User Requirements for the Energy Research Computational Community

Executive Committee Energy Research Scientific Computing User Group (ExERSUG)

Editors:
Ricky A. Kendall, principal editor
Bastiaan J. Braams

July 13, 1998

Contents

List of Figures

Executive Summary

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

A large shared computing center continues to be essential to the scientific mission of the Office of Energy Research as it provides the researchers with resources that could not be supplied locally. We document this role (Chapter 3) by providing samples of computational applications performed at NERSC. 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 Computational and Technology Research.

The computing hardware normally needed at the forefront of scientific research has evolved from the "vector supercomputers" of the late 1970's and 1980's to presently a mixture of massively parallel computers employing distributed memory, symmetric multiprocessors employing shared memory, and vector multiprocessors. We describe (Chapter 4) these architectures and their anticipated near-term evolution, and document the importance of each of these architectures to the computational needs of Energy Research. The massively parallel machines, based on commercially mass produced chips, can offer extremely large memory at the most attractive price-performance ratio. The vector multiprocessors still offer the fastest turn-around time for many of the scientific codes. Symmetric multiprocessors (SMPs) and clusters of SMPs are developing to provide the same advantages of massively parallel machines and provide more flexibility in use at the same time. 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 for Energy Research. Nothing useful happens in a computer without good software, and we document (still Chapter 4) our requirements for base software, for mathematical and applications libraries, and especially for a unified production environment that supports distributed computing across the high performance computing centers, the Energy Research laboratories, and affiliated university groups. We also discuss research and services at the computing centers.

Based on the presentation of computing endeavors within Energy Research, and of the hardware and software environment at the Centers and as available commercially now or anticipated, we make these specific recommendations:[1]

[1] c.f., Chapter 5

- Expand the high performance computing resources available at NERSC. The procurement (already underway) must ensure an appropriate balance of computational equipment to support the wide range of large-scale applications that use the current facilities.

- Focus efforts to convert appropriate applications to parallel systems. To meet this requirement, NERSC must have a highly qualified and diverse parallel computational science and applications group to work with the core Energy Research computational efforts.

- Strengthen the networks. Reliable, robust and fast access to the computational resources is essential.

- Encourage collaborations and integration among ER physical, computational, and computer science programs. NERSC can act as a collaboratory meeting place for Energy Research scientific efforts to form "virtual laboratory groups."

- Encourage software development for the advancement of Energy Research missions. NERSC should serve as a facilitator and where appropriate as partner in these activities.

- Encourage large-scale applications by enhancing the effective use of the Capability computing systems. Supercomputers should be used primarily for the applications that require the specific and special capabilities of the resource. NERSC should support both Capability and Capacity computing, and is in a unique position to foster the development of a unified production environment for distributed computing.

- Encourage expansion of data intensive computing capabilities. This is an emerging area for many experimental efforts collecting, storing, refining and analyzing large (multi-terabyte) data sets.

User Requirements for the Energy Research Computational Community

1 Background and Introduction

The Energy Research Supercomputer Users Group (ERSUG) encompasses all investigators utilizing the High Performance Computing Access Center (HPCAC) resources of the Department of Energy Office of Energy Research. At the June 1995 meeting held at the National Energy Research Scientific Computing (NERSC) facility, the ERSUG executive committee (ExERSUG) determined the need to reassess the role of computational science in Energy Research (ER) programs. The continuing rapid development of the computational science fields, computer technology (both hardware and software), and the fiscal realities of the 1990s demand frequent periodic review of user requirements and the role that computational science should play in meeting ER program commitments.

It is clear from the current environment at NERSC and at the National Science Foundation (NSF) funded supercomputer centers that the demand for computational resources greatly exceeds the supply of those resources. The CRAY-C90 was a saturated resource only a few days after it was made available to the ER community at NERSC. The advent of Massively Parallel Processing (MPP) supercomputers and Symmetric Multiprocessing (SMP) systems at supercomputer sites have increased both the capabilities and capacity of supercomputer centers but have only marginally met the aggregate demand for computational resources.

When the last user requirements update was written, 1991, parallel supercomputers, in particular MPPs, had just started becoming viable alternatives to vector supercomputers for certain classes of problems. The efforts of the computational sciences and computer sciences community have positioned MPPs as full partners with traditional vector parallel architectures as evidenced by offerings from the vendors such as SGI/CRI and IBM. The explosive growth of workstation processor technology coupled with specialized communication designs have been the mainstay of the hardware technologies in these MPP systems. MPP architectures based on workstation processor technology coupled with standard bus based shared memory systems (i.e., the SMP workstation) have been approaching the throughput performance of traditional vector supercomputers from the other end of the parallel computing spectrum. The evolution of MPP supercomputers by the vendor community benefits from both the fully distributed technology and the cluster of SMP approach. Currently the most popular MPP supercomputers are of the fully distributed class, SGI/CRI T3E, IBM SP. The clustered SMP approach is represented by the SGI Power Challenge Array, Origin-2000, SUN HPC-10000 and the Convex Exemplar.

The difficulty in using MPP systems still lies in adapting the relevant application software to efficiently exploit the distributed resources available (CPU, memory bandwidth and capacity, I/O). The key issue for developing efficient MPP applications is the access to the 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. Both the fully distributed and clustered SMP MPPs require labor intensive re-engineering or redesign and re-implementation of algorithms in most fields of computational science. This labor intensive trend is very much like the transition that occurred in the computational science fields when vector supercomputers were first introduced in the late 1970s. The relatively small local memory (e.g. 256 MB) available on the fully distributed MPPs poses special problems for many applications that are partially alleviated in the cluster-of-SMPs architecture.

There are at least four underlying components to the computational resources used by ER 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 ER community.

The ER community has an interest in all of these components. ER funded programs are vitally interested in and contribute to new high-performance computing technology. The ER community defines and does research in "Grand Challenge" problems that require the development and integration of both new production computing resources and advanced software technologies and algorithms. The Energy Sciences Network is the electronic nervous system that connects ER scientists to the world-wide scientific community and to distributed computational resources that allow ER scientists to complete their programmatic missions.

In this document we will classify the user community of the ER HPCAC resources, focused on the NERSC facility (section 2), delineate a subset of the important scientific areas that require these resources (section 3), and describe the mix of the centralized and local resources (both current and future technology) that are necessary to the vitality of the ER funded research programs (section 4). We will not attempt to measure the actual ER programmatic needs in terms of hours of computer time since they could never be completely met but we do intend to make recommendations (section 5) that will facilitate more effective utilization of current and future resources in the realm of fixed or declining budget scenarios of the 1990s.

2 Energy Research User Community

There are four broad classes of users and projects that utilize the National Energy Research Scientific Computing facility. Grand Challenge projects are characterized by the need to use a significant fraction of the largest computational resources available. Large scale production projects require large computational resources, but do not require to use these resources in dedicated mode. Medium scale projects are those that require computational resources that exceed those available at many local sites but that are modest on the scale of the High Performance Computing Access Center (HPCAC), and they do not require the use of these resources for very long. Small scale projects are those that require no more than a few dedicated high-end workstations for their completion.

The boundary between Grand Challenge projects and large scale production projects is fuzzy at best. Often a relatively small sequence of Grand Challenge computations may be employed to validate a model at the highest achievable accuracy, while a large sequence of production calculations at lower resolution will provide the practical output of the scientific endeavor. Each ER research project may have components that fit into all of the categories, and every ER discipline has projects in each of these classes. It is our goal to utilize these defined classes of users and projects to better determine the aggregate computational resources required by the High Performance Computing Access Center (HPCAC).

2.1 Computing

Grand Challenge (GC) projects require access to a substantial fraction of the top-end high performance computing systems, e.g. an entire Cray C-90, or 256 nodes or more of a Cray T3E or IBM SP2, often in dedicated mode. It can be the demand for large memory (10-100 GB or more) or for high speed (100-300 Gflops or more) that makes an application into a Grand Challenge, though the first demand normally implies the second. Grand Challenge projects arise, for example, in lattice field theory, molecular dynamics, quantum chemistry and quantum solid state, and fluid and plasma turbulence.

Large scale production projects require long-time access to either a top-end high performance computer system or a large cluster of high performance systems. An individual computation will be of the scale that can be executed on current high-end shared-memory parallel machines (memory requirement normally below 10 GB) or on a relatively small segment of an MPP, and the main requirement on the hardware configuration is that it provides fast throughput. The access need to this system may be intermittent or constant.

Medium scale projects have been the bread-and-butter of NERSC's operation; with the current spread of high-end workstations and small SMPs in the laboratories we would classify as medium-scale those projects that require extensive use of a small multiprocessor; many ER scientists would not now have good access to such a machine and many that do have access find their productivity enhanced if they can use, intermittently, a high-end computational facility such as the NERSC C90 or its expected Origin-2000.

Small scale projects may utilize the resources of the High Performance Computing Access Center for a variety of reasons: the required hardware or software is not available locally, the center is convenient for remote collaborations, the center's consulting facilities or charging policies make it attractive. These projects can provide the fill for gaps in effective management of larger scale jobs.

2.2 Services

The Grand Challenge and the large-scale-production users are expected to be computer literate. The HPCAC staff will likely know these users because they have worked together to solve production and development problems and to optimize the code for the hardware configurations available through the HPCAC; in the case of the Grand Challenge applications such an activity may be mandated by the Center. Also the users and NERSC staff may have worked together to integrate visualization tools with the applications code.

Medium scale and small scale users may rely heavily on third party applications and discipline specific applications. They are larger in number than the GC and large scale production users, and their interaction with the HPCAC staff will be as great or greater than the former two classes. The level of computer literacy among these users will span the full gamut.

As users develop projects they are likely to migrate up from the Small-scale to Mid-scale and Grand Challenge projects. Additionally, these divisions are somewhat arbitrary and there is a "gray" area at the transition between these areas. For example, a small-scale user might migrate his application to the MPP starting with a modest initial allocation. Once his/her application scales and is efficient the next years allocation would move to a medium or large scale project allowing scientific testing and validation of the models in the application. If the models provide a solution to a grand-challenge-scale project then that transition should also be forthcoming.

3 Representative Energy Research Scientific Computing Efforts

The Department of Energy, in partnership with it's stake-holders, is entrusted to contribute to the welfare of the Nation by providing the technical information and scientific and educational foundation for technology, policy, and institutional leadership necessary to achieve efficiency in energy use, diversity in energy sources, a more productive and competitive economy, improved environmental quality, and a secure national defense.

The programs of the Office of Energy Research fund basic research in order to advance the fundamental science knowledge base, as well as train future scientists. Research of this type has been supported by the Department of Energy and its predecessors for over 40 years and includes research in the natural and physical sciences; including *basic energy sciences* (research in materials and chemical sciences, engineering and geosciences, and energy biosciences), *magnetic fusion energy*, *health and environmental research*, *high energy and nuclear physics*, computational and technology research in *mathematical, informational, and computational sciences*. These research programs build the science and technology base for our Nation's energy security and scientific leadership.

Access to a broad variety of facilities and laboratories is essential to success in our research efforts. The facilities include large accelerators, experimental detectors, and reactors, high-precision instruments, synchrotrons, massively parallel computers, high-capacity networks, and high-resolution microscopes. These facilities are used by a broad base of user groups from universities and industries all over the country and abroad.

The knowledge, techniques, and instruments developed in the diverse programs of the Office of Energy Research contribute to U.S. economic and national defense security, and are used in therapeutic and diagnostic medical applications, electricity generation, and an array of applications in industry and other scientific areas.

Among the major facilities available to Energy Research scientists are the National Energy Research Scientific Computing (NERSC) facility and the High Performance Computing Research Centers sited at Oak Ridge National Laboratory, Argonne National Laboratory, and Los Alamos National Laboratory. The High Performance Computing Research Centers were established to provide state-of-the-art resources for Grand Challenge computing. NERSC also provides significant resources to the Grand Challenge effort but NERSC is also the principal supplier of production high-performance computing and networking services to the nationwide energy research community.

In the following sections we describe representative research programs that utilize the computational resources available within the ER system. The focus is on research done at the NERSC facility and the material included here is only a small fraction of the high quality science enabled by the facility.

3.1 High Energy and Nuclear Physics

The Office of High Energy and Nuclear Physics is a $1 billion per year basic research program that provides new insights into the nature of energy and matter and operates large world-class scientific facilities for the nation and the international community. High energy and nuclear physics research is conducted by over 3,000 researchers and over 1,000 graduate students from over 100 universities and the National Laboratories.

High Energy and Nuclear Physics are concerned with understanding of matter and space-time at the fundamental level. There are many unanswered questions:

- Why is the electron lighter than the proton?
- Why is there more matter than antimatter in the universe?
- What was matter like in the early moments of the universe?
- Why are the weak interactions weak?
- Why are there fundamental particles that play no role in normal matter?
- Is there supersymmetry?
- What states of matter exist under conditions very unlike those around us?

The answers to these questions will reveal the fundamental building blocks of nature and reveal the complex laws which govern their behavior. These fundamental building blocks or constituents are:

- Quarks: There are six known quarks (up, down, strange, charm, bottom, top) and their antiparticles.
- Leptons: The electron and the heavier muon and tau, three different kinds of neutrino, and the antiparticles of all of these.
- Gauge mesons: The familiar photon, the W and Z mesons, and the gluons.

The standard model of elementary particle interactions provides a detailed description of the components of matter and how they interact. Strongly interacting particles such as protons and neutrons are bound states of quarks and gluons. The dynamics of this binding is described by quantum chromodynamics. The only technique available for computing with this theory on the scales of nuclear binding is lattice gauge theory.

In the search for new particles or symmetries, high energy physics experiments study the fragments emitted from extremely energetic collisions of relatively simple particles (electrons, protons, anti-protons). Planned experiments of this type that will be in operation in the early years of the next decade at several facilities include BaBar, using electron positron collisions at the Stanford Linear Accelerator Center, and CDF and D0 using the proton-antiproton collisions at Fermi National Accelerator Laboratory in Illinois. In 2005, ATLAS and CMS are expected to commence operation at CERN in Geneva using proton-proton collisions. The BaBar experiment is designed to study in detail the properties of b-mesons with the aim of understanding the breaking of matter-antimatter symmetry. The others are aimed at understanding the mechanisms by which particle masses are generated, and exposing the fundamental symmetries of nature.

In the search for new states of strongly interacting matter under extreme conditions of energy density or angular momentum, nuclear physics experiments are carried out at a variety of both high and low energy accelerators. A new nuclear physics facility, the Relativistic Heavy Ion Collider (RHIC), is scheduled to start operation at Brookhaven National Lab in 1999. Here, experiments including PHENIX and STAR will use very high energy colliding beams of heavy nuclei to study the properties of nuclear matter under extreme energy density, replicating the environment of the early universe. A new detector facility, gamma ray energy tracking array (GRETA), is being developed to operate at various low energy nuclear physics accelerators, to investigate the structure of cold nuclei at very high spin states.

The last few years have witnessed an explosion in the complexity of the experiments being carried out in high energy and nuclear physics, both in the volume of data produced and in the complexity of the analysis required to extract the physics from them. The next generation of experiments will produce data sets up to 1000 times that or previous experiments.

ATLAS at LHC

ATLAS is one of the two multipurpose detectors under construction to operate at the LHC at CERN. Proton-proton collisions will be used at 14 TeV center of mass energy. The goal of ATLAS is to provide data that will enable us to determine mechanism responsible for the generation of mass. This hinges on the ability to extract rare events and subtle effects in huge data volumes. Understanding the underlying physics requires detailed simulations. Hence, the simulated datasets will be comparable in size to the experimental ones. Although the amount of CPU required to perform those simulation is large (for example of the order of 2 million MIPS for ATLAS alone), the most pressing problem is data management. We need a distributed Peta-byte-scale data storage facility and sufficient networking bandwidth to guarantee the effective use of this data by the physicists distributed globally. We need data transfer rates of the order of 100 Mbits/s at the university level and tens of Gbits/s in and from the data-store.

The PDSF facility is being used by the Atlas collaboration to simulate the response of the detector to jets. So far 321000 events have been generated using the PDSF system.[2] The data are sent back to CERN for analysis.

STAR (Solenoidal Tracker at the Relativistic Heavy Ion Collider)

---

[2]c.f., http://atlasinfo.cern.ch/Atlas/GROUPS/SOFTWARE/HELP/jet_production.html

The STAR experiment is one of two major experimental facilities being constructed at the Relativistic Heavy Ion Collider (RHIC), which will begin operation in October 1999 with the collision of counter-circulating beams of gold ions, each having an energy of 100 GeV/nucleon. Head-on ("central") collisions of these gold ions will produce a system with a total energy in the center of mass of 3.9 TeV, spread over the volume of the colliding system. Such violent collisions will produce final states of a complexity unprecedented in accelerator-based nuclear and high energy physics. The final state will consist of about 10,000 particles, mostly pions of energy below 1 GeV. Analysis of this enormously complex final state will yield signals of the conditions at the hottest and densest moment of the collision and the system's subsequent decay, and possibly signs of the fleeting phase transitions from confined to de-confined matter and back again. STAR aims to track and measure a large fraction of these particles, perhaps up to 4000 per event, which will be recorded at a rate of one collision per second. Buried among this huge number of "soft" pions are other, rarer signals, such as high transverse momentum jets and electron-positron pairs, which are uniquely sensitive to the hot, dense early phase of the collision, and which STAR will also detect and analyze.

The data from a central gold-gold collision at STAR is expected to have a volume between 12 and 20 MB (in the case of a strong first order phase transition, events could be much larger), leading to a total data volume (after additional processing) of about 300 TByte produced per year. Estimates of the CPU power needed to analyze these data, together with the simulations to understand instrumental effects, range between 10 and 15 kSPECint95. Efficient and organized access to these data sets requires new developments in data handling and databases.

The required computing power and infrastructure to analyze this enormous amount of data is not expected to be fully satisfied by the RHIC Computing Center. NERSC plays an essential role in the analysis of STAR data, both in terms of capacity and capability, the latter being perhaps especially important, since NERSC brings many years of experience in the handling and analysis of very large data sets.

GRETA

GAMMASPHERE is one of the world's two advanced gamma-ray spectrometers which have continued the marked advance in nuclear structure physics seen in the last decade. Detector arrays of greater sensitivity such as GRETA, which is currently being prototyped at Berkeley Lab, promise to open exciting new vistas in nuclear physics. However to take full advantage of current and proposed spectrometers, new computing capabilities are required.

A major focus of GAMMASPHERE has been to investigate the behavior and response of nuclear systems when they are subjected to extremely high angular momentum values. This nuclear rotation can be as fast as a hundred billion billion revolutions per second. The crucial nuclear structure information is contained in the details of the gamma-ray de-excitation flash when 30 or more gamma-rays are emitted. A number of preferred pathways in the de-excitation process can occur. They relate to favorable arrangements of protons and neutrons which often can be associated with specific symmetries or nuclear shapes.

The increase in sensitivity to weak pathways in current and planned detectors is realized by spreading the less correlated background events over a space of increased dimensionality. For instruments such as GRETA, the optimal dimensionality for data analysis is 7-8. Large, specialized databases are required to store and query the data from these experiments. Given a typical datasets consists of 10**9 to 10**10 events with 50 such experiments being conducted each year, a data store of approximately 2TByte is required. NERSC is in a unique position to provide such a super-computing level data store with the computational resources to query such databases efficiently. With the inclusion of proper network resources, such a facility would be a valuable resource to the low-energy nuclear physics community throughout DOE laboratories and the university community. Furthermore, access to high quality scientific visualization services (as described in 4) will be important in the analysis of these high-dimensional datasets.

3.1.1 Lattice Field Theory

Doug Toussaint, Physics Department, University of Arizona

Elementary particle physics is concerned with the fundamental questions "what are things made of" and "what holds them together". Current theories of fundamental forces are theories of fields, objects which are functions of space and time. Quantum field theories of the strong, weak and electromagnetic interaction, namely quantum chromodynamics (QCD), and the Weinberg-Salam model, are generally accepted, while quantum theory of gravity is at best in its infancy. All of these theories involve local invariance principles, or "gauge symmetries".

Quantum chromodynamics describes the strong interactions between quarks; these interactions bind quarks together to form protons, neutrons, and other particles known collectively as hadrons. The strong nuclear force is believed to be a consequence of the much stronger forces of QCD. The other half of the Standard Model, the Weinberg-Salam model, unifies the weak force with electromagnetism. The Weinberg-Salam model predicts the existence of the W, the Z, and the Higgs particle. Although very successful, the Standard Model is incomplete: a true "Theory of Everything" still requires much theoretical and experimental work. Furthermore, many predictions of the Standard Model will be extracted only with considerable effort.

An important tool for this work is lattice gauge theory, introduced by Kenneth Wilson in 1975, in which continuous space-time is replaced by a four dimensional lattice. Wilson's work made possible the use of numerical simulation in the study of the Standard Model. The introduction of a space-time lattice is an approximation; to represent reality accurately, the lattice spacing must be small and the overall lattice size large compared to the scale of the problem studied.

Numerical simulation has been used extensively in the study of Quantum Chromodynamics. QCD is a very successful theory, explaining many results from high energy scattering experiments. However, QCD is also a very complicated theory and defies analytic solution for the fundamental properties of strongly interacting particles, such as their masses. The basic problem is the fact that the strong interactions, which are one of four basic forces of nature, are "strong". This means that they entail highly non-perturbative and non-linear interactions that are not amenable to analytical methods. Over the past decade, computer simulations have proven to be valuable theoretical tools for the study of QCD and other quantum field theories. Many properties of strongly interacting particles have been studied, including the following:

- **Particle Masses:** Hundreds of strongly interacting particles and their masses are, in principle, completely determined by the handful of parameters that define QCD. The explanation of these masses is the first important step in a successful first-principles analysis of the properties of QCD. There is now strong evidence that lattice calculations in QCD give the correct hadron spectrum when the lattice spacing and quark masses are made small enough.

- **Glueball Spectroscopy:** QCD firmly suggests the existence of strongly interacting particles that are not made of quarks. These particles, known as glueballs, are bound states of the gluons that bind quarks. The experimental situation is confused, with several glueball candidates but no definitive evidence. Lattice studies of QCD give important information on glueball properties.

- **Heavy Quark Potential:** The potential between heavy quark-antiquark pairs is important for understanding the spectroscopy of heavy quarks such as the bottom ("b") and top quark.

- **QCD Thermodynamics:** QCD appears to exhibit a phase transition at high temperature to a new form of matter in which quarks are no longer confined inside the proton and the neutron. Such a phase was presumably present during the very early universe. More immediately, experiments at RHIC, the "Relativistic Heavy Ion Collider" under construction at Brookhaven National Laboratory, and at CERN are expected to create short lived "fireballs" of matter in this new phase. Theoretical calculations are important to understand the properties of this phase, and to help interpret these experiments.

- **Hadron Form Factors:** Recent simulations have measured the electromagnetic structure functions of some strongly interacting particles. These results are consistent with experiments, within acceptable error limits.

- **Weak interaction matrix elements:** The past several years have seen slow but steady progress in the calculation of weak interaction matrix elements, which determine the effect of the strong interaction on the weak interaction. These calculations aim, for example, to explain the $\Delta I = 1/2$ rule in Kaon decay and to predict the charge-parity (CP) violation parameters. The purely theoretical background needed for these calculation is formidable, as are the simulations themselves and the data analysis.

- **B meson properties:** A specific set of weak matrix elements which have attracted a lot of attention during the last few years are those involving mesons containing a b quark, the second heaviest of the six known quarks. These mesons are quite massive (~ 5 GeV) and are considered to be the best candidates for elucidating the extremely important phenomena of CP-violation. The phenomena of CP violation is intimately related to the observed baryon asymmetry in the universe. The B-factories being constructed at SLAC and at KEK (Japan) will produce a large number of B-mesons, about 10^8/year, which will be used to study the properties of B-meson especially to view the CP violation phenomena. Interpretation of these experiments will require the input from computer calculations that we are doing. The methodology that we have developed allows us to calculate many quantities with a precision of about 25%–30%. Cost-effective use of the forthcoming results from the B-factory experiments demands that computational methods such as ours must continue to further improve the accuracy in our calculations. With increased computing power and improvements that are now underway, we believe that we can increase our accuracy for calculating many of the important entities to a precision of about 10% in the next few years. This should be very helpful in a confrontation of the predictions of the Standard Model with results from experiments at the B-factories. Amarjit Soni, working at NERSC, and Claude Bernard and the MILC collaboration, working on the Paragon at ORNL, have been pursuing these studies on the DOE supercomputers.

In addition to studies of QCD, supercomputers are essential to the study of other theories. The simulation of the Weinberg-Salam model is of special interest, notably the simulation of the Higgs particle. Numerical simulation of the Higgs sector have been used to bound the energy region where experimentalists must look to find the Higgs particle. Also, simulations allow us to increase our understanding and test our ideas by simulating models other than the real world. We may wish to understand the effects of changing the number of kinds of quarks, or to isolate a phenomenon such as the breaking of chiral symmetry by studying simpler models, often in fewer than four space-time dimensions. For example, John Kogut and his collaborators have used DOE supercomputer facilities to study chiral symmetry breaking using a simpler interaction than the eight kinds of gluons required in QCD.

Lattice QCD simulations make large computational demands. The necessary computational power and memory are available only on massively parallel or vector supercomputers. Lattice theorists have made aggressive use of the newest parallel machines and even built specially designed machines, to do their calculations.

