SciDAC-sponsored TeraScale Supernova Initiative led by A. Mezzacappa (ORNL).

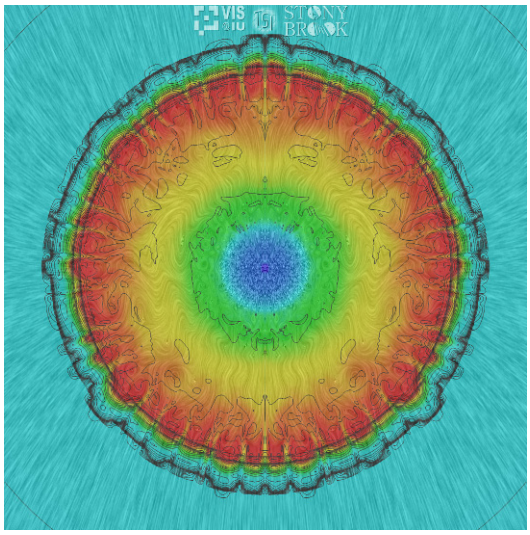

Figure 41. Visualization by E. Bachta and P. Baker (IU) of the development of turbulence in the stellar core and isocontours of neutrino "optical depth" in a two-dimensional supernova simulation performed by F. D. Swesty and E. Myra (SUNYSB) as part of the SciDAC-sponsored TeraScale Supernova Initiative.

Thermonuclear Supernovae

Unlike their core collapse cousins, thermonuclear supernovae signal the death of a star not much larger than our sun. The star, most likely a degenerate carbon-oxygen white dwarf, is in a binary system and has been accreting matter from a nearby companion and growing in mass for several million years. The star can entirely support itself by electron degeneracy pressure up to the Chandrasekhar mass (about 1.4 times the mass of our sun). Within one percent of this mass, a unique conspiracy in nuclear physics occurs such that the density and temperature at the center reach the point where the carbon ignites in a runaway thermonuclear explosion. The entire star is destroyed in this explosion, with material reaching velocities greater than a tenth of the speed of light. Over one-third of the star is burned to radioactive nickel, which then decays to cobalt, and the cobalt to iron. The energy released from this chain of radioactive decay powers the optical light from the supernova for several months. At its peak, the supernova can be as bright as the light from all the rest of the stars in the galaxy where it was born.

Because these supernovae, known as Type Ia supernovae, explode in a very similar manner every time, it is not surprising that their peak brightnesses are very similar. Over the past decade astronomers have been exploiting this fact, using a correlation between the shape of the optical lightcurve and the peak brightness to measure distances to better than 10% to these objects. Since Type Ia supernovae are so bright they can be seen across most of the observable Universe, distances to them along with their redshifts can be used to measure the expansion history of the Universe and constrain the cosmological parameters.

For the past several years at NERSC, researchers have been modeling the spectra of supernovae of all types. The goal of this work is to better understand the physics of supernovae and, for the Type Ia's, to improve their utility as cosmological tools. Spectrum synthesis is one of the only ways to validate or falsify detailed hydrodynamic models of the explosion event. During the past year, researchers have concentrated on the 3-D spectrum synthesis of Type Ia supernovae in order to probe the explosion mechanism behind these objects. Unlike core collapse supernovae, the progenitors of these supernovae have never been observed. Thus modeling may be the only way to gain insight on how Type Ia supernovae explode.

Kasen and Plewa have considered some of the spectral and polarimetric signatures of the gravitationally confined detonation scenario for Type Ia supernova explosions.[47] In that model, material produced by an off-center deflagration forms a metal-rich atmosphere above the white dwarf surface. Using hydrodynamical simulations, they show that this atmosphere is compressed and accelerated during the subsequent interaction with the supernova ejecta, eventually leading to the formation of a high-velocity pancake of metal-rich material that is geometrically detached from the bulk of the ejecta. When observed at the epochs near maximum light, this absorbing pancake produces a highly blueshifted and polarized calcium IR triplet absorption feature similar to that observed in several Type Ia supernovae (Figure 42).

