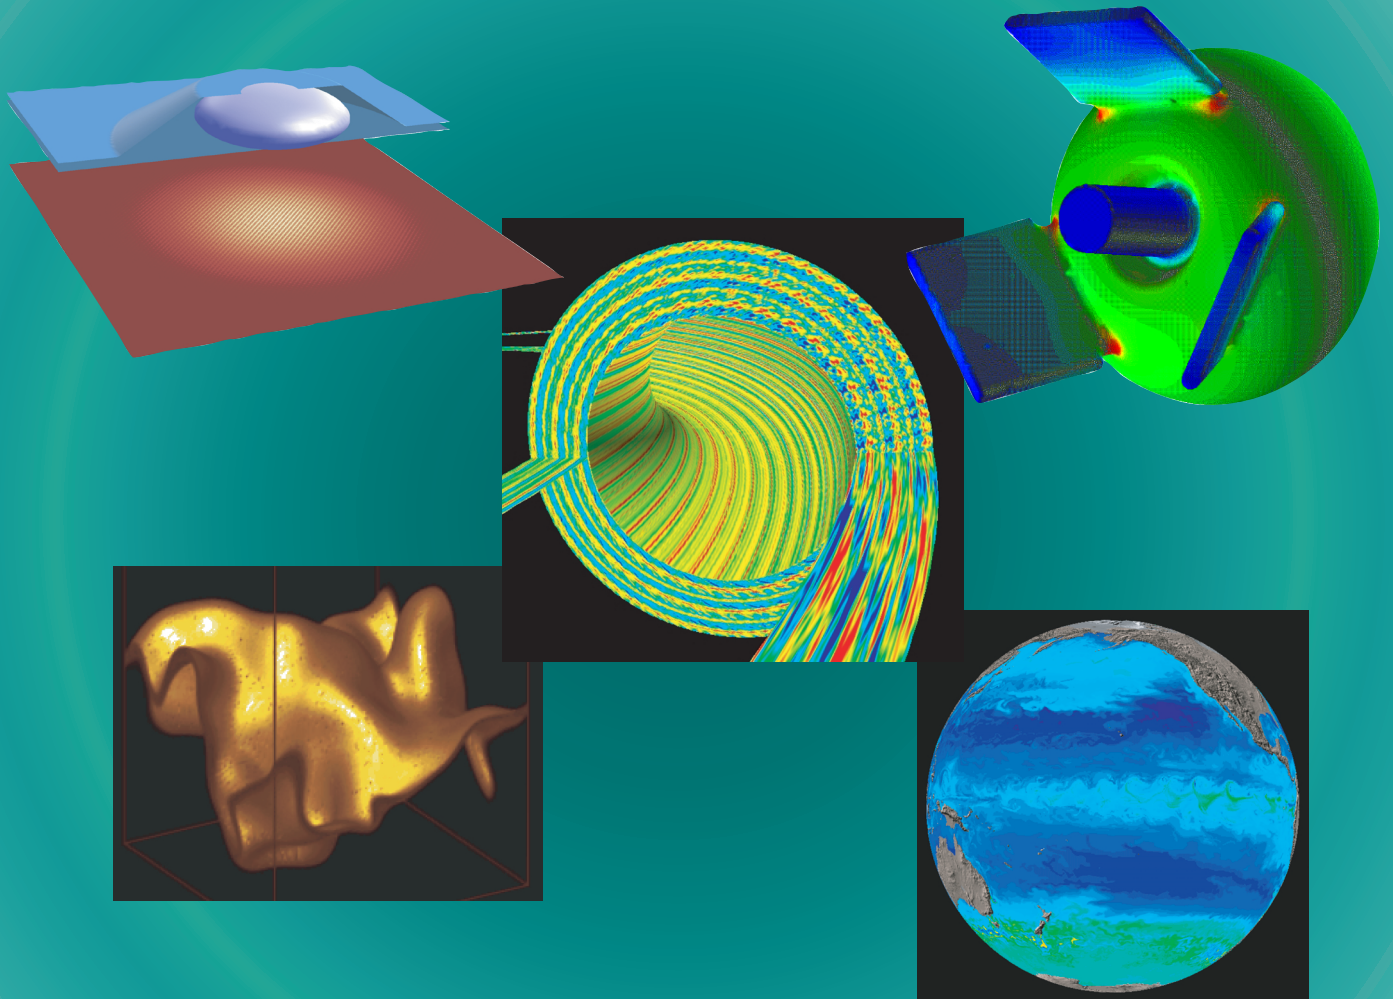# DOE Greenbook – Needs and Directions in High-Performance Computing for the Office of Science

A Report from the NERSC User Group
April 2002

**Prepared for the United States
Department of Energy
Office of Science**

Lawrence Livermore National Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

This report has been reproduced directly from the best available copy.

Available electronically at http://www.doe.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
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/ordering.htm

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
http://www.llnl.gov/tid/Library.html

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

A Report from the NERSC User Group

April 2002

Compiled by
Douglas A. Rotman
Paul Harding

Prepared for the United States Department of Energy Office of Science

Lawrence Livermore National Laboratory

Contents

Executive Summary

Computational Challenges

The purpose of this report is to document the computational science being done at the Office of Science High-Performance Computing Centers especially the National Energy Research Scientific Computing Center (NERSC) of the U.S. Department of Energy (DOE). NERSC and the other centers provide scientists from the Office of Science with a very-high-level shared computing facility, including massively parallel computers, a cluster of vector multiprocessors, an experimental cluster of workstations, a mass storage system, visualization facilities, a user support organization, networking, and software. This report includes discussions of the ways in which NERSC should evolve in order to meet the needs of DOE's scientists to continue to do forefront science.

Simulation has gained equal footing to experiments and theory in the triad of scientific process. A large shared computing center continues to be essential to the scientific mission of the Office of Science as it provides the researchers with resources that could not be supplied locally. We provide examples of computational challenges and opportunities that will guide activities at NERSC (and the other Office of Science computing centers) over the next few years and provide insights into the computational needs of these activities. Brief descriptions of current and/or future scientific efforts are presented from each of the five major offices within the Office of Energy Research: High Energy and Nuclear Physics, Basic Energy Sciences, Fusion Energy Science, Biological and Environmental Research, and Advanced Scientific Computing Research.

During the 1970s and 1980s much of the work termed as super-computing was done using vector supercomputers (i.e., CDC and the Cray series of machines, Cray 1, XMP, YMP, C90, and T90). During the 1990s, massively parallel computers became equal partners to these vector computers, and by the end of the decade few very large vector machines remained operational within the United States. Massively parallel computing is now approaching a level of maturity that allows for long-running and very large production simulations. Moreover, a majority of scientific models have now been transitioned to parallel computing and many teams

are now developing and running applications primarily on large parallel systems. While there remain some performance issues and the need for additional software tool development, we believe parallel computing is now the primary production computing paradigm. Software and algorithms will continue to play an integral role in better utilization of parallel systems. We discuss algorithm needs for upcoming science research and the use of distributed computational systems across the DOE (and non-DOE) complex for enhanced computational resources. We discuss platform configuration needs to align with computational requirements. Mass storage, visualization facilities, and networks, which are essential complements to raw CPU power, are briefly described with a discussion of the future needs in these areas.

Recommendations

- On the basis of the computing activities within the Office of Science presented in the section entitled Computational Resources for the Office of Science, we make these specific recommendations:

- Expand the high-performance computing resources available at NERSC. The upcoming procurement must ensure an appropriate balance of computational equipment to support the wide range of large-scale applications involving production computing and development activities.

- Encourage the continued improvement of algorithms, software (including middleware), and database technology for improved performance on parallel platforms. More pro-active support for implementation of NERSC (and other DOE) software tools into Office of Science application codes may be necessary.

- Strengthen the networks. Reliable, robust and fast access to the computational resources is essential. This is especially true for increased use of Grid Computing.

- Coordinate NERSC activities with SciDAC, ISIC, and other Office of Science compute centers.

- Identify mechanisms to provide larger-scale production facilities. While model development and physics/chemistry/biological improvements will always be moving forward, many Office of Science applications are now ready for large-scale production computing on NERSC parallel platforms. Future simulations will

involve long-running simulations requiring hundreds of hours of continuous simulation time. Current computer system formula tions do not lend themselves to these types of simulation; this needs to be addressed.

Appendix A provides important information on computer architecture and hardware issues that are crucial to the continued success of NERSC. Issues addressed are memory performance/bandwidth, disk and I/O systems, data storage and transfer, and visualization.

Introduction

The NERSC Users Group (NUG) encompasses all investigators utilizing the NERSC computational and storage resources of the Department of Energy Office of Science facility. At the February 2001 meeting held at the National Energy Research Scientific Computing (NERSC) facility, the NUG executive committee (NUGEX) began the process to assess the role of computational science and determine the computational needs in future Office of Science (OS) programs. The continuing rapid development of the computational science fields and computer technology (both hardware and software) suggest frequent periodic review of user requirements and the role that computational science should play in meeting OS program commitments.

Over the last decade, NERSC (and many other supercomputer centers) have transitioned from a center based on vector supercomputers to one almost entirely dedicated to massively parallel platforms (MPPs). Users have had to learn and transform their application codes to make use of these parallel computers. NERSC computer time requests suggest that a vast majority of NERSC users have accomplished this transition and are ready for production parallel computing. Tools for debugging, mathematical toolsets, and robust communication software have enabled this transition. The large memory and CPU power of these parallel machines are allowing simulations at resolutions, timescales, and levels of realism in physics that were never before possible.

Difficulties and performance issues in using MPP systems remain linked to the access of non-uniform memory: cache, local, and remote memory. This issue includes both the speed of access and the methods of access to the memory architecture. Optimized mathematical tools to perform standard functions on parallel machines are available. Users should be encouraged to make heavy use of those tools to enhance productivity and system performance.

There are at least four underlying components to the computational resources used by OS researchers.

- High-Performance Computing Technology

- Advanced Software Technology and Algorithms

- Energy Sciences Network

- Basic Research and Human Resources

In addition to the availability from the vendor community, these components determine the implementation and direction of the development of the supercomputing resources for the OS community.

In this document we will identify scientific and computational needs from across the five Office of Science organizations: High Energy and Nuclear Physics, Basic Energy Sciences, Fusion Energy Science, Biological and Environmental Research, and Advanced Scientific Computing Research. We will also delineate the current suite of NERSC computational and human resources. Finally, we will provide a set of recommendations that will guide the utilization of current and future computational resources at the DOE NERSC.

The NERSC User Community

The NERSC User Community (NUC) represents a complex and diverse set of scientists that have computational science requirements that span the full gamut of algorithms, techniques, and applicability to specific hardware systems. NUC applications range from nearly fully vectorized suites to highly optimized vector-parallel codes to codes that are candidates for the Gordon-Bell award for parallel performance on MPP systems. NUC users range from the highly skilled computational scientists who can adapt their codes to get the highest performance available on a given computational resource to users who simply want to compile and run. This provides a daunting challenge to the NERSC staff: supplying hardware, software, and infrastructure that can meet the user community's needs. In this document, we hope to capture the salient components of the user community's requirements for computational resources, both currently and in the foreseeable future. The fact that representative NUC applications have and will be used in the evaluation of new NERSC procurements is a measure of NERSC's commitment to its user community.

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 simulations, 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. 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 in the form of massively parallel architectures.

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 start to blur. For example, as material surfaces become more accurate and defects are well defined, chemical reactions on those surfaces become not only interesting topics of research, but also computationally feasible topics of research. In each of the following sections, general information concerning computations in each subprogram is given, some specific examples are given, and the computational challenges are explored.

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 material science today are performing computations that cross length and timescales, including atomistic to nanoscale to macroscale computations, and creating materials by

design, including the construction of multi-component molecular devices and machines having desired properties. Both of these immense challenge areas will require much theoretical and algorithmic development; however, they will also require larger computational resources both to examine the physical systems and to describe the reactions or properties of those systems.

Some of the applications that are beyond the current computational capability, but could be addressed with a two-to-three-order increase in computational power are:

- Patterned arrays of quantum dots, with system sizes greater than 20 million atoms, using a combination of plane-wave empirical pseudopotential (EPM/ESCAN) and the Linear Combination of Bulk Bands technique (Figure 1). The current practical system size is around 0.5—3 million atoms. One could investigate quantum dot devices, including optical and transport phenomena, with such calculations.
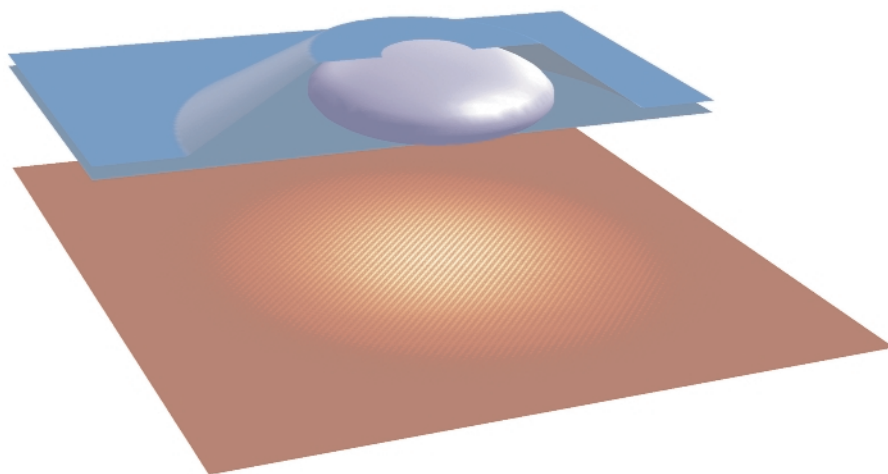
- Microstructured alloys, fully including the effects of ordering and clustering of the constituents and impurities. Currently these phenomena must be investigated separately. A concurrent investigation of these structural phenomena would allow for greater accuracy in property prediction and give detailed insight for materials design and optimization.

- Large-scale structural relaxation of defects in crystals that allow for the characterization of, e.g., shallow impurity states originated from a point defect in semiconductor crystals very accurately.

- Determination of large-scale magnetic properties which would allow one to examine the local and nonlocal phenomena, including interactions of distinct magnetic domains.

- Simulation of much more realistic systems with photonic band gap codes. This would include three-dimensional (3D) systems with disorder.

- Full-scale molecular modeling design of polymeric systems and the first fully atomistic simulations of substantial polymer melts, to predict their rheology and blend immiscibility (Figure 2).

One nanodot consists of thousands of copper oxide molecules.

Figure 2.
Molecular modeling.

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 Department'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.

As with the Materials Sciences program, Chemical Sciences faces several challenges that will require orders-of-magnitude increase in computational capabilities and considerable effort in algorithm development. Some of these 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), 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 (Figure 3).

Figure 3.
Structure of $UO_2(CO_3)_3^-$ (uranium—light blue; carbon—gray; oxygen—red). Image courtesy of Wibe deJong, Pacific Northwest National Laboratory.

9

Some of the applications that are beyond the current computational capability, but could be addressed with a two-to-three-order increase in computational power are:

- Very accurate molecular calculations on large systems:

  - Calculations that start to cross into the nanoscale regime. This includes the examination of nanoscale wires, tubes and boxes. The information gained from these calculations is not only valuable for the properties and structures that are determined, but they also produce useful information to be used by less accurate methods such as molecular mechanics and structural simulations.
  - Highly accurate energetics (both thermodynamics and kinetics) for the examination of chemical reactions of large, realistic systems. This may also involve the inclusion of solvent effects since most reactions occur in solution. The results of these calculations could have a large impact in several key areas such as combustion (see below) and climate modeling.
  - Calculations involving relativistic effects. These include light element systems where extreme accuracy is needed or heavy element systems where the problem cannot be addressed qualitatively without the inclusion of relativity. Not only are stru-c tures and ground state energetics important, but also properties such as infrared frequencies, nuclear magnetic resonance shifts, and optical activity are required for the computations to relate to the experimental results.

- Much more detailed and accurate direct numerical simulations (DNS) of flames:

  - Larger physical sizes could be examined, such as a few decimeters in two-dimensional (2D) simulations and a few centimeters in 3D simulations.
  - The inclusion of more of the chemical mechanisms of the flame. Current DNS models use very simplified chemical models because it is simply too expensive to use realistic models.
  - More geometrical complexity to model realistic flame conditions.
  - New capabilities, such as advanced modules to describe radiative heat transfer as well as condensed-phase physics (soot particles, liquid fuel droplets), could be used. These physics models are required to enable detailed comparisons with fundamental experimental studies (laminar and turbulent).

10

Geosciences

The Geosciences Research 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 provide 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 minrals with and without absorbates. Again, 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.

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 (Figure 4).

Figure 4.
Electrostatic potential of the LPS membrane of *Pseudomonas aeruginosa*. Image courtesy of T.P. Straatsma, Pacific Northwest National Laboratory.
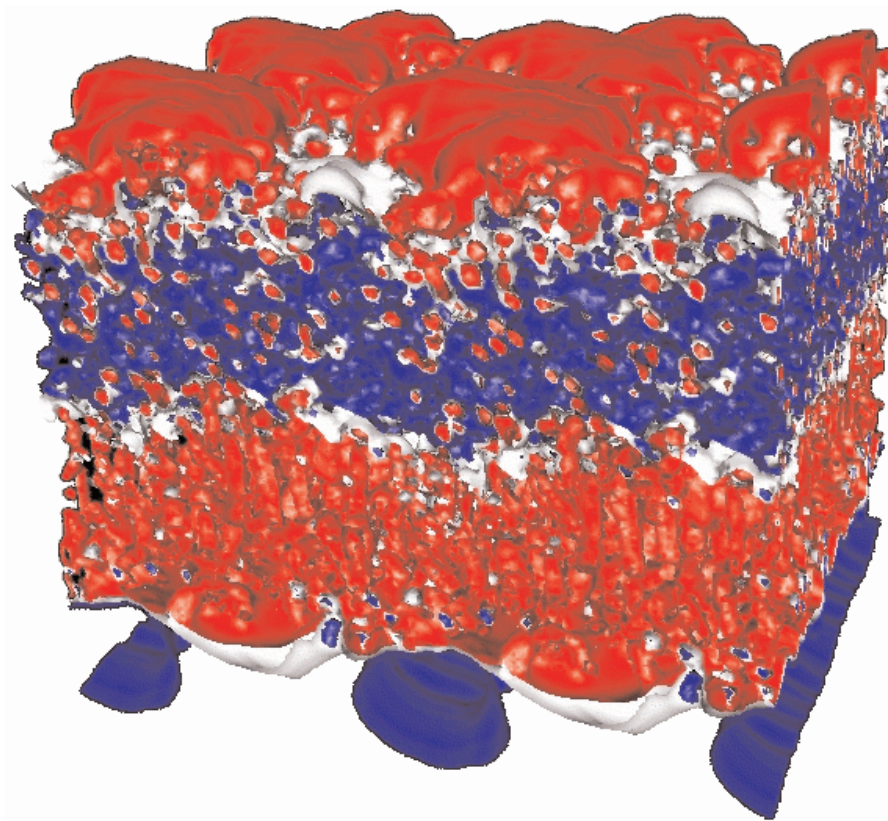
Some of the applications that are beyond the current computational capability, but could be addressed with a two-to-three-order increase in computational power are:

• Simulations of several hundred million atoms moving a million time steps (nanosecond) could be accomplished in a more routine fashion. Currently, these calculations are either beyond the current capability or are accomplished in a grand challenge fashion (large allocations of computer time to a select number of researchers). Simulations of this size allow for:

  – Larger membrane sections to simulate more realistic systems.
  – Examination of phenomena that only occur on long timescales such as ion migration, catalytic events, and photosynthetic events.

- Simulations with more detailed interactions can be performed. These interactions allow the accuracy of simulations to be improved and new phenomena to be examined.

Predominately, scalability at the hardware level of algorithms used for BES research is limited by the amount of memory per processor and the latency to that memory, communication bandwidth and latency, and CPU speed (the particular limiting factor is dependent on the algorithm used). At the software level, many of the algorithms depend on parallel eigensolvers (e.g., Kohn-Sham equations for density functional calculations, Hartree-Fock computations), matrix inversion, and Fast Fourier Transforms (FFTs e.g., classical or tight binding molecular dynamics and Car-Parrinello methods), all of which are currently limited in their scalability. To take advantage of the enlarged computational resources, new algorithms need to be developed to mitigate the current scalability problems.

Some of the other challenges to the use of the massively parallel platforms are the analysis of tens of gigabytes of data and the transfer of this data to macroscale calculations such as those of structural mechanics. This transfer of appropriate data and information across computational scales is a very active area of research and will require fast communication networks between DOE laboratories and collaborating institutions.

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 the Human Genome, Structural Biology, and Microbial Genome research. In close cooperation with improved technologies, NERSC computational capabilities are focused on mapping, sequencing, and information management, toward characterizing the molecular nature of the human genome and the determination of 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 an 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.