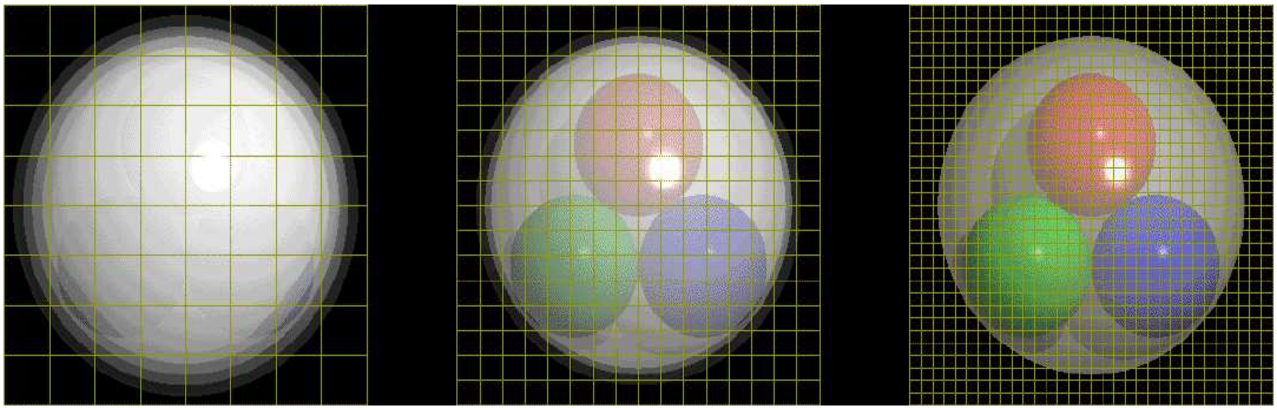

Lattice Size:	8^4	16^4	32^4
Time:	5 node-minutes	5 node-hours per simulation unit	300 node-hours

Figure 1: Lattice gauge theory derives its name from the four-dimensional space-time lattice used to simulate the behavior of elementary particles. This drawing illustrates the need for finer lattice spacing in resolving the quark structure inside the proton. Listed below each lattice are the times required for a single step of the computer code on a modern parallel supercomputer (e.g., the CRI T3D). Thousands of steps are required for a complete simulation.

3.1.2 Ab initio calculations of nuclear systems with realistic forces

J. Carlson (Theory Division, Los Alamos National Laboratory, Los Alamos, NM 87545), J. L. Forest (CEBAF Theory Group, Newport News, VA 23606), V. R. Pandharipande (Department of Physics, University of Illinois, Urbana, IL 61801), S. C. Pieper (Physics Division, Argonne National Laboratory, Argonne, IL 60439), D. G. Ravenhall (Department of Physics, University of Illinois, Urbana, IL 61801), R. Schiavilla (CEBAF Theory Group, Newport News, VA 23606 and Department of Physics, Old Dominion University, Norfolk, VA 23529), and R. B. Wiringa (Physics Division, Argonne National Laboratory, Argonne, IL 60439).

The long-term goal of this group is to understand the stability, structure, and reactions of nuclei as a consequence of the interactions between individual nucleons. We seek a consistent description of nuclear systems ranging in size from the deuteron to neutron stars using a single Hamiltonian. To achieve this we must develop and refine both the nuclear Hamiltonian and the many-body computational techniques used to compute properties of nuclei from the Hamiltonian.

Realistic nuclear forces, that accurately describe two-nucleon scattering and bound states, are very complicated because they are strongly dependent on the relative orientations of the nucleon spins, positions, and momenta, and on the nuclear isospin (which differentiates between protons and neutrons). We express this by using non-central operators such as tensor and spin-orbit operators, with and with out isospin operators. We have constructed several force models over time; our most recent, the Argonne v_{18} nucleon-nucleon potential, accurately fits over 4,300 NN scattering data with a χ^2 near 1, but requires eighteen operator components. There is strong evidence for many-nucleon forces and special relativity can also be important. Solving the many-nucleon Schrödinger equation is consequently a very challenging theoretical problem.

Our quantum Monte Carlo (QMC) calculations of light nuclei use sophisticated variational trial functions, and then systematically improve on them with Green's function propagation (GFMC) to approach the true ground state. These calculations are very computer intensive and have always been limited by the available resources. In 1987 we made the first calculations of ^{4}He that were accurate to $\sim$1% for the given Hamiltonian. In 1995 such calculations of ^{6}Li became possible and at present we can do eight-nucleon systems. Variational cluster Monte Carlo calculations are used for larger closed-shell nuclei, and variational chain summation methods for nuclear and neutron matter. These are necessarily less accurate than the GFMC calculations of light nuclei and we are developing a cluster GFMC method for heavier nuclei.

Our recent progress in both the light and closed-shell nuclei has been made possible by massively-parallel super computers. The programs are written in Fortran-90 and use MPI. Because the kernels cannot be expressed in terms of BLAS's or other standard subroutines, care has been taken to optimize the Fortran for the computers being used; at present this means cache-based machines. The size of the cache is critical for good per-processor speed. Communication rates are very low and with load-balancing we achieve parallel efficiencies better than 95%. Per-processor memory sizes of up to 200 Mbyte are currently used. We use both the IBM SP at Argonne and the Cray T3E at NERSC, and achieve total job speeds of 8 GFLOPS on 64 nodes of the IBM SP (RS6000/593 processors).

Some of our recent results are itemized below; work in all of these areas is continuing.

- QMC has been used to calculate the ground states and many low-lying excited states for all $A \leq 8$ nuclei, demonstrating for the first time the microscopic origin of nuclear shell structure (figure 2).
- Our variational wave functions have been used to explain both elastic and inelastic electromagnetic form factors observed in electron-scattering experiments without the need of effective charges.
- The strong nuclear tensor force has been shown to produce novel toroidal correlations in nuclei (figure 3); several experiments have been proposed to search for these structures.
- The effects of special relativity on nuclear forces and nuclear binding have been studied in the framework of relativistic quantum mechanics for light nuclei.
- QMC has been used to study neutron drops, collections of neutrons bound by an external well, thus providing benchmarks for more schematic methods used to study nuclei far from stability and nuclei in the crusts of neutron stars.
- Studies are in progress for several radiative and weak capture reactions of astrophysical interest, including: ${}^1\mathrm{H}(p,e^+\nu_e){}^2\mathrm{H}$, ${}^3\mathrm{He}(p,e^+\nu_e){}^4\mathrm{He}$, ${}^4\mathrm{He}(d,\gamma){}^6\mathrm{Li}$, and ${}^4\mathrm{He}({}^3\mathrm{He},\gamma){}^7\mathrm{Be}$.
- The QMC wave functions are being used to study the feasibility of constructing polarized ${}^6\mathrm{Li}$ and ${}^7\mathrm{Li}$ targets for electron-scattering experiments.

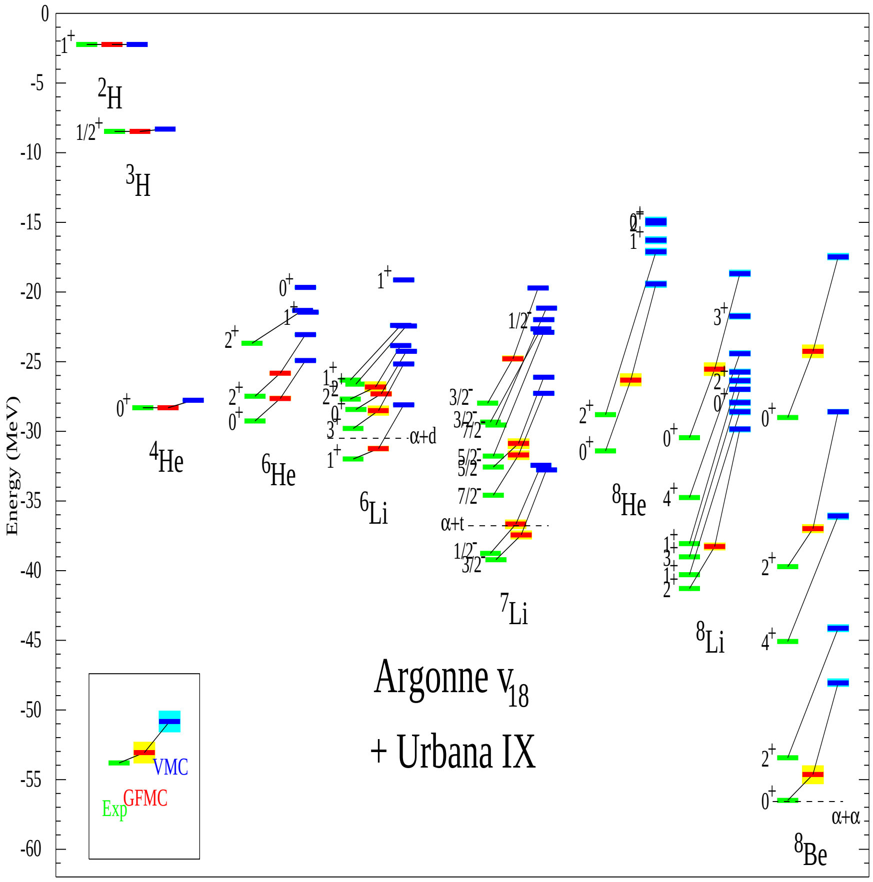

Figure 2: Spectra for $A = 6, 8$ nuclei. For each nucleus, experimental, Green's function Monte Carlo (GFMC), and variational Monte Carlo (VMC) energies are shown.

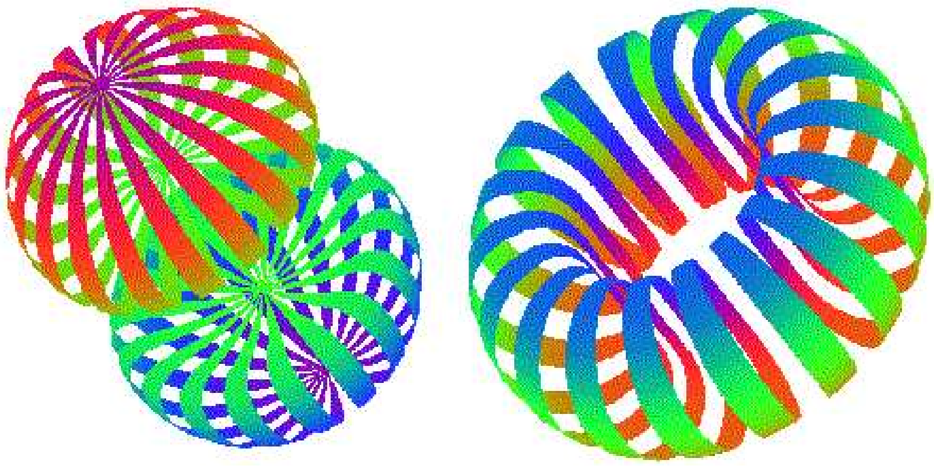

Figure 3: Constant density surfaces for a polarized deuteron in the $M_d = \pm 1$ (left) and $M_d = 0$ (right) states.

3.1.3 Grand Challenge on High-Energy and Nuclear Physics Data Access

Douglas Olson, Point of Contact, Lawrence Berkeley National Laboratory

Research Objective The purpose of this project is to develop techniques and tools that will enable efficient access to the massive datasets of modern High-Energy and Nuclear Physics (HENP) experiments. The prototype development is being carried out in conjunction with the computing facility and experiments at the Relativistic Heavy ion Collider (RHIC) beginning in late 1999. Access to these 100+TByte datasets by hundreds of scientists, in conjunction with carrying out the large-scale computations necessary to refine and reduce them to the essential physical properties buried within, is one of the forefront problems of high-performance computing today. By capitalizing on recent advances in storage systems (HPSS), object database technology (Objectivity Inc.), high-performance scientific computing (at NERSC) and expertise from several institutions across the U.S. we are able to address this important problem affecting fundamental science.

Computational Approach The principle approach to this problem is to cast it in the form of an extremely large hierarchical collection of objects (some persistent and some transient) that experience has shown is very well matched to this event-based experimental data. Modern object-database technology allows us to address the issue of physical storage layout separately from the logical relationships between the objects. The primary issue for physical storage is to organize the data so that it is stored according to how it is accessed rather than by how it is generated.

Coupling large robotic tape storage systems to object databases is a very recent development and there are a number of issues that need to be addressed to enable applying the required amount of high-performance computing cycles to these massive object databases. Addressing these issues is a principle area of research and development this project that is enabled by use of the newly installed HPSS system at NERSC.

Accomplishments A software architecture for access to these massive datasets has been developed and is currently in prototype development. It consists of storage manager components which provide a central coordination service and application code client side components that interface with the storage manager. The storage manager components consist of an index, query estimation, query execution and cache management of the staging disk. The client side components consist of a query interface and an order optimized iterator. Collectively these components permit querying a massive dataset consisting of many files resident on tape, parallel processing of data a file (or set of files) at a time as it is staged to disk, by many users simultaneously.

In order to produce a simulated dataset the NERSC T3E is used to calculate theoretical predictions of what may occur in the relativistic nuclear matter collisions at RHIC. An example of one such study is described below.

Anisotropy in Relativistic Heavy-Ion Collisions We have learned from current experiments studying the collision of nuclei at high energy that there is a large amount of re-scattering of produced secondaries. The Relativistic Quantum Molecular Dynamics (RQMD) cascade model shows that re-scattering reduces the pion multiplicity by 8% at CERN SPS energy. An interesting consequence of this is that a non-spherical, high energy density domain, which can be formed at early stage of a heavy ion collision due to fluctuations in energy deposition, may produce an anisotropy in pion emission at freeze-out.

To investigate this possibility, we used RQMD model to simulate head-on Pb+Pb collisions at SPS energy (158 GeV/nucleon). The resulting freeze-out pion distribution is azimuthally asymmetric (c.f., figure 4; there are fewer pions emitted along the major axis of the ellipse, in agreement with the picture that pion yield is reduced by re-scattering.

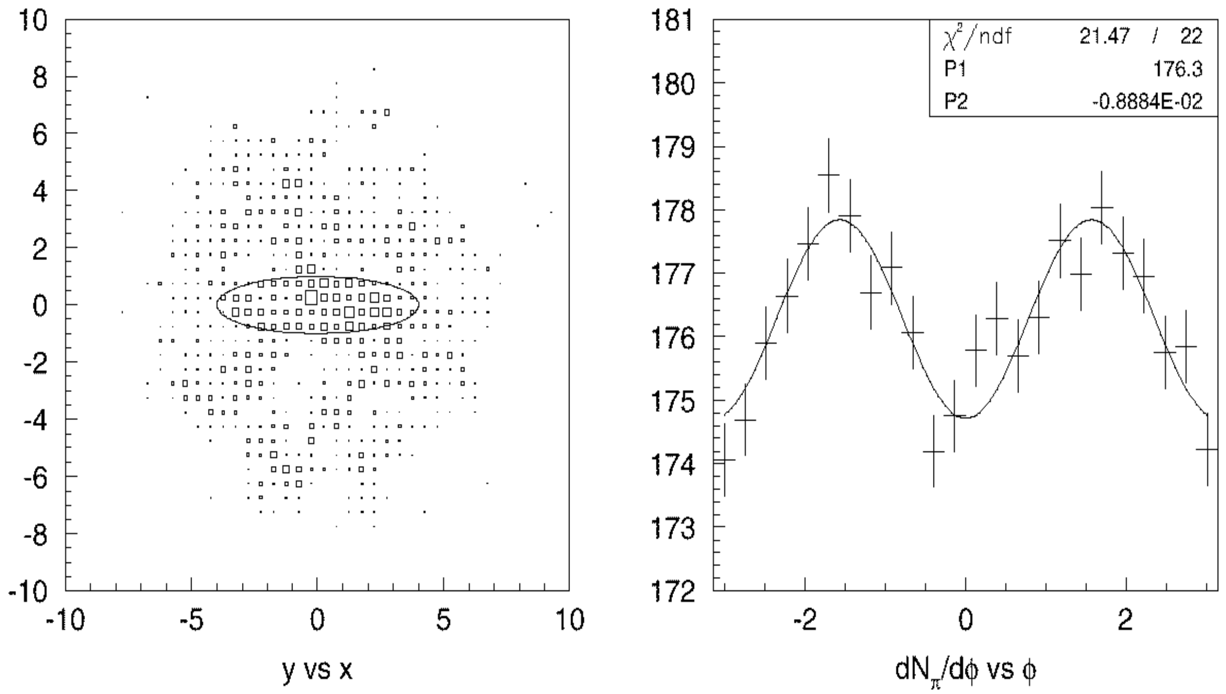

Figure 4: Relativistic Quantum Molecular Dynamics (RQMD) simulation of head-on Pb+Pb collisions.

The RQMD model, or any typical heavy ion collision model, is cpu intensive. Running an RQMD central Pb+Pb event at SPS energy takes about 15 minutes on the NERSC CRAY T3E-900 (peak 0.9 GFlops/cpu, or 70 SPECint95); longer for special events as studied in the present work. For a typical study of special events, one needs on the order of 10K events (or $\sim$ 100 days of processing). Therefore many parallel processors are needed for the calculation to have a reasonable turn-around. Moreover, the current study deals with only one configuration, i.e., a single high density domain with fixed size and shape at the center in every event. It remains a future task to study the location, size and shape dependences of the anisotropy.[3]

[3]See http://www-rnc.lbl.gov/GC/ for additional accomplishments and results of our research efforts and tests.

Significance The direct significance of this work is to enhance the capabilities of large-scale experimental high-energy and nuclear physics to achieve their full scientific potential in exploring new regions of the physical world.[4]

3.1.4 Overview of Accelerator Physics and High Performance Computing

William B. Herrmannsfeldt, SLAC, Stanford University, Stanford, CA 94309

Computational Science is applied to the field of Accelerator Physics to enable fundamental advances in theory and technology leading to new capabilities in particle accelerators.

Importance of Accelerator Physics theory and technology

Experiments based on Accelerator Physics theory and technology have made remarkable discoveries about the basic nature of matter. These discoveries include nuclear structure, the behavior of nuclear matter, quark dynamics, the nature of elementary particles and the fundamental forces, unified field theories, and cosmology. Future possible discoveries include quark plasmas, supersymmetric counterparts to the known existing particles, and the fundamental origin of mass.

At the same time, Accelerator Physics theory and technology have made substantial contributions to other branches of science and technology. Major existing applications include electron microscopy, microprobes, charged-particle-beam lithography, ion implantation, isotope production, particle beams for precision irradiation therapy, superconducting magnets and medical magnetic resonance imaging, neutral-beam heating of plasmas, synchrotron light sources, x-ray lithography, and free-electron lasers. Just this past year synchrotron light was used in the field of biology to determine how meters of DNA can be coiled and managed in cells, to determine the long-sought structure of bacteriorhodopsin, and to determine the largest x-ray crystal structure to date, that of the blue tongue virus made of more that 1000 separate proteins. Indeed, because of their great utility, some 26 new synchrotron light sources are anticipated worldwide.

Possible future applications of Accelerator Physics theory and technology include intense beams for inertial fusion, the production of tritium, the production of nuclear fuels, the transmutation of nuclear waste, and high-speed and high-resolution proton radiography for nuclear stockpile stewardship.

The principal new facility planned for the FY1999 budget by the US Department of Energy is the $1.3B Spallation Neutron Source. This collaborative project, involving several DOE laboratories, is based entirely on Accelerator Physics supported by modeling and simulation of the components and the entire facility.

The Role of Computation In Accelerator Physics

[4] c.f., Relativistic Heavy Ion Collider (RHIC) http://www.rhic.bnl.gov/ and HENP Grand Challenge Project http://www-rnc.lbl.gov/GC/

Accelerator Physics theory and technology make extensive use of modern numerical methods and computational facilities. First, extensive computation of electric and magnetic fields is required to design high quality magnets and high performance RF components. These computations involve fast and highly accurate 3-dimensional Maxwell solvers based both on mesh and finite-element methods.

Second, extensive computation is required to study the nature and long-term stability of single-particle orbits moving in these electromagnetic fields. For example, determining the stability of protons in the CERN Large Hadron Collider currently under design is equivalent in effort to determining the stability of the solar system since the Big Bang. Significant advances have been made in the study of single-particle orbits including the use of truncated power series algebra, Lie algebraic methods, symplectic integraters, normal forms, and the extensive use of maps. These methods, which require both extensive symbolic and numerical computation, have led to substantial improvements in both linear and circular accelerators.

Finally, extensive computation is required to treat both random effects and collective effects involving large numbers of particles. Random effects include synchrotron radiation reaction, RF noise, power supply noise, and ground motion. Collective effects include those arising from electromagnetic wake-fields produced by particles as they pass through various structures, space-charge effects in high current machines, laser-plasma-beam effects in laser or charged-particle driven acceleration in plasmas, and beam-beam effects in both linear and circular colliders. The current calculation of these effects often involves the use of specially written particle-in-cell codes.

Future Computational Needs

In order to realize the promise of future applications of Accelerator Physics theory and technology described above, and in order to realize possible new advances in Accelerator Physics such as laser or particle-driven acceleration in plasmas, acceleration in crystal lattices, and ultra-cold crystal beams, new computational resources will be required. The computation of electromagnetic fields and the computation of a limited number of single-particle orbits in these fields can be carried out by making extensive use of current computer hardware. However, in addition to computing a small number of single-particle orbits, in many cases it would be highly desirable to simulate the behavior of tens of thousands of noninteracting particles in order to make comparisons between theoretical predictions and the behavior of beams in existing accelerators. Such calculations are currently impossible, but are within reach of the next generation of large-scale (100 TByte and 100 Tflop) parallel computers. It is anticipated that the treatment of random effects, which are important for existing electron machines and the Large Hadron Collider, could also be carried out using such computers. By contrast, it is anticipated that a completely satisfactory treatment of long-term many-particle collective effects awaits the advent of computers with Peta-FLOPS speed. Finally, in additional to computational power, advances are required in visualization. For example, in understanding even single-particle phenomena, it would be very helpful to have ways of visualizing 4-dimensional (and even 6-dimensional) phase spaces.

Two Significant Problems

The requirements to proceed to the next level of sophistication in beam dynamics simulations can be represented by considerations of the so-called strong-strong beam-beam interaction. A sample distribution of several million particles in each beam must be tracked for many damping times, or at least 100,000 turns. At each turn, the fields from the beam distribution (which is highly elliptical and can contain important tail regions) must be found from the distribution and then the particles from the two counter-moving beams propagated through the fields. To approach this problem successfully would require a factor of 1000 more cpu power and the development of adequate and tested algorithms.

The complex physical phenomena that occur in the interaction region of an electron-positron linear collider can only be approximately modeled using present capabilities. Achieving the required luminosity for a future TeV linear collider requires on the order of 10^{10} particles per bunch in a collision area of roughly 5 by 400 nanometers. To successfully explore the physics from such an interaction requires separating the subnuclear particle production from the background created by the electromagnetic interactions of the colliding beams. Detailed simulation of these phenomena requires orders of magnitude greater computational power than exists in present facilities.

Summary

Compared to the beam dynamics in an accelerator or storage ring, the simulation capability that has been developed is still simplified and approximate. To understand many experimental observations, such as the beam-beam limit, beam lifetime, and particle backgrounds in detectors, we need the development of more sophisticated computational tools and better algorithms. We have made only the first crucial steps toward the goal of understanding and improving particle accelerators using computational science.

3.1.5 The DOE Grand Challenge in Computational Accelerator Physics

Robert Ryne, Salman Habib, Ji Qiang,
Los Alamos National Laboratory
Los Alamos, New Mexico

and

Kwok Ko, Zenghai Li, Brian McCandless, Wanjun Mi, Cho-Kuen Ng,
Mikhail Saparov, Vinay Srinivas, Yong Sun, Xiaowei Zhan
Stanford Linear Accelerator Center
Stanford University,
Stanford, CA 94309

Introduction[5]

The previous section provided an overview of Accelerator Physics and the crucial role of High Performance Computing to that field. In this section we will discuss progress-to-date in a project now underway that is using high performance computing to advance the state-of-the-art in accelerator modeling.

In 1997 the U.S. Department of Energy initiated a Grand Challenge in Computational Accelerator Physics. The primary goal of this project is to develop a new generation of accelerator modeling tools for HPC platforms, and to apply them to problems of national importance including those mentioned above. The new tools will enable the simulation of problems 3 to 4 orders of magnitude larger than has ever been done before; additionally the use of algorithms and software optimized for the HPC environment will make it possible to obtain results quickly and with very high accuracy. Significant progress has already been made in regard to parallel beam dynamics calculations for proton linac design and parallel wake-field calculations for linear collider design.[6] The tools are being developed by a collaboration involving U.S. national laboratories (LANL and SLAC), universities (Stanford and UCLA), and high performance computing and communications research centers (ACL and NERSC).

Computational Accelerator Physics

Computational accelerator physics consists largely of two main areas: beam dynamics and electromagnetics. In the area of beam dynamics, our goal is to develop an HPC capability to model the acceleration and propagation of an intense charged particle beam as it passes through the various structures in an accelerator complex. In the area of electromagnetics, the goal is to develop a capability to calculate the electromagnetic fields in very large, highly complex structures. These two main areas are described in more detail below.

Beam Dynamics

Many systems involving intense charged particle beams can be described by the Vlasov-Poisson equations. There are two main approaches to solving these equations: particle simulation techniques and direct methods.