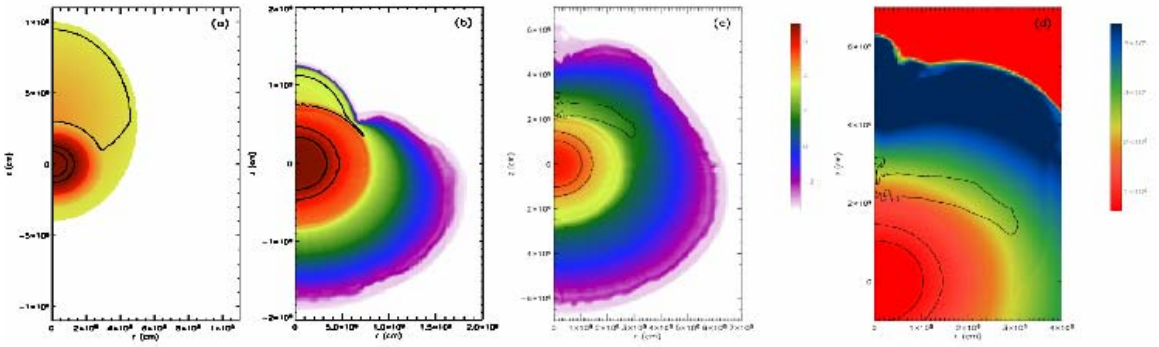

Figure 42. Hydrodynamical simulation of the model supernova ejecta interacting with the extended atmosphere. Panels a–c show the density in log scale; panel d depicts the magnitude of the velocity. The density scale is shown by the color bar in panel c, while the velocity scale is shown by the color bar in panel d. The black contour line corresponds to a calcium number abundance of 0.01. (a) Density distribution at the beginning of the simulation (t = 0 s). (b) Density distribution at t = 0.3 s. Notice the remarkable deformation of the calcium-rich material in the upper part of the computational domain. (c) Density distribution at the final time (t = 1.24 s). Notice that the calcium-rich region has been strongly compressed into a pancake-like structure. (d): Velocity magnitude at the final time. The calcium-rich absorber is seen moving at velocity ~21,000 km/s with a substantial velocity gradient across the structure.

The synthetic spectrum of the above model at a time when the supernova is near maximum light has been obtained after homologously expanding the ejecta to 20 days (Figure 43). Given the enhanced abundance of intermediate mass elements and relatively low

[47] D. Kasen and T. Plewa, "Spectral signatures of gravitationally confined thermonuclear supernova explosions," Astrophys. J. Lett. **622**, L41 (2005).

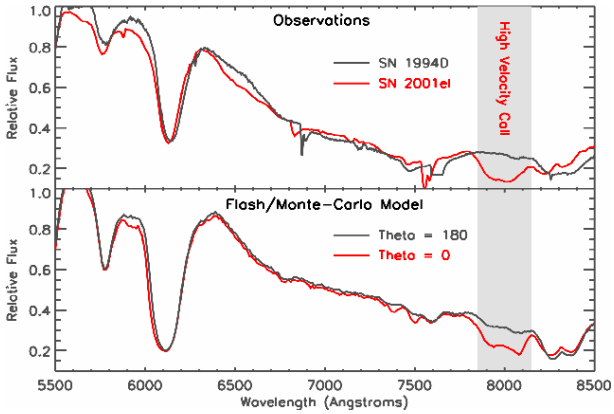

Figure 43. Comparison of the synthetic model spectra to Type Ia supernova observations. Top panel: Spectral observations of two Type Ia SNe near maximum light. SN 2001el (red line) shows a strong HV calcium absorption at 8000 Å, while SN 1994D (black line) does not. Bottom panel: Synthetic model spectra at 20 days. A HV calcium absorption is clearly seen when looking straight down on the calcium-rich pancake ($\theta = 0°$, red line), whereas none is seen from the opposite side ($\theta = 180°$, black line).