Genomes to Life

Having the complete DNA sequences of genomes for organisms ranging from humans to mice to microbes now brings us to perhaps the greatest scientific frontier in history. The aspiration of the biology community for the 21st century is to build from the foundation of whole-genome sequences a new, comprehensive, and profound understanding of complex living systems.

This objective can be achieved only by joining revolutionary technologies for systems-level and computational biology. A central goal of the Genomes to Life program is to establish, within a decade, a national infrastructure to transform the tremendous outpouring of data and concepts into a new computationally based biology. The offices of Biological and Environmental Research and Advanced Scientific Computing Research have formed a strategic alliance to meet this grand challenge.

Built on the continuing successes of international genome-sequencing projects, the Genomes to Life program will take the logical next step: a quest to understand the composition and function of the biochemical networks and pathways that carry out the essential processes of living organisms. This roadmap sets forth an aggressive 10-year plan designed to exploit high-throughput genomic strategies and centered around four major goals:

- Identify and characterize the molecular machines of life the multiprotein complexes that execute cellular functions and govern cell form.

- Characterize gene regulatory networks.

- Characterize the functional repertoire of complex microbial communities in their natural environments at the molecular level.

- Develop the computational methods and capabilities to advance understanding of complex biological systems and predict their behavior.

The Machines of Life
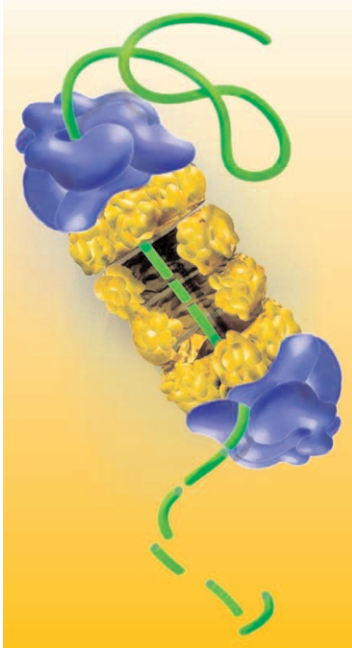
Proteins carry out almost all of life's essential processes, often working in highly complex multicomponent assemblies that sometimes include other types of macromolecules such as DNA or RNA. These machines typically work together in functional networks called pathways that underlie the dynamic life of a cell as they execute important metabolic functions, mediate information flow within and among cells, and build cellular structures. Breaking down unneeded proteins a task equal in importance to synthesizing new proteins is accomplished by the orderly action of several multiprotein complexes. At the heart of this process is a multiprotein complex called the proteasome (Figure 5).

An entirely different class of molecular machines functions as motors, converting chemical energy into mechanical motion, both linear and rotary. The dynein motor (Figure 6), a cellular complex believed to be composed of 12 distinct protein parts, performs fundamental transportation tasks critical to the cell; defects in its structure can prove fatal. This machine converts chemical energy stored in an ATP molecule into mechanical energy that moves material though the cell along slender filaments called microtubules. One of the dynein motor's most important functions occurs during cell division, when it helps move chromosomes into proper position.

The Human Genome Project taught that evolutionary improvement in existing technologies (e.g., DNA sequencing) can have a revolutionary impact on science. The systems approach taken by the Genomes to Life program dictates that existing technologies must evolve to a high-throughput capability. In addition, revolutionary technologies need to be developed, incorporating new modes of robotics and automation as well as advanced information and computing technologies. The following is a list of some key computational technologies for informatics, modeling, and simulation:

- Algorithms for genome assembly and annotation and for bioinformatics to measure protein expression and interactions.

- Standardized formats, databases, and visualization methods for complex biological data sets, including expression profiles and protein—protein interaction data.

- Molecular modeling methods for long-timescale, low-energy macromolecular interactions and for prediction of chemical reaction paths in enzyme active sites.

Figure 5.
A multiprotein complex called a proteasome breaks down unneeded proteins.

Figure 6.
The dynein motor, a cellular complex, converts chemical energy into mechanical energy in cells.

- Methods for automated collection and integration of biological data for cell-level metabolic network analysis or pathway model ing; improved methods for simulation, analysis, and visualization of complex biological pathways; and methods for prediction of emergent functional capabilities of microbial communities.

Goals for computing capabilities in the Genomes programs are:

1. **Identify and characterize the molecular machines of life the multiprotein complexes that execute cellular functions and govern cell form.** The identification and characterization of the multiprotein machines of life involve substantial computational demands, ranging from sophisticated data analysis to atomic-scale simulations of protein interactions. Meeting these needs will require the development of new algorithms and databases and the use of very-high-performance computers. Following are the initial specific computational aims to be addressed for Goal 1:

   - Improve bioinformatics methods needed to deconvolute mass spectroscopic protein-expression data, including isotopic label ing ratios, and to handle massive amounts of protein chip expression data.

   - Adapt and develop databases and analysis tools for integrating experimental data on protein complexes measured with differ-ent methods under varied conditions.

   - Develop algorithms for integration of diverse biological data bases including transcriptome and proteome measurements, as well as functional and structural annotations of protein-sequence data.

   - Develop modeling capabilities for simulating multiprotein machines and for predicting the behavior of protein machines in cell networks and pathways.

2. **Characterize gene regulatory networks.** The task of elucidating and adapting the circuitry of gene regulatory networks will be dependent on computational methods to identify and characterize regulatory sequences and computer models of the regulatory net works. The computational research tasks for Goal 2 involve developing methods to do the following:

   - Extract regulatory elements, including operon and regulon sequences, using sequence-level comparative genomics.

- Simulate regulatory networks using both nondynamical models of regulatory capabilities and dynamical models of regulatory kinetics.

- Predict the behavior of modified or redesigned gene regulatory networks.

3. **Characterize the functional repertoire of complex microbial communities in their natural environments at the molecular level.** There are many computational challenges to characterizing the composition and functional capability of microbial communities. New algorithms for DNA sequence assembly and annotation will be required to analyze the multiorganism sequence data, and new modeling methods will be required to predict the behavior of microbial communities. The computational research tasks will be to develop methods to:

- Deconvolute mixtures of genomes sampled in the environment and identify individual organisms.

- Facilitate multiple-organism shotgun-sequence assembly.

- Improve comparative approaches to microbial sequence annotation and gene finding and use them to assign functions to genes where possible.

- Accomplish pathway reconstruction from sequenced or partially sequenced genomes and evaluate the combined metabolic capabilities of heterogeneous microbial populations.

- Integrate regulatory-network, pathway, and expression data into integrated models of microbial community function.

4. **Develop the computational methods and capabilities to advance understanding of complex biological systems and predict their behavior.** A long-term goal of the computational-modeling section of Genomes to Life is to develop the next generation of methods for simulating cellular behavior and pathways. Other goals are to create molecular-modeling and bioinformatics tools for studying multi-protein complexes, along with new computational methods to explore the functional diversity of microbes. The following are the specific aims:

- Aim 1. Develop methods for high-throughput automated genome assembly and annotation.

- Aim 2. Develop computational tools to support high-throughput experimental measurements of protein-protein interactions and protein-expression profiles.

- Aim 3. Develop predictive models of microbial behavior using metabolic-network analysis and kinetic models of biochemical pathways.

- Aim 4. Develop and apply advanced molecular and structural modeling methods for biological systems.

- Aim 5. Develop the groundwork for large-scale biological computing infrastructure and applications.

Atmosphere—Ocean Dynamics and Global Change

Understanding the response of the Earth's climate system to both natural and anthropogenic (man-made) stimuli, together with understanding the natural variability of the climate system, is part of the science base required for rational formulation of energy policy. One of the most important public policy environmental issues concerns the effects of increasing greenhouse gases and aerosols on the climate system. The use of numerical models in climate prediction has already been shown very useful. However, climate models need to be substantially improved. This is being done by improving our knowledge base through field and satellite observations and by model process and algorithm improvements that complement advances in computation power.

A comprehensive climate system model must describe the behavior and interactions of Earth's atmospheric, oceanic, cryospheric, geochemical, and biospheric systems. At the core of such comprehensive earth simulation systems are atmosphere—ocean general circulation models (AOGCMs) that simulate the behavior and motions of the global atmosphere, oceans, and sea ice. These models are computationally demanding because adequate realism requires AOGCMs to discretize the planet with large numbers of grid cells and march forward in time in relatively small steps.

The climate modeling effort within the DOE community is currently focused on the development and use of the Community Climate System Model (CCSM), which is jointly supported by the National Science Foundation (NSF), DOE, and the National Aeronautics and Space Administration (NASA). Predecessor versions of the CCSM's

atmospheric model (the CCM model series) were developed at the National Center for Atmospheric Research (NCAR) for use by the academic community and executed mostly on vector supercomputers. NCAR and the associated NSF science community play a leading role in guiding the climate science aspects of the CCSM. DOE labs play a strong role in component model and algorithm development, and all aspects of terascale computation.

A DOE-supported complementary activity to the CCSM has been the Parallel Climate Model (PCM). The PCM uses some of the same components as the CCSM, but is targeted toward running applications on DOE's larger distributed memory computers. The PCM and CCSM model systems are currently being upgraded and improved for high-performance computers. The two models will be merged into a single CCSM modeling system over the next year or so. The current version of the PCM is the PCM-CSM-Transitional Model (PCTM). As the name implies, the PCTM represents a transition to the merged model.

A typical production run of an AOGCM of the CCSM/PCM class uses an atmospheric model with a horizontal spatial resolution of 200—300 km and 18—30 vertical levels from the surface through the middle stratosphere. Typical resolutions of the ocean model component range from less than 100 km in the 2/3 degree model used in PCTM to several hundred kilometers in the coarser-resolution ocean components of the NCAR CSM-1 (the precursor to the CCSM). Production runs of 100 to 300 simulated years are typical. A standard prediction run might be defined as one starting at a pre-industrial year of around 1870 and running forward until the end of the 21st century.

The computational requirements for running climate models are increasing. Not only are the models themselves becoming more demanding as resolution and modeled processes are improved, but the number and length of runs is also increasing. The proliferation of Intergovernmental Panel on Climate Change (IPCC) greenhouse forcing scenarios combined with the need to understand the various uncertain climate forcings of the past few centuries has resulted in a large number of potential production runs. This is further compounded by the desirability of doing ensemble runs in order to quantify uncertainty and increase the statistical significance of results.

Looking toward the future of climate modeling, we can anticipate a fundamental change in how climate prediction simulations will be

performed. Knowledge of the future concentration of atmospheric carbon dioxide (CO_2) is required to predict future climate change. However, at present, atmospheric CO_2 concentrations are *prescribed*, not simulated, in almost all comprehensive climate models. The CO_2 concentrations used in the simulations are taken from hypothetical scenarios, simple extrapolations, or simplified offline models. Credible simulations of the entire climate system, however, need to *predict* time-evolving atmospheric greenhouse forcing using anthropogenic emissions as the fundamental input.

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.

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 19th 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.

As one example, consider that the long-term uptake of carbon dioxide by land plants can be perturbed by changes in climate.

Recent simulations have shown that interannual climate variations during the past few decades probably resulted in large changes in net carbon uptake by the land biosphere (Figure 7). Even the *sign* of the uptake can vary from year to year. In a globally warmed future, the response of the land biosphere is uncertain, but it has the potential to play a large role in determining how much CO_2 remains in the atmosphere. Added to this is the uncertain direct effect of extra CO_2 on plant growth, the so-called fertilization effect.

Figure 7.
Correlation between inter-annual variability of net carbon uptake and annual mean temperature and precip-itation. Results are from an ensemble of 16 runs of 13 years, each using the CCM3 climate model coupled to the IBIS2 dynamic terres-trial biosphere model.
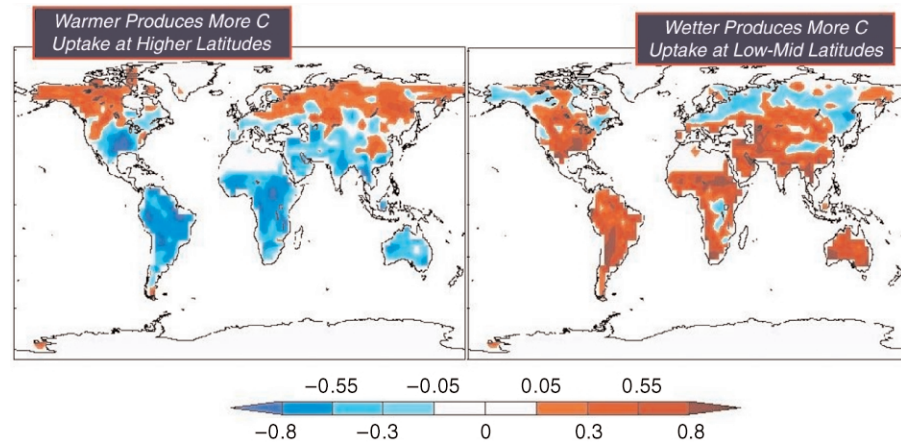
Work on including the terrestrial and oceanic components of the carbon cycle in climate models is taking place within the NSF—DOE modeling community and beyond. The next few years should see a rapid growth of modeling capability, resulting in simulations using comprehensive coupled climate and carbon cycle models. Computational requirements of such models are uncertain, but they may not require a substantial increase over the demands of the core AOGCM if the first generation of land and ocean biogeochemical models are somewhat restrained in the number of included processes and prognostic variables. One concern, however, will be the need to spin-up or equilibrate the biogeochemical components of climate—carbon models. The terrestrial biosphere (e.g., trees) can takes hundreds of years to equilibrate. Similar equilibration for oceanic carbon fields can require thousands of simulated years.

Another source of increased computational demand is the continued drive to run climate simulations at higher spatial resolution. Current operational resolutions are adequate to resolve most of the funda-mental features of the global atmospheric circulation, but are inade-quate for detailed analysis of climate change impacts. Further statis-tical or physically based downscaling of global model results is generally required to increase the value of climate model results for decisionmaking. Using the results of coarse-grid global models as

input, it is possible to downscale by running limited-domain high-resolution models, or high-resolution global models in a two-stage process (Figure 8). High-resolution runs of global models currently strain the largest computers but will become the standard over the coming decade as computer power continues to increase.
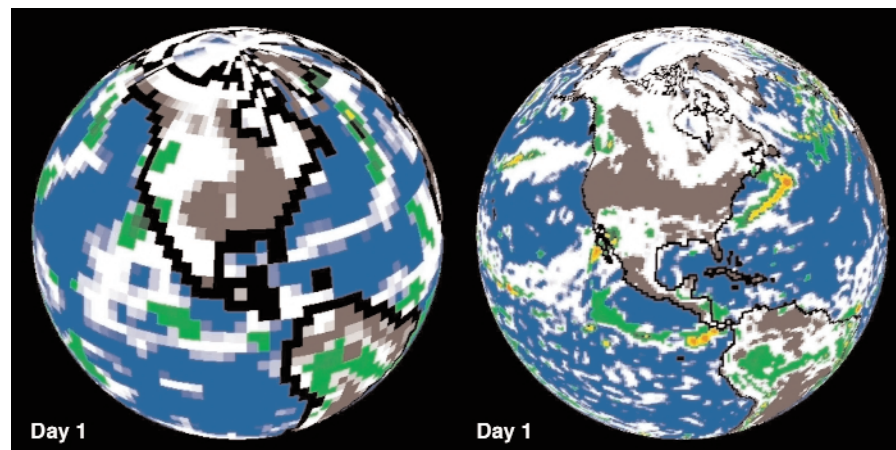
Figure 8.
A comparison of "snapshots" of weather simulated by climate models having different horizontal resolutions. Left: a standard resolution of about 280 km. Right: an experimental high resolution of about 80 km.

Atmospheric Chemistry and Aerosols

The goal of atmospheric chemistry research at DOE is to understand the impact of energy use emissions on the distribution of chemical species in the atmosphere, especially on the distribution of ozone. The use of fossil fuel for energy production creates combustion products of carbon dioxide, carbon monoxide, sulfur dioxide, soot, oxides of nitrogen, and many other byproducts; most of these species are involved with important catalytic cycles in the production or loss of ozone. Ozone is an important radiative species in both the stratosphere and troposphere, especially in the tropopause region the region separating the troposphere and stratosphere. Ozone in the stratosphere is beneficial, as it absorbs a significant fraction of the sun's shorter-wavelength radiation. Ozone in the troposphere is a pollutant (a respiratory irritant in humans and damages crops, vegetation, and many materials) that absorbs both incoming solar radiation and outgoing longwave radiation. Ozone is also an important part of the oxidizing capacity of the atmosphere through its photolysis that leads to the hydroxyl radical (OH). Since reaction with OH is a major sink of many atmospheric species, its concentration can control the distributions of many radiatively important species.

Fossil fuel combustion and biomass burning have also led to an increase in aerosol concentrations. Aerosols may tend to cool the Earth by directly reflecting incoming solar radiation. Additionally,

aerosols cause an indirect effect, by acting as cloud condensation nuclei, leading to clouds with more (but smaller and more reflective) droplets, thereby increasing albedo and further cooling the atmosphere. There is a tradeoff: greenhouse gases tend to warm the atmosphere while aerosols tend to cool it. While CO_2 is a dominant greenhouse gas, combined, the non-CO_2 greenhouse gases are also very important to the radiative balance of the atmosphere, and atmospheric concentrations of many of these non-CO_2 gases (e.g., methane) are growing very quickly.

In addition to observations and laboratory studies, atmospheric chemistry models have always played an important role in understanding the chemical state of the atmosphere. Chemistry models include many prognostic variables on large grids, and within each grid point, a large set of coupled stiff ordinary differential equations are solved. Solving for a time-evolving distribution of atmospheric species is computationally intensive, requiring large computing platforms. In the future, our computing needs will continue to grow to include (1) a more comprehensive numerical description of atmospheric chemistry and physics, (2) enhanced interaction with land and ocean biogeochemical models, and (3) interactive links between chemistry and climate.

Comprehensive Atmospheric Chemistry Models

Observations and laboratory studies are driving models to include more chemical species and involved physical and chemical processes, as well as higher-resolution studies of larger regions of the atmosphere. Studies continue to show the importance of detailed surface and aqueous chemistry on the surfaces of aerosols and water droplets to ozone distributions. Including these processes presents a challenge to scientists from not only a computational perspective but also from the required algorithms to accurately account for these reactions. Observations and numerical simulations have also shown the importance of linking the chemistry and physics of the troposphere with that of the stratosphere. Moreover, the future may also include the need to include mesospheric chemistry and physics. Ozone in the troposphere arises from both *in-situ* photochemical production and transport from the stratosphere. Within the troposphere, ozone is formed in-situ when carbon monoxide (CO), methane (CH_4), and non-methane hydrocarbons (NMHCs) react in the presence of nitrogen oxides ($NO_x = NO + NO_2$) and sunlight. The photochemistry of the stratosphere differs significantly from that in the troposphere. Within the stratosphere, ozone formation is

initiated by the photolysis of O_2. Stratospheric ozone may be destroyed via catalytic reactions with NO, H, OH, Cl, and Br.