In the particle simulation approach, the beam distribution function is represented by a number of macroparticles, typically 10's to 100's of millions in a large scale simulation. Often the single particle equations of motion are derived from a Hamiltonian which includes both externally applied fields and a mean field due to the beam's space charge: $H = H_{ext} + H_{sc}$. Such a form is ideally suited for the application of symplectic split-operator methods.[7] Consider a Hamiltonian that can be written as a sum of two terms, $H = H_1 + H_2$, that are each exactly solvable, *i.e.*, suppose we can obtain the mapping $\mathcal{M}_1$ corresponding to H_1 and

[5]Work supported by the U.S. Department of Energy, Division of Mathematical, Information, and Computational Sciences, Division of High Energy Physics, and Office of Defense Programs.

[6]c.f., R. Ryne and S. Habib, *AIP Conf. Proc.* **391**, p. 377 (1997) and X. Zhan and K. Ko, *AIP Conf. Proc.* **391**, p. 389 (1997).

[7]E. Forest *et al.*, *Phys. Lett. A* **158**, p. 99 (1991).

the mapping $\mathcal{M}_2$ corresponding to H_2. Then the following algorithm is accurate through 2nd order:

$$\mathcal{M}(\tau) = \mathcal{M}_1(\tau/2)\ \mathcal{M}_2(\tau)\ \mathcal{M}_1(\tau/2)\ , \tag{1}$$

where τ denotes a time step. If H_1 depends only on momenta and H_2 depends only on positions, then this algorithm is the same as the well-known leap-frog algorithm. However, the split-operator approach provides a powerful framework capable of dealing with far more complicated Hamiltonians often encountered in accelerator physics, where the Hamiltonian is usually approximated by a high-order polynomial in the phase space variables.

Besides being able to treat Hamiltonians with more terms, this approach is easily generalized to higher order accuracy in time. A well-known fourth-order algorithm is due to Forest and Ruth,[8] and an arbitrary order scheme was derived by Yoshida.[9] There are also implicit symplectic methods based on this approach that do not require the Hamiltonian to be split into a sum of exactly solvable pieces.[6] Finally, time-dependent systems are easily treated by "extending the phase space".[7]

Unlike some split-operator treatments that separate the Hamiltonian into terms involving only positions and only momenta, our particle simulations separate the Hamiltonian into terms involving the external fields and terms involving the self fields. The external fields are treated using well-established techniques from magnetic optics. One advantage of this is that it enables us to take large time steps, since the dynamics due to external fields is usually dominated by a linear map which is easily obtained analytically or numerically. To treat the self fields, we use a 3D Particle-In-Cell (PIC) approach with area weighting. Open boundary condition are treated using the method of Hockney.[10] Charge deposition and field interpolation are parallelized using the method of Ferrell and Bertschinger.[11]

Unlike particle-based methods, direct solvers represent the distribution function on a phase space grid. A Vlasov-Poisson code solves the equation:

$$\frac{\partial f}{\partial t} + (\mathbf{p}\cdot\partial_{\mathbf{x}})f - (\nabla V\cdot\partial_{\mathbf{p}})f = 0, \tag{2}$$

where $f(\mathbf{z},t)$ is the phase space distribution function, and where $\mathbf{z}$ denotes the phase space variables. The potential V is a sum of external and space charge contributions. To solve this equation, the approach we use is a spectral method combined with the above-mentioned split-operator techniques. For example, a second-order algorithm for evolving the distribution function is:

$$f(\mathbf{z},\tau) = \mathrm{e}^{-\frac{\tau}{2}\mathbf{p}\cdot\partial_{\mathbf{x}}}\ \mathrm{e}^{\tau\nabla V\cdot\partial_{\mathbf{p}}}\ \mathrm{e}^{-\frac{\tau}{2}\mathbf{p}\cdot\partial_{\mathbf{x}}} f(\mathbf{z},0)\ . \tag{3}$$

[8]E. Forest and R. Ruth, *Physica D* **43**, p. 105 (1990)

[9]H. Yoshida, *Phys. Lett. A* **150**, p. 262 (1990).

[10]R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles*, (Adam Hilger, NY, 1988).

[11]R. Ferrell and E. Bertschinger, *Int. J. Mod. Phys. C* **5**, p. 933 (1994).

Stochastic corrections to Vlasov-Poisson evolution occur due to particle collisions and noise in external fields. To treat these effects requires the addition of damping and Langevin forces in the PIC simulations which corresponds to solving the Fokker-Planck equation for the distribution function.

Electromagnetics

The new set of electromagnetic tools that is targeted by the Grand Challenge differentiates itself from existing available codes in several respects. One is the use of irregular grids for better geometry description. Another is the capability for adaptive mesh refinement to improve accuracy. The last important distinction is the implementation on multi-processing platforms to accommodate very large scale simulations. These combined features will provide the high resolution required to accurately model large complex structures, such as the damped detuned structure (DDS) which has been specifically designed to suppress wake-fields in the NLC linac.

Included in this tool set are solvers in both the time and frequency domains. The time domain solver is useful for computing wake-fields and other transient effects, such as bunch heating, that are associated with beam passage through an accelerator component. The frequency domain solver calculates normal modes and is needed for designing and analyzing accelerating cavities. Another solver of practical interest deals with the matching of RF components by finding the scattering parameters at the input and output ports. All these modules are being developed with either the finite element or finite volume formulation that is based on an unstructured grid.

It is a challenging task to implement adaptive mesh refinement on parallel computers when unstructured grids are involved. The procedure entails the iteration of a mesh manipulation step and a numerical solution step which provides the feedback as to where mesh quality has to be improved. On parallel architectures this requires a domain decomposition strategy that preserves load balancing from one refinement to the next.

Accomplishments

In 1997 we developed a new 3D beam dynamics code called IMPACT, which is based on split-operator techniques. IMPACT (Integrated-Map and Particle Accelerator Tracking code) has an accurate and efficient treatment of RF accelerating gaps, obtained by numerical integration of the gap transfer map rather than integration of single particle trajectories. The code is especially useful for modeling superconducting proton linacs, where there are only a few types of accelerating cavities. The code has been parallelized and runs on the T3E at NERSC and the Origin 2000 system at the ACL. A second code, called HALO, has been developed specifically for beam halo studies. This code was developed in collaboration with the University of Maryland, and includes a new 3D beam equilibrium model, based on analytical work of the Maryland group, which helps isolate halo growth mechanisms. Finally, in 1997 we also developed a parallel version of a code called LINAC, which is being used to support the Accelerator Production of Tritium project. This is now the primary code for performing large-scale beam halo simulations for the project.

The huge amount of data in a high-resolution beam dynamics simulation, coupled with the fact that we are often interested in the very small fraction of the particles in the halo, necessitates the use of internal data analysis in our codes prior to storing simulation results. We have developed parallel algorithms that drastically reduce the amount of data needed to visualize the halo. An example is shown in Fig. 5, which is the result of a HALO simulation. The system being modeled is a spheroidal bunch which develops a halo due to improper matching into the beam-line.

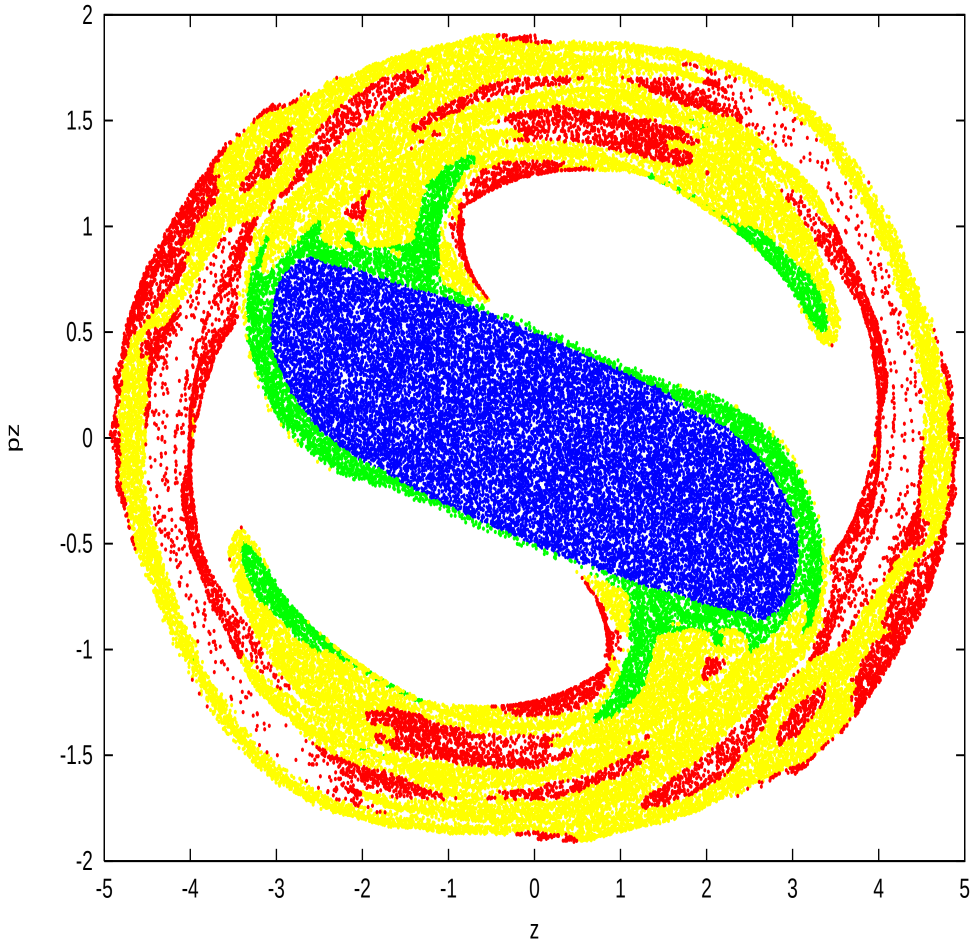

Figure 5: Density in longitudinal phase space, from a 25 million particle simulation of a mismatched beam.

The simulation used 25 million particles and a 256x256x256 grid for the Poisson solver. The figure shows the longitudinal phase space, (z-pz), after the halo has formed. It contains approximately 100,000 particles grouped in four regions. The density varies from 1 at the beam center to $\frac{1}{10}$ at the edge of the central region. Similarly, the density equals $\frac{1}{100}$, $\frac{1}{1000}$, and $\frac{1}{10000}$ at the edge of the remaining three regions. From the core to the halo, these regions appear as blue, green, yellow, and red, when shown in color. If we had simply plotted the same number of particles chosen randomly from the 25 million in the simulation, the halo would barely be visible and show almost no structure.

We have also developed two- and three-dimensional eigenmode solvers for computing normal modes in accelerator cavities. These codes are based on the finite element method using quadratic elements and have mesh refinement capabilities. A parallel version of the 2D module has been used to calculate the dipole wake-fields in a tapered accelerator section (in this case, a detuned structure) consisting of 206 individual cavities. Fig. 6 shows some of the modes in the section. The horizontal axis denotes the axis of the structure, and every fifth mode, from 1 to 206, is shown from the bottom to the top of figure 6. These are the detuned modes which lead to reduced transverse wake-fields due to de-coherence. Work is in progress to model the Damped Detuned Structure (DDS) which incorporates damping by coupling each cell to external waveguides.

3.1.6 Non-equilibrium Quantum Field Theory

Salman Habib
T-8, Elementary Particles and Field Theory, Theoretical Division
Los Alamos National Laboratory.

In recent years there has been considerable progress in the relatively new area of non-equilibrium quantum field theory on several different fronts. Theoretical advances include systematic techniques (such as the large-N expansion) for incorporating higher point correlations into the causal evolution of quantum fields from well-defined initial data, first-principles derivation of transport coefficients and effective kinetic theory at arbitrary temperatures in weakly coupled scalar theories, analysis of the effects of damping in gauge theories and the associated construction of effective theories for the low frequency modes, and systematic studies of de-coherence, effective irreversibility, and the quantum-to-classical transition. Coupled with rapid advances in high performance computing, quantitative attacks are now possible on a host of outstanding, but until recently quite intractable, problems in quantum field theories under extreme non-equilibrium conditions. In addition to the many specific applications, the theoretical aspects of non-equilibrium QFT involve a fascinating interplay of quantum dynamics and coherence, statistical mechanics, kinetic theory, and thermodynamics, whose practical solution will present novel computational challenges.

Specific problems that require non-equilibrium quantum field theoretic techniques for their solution include such diverse examples as topological transitions and electroweak baryogenesis, dynamics of the quark-gluon plasma and disoriented chiral condensates, effective field theories and the real-time renormalization group, and thermalization and reheating in the early universe. In addition there are applications of these methods in condensed matter physics to areas such as dynamics of phase transitions and transport of nonlinear coherent structures.

The main focus of our research group has been on the development of field theory and numerical methods for attacking time evolution problems in quantum field theory such as the back-reaction of quantum fields on classical fields as well as the dynamics of non-equilibrium phase transitions. One set of problems is related to experimental programs planned for RHIC and Fermi-lab: The processes we are studying are the production of unusual muon pairs from a time evolving quark gluon plasma produced following a relativistic heavy ion collision and the production of Centauro events (large regions with no neutral pions) from disoriented chiral condensates produced in the non-equilibrium chiral phase transition expected during the time evolution of the plasma. A second set of problems is related to fundamental issues in field theoretic dynamical systems such as the quantum-to-classical transition, the dynamics of second order phase transitions and of nonlinear coherent structures, and the thermalization of quantum fields in initial value problems. A final set of problems connects with the fields of cavity QED and atomic optics both of which have become excellent testing grounds for fundamental ideas in non-equilibrium quantum dynamics. We are providing theoretical support for experiments being carried out at Caltech and at the University of Texas, Austin.

The study of non-equilibrium field systems poses tremendous numerical challenges. Computational difficulties such as multiple length and time scales with very large dynamic range, nonlocal kernels, inclusion of fluctuations, and extreme storage requirements are typically encountered. Because of these barriers, real progress has been possible only with the advent of massively parallel computing platforms, such as the T3E now at NERSC. Looking to the near future, non-equilibrium field theory codes will require Teraflop performance as well as memory requirements in the low TBytes.

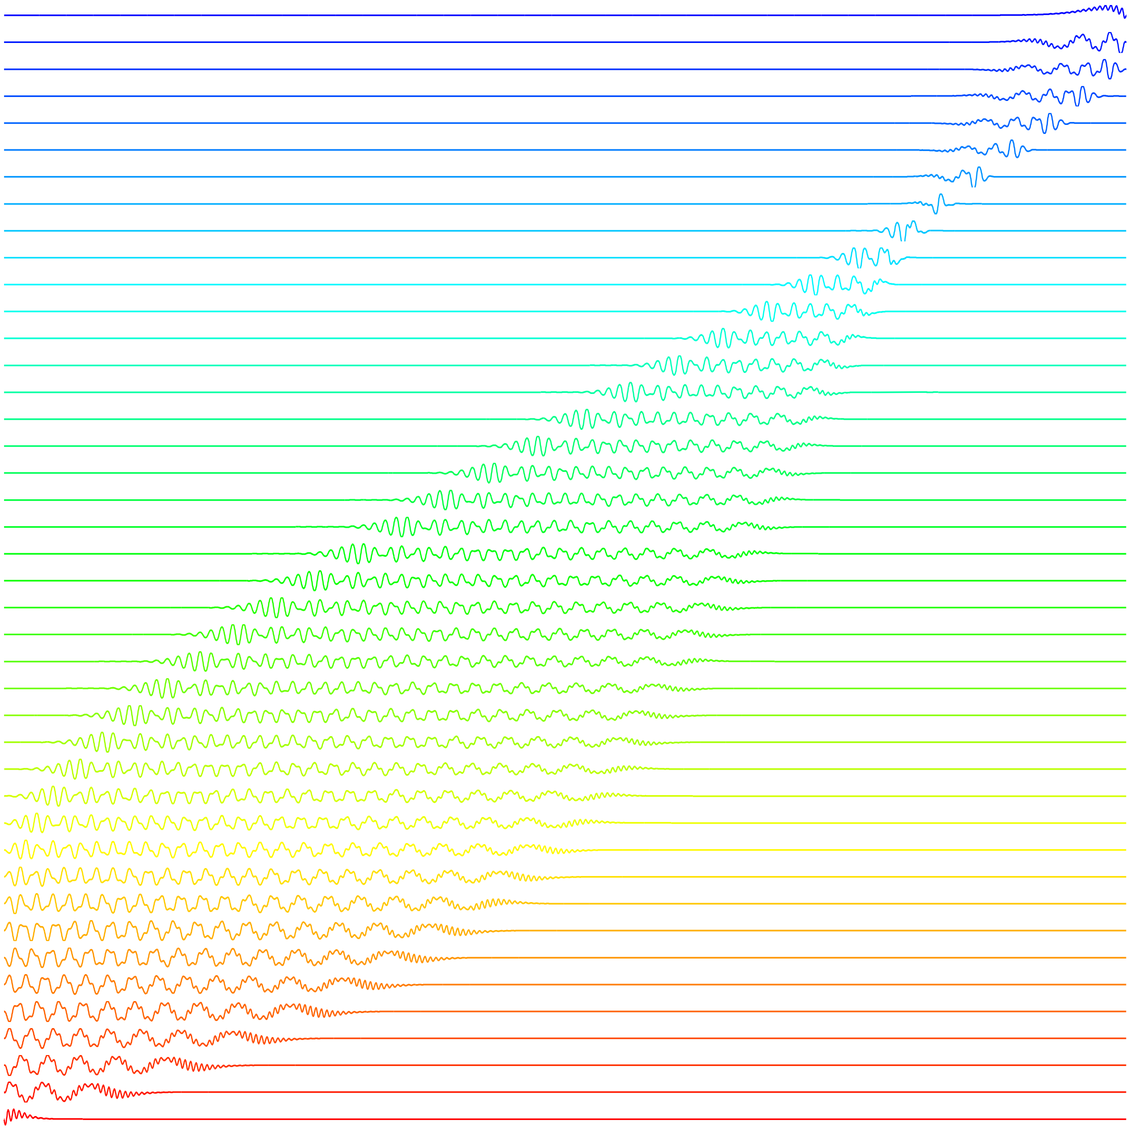

Figure 6: Modes in a detuned structure, from a simulation of an entire 206-cell section.

3.2 Basic Energy Sciences

The Office of Basic Energy Sciences supports research to advance the scientific and technical knowledge and skills needed to develop and use new and existing energy resources in an economically viable and environmentally sound manner. Areas of interest are:

- Materials Sciences: The objective of the Materials Sciences 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 comprised of the subfields metallurgy, ceramics, solid state physics, materials chemistry, and related disciplines where the emphasis is on the science of materials.

- Chemical Sciences: The objective of this program is to expand, through support of basic research, knowledge of various areas of chemistry, physics and chemical engineering with a goal of contributing to new or improved processes for developing and using domestic energy resources in an efficient and environmentally sound manner.

- 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.

- Geosciences: The Geosciences Research Program focuses on basic, or fundamental, research dealing with the response of natural rocks, minerals, and fluids to natural or anthropogenic perturbations. The emphasis is on processes active in the near-surface region of the earth's crust which modify and alter the physical and chemical state of earth materials.

- Engineering: The major goals of the Engineering Research Program are to extend the body of knowledge underlying current energy engineering practice and to broaden the technical and conceptual base for solving future engineering problems in energy technologies.

3.2.1 Multi-scale Materials Simulations on Massively Parallel Machines

Martina Bachlechner, Rajiv K. Kalia, Aiichiro Nakano, Andrey Omeltchenko,
Kenji Tsuruta, and Priya Vashishta
Concurrent Computing Laboratory for Materials Simulations,
Louisiana State University.

In recent years we have witnessed rapid progress in large-scale atomistic simulations, highly efficient algorithms for massively parallel machines, and immersive and interactive virtual environments for analyzing and controlling simulations in real time. As a result of these advances, simulations can now reliably predict properties of materials in advance of fabrication and are thus capable of providing vital input to experimental efforts.

In the past few years, the LSU group has carried out multi-million particle molecular dynamics (MD) simulations to investigate properties and processes in novel ceramic materials, atomic-level stresses in silicon/silicon nitride nanopixels, hypervelocity impact damage in thermal-barrier coatings, and nanoindentation in ceramics. A couple of these simulations are presented below.

Dynamic Fracture in Nanophase Ceramics

Advanced ceramics such as Si_3N_4, SiC, SiO_2, and Al_2O_3 are highly desirable materials for applications requiring extreme operating conditions. Light weight, elevated melting temperatures, high strengths, and wear and corrosion resistance make them very attractive for applications in aeronautics, automotive, electronics, and advanced manufacturing industries. The only serious drawback of ceramics is that they are brittle at low to moderately high temperatures.

In recent years a great deal of progress has been made in the synthesis of ceramics that are much more ductile (and metals that are harder) than conventional coarse-grained materials. These so called nanophase materials are fabricated by in-situ consolidation of nanometer size clusters. Despite a great deal of research, many perplexing questions concerning nanophase ceramics remain unanswered. Experiments have yet to provide information regarding the morphology and dynamics of cracks or why nanophase materials fracture much more gracefully than crystalline ceramics with the same chemical constituents.

Using large-scale MD simulations, we have investigated the morphological and dynamic aspects of fracture in nanophase Si_3N_4. We find that the critical strain at which the nanophase system fractures is enormous compared to the strain crystalline Si_3N_4 can sustain. This is due to: i) plastic deformation in interfacial regions; ii) deflection and arrest of cracks by nanoclusters; and iii) multiple crack branching. None of these mechanisms are operative in the crystalline system. This demonstrates the dramatic effect of nanostructures on fracture in nanophase ceramics.

Simulation of Internal Stresses at Semiconductor/Dielectric Interfaces

In silicon and gallium arsenide semiconductor technology, the use of SiO_2 and Si_3N_4 thin films as dielectrics, encapsulation layers or mask layer is ubiquitous. One of the main issues concerning as-grown processed semiconductor devices is stress/strain distributions in dielectric layers and semiconductor structures as a function of the pixel size. Device pixellation and isolation create edge and corner stresses that are highly inhomogeneous, evolve with the device processing step (and hence each processing step can be impacted by the preceding steps), and in the final configuration, impact the electronic behavior and long-term reliability of the device.

Historically, stress related problems in semiconductor (mostly Si) technology have been dealt with in terms of the classical linear elasticity theory of continuum bodies with the aid of the finite-element method (FEM). However, as the minimum feature size in Si integrated circuit (IC) technology approaches the nanometer regime, the applicability of the FEM to examining stress (particularly at the edges and corners) becomes highly questionable. The impact of the discrete atomic nature of the system on such stresses and of such stresses, in turn, on stabilizing local atomic configuration, needs to be self-consistently examined.

Recently we have performed 10-million-atom MD simulations to investigate the atomic level stress/strain distributions in $Si/Si/Si_3N_4$ nanopixels, see Figures 7 and 8. These simulations provide insight into the nature and degree of stresses in mesas as a function of temperature, density and stoichiometry of the dielectric films.

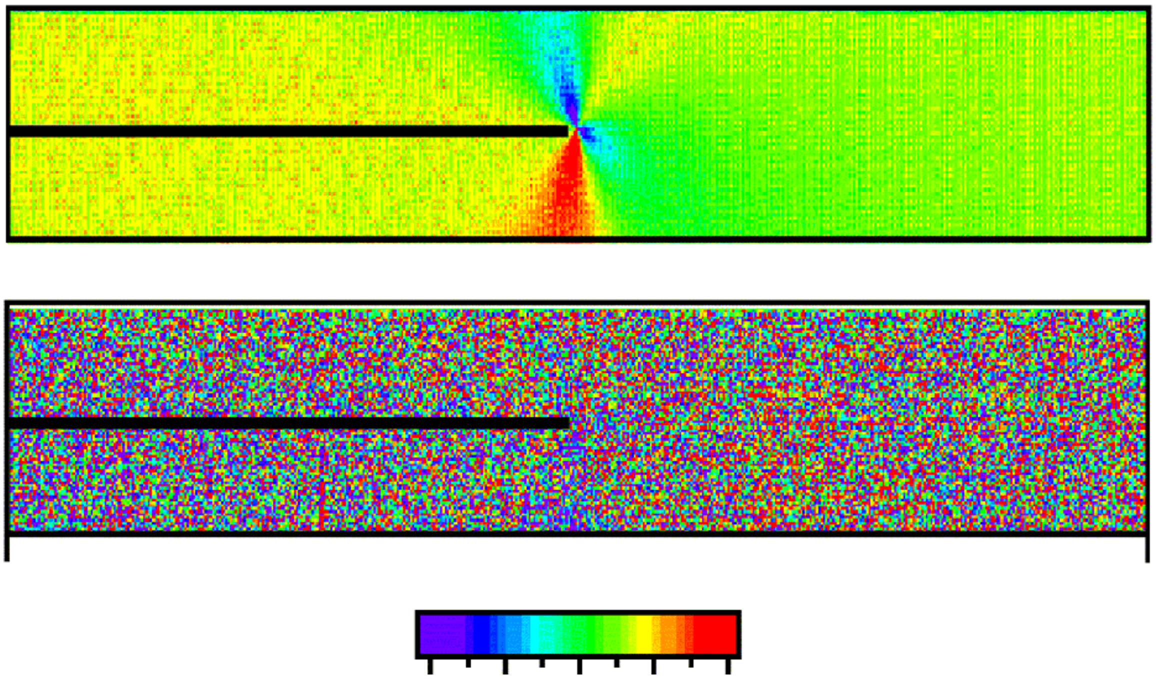

Figure 7: Results of a 30-million-atom MD simulation for shear stress distributions in SiC at 300 K (top) and 1500 K (bottom).