temperature, the pancake is opaque in the Ca II IR triplet lines. For a viewing angle in which the observer looks directly down upon the pancake ($\theta = 0°$), the pancake obscures the supernova photosphere, creating a broad and highly blueshifted high-velocity (HV) calcium absorption feature near 8000 Å. The HV feature seen in the model compares well to that observed in SN 2001el. For larger viewing angles, the pancake obscures less (or none) of the photosphere, and the HV feature is weaker or absent in the model spectrum. This orientation effect may explain why in some supernovae, such as SN 1994D, a HV calcium feature at maximum light is seen only weakly or not at all. As this hydrodynamical model is the first one which naturally goes from a deflagration to a detonation, which many believe is a key aspect for a successful Type Ia explosion, observations which constrain the geometry of a Type Ia supernovae may be the critical for understanding why they explode and how far they can be pushed as cosmological probes.

Terascale supernova simulations require a terascale applied mathematics and computer science infrastructure, including (a) solvers for terascale algebraic equations at the heart of the solution of the neutrino transport equations; (b) parallel I/O scalable to thousands of processors; (c) workflow management and automation; (d) visualization of terabytes of data; and (e) networking at hundreds of Mbps to Gbps to enable the research of a nationally distributed scientific team. The demands will escalate rapidly over the next few years. Three-dimensional supernova simulations with realistic three-dimensional multifrequency, multiangle neutrino transport will require sustained computational speeds in excess of one hundred teraflops. Moreover, these simulations will require platforms with significant memory bandwidth, memory per processor, and speed in processing global communications. In addition, they will generate hundreds of terabytes of data per simulation over a period of months, pushing the data management (e.g., parallel I/O) and analysis, networking, and visualization requirements yet another order of magnitude further.

3.6.3 Nuclear Matter at Extreme Energy Density

The first three runs at RHIC have allowed STAR to accumulate a wealth of data from $p + p$, d + Au, and Au + Au collisions on event-by-event observables; inclusive spectra for strange, nonstrange, and charmed mesons, baryons, and resonances; and the spectra of leading hadrons from hard-scattered partons. These studies have focused on investigating

strongly interacting matter at high energy density and the search for the quark-gluon plasma. These results have led quickly to the discovery of qualitatively new behavior, not previously observed, indicating the formation of a strongly interacting dissipative medium in central Au + Au collisions at RHIC. This discovery, observed by the other RHIC experiments as well, has been published in Physical Review Letters.[48]

Comparison of inclusive yields of high-p_T charged particles and the correlation of back-to-back leading charged hadrons has shown (Figure 44) that inclusive high-p_T particle production and "away side" leading hadrons are strongly suppressed in central Au + Au collisions relative to the yields in $p + p$, peripheral Au + Au, and central d + Au collisions. These studies show that the strong suppression (a factor of 5) observed in central Au + Au collisions is due to final state interactions in a dense, dissipative medium produced during the collision and not to the initial state wave function of the Au nucleus. The picture which emerges is that energetic partons traversing the medium produced in head-on Au + Au collisions lose sufficient energy that only those streaming outward from near the surface of the system produce the jets observed.

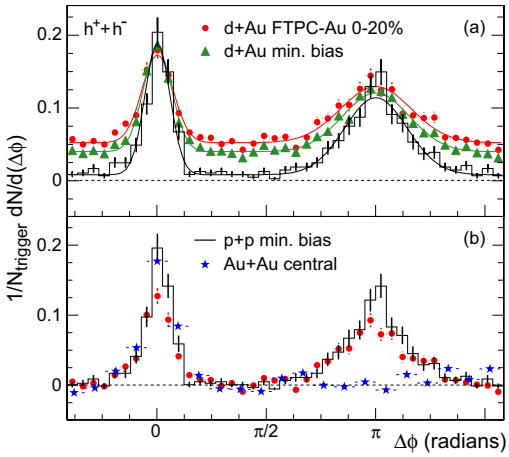

Figure 44. (a) Efficiency corrected two-particle azimuthal distributions for minimum bias and central d + Au collisions, and for $p + p$ collisions. (b) Comparison of two-particle azimuthal distributions for central d + Au collisions to those seen in $p + p$ and central Au + Au collisions. The respective pedestals have been subtracted.