In-situ production and loss processes in the stratosphere and troposphere are crucial to understanding ozone distributions. Additionally, however, other processes, such as transport of stratospheric ozone into the troposphere, are important in understanding ozone distributions. Stratospheric transport arises from tropopause folding events, the decay of cutoff weather systems, and stratospheric streamers. These transport events vary regionally and seasonally. Figure 9 shows an isosurface of 100 ppbv ozone from a chemistry simulation involving both the troposphere and stratosphere the value of 100 ppbv being representative of lower-stratospheric ozone. This figure presents evidence of the transport of lower-stratospheric ozone far into the troposphere and suggests that understanding tropospheric ozone concentrations and trends requires detailed knowledge of stratospheric ozone and its trends.
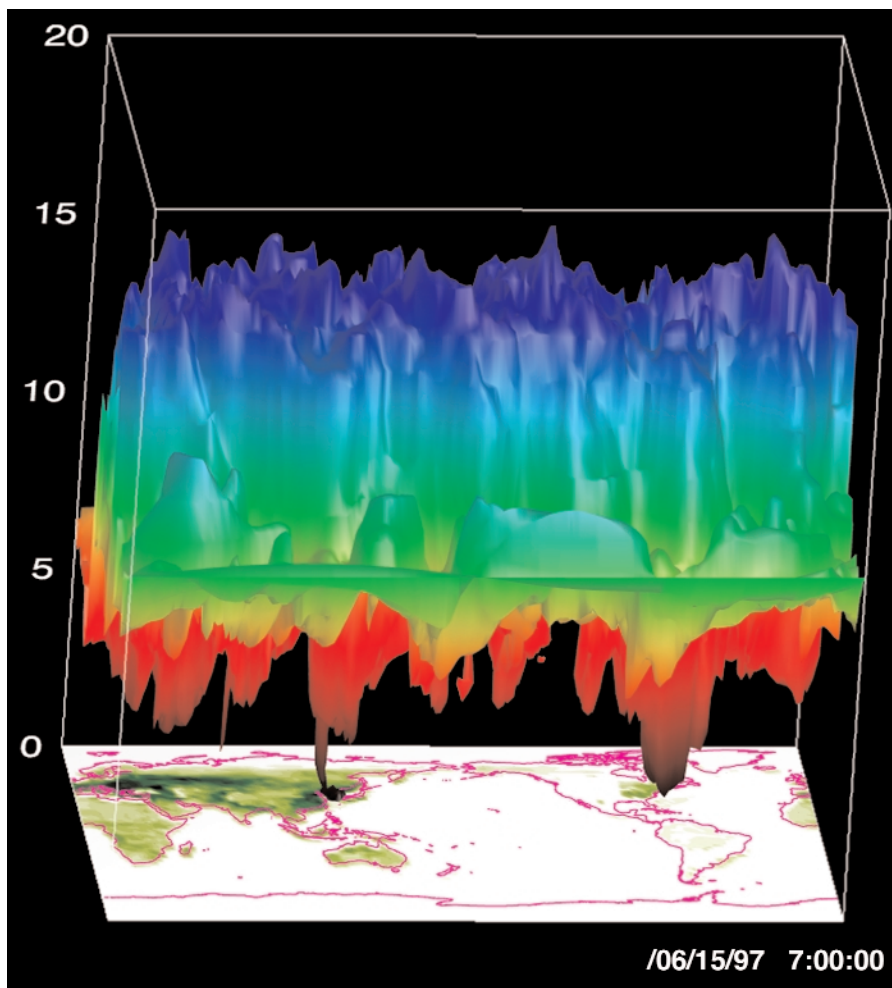
Figure 9.
A 100-ppb isosurface of ozone showing transport of Stratospheric ozone into the troposphere.

Coupling Atmospheric Chemistry with Land and Ocean Biogeochemistry

The distribution of atmospheric chemical species is integrally tied to emissions from the land and ocean. These surface emissions are controlled by biogeochemical cycles which themselves are determined by concentrations of chemical species, sunlight levels, vegetation type, and other meteorological variables. Surface-emitted species that play important roles in atmospheric catalytic cycles include acetone, isoprene, monoterpenes, dimethyl sulfide, and many other volatile organic compounds. Especially important to tropospheric ozone (and regional air quality) is the emission of isoprene because of its large emission flux from vegetation. Many chemistry models currently emit these biogenic species from static databases of assumed monthly averaged quantities. In the coming years we will directly link biogeo-chemical models to atmospheric chemistry models. Only through this linkage can we accurately incorporate local meteorological conditions, recent history of those meteorological conditions (i.e., emissions differ in rainy as opposed to dry conditions), clouds and available sunlight, and vegetation type. All these conditions may (and will) vary depending on global climate and climate change.

The biology in the upper oceans will also play an important role in atmospheric chemistry studies. Upper-ocean biology includes complex processes and interactions involving iron, nitrates, sulfur, ammonium, bacteria, zooplankton, phytoplankton, and sunlight. Ocean dynamics play an important role through the pumping of nutrients into the biologically active upper-ocean area. Changes in ocean dynamics could affect emissions and atmospheric chemistry. Figure 10 shows chlorophyll distributions across the Pacific basin before and after the onset of El Niño. Such a change in chlorophyll will change dimethyl sulfide concentrations and ocean emissions, which in turn impact the level of tropospheric sulfate aerosols and ozone.

Figure 10.
Simulated chlorophyll distributions in the Pacific Ocean before (left) and during (right) an El Niño event.
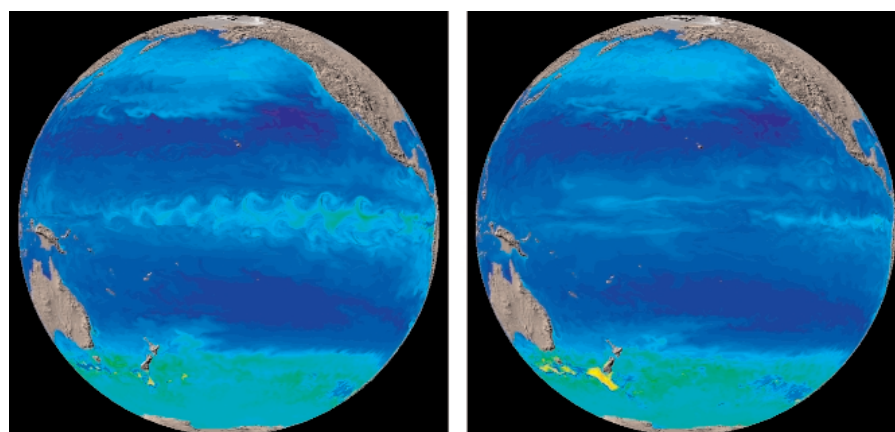
Coupled Chemistry/Climate Simulations Including Aerosols

Since pre-industrial times, the concentrations of a number of key greenhouse gases such as carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), and tropospheric ozone (O_3) have increased. Simultaneously, concentrations of stratospheric ozone have been decreasing. Although CO_2 is a dominant greenhouse gas, the non-CO_2 greenhouse gases combined are also very important to the radiative balance of the atmosphere. Changes in these chemical species change the radiative structure of the atmosphere, which in turn changes atmospheric dynamics that, completing the feedback loop, change the distribution of species. Atmospheric chemistry and climate are strongly coupled in the region near the tropopause; this is especially true for ozone and its precursors. Modeling studies will increasingly couple climate simulations with evolving chemistry (and, as discussed in the previous section, to biogeochemical models).

In addition to non-CO_2 gases, aerosols are very important to climate and climate change. Direct effects of aerosols (reflecting incoming solar radiation) have been included in some climate studies. However, indirect effects (changing cloud characteristics) are more complex and less understood. Indirect effects depend on aerosol size distributions determined by complex aerosol microphysical processes, such as nucleation, sedimentation, coagulation, and condensation. Atmospheric chemistry models are including those processes to better determine aerosol distributions so that climate effects can be studied.

Office of Fusion Energy Sciences

Fusion energy research is concerned with the physics of hot, fully ionized gas (plasma) in a magnetic field, and with the design and optimization of magnetic plasma confinement systems for fusion experiments. Also included in the Fusion Energy Science portfolio is the physics of intense ion beams for inertial fusion energy. Fusion energy science has a prominent history of advanced computation, and in fact, NERSC traces its roots to the Magnetic Fusion Energy Computer Center of the early 1970s.

A magnetically confined plasma is an extraordinarily complicated medium by any standards of scientific computing. 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 magnetohydrodynamic (MHD) phenomena, to many seconds for slow plasma evolution. The magnetic field is itself a dynamical quantity of interest, and the electric field and magnetic field perturbations created by the plasma provide global coupling terms in the plasma evolution equations. The confined plasma supports a large variety of electromagnetic waves, of interest for plasma heating and current sustainment and for plasma diagnostics, which may have widely disparate wavelengths and yet be tightly coupled through mode conversion processes. Plasma-wall interaction brings in neutral particle dynamics, impurities, and associated atomic and molecular processes.

In fact, plasma is the most common state of matter in the universe, and work done under OFES auspices has many connections to astrophysics, plasma processing of materials, particle accelerators, and other scientific fields.

Plasma and beam modeling has progressed enormously in the last decade. We now have massively parallel computer codes that are allowing us to extend the frontiers of our understanding in most of the key areas in fusion energy science. These codes also provide us with new capabilities to undertake detailed analysis of complex experimental results, and to design new experiments that are tailored

to meet their scientific objectives. The key areas being targeted
include:

- 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 studies 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.

- Electromagnetic wave/plasma interaction the study of the propa-
  gation and absorption of electromagnetic waves in plasma.

- 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 onto the target.

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 experimen-
tally 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 subse-
quently as the Plasma Microturbulence Project). These models have
been extensively benchmarked against independent linear calcula-
tions 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 and General Atomics DIII-D. 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 gyro-averaging the equations). The three models are gyrokinetic particle codes (Lagrangian description), gyrofluid codes with Landau closure, and gyrokinetic continuum codes (Eulerian description). All three models run efficiently on the NERSC T3E and SP2 and other massively parallel platforms. By developing and applying three different approaches to microturbulence simulation, the magnetic fusion community has been able to carefully explore the comparative computational efficiencies of the three approaches and perform important code crosschecks on the nonlinear simulations. Because the three 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 with benefits ensuing to other scientific disciplines faced with similar computational challenges.

We have 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. We have experience in being able to make successful use of new platforms at NERSC and to fully utilize NERSC resources. To date our 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 annular region of a tokamak experiment, albeit with realistic parameters (Figure 11), 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.
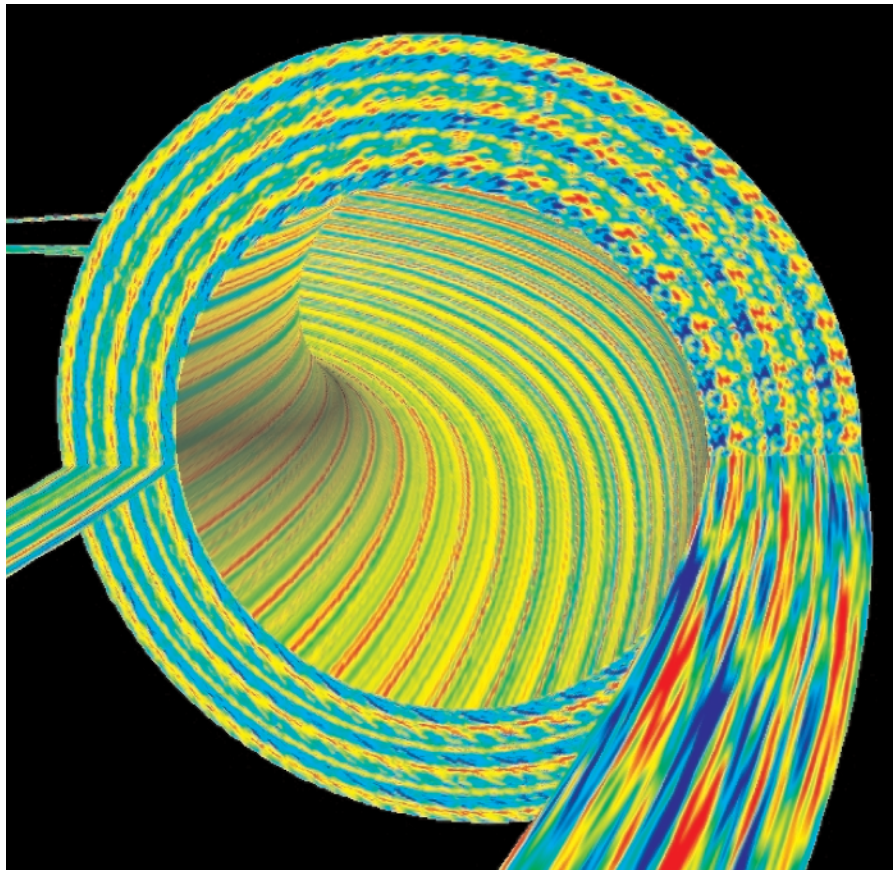
Figure 11.
Flux-tube gyrokinetic simulation of ion-temperature-gradient drift-wave microturbulence. Contours of density fluctuations. Image courtesy of A. Dimits, LLNL.

Our projections show that full-device simulations of ITG turbulence in an ignition-scale magnetic fusion device are ambitious but achievable on a 10-teraflop 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. The 10-teraflop computer will also make easier the inclusion of the more complete and better physics models that we have developed that include, for example, kinetic electron effects and electromagnetic coupling of the drift waves to kinetic shear-Alfven waves that modify the microinstabilities in the plasma at finite plasma pressures.

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

31

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 fusion devices. The next generation of burning plasma fusion devices will need to operate near their stability limits. These limits are set by an operational window 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.

The programmatic importance of this work is high. The extension of the MHD model and the higher numerical resolution will greatly improve our ability for evaluating and understanding the essential nonlinear mechanisms that set pressure limits in tokamaks and other plasma confinement devices. In addition to the programmatic benefits, this program 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.

A measure of the difficulty of MHD simulations is the *Lundquist number*, S, which is the ratio of typical resistive diffusion times to typical MHD wave transient times. Present computer power restricts us to values of S much lower than those typical of present-day and proposed experiments. An increase in computing power and improved algorithms will allow 3D nonlinear extended-MHD

simulations to include adequate space and time resolution to approach realistic values of the Lundquist number S under fusion conditions, and will allow solving the more complete and more computationally demanding extended-MHD models. For a model of a given mathematical complexity, the computational requirements for 3D extended-MHD scale like S^p where p = 1—2, depending on how effective the implicit time-stepping and adoptive mesh refinement are. Therefore, a factor of 100 increase in compute power alone will allow at least an order of magnitude increase in simulation values, and possibly two orders. Algorithmic advances and better single-processor optimization may lead to another order of magnitude. Thus, we can expect the fluid models, now limited by computational resources to compute at S values in the range of 10^5 to 10^6, to be able to approach simulating plasmas with fusion-relevant values of 10^7 to 10^8. The increase in computing resources will have an equally dramatic effect on enabling the codes utilizing the physical models that involve particle closures to resolve velocity space. As we are able to exercise these computationally demanding models with better resolution and for longer times, we will develop better understanding of the physical differences between fluid and particle closures, which we can expect to lead to improved and more efficient models (Figure 12).

Figure 12.
Time sequence of a simulated "reconnection event," or sawtooth oscillation, in the National Spherical Torus Experiment (NSTX) at PPPL. Red and green isosurfaces of constant pressure are shown. Some frames also show select magnetic field lines before, during, and after the recon-nection process. Initial (red) high-pressure region has been expelled from the center, and replaced by (green) lower pressure region in this sponta-neous self-regulating event. Image courtesy of W.Park, PPPL.
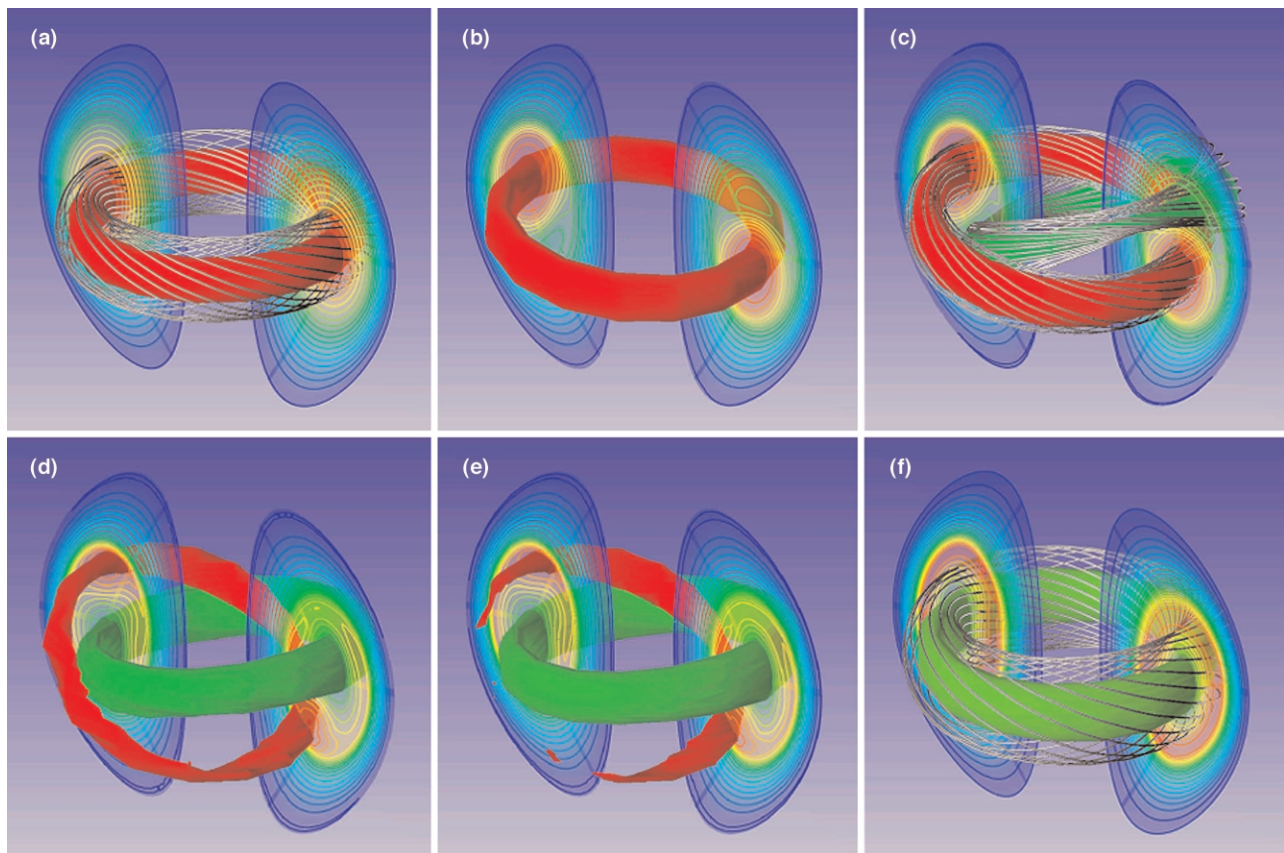
Stellarator Physics

Stellarators are 3D 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. Stellarator physics and simulation are of increasing interest in the U.S. Fusion program due to the possibility that new compact (Figure 13) experimental facilities will be built at several U.S. research centers. 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. 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 13.
(a) and (b) Top and side views of a QPS (quasi-poloidal stellarator) compact stellarator magnetic flux surfaces and filamentary magnet coils (color contours show magnetic field strength (in Tesla); (c) and (d) Top and side views of the NCSX (National Compact Stellarator Experiment—a quasi-toroidal device) magnetic flux surfaces. (Image courtesy of D.Spong, ORNL.)
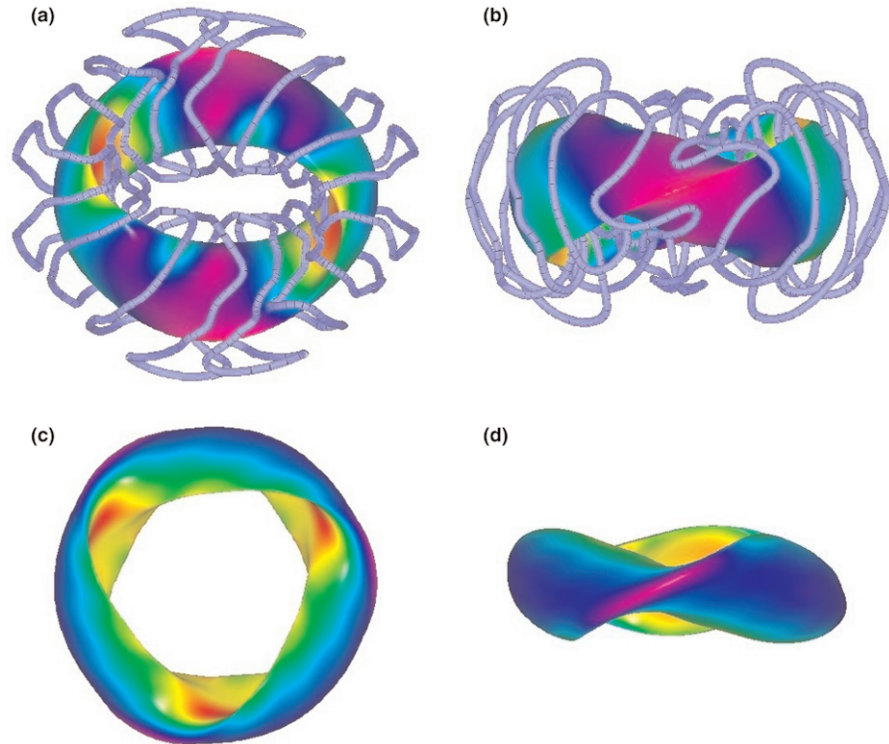
Stellarator Optimization and Design. There are two primary components in a stellarator design efort: the physics design, which varies the plasma shape to enhance selected physics criteria; and the coil design, which synthesizes magnetic coil geometries that will produce the optimized plasma shape. Both of these calculations involve a nonlinear optimization; this calculation has been adapted to parallel architectures by using a bank—queuing model (since the individual parallel tasks are generally of unequal length) with MPI

language calls for inter-processor communication. The physics optimization typically varies 40—60 shape and plasma profile parameters to optimize plasma transport, stability, and equilibrium. The coil optimization uses ~300 coil shape parameters to minimize the normal component of magnetic field on the outer plasma magnetic flux surface. A typical design involves multiple runs of the optimization code coupled with human mid-course corrections (e.g., adjustment of the optimization weights) to steer it into the desired region of design space. Once a design is obtained, it is evaluated in further depth with respect to its physics and engineering characteristics. As this process may go through many iterations, it has been essential to have access to massively parallel computing platforms which allow each isolated optimization run to be performed over a period of a few hours.