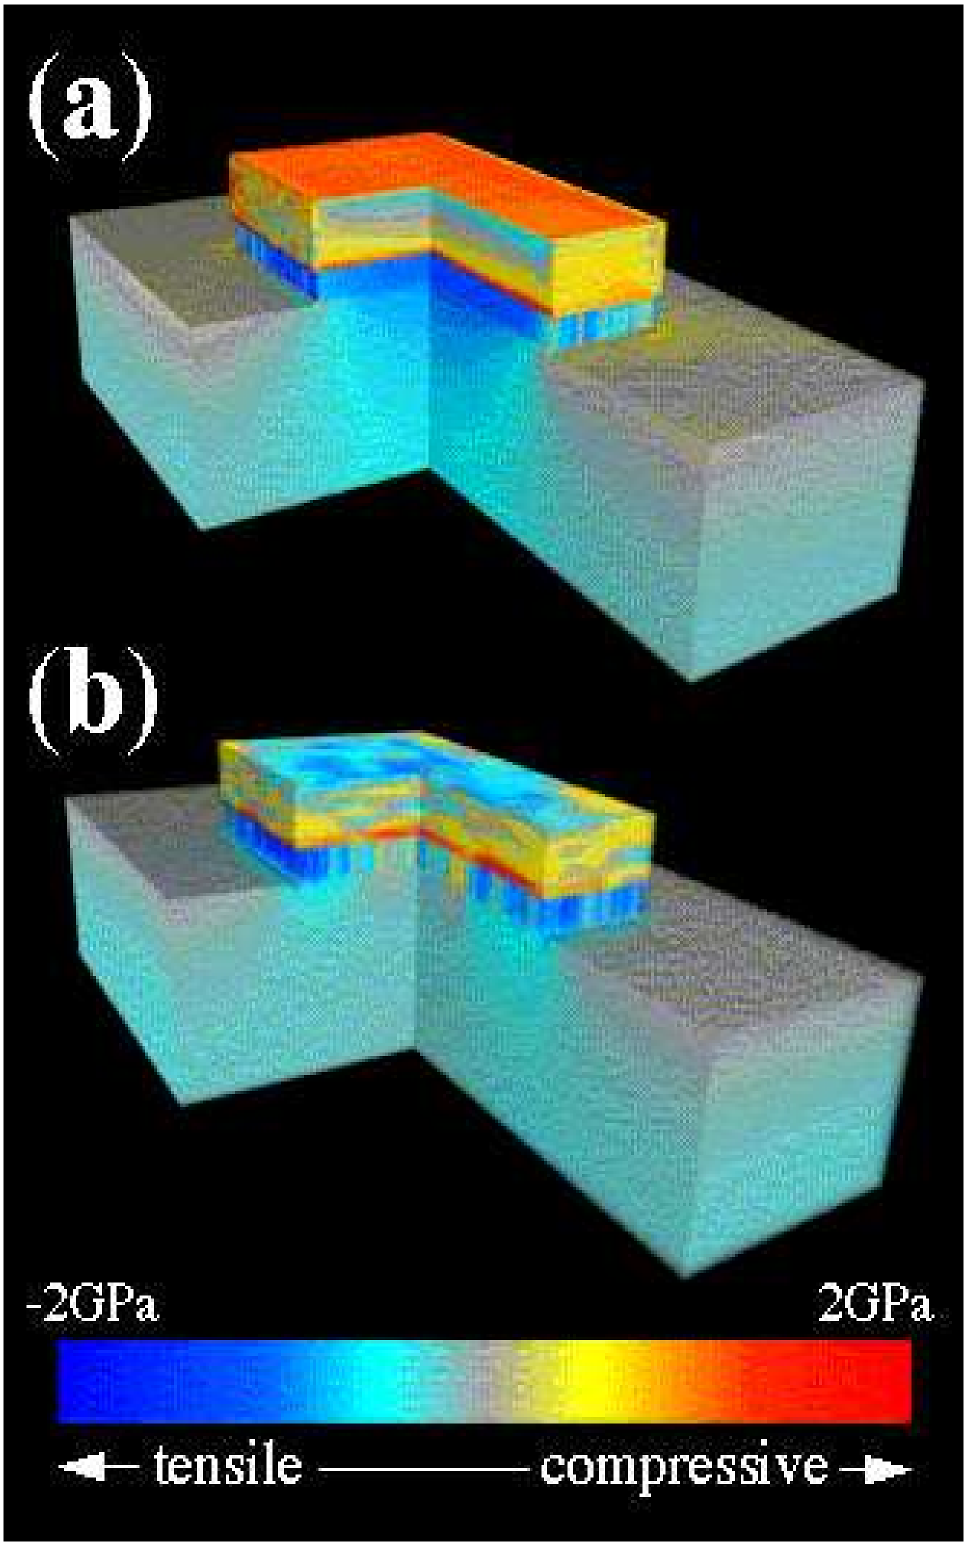

Figure 8: Stress distributions in (a) crystalline and (b) amorphous $Si/Si/Si_3N_4$ nanopixels.

Multi-scale Algorithms for Materials Simulations

Large-scale simulations discussed above have been performed with:

- Space-time multi-resolution MD based on the O(N) fast multipole method and a multiple time-scale scheme;
- Fuzzy-clustering approach to hierarchical dynamics including quaternion-based rigid-body dynamics, symplectic implicit integration, and normal-mode analysis;
- Wavelet-based adaptive load balancing scheme using curvilinear-coordinate transformation;
- Multilevel preconditioned conjugate gradient method for variable-charge MD.

Efforts are underway to develop a novel multi-scale simulation approach which will combine the electronic structure, MD, and finite-element methods (Fig. 9) to study the mechanical behavior of materials at mesoscopic length scales.

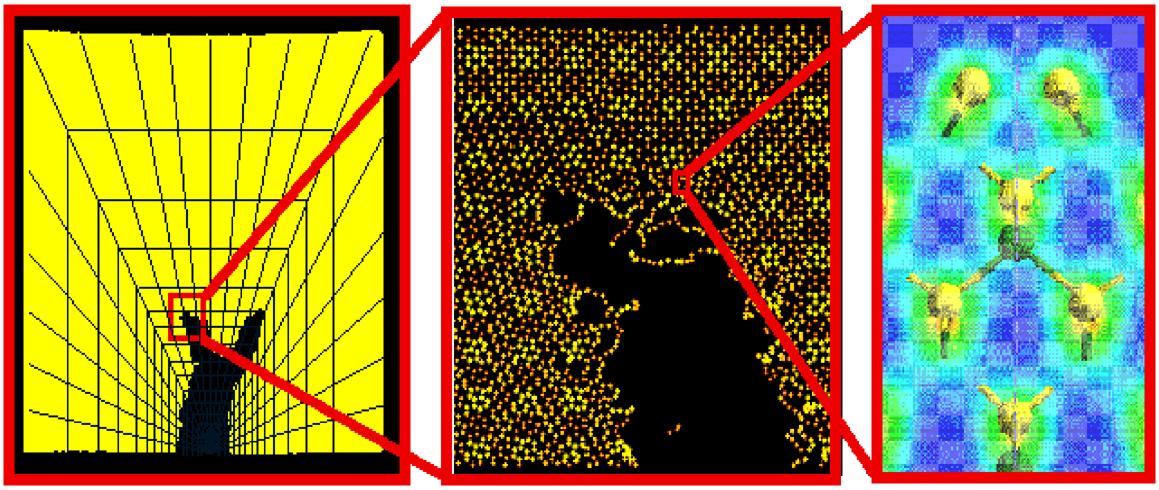

Figure 9: Schematic of a multi-scale simulation approach which combines electronic structure method with molecular dynamics (MD) to describe properly bond breakage and bond formation phenomena. The MD, in turn, is embedded in the finite-element method to take into account nonlinearities in the system.

3.2.2 First Principles Simulation of Materials Properties

G. Malcolm Stocks, Oak Ridge National Laboratory
Bruce N. Harmon, Ames Laboratory.
James W. Davenport, Brookhaven National Laboratory

The central objectives of the Partnership in Computational Sciences (PICS) "First Principles Simulation of Materials Properties" grand challenge project are the development of new scalable parallel first principles methods for large systems, the application of them to problems that are beyond the limits of present methods, and the demonstration that massively parallel supercomputers can play an essential role in enabling their development. We have exceeded our initial targets of being able to accurately treat large collections of atoms (>1,000,000 for classical potentials, >10,000 for tight binding molecular dynamics, and >500 for ab-initio methods) in both static and dynamical applications. The long range goal is to provide theoretical input that can guide the discovery and development of advanced materials and to make a new generation of materials simulation methods available to industrial, government and academic researchers.

A key to designing materials for structural, magnetic, optic, electrical, and high temperature applications is the understanding and ability to control the synthesis and processing at the atomic level. As many of the crucial macroscopic properties of materials actually depend on defects, clusters, and microscopic structures involving hundreds to thousands of atoms, it is only with the availability of modern high performance computing, provided by the HPCCP, that first principles modeling of these important structures and related materials properties can now be undertaken. To this end we have developed, tested and are now applying a hierarchy of increasingly accurate and computationally intensive techniques:

- classical molecular dynamics
- tight binding molecular dynamics
- ab initio methods.

Parallel computing combined with new algorithms that scale linearly with the number of atoms are being used to efficiently calculate the electronic structure and quantum mechanical forces for systems of up to 500 atoms for ab initio methods and up to 10,000 atoms for tight binding molecular dynamics using the Paragon XP/S 35.

The melting and pressure/temperature phase diagram of carbon is an example where conditions are too extreme for laboratory experiments, but where accurate molecular dynamics simulations are leading to new insights for understanding natural and artificial diamond synthesis. Figure 10 shows a snapshot of 512 carbon atoms in a diamond lattice in the process of melting ($T > 4000K$). In the picture, the red atoms indicate four-fold bonded (diamond-like) atoms, the blue atoms indicate three-fold (graphite-like) bonded atoms, and there are also a number of two-fold and five-fold coordinated atoms. The large number of three-fold atoms is an indication that the liquid phase is less dense than the four-fold diamond phase. This is in contrast to silicon, in which the liquid phase has a higher average coordination than in the diamond structure. By running such simulations for the coexistence conditions of the solid and liquid phases the melting temperature of diamond (as a function of pressure) are determined.

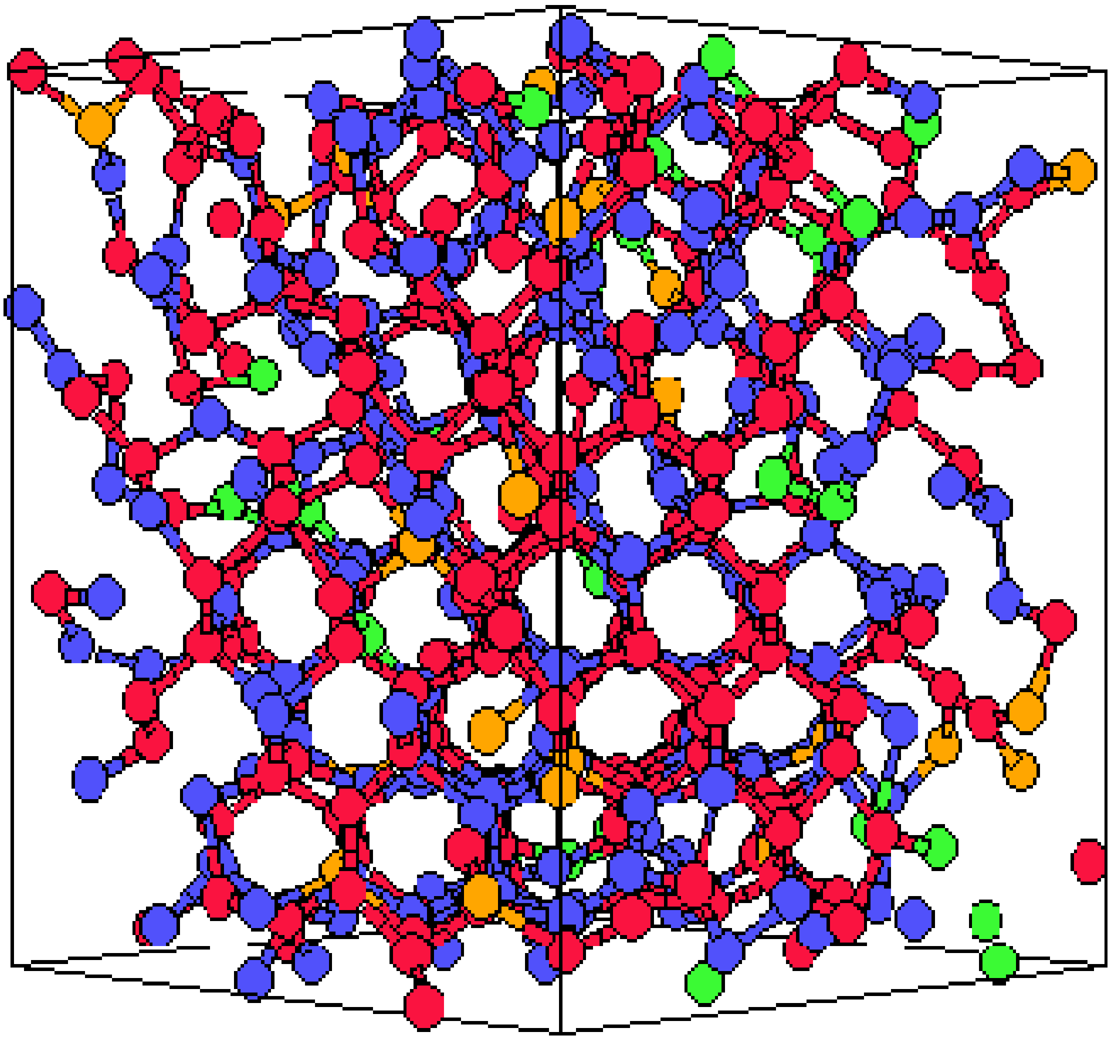

Figure 10: Melting of Carbon in a Diamond lattice.

Magnetic alloys are at the heart of a wide range of technological applications from the oldest of structural materials to the next generation of data storage and retrieval devices. However, a detailed microscopic understanding of alloy magnetism is lacking, hindering further development of these technologies. Using a new ab initio method we have been able to study, for the first time, the nature of the magnetic state in disordered alloys. In figure 11 Ni (large blue spheres) and Cu (small red spheres) atoms occupy the lattice sites of a 256-atom/unit cell model of a Ni rich disordered NiCu alloy. The local Ni-site magnetic moment is distributed inhomogeneously, varying from a minimum of approximately 0.29 Bohr magnetons (short blue arrows) to a maximum of approximately 0.6 Bohr magnetons (long red arrows). Interestingly, the magnetic moment on a Ni-site correlates with the total magnetic moment on the nearest neighbor shell of atoms surrounding it: large red arrows tend to be surrounded by other reddish arrows, while small blue arrows are surrounded by either Cu sites having no moment or other blue arrows. The results of these simulations are being used to re-interpret results of neutron scattering measurements of magnetic correlations in these alloys and to provide new insights into the properties of magnetic alloys.

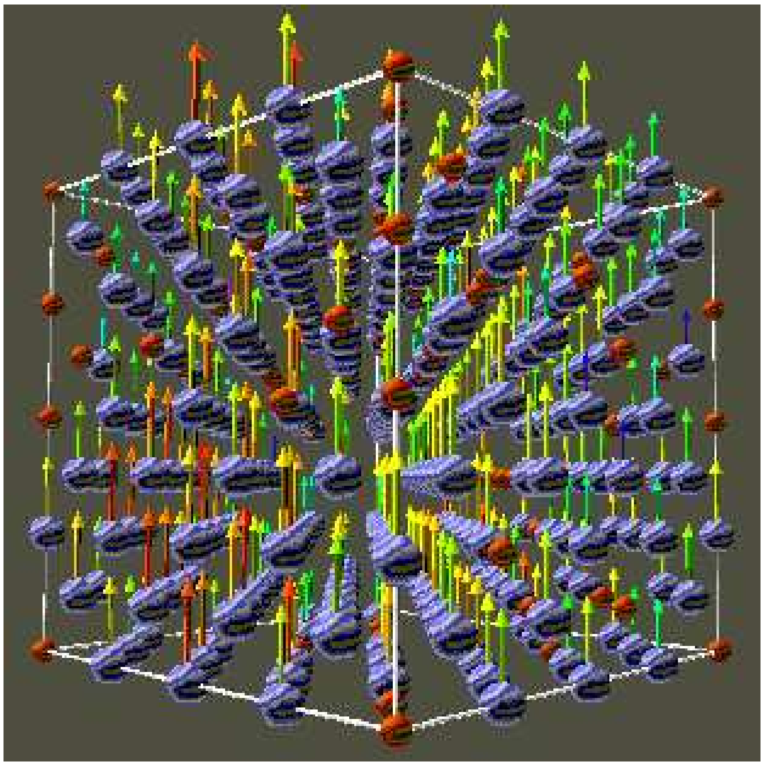

Figure 11: Magnetic state of a NiCu disordered alloy.

What is needed for Atomistic Simulations of Large-Scale Defects:

This effort has been utilizing parallel computers for a number of years, and will soon be applying Teraflop computational speeds (and Gigabytes of memory) to outstanding problems in materials science. The properties of real materials are often dominated by defects. The strength of most metals, for example, is determined by the size and properties of the boundaries between differently oriented crystallites ("grains"). The Ames Laboratory Condensed Matter Theory Group is currently studying the structure and energetics of grain boundaries in diamond (for applications to thin-film growth), silicon (with applications to semiconductor properties), and in FCC metals (for understanding the strength and ductility of metals and alloys). The determination of these structures is very difficult, due to the many possible atomic arrangements near the grain boundaries. By combining the strengths of traditional techniques (such as molecular dynamic simulations, implemented on parallel architectures) with novel "genetic algorithms" that help systematically search different possible structures, we are able to efficiently explore and locate the lowest energy structures. The genetic algorithm provides a natural bridge to massively parallel machines, with multiple structures being simultaneously optimized. This requires the capability to perform 10-100 simulations concurrently, with each simulation requiring 10,000 atoms or more.

This project was accepted for part of the High Performance Computing Challenge at Supercomputing '95. Unlike many algorithms, our approach is not limited to a single computer, or to a single set of nodes on a parallel computer. Different initial structures may be optimized on different machines, with the different machines sending back optimized atomic structures. In this way, we can use many of the largest computers in the U.S., connected via ultra-fast communication lines, to combine all of the computational power for one project. We believe that our novel approach to distributing the problem will allow us to achieve Teraflop computing speeds on a real materials problem.

3.2.3 Molecular Science in the Environment

Molecular Theory Group, Theory Modeling and Simulation,
Environmental Molecular Sciences Laboratory,
Pacific Northwest National Laboratory

Introduction:

Recent technological breakthroughs in the development of high end computer architectures[12] combined with new and novel approaches in high performance software methods, have provided the computational chemistry/physics community with the resources needed to address progressively more complex and difficult scientific problems. The presently available, and rapidly evolving, state-of-the-art hardware/software tools allows scientists to expand the scope and complexity of the systems and processes under investigation. In many cases,

[12] massively parallel processors and symmetric multiprocessors

this has allowed us to tackle "real world" problems that were previously computationally intractable. It is now becoming possible to understand how molecular level events influence important processes in the environment, materials design, biological organisms, and many other systems. These technological breakthroughs require us to expand the vision of our scientific inquiry beyond traditional boundaries and present us with a tremendously exciting and scientifically challenging environment.

Scientific inquires must keep pace with the rapidly advancing computer industry. New scientific disciplines such as molecular environmental science and molecular biochemistry are, in part, products of the recent "technological revolution." Major breakthroughs in these and related fields will be possible if we can understand the underlying molecular processes that control our world. In order to realize these goals, it is imperative that the scientific research community continues to forge strong links with computer developers and computer scientists. The crucial link between these areas is the development of high performance software that can utilize the hardware's full potential. A symbiotic relationship must exist between the theoretical chemical physics community and computer scientists and designers in order to continue to expand our scientific horizons.

Theoretical Chemical Physics at PNNL:

Rapidly evolving high performance computer technologies allow us to probe complex problems in molecular chemistry/physics at many different levels. One approach is to use and develop these resources in order to examine the properties of "large length scale" systems - e.g. (10-1000+) atom clusters, nanoparticles, liquids, solids and homogeneous and heterogeneous interfaces. Understanding the underlying physics and chemistry of these systems provides us with insight into the behavior of many environmentally and industrially important processes including, the transport of contaminants through soils, the separation of heavy metals and radionuclides from liquid hazardous wastes, the development of efficient and more cost effective catalysts, the effect of heterogeneous reaction chemistry on atmospheric and terrestrial cycles, the development of bioremediation methods. These studies employ a range of methods from ab initio quantum mechanics to classical simulation techniques. The level of theory used to study a particular system depends on the nature of the problem, the size and complexity of the system, and the availability of sufficient computational resources (both hardware and software).

An alternate approach employs state-of-the-art computational resources to study complex electronic process (e.g. basis set development, reaction rate theory, electron-electron correlation effects, many body theories) in "small length scale" molecular or atomic systems. These scientific inquires are the basis for understanding the fundamental properties of molecules and their extended assemblages, including important information about interaction energies and reaction rates. This information is essential for developing quantitatively precise models that can predict the outcome of complex industrial and environmental processes. Highly accurate theoretical models ("computer experiments") are often more cost effective and efficient than the corresponding laboratory or field studies and hence, are a fundamental component of research in fields ranging from environmental to biomedical to materials research.

Research in the Molecular Theory and Modeling program focuses on understanding the role that molecular processes play in condensed phases, characteristic of natural and contaminated environments. This research program integrates ab initio studies of fundamental molecular processes in model systems with modeling of the complex molecular systems found in the environment. This integrated approach is computationally demanding and relies heavily on the accessibility and utility of high-end computational resources.

The objectives of this research program are:

1. to apply available theoretical techniques to the study of fundamental molecular-level processes;

2. to extend current state-of-the-art methods to treat progressively more complex systems;

3. to use molecular-level data to develop models of processes occurring in large complex systems;

4. to integrate complimentary theoretical approaches for examining the chemistry and physics of multi-species, multi-phase systems; and

5. to link theory and experiment through collaborative studies.

Applications in this program focus on four overlapping research areas: *cluster chemistry*, *solution-phase chemistry*, *separations chemistry*, and *chemistry at interfaces*. A common element in each of these areas is the need to understand molecular interactions in aqueous solutions or at environmentally-important interfaces such as the interfaces between aqueous solutions and air, other liquids (e.g., organic solutions), minerals, and glasses. The development of new theoretical methods for treating these problems is concentrated in the areas of electronic structure method, reaction rate theory, and classical model development. The range and extent of these theoretical efforts are closely coupled to advances made in high-end computer architectures.

Examples from the four research areas - cluster chemistry, solution-phase chemistry, separations chemistry and interface chemistry - are described below. These examples are meant to illustrate the importance of high-end computing for solving complex problems in environmental molecular research.

Cluster Chemistry

Research in cluster chemistry is focused on the structure and properties of aqueous clusters and on the energetics and dynamics of molecular processes involving such clusters. Molecular clusters offer a unique opportunity to examine the transition from the gaseous phase to condensed phases and to enhance our understanding of the intermolecular interactions between solutes and water. In addition, to the extent that clusters model the solution phase, studies of molecular clusters provide an opportunity to ascertain the effects of solvation on chemical reactions: comparison of the pathways, energetics and dynamics of the "solvated" reaction with the gas-phase reaction will lead to insights into the detailed role that the solvent plays in altering the mechanism and rates of reactions in solution.

The methodologies developed to model gas-phase molecules and molecular processes are being applied to the study of the aqueous clusters. However, solution of the electronic and nuclear Schrödinger equation for aqueous clusters poses two challenges. First, the number of atoms involved in the calculations increases with the size of the cluster and accurate calculations rapidly become intractable. Second, the resulting potential energy surface possesses multiple minima that complicates the characterization and representation of the potential surface and subsequent dynamical studies. The primary focus of work in this area is the understanding of the structure and energetics of aqueous clusters and the development of improved interaction models for these systems.

Ab Initio Studies of $M^+(H_2O)_n$ and $M^{2+}(H_2O)_n$ Clusters

Reliable theoretical data on the structure and energetics of small water clusters containing alkali or alkaline earth ions (Li^+ through Cs^+ and Mg^{2+} through Ra^{2+}) are of interest due to the widespread appearance of such systems in nature and the difficulty of obtaining this information experimentally. Furthermore, the data is of interest because the observed binding preferences of crown ethers for particular cations in aqueous solution involves a competition between the binding of the ion to water and to 18-crown-6. The ion selectivity properties of crown ether complexes makes them unique materials for understanding enzyme chemistry and designing ion-separation techniques. This process can be modeled by calculating the energetics of the exchange reactions:

$M^+(H2O)_n$ + M'^+(18-crown-6) $\Rightarrow$ $M'^+(H2O)_n$ + M+(18-crown-6)

$M^{2+}(H2O)_n$ + M'^{2+}(18-crown-6) $\Rightarrow$ $M'^{2+}(H2O)_n$ + M^{2+}(18-crown-6)

This competition between water-ion and crown ether-ion interactions is most easily determined theoretically because the synthesis of these materials is very complex.