These and other results from the initial survey performed by STAR have shown that the matter being produced exhibits features qualitatively different from those observed before in collisions of heavy nuclei. The following picture emerges.

The system is highly dynamic and the evolution is fast; characteristic features include:

- Transverse expansion with an average velocity of ~0.55 c.
- A large degree of anisotropic flow (v_2) suggesting hydrodynamic expansion and high pressure at early times in the collision history.
- The duration of hadronic particle emission appears to be very short.
- Near side correlations show that the fragmentation properties of the observed jets are the same in $p + p$ collisions and all centralities for Au + Au collisions, indicating that the fragmentation in Au + Au collisions occurs in vacuum.

The produced matter is opaque, exhibiting:

- Persistence of the saturation of v_2 at high p_T.

[48] J. Adams et al. (STAR Collaboration), "Evidence from d + Au measurements for final-state suppression of high-p_T hadrons in Au + Au collisions at RHIC," Phys. Rev. Lett. **91**, 072304 (2003).

- Suppression of high-p_T particle yields relative to binary scaled p + p.
- Suppression of away-side leading particles from jets.
- Large-scale correlations of net charge, total charge, and p_T.

Statistical models describe the final state well, as indicated by:

- Excellent fits to particle ratio data with equilibrium thermal models.
- Excellent fits to flow data with hydrodynamic models that assume equilibrated systems.
- Chemical freeze-out at about 175 MeV; thermal freeze-out at about 100 MeV.

The research programs in the next few years are planning to have focused studies to fully elucidate the fundamental properties of the new state of strongly interacting matter produced in RHIC collisions. These measurements will decisively establish the relation of this matter to the quark-gluon plasma state predicted to result from heating the QCD vacuum. They will further allow STAR to test and extend QCD theory and its predictions regarding the behavior of bulk color-deconfined matter.

Studies key to this program will include:

- Measurement of the gluon density of the plasma using direct-photon tagged jets.
- Measurement of flavor tagged jets to test perturbative QCD predictions of the quark mass dependence of partonic energy loss.
- Measurement of spectra and yields for the upsilon family of states to place significant constraints on the temperature in the initial stage of the collisions.
- Studying partonic collectivity by measuring bulk physics properties (e.g., spectra, elliptic flow, particle ratios, non-identical particle correlations) for particles and resonances containing light, strange, and charmed quarks.
- Detailed unfolding of large- and small-scale fluctuations and correlations for identified particles to map the dynamics and evolution of the produced matter.
- Studying the effects of chiral symmetry restoration via leptonic decays of hadronic resonances in-medium.
- Direct photon spectra via gamma-gamma HBT to provide information on the temperature and lifetime of the early time partonic and later stage hadronic phases using a penetrating probe.

The computational aspects of all these experimental studies have been and continue to be carried out primarily at two main computing facilities, the RHIC Computing Facility (RCF) at BNL and at NERSC. The RCF serves the role of raw data storage, first pass event reconstruction, and a significant share of the individual physics analyses. The NERSC role utilizes PDSF for production detector efficiency and simulation computations and also a significant share of the individual physics analyses, and utilizes HPSS in support of these very data-intensive computations.

STAR and PDSF have collaborated with several SciDAC projects (Particle Physics Data Grid, DOE Science Grid, Scientific Data Management Center) in recent years to establish an effective data grid connection between BNL and NERSC. As the volume and complexity of data have increased greatly, the amount of human effort in transferring data has been significantly reduced, as well as the failure rate, because of this extremely useful data grid effort. One of the key elements to the success of this effort is the participation of the Scientific Data Management group at Berkeley Lab, in close proximity to PDSF and STAR personnel.

Looking forward over the next few years, it is expected that the existing allocation model and procedures will continue. For PDSF users, this means the program agencies will fund the level of CPU and disk storage space needed at PDSF and they will continue to participate in the process for SRU (storage resource unit) allocations for HPSS. Beyond basic resource capacity needs, however, it is important that the data grid functionality continues to be enhanced and made easily available to all PDSF users. The most effective way of accomplishing this is to strengthen the participation of NERSC in the Open Science Grid.