Stellarator Physics. Interest in developing better tools for understanding and simulating stellarator plasma physics is motivated both by the design efforts described above, as well as by several large planned and existing $1 billion-class stellarator experiments (the LHD device in Japan and the W7-X device in Germany) in the world fusion program. Areas that are under development include: plasma equilibrium, transport, heating, and stability. Due to the three—dimensional nature of stellarators, these calculations can easily require several orders of magnitude more computational time than those for an equivalent axially symmetric tokamak device. Although much of the plasma physics in stellarators can be analyzed using similar equations and models as developed for tokamaks, the addition of an extra dimension to the equilibrium leads to a much greater emphasis on performance and parallel scalability.

The accurate calculation of plasma equilibria in stellarators is basic to all the other physics models. It is also necessary in order to reconstruct experimental equilibrium characteristics from magnetic measurements external to the plasma. In general, mathematically smoothly varying equilibria do not exist for three-dimensional configurations and some degree of sheet current and magnetic island formation will always be present near rational magnetic surfaces. Including these features demands high computational resolution and large memory requirements.

In the area of transport, both the confinement of the thermal plasma and of energetic particle components used to heat the plasma is of interest. Particle simulation models have been developed that have good scaling on parallel computers, but can be difficult to apply to very low collisionality due to the long simulation times required to attain a collisional equilibrium. Direct solutions of the continuum

kinetic equation can address this issue, but provide their own challenges due to the high dimensionality (3 spatial and 2 velocity dimensions) involved and sharp boundary layers in the distribution function at low collisionality.

Plasma heating efficiency in stellarators using neutral beams is generally analyzed using particle simulation methods. RF (radio frequency) heating is also used extensively in stellarator experiments due to its lower cost and greater flexibility. New computationally intensive tools are under development to calculate the RF wave field and its plasma absorption for 3-dimensional systems.

Stability in compact stellarators will encompass large-scale current-driven modes, short-scale-length pressure-driven ballooning modes, energetic-particle-driven instabilities, and plasma microturbulence. Although several linear variational stability codes for addressing current-driven and ballooning instabilities are already developed (and coupled into the optimization efforts), relatively little has been done in the area of nonlinear MHD evolution and microturbulence simulation codes. As low-aspect-ratio experiments become operational, there will be increasing interest in having the latter two types of codes available for interpreting these experiments.

Optimized compact stellarators are a recent innovation in magnetic fusion research with the potential to lead to an attractive, cost-effective path to fusion power. Their development was not possible until recently due to the requirements for high-performance parallel computing and the creation of algorithms that efficiently utilize parallel computers. The anticipated construction of experiments of this type will lead to continued and increasing reliance on numerical simulation in order to optimize their performance and interpret the observed plasma physics.

Electromagnetic Wave/Plasma Interactions

Electromagnetic waves play a fundamental role in the dynamics of plasmas ranging from the solar corona, to planetary magnetospheres, to the Earth's ionosphere, and to laboratory devices. However, it is in the complex media that arise in fusion-relevant laboratory plasmas that understanding of wave dynamics poses perhaps the most significant scientific challenge of any area of wave physics. The plasmas in fusion devices are bounded and inhomogeneous, having essential variations in at least two spatial dimensions and in many cases three dimensions. Due to the confining magnetic field, these plasmas are also anisotropic. Within these plasmas there can be many distinct wave modes propagating in a single plasma region

at a given frequency. The different wave modes can coexist at widely separated wavelength scales, or can interact when spatial variations of the plasma bring the wavelengths close together a process called mode conversion. There are numerous mechanisms by which the waves can be absorbed collisional damping, collisionless Landau and cyclotron damping, damping by collisionless stochastic particle interactions. The waves can produce nonlinear modifications of the plasma medium, thereby altering the wave propagation characteristics and changing the macroscopic properties of the plasma, which in turn affect other processes occurring in the plasma such as stability and transport. It is important to the future advancement of wave-plasma interaction physics that all of these wave and plasma effects be self-consistently analyzed in the appropriate, multi-dimensional geometry.

Although complex plasma wave processes are intrinsically interesting, understanding these processes has practical significance since externally introduced waves constitute an important and flexible tool for penetrating a high-temperature plasma and for controlling it. Recent assessments of the fusion program have emphasized the programmatic need to develop techniques to control the various interacting processes at work in fusion plasmas and have recognized the potential for electromagnetic waves to affect much of this control. A collection of codes has been developed which can approximately model each of the three major processes of interest: coupling of wave power into the plasma; propagation of the excited waves from the plasma edge into the core regions and subsequent absorption of the wave power by the plasma; and longer-time-scale modifications of the plasma distribution function. There has been considerable effort in benchmarking these codes against each other and against experimental observations, so that there is good confidence in the underlying system of equations. Nevertheless, a fully comprehensive, quantitatively accurate model that treats all of these processes self-consistently in sufficient detail has yet to be developed. To a great extent, progress has been impeded by present computational limitations on execution speed and memory size. With the advent of the new generation of massively parallel computers, there is the opportunity to develop robust predictive codes that can self-consistently incorporate the critical processes on different scale lengths.

The fundamental difficulty can be traced to the necessity of retaining multiple physical processes that are significant on widely disparate spatial and temporal scale lengths. In essence, constraints on the maximum memory size available and required run time limit the extent to which the various physical processes can be incorpo-

rated and self-consistently linked in the numerical models. As a specific example, mode conversion of long-wavelength electromagnetic waves (such as fast magnetosonic waves) to short-wavelength electrostatic waves (such as ion or electron Bernstein waves) is currently understood in detail only in one-dimensional models. However, in fusion applications as well as in space plasmas, the wave propagation and mode conversion processes occur in plasma equilibria that vary in two or more dimensions. In an idealized one-dimensional model, all plasma properties vary only in a single direction (say the x direction), so that as the wave propagates from one plasma region, if it undergoes mode conversion, it will do so along an infinite x = constant surface, and it must cross this surface as it propagates in x. If the plasma properties vary in two or more dimensions the conversion region may have a more complicated geometry. This is analogous to the topological properties of a closed curve. If considered as part of a 2D surface, to get from inside to outside of the closed curve it is necessary to cross the boundary, whereas if the curve is embedded in a higher dimensional space there are richer possibilities one can jump over the boundary.

With the present technology, it is difficult to include non-local effects in the plasma dielectric that arise from 2D variations in equilibrium or to correctly resolve the short-wavelength modes in two or three spatial dimensions. Making a restrictive approximation that the ratios of the ion gyroradius to the wavelength and to the background plasma equilibrium gradient scale lengths are small, one can obtain 2D differential models for the plasma dielectric response function that include cyclotron resonances only up to the second harmonic. These models ignore non-local effects on the plasma dielectric response function arising from the finite orbit size of the plasma particles, from trapped versus passing particle distributions, from 2D gradients and curvature in the equilibrium magnetic field, and absorption and mode conversion effects at higher cyclotron harmonics. Even within the differential model for the plasma dielectric response in 2D, the run times for the codes are prohibitively long and it is difficult to resolve the modes even on the largest computers currently available.

Our community is increasingly committed to using massively parallel processing to address the key physics issues described above. As an example (Figure 14), we have recently developed an alternate full-wave model, based on spectral rather than finite difference or finite element methods. This model, AORSA (or all orders spectral algorithm), takes advantage of new computational

techniques for massively parallel computers to solve the integral form of the wave equation in multiple dimensions without any assumption regarding the smallness of the ion Larmor radius relative to the perpendicular wavelength.

(a) Re (E_η)

(b) Blowup

$k_\perp \rho = 1.2$
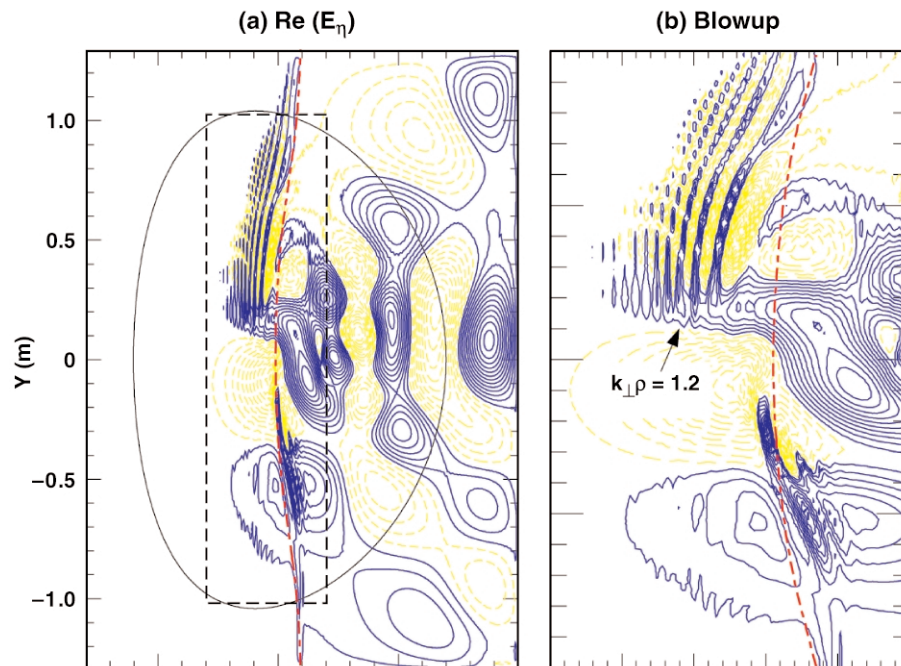
Figure 14.
All-orders spectral solution for mode conversion in DIII-D including the poloidal magnetic field. The dashed curve shows the mode-conversion layer from 1D cold plasma theory. (Image courtesy of D.Batchelor, ORNL.)

Intense Ion Beams for Inertial Fusion Energy

The Heavy Ion Fusion (HIF) program's principal mission is to develop the body of knowledge needed for Inertial Fusion Energy (IFE) to realize its promise. Heavy-ion-beam-driven IFE is the principal alternate approach to fusion energy; it has a number of favorable attributes and is very different from magnetic confinement. In Heavy Ion Fusion, intense beams of heavy ions (with masses in the range 100—200 AMU) will be accelerated to multi-GeV kinetic energies (several megajoules total), temporally compressed to durations of ~10 ns, and focused onto a series of small (few-mm) targets, each containing a spherical capsule of fusion fuel. The capsules are compressed and heated to a point where fusion ignition and a propagating burn occur. The energy so produced is captured and used to heat a working fluid, and electricity is produced using conventional steam turbine generators. Heavy ion drivers are attractive for this purpose because of their efficiency and because the final focusing onto the target is achieved by magnetic lenses which can be made robust to the effects of the

target explosions, which must repeat at rates of the order of 5 Hz. While such a system is many years from fruition, ongoing experiments in the United States (under OFES support) are developing the novel and challenging intense-beam physics needed for its realization.

The intense ion beams that will drive HIF targets are nonneutral plasmas and exhibit collective, nonlinear dynamics that must be understood using the kinetic models of plasma physics. This beam physics is both rich and subtle: a wide range in spatial and temporal scales is involved, and effects associated with both instabilities and non-ideal processes must be understood. Ion beams have a long memory, and initialization of a beam at mid-system with an idealized particle distribution introduces uncertainties; thus, a key goal is to develop, and to extensively use, an integrated and detailed source-to-target HIF beam simulation capability. The major issues include:

Long-term evolution of space-charge-dominated beams. In the driver, the array of beams is accelerated by inductive electric fields, and is confined by applied focusing fields. The dynamics are space-charge dominated, that is, governed by a balance between applied fields and space charge forces, and 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 necessity to focus the beams ultimately onto a small (few millimeters) focal spot on the fusion target. 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 (especially the injection of stray ions, electrons, and neutrals into the beamline), 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. Collective beam interactions with stray electrons in the accelerator and transport lines must be understood quantitatively. This area is computationally challenging because of

the ratio between the timescales for electron motion and those for electron build-up; the need to efficiently gather/scatter and communicate multi-species information for ionization and surface-physics processes; and the needs for efficient dynamic load balancing and perhaps an adaptive mesh.

Beam interactions with fusion chamber environment. 3D simulations of the propagation of the cluster of beams through the final focusing optics, and onward through the fusion 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 photo-ionization of the beam ions and background gas. Multiple-beam interactions near the target will be one important 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 timescales not essential to the physics) or a magneto-inductive (Darwin) model that eliminates light waves from the description is also needed.

Representative images from simulations of a variety of problems are shown in Figure 15. Such HIF beam simulations at the current state of the art are playing an important role in the inertial fusion energy program. 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-beam effects in both the driver and fusion chamber must be understood, the research effort stands to benefit greatly from a qualitative increase in available computational resources. PIC simulation of a single beam in a driver-scale accelerator will require on the order of 100 million particles, at least 100,000 time steps and a computational grid of at least 128×128×4096 mesh points; the minimal total time required for such a driver simulation is estimated to be about six hours on a 10-teraflop machine. High-fidelity simulations of driver beams, with detailed applied-field models, magneto-inductive effects, and ultimately multiple-beam interactions, will require significantly more resources. In the fusion chamber, the requirements are

comparably challenging. We estimate a wall-clock time of order one day on a 10-teraflop system for a fully electromagnetic 3D chamber simulation with 16 interacting ion beams. This simulation will include the detailed physics of the beams interactions with the fusion chamber plasma and the target; ultimately the full ~100-beam cluster will need to be simulated.
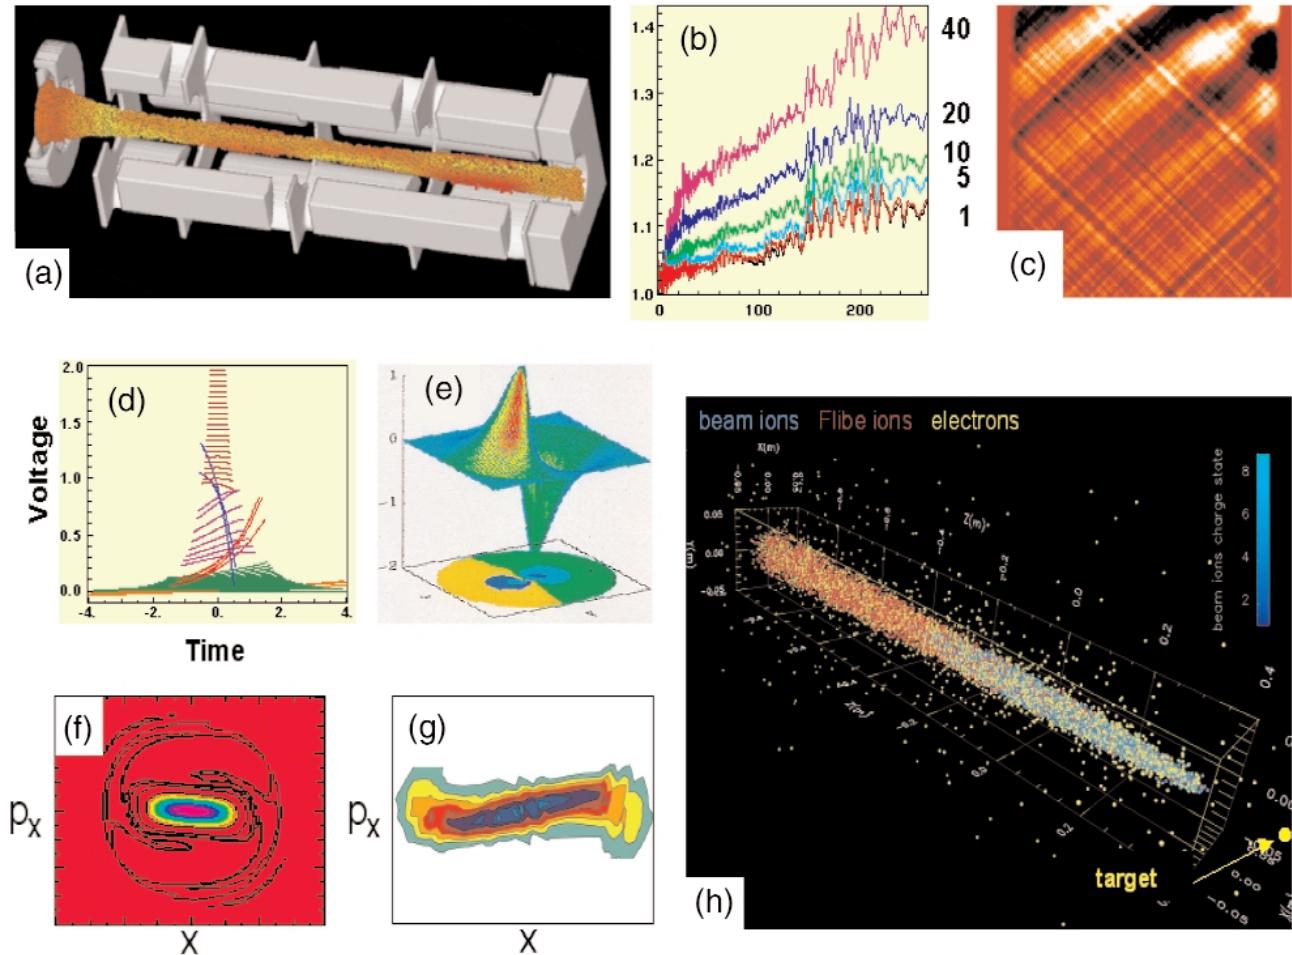
Figure 15.
Representative output from HIF beam simulations: (a) WARP3d particle-in-cell simulation of space-charge-limited emission off a curved surface, and acceleration in a 3D structure, including sub-grid-scale placement of conductor boundaries (cut-cell method); (b) WARPxy particle-in-cell study of beam emittance at mid-pulse versus time in an imperfectly aligned beamline, for five different intervals between applications of steering; (c) WARP3d study of longitudinal waves on beam, driven unstable by impedance of accelerating structures; (d) accelerating waveforms for a possible future experimental accelerator, for use in WARP3d simulations; (e) BEST nonlinear-perturbative simulation of unstable electron-ion two-stream mode in a beamline; (f) semi-Lagrangian Vlasov simulation of beam halo generation due to anharmonic focusing fields, using prototype model in WARP-SLV; (g) distorted beam phase space in final focusing, as simulated using WARPxy; (h) BPIC simulation of beam transit through fusion chamber environment and onto the target. (Images courtesy of A. Friedman, Heavy Ion Fusion Virtual National Laboratory, a partnership among LBNL, LLNL, and PPPL.)

42

Simulation of Fast Ignition for Inertial Fusion Energy

One of the most promising, albeit high risk, IFE concepts is that of the Fast Ignitor (FI). While both Direct and Indirect Drive Laser Fusion appear feasible, a critical challenge is robust formation of the central hot spot for ignition and propagating burn. If the pellet shell does not implode uniformly with initial masses smooth to a few hundred and implosion drives symmetric to 1—2% it is predicted that hydrodynamic instabilities will destroy the pellet before hot-spot ignition conditions can be realized. The necessary stability requirements cascade into numerous constraints on the design and economics of IFE reactors (i.e., power balance for the driver beams; reactor chamber first wall protection; target injection). However, it has been experimentally demonstrated that existing laser systems can compress spherical laser pellets to densities that are high enough for fusion to occur. What is required is a match to ignite the compressed fuel, in order to achieve an adequate areal density, density and temperature product for cost-effective fusion. In the fast ignitor concept, the heating energy is envisioned to be provided in a separate step, by a very-high-intensity, tightly focused laser beam which penetrates the plasma corona surrounding the pre-compressed pellet, and deposits energy sufficient to spark the cold fuel to ignition. This theoretical concept is made practically attractive by: (1) the recognition/ demonstration that efficient pellet compression to high densities can be achieved with relatively crude drivers and targets; and (2) the ongoing development of potential ignitor beams: short-pulse (<10 psec), high-brightness (focused irradiances $>10^{19}$ W/ cm^2), high-power petawatt lasers. Calculations project that FI pellets can produce gains comparable to those achievable with central hot-spot ignition, with sub-megajoule drive energies, enabling very flexible reactor concepts.