Geometries, binding energies, and enthalpies of molecular clusters incorporating a single metal ion and up through 8 waters were determined using basis sets and levels of theory comparable to those employed in a parallel study of ion/18-crown-6 systems. Most calculations were done with the polarized 6-31+G* basis set and second order perturbation theory. In order to estimate the inherent accuracy in the MP2/6-31+G* level of theory, some of the same clusters were re-examined with more elaborate theoretical methods. This calibration study involved the use of the correlation consistent sequence of basis sets to estimate the complete basis set limit, coupled with higher order correlation methods (e.g. MP4 and CCSD(T)). Representative results for Li^+(H2O) and Na^+(H2O) are shown in Figure 12.

Excellent agreement is observed between counterpoise (CP) corrected 6-31+G* results and the best theoretical values. Slightly poorer agreement is seen in the metal-oxygen distance because of the lack of a basis set superposition correction. In general the effects of cor-

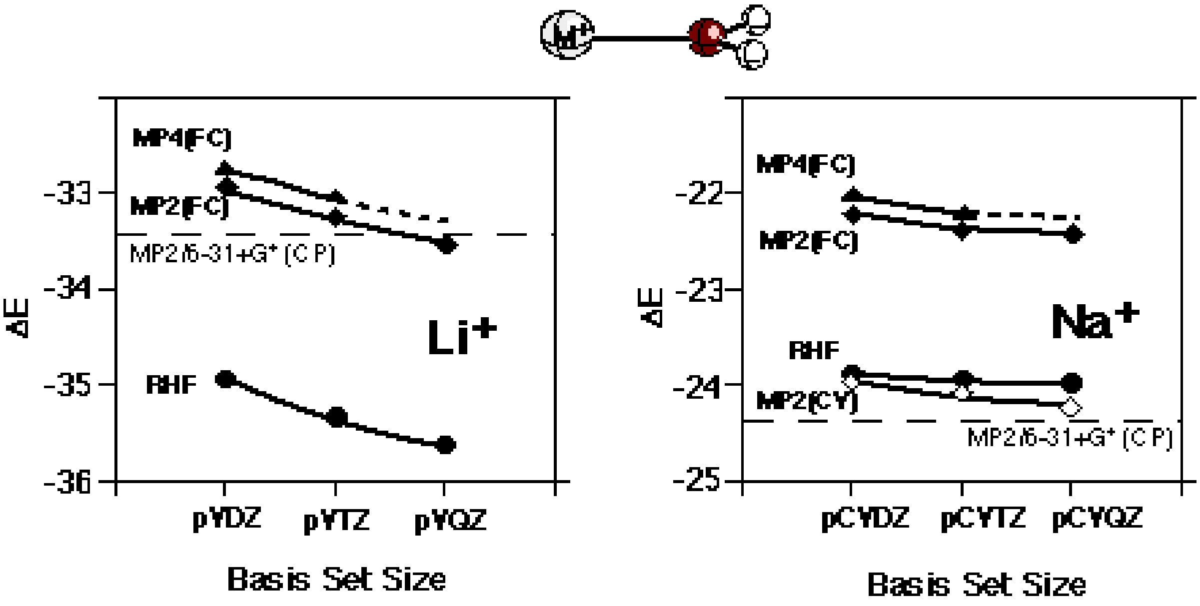

Figure 12: The binding energy (in kcal/mol) of $Li^+(H_2O)$ and $Na^+(H_2O)$ as a function of basis set size.

relation recovery on the basic cation-water distance and binding energy is small (~0.03 Åand ~2 kcal/mol, respectively). However, as the clusters increase in size and the potential for multiple water-water interactions increases, the importance of correlation effects also increase.

Due to the time consuming nature of large-basis-set, correlated calculations, we are limited in our ability to carry out the highest quality calculation to clusters containing no more than five or six waters. Such calculations require hundreds of hours on conventional supercomputers like the Cray C90. For example, a single point MP2 calculation on $K^+(H2O)_6$ using a Dunning aug-cc-pVDZ basis set (277 functions) takes about 8 CPU hours using 30 Mwords of memory and 2 GB of scratch disk space on a C90. Improved algorithms for optimizing the structure of floppy systems and computationally cheaper theoretical methods are both needed, along with new computer architectures, such as massively parallel machines. We are currently investigating density functional theory with non-local gradient corrections and the "Resolution of the Identify MP2" method.

Solution Phase Chemistry

Solution-phase chemistry plays a central role in the Molecular Theory and Modeling program. The effort is focused on the structure, energetics and dynamics of molecular processes in aqueous solutions. This work is the basis for an enhanced understanding of molecular processes occurring in contaminated and natural ground-water and in complex waste solutions. This core program in solution chemistry is designed to further the understanding of the fundamental molecular processes occurring in solution (e.g., solvation and reaction), providing a basis for the development of accurate techniques for simulating the behavior of complex waste solutions. This work includes studies of ionic solvation, supercritical fluids, and solvent effects on chemical reactions.

Molecular dynamics simulations of liquid-liquid interface with polarizable potential models

An understanding of ion transport mechanisms across the interface of two immiscible liquids is fundamental to electrochemistry and surfactant chemistry. It is also closely related to environmental problems such as separations chemistry that is performed in binary solvent systems (i.e., organic and aqueous phases) and interactions of contaminated organic solvents with ground water. We begin our effort in this area by examining the equilibrium properties of the CCl_4/H_2O liquid-liquid interface using molecular dynamics (MD) simulation techniques. It is well known that some of the properties of H_2O and CCl_4 near the CCl_4/H_2O interface are significantly different from their respective bulk properties. For instance, differences in the hydrogen bonding network and/or orientation preferred by the water molecules at the liquid/vapor interface results in a change in the dipole moment from 2.6 - 2.8 D in the bulk to 2.2 - 2.3 D at the interface. Thus, it is instructive to include the non-additive effects in describing solvation properties at the interface.

Computer simulations of liquid-liquid interfaces have been reported by Benjamin and coworkers on the structure, dynamics and conformational equilibria at the water/1,2-dichloroethane liquid-liquid interface and subsequently the mechanism of ion transport across this interface. Work has also been done on the transport of water molecules through a lipid membrane and on the water/hexane liquid-liquid interface by Berendsen and coworkers, and on the problem of solute transfer across a liquid/liquid interface by Hayoun et al. All of these studies of the liquid-liquid interfaces employ pair-wise additive models. Thus, the effects of including polarization on the liquid-liquid interfacial equilibrium properties are not known. We report here a molecular dynamics study of the CCl_4/H_2O liquid-liquid interface using polarizable potential models. This particular interface is of great environmental importance. Due to its usage in the processing of the radioactive materials, CCl_4 contributes significantly to the problem of nuclear waste remediation. A detailed investigation of the CCl_4/H_2O liquid-liquid interface at a molecular level should aid in solving these environmental problems. In this work, we are able to describe the electrical properties (i.e. the induced and the total dipole moments) of the CCl_4 and H_2O molecules as a function of their distance normal to the interface. In Figure 13, we present the computed CCl_4/H_2O interface density profile averaged over a 100 ps simulation as a function of the z axis normal to the interface.

The density profiles are calculated using slab thickness of 1Å parallel to the XY plane. After closely examining the Figure 13, we observe that the interface is not molecularly sharp, and that the density profile of H_2O indicates no layering and it is similar to that obtained for the water liquid-vapor interface using our polarizable potential model. The individual density profiles of H and O atoms are also identical to the corresponding to that of H_2O. The density profile of CCl_4, on the other hand, exhibits some oscillations. This is maybe due to the small number of CCl_4 molecules used as compared to that of water. The simulations with a significantly longer trajectory and larger system size are in progress to evaluate the

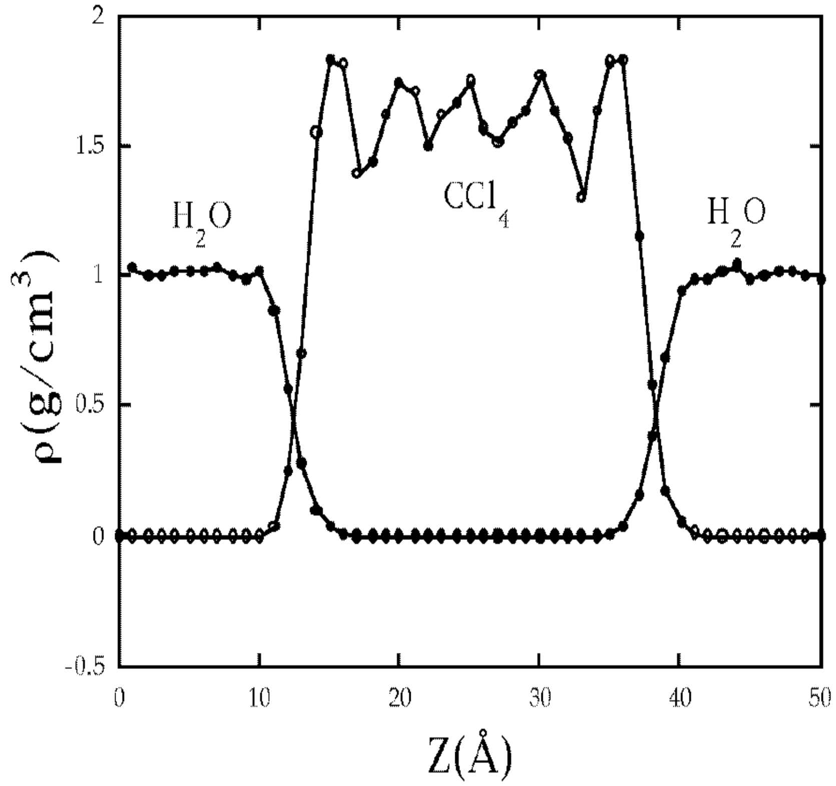

Figure 13: The density profiles of water and carbon tetrachloride at the room temperature from a 100 ps MD simulation.

origin of these oscillations. During the dynamic simulation, the interfacial fluctuations are monitored by computing the interfacial width as a function of time (the results are not shown here). We find that the width is oscillating, which indicates that there are some penetrations and retreats by water and carbon tetrachloride molecules between the interface regions.

To understand the effect of the polarization effects on the electrical properties of water molecules near the interface, we calculate the total dipole moments of the water molecules as a function of Z axis normal to the interface, and the results are shown in Figure 14.

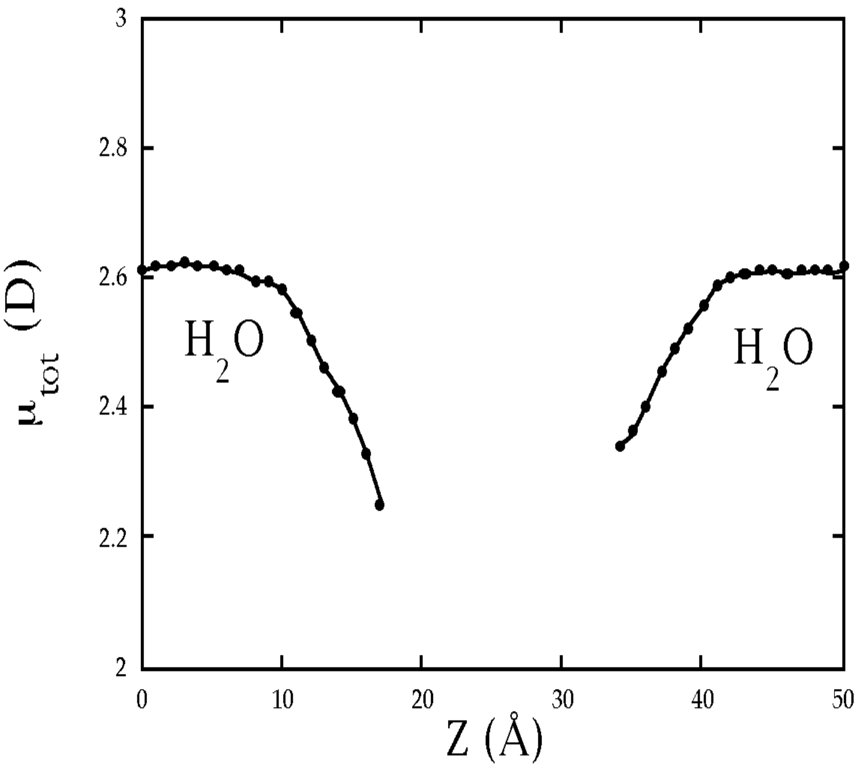

Figure 14: The total dipole moment of water molecules as a function of the Z-axis.

The dipole moments are computed using slab thickness of 1Å parallel to the XY plane. Examining this figure, several observations are in order:

1. the water molecules far from the interface (i.e., Z =1 - 10Å and 40 - 50Å) have the average dipole moment of 2.6 D, which is very close to the values obtained from bulk simulations,

2. the dipole moments of water molecules residing within the interface ($\sim$ 5Å) decrease monotonically as water approaches the interface and are close to their gas phase values.

This result is physically reasonable, and due to the changes in the hydrogen bonding patterns and/or orientational directions of water molecules near the interface. An in-depth analysis of the water hydrogen bonding and other relevant properties is underway. To further characterize the electrical properties near the interface, we show in Figure 15 the induced dipole moments of both water and carbon tetrachloride molecules.

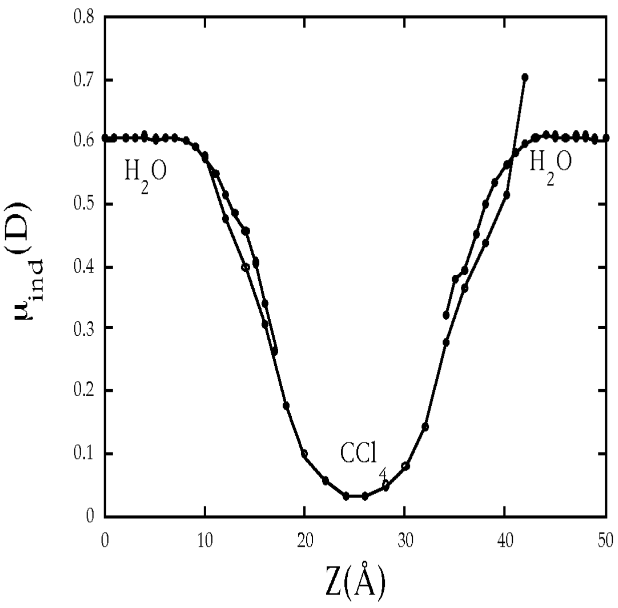

Figure 15: The induced dipole moments of water and carbon tetrachloride molecules as a function of the Z-axis normal to the interface.

As expected, the average induced dipole moment of water molecules near the interface are rather small compared to water molecules in bulk-like regions. The induced dipole moment of carbon tetrachloride molecules near the interface regions significantly increase, which may be explained by the presence of the local electric field induced by the water molecules at the interface.The static orientations of the CCl_4 - H_2O near the interface are examined via the angular distribution functions as a function of C - O separation, and are presented in Figure 16.

The angle is defined between the intramolecular C - Cl bond and the vector connecting the oxygen and carbon atom. When the C - O distance is around 3Å, the computed angular distribution functions show dramatic deviations from the sine curve that describes a uniform distribution. Two peaks centered at 72° and 170° are observed for a C - O distance less than 4Å, while the probability of finding CCl_4 in other orientations with respect to H_2O is almost vanished. Integrating over the area under these two peaks gives a ratio of magnitude of roughly 3:1.

This implies that each CCl_4 molecule has three of its chlorine atoms facing O, while the fourth chlorine atom sits on the opposite side of the central carbon atom away from O in an almost linear arrangement. This result suggests that water molecules induce a strong orientational order on the CCl_4 molecules near the interface, which is consistent

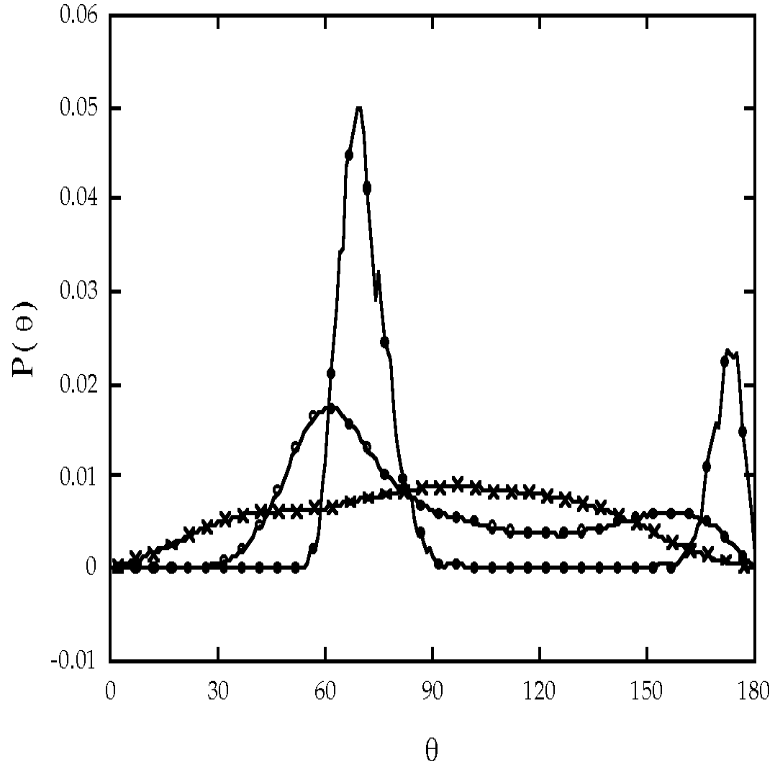

Figure 16: Probability distributions for the angle between the intramolecular C - Cl bond and the vector pointing from C to the O atom as a function of the O - C separation. The data points correspond to O - C distances between 0-3 (filled circles), 3-4 (open circles), 4-5 (crosses), respectively. A uniform angular distribution function is described by a sine function.

with the increasing of the induced dipole moment of the CCl_4 molecules. As the C - O separation becomes larger (i.e., >5Å), the angular distribution curve behaves as a normal sine function, indicating the orientational order has been smeared out and the influences of water molecules on the carbon tetrachloride molecules have diminished. To our knowledge, this work is the first MD simulation on the liquid-liquid interfacial equilibrium properties that explicitly includes non-additive effects. As expected, the interface is not molecularly sharp. The density of water indicates a uniform distribution while that corresponding for the CCl_4 exhibits some oscillations. This oscillatory behavior deserves further investigation which is in progress.

Separations Chemistry

Research in separations chemistry concentrates on the structure and energetics of ion-ligand complexes (such as crown ethers) and the dynamics of complex formation in aqueous solutions. The goal of this effort is to discover the factors that control the selectivity and efficiency of ligating species important in the pretreatment of nuclear and chemical wastes.

To date, this effort has focused on the crown ethers. Crown ethers have drawn much experimental and theoretical interest since they were first described by Pederson in 1967. These cyclic polyethers can form remarkably stable complexes with alkali- and alkaline-earth metal cations. Furthermore, the selectivity of this binding is sensitive to macrocycle ring size, ether-heteroatom, and macrocycle substituents. The ion specificity of crown-ethers makes them attractive as models for understanding enzyme chemistry as well as ligands useful to aid cleanup of mixed wastes that contain radionuclides such as cesium (Cs^+) and strontium (Sr^{2+}), and transuranic ions.

Combined Ab Initio and Experimental Studies of Small Ether/Cation Clusters

The only previous direct experimental measurement of the gas-phase binding energy of alkali metal cations and 18-crown-6 found values considerably smaller than our ab initio results. For example, their ion-cyclotron-resonance binding enthalpy for K^+:18-crown-6 at 350 K was -40 $\pm$ 8 kcal/mol, compared to an MP2 value of -72 kcal/mol. To shed light on underlying sources of error in the theoretical numbers, a joint experimental/theoretical investigation was begun on the binding energies of alkali cations with small ethers (e.g., dimethyl ether, dimethoxy ethane and 12-crown-4). While the results of that study are still in a preliminary stage, to date there have been no instances of deviations greater than $\sim$2 kcal/mol. The theoretical approach employed the same basis sets and levels of theory which were used in the earlier crown ether study. As seen in Figure 17 the counterpoise corrected MP2/6-31+G* binding energy is in remarkable agreement with the apparent complete basis set limit.

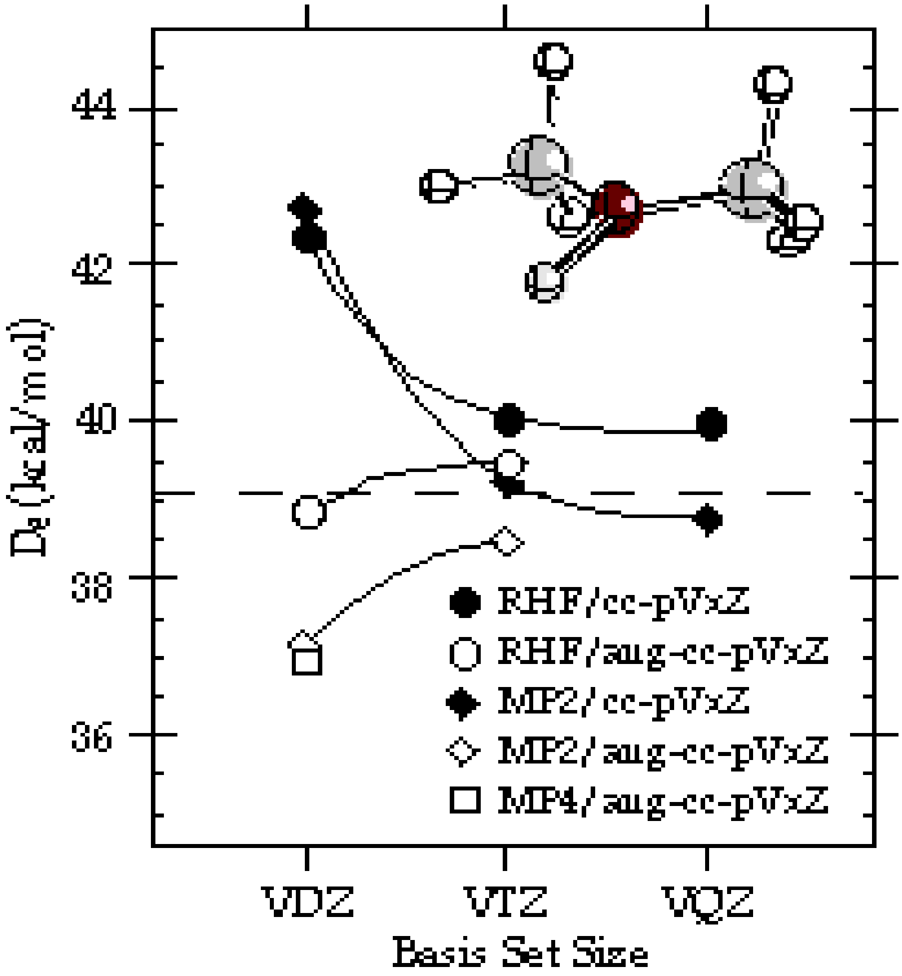

Figure 17: Convergence of the Li^+:dimethyl ether binding energy as a function of the basis set completeness and degree of correlation recovery. The dashed line represents the value obtained at the MP2/6-31+G* level, including counterpoint correction.

Mechanism and Thermodynamics of Ion Selectivity in Aqueous Solutions of 18-Crown-6 Ether: A Molecular Dynamics Study

We have performed an extensive classical molecular dynamics simulation to examine the mechanism and thermodynamics for ion selectivity in aqueous solutions of 18-crown-6 ether. 18-crown-6 is one of the most extensive studied crown systems and the abundance of the experimental data for these systems provide a good prototype model to verify our approach. To the best of our knowledge, the present work is the first application of the potential of mean force approach to the evaluation of ion selectivity by crown ethers in aqueous solutions. Our simulations differ from previous work in that the free energies of complexation are computed as a function of the ion-crown separation, or potential of mean force (PMF). This makes the current study valuable for several reasons. First, the computed PMF allows us to examine the minima positions of the complexes and they can be related directly to the x-ray data. The second valuable aspect is that we can compute the absolute binding energies by integrating over the PMFs, the resultants can be directly compared with the corresponding experimental measurements. The third advantage of our approach is that we can examine the role of solvent effects along the reaction coordinate (i.e., contact pair or solvent separated pair).

The free energy profiles or PMF for M^+:18-crown-6 ($M^+ = K^+$, Na^+, Rb^+, Cs^+) complexation are computed via the free energy perturbation method. Counterion effects are considered and all long-range interactions are computed via the Ewald summation method. The resultant PMFs indicate that free-energy minima for the K^+ and Na^+ ions are located at the crown ether center-of-mass as shown in Figure 18(a). Potassium is selected over sodium because of the greater free energy penalty associated with displacing water molecules from sodium as it approaches the crown ether. This leads to a substantial increase in the activation free energy and a decrease in the binding free energy for Na^+:18-crown-6 complexation as compared to K^+:18-crown-6. The selection of potassium over rubidium and cesium is due to the cation-size relative to the crown-ether cavity. The calculated PMFs for all the ions are shown in Figure 18(a). We computed binding free energies by integrating over the PMFs and the results are summarized in Figure 18(b). They are slightly underestimated, although, they follow the sequence $K^+ > Rb^+ > Cs^+ > Na^+$ reported experimentally. The shortcoming in binding free energies can be overcome by improving the cation-crown ether interaction potential and/or taking to account the non-additive effects in these molecular systems. Studies on these and the related questions are in progress.