The two primary NERSC resource capacities used are HPSS (SRUs) and network bandwidth to ESnet. It is anticipated that these will continue to grow at a rate of about 50% per year for the next five years. There are tentative plans for a RHIC II accelerator that may significantly increase the data needs around 2010.

4. Computational Resources at NERSC

The resources furnished by the NERSC High-Performance Computing Facility (HPCF) are the focal point of high-performance production computing for the DOE Office of Science scientific programs. These resources include an assortment of hardware, a large ensemble of installed software, network connectivity to sites that use the facility, and user support services.

4.1 Hardware

The NERSC HPCF now provides all of its computational resources via massively parallel processing (MPP) systems, including cluster systems, and specialized compute servers for visualization and mathematical software.

NERSC supports a diverse workload, but one which tends increasingly towards capability computing, which can require a large fraction of an entire computing resource

4.1.1 Massively Parallel Supercomputers

Distributed-memory massively parallel processor (MPP) systems are now common in scientific computing. For suitable problems, these systems are capable of very high computational performance, and of addressing capability computing problems that require very large numbers of numerical calculations and very large amounts of memory. NERSC has one large MPP system in production, and is actively preparing to evolve along with the MPP and cluster systems available in the open market.

The IBM SP (seaborg.nersc.gov) is one of the larger unclassified systems in existence. First installed in September 2001 and doubled in size in 2003, it is now composed of 380 compute nodes (6080 processors), with over 7 TB of memory and 44 TB of disk. This system has a peak performance of 1.5 Gflops per processor, for a theoretical peak of over 9.1 Tflops. Another system is being investigated that, when put into production, is expected to deliver over 2.8 Tflops and contain over 1.2 TB of memory.

4.1.2 Clusters

In the area of cluster computing, NERSC supports the Parallel Distributed Systems Facility (PDSF), a networked distributed computing environment used to meet the detector simulation and data analysis requirements of large-scale high energy physics and nuclear science investigations. PDSF is composed of 344 compute nodes with dual Pentium III or AMD 2100/2600 processors.

Knowledge gained from NERSC's past cluster work is being used in studying possible future systems.

4.1.3 High Performance Storage Systems

HPSS is a modern, flexible, performance-oriented mass storage system, designed and developed by a consortium of government and commercial entities. It is used at NERSC for data storage, archives, and system backups. The current configuration consists of two libraries, totaling over 8.8 PB of tape capacity, served by eight STK tape robots, 50 TB of disk cache, and multiple data movers and meta-data servers. The theoretical throughput is 3.2 GB/s.

4.1.4 Servers

NERSC operates servers that provide direct and indirect services to staff and users. The two most visible are a 12-processor SGI Onyx 3400 (escher.nersc.gov) for visualization processing and a four-node, eight-processor Linux system of AMD Athlon processors (newton.nersc.gov) for symbolic mathematics computing. A new SGI Altix system is being configured, which will replace the current visualization server in the next few months. A CVS repository server is also maintained to support software collaboration among groups of NERSC users.

NERSC also operates servers that are less directly visible, but nonetheless essential for the secure but accessible utilization of NERSC's resources. These systems support security applications (monitoring and intrusion detection), addressing and authentication services (LDAP), e-mail, account support databases (NIM), and Web-based documentation and information services.

4.2 Grid Computing

The Grid promotes a standard large-scale science environment that spans many different projects and provides the needed infrastructure, data transfer and storage resources, and computing services in a relatively uniform and supportable way. All NERSC production computing and storage systems are now Grid-enabled, and can be accessed by client software based on the Globus Toolkit. These capabilities are in daily use, especially in PDSF, where job submission and data transfers are routinely accomplished by Grid client-ware, and in HPSS, which has Grid-based data transfer capabilities.

4.3 Software

The software infrastructure is what a user truly interacts with, and it determines to a large extent the usability and security of any particular system. This complex includes distributed computing tools, operating systems, networking and access software, compilers and debuggers, performance tools, software libraries, job control and queuing systems, and third-party applications that are of importance to many users in the community.