Full-scale modeling of the fast ignitor concept will require a birth-to-burn FI laser-matter interaction numerical simulation capability. Codes currently exist which well simulate the development phase, and the final pellet implosion phase; and, codes also exist which simulate on a small spatial scale the fundamental interaction of a high-intensity petawatt laser beam with a solid or pre-heated laser target. To bridge the regimes from the particle-dominated petawatt-laser matter interaction, through the fluid-dominated compression and finally ignition, and into the novel area of asymmetric propa-gating burn, will require significant basic plasma physics research in areas of parameter space for which fluid models are generally inapplicable and validated hybrid models do not yet exist. Current FI computational research thus necessarily comprises a fully kinetic

module that can model laser absorption, transport of energy, and heating of the ion background. To achieve FI, severe constraints have to be met which at present can only be modeled by PIC techniques. The PIC numerical scheme is a finite element method for solving the Vlasov equation, a method that shows excellent scalability on massively parallel computing platforms. However, despite recent progress in parallel computing, the kinetic PIC part only of a reasonable FI simulation would take more than ten days on a computing platform that is three orders of magnitude larger than the T3E. This is a multiple tera-scale application. Full FI simulations, with hydrodynamic, hybrid, and kinetic aspects, will extend significantly beyond.

Although multiple tera-scale simulations are out of scope at present, single terascale simulations with existing PIC codes can help to understand important processes relevant for FI. The most prominent ones concern transport of laser light and particles in bounded plasma, as illustrated, e.g., in Figure 16. It is important to learn how laser-generated electron spectra look, which major energy groups they contain, and for which specific phenomena each group is responsible. For example, we know that the least energetic electrons are inhibited by extremely large magnetic fields. Hence, their range is limited to the plasma around the critical layer where the laser deposits its energy. Faster electrons have a larger range. They tend to propagate in magnetized current filaments while the most energetic electrons seem to follow ballistic trajectories which obtain their initial parameters at the critical density of the plasma. We need to learn how the absorbed laser energy is distributed among all accelerated electrons for obvious reasons. In the context of FI the only sensible definition of energy deposition is the rate of energy gain of the ions that compose the thermonuclear fuel. It may be clear how to model ion heating in a diffusive plasma but it is far from clear how much energy background ions gain in a non-equilibrium non-diffusive laser plasma. Furthermore, plasma boundaries are known to have a sizable impact on transport and absorption. Because of ambipolar fields present at short gradient length boundaries, energetic electrons can be guided. This may impair FI efficiency. At present, numerous experiments are being carried out with an aim to better understand laser absorption, particle transport, and acceleration. Terascale computing with existing, fully parallel PIC codes can significantly support these experimental efforts, and will, further, enable the efficient survey of a wide range of possible operating parameters relevant for fast ignitor IFE.
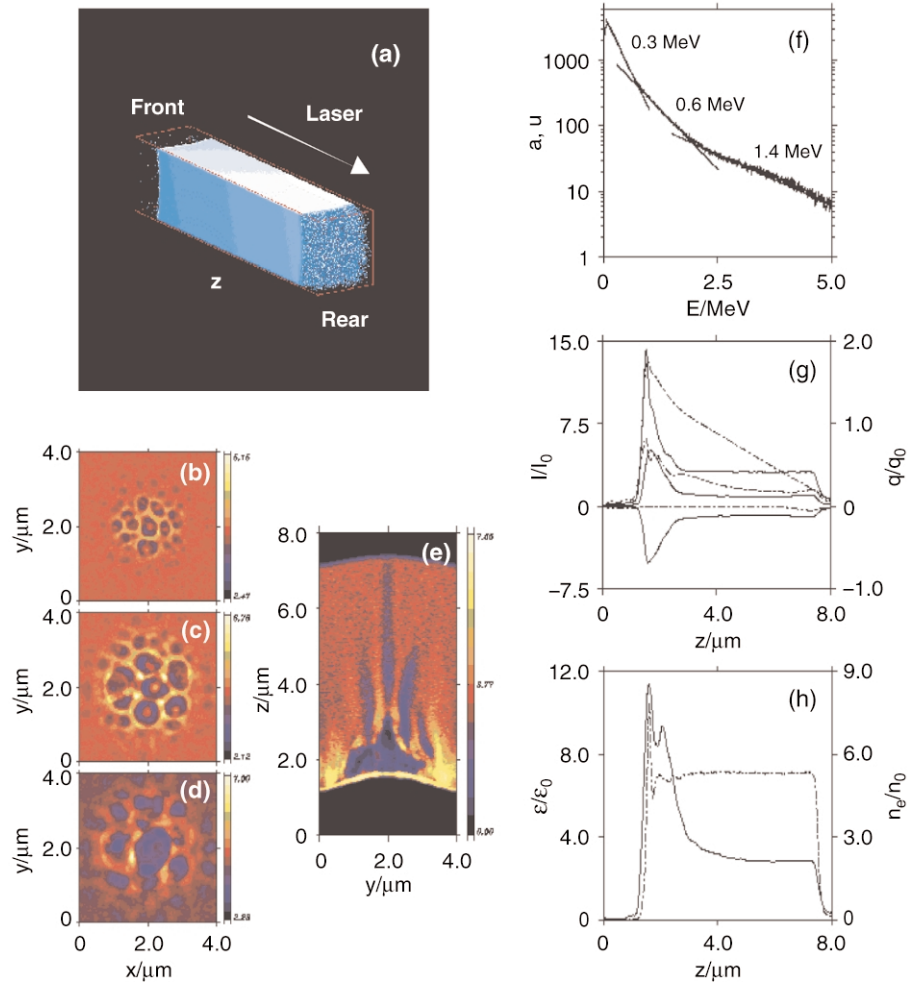
Figure 16.
3D particle-in-cell simulation of an intense laser beam interacting with a short-gradient-length plasma slab which is many times over critical. Plot (a) shows the electron density in 3D. Fast electrons are coming off the rear surface. Within the plasma electrons form magnetized current filaments that lead to ion density filaments as shown in plots (b)–(e). The laser-generated electron spectrum shows three distinct energy groups plotted in (f). Electrons with temperatures of 0.3 MeV belong to the return flow. The filaments are formed by electrons with 0.6 MeV temperature while the fast electrons ballistically cross the plasma slab. Plot (g) shows the total current and energy flows separated in flows and return flows. Plot (h) shows charge and energy densities. For brevity, units are not stated.

Office of High Energy and Nuclear Physics

Research in the High Energy and Nuclear Physics (HENP) division falls under two major subfields. 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. 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. Research efforts focus on, broadly speaking, understanding origin of mater and the way it interacts.

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 HENP research taking place at NERSC. These are computational studies of quantum chromodynamics (QCD), nuclear structure, astrophysics, simulations of accelerator design, and examples from experimental particle and nuclear physics, including the quark-gluon-plasma phase transition, the solar neutrino abundances, and studies of CP violations in Lambda and Xi hyperon decays as well as in charged kaon decays.

Lattice QCD

Most of the computational effort in lattice gauge theory is devoted to the study of QCD, which is the theory that describes the interactions of quarks and gluons. These particles are bound together by the strong force to make up protons, neutrons, and many other particles, and residual strong interactions among these particles are the starting point for nuclear physics. QCD is very successful at predicting results of experiments at many GeV, where perturbative calculations can be reliable. However, at energy scales of a few hundred MeV to a few GeV, or length scales around one fermi, QCD is a strongly interacting nonlinear theory, and most of the hadronic properties cannot be calculated analytically.

Lattice QCD calculations have several goals. First, we would like to understand how the simple fundamental theory leads to fundamental properties like the masses and decay rates of hadrons, or their internal structures. Secondly, parameters of the standard model, namely quark masses and CKM matrix elements, are explored in hadronic experiments. To extract the fundamental parameters, the matrix elements of the corresponding operators in the initial and

final hadronic states must be known, and many of these matrix elements are best calculated using lattice simulations. Finally, the early universe contained, and neutron star interiors may still contain, matter with temperatures and densities extreme enough that hadronic matter enters new phases, and these phases are studied through lattice simulations. Recent and ongoing work at NERSC addresses all of these aspects of QCD. Among the problems studied are:

Masses and interactions of the light hadrons. Computations of light hadron properties are both a testing ground for evaluating lattice methods and a subject that can give insight into the workings of QCD. Many groups have worked on these computations over the years. The MILC collaboration has recently used NERSC facilities for exploratory studies of the light hadron spectrum using an improved lattice discretization that greatly reduces some of the troublesome lattice artifacts. In addition to the masses, lattice QCD can be used to compute matrix elements that govern interactions and decays of these particles. Recently promising techniques, domain wall quarks and the closely related overlap action, were developed. These techniques may make it possible to compute some delicate quantities that have previously been inaccessible.

Nature of the QCD vacuum. NERSC facilities are used to investigate the role of instantons, which are topologically non-trivial field configurations that are thought to be important in the QCD path integral. Such studies may lead to an explanation of how quarks and gluons become hidden inside colorless hadrons.

Heavy quark physics. Quark masses and mixing angles, the CKM matrix elements, are free parameters in the standard model. As such, their values may well give insight into physics beyond the standard model. While the relevant experiments for the light quarks were done long ago, measurement of these parameters for the heavy b and t quarks is of intense experimental interest, as are the theoretical calculations which support these efforts. Several current projects at NERSC involve these problems. One group is working to compute renormalization constants for $O(a)$ improved lattice QCD, and to use these constants to extract fundamental parameters. Another group is exploring the use of an anisotropic lattice to study hadrons containing these heavy quarks.

QCD at high temperature and density. The collective behavior of large numbers of quarks and gluons at extreme temperature or density is an intense computational challenge. Again, several important projects are under way at NERSC. Near possible second-order phase transitions, QCD's behavior should have universal

properties determined by the symmetries of the order parameter. This idea can be tested, albeit with practical difficulties, by dialing quantities such as quark masses, number of flavors, or even the gauge group. Several projects at NERSC are exploring these questions, and developing new lattice techniques to get around practical obstacles. At the other extreme, the MILC collaboration is using NERSC machines in an attempt to push calculations closer to the real world, using three quark flavors with an improved action, in hopes that theoretical predictions can be tested against experimental results from RHIC or the LHC.

Lattice QCD simulations make large computational demands. The necessary computational power and memory are available only on massively parallel supercomputers or on specially built machines. The steady growth in available computer power, coupled with progress in algorithms, has steadily increased the number of problems accessible to simulations, and improved the reliability of results for problems that have been studied for some time. Gaining better resolution on the numerical lattice is one of the major challenges of lattice QCD. We illustrate the problem of lattice resolution by investigating the famous painting shown below. At the lowest resolution, of 9 by 12 pixels, most people would not recognize this work. At 12 by 16, it is still obscure. At the next finest resolution, people recognize the painting as being the Mona Lisa. With two more refinements of the resolution, you can begin to see her eyes and smile, but only at the finest resolution does her beauty really show (Figure 17).

Figure 17.
Resolution is a key component to understanding lattice QCD results.

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 like 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.

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. 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, Jefferson Laboratory), and to compute astrophysical reaction rates, many of which cannot be measured in the laboratory.

This year the Green's Function Monte Carlo (GFMC) group finished a study of new three-nucleon potential terms. Previously the GFMC group demonstrated that the Hamiltonian that has been used successfully for more than a decade in studies of s-shell nucleon 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 Green's function and variational (VMC) Monte Carlo 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 (Figure 18).

Figure 18.
GFMC calculations of mass 4 through 10 systems.
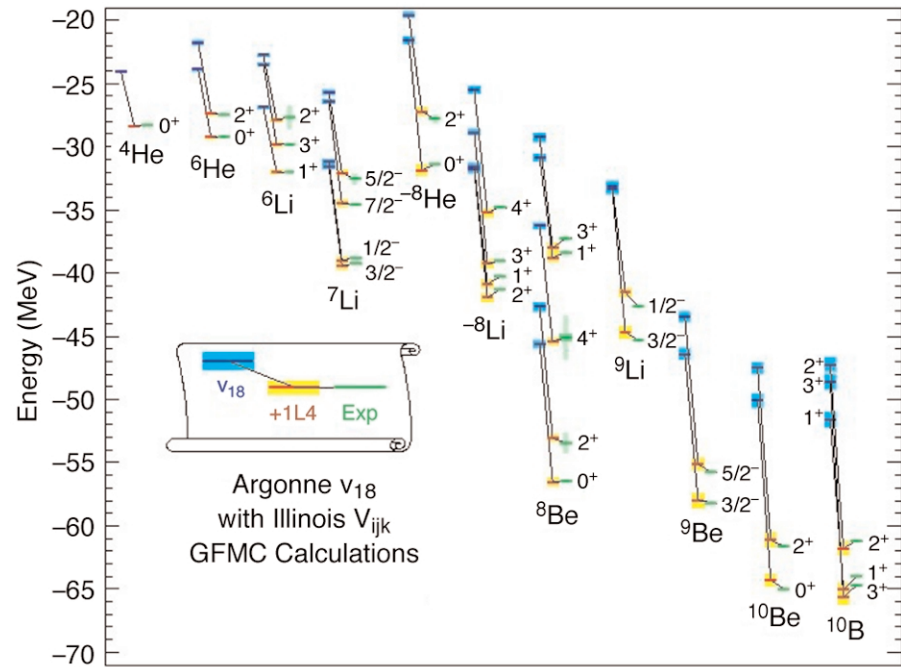
Nuclear structure studies in medium mass nuclei require an implementation of Auxiliary Field Monte Carlo (AFMC) techniques to calculate ground state and thermal properties of these systems. The calculation of the heat capacity of a finite interacting nuclear system including correlations is a difficult 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. In particular, the thermal energy $E(T)$ can be calculated versus temperature T, and the heat capacity can be obtained by taking the numerical derivative of E with respect to T. Recent calculations of the heat capacity of isotopes of iron nuclei using NERSC computing resources have shown that a residual signature of a pairing phase transition exists in these finite nuclear systems, and that the transition becomes stronger with increasing neutron number.
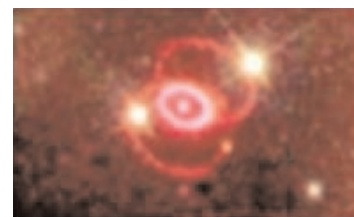
Electron capture on nuclei is an important process that occurs during the early stages of the evolution of core-collapse supernovae. These weak captures serve to deleptonize the core of the massive star and determine the final electron fraction within the core. Therefore they set the size of the homologous 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 require also significant computational resources. Although the iron region has been sufficiently 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. During the last year investigations began of the neutron-rich germanium nuclei using AFMC techniques to investigate the electron capture rates in these systems. AFMC calculations have also been extended to other quantum many-body systems such as quantum dots.

Astrophysics

We highlight here two astrophysical problems that are currently being pursued using NERSC computational resources. The ultraviolet spectra of nearby Type Ia supernovae (roughly, those less than a billion light-years away) are used to determine what effects their evolutionary history may have on their use as standard candles to make fundamental cosmological measurements. Since all Type Ia supernovae have nearly the same brightness, they have been used to compare the distance and redshift of the galaxies in which they occur, and thus to measure the expansion rate of the universe. The conclusion that the universe is accelerating, reached in 1998 by the Supernova Cosmology Project and colleagues in the High-Z Supernova Search Team, was based on comparing dozens of Type Ia supernovae spanning a huge range of redshifts (Figure 19).

Figure 19.
Measurements of supernova like SN1987a are used to measure the expansion rate of the universe.

Type Ia supernovae are similar because their progenitor stars are white dwarfs that accrete mass from their partners in a binary system, exploding in a thermonuclear blast when the mass exceeds a specific limit, 1.4 times the mass of our sun. However, very distant supernovae those up to 8 billion light-years away exploded when the universe was only half its present age, and their progenitor stars may have been poor in elements heavier than hydrogen and helium. Hydrogen and helium and a little bit of lithium were made in the big bang, but heavier elements were made later in stars and supernovae. Even in the earlier universe the age of individual galaxies varied greatly, but since nearby supernovae occur in an older universe, some of their progenitors may be richer in metals.

The ratio of heavier elements to hydrogen and helium is known as metallicity. The IBM SP supercomputer at NERSC was used to develop a theory of how metallicity should affect the spectra and light curves (patterns of rising and falling brightness) of supernovae, investigating whether differences between nearby and distant supernovae might lead to systematic errors in comparing their redshifts and magnitudes.

These studies predicted what the spectrum should look like if a supernova's progenitor star is much poorer or much richer than the sun in iron, cobalt, and nickel, and concluded that distinct differences would be evident in the ultraviolet part of the spectrum, although there should be virtually no difference in the optical spectrum. Since observations are based on optical spectrum measurements, the expanding universe appears to be a reality.

The search for the explosion mechanism of core-collapse supernovae and the computation of the nucleosynthesis in these spectacular stellar explosions is one of the important and challenging problems in computational nuclear astrophysics pursued at NERSC. Core-collapse supernovae are the among most energetic explosions in the cosmos, releasing tremendous amounts of energy in the form of neutrinos of all flavors, disrupting stars more massive than ten Suns and disseminating and producing many of the elements in the periodic table, without which life as we know it would not exist. They are a nexus for nuclear physics, particle physics, fluid dynamics, radiation transport, and general relativity, and serve as cosmic laboratories for matter at extremes of density, temperature, and neutronization that cannot be produced in terrestrial laboratories and physics beyond the Standard Model. With the computing power afforded by vector and massively parallel supercomputers at NERSC, we are presented with a unique opportunity to finally solve one of Nature's more important problems, and in so doing, take a big step toward understanding how life as we know it became possible.

At the heart of the solution of the multidimensional integro-partitial differential equations governing the radiation hydrodynamics of supernova explosions is the solution of large, sparse linear systems of equations involving as many as 10^{11} unknowns and requiring up to tens of terabytes of memory. These systems must be solved multiple times per time step for the order of $10^{4\text{-}5}$ timesteps. The supernova application is a multiscale, multidimensional, multiphysics application. One must come to grips with the convective environment in the wake of the supernova shock, the disparate

timescales associated with the neutrino-matter interactions and the explosion, and the disparate spatial scales associated with the neutrino mean-free path in the proto-neutron star and the evolution of the protoneutron star itself.

Accelerator Design

Particle accelerators are among the largest, most complex, and most important scientific instruments in the world. They have enabled a wealth of advances in applied science and technology, many of which have huge economic consequences and many of which are greatly beneficial to society. They are also critical to research in the basic sciences (such as high energy physics, nuclear physics, materials science, chemistry, and biology). In particular, accelerators are the most versatile and powerful tools for exploring the elementary particles and fields of the universe. Experiments associated with high-energy accelerators led to some of the most remarkable discoveries of the 20th century. Near-term experiments are likely to be just as exciting, if not more so, with the possible discovery of new physics beyond the Standard Model, such as supersymmetry and its associated implications for a radical new geometry of space—time, which will fundamentally change our view of the universe.

Given the great importance of particle accelerators, it is imperative that the most advanced computing technologies be used for their design, optimization, commissioning, and operation. This is especially important for the design of the next generation of particle accelerators. In fact, all near- and far-future accelerator designs have very challenging modeling requirements that require high-performance computing (HPC), as will now be described:

High intensity proton drivers needed for conventional neutrino Superbeams, neutrino factories, and muon colliders require precise predictions of the effects of space charge. This need is shared by currently operating proton drivers, like the FNAL Booster and the BNL AGS, which are experiencing significant losses, currently attributed to space-charge effects at injection. The losses at the FNAL Booster are currently the biggest issue for the success of the near-future FNAL program (RunII+neutrino program). Due to the nature of this type of problem which involves long (high aspect ratio) bunches propagating for thousands of turns including space-charge effects and wakefield effects a full 3D simulation is prohibitive using the current algorithms and existing multi-processor hardware. Simulations using roughly 100 processors have been estimated to require 1 year of computer time.

Next-generation linear colliders require demanding computer simulations in regard to both electromagnetic and beam dynamics modeling. For example, extremely complicated 3D electromagnetic structures for the NLC must be modeled and analyzed with greater speed, accuracy, and confidence than has previously been possible. Presently popular serial electromagnetics codes are inefficient in handling complex geometric shapes, or are limited in their ability to solve large-scale problems. However, the recent development of parallel eigenmode and time-domain codes has already increased our modeling capabilities by roughly three orders of magnitude. In addition to modeling electromagnetic components, HPC capabilities are needed to model beam dynamics in linear colliders. For example, in both the NLC and TESLA designs, the accurate treatment of space-charge effects and other collective effects is important to predicting the beam's behavior in the damping rings. In order to validate the basic operational characteristics of these machines, the linac and beam delivery systems need to be modeled including component fluctuations, tuning, and feedback systems. Such simulations are impossible on serial computers, where the execution time to run one such code with the desired accuracy has been estimated to be 1 year per processor.

Very large hadron colliders like the VLHC require HPC capabilities in areas such as long-term tracking to predict dynamic aperture, self-consistent simulations of beam—beam effects in the strong—strong regime, predicting the thresholds for instabilities (such as the electron-cloud, resistive wall, and transverse mode coupling instability), and the simulation of beam/material interactions (e.g., energy deposition from collision byproducts) that address safety and environmental issues. In addition to these conventional requirements, there are also operational ones involving the use of HPC to develop orbit correction algorithms, alignment procedures, etc., that are challenging due to the size of the machine, the large amount of diagnostic data, and the short period of time in which the analysis has to be performed. Here accelerator simulation is used in a similar way to HEP experiment simulation: accurate modeling of the machine and diagnostics are used to develop and optimize analysis algorithms, which are then used to optimize machine operation. Like a linear collider, full system simulations of the VLHC including beam dynamics and feedback systems are needed to verify operational characteristics of the proposed design.