(a)

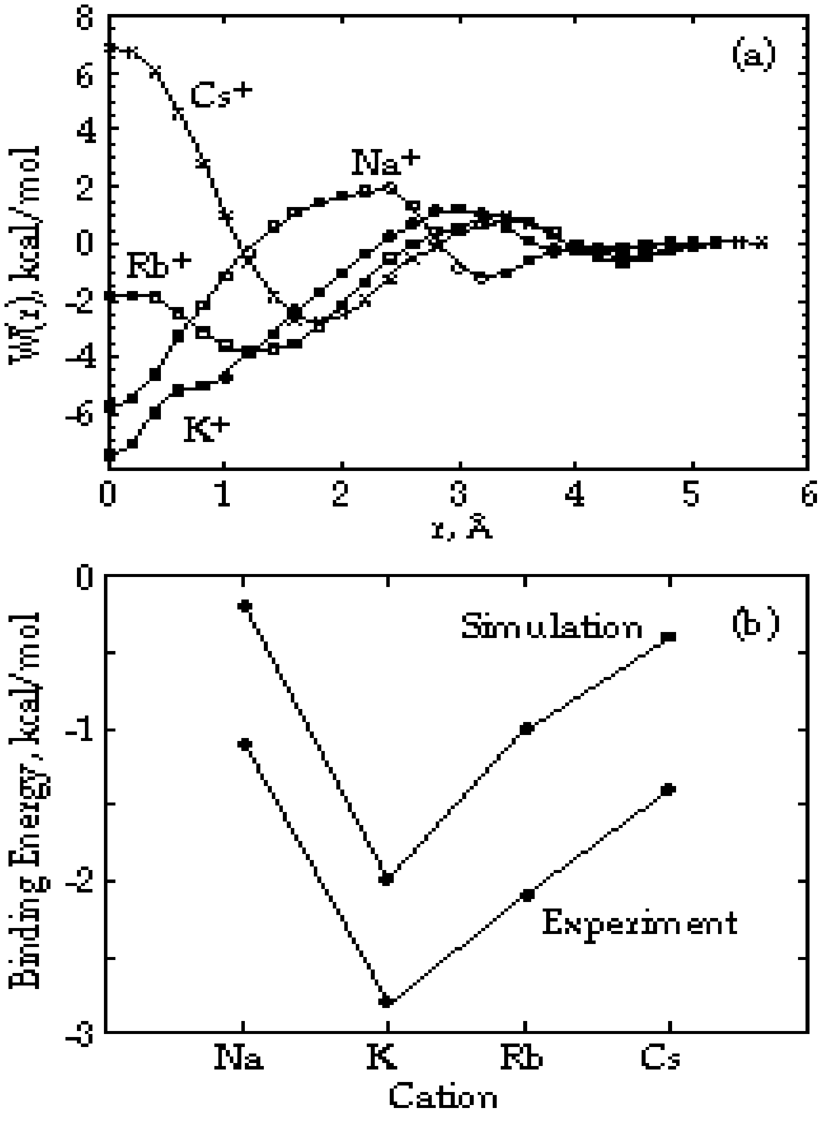

(b)

Figure 18: (a) The potential of mean force for the association of alkali cations to 18-crown-6 ether in water near 300 K. (b) Calculated binding energies of the cations to 18-crown-6 ether in water near 300 K.

Chemistry at Interfaces

Research on the chemistry at interfaces includes studies of aqueous interfaces with air (or vapor), other liquids (e.g., carbon tetrachloride), crystalline solids (e.g., mineral oxides, clays, found in soils), and glasses. The studies of air-vapor interfaces are important to understanding heterogeneous processes that occur in the atmosphere. The studies of liquid-liquid interfaces is important to understanding the transport of molecules and ions between two immiscible liquids, which is relevant in a number of field such as electrochemistry, separations, and the interaction of contaminated organic solvents with ground water. The studies of aqueous-mineral interfaces is important to understanding the binding of contaminants to the soil. The studies of glass-water interfaces provides information on the dissolution and degradation of materials that are being proposed for long term storage of nuclear wastes.

Heterogeneous and interfacial systems are exceptionally challenging systems to study theoretically. The complex chemistry and physics occurring in these systems over large length and time scales requires extensive compute power to probe and predict experimental observables. To accomplish this task many new and inventive theoretical methodologies must be used and developed. Integrated methods that exploit the accuracy of ab initio quantum mechanical theories with the speed and time-dependent properties of classical mechanics are essential for examining complex interfacial phenomena. The development of efficient methods of this type is currently an area of active research that requires high end computer platforms. Future work in this field will utilize these compute resources to address more complex, "real world" problems.

Physics and Chemistry of Structurally Simple Oxides

The chemistry of naturally occurring and/or industrially important interfaces often involves dynamical processes of structurally and electronically complex materials. In particular, it is necessary to determine the properties of realistic surface sites (containing point defects, steps, kinks) to quantitatively predict surface chemistry. Magnesium oxide (MgO) is a structurally simple cubic oxide material whose interfaces play an important role in subsurface geochemistry, catalysis and materials design. The bulk, surface and interface properties of MgO are being studied to understand the behavior of MgO and to derive models that can be used to model processes involving more complex oxides. The composition (number of electrons) and structure (high symmetry) make this a "computationally practical" system that can be studied in detail using periodic ab initio electronic structure theory.

Research in this area probed several different problems. The electronic structure investigations of the MgO/water interface were extended to examine the effect of defect sites on the surface chemistry. Water is ever-present in the environment and can have a variety of interactions with exposed oxide surfaces, each of which will have a different effect on soil chemistry and on the transport of other chemicals through soils. Therefore, understanding the processes of water adsorption and chemical dissociation on oxide surfaces is fundamental to modeling subsurface transport phenomena. Ab initio periodic Hartree Fock (PHF) theory (including correlation corrections) was used to study the chemistry of water on clean, defected and strained MgO (001).

Structural defects were constructed containing both corner (3-coordinated) and edge (4-coordinated) sites. Isotropic distortions of the lattice probed the properties of strained MgO. The ab initio energetics revealed that molecular water binds weakly (~10-15 kcal/mol) to flat MgO (100) and more strongly (17-25 kcal/mol) to the defected sites. Mulliken populations show that a fundamental change occurs in the surface oxygen sites as the coordination number of these sites decreases while the magnesium sites remain relatively unaffected by changes in their local environment. The reduction of the Madelung field surrounding the oxygen sites results in the migration of small amounts of charge from the surface oxygens to oxygens residing in the bulk material leaving the oxygen defect sites capable of forming more strongly covalent bonds than sites on the perfect surface. Chemidissociation - resulting in surface hydroxylation - is found only to occur on surface defect sites and not on the equilibrium structure flat surface. However, chemidissociation becomes more energetically favorable as the flat lattice is isotropically distorted (becoming exothermic with ~5-8% dilation of the lattice constant). This data indicates that the surface chemistry of MgO could be adjusted if stable thin films of the material could be prepared under strained conditions. The dramatically different chemistry on defected, clean and strained surfaces illustrates, quantitatively, the importance of surface structure in regulating interfacial phenomena. Analytic expressions for the water/MgO (001) interaction potential were derived from the quantum mechanical data and are currently being used in classical simulations to examine the dynamical properties of the interface.

A new methodology was also developed to derive electrostatic models of oxide interfaces. These models are constructed from the energetics of an adsorbate - described by a multipolar expansion of the charge density - interacting with the quantum mechanical electric fields generated by the surface.

$$V_{classical}(\vec{r}) = q^2\phi^1(\vec{r}) - \vec{p}^2\cdot\vec{E}^1(\vec{r}) - \frac{1}{6}Q^2_{ij}\frac{\partial E_j}{\partial r_i} + \cdots \tag{4}$$

The terms on the right hand side of this expression represent the interaction of a charged species with the surface potential, a permanent dipole with the field, and a molecular quadrapole with the field gradient. The field generated by the surface and the moments of the adsorbate are computed quantum mechanically and the adsorbate-surface interaction is deduced from the classical electrostatic interaction (eq. 4). This electrostatic model produces a reliable computationally "trivial" method for examining the properties of a range of interfaces. The results have been validated by comparing the classical data to quantum mechanical calculations on several systems. Excellent agreement was found between the computed electrostatic energetics and the quantum mechanical values for HCl, H_2O and NH_3 on MgO (see Figure 19). This approach provides us with a quantitative interpretation of the interfacial bonding interactions that allow us to predict the properties of more complex, low symmetry materials.

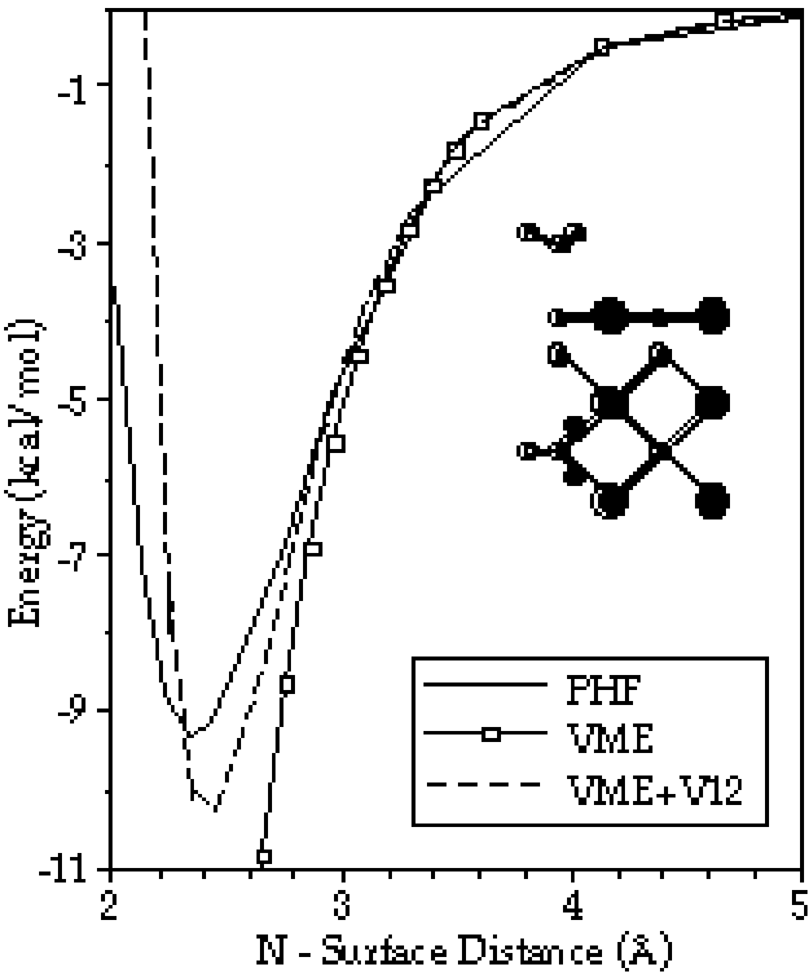

Figure 19: Comparison of ab initio and electrostatic interaction energies for NH_3 on MgO (001). NH_3 is bound over the magnesium surface site with the nitrogen end pointing towards the surface. The repulsive potential (V12) is added to the purely electrostatic interaction potential (VME) to describe the classical electrostatic binding energetics. The surface oxygens are represented by bigger radii than the surface magnesiums.

3.2.4 Atomistic Simulation of Ceramic/Metal Interfaces

R. Benedek[13] and D. N. Seidman[14], Northwestern Univ., M. Minkoff[15], ANL, and L. H. Yang[16], LLNL

Ceramic/metal interfaces are prominent in such applications as metal/matrix composites, electronic packaging, catalytic systems, and high-temperature alloy coatings. An understanding of interface structure and chemistry on an atomic scale would provide valuable guidance for optimizing performance and lifetime of the material. Owing to the complexity of ceramic/metal interfaces, their atomistic modeling has been limited. This report summarizes our recent progress in the development of realistic atomistic simulations for ceramic/metal interfaces.

Interatomic Potentials for Ceramic/Metal Interfaces: Interatomic potential models that describe either bulk metals (e.g., the embedded-atom method) or ionic crystals (e.g., the Born model, or the shell model) are reasonably accurate for many materials and have been employed to simulate grain boundaries in homogeneous (metallic or ionic) systems. The lack of a suitable model for the interatomic forces at ceramic/metal interfaces, however, has been a significant obstacle to atomistic simulations for such heterophase systems. Before such force models can be devised, the microscopic nature of the adhesion mechanism must be known. The results of ab initio local-density-functional-theory electronic structure calculations, as well as experiments, indicate that even for relatively weakly bonded ceramic/metal interfaces, such as MgO/Ag (100), chemical bonding at the interface dominates over weaker van der Waals and image-force contributions to adhesion. Our efforts have focused on the "polar" interface MgO/Cu (111), in which the ceramic consists of purely anion and purely cation layers parallel to the interface, and for which chemical bonding is even more dominant than for the MgO/Ag (100) [or MgO/Cu (100)] interface. An important implication of the dominance of chemical bonding is that forces are relatively short range, a feature advantageous to atomistic simulations.

MgO/Cu (111) Interface: MgO/Cu (111) was selected as an initial target system for our simulations for several reasons. This interface is well characterized experimentally by both atom-probe-field-ion microscopy (APFIM) and high-resolution electron microscopy (HREM) measurements. The APFIM measurements show chemically sharp interfaces between MgO precipitates and the matrix in internally oxidized Cu specimens; this means that essentially no atomic intermixing occurs between the MgO (bi)layers on one side of the interface and the Cu layers on the other. The HREM measurements show "coherent interface" regions

[13]Supported by the Office of Basic Energy Sciences, U. S. Department of Energy under grant No. DE FG02ER45403/06.

[14]Supported by the Office of Basic Energy Sciences, U. S. Department of Energy under grant No. DE FG02ER45403/06.

[15]Supported by the Mathematical, Information,and Computational Sciences Division of the Office of Computational and Technology Research, U.S. Department of Energy under contract W-31-109-ENG-38.

[16]Supported by the U.S. Department of Energy under Contract W-7405-ENG-48.

(where there is a one-to-one relation between the Cu atoms on one side of the interface and O [or Mg] atoms on the other) interrupted by misfit dislocations. Misfit dislocations are a characteristic feature of ceramic/metal interfaces and result from the mismatch between the ceramic and metal lattice constants. Since the mismatch for MgOCu (111) is large (about 14%), the misfit dislocations are relatively dense, with approximately one dislocation for every seven atoms in a row. This high dislocation density is favorable for simulation, since it reduces the cell size required to encompass the characteristic dislocation structures in molecular dynamics or Monte Carlo simulations. MgO/Cu (111) is the only orientation that is observed in the APFIM observations of internally oxidized Cu(Mg) specimens, which suggests that this orientation is thermodynamically the most stable.

Approach: The most reliable framework for the calculation of cohesive energies in condensed matter is local-density-functional theory (LDFT), within which single-electron orbitals are determined self-consistently, and then employed to calculate the total electronic density and the total internal energy. Within this framework, the predominance of chemical bonding at some prototype ceramic/metal interfaces was established. The large numerical effort in LDFT calculations, however, restricts their application to cells with at most a few dozen atoms. Even systems with relatively small spacing of misfit dislocations require hundreds (or thousands) of atoms, for which LDFT calculations would not currently be feasible. Since misfit is an important feature of ceramic/metal interfaces, we pursue a different approach that will enable atomistic simulations on a sufficiently large scale to treat misfit dislocations. Specifically, we employ LDFT total-energy calculations for selected interface configurations (without misfit) to establish a database from which effective interface potentials may be determined. The resultant interface potentials, in conjunction with appropriate potentials for the bulk metal and bulk ceramic, will then be employed in large scale molecular dynamics or Monte Carlo simulations.

Results: LDFT calculations for MgO/Cu (111) are performed for a unit cell with the Cu lattice parameter parallel to the interface stretched to the observed value for MgO, so that no misfit is present (this device is employed here for the purpose of determining the interatomic forces; however, one can also view these calculations as a representation of the coherent regions of the interface). The results presented here correspond to the oxygen "termination", which means that the (111) layer on the MgO side of the interface consists of O rather than Mg atoms. The equilibrium separation of the Cu and the O layers at the interface is that at which chemical bonding is maximized, which occurs when the total internal energy is minimized. The total energy as a function of the interface separation is referred to as the adhesive-energy curve. Besides the interface separation, two other significant degrees of freedom must be specified, namely the translation parallel to the interface of the Cu interface layer relative to the O interface layer. Thus a different adhesive-energy curve exists for each choice of the parallel translation vector. Two limiting cases are (1) the "hollow" position, in which the Cu atoms lie equidistant from three nearest-neighbor O atoms at the interface, and (2) the "on top" position where the Cu atoms are directly opposite the O atoms at the interface. Adhesive energy curves were performed for these two configurations, as well as several others, to establish the general pattern of behavior. The calculations were performed within the plane-wave-pseudopotential representation of LDFT, and the electron orbitals were obtained with a "band-by-band" conjugate gradient algorithm. Unit-cells with 9 atoms (3MgO and 3 Cu layers) were employed. A converged result for each configuration (specified by a given interface separation and parallel translation vector) required 10-20 CPU-hours on the NERSC C-90.

The results of these calculations are shown in figure 20. The inset to figure 20 shows the relative shift of the Cu interface layer relative to the O interface layer. The filled inverted triangle corresponds to the hollow site, and the open inverted triangle to the on-top position. An expected feature of the results is that at large enough interface separations, z, the adhesive energy curves essentially all coincide, and the results are insensitive to the parallel translation vector. For smaller values of z, the adhesive energy curve for the hollow position shows the largest binding and that for the on-top position the least. The adhesive energy for the hollow position can be regarded as defining a universal envelope curve, to which the others approach asymptotically at large z. This envelope curve can be fit accurately to an analytical form known as the "Universal binding energy" curve,

$$E1 = Eu \;=\; -Eo(1+a)exp(-a), \tag{5}$$

where a is linearly related to the interface separation z. The critical issue is how to describe the deviations from the universal curve as a function of the parallel translation vector. We have found that a simple analytical representation gives an excellent fit to essentially all the LDFT total energy calculations. The adhesive energy is expressed as the sum of two terms

$$E \;=\; E1 + E2, \tag{6}$$

where the additional term

$$E2 \;=\; Ae^{-Br}fc(r), \tag{7}$$

represents a summation of Born-Mayer-like short-range repulsive contributions between O and Cu atoms at separations r less than about 1.6 Å(the "ideal" bond length, corresponding to the minimum in Eu, is 2.1 Å). In (7) fc(r) is a Tersoff cutoff function, which accounts for the faster than exponential approach to the universal curve Eu exhibited by the family of curves shown in figure (20). The continuous curves in the figure represent the predictions of (6). The close agreement between the LDFT calculated energies and the analytical model is striking.

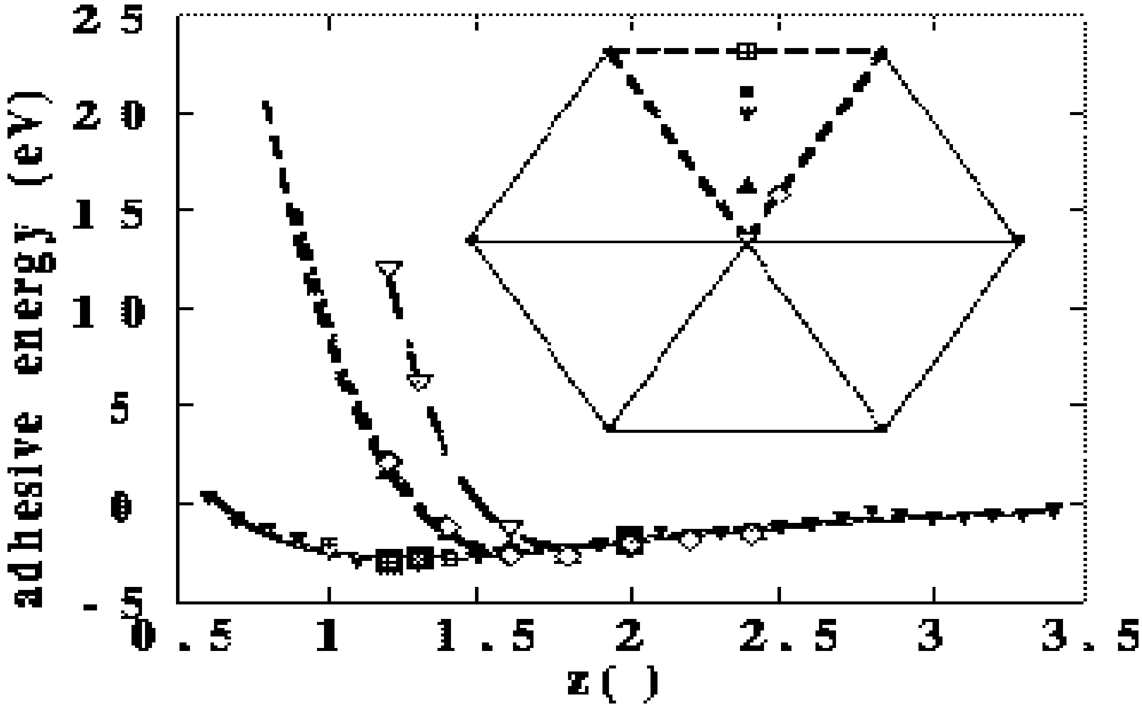

Figure 20: Results of local-density-function theory adhesive energy calculations for oxygen-terminated MgO/Cu (111) interface. The symbols represent the LDFT calculations and the continuous curves are the predictions of the model, E = E1 + E2 (see text), as a function of interface separation z. The inset shows relative position of interface Cu and O layers parallel to the interface. The triangular lattice of the (111) layers has six-fold rotational symmetry. Filled inverted triangle represents hollow site, and open inverted triangle represents on-top position.

It is perhaps surprising that the bonding (i.e., negative) contributions to E, which occur only in E1, do not depend explicitly on CuO bond lengths and angles, as one would expect for ionic or covalent bonding, but only on the interface separation z. This may signify that the interface oxygen layer has become effectively "metallized" as a result of strong hybridization with the interface Cu layer, so that bond lengths and angles are less important than total "volume" (proportional to z).

The analytical expression given by (6) provides a tractable model for large-scale atomistic simulation, by either molecular dynamics or Monte Carlo, of the MgO/Cu (111) interface. Efforts are now underway to implement (6) into an existing parallel MD code, for processing on the T3D at NERSC. In our initial calculations, the MgO layers, which are elastically much more rigid than Cu, will be frozen, and only the metal atoms will be allowed to relax. Functionally, the interface simulation then reduces to the simulation of a metal (interacting via an embedded atom model potential), subject to the boundary condition specified by (6).

Utilizing the present approach, predictions of the atomistic structure of MgO/Cu(111) will be made, including lattice parameter misfit and the consequent misfit dislocation array. Results of these simulations will be compared with HREM observations, to test the validity of the model. Upon completing these preliminary tests, we will extend the simulations to address more detailed properties, such as the interactions and segregation of solutes to metal/ceramic interfaces, and mechanical properties.

3.2.5 Large-scale Atomic Structure Calculations

Charlotte Froese Fischer
Vanderbilt University,
Nashville, TN 37235

Introduction: Atomic data is needed in many areas of scientific endeavor, the development of controlled thermonuclear fusion being one example. Energy level structure, autoionization rates, cross-sections for excitation, ionization, charge exchange, and recombination all enter into model calculations. In geophysical studies of the atmosphere, for example, the emission features of atomic oxygen associated with branching ratios for decay from $2p^3 3s\ ^3S$ state are related to the abundance of oxygen in the thermosphere. Astrophysics has a longstanding need for large amounts of data. A successful "opacity" project for generating such data has produced vast amounts of data, but all calculations were all performed in the LS approximation and do not include relativistic effects or nuclear effects. Heavy elements also are gaining in importance, not only as a test of theories: data for rare earth elements, for example, are needed for the development of high intensity discharge lamps, and data for other elements are relevant to environmental clean-up problems.

Atomic structure is an area where accuracy is extremely important. In spectroscopy, observations are related to energy differences. It is these differences that need to be predicted to spectroscopic accuracy, which, in many applications, may be defined as a fraction of a

cm^{-1}. This translates to an error of less than 4.5×10^{-6} au. in an energy difference. Even for relatively small systems, this requirement may present a considerable challenge. A number of physical effects may need to be included to achieve the desired accuracy.

Achievements: The inclusion of correlation in the motion of electrons in many-electron systems is computationally challenging. In light atoms, using a non-relativistic formalism together with a parallel version [17] of our atomic structure package, called MCHF a full-correlation study is feasible, put the magnitude of the problem increases exponentially with the number of electrons. But many atomic processes, such as transition probabilities, are outer-electron phenomena. Through the use of systematic methods within a series of models, we are able to determine properties along with estimates of uncertainties. This is illustrated in figure 21 where results are presented for an intercombination line in C III that has been of great interest in astrophysics. Note the final error bar for the calculation.