4.3.1 System Software

The operating system and supporting software layers must be a robust software suite or the underlying hardware will not be well utilized. The operating system must allow users to manage files, run interactive and batch jobs, import and export data, compile and develop programs effectively, and support significant third-party applications. Network access software is a component of this, and its components also protect against unauthorized access and compromise of proprietary information. NERSC supports several operating systems on MPP and server systems, and interacts with vendors on their maintenance and development.

Standard programming languages and debugging and performance tools are essential components of any high-performance computing resource. The following languages are supported, some by multiple compilers, as part of a fully integrated, interoperable environment for compilation and debugging: Fortran 77/9x, C, C++, Java, assembler, and various shell and scripting languages. Tools to monitor and diagnose performance problems are also provided.

4.3.2 Basic Scientific Libraries and Environments

Communication libraries, such as MPI, are available and are augmented by tools to make their use amenable to debugging and performance tuning. Various I/O libraries such as netCDF, HDF, etc., support code portability and data exchange.

Finally, graphics and visualization libraries are provided on MPP and server systems for post-processing and transformation of numerical output into static and dynamic graphics of many types.

4.3.3 Scientific Application Codes and Systems

NERSC supports a wide variety of third-party software applications from the independent software vendor community: for example, computational chemistry programs such as Gaussian, NWChem, Amber, and Molpro; mathematical applications such as Mathematica, Maple, and Matlab; and visualization applications such as AVS and Ensight. NERSC provides ongoing licensing and support of commercial applications, and staff liaisons to their vendors.

4.4 Human Resources

NERSC provides effective, high-quality client support to assist users in accomplishing their scientific work. This support includes provision of servers, tools, information systems, consulting services, and training. NERSC staffers continuously develop their skills in new areas, such as Gridware, to keep up with emerging trends.

4.4.1 User Support Services

The NERSC service architecture aims to ensure a response to client problems within four working hours and resolution of 95% of problems within two working days. In addition to basic Help Desk services, NERSC offers some direct scientific support to the DOE scientific community in their computational and data management efforts. Such support includes technical advice in system usage; scientific advice on algorithmic, programming, and application problems; and support for the development of new applications, tools, and utilities. NERSC also provides training and instruction in the use of technical resources, and maintains on-line documentation for all NERSC systems.

NERSC also provides support and tools for user account initialization, monitoring, management, and security for staff and users and their allocations. Much of this is provided by state-of-the-art tools and databases with continuously advancing capabilities.

Acknowledgements

The contributors and editors gratefully acknowledge the cooperation and support of DOE SC personnel and the NERSC staff in the preparation of this report. In particular, we thank F. Verdier and J. Hules for their assistance.

NERSC Users Group Executive Committee

David Dean, Chair
Stephane Ethier, Vice Chair

Office of Advanced Scientific Computing Research
David Bernholdt, Oak Ridge National Laboratory
David Keyes, Old Dominion University
David Turner, Ames Laboratory

Office of Basic Energy Sciences
Feng Liu, University of Utah
G. Malcolm Stocks, Oak Ridge National Laboratory
Theresa L. Windus, Pacific Northwest National Laboratory

Office of Biological and Environmental Research
Michael Colvin, University of California, Merced, and Lawrence Livermore National Laboratory
Brian Hingerty, Oak Ridge National Laboratory
Vincent Wayland, National Center for Atmospheric Research

Office of Fusion Energy Sciences
Stephane Ethier, Princeton Plasma Physics Laboratory
Alex Friedman, Lawrence Livermore and Lawrence Berkeley National Laboratories
Carl Sovinec, University of Wisconsin, Madison

Office of High Energy Physics
Olga Barranikova, Purdue University
Warren Mori, University of California, Los Angeles
Donald Sinclair, Argonne National Laboratory

Office of Nuclear Physics
David Dean, Oak Ridge National Laboratory
Patrick Decowski, Lawrence Berkeley National Laboratory
Douglas L. Olson, Lawrence Berkeley National Laboratory

Members at Large
Yuen-Dat Chan, Lawrence Berkeley National Laboratory
Ricky A. Kendall, Ames Laboratory, Iowa State University
Donald A. Spong, Oak Ridge National Laboratory
Douglas Swesty, State University of New York at Stony Brook