Neutrino source/muon colliders present unique modeling challenges due to the fact that they involve ionization cooling. Ionization cooling requires accurate modeling of muon/matter inter-

actions, especially energy loss and multiple scattering. There are a few codes that share the physics description of the above processes borrowed (or directly implemented) from HEP modeling packages. These codes are very slow, prohibiting accurate simultaneous optimization of the sub-systems of the design, although in many cases both performance and cost of these sub-systems are dependent on each other. In addition, for high-intensity muon colliders space charge effects are crucial at the final stages of cooling. In both cases HPC is needed. An initial effort to embed cooling simulation capability in an HPC beam dynamics code has been successful, providing a good base for further development.

Besides the design of next-generation accelerator complexes, HPC is also needed, in concert with theory and experiment, to explore and develop novel methods of acceleration like *plasma-based and laser-based acceleration* techniques. Using these techniques, extremely high gradients (up to 100 GV/m) have been measured over short distances in the laboratory. The challenge is to control and stage high-gradient sections so that one can produce high-quality, high-energy beams in a less costly, more compact configuration than would be impossible using conventional technology. Beyond applications to HEP, such compact accelerators would have huge consequences in other areas of basic and applied science, industry, and medicine. However, modeling these complex systems requires solving the 3D coupled Maxwell/Vlasov equations. Given that the phenomena involve multiple length and timescales (a situation that is particularly challenging when the laser wavelength must be resolved), 3D simulations can only be performed using HPC resources. As an example, the simulation of a 1-GeV plasma accelerator stage using a fully explicit PIC code has been estimated to require 10,000 to 100,000 CPU hours for a single run.

The accelerator community has made significant advances in HPC since the start of the 1990s. For example, parallel electromagnetic modeling codes under a DOE/HENP Grand Challenge have been applied to several major accelerator projects, including PEP-II, NLC, ALS, SNS, and APT. In regard to NLC, simulations have been used to improve the design of accelerating cavities, resulting in a projected cost saving of order $100 million. In regard to multi-particle beam dynamics, parallel codes have been developed that have been applied to several projects involving high-intensity beams, such as SNS, APT, and the CERN SPL. And in regard to plasma accelerators, parallel codes have been developed that have been applied to experiments such as the SLAC E-157 experiment. Combined with theory and experiment, such a modeling capability

will help forge a path toward the control and staging of advanced accelerator sections. Figure 20 shows parallel simulation results from the three areas just described: electromagnetic systems simulation, beam systems simulation, and laser/plasma accelerators.

Figure 20. Parallel simulation results: (left) Beam halo studies using a parallel particle-in-cell code; (middle) Electromagnetic/thermal calculation of wall heating in a PEP-II cavity; (right) comparison of experiment (top) and simulation (bottom) for SLAC experiment E-157.
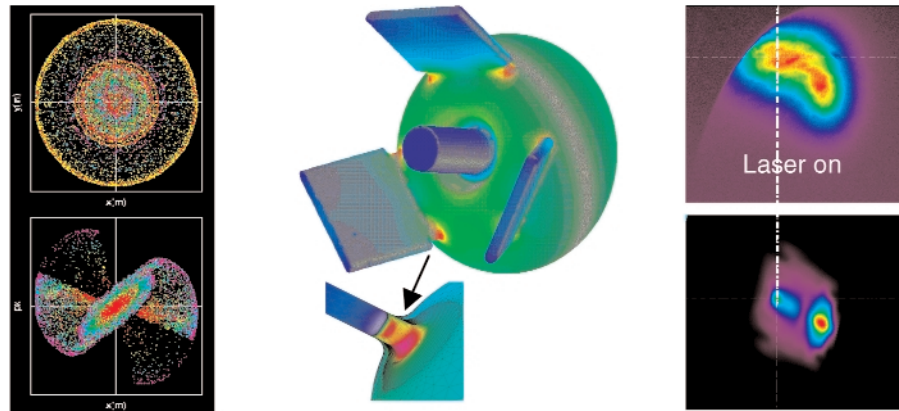
Although the accelerator community has made great progress in developing and using HPC, the earlier paragraphs of this section show that many challenges remain for which large-scale simulation is a necessity. Recognizing the challenges posed by these and other projects, a SciDAC (Scientific Discovery Through Advanced Computing) project on 21st Century Accelerator Simulation was approved in mid-2001. The primary objective of this national R&D effort is to establish a comprehensive terascale simulation capability for the U.S. Particle Accelerator Community. The success of this effort, which is supported by both HENP and ASCR, will involve close collaboration of accelerator physicists with applied mathematicians, numerical analysts, and computer scientists to develop new theoretical formulations and new algorithms capable of high performance and scalability on massively parallel systems. In particular, the accelerator community will utilize HPC tools for mesh generation, mesh refinement, particle/mesh methods, multi-level PDE solvers, eigensolvers, performance optimization, software component integration, and visualization. Many of these tools will be developed in the SciDAC Integrated Software Infrastructure Centers. Code verification and validation will require collaboration of code developers working with researchers performing controlled, well-instrumented experiments.

The accelerator community is well positioned to develop a comprehensive terascale capability that will utilize the latest advances in HPC technologies. Such a capability will help insure the success of future accelerators, by facilitating design decisions aimed at controlling and reducing cost, reducing risk, and optimizing

performance. The use of terascale simulation, combined with theory and experiment, will provide greater understanding of the complex, nonlinear, multi-scale, and many-body phenomena encountered at the frontier of accelerator technology.

Examples from Experimental Particle and Nuclear Physics

In the last few years NERSC, in cooperation with the HENP Office, has developed another style of production computing facility that is particularly suited to the embarrassingly parallel and data intensive applications of experimental high energy and nuclear physics. This is the PDSF facility, a loosely coupled cluster of Linux processors with a significant amount of disk space. This disk effectively serves the role of cache space between the processors and the archival storage facility (HPSS). At present this facility is actively used by about a dozen collaborations, with the STAR experiment at RHIC being the largest. Most of these physics experiments have adopted a grid-like computing model even before the term grid was fashionable and consequently the PDSF facility and the groups using it are the leading edge of the datagrid developments at NERSC. The role that PDSF plays for these collaborations ranges from the primary computing facility for some of the smaller collaborations to that of regional center and remote simulation and analysis facility for the larger ones. The combination of production computing infrastructure provided by NERSC, including the massive HPSS storage system and connectivity via ESnet, makes PDSF a unique and extremely effective shared facility for the computations of many HENP experiments.

These experiments are focused on an assortment of measurements that are helping to elucidate the nature and properties of matter and energy in a perfectly complementary fashion to the theoretical efforts that construct analytic and numerical models to describe, and occasionally predict, these fundamental properties of the physical world. A few of these experiments are described below.

The SNO experiment (Sudbury Neutrino Observatory) is observing neutrinos produced in the Sun as part of the effort to understand why the number of electron neutrinos observed on Earth are significantly less than the number generated by solar models. A recent result, in which data analysis at NERSC played a crucial role, is that electron neutrinos undergo a transformation to other types in the time it takes to reach the Earth, so the number of electron neutrinos reaching the earth is less than the number produced. This discovery

also is direct evidence that neutrinos have a rest mass, a significant puzzle in its own right.

One of the most prominent and yet puzzling features of the universe is that it is composed primarily of matter, and not anti-matter. It is thought that the resolution of this puzzle lies in some asymmetries in the fundamental properties of matter so that after so many billions of years we have one type and not the other. One of the experimental groups (Fermilab experiment E871) is carrying out data analysis at NERSC searching for asymmetries between matter and anti-matter observed in the decays of Lambda and Xi hyperons as well as charged kaons. Besides the exciting nature of the physics questions being addressed, this experiment has the distinction of having recorded the largest amount of data for any fixed target experiment in 1999 120 terabytes of data, and 231 billion events. Analysis of this data is ongoing at PDSF.

One of the early experiments involved at NERSC/PDSF is the AGS experiment E895 at Brookhaven National Laboratory. This relativistic heavy-ion experiment, which concluded its data taking in the late 1990s, is nearly finished with analysis of all the data taken and is unique in its nearly exclusive measurements of nuclear collisions in the energy range spanning the transition from nuclear and hadronic matter (2 GeV/nucleon) to the range where partonic properties of nucleons start to become noticeable (8+ GeV/nucleon). Results from this experiment include measurements of the size of the interaction region as determined from particle correlation studies as well as collective flow of various types of particles yielding information on the properties of hadronic matter. This has yielded constraints on the equation of state of nuclear and hadronic matter.

At present, the STAR experiment at the Relativistic Heavy Ion Collider (RHIC) accelerator at Brookhaven National Laboratory is the largest user (and contributor) to the PDSF facility. The first physics data was acquired by STAR in 2000 with colliding beams of gold nuclei at 65 GeV/nucleon ($\sqrt{s}_{nn}$ = 130 GeV) and in 2001 the full energy of $\sqrt{s}_{nn}$ = 200 GeV was achieved. The 2000 data served as something of a shakedown run for STAR, with about 1 million good events analyzed for physics results, and the 2001 run achieved about 10 times greater statistics. NERSC contributes to the overall computational process by providing resources for the majority of the simulation studies (model calculations as well as detector acceptance) and a large amount of final-stage physics analysis. At

time of this writing the analysis of the 2000 data is mature; many physics topics have published results, while the analysis of the 2001 data are just getting under way. A number of the published results from the 2000 data serve to illustrate the type of science to which NERSC is contributing. The overall goal of the RHIC heavy-ion program is the measurement and characterization of a new form of matter called the Quark Gluon Plasma (QGP), in which the hadronic structure of nuclear matter has been dissolved into a soup of quarks and gluons. Actually achieving the understanding of properties of the QGP will be derived from very many snapshots of individual measurements of the type that resulted from analysis of the 2000 data taken by STAR (described below).

The first physics results published by STAR were on the measurement of elliptical flow of charged particles at mid-rapidity. Elliptical flow is a characteristic of one of the matter properties and it was deduced from this measurement that there is a higher degree of thermalization in these collisions of gold nuclei at $\sqrt{s_{nn}} =$ 130 GeV than was expected based on measurements at lower energies. When the data were analyzed for flow of particles identified by species (π, K, p), even stronger confirming evidence was observed that hydrodynamic flow of matter is occurring in these collisions. Since there is no anti-matter in the initial state of the collisions of nuclei, a measurement of the anti-particles in the final state provides information on what occurs during the collision. It was observed that the number of anti-protons at mid-rapidity is about 2/3 of the number of protons indicating that even though most of the matter at mid-rapidity is created in the collision there is still a considerable excess of baryons over anti-baryons. Two-particle intensity interferometry has been used for years to infer the geometrical properties of the late stages of these systems. A measurement of two-pion correlations indicates that the size of the colliding system at these energies is similar to that at lower energies and is not significantly larger, as some models of a phase transition have suggested. One of the interesting things that can be observed while looking into the new energy density regime produced in these collisions is to simply count the number of particles produced. It was observed that the number of negatively charged hadrons varies significantly with transverse momentum as compared with that observed in proton—proton collisions, while the longitudinal distribution is quite constant in the range $|\eta| < 1$. It is also observed that the number of anti-protons produced is approximately proportional to the negative hadron production (Figure 21).
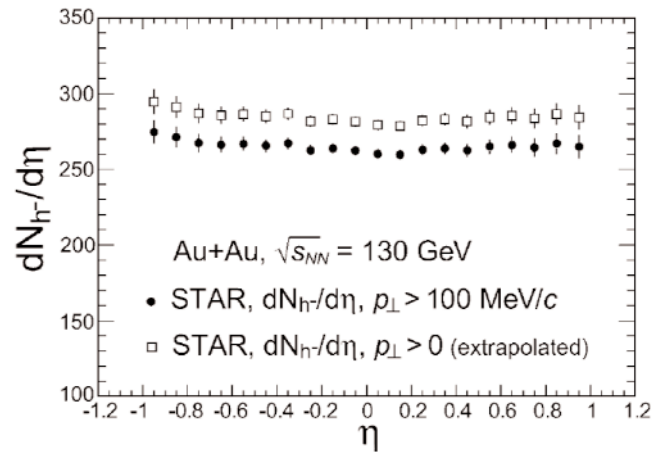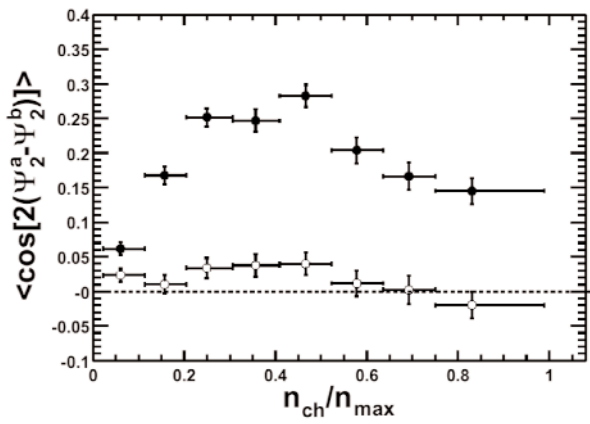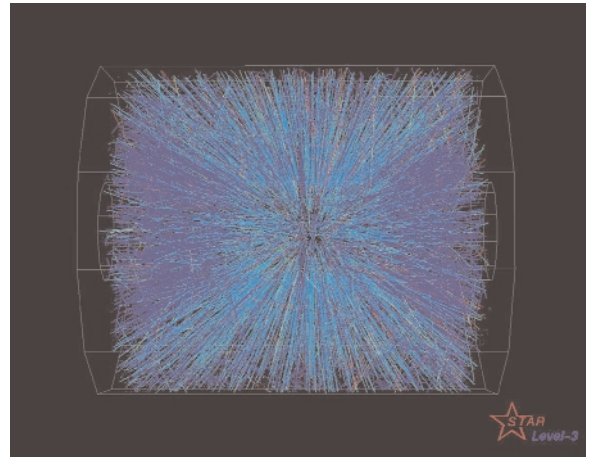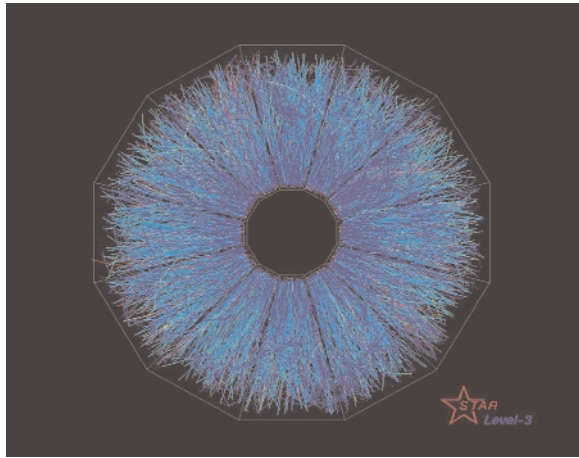
Figure 21.

(Left top), end view of a collision of two gold nuclei at $\sqrt{s_{nn}}$ = 130 GeV.

Bottom, a measurement of azimuthal asymmetry in events like that shown above.
The measured asymmetry indicates a hydrodynamic flow of matter in these collisions.
(Right top), side view of a collision of two gold nuclei at $\sqrt{s_{nn}}$ = 130 GeV.

Bottom, a measurement of the longitudinal distribution of negatively charged hadron,
indicating that the pseudorapidity distribution is relatively constant.

Office of Advanced Scientific Computing Research

The Office of Advanced Scientific Computing Research (OASCR) is composed of three divisions:

- **Mathematical, Information, and Computational Sciences Division**

- **Technology Research Division**

- **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 the Department of Energy. 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 Plans 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.

The program of OASCR provides enabling technology to the NERSC user-community to allow the types of simulation conducted elsewhere in this report. Thus the focus ranges from new mathematical and numerical methods to issues of software libraries,

innovative uses of resources (both hardware and software) such as grid-computing, and new software tools for use on current and future hardware environments. Below we provide some illustrations of the approaches that are being developed and their utility. *The Aggregate Remote Memory Copy Interface (ARMCI) library* provides a library for managing arrays in a distributed parallel computing environment. This library is currently being applied in Office of Science—supported simulation studies, particularly, although not exclusively, in computational chemistry simulations. *The Advanced Computational Testing and Simulation (ACTS) toolkit* contains several numerical libraries for use in parallel computation. These libraries address the areas of linear algebra, grid generation, and nonlinear iterative techniques and provide the user with a portable software capability in conducting simulations. The individual components of the *ACTS toolkit* are continually being upgraded responding to new user requirements and changes in hardware environments. In terms of future directions to enable users to invoke combinations of simulation software, *the Common Components Architecture (CCA) library* is being developed for use, in particular, in the SciDAC program. This library is part of the Center for Component Technology for Terascale Simulation Software which is one of the SciDAC Integrated Software Infrastructure Centers (ISIC). Finally, in the area of grid-computation, one of the SciDAC National Collaboratories and High-Performance Networking projects is the DOE Science Grid: Enabling and Deploying SciDAC Collaboratory Software Environment . This collaboratory involves the development of a grid-computing capability in the evolving *DOE Science Grid* and thus enables distributed internet-based computation. Finally, a demonstration of advanced computational modeling in *Combustion* demonstrates the use of advanced mathematical and numerical modeling techniques when applied to combustion modeling.

Development of ARMCI: A Portable Communication Run-time System for Advanced Libraries and Programming Models

Evolution in programming methodologies for massively parallel computers has been driven by advances in computer technology and increased performance requirements of scalable applications. Specifically, attempts to address the growing gap between raw microprocessor power and memory performance became the leading source of complexity in managing distributed data in scientific applications. A hybrid OpenMP/MPI model is an example approach to manage memory hierarchy in clustered SMP systems and can lead to substantial increase of complexity in the programming model.

The tradeoffs between performance and complexity require development of novel programming methodologies, tools, and frameworks.

ARMCI was created to support development efforts in the area of management of distributed data structures and application domain programming frameworks. Such efforts require efficient and flexible mechanisms for accessing remote data structures, for example, remote memory copy and distributed synchronization and mutual exclusion operations. As a portable run-time system, ARMCI aims to exploit platform-specific resources and interfaces to provide the highest possible performance of remote memory operations that are needed by higher/user level libraries and tools. To deliver performance to the applications, ARMCI advanced interfaces (multi-strided and generalized *iovec* I/O formats) have been designed to exploit the information about the data layout and structures at the lowest level. They allow avoidance of performance-degrading factors such as redundant memory copies and data reorganization that otherwise would be involved in the communication.

ARMCI offers general-purpose, portable, and efficient remote memory copy operations (one-sided communication) optimized for noncontiguous (striped, scatter/gather, I/O vector) data transfers. Atomic and mutual exclusion operations are available to support management of shared distributed data structures. Development of this system has been driven by the need to support the global-address space communication model in the context of distributed regular or irregular distributed data structures including multidimensional arrays, communication libraries, and compiler run-time systems. Its applications range from the Global Arrays shared distributed memory programming toolkit, the Generalized Portable SHMEM (GPSHMEM) library, the Adlib compiler run-time system, and direct use in scientific applications such as DDI layer of the GAMESS chemistry application. At NERSC, a major user of ARMCI is NWChem (NorthWest Chemistry), a large (1 million lines of code) package developed on top of Global Arrays.