New Opportunities: As accuracy is increased, or heavier atomic systems are considered, it becomes desirable to use a fully relativistic formalism that includes the effect of the finite size of the nucleus. In heavy atoms the latter is an extremely important effect. In our work to date, we have modified the General Relativistic Atomic Structure Package (GRASP) developed at Oxford under the direction of I.P. Grant. In particular, dynamic memory allocation has been included along with an eigenvalue solver relying only on matrix vector multiplication so that sparse matrix methods may be used. Though the methodology of this code is quite different, its structure is quite similar to that of MCHF. In fact, the recently published GRASP92 [18], has been modularized in a similar fashion. We believe this code should be modified for massively parallel computing, as supported for the Cray T3D.

With this implementation, our methodology could be extended to heavier systems. In such systems relativistic effects and correlation cannot be separated. Rapid progress could also me make on a collaboration with G. Malli for the study the energy levels structures of the transactinides where relativistic correlation plays an extremely important role.

[17] A. Stathopoulos, A. Ynnerman, and C. Froese Fischer, "A PVM implementation of the MCHF atomic structure package", *Int'l J. Supercomputer Appl. and High Performance Computing* (in press).

[18] F.A. Parpia, C. Froese Fischer, and I.P. Grant, "GRASP92: A package for large-scale relativistic atomic structure calculations", *Comput. Phys. Commun.* (in Press).

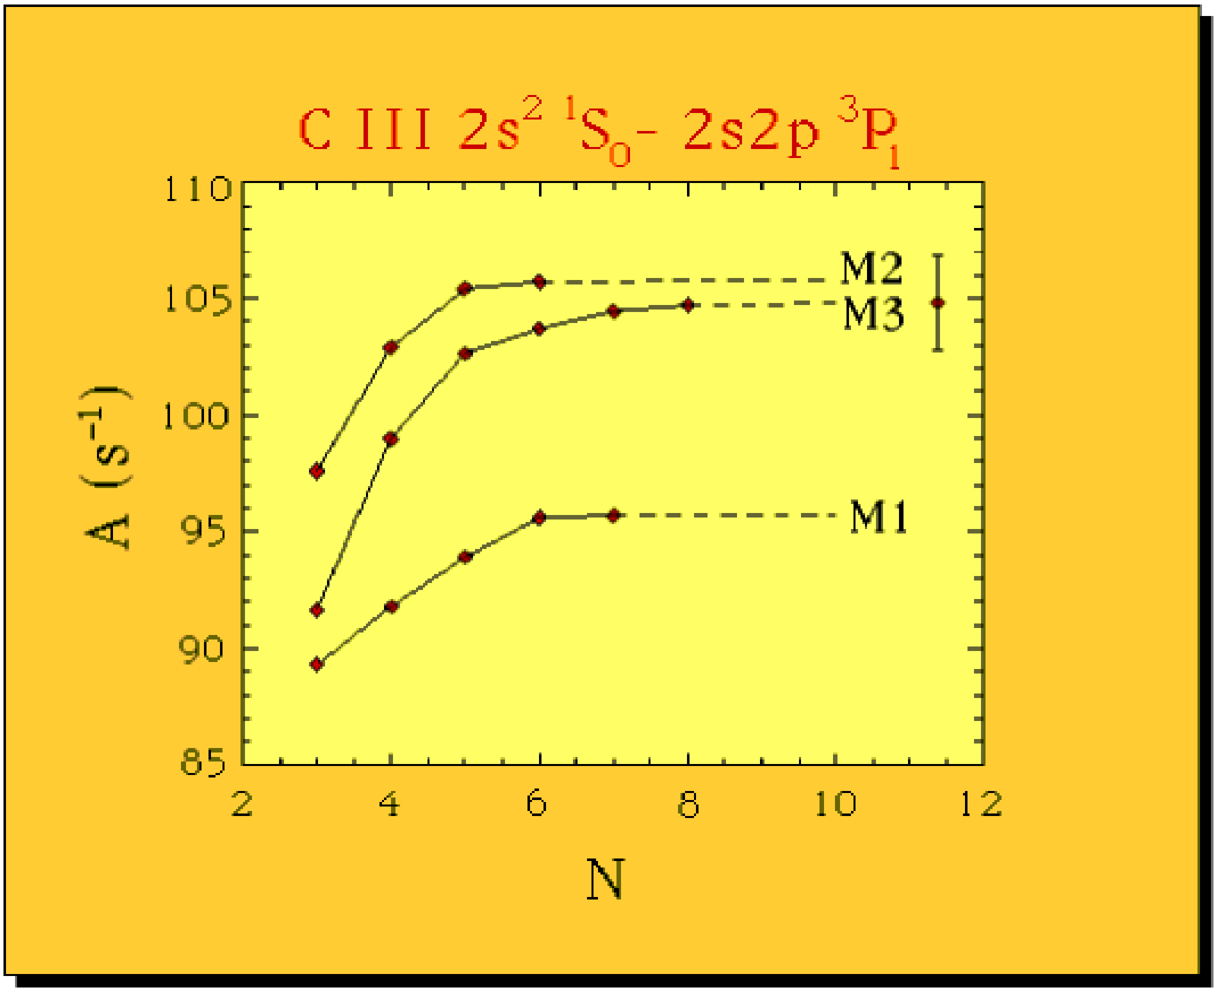

Figure 21: Plot showing the transition rate trend with respect to increasing size of the calculation within models (M1, M2, M3) including more and more correlation.

3.3 Fusion Energy

The mission of the U.S. Fusion Energy Science Program is to acquire the knowledge base needed for an economically and environmentally attractive fusion energy source.

Overview of Fusion and Plasma Science Computing:

B. Braams (NYU), V. Chan (GA), W. Houlberg (ORNL), S. C. Jardin (PPPL), C. Karney (PPPL), W. Nevins (LLNL), A. Turnbull (GA) D.Schnack (SAIC)

Plasma, sometimes called the 4th state of matter, is a gas of such high temperature that the atoms have Dis-associated into their component parts of electrons and charged nuclei. Plasmas play a key role in controlled fusion and in space physics, as well as in other applications and disciplines such as material processing, particle accelerators, and astrophysics. Many of these applications involve plasmas in strong electric and magnetic fields. To simulate (or model) the behavior of plasmas, a system of mathematical equations and boundary conditions is formulated that describes the plasma and the electromagnetic field. Different mathematical models arise from different approximation techniques and hence have different regions of validity. For example the ideal MHD equations provide a simplified set of equations suitable for describing a wide range of global plasma instabilities, but these completely neglect all dissipative effects, which are not important for this class of global, fast time scale instabilities described well by ideal MHD. Other classification schemes sometimes used are that of microscopic and macroscopic models, which are appropriate for describing the small-scale and the large-scale features, respectively, and that of classification according to the typical time scales of the plasma phenomena being addressed.

As an illustration of the richness of this field, we show in figure 22 a range of time scales spanning fifteen orders of magnitude that are of relevance to a modern magnetic fusion experiment. On that figure, we also indicate many of the equation sets that are appropriate to describe phenomena occurring on those time scales, and some of the computational plasma models that have been developed to address these. These models include sets of linear and nonlinear equations, and both fluid and particle-in-cell based representations. Within each time scale category, plasma models may also be distinguished by their physical and corresponding mathematical complexity, and by their need to be supplemented with experimentally measured parameters. Because of the complexity of the physical model, the geometry, the boundary conditions, and/or the nonlinear phenomena, these equations can only be solved on the most powerful computers.

During the last few decades, remarkable progress has been made in measuring, understanding and in predicting many aspects of plasma behavior in the laboratory as well as in space. Many simulation models, such as those referred to here, have been developed that accurately describe many facets of plasma behavior. In magnetized fusion physics, the tokamak configuration has emerged as the leading candidate to serve as the embodiment of the first fusion-powered electrical power plant. Because of it's importance to the program, an extensive system of tokamak specific simulation and analysis codes have been developed by the fusion community. We illustrate this system, which consists of over 50 separate stand-alone codes in figure 23. A typical project to optimize the design of a new fusion device would use most if not all of these codes to define the configuration, analyze its stability and confinement properties, develop heating and current-sustainment methods, and analyze potentially destructive failure modes.

The system of codes shown in this figure illustrate the impressive achievements that have been made, but they also hint at some of the needs for improvement. Present emphasis is on extending the range of validity of existing simulation models so that realistic fusion parameters can be addressed, and also in developing models that span wider ranges in time and space scales for more comprehensive simulations of complex phenomena. This direction would eventually lead to a complete physics based predictive simulation model of fusion and space plasma behavior, and there is reason to believe that such comprehensive simulation capabilities are now within reach. The creation of such fundamental simulation models would have dramatic implications for fusion, as well as for related fields such as space sciences.

In fusion, more comprehensive simulation models would provide better guidance for the design and operation of experiments. In this new paradigm, experiment and theory will work in tandem in the validation of the fundamental principles, thereby providing the high scientific confidence required to design and construct the next generation of experimental devices. A credible-physics-based simulation model also holds forth the promise of enabling the discovery of a breakthrough advance in fusion science. An example would be one that would radically increase the obtainable fusion core power density and hence the economic attractiveness of a fusion power plant.

It now appears that a complete physics-based predictive simulation model could result from an appropriate combination and extension of 3D simulation models that now exist. This integrated modeling would involve solving the appropriate equations on both the microscopic and macroscopic spatial and temporal scales and developing efficient algorithms for coupling the simulations. On the microscopic scales, the Numerical Tokamak Turbulence Grand Challenge has produced near first-principles 3D gyrokinetic and gyrofluid simulation models that are now making realistic predictions of the turbulent transport of plasma and energy in localized regions and in simplified geometry. These models need to be extended by increasing the order of the fundamental mathematical equations and by increasing the generality and the spatial extent of the geometry they are applied to, including extending them to the plasma edge. In the integrated modeling, the geometry for the microscopic model would be defined by the macroscopic model that would be solved concurrently. The need for effective coupling algorithms can be illustrated by a simple example that shows 3D micro-turbulence calculations cannot be applied directly to macroscopic problems in the near future. A 10 msec full torus simulation with 36,000 modes would require 50 Terabytes of fast memory, generate 27 Petabytes of data, and take 80 hours of CPU time on a Teraflop computer.

On the macroscopic scale, great progress has been made in solving the 3D magneto-hydrodynamic (MHD) equations accurately, including the effects of plasma resistivity. As viscous hydrodynamics is characterized by the Reynolds number R, resistive MHD is characterized by a dissipation parameter S called the magnetic Reynolds number. Present MHD simulations are generally limited to S values of 10^4, whereas values of 10^8 would be more typical of fusion plasmas. Also, today's macroscopic codes normally utilize relatively simple idealized boundary conditions, and need to be supplied with empirical "transport coefficients" to describe the sub-grid scale turbulent phenomena. The implementation of more realistic boundary conditions is relatively straightforward and would open up a whole host of new capabilities. The integrated modeling would in effect self- consistently determine these transport coefficients from the microscopic simulation model.

Other major modeling and simulation issues include the development of optimization methods to design compact helical or toroidally symmetric plasmas that simultaneously have good confinement, MHD pressure limits, and position control characteristics with a realizable set of external coils and acceptable plasma access. Several different approaches to this design problem have begun in the last few years and are leading to a new generation of proposals for advanced confinement experiments that could significantly enhance the attractiveness of a fusion power plant.

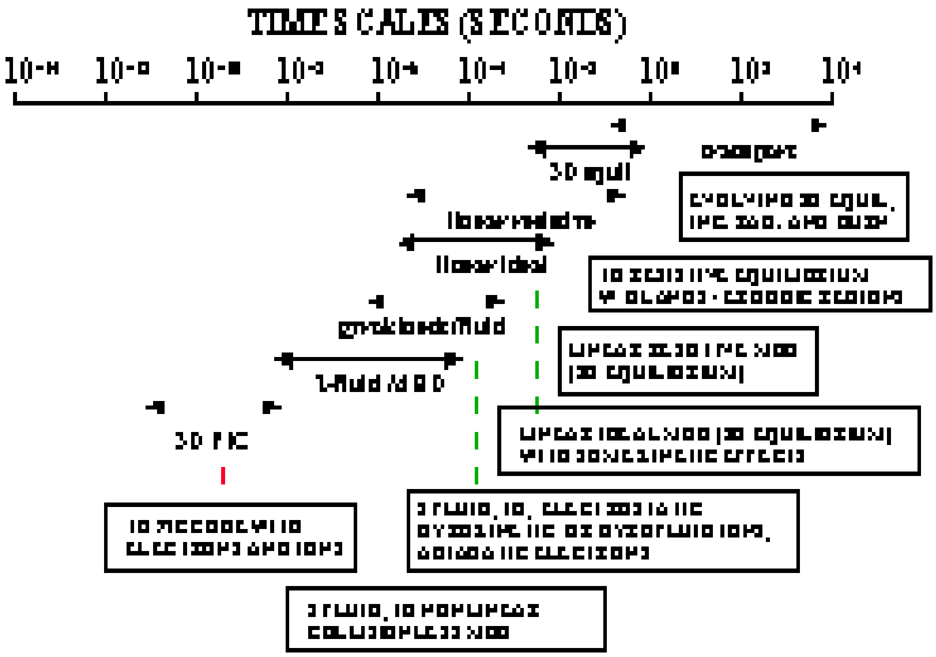

Figure 22: Time scales covered by some of the major plasma physics codes as applied to a major magnetic fusion experimental device.

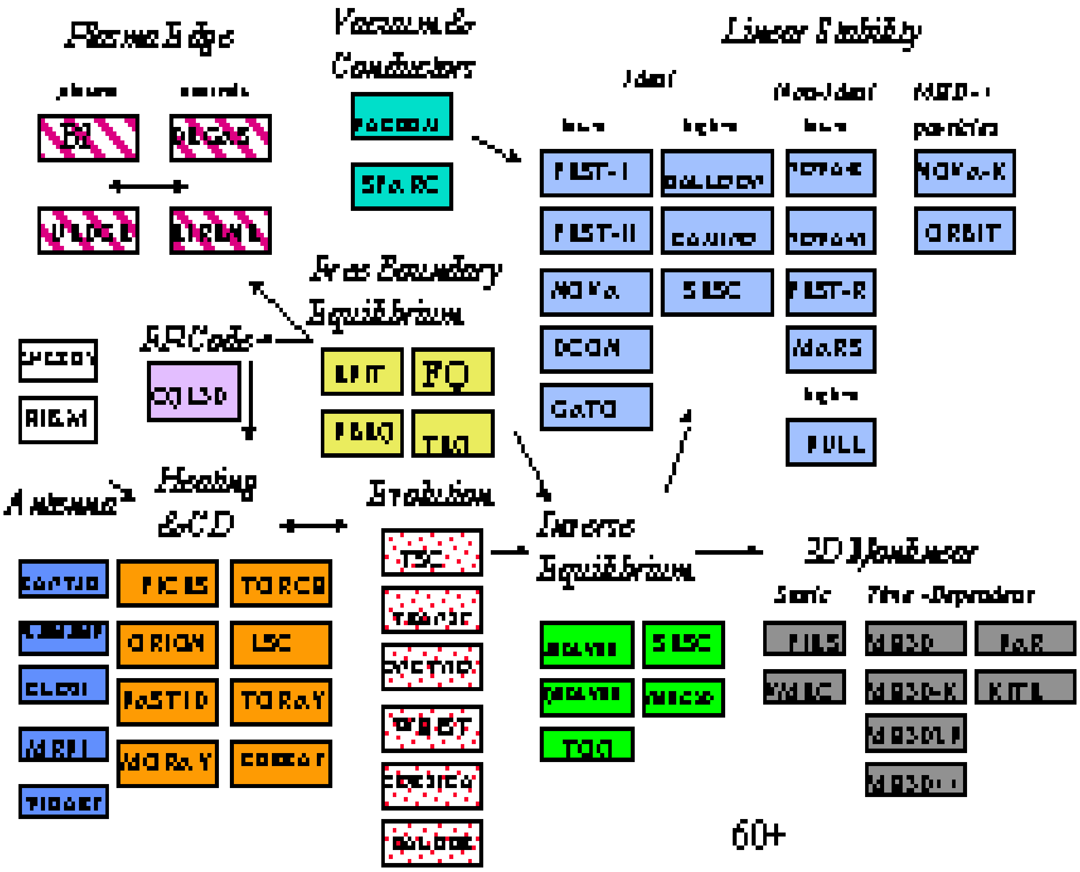

Figure 23: Some of the most widely used codes for Tokamak Physics Design and Analysis Computation (TOPDAC)

3.3.1 3d MHD Calculations in Toroidal Geometry

S. C. Jardin, D. Monticello, W. Park and A. Reiman
Princeton Plasma Physics Laboratory
Princeton NJ 08543

High-β Disruption Modeling: In the record-setting deuterium-tritium discharges in TFTR, performance is limited by magnetohydrodynamic (MHD) activity, the most dangerous type being the major disruption. This event causes the loss of the confined thermal energy during the thermal quench phase and a subsequent loss of plasma current during the current quench phase. Ballooning modes are expected in plasmas with pressure gradients exceeding a threshold set by the structure of the confining magnetic field. Recently, ballooning modes were clearly observed on TFTR for the first time.[19]

The first dynamic calculation of a non-axisymmetric ballooning instability was performed at PPPL[20] using the 3D nonlinear MHD initial value code, MH3D. A 2D equilibrium and its low-n linear ballooning mode were first calculated using the experimental pressure and current profiles prior to the onset of a disruption. Nonlinear evolution of the mode results in a local pressure steepening, which in turn drives a linear toroidally localized high-n ballooning mode. Nonlinearly, the mode becomes even more localized producing a local pressure bulge which destroys the flux surfaces, resulting in a thermal quench. Figure 24(a) shows the simulated Electron Cyclotron Emission (ECE) signal obtained. The calculated signals are in good agreement with the experimental ECE data shown in Fig. 24(b).

Three-Dimensional Equilibrium: Stellarator configurations, which are receiving substantial financial support and construction money in Japan and Europe, are intrinsically non-axisymmetric three-dimensional plasma confinement systems. Because of the lack of symmetry, good flux surfaces (KAM surfaces of the magnetic field line trajectories) are not guaranteed. Non-integrable field line trajectories may wander through a volume, and cause loss of confinement. The magnetic fields produced by stellarator coils are designed to minimize this effect for a low pressure plasma, but the self-consistent plasma currents that arise with increasing pressure will in general degrade the flux surfaces. The three-dimensional equilibrium PIES code[21] has been used in several studies of this effect, which have aimed to understand the underlying physics, and to assess the significance of the effect for several proposed stellarator devices. PIES is one of only two codes in the world that are capable of addressing this key physics issue, and PIES is the only such code that can handle net plasma currents, and is also the only such code that can realistically specify the equilibrium through the currents in the external coils. Figure 25 shows the output of the PIES code for a reactor proposal, MHH, based on the Stellarator concept. At $<\beta>= 2.35\%$, Fig. 25, shows magnetic islands developing over much of the reactor's cross-section.

---
[19] E. D. Fredrickson, et al, *Phys. of Plasmas* **2**,4216 (1995)
[20] W. Park, et al, *Phys. Rev. Lett.* **75**,1763 (1995)
[21] A. Reiman and H. Greenside, *Comp. Phys. Com.* **43** (1986)

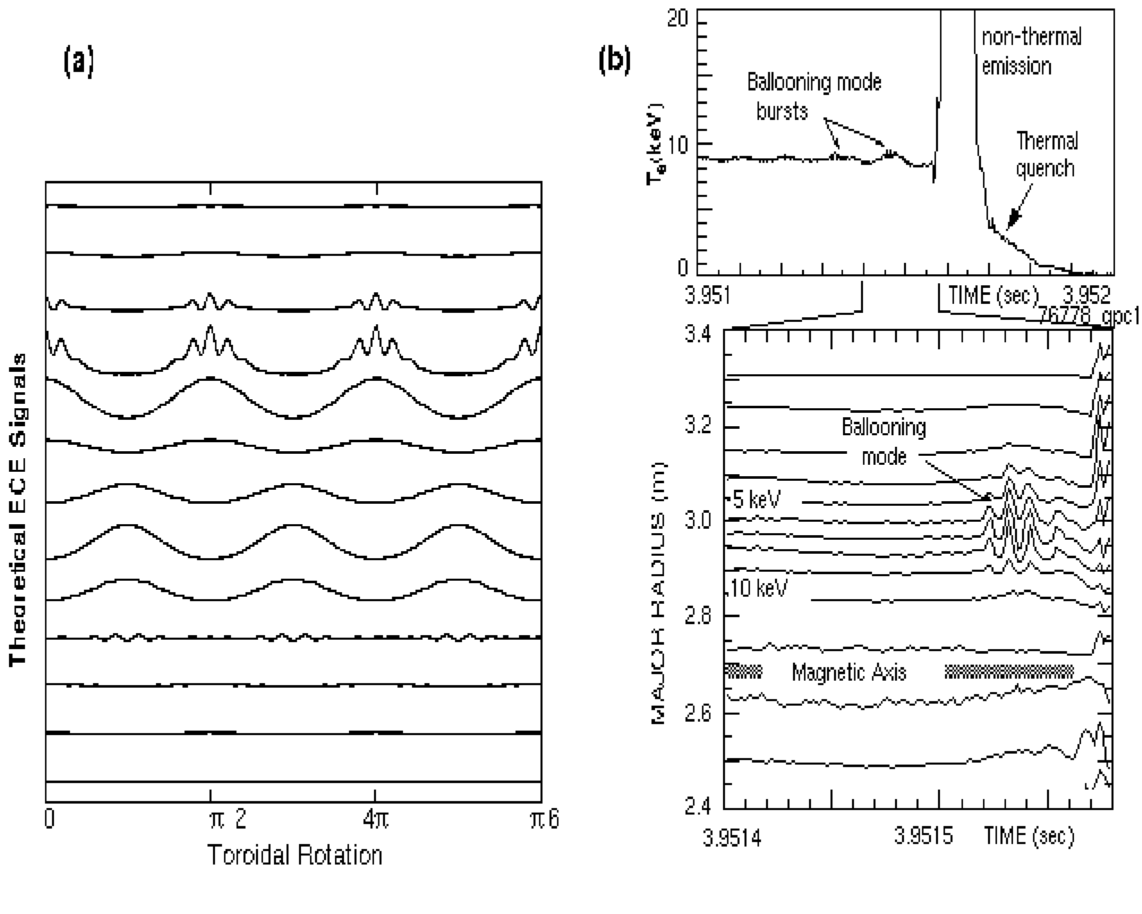

Figure 24: ECE signals; (a)theoretical and (b)experimental.