Acronyms

ACPI	Advanced Climate Prediction Initiative
ADF	Advanced Data Format (a hierarchical database system)
AFMC	auxiliary field Monte Carlo
Al	aluminum
Alcator C-Mod	a tokamak experiment
AMD	Advanced Micro Devices
AMO	atomic, molecular, and optical
AMR	adaptive mesh refinement
ANIR	Advanced Networking Infrastructure and Research
ASCI	Accelerated Strategic Computing Initiative
ASCR	Office of Advanced Scientific Computing Research
AST	Accelerator Science and Technology
Au	gold
AVS	Advanced Visual Systems
BES	Office of Basic Energy Sciences
BEST	Beam Equilibrium Stability and Transport code
BNL	Brookhaven National Laboratory
BOOMERanG	Balloon Observations of Millimetric Extra-galactic Radiation Anisotropy and Geophysics
C	carbon
CAS3D	a linear variational stability code
CCA	Common Component Architecture
CCP	Climate Change and Prediction
CCSM3	Community Climate System Model, Version 3
CKM	Cabibbo-Kobayashi-Maskawa
CMB	cosmic microwave background
COBE	Cosmic Background Explorer satellite
COBRA	a linear variational stability code
CORBA	Common Object Request Broker Architecture
CPU	central processing unit
CS	computer science
Cu	copper
CVS	Concurrent Versions System
DELTA5D	a particle simulation code
DFT	density functional theory
DIII-D	a tokamak experiment at General Atomics
DiVA	Distributed Visualization Architecture
DKES	a Eulerian localized transport code
DNA	deoxyribonucleic acid
DNS	direct numerical simulation
DOE	U.S. Department of Energy
DT	deuterium-tritium
DUSEL	Deep Underground Science and Engineering Laboratory
EM	electromagnetic

eV	electron volt
FAU	Florida Atlantic University
FES	Office of Fusion Energy Sciences
FPGA	field-programmable gate array
GASNet	Global-Address Space Networking
GB	gigabyte
Gbps	gigabits per second
Ge	germanium
GeV	giga-electron-volt
Gflops	gigaflops
GFMC	Green's function Monte Carlo
GGA	generalized gradient approximation
GPCR	G protein coupled receptor
GPFS	Global (*or* General) Parallel File System
GPU	graphics processing unit
GtC	gigatonnes of carbon
GTC	Gyrokinetic Toroidal Code
GTL	Genomes to Life
GUPFS	Global Unified Parallel File System
GYRO	a gyrokinetic fusion code
H	hydrogen
HBT	heterojunction bipolar transistor
HCCI	homogeneous charge compression ignition
HCX	High-Current Experiment
HDF5	a library and file format for storing scientific data
HECRTF	High End Computing Revitalization Task Force
HEDP	high energy density physics
HEP	Office of High Energy Physics
HIF	heavy ion fusion
HPC	high performance computing
HPCF	High-Performance Computing Facility
HPSS	High Performance Storage System
HSX	Helically Symmetric eXperiment
HV	high-velocity
ICF	inertial confinement fusion
IMPACT	a PIC code for studying high-intensity beam dynamics
INCCA	Integrated Climate and Carbon model
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
I/O	input/output
IP	Internet Protocol
IPCC AR4	Intergovernmental Panel on Climate Change Fourth Assessment Report
IPM	Integrated Performance Monitoring
IR	infrared
ISIC	Integrated Software Infrastructure Center