ARMCI is available across a wide range of computer platforms including massively parallel processors such as the IBM SP, Cray T3E, Fujitsu VPP, to symmetric multiprocessor servers from HP, IBM, SGI, Sun, and clusters of workstations. It exploits a range of vendor-specific communication capabilities, including the LAPI library on the IBM SP, VIA on the Giganet/Emulex clusters, GM on the Myrinet clusters. ARMCI tries to combine strengths of multiple protocols and techniques to maximize performance. For example, on

the IBM SP ARMCI uses: LAPI active messages, threads, LAPI remote memory copy, and shared memory to deliver better performance than the IBM LAPI communication library (Figure 22).
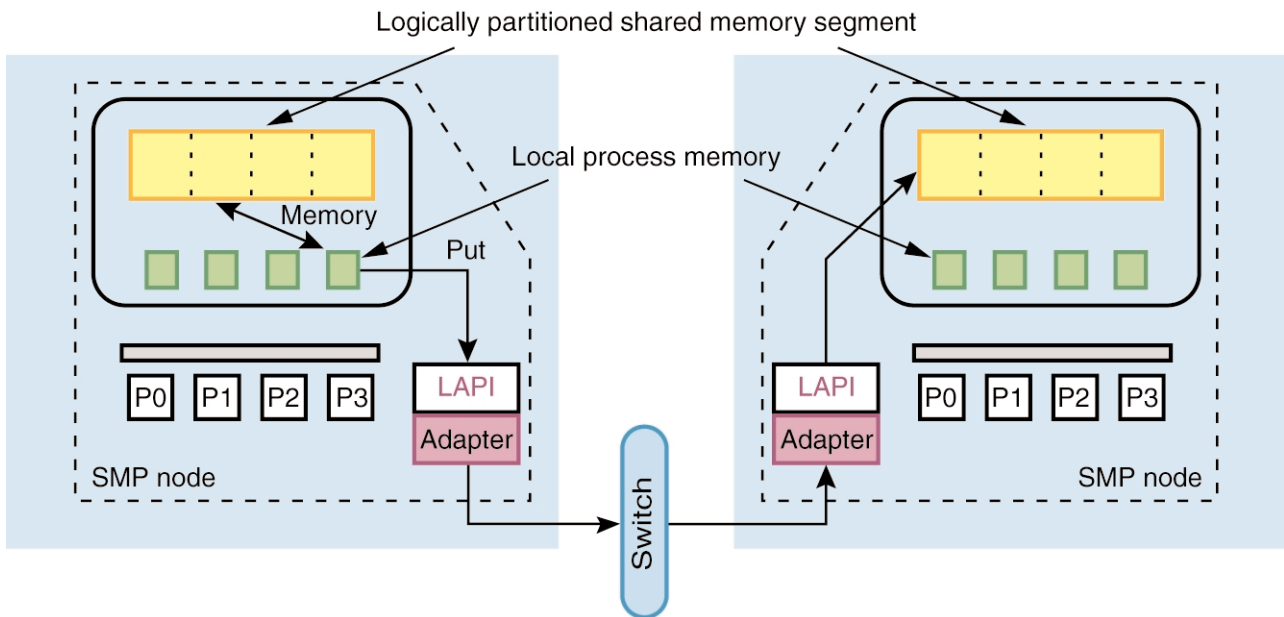
Figure 22. ARMCI combines the strengths of multiple protocols and techniques to maximize performance.
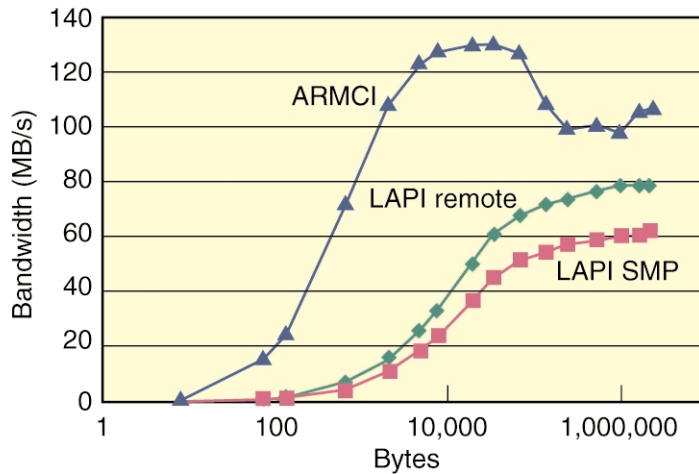
An important feature of multiprotocols in ARMCI is that they are not exposed to the applications and thus they do not lead to increased complexity of the programming model.

The Advanced Computational Testing and Simulation Toolkit

The Advanced Computational Testing and Simulation (ACTS) Toolkit comprises a set of software tools that make it easier for programmers to write high-performance scientific applications for

today's computers. The tools were developed at DOE Labs and universities, and are mostly libraries (Fortran or C libraries, or C++ class libraries). They were primarily designed to provide code portability and high performance on existing and emerging computer platforms.

Based on their functionality, the tools fall into four broad categories: support for numerical algorithms, support for code development, support for code execution, and support for library development. The numerical tools implement numerical methods and algorithms, and include sparse linear system solvers, ODE solvers, Optimization Solvers, etc. The ones under the support for code development category provide infrastructure that manages some of the complexity of distributed programming (such as distributing arrays, communicating boundary information, etc.) but do not actually implement numerical methods. Execution support is a catch-all category for application-level tools; these tools include performance analysis and remote visualization support. Library support tools provide an infrastructure for tool developers and probably will not be used or seen directly in scientific applications.

ACTS tools fall into a category the National Energy Research Scientific Computing center (NERSC) labels experimental software (ES). This is software from the research community that is not supported by vendors. NERSC has developed policies for supporting ES, but because of issues relating to long-term tool maintenance and development, ES support is generally limited in scope. The ACTS project has substantially increased ES support at NERSC.

Making ACTS tools more effective in solving DOE scientific problems matches NERSC's goals of making top-quality computing services available to its users. The unique scenario provided under the ACTS project facilitates a number of opportunities to take a leap forward in the use and acceptance of ES. The ACTS project brings together software development from DOE laboratories, in some cases in collaboration with universities, to the hands of a scientific community that is broader than the community from a single DOE laboratory. Consequently, ACTS has promoted reusability over duplication of efforts and tool interoperability over hard-wired and intrusive library interfaces that serve only a specific project.
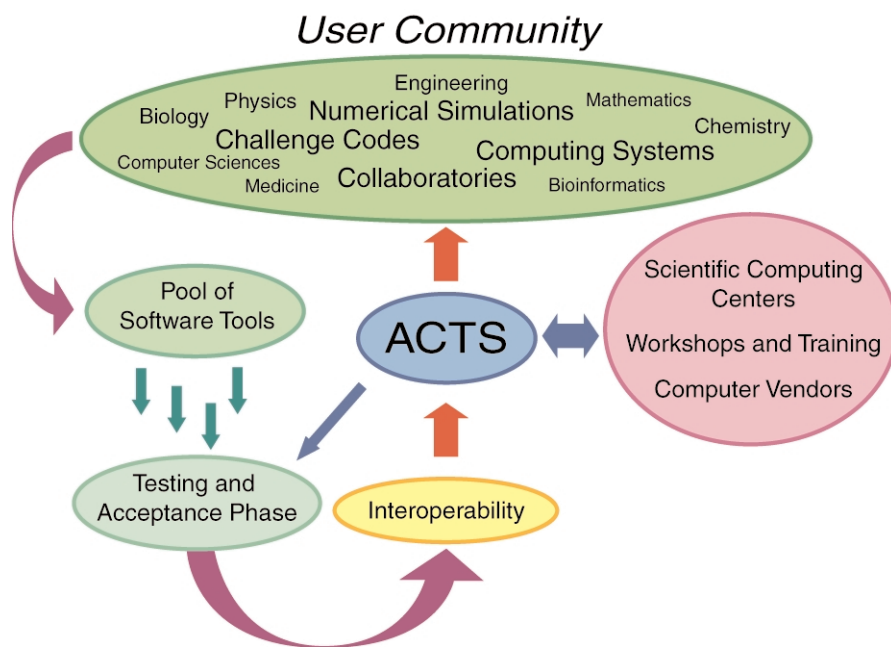
All the potential provided by the tools cannot be achieved, however, if the tools are not used effectively or not used at all. Therefore, it is necessary to continue working on improving the usability, accessi-

bility, and acceptance of ACTS tools. This goal can be achieved by means of technical support (by providing specialized consulting services for new and advanced users, providing documentation and training for ACTS tools, identifying potential users, and promoting ACTS tools where appropriate), an independent evaluation of the tools (addressing the performance, robustness, functionality, and overall usability of ACTS tools, so that users can make informed choices and developers can improve the tools if necessary) and increased availability and stability of current versions of ACTS tools on NERSC platforms. Additional information about the status of the tools can be found at http://acts.nersc.gov.

The primary beneficiaries of ACTS tools are developers of scientific applications. Many areas of scientific computing are covered by ACTS Toolkit components, and can potentially make use of them. Furthermore, the ACTS Toolkit by its evolutionary nature will continue to embark on emerging areas of computational science. Developing software for large-scale systems and obtaining high performance is inherently more complex than the development of serial software and significantly more expensive to implement. In the past, many smaller projects have been unable to access the benefits of parallel computing because they lack the resources to port their computational codes from serial to parallel implementations. Moreover, as the problems being solved become more complex, the challenges in building parallel programs become even greater. The goal of the ACTS is to alleviate this situation, and DOE has made a strong commitment to the success of this project.

ACTS is already bringing tools to higher acceptance levels among computational scientists and institutions and providing the necessary means for the individual tool projects to interact with more users and gradually mature, becoming more robust and portable to state-of-the-art high-performance computing environments. As depicted in Figure 23, the ACTS Toolkit project can serve as a main resource for the distribution of quality software to a broad computational community. ACTS will continue to leverage the acceptance of experimental software inside the scientific computing community and promote their fast integration into the computer-vendors-supported software libraries. ACTS is exploring mechanisms for component software, library, and language interoperability to facilitate the use of existing software technologies with future ones. In addition, ACTS uses workshops, conferences, and training sessions for the dissemination of knowledge and technical expertise. These activities not only increase the visibility of the ACTS tools but also bring new users and applications that further enrich toolkit's and tool developers' expertise.

New Approaches to the Development of High-Performance Simulation Software

As the capabilities of high-performance computers increase, expectations of the speed, fidelity, and complexity of the computational simulations that run on them also increase. This is reflected in the increasing size and complexity of state-of-the-art simulation software. For example, the NWChem parallel computational chemistry package, the development of which is lead by PNNL, is now seven years old, is approaching one million lines of source code, and represents an investment of more than 100 person-years of effort by more than 40 people. Software of this scale, combined with the pace of evolution dictated by both the hardware environment and scientific considerations, can present a serious challenge to the development team that must be handled with extreme care.

Increasing software complexity is not just a concern for high-performance computing but is shared with commodity desktop computing as well. In this area the adoption of component architectures has gone a long way toward managing this complexity. The Center for Component Technology in Terascale Simulation Software, which includes participants from six national laboratories and two universities, is founded on the idea that component architectures for HPC will do for high-fidelity simulations what they have done for the commodity world. The CCTTSS goal is to investigate component models appropriate to HP parallel computing and to integrate existing commodity component architectures (e.g.,

JavaBeans, CCM) into the high-performance arena. Although the CCTTSS recognizes that special component mechanisms are required for parallel processes and data structures, it is viewed as an extension or augmentation to existing component models. In the CCTTSS's view, the only way that HP simulations can move beyond the few-person-per-code kitchen operation that they are today is to adopt a component architecture for parallel HPC, and leverage what is being done in the commodity component world.

Using a component technology such as CCA, applications will be composed from components of varying granularity and loaded into a framework for execution. A rich library of components will be available from many developers, providing familiar capabilities such as numerical libraries, distributed computing environments, or solvers in a fully inter-operable language-independent form. The component environment also makes it easier to develop some capabilities that are not widely used at present. This includes support for parallel data redistribution between parallel components, or run-time computational visualization and steering. A large repository of plug and play components provides the building blocks for the creation of large-scale simulation applications, allowing the application developer to focus on the components corresponding to their scientific domains. The developers are able to reliably connect their components together with components from other domains created by experts in those areas, as needed to solve a larger problem.

One example scenario involves a group of combustion scientists who plan to develop a computational facility for parallel distributed-memory chemically reacting flow simulations. As is commonly the case in this computational domain, their current approach uses an explicit time-marching algorithm. Implicit methods enable the use of longer time steps and therefore can achieve faster solutions. However, implicit methods are not often used in combustion simulations because they require advanced parallel linear algebra capabilities that are beyond the scope of the combustion scientists primary interests and expertise. In order to introduce an existing linear equation solver component to deal with thin boundary layers, the scientists decide to reformulate their existing simulation code as a collection of components. Their principal task in restructuring their application is identifying the interfaces that each component will provide and use, where an interface is a set of function calls that together encapsulate the functionality of the software. The researchers decide that their application can be split into two components: a multi-scale time integrator and a combustion physics component. In addition, they decide that the linear solver will

interface with the combustion physics component. They choose to use a solver interface defined by the Equation Solver Interface (ESI) Forum, whose solvers have been encapsulated as CCA components. By using ESI components, the combustion scientists will be able to experiment easily with multiple solution strategies and upgrade to new solver approaches, as they become available.

Since CCA will, by design, impose a minimal performance cost on applications, the use of component-based software development techniques should not directly impact the computational require-ments of existing Office of Science applications. However, we can anticipate two types of indirect impacts. First, CCA is intended to facilitate more rapid development of even larger and more complex simulation applications that would be considered impractical to develop without component technologies. Clearly, such a trend would accelerate the growth of computational demands in an environment like NERSC. Second, the concepts of community and communication become more important. An HPC center, like NERSC, would want to provide its users with the richest possible set of components for their use, which implies the need to be strongly connected to distributed component repositories, from which users would load components to compose their applications. Insofar as HPC centers are a nexus for large numbers of users with common interests, they can play an significant role in encouraging the creation of communities for the development of components in a given domain, and encouraging users to make new components available (Figure 24).
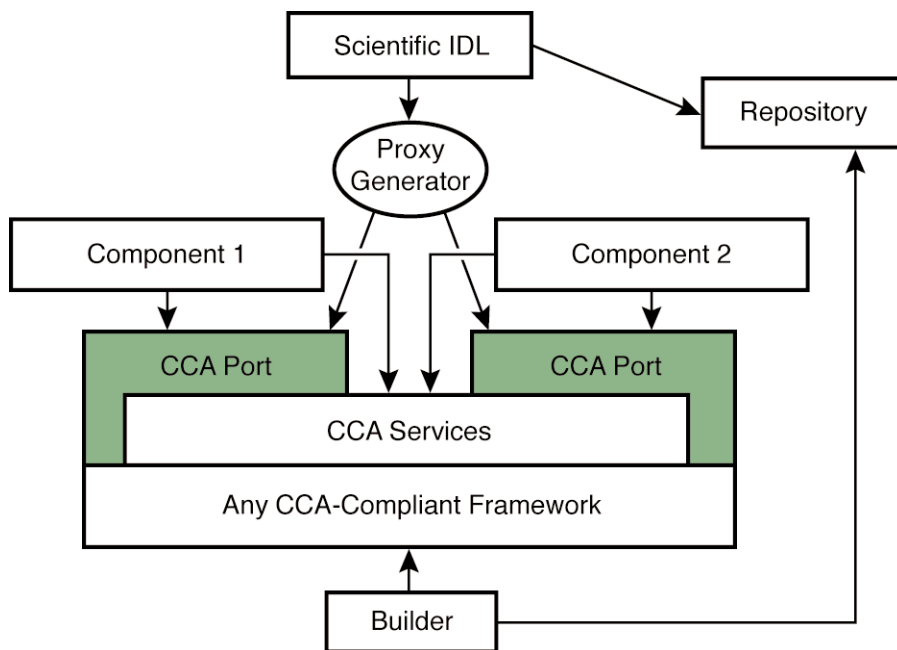
Figure 24.
Using a component technology such as CCA, applications will be composed from components of varying granularity and loaded into a framework for execution.

Development of DOE Science Grid

DOE Office of Science laboratories operate a wide range of unique resources, from light sources to supercomputers and petabyte storage systems, that are intended to be used by a large distributed user community. The laboratories geographically distributed staff are frequently faced with scientific and engineering problems of great complexity, the solution of which requires the creation and effective operation of large multidisciplinary teams. The problems to be addressed are large and highly challenging, often greatly exceeding the limits of traditional computing and information systems approaches.

Recognizing these challenges, the SciDAC call for proposals in National Collaboratories and High-Performance Networks, echoing the document Scientific Discovery through Advanced Computing, calls for the creation of a *Collaboratory Software Environment* to enable geographically separated scientists to effectively work together as a team and to facilitate remote access to both facilities and data.

A Collaboratory Software Environment must, by definition, span multiple institutions and link numerous, geographically distributed resources of different types. These characteristics introduce two unique challenges that have, until recently, made the development of collaboratory applications very difficult. First, the diverse resources involved typically feature vendor- and site-specific software and management mechanisms; writing code that can deal with the variety of different mechanisms that arise in even a small collaboratory can become prohibitively difficult. Second, the sheer scale of the resources involved, and the fact that they span multiple administrative domains, make such fundamental issues as resource discovery, authentication, and fault detection very difficult often beyond the abilities of a typical collaboratory application developer.

A new class of distributed infrastructure called Grids address these two challenges directly by (a) defining and deploying standard protocols and services that provide a uniform look and feel for a wide variety of computing and data resources, and (b) providing global services that provide essential resource discovery, authentication, and fault-detection capabilities. Together, these capabilities reduce the barriers to the large-scale coordinated sharing and use of distributed resources that is at the heart of collaboratory applications. Grid capabilities are thus fundamental to the efficient construction, management, and use of widely distributed application systems; human collaboration and remote access to, and operation of, scientific and engineering instrumentation systems; and the management and

securing of the computing and data infrastructure. For these reasons, emerging Grid technologies have been adopted by major collaboratory and supercomputer center access projects worldwide, such as the NSF National Technology Grid, NASA's Information Power Grid, the European Union's Data Grid, the NSF's Network for Earthquake Engineering Simulation Grid, and the DOE ASCI DISCOM project. The Global Grid Forum (www.gridforum.org) provides an international coordinating and standards definition body, analogous to the Internet's IETF.

The creation of a DOE-wide distributed computing infrastructure, or *DOE Science Grid* represents an initiative of similar ambition and, we believe, opportunity for impact (Figure 25). The DOE Science Grid is designed to reduce or eliminate barriers to the coordinated use of DOE resources, regardless of the physical location of those resources and the users that access them. For the first time, we are providing a persistent and supported set of Grid infrastructure and deployable services in the DOE community. This infrastructure, as we explain below, will provide a range of new Grid Services addressing issues of resource discovery, security, resource management, instrumentation, and the like. These services will in turn be used to create a range of innovative Grid tools targeting specific application classes. Our goal in creating this infrastructure and tools will be to enable innovative approaches to scientific computing based on such concepts as secure remote access to online facilities, distance collaboration, shared petabyte datasets, and large-scale distributed computation; the eventual outcome should be revolutionary changes in a wide range of scientific disciplines across DOE.
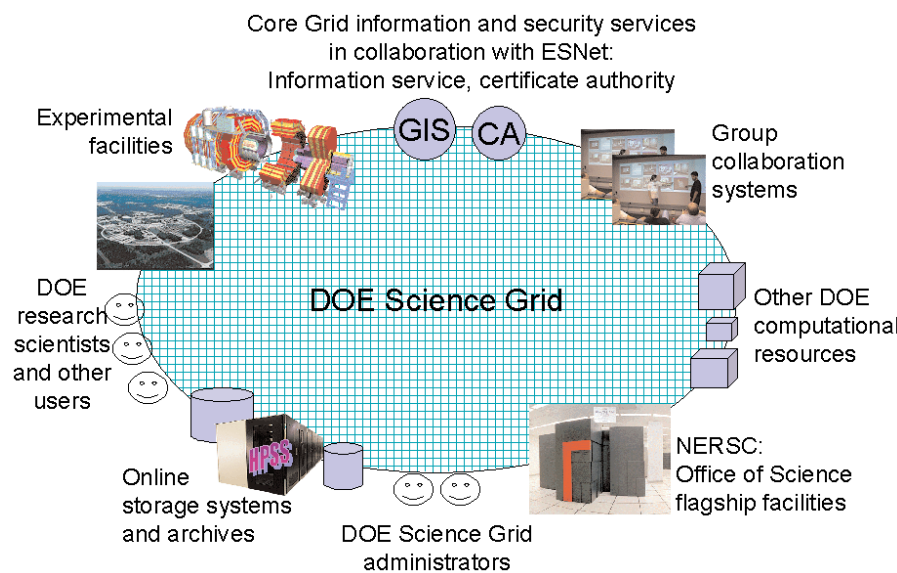
Figure 25.
The DOE Science Grid is designed to reduce or eliminate barriers to the coordinated use of DOE resources.

Combustion Simulation

Most of the energy required for transportation and for stationary power generation is obtained by combusting fossil fuels. In spite of its fundamental scientific and technological importance our knowledge of the combustion process is incomplete. Recent progress in simulation methodologies and new, high-performance parallel architectures have made it is possible to perform detailed simulations of multidimensional combustion phenomena using comprehensive kinetics mechanisms. Our goal is to model turbulent-chemistry interactions with sufficient fidelity to predict not only the basic energetics but also the more detailed aspects of the chemistry such as the formation and destruction of pollutants.

These efforts build on major successes of the DOE research program in scientific computing over the last two decades on the development of highly accurate numerical methods for solving fluid mechanics problems. These computational technologies include high-resolution finite-difference methods for problems such as compressible and incompressible flow. These numerical methods have been coupled to a variety of adaptive techniques that focus computational effort where it is required. These adaptive techniques include local adaptive mesh refinement, in which regions of the finite-difference grid are automatically refined or coarsened as the computation progresses, and front tracking, in which specialized algorithms are used to explicitly model the propagation of thin fronts, such as shock waves, flame fronts, or material interfaces, through the computational domain.

This advanced computational methodology is currently being applied to analyze the behavior of a pre-mixed turbulent methane flame in three dimensions. The simulations use an adaptive time-dependent low-Mach-number combustion algorithm based on a second-order projection formulation that conserves both species mass and total enthalpy. The species and enthalpy equations are treated using an operator-split approach that incorporates stiff integration techniques for modeling detailed chemical kinetics. The methodology also incorporates a mixture model for differ-ential diffusion. This methodology is parallelized using the BoxLib framework that provides support for hierarchical adaptive methods. The framework performs the data distribution and dynamical load balancing as grids in the adaptive hierarchy are created and destroyed.

We are currently modeling methane chemistry and transport using the DRM-19 (19-species, 84-reaction) mechanism derived from the GRIMech-1.2 along with its associated thermodynamics and transport databases. We have simulated a lean flame with equivalence ratio 0.8. The simulations allow us to examine the basic structure of the flame, including turbulent flame speed and flame surface area. We are also investigating the effect of turbulent flame interaction on the flame chemistry. Figure 26 depicts the flame surface during the simulation. The computation requires approximately 32 gigabytes of memory. Each time step requires approximately 2.5 hours of wall clock time using 128 CPUs on the IBM SP at NERSC and generates 4 gigabytes of output for analysis. Thus, the overall computation which requires 400 time steps will require approximately 125,000 CPU hours and generate over a terabyte of data.

The current model uses a relatively simple representation of methane; more comprehensive models, particularly those that include pollutant chemistry are 4–5 times larger. Furthermore, the physical domains we are able to simulate with a high degree of chemical fidelity are severely limited. Terascale computing will enable us to relax the restrictions on our current models. With our current algorithmic technology we are well positioned to exploit terascale computing resources.

Figure 26. Combustion simulation of turbulent flame, showing surface of maximum heat release.

Computational Resources for the Office of Science

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

Hardware

Over the past five years the NERSC HPCF has transitioned from providing the majority of computational resources via parallel vector (PVP) supercomputers to providing over 95% of computational resources via massively parallel processing (MPP) systems. The HPCF also provides several cluster computing environments, and specialized compute servers for visualization and mathematical software.

NERSC supports a diverse workload. The hardware must support capability computing, which requires a large fraction of an entire computing resource, to a much smaller amount of related capacity computing, comparable to running on a high-end desktop system.

Massively Parallel Supercomputers

Distributed-memory Massively Parallel Processor (MPP) systems have become increasingly dominant in supercomputing over the last five years. For suitable problems, these systems are capable of much higher computational performance, and of addressing capability computing problems that cannot be addressed with other hardware. NERSC has two MPP systems in production. The IBM SP (*seaborg.nersc.gov*) is one of the largest configurations available anywhere. The system, installed in September 2001, is composed of 184 compute nodes (2944 processors), with over 4 TB of memory and 20 TB of disk. This system has a peak performance of 1.5 Gflops per processor for a theoretical peak of nearly 4.5 teraflops.

Also installed is a Cray T3E (*mcurie.nersc.gov*) with 644 processors, 165 GB of memory, and 1.5 TB of disk. This system has a peak performance of slightly less than 0.5 teraflops.

Cluster and Grid Computing

In the area of cluster computing, NERSC supports 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 (HEP) and Nuclear Science (NS) investigations. PDSF is composed of some 150 compute nodes with 281 Pentium II or II processors. In addition, the NERSC PC Cluster Project runs two clusters for software development and testing: a 36-node Pentium II cluster connected by switched fast Ethernet, and a 12-node Compaq Alpha EV6 cluster connected by Myrinet and Gigabit Ethernet.

With the advent of Grids infrastructure for using and managing widely distributed computing and data resources in the science environment there is now an opportunity to provide a standard large-scale science environment that spans many different projects and provides the infrastructure and computing services needed in a relatively uniform and supportable way.

Parallel Vector Supercomputers

In the arena of traditional vector supercomputing, NERSC supports three Cray SV-1 Parallel Vector Processor (PVP) systems with a total of 64 processors. The cluster has a total of 32 GB of memory and approximately 1 TB of disk. With the peak performance of each SV-1 processor at 1.2 Gflops, the peak total performance of the system is just over 75 Gflops.

High-Performance Storage System

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, archive, and system backups. The current configuration consists of over 1.3 PB tape capacity serviced by 8 STK tape robots, 10 TB of disk cache, and multiple data-movers and meta-data servers.

Other Servers

NERSC provides two specialized servers; an 8 processor SGI Onyx 2 (*escher.nersc.gov*) for visualization needs and a 2 processor Sun UltraSparc II (*newton.nersc.gov*) for symbolic manipulation software needs. In the case of the visualization server, many common visualization tools are installed, and it is connected to the NERSC Visualization Group's ImmersaDesk.

Software

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

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 jobs interactively and in a batch system, import and export data, compile and develop programs effectively, and support significant third-party applications.

Standard programming languages, debugging and performance tools are an essential component of any high-performance computing resource. The following languages need to be supported as part of a fully integrated, interoperable environment for compilation and debugging: Fortran 90, C, C++, and assembler. Tools to monitor and diagnose performance problems are also required.

Basic Scientific Libraries and Environments

Scientific and mathematical libraries are crucial to portable and efficient software development. NERSC supports most standard serial and parallel libraries such as: BLAS, LAPACK, and ScaLAPACK.

In addition, communication libraries, such as MPI, should be efficient and amenable to debugging and performance tuning.

Finally, various I/O libraries such as netCDF, HDF, etc., should be installed and supported to facilitate code portability and data exchange.

Scientific Application Codes and Systems

NERSC supports a wide variety of third-party software applications and libraries from the independent software vendor community: for example, the computational chemistry programs Gaussian and NWChem, and the mathematical libraries NAG and IMSL. The

continued support of these applications and the staff to maintain and provide consulting services is required to allow the NERSC user community to conduct their science.

Human Resources

Providing effective, high-quality client support is necessary for users to accomplish scientific work. NERSC must assist users in being as productive as possible by providing systems, tools, information, consulting services and training.

User Support Services

The NERSC service architecture aims to ensure a response to client problems within four working hours, and timely resolution for 95% of problems within two working days. In addition to basic Help Desk services, NERSC offers direct scientific support in order to help the DOE scientific community become more productive in their computational and data management work. Such support includes such things as assisting clients in using the systems in the most effective and efficient manner, solving client application problems, and supporting scientific application development tools and utilities. NERSC provides advanced training and client instruction in the use of the latest technology. NERSC staff develops skills in new areas, and shares them with clients by creating, updating, and presenting all the relevant training information related to NERSC systems.

NERSC provides user account initialization and security and tools for Principal Investigators and account managers to allocate their available resources, query the current state of their resource usage, and plan their own resource utilization.

Specific Recommendations

Expand the high-performance computing resources available at NERSC. The upcoming procurement must ensure an appropriate balance of computational equipment to support the wide range of large-scale applications involving production computing and development activities.

Encourage the continued improvement of algorithms, software (including middleware), and database technology for improved performance on parallel platforms. More pro-active support for implementation of NERSC (and other DOE) software tools into Office of Science application codes may be necessary.

Strengthen the networks. Reliable, robust and fast access to the computational resources is essential. This is especially true for increased use of Grid Computing.

Coordinate NERSC activities with SciDAC, ISIC, and other Office of Science compute centers.

Identify mechanisms to provide larger-scale production facilities. While model development and physics/chemistry/biological improvements will always be moving forward, many Office of Science applications are now ready for large-scale production computing on NERSC parallel platforms. Future simulations will involve long-running simulations requiring hundreds of hours of continuous simulation time. Current computer system formulations do not lend themselves to these types of simulation; this needs to be addressed.

Appendix A:
Computational Platform Specific Needs and Recommendations

Software and Algorithms

As our machines evolve, so to must the algorithmic developments that utilize a given platform. Codes must be optimized for efficient use of the new computer architecture. In nearly every upgrade of the NERSC facility, we have seen a need to partially reconfigure codes in order to maintain performance standards. This trend will likely continue into the future. It is therefore essential that NERSC maintain a quality team of consultants who can help the user community with obtaining science from various codes in an efficient manner when the new computing facility becomes operational.

Algorithmic requirements are substantial across the whole of the Department of Energy and some of these needs are being addressed by components of the SciDAC program. Those not addressed by SciDAC remain the purview of the respective offices within the Office of Science. We note here only that the combined efforts of programs like SciDAC and the combination of teams composed of physical science specialists, computer scientists, and mathematicians are needed to make real progress in each of the physical science areas. A large part of the computational community is involved in developing and exploring new and improved numerical algorithms to solve classes of science problems. These endeavors often require several years of intense work before large dividends become visible. Toward this end, the consulting services at NERSC are essential to the efficient use of the massively parallel machines and toward improving the algorithms used.

Continued support of and improvement in compiler technology is very important to both scientific efforts and the NERSC mission. Given the large financial investment that NERSC is making in hardware upgrades it is vital that the compiler technology be able to support high-usage codes at the level of 10—25% of the theoretical peak speed of the machine. Anything less is wasted money. By investing in performance support personnel, NERSC will continue to play a major role in helping users achieve good performance on the new platform. High performance means that more science and engineering are performed per dollar and more cycles are available for NERSC users.

Major developments in physical modeling continue to push technology. This is true in all areas of research mentioned in this Greenbook. Whether we are discussing new quantum Monte Carlo algorithms to alleviate the Monte Carlo sign problem, or multibeam models in intense ion beam physics, or an inclusive coding environment such as NWChem, many scientists are involved in theoretical developments that push our computational hardware. This trend will continue indefinitely and requires us to think in various ways concerning machine optimization.

Hardware

A common theme that runs through BER, BES, FES, and HENP is the need for better performance of memory bandwidth. The processing power has increased phenomenally over the last few years but access to data in memory has not kept up. Hence it has become increasingly difficult to utilize the exceptional speeds of the compute CPUs on a given platform.

Recent hardware offerings that have come from the vendor community have been following this somewhat disturbing trend. Although the processing power follows Moore's law (i.e., a doubling of processor power every 18—24 months), the other subsystems of the parallel supercomputer do not. So, while the overall performance of application code suites and algorithms has improved with each generation of computational resources, the codes have shown that the bottleneck in achieving the best performance is, generally, the latency of the memory and communication subsystems of the parallel supercomputer. With the trend toward multiple processors per node (SMPs), these bottlenecks are even more severe. Therefore, scientists need systems that are well balanced in all respects processing power, memory speed and latency, communication speed and latency, and disk speed and latency.

External to the processors, all computer designers must confront the increasing gulf between the speed of the processors and that of the memory subsystem. Since the early days of Cray Research, Inc., it has been recognized that the performance of many scientific applications is determined by the performance of the memory subsystem as much as by the speed of the floating-point unit. In fact, beginning with the Cray X-MP, the balance between memory bandwidth (B, in bytes/sec) and floating point performance (F, in flop/sec) became an explicit design element in Cray supercomputers. The later Cray models boasted exemplary memory bandwidths. For example, in the

Cray T-90, the last vector supercomputer from CRI, the B/F ratio was over 15! This is to be compared to a B/F ratio of 0.67—1.33 for Nighthawk-II nodes used in the ASCI White system and 0.98 for the ES40 nodes used in the AlphaServer SC systems. For the IBM machine, the lower figure corresponds to the case where all of the processors are simultaneously requesting data from memory, whereas the higher figure refers to the bandwidth when only one processor is requesting data. To put this in context, for a DAXPY operation, $y_i = ax_i + y_i$, a common arithmetic operation in scientific codes, to execute at full speed out of memory requires that B/F > 12.

Modern microprocessors attempt to solve the memory performance problem by using caches. Caches are built from fast, but expensive memory. Both the IBM Nighthawk-II and Compaq ES40 nodes, for example, have an 8-Mbyte L2 cache for each microprocessor. The B/F ratio for the IBM L2 cache is 5.3, an improvement of as much as a factor of 8 over retrieving data from main memory. Clearly, if an algorithm can be structured so that the data are in cache when needed, the performance of the microprocessor will improve substantially. For a small number of scientific applications, it may be possible to keep the data in cache. For others, the key to a successful cache strategy is data reuse. In addition, high efficiency can be achieved if sufficient operations are performed for each memory reference, e.g., evaluation of an exponential or square root.

A good cache helps on certain sizes of problems, but for the larger problems that may be investigated in the near future, most of the memory accesses miss the cache. As an example, on the SP processors at SDSC a conjugate gradient solution on last year's problem size runs at 140—170 Mflops per processor, but the same solution on a larger lattice size runs at only 70 Mflops/processor since that problem no longer obtains a good fraction of cache hits. This problem is endemic to many applications and therefore should be a major performance issue as NERSC begins the next-generation procurement process.

The trend toward clusters of SMP nodes has mixed implications for the performance of applications on parallel supercomputers. Larger SMP nodes allow more expensive, and hence higher-performance, network cards or interconnects to be used since their cost can be amortized over all of the processors in the node. And, as long as the processes on a node do not saturate the network interface, a performance gain will be realized. However, communication-intensive applications will fare worse on parallel supercomputers based on large SMP clusters. Consider, for example, the latency and

bandwidth of the interconnect on the IBM SP series. The IBM SP-1 switch had a bandwidth of just 8 Mbytes/s and a message-passing latency of about 80 microsecs. The switch on the next generation of parallel supercomputer from IBM, the SP-2, delivered about 80 Mbytes/s with a latency of 30 microsecs. The most recent switch, the Colony switch in ASCI White, delivers a uni-directional bandwidth of 500 Mbytes/s (twice this in two directions) and about 15 microsecs latency. So bandwidth and latency have improved dramatically. However, over the same period of time, processors have also been getting much faster and, now, with the move to SMP clusters, more processors are being bundled into a node.

Above we used the ratio of memory bandwidth to floating-point speed (B/F) to characterize the balance in computational nodes. The same may be used for the network, denoted as netB/F. The five-year-old IBM SP-2 at PNNL/EMSL has a single processor node running at 120 MHz with a peak performance of 480 Mflops. Thus, the netB/F ratio is 0.17. The 16-processor 375-MHz Nighthawk-II nodes have a peak performance of 24 Gflops, which with the Colony switch yields a netB/F of 0.02, assuming that all 16 processes are sending data, or 0.33 if only one processor is sending data. Thus, despite the substantial increase in the communications bandwidth of the Colony switch, on average, the communications-to-compute balance in ASCI White is worse than it was in the IBM SP-2. The implication with respect to the parallel scalability of jobs within and between nodes is clear the scientific code developer must maximize computation within the node and minimize communication among the nodes. This trend, which is not specific to IBM, shows no sign of abating.

Disk drives, just like memory, are increasing dramatically in capacity but not as rapidly in speed. In the past, out-of-core solvers were common in scientific calculations, and in electronic structure calculations. For example, it was once routine to compute all of the integrals once and store them on disk for repeated use later in the calculation. The first parallel supercomputers had just a few processors in each node and, thus, there was room to include several local disk drives. Such nodes were good building blocks for scalable disk subsystems, i.e., disk subsystems whose bandwidth scaled with the number of processors. Examples of machines like this were the IBM SP-2 and Compaq Alpha clusters, and traditional out-of-core techniques were very efficient on these machines. General parallel file systems were often layered on top of these distributed disk systems, although fault-tolerance and interference with running jobs were problems.

The trend toward large SMP nodes means there is relatively little room in the nodes for disks. So, general-purpose subsystems must be used for I/O. Advantages of this approach are that fault-tolerance can be built in, and the parallel file system is separated from the compute nodes so both deliver more reliable performance. The downsides are that such a system can never deliver the raw bandwidth of disks local to each node, I/O consumes network bandwidth, and there is the additional overhead of the parallel file system software. Experience over the years with general parallel file systems (GPFS) from many vendors has not been positive. They have been both slow and unstable. However, the latest version of GPFS from IBM seems to be both fast and robust when supported by the appropriate hardware and used by large parallel applications.

Data storage is also linked to movement of data from one platform to another. Many simulations generate significant amounts of data (on the terabyte scale) and storing, moving, and analyzing these data sets becomes an issue. As a consequence of the inability to rapidly investigate large data sets generated by various simulations, we face a loss of scientific understanding of our results. We need the ability to easily manage extremely large data sets in a way that makes it possible to squeeze more science out of our simulation output. This poses increased requirements for not only long-term data storage, but also for short-term data storage. The latter is especially important for data analysis. We also need the ability to move these data sets between various DOE and NSF HPC centers in order to make optimal use of machine time and visualization capabilities. This need translates into continued expansion of high-bandwidth pipes between these facilities. However, it is possible to ameliorate this situation by growing NERSC in another area: visualization support. Increased capability for scientific visualization support at NERSC would reduce the need for movement of large data sets. It would also aid the NERSC users in the scientific discovery process, thereby further enhancing efficiency and productivity at NERSC.

Appendix B:
Contributors and Contact Information

OBES

Theresa Windus, Greenbook OBES coordinator
Pacific Northwest National Laboratory
509-372-6410
theresa.windus@pnl.gov

Jacqueline Chen
Sandia National Laboratories, California
925-294-2586
jhchen@sandia.gov

Marvin Cohen
Department of Physics
University of California, Berkeley
510-642-4753
cohen@jungle.berkeley.edu

Bruce Harmon
Ames Laboratory
515-294-5772
harmon@ameslab.gov

Robert Harrison
Pacific Northwest National Laboratory
509-375-2037
robert.harrison@pnl.gov

Malcolm Stocks
Oak Ridge National Laboratory
865-574-5163
gms@ornl.gov

Alex Zunger
National Renewable Energy Laboratory
303-384-6432
alex_zunger@nrel.gov

OBER

Doug Rotman, Greenbook OBER Coordinator
Atmospheric Science Division
Lawrence Livermore National Laboratory
925-422-7746
drotman@llnl.gov

Scott Elliot
Earth and Environmental Sciences Division
Los Alamos National Laboratory
505-667-5092
sme@lanl.gov

Brian Hingerty
Life Sciences Division
Oak Ridge National Laboratory
865-574-0844
hingertybe@ornl.gov

Betty Mansfield
Life Sciences Division
Oak Ridge National Laboratory
865-576-6669
bkq@ornl.gov

Starley Thompson
Atmospheric Science Division
Lawrence Livermore National Laboratory
925-423-9923
thompson59@llnl.gov

Warren Washington
Climate and Global Dynamics Division
National Center for Atmospheric Research
303-497-1321
wmw@ucar.edu

OFES

Alex Friedman, Greenbook OFES Coordinator
Accelerator and Fusion Research
Lawrence Livermore National Laboratory and
Lawrence Berkeley National Laboratory
510-486-5592
afriedman@lbl.gov

Don Batchelor
Fusion Energy Division
Oak Ridge National Laboratory
865-574-1288
batchelordb@ornl.gov

Bastiaan Braams
Courant Institute of Mathematical Sciences
New York University
212-998-3276
braams@courant.nyu.edu

Bruce Cohen
Fusion Energy Program
Lawrence Livermore National Laboratory
925-422-9823
bcohen@llnl.gov

Jill Dahlburg
General Atomics Corp.
858-455-3571
jill.dahlburg@gat.com

Steve Jardin
Princeton Plasma Physics Laboratory
Princeton University
609-243-2635
jardin@pppl.gov

Don Spong
Fusion Energy Division
Oak Ridge National Laboratory
865-574-1304
spongda@ornl.gov

HENP

David J. Dean, Greenbook HENP Coordinator
Physics Division
Oak Ridge National Laboratory
865-576-5229
deandj@ornl.gov

Kwok Ko
Stanford Linear Accelerator Center
Stanford University
650-926-2349
kwok@slac.stanford.edu

Doug Olson
Nuclear Science Division
Lawrence Berkeley National Laboratory
510-486-4567
dlolson@lbl.gov

Rob Ryne
Accelerator and Fusion Research Division
Lawrence Berkeley National Laboratory
510-486-7639
rdryne@lbl.gov

Doug Toussaint
Physics Department
University of Arizona
520-621-2554
doug@physics.arizona.edu

OASCR

Mike Minkoff, Greenbook OASCR Coordinator
Mathematics and Computer Sciences Division
Argonne National Laboratory
630-252-7234
minkoff@mcs.anl.gov

David Bernholdt
Oak Ridge National Laboratory
865-574-3147
bernholdtde@ornl.gov

Phil Colella
Applied Numerical Algorithms Group
Lawrence Berkeley National Laboratory
510-486-5412
pcolella@lbl.gov

Ian Foster
Mathematics and Computer Sciences Division
Argonne National Laboratory
630-252-4619
foster@mcs.anl.gov

Jarek Nieplocha
Pacific Northwest National Laboratory
509-372-4469
jarek.nieplocha@pnl.gov