ITER	an international fusion experiment
IU	Indiana University
JET	Joint European Torus (a tokamak experiment)
K	kelvin
kcal	kilocalorie
keV	kilo-electron-volt
LBNL	Lawrence Berkeley National Laboratory
LDAP	Lightweight Directory Access Protocol
LLNL	Lawrence Livermore National Laboratory
LN	Logistical Networking
LSP	a large-scale electromagnetic plasma simulation PIC code
M3D	a resistive MHD and hybrid fluid-particle code
MADmap	a cosmic microwave background map-making code
Mbps	megabits per second
MDSplus	software for scientific data acquisition, storage, and management
MeV	mega-electron-volt
Mg	magnesium
MHD	magnetohydrodymamic
MICS	Mathematical, Information, and Computational Sciences Division
MIMD	multiple instruction stream, multiple data stream
MM	molecular mechanics
mol	mole
MPI	Message Passing Interface
MPICH	a portable implementation of MPI
MPP	massively parallel processing (*or* platform)
MRMP2	Multi-reference Møller-Plesset perturbation theory
M-VIA	Modular VIA
MVICH	an MPICH-based implementation of MPI for VIA
Na	sodium
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NCSA	National Center for Supercomputing Applications
NCSM	no-core-shell model
NCSU	North Carolina State University
NCSX	National Compact Stellarator Experiment
NDCX	Neutralized Drift Compression Experiment
NERSC	National Energy Research Scientific Computing Center
NetCDF	Network Common Data Form (an interface and library for data access)
Ni	nickel
NIF	National Ignition Facility
NIM	NERSC Information Management
NP	Office of Nuclear Physics
NSA	near-surface alloy
NSF	National Science Foundation
NSTX	National Spherical Torus Experiment
NUG	NERSC Users Group

NUGEX	NERSC Users Group Executive Committee
O	oxygen
OBER	Office of Biological and Environmental Research
ONIOM	Our own N-layered Integrated molecular Orbital and molecular Mechanics
ONIOM-EE	ONIOM electronic-embedding
ORBIT	a particle simulation code
ORNL	Oak Ridge National Laboratory
OSIRIS	a PIC code for modeling intense beam-plasma interactions
PB	petabyte
PCM	Parallel Climate Model
PCMDI	Program for Climate Model Diagnosis and Intercomparison
PDSF	Parallel Distributed Systems Facility
PEP-II	a dual-ring particle accelerator
PIC	particle-in-cell
PIES	a code that calculates three-dimensional MHD equilibria
pNetCDF	Parallel NetCDF
POE+	a performance data collection tool
ppmv	parts per million by volume
PPPL	Princeton Plasma Physics Laboratory
QCD	quantum chromodynamics
QCDOC	QCD On a Chip (a computer designed for QCD calculations)
QD	quantum dots
QM	quantum mechanics
QMC	quantum Monte Carlo
QPS	Quasi-Poloidal Stellarator
RCF	RHIC Computing Facility
Re	rhenium
RF	radio frequency
RHIC	Relativistic Heavy Ion Collider
RIA	Rare Isotope Accelerator
Ru	ruthenium
SAPP	Scientific Application Pilot Program
SC	Office of Science
SC	an annual supercomputing and networking conference
SciDAC	Scientific Discovery through Advanced Computing
SCIRun2	a parallel component framework
SDSC	San Diego Supercomputer Center
SGI	Silicon Graphics, Inc.
Si	silicon
SLAC	Stanford Linear Accelerator Center
SPMD	single-program, multiple-data
SRES	Special Report on Emissions Scenarios
SRM	Storage Resource Manager
SRU	storage resource unit
STAR	Solenoidal Tracker at RHIC

STELLGAP	a stellarator simulation code
STELLOPT	a stellarator simulation code
STI	Sony, Toshiba, and IBM
STK	Storage Technology Corp.
SUNYSB	State University of New York, Stony Brook
TB	terabyte
TERPSICHORE	a linear variational stability code
Tflops	teraflops
TFTR	Tokamak Fusion Test Reactor
TJNAF	Thomas Jefferson National Accelerator Facility
TM	transition metal
UCSD	University of California, San Diego
UPC	Unified Parallel C (programming language)
VIA	Virtual Interface Architecture
VMC	variational Monte Carlo
VMEC	a code that solves three-dimensional equilibrium equations
VORPAL	a relativistic, hybrid plasma and beam simulation code
WAN	wide-area network
WARP3D	a code for the static and dynamic nonlinear analysis of solids
WARP-SLV	a semi-Lagrangian Vlasov solver
WDM	warm dense matter
WMAP	Wilkinson Microwave Anisotropy Probe
Zn	zinc
Zr	zirconium