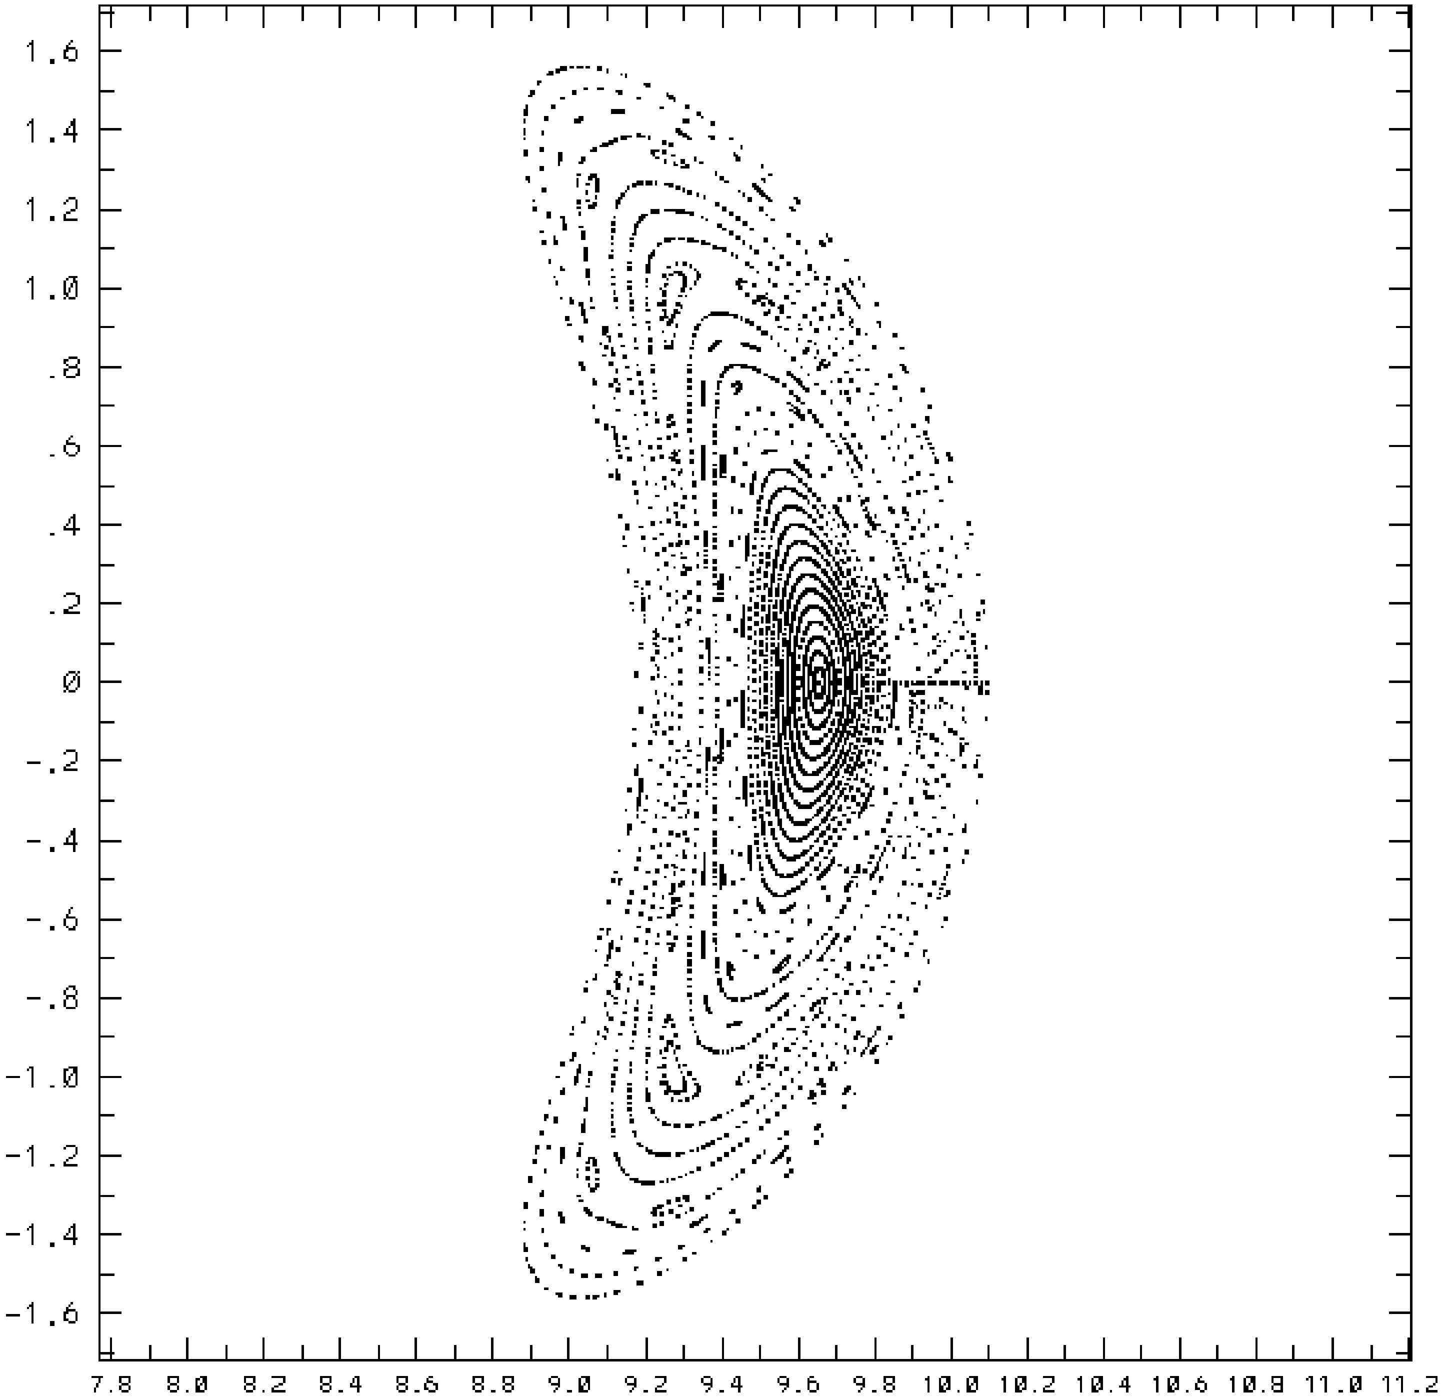

Figure 25: Poincare plot of magnetic field for the MHH Stellarator

3.3.2 Numerical Tokamak Turbulence Grand Challenge Project

B.I. Cohen, Principal Investigator, Lawrence Livermore National Laboratory
J.M. Dawson, Scientific Leader, Univ. of California at Los Angeles
J.V.W. Reynders, Co-principal Investigator for Infrastructure, Los Alamos National Laboratory
V.K. Decyk, University of California at Los Angeles
W.D. Dorland, University of Texas at Austin
G.W. Hammett, Princeton Plasma Physics Laboratory
G.D. Kerbel, Lawrence Livermore National Laboratory
J.N. Leboeuf, Oak Ridge National Laboratory
W.W. Lee, Princeton Plasma Physics Laboratory
S.E. Parker, University of Colorado, Boulder
R.E. Waltz, General Atomics

Research Objectives: The primary research objective of the Numerical Tokamak Turbulence Project (NTTP) is to develop a predictive ability in modeling turbulent transport due to drift-type instabilities in the core of tokamak fusion experiments through the use of three-dimensional kinetic and fluid simulations and the derivation of reduced models.

Computational Approach: Two classes of three-dimensional initial-value simulation algorithms, gyrokinetic (GK) and gyro-Landau-fluid (GLF), are being applied to the simulation of tokamak turbulent core transport. The GK simulations are based on particle-in-cell (PIC) methods for the self-consistent solution of Poisson's equation (reduced to a quasi-neutrality relation) and plasma equations of motion, and the use of domain decomposition methods to run efficiently in parallel on the T3E and other parallel computers. The GLF algorithm is based on an alternative solution of the fundamental GK and quasi-neutrality equations in which fluid moment equations are solved instead of particle equations. The GLF simulations have been performed on both massively parallel computers and parallel vector computers, particularly the C90 at NERSC. Both flux-tube, i.e., toroidal annulus, and global toroidal GK and GLF simulations are being performed to study tokamak turbulence.

Accomplishments: Because of its relevance to the International Thermonuclear Experimental Reactor (ITER), detailed parameter studies have been conducted with our GK and GLF simulations addressing discharge #81499 in the General Atomics DIII-D tokamak as a base case. Careful comparisons have been made between the results of the GLF and GK simulations of shot 81499 and variants to determine parametric dependences and points of agreement and disagreement between the simulation algorithms (the Cyclone study). Elsewhere GLF and GK simulation results have been compared to experimental findings with good agreement in some cases. As part of the Cyclone study, a GK flux-tube PIC simulation has been the object of a vigorous study of its convergence properties with results so far suggesting that statistical convergence is being obtained. Although the differences between the results for the turbulent ion thermal diffusivity in GK and GLF simulations are not large (GK results are often a factor of 2 to 3 lower than GLF results), the implications of these differences for the projected fusion performance of ITER can be significant. The NTTP simulations have also demonstrated the importance of flow shear and negative central magnetic shear in reducing drift-wave turbulence in tokamaks as observed in experiments. GK simulations have verified the improved confinement in the TEXTOR tokamak due to a small percentage of Ne impurity which has a stabilizing effect on ion-temperature-gradient turbulence. Many NTTP codes have been successfully ported to the T3E in the last year, and a new global GLF code is yielding physics results (c.f., 26 and 27). Progress has been made in including realistic geometrical effects (for example, shaped cross sections and magnetic X-points needed to model plasmas with divertors) and realistic equilibrium plasma profiles in the simulations.

Significance: The NTTP simulations are being used to produce linear and nonlinear calculations of drift-type instabilities in realistic tokamak equilibria, which are leading to a deeper understanding of anomalous transport in current experiments and to improving their performance. This simulation work builds a bridge between theory and experiment, and is providing a basis for reduced transport models intended to fit current experimental databases and from which it is hoped that performance in future experiments can be reliably predicted and optimized. As controlling the energy transport has significant leverage on the performance, size, and cost of fusion experiments like ITER, reliable NTTP simulations can lead to significant cost savings and improved performance in future experiments.

Recent Publications:

- "Stabilization of Ion-Temperature-Gradient-Driven Tokamak Modes by Magnetic-Field Gradient Reversal," M. Fivas, T. Tran, K. Appert, J. Vaclavik and S. Parker, Phys. Rev. Lett. 78, 3471 (1997).
- "Implicit, partially linearized, electromagnetic particle simulation of plasma drift-wave turbulence," B.I. Cohen and A.M. Dimits, Phys. Rev. E 56, 2151 (1997).
- "Nonlinear Kinetic-Fluid Equations," N. Mattor and S.E. Parker, to appear November 1997 in Physical Review Letters.
- "Simulating Plasma Turbulence in Tokamaks," J. Kepner, S. Parker and V. Decyk, SIAM News 30, 1 (1997).
- "A gyro-Landau-fluid transport model", R. E. Waltz, G.M. Staebler, W. Dorland, G.W. Hammett, M. Kotschenreuther, and J.A. Konings, Phys. Plasmas 4, 2482 (1997).
- "Plasma turbulence calculations on the CRAY T3E", V. E. Lynch, J. N. Leboeuf, B. A. Carreras, J. D. Alvarez, and L. Garcia, Supercomputing '97, San Jose, California, November 15-21, 1997. http://www.supercomp.org/sc97/proceedings/TECH/LYNCH/INDEX.HTM
- "Theory and simulation of rotational shear stabilization of turbulence", R.E. Waltz, R.L.Dewar, and X. Garbet, '97 APS invited talk to appear in Phys. of Plasmas.

Other information at these WWW URLs:

- Numerical Tokamak Turbulence Project
  http://www.acl.lanl.gov/GrandChal/Tok/tokamak.html
- Cyclone
  http://www.er.doe.gov/production/cyclone/

Summary of Computing Resources Requested by the Numerical Tokamak Grand Challenge

		FY97	FY98	FY99
NERSC T3E	T3E node hours (10^5)	2.5	3.5	4.5
	J-90 node hours (10^5)	0.04	0.04	0.04
	memory (GB per job)	1-60	1-60	1-60
	disk (TB)	0.2	0.2	0.2
	tertiary storage (TB)	6	6	6
	Onyx hours/yr	200	200	200
ACL Cluster of SMPs	node hours (10^5)	3	7	9
	memory (GB hr per job)	10^3	$2x10^3$	$3x10^3$
	disk (TB)	0.3	0.5	1.5
	tertiary storage (TB)	6	6	6
	Onyx hours/yr	200	200	200
Paragons at ORNL	processor hours(10^5)	0.055	0.15	0.3

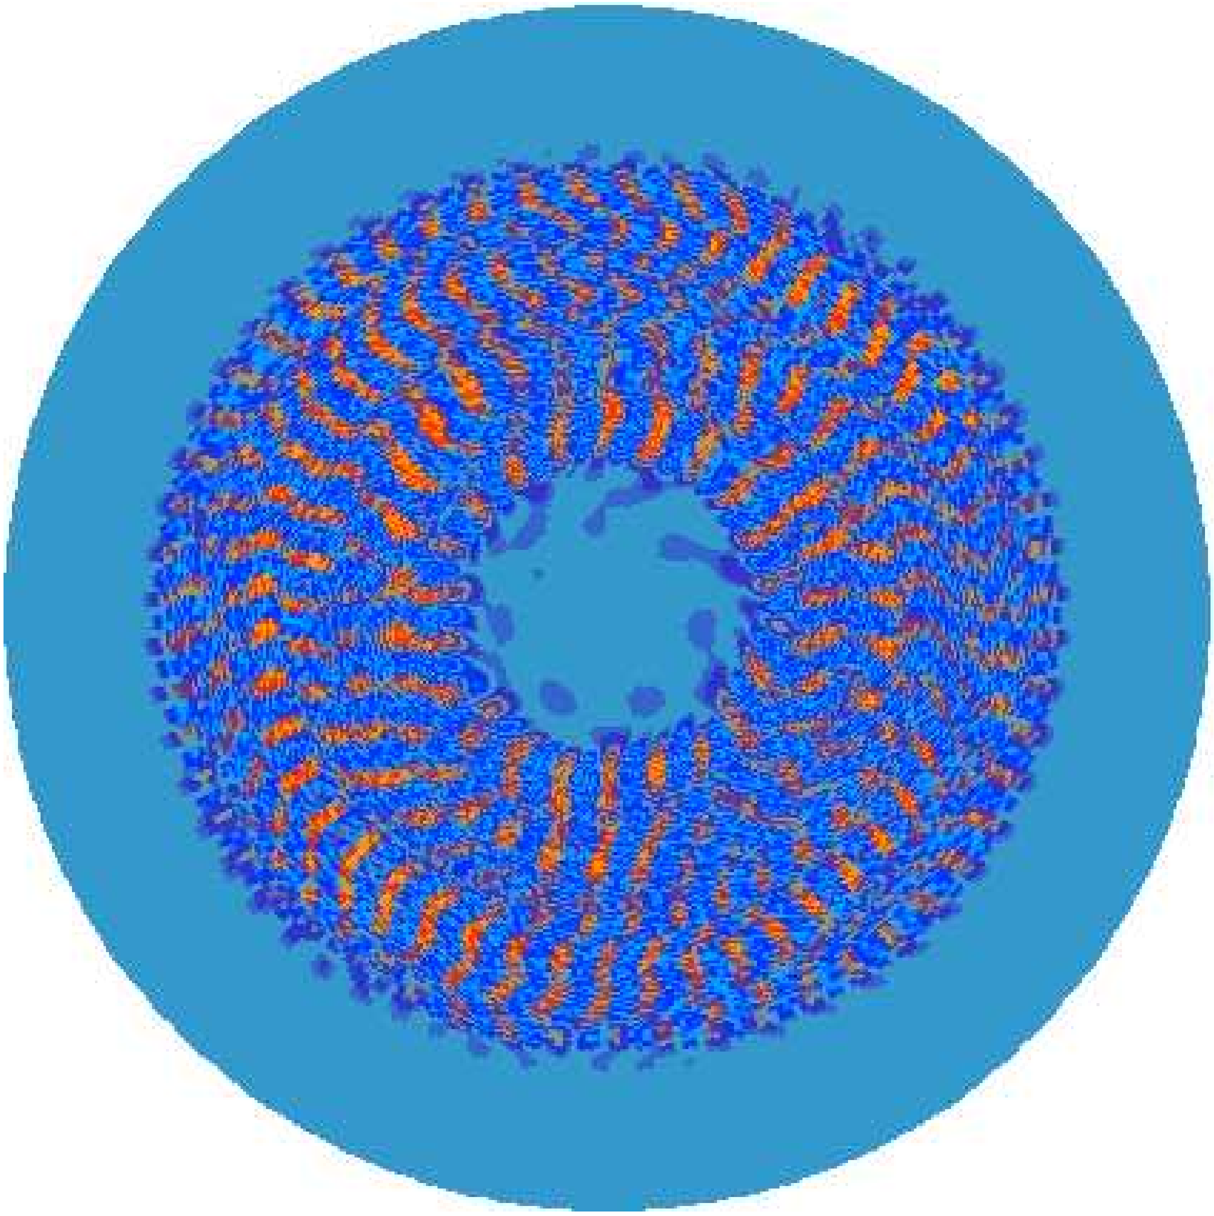

Figure 26: Full-cross-section simulation of ion-temperature-gradient-driven turbulence. V. Lynch and J.-N. Leboeuf, ORNL

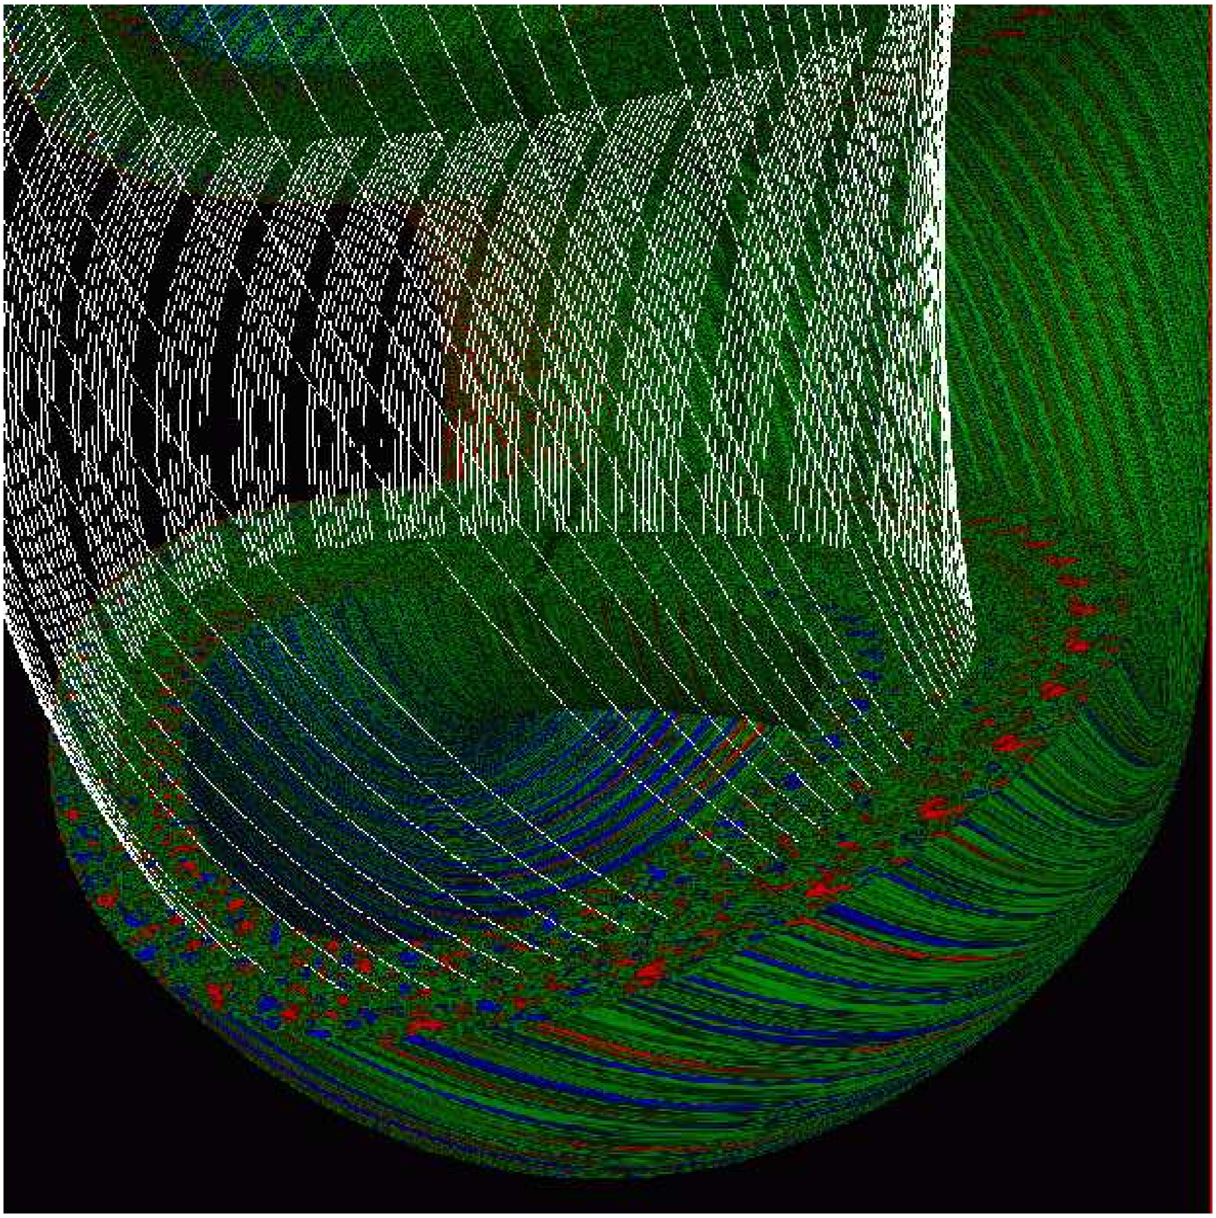

Figure 27: Gyrofluid simulation of ion-temperature-gradient turbulence in a torus mapped onto a DIII-D equilibrium with triangular cross-section. G. Kerbel, LLNL

3.3.3 3D Simulation of Heavy-Ion Beams for Inertial Confinement Fusion Drivers

David P. Grote and Alex Friedman
Heavy Ion Fusion Group, LLNL

Heavy-ion induction accelerators are being developed in the United States as the principal candidate "drivers" for inertial confinement fusion (ICF) power production. They also represent an attractive driver option for a high-yield microfusion research facility. Heavy-ion drivers are attractive in both cases because of favorable reliability, repetition rate, efficiency, lifetime, cost, and the favorable fusion chamber geometry they allow.[22]

The requirements for high-gain ignition of the fusion targets set the driver parameters. Current concepts employ about 10 beams of singly charged ions with masses of 100 to 200 amu and energies of 1 to 10 GeV, depositing a total of about 5 MJ on the target. Each beam at the target has a current of 1 to 10 kA and a main pulse duration of 10 ns (preceded by a longer, less intense "foot"). Focusing the beam from a distance of 3 to 10 m onto the target with a focal-spot radius of about 3 mm requires that the transverse temperature of the beam be kept below 1 keV.

The space-charge-dominated beams in an induction accelerator are effectively non-neutral plasmas; therefore, computational modeling of these beams uses techniques related to those used in both the accelerator and plasma physics communities. Because the beam resides in the accelerator for relatively few plasma oscillation periods, particle-in-cell (PIC) simulation techniques are especially effective and have proved valuable in the design and analysis of both ongoing experiments and future machines and in the study of basic physics issues. The WARP code includes both fully 3D (WARP3d) [23] and axisymmetric (WARPrz) [24] electrostatic PIC simulation models. Studies of an electrostatic quadrupole (ESQ) injector and a small recirculating induction accelerator with WARP3d are described in this article.

WARP3d uses a number of novel techniques to improve accuracy and efficiency. These include a capability for subgrid-scale placement of internal conductor boundaries and a technique for rapidly but accurately following particles through a long sequence of sharp-edged accelerator "lattice elements," using a relatively small number of time steps. WARP3d can simulate "bent" accelerator lattices, which iss how the code gets its name.

Electrostatic Quadrupole Injector

The goal of the Heavy Ion Fusion Accelerator program being pursued at LBNL, in collaboration with LLNL, is to address critical beam physics and technology issues of a heavy-ion

[22] R. O. Bangerter, "The Induction Approach to Heavy-Ion Inertial Fusion: Accelerator and Target Considerations," *Il Nuovo Cimento* **106 A**, no. 11, 1445(1994).

[23] A. Friedman, D. P. Grote, and I. Haber, "Three-Dimensional Particle Simulation of Heavy-Ion Fusion Beams," *Physics of Fluids B* **4**, no. 7, 2203 (1992).

[24] D. A. Callahan, A. B. Langdon, A. Friedman, and I. Haber, "Longitudinal Beam Dynamics for Heavy Ion Fusion," *Proceedings 1993 IEEE/APS Particle Accelerator Conference, IEEE Cat. No. 93CH3279-7*, 730, (1993)

fusion accelerator using beams of full driver-scale radius and line-charge density. [25] The injector developed as a part of that program uses a sequence of ESQ lenses with a superposed voltage gradient along the axis. The net effect is to both confine the beam transversely and accelerate it. The beam is emitted from a hot-plate source and is accelerated through a diode section into the main ESQ region. Figure 28 shows the layout of the injector.

A major issue in the injector design is degradation of beam quality (brightness) by the nonlinear multipole components of the focusing fields and the so-called "energy effect." The energy effect arises from the axial velocity spread of particles in the ESQ, which is due to the transverse variation of the potential. This effect is present in all electrostatic focusing systems, but is generally insignificant since the focusing potentials are a small fraction of the beam energy. However, at the low-energy end of an ESQ injector, the focusing potential is about half the beam energy. Proper treatment of the beam dynamics requires detailed 3D simulation with a realistic lattice model.

For design of the injector optics, the steady-state beam behavior was of primary interest. Runs were made in a quasi-steady-state mode, taking several particle time steps between each field solution and continuing until the injector was filled with beam and the system converged to a time-independent state. These 3D simulations of the injector required less than three minutes on a single processor of the NERSC CRAY Y-MP C90. Taking advantage of the ESQ's two-plane symmetry, the dimensions of the field grid were typically 40x40x348. Typically, 100,000 particles and 500 time steps were used. The savings made by this approach are apparent when compared to fully time-dependent simulations of transient behavior which require several hours of C90 CPU time.

These simulations led to an understanding of the mechanisms of beam quality degradation in an ESQ, and a number of corrective measures were identified. The most effective of these was increasing the diode voltage, which directly reduces the energy effect, decreases the size of the beam, and so reduces the influence of nonlinear fields.

Using WARP3d simulations as a guide, a scaled-down version of the injector was built at LBNL as a proof-of-principle experiment. Figure 29 shows an example of the excellent agreement between the experiment and simulation results.

The final design of the full-size ESQ injector was also carried out with the help of WARP3d simulations. The completed injector exceeds its design requirements. Figure 30 shows phase space representations of the simulations and experiment.

[25] T. Fessenden, R. Bangerter, D. Berners, J. Chew, S. Eylon, A. Faltens, W. Fawley, C. Fong, M. Fong, K. Hahn, E. Henestroza, D. Judd, E. Lee, C. Lionberger, S. Mukherjee, C. Peters, C. Pike, G. Raymond, L. Reginato, H. Rutkowski, P. Seidl, L. Smith, D. Vanecek, S. Yu, F. Deadrick, A. Friedman, L. Griffith, D. Hewett, M. Newton, and H. Shay, "ILSE, The Next Step toward a Heavy Ion Induction Accelerator for Inertial Fusion Energy," in *Proceedings of the 14th International Conference on Plasma Physics and Controlled Nuclear Fusion Research, 1992 (Vienna: International Atomic Energy Agency, 1993)* **3**, 89 (1993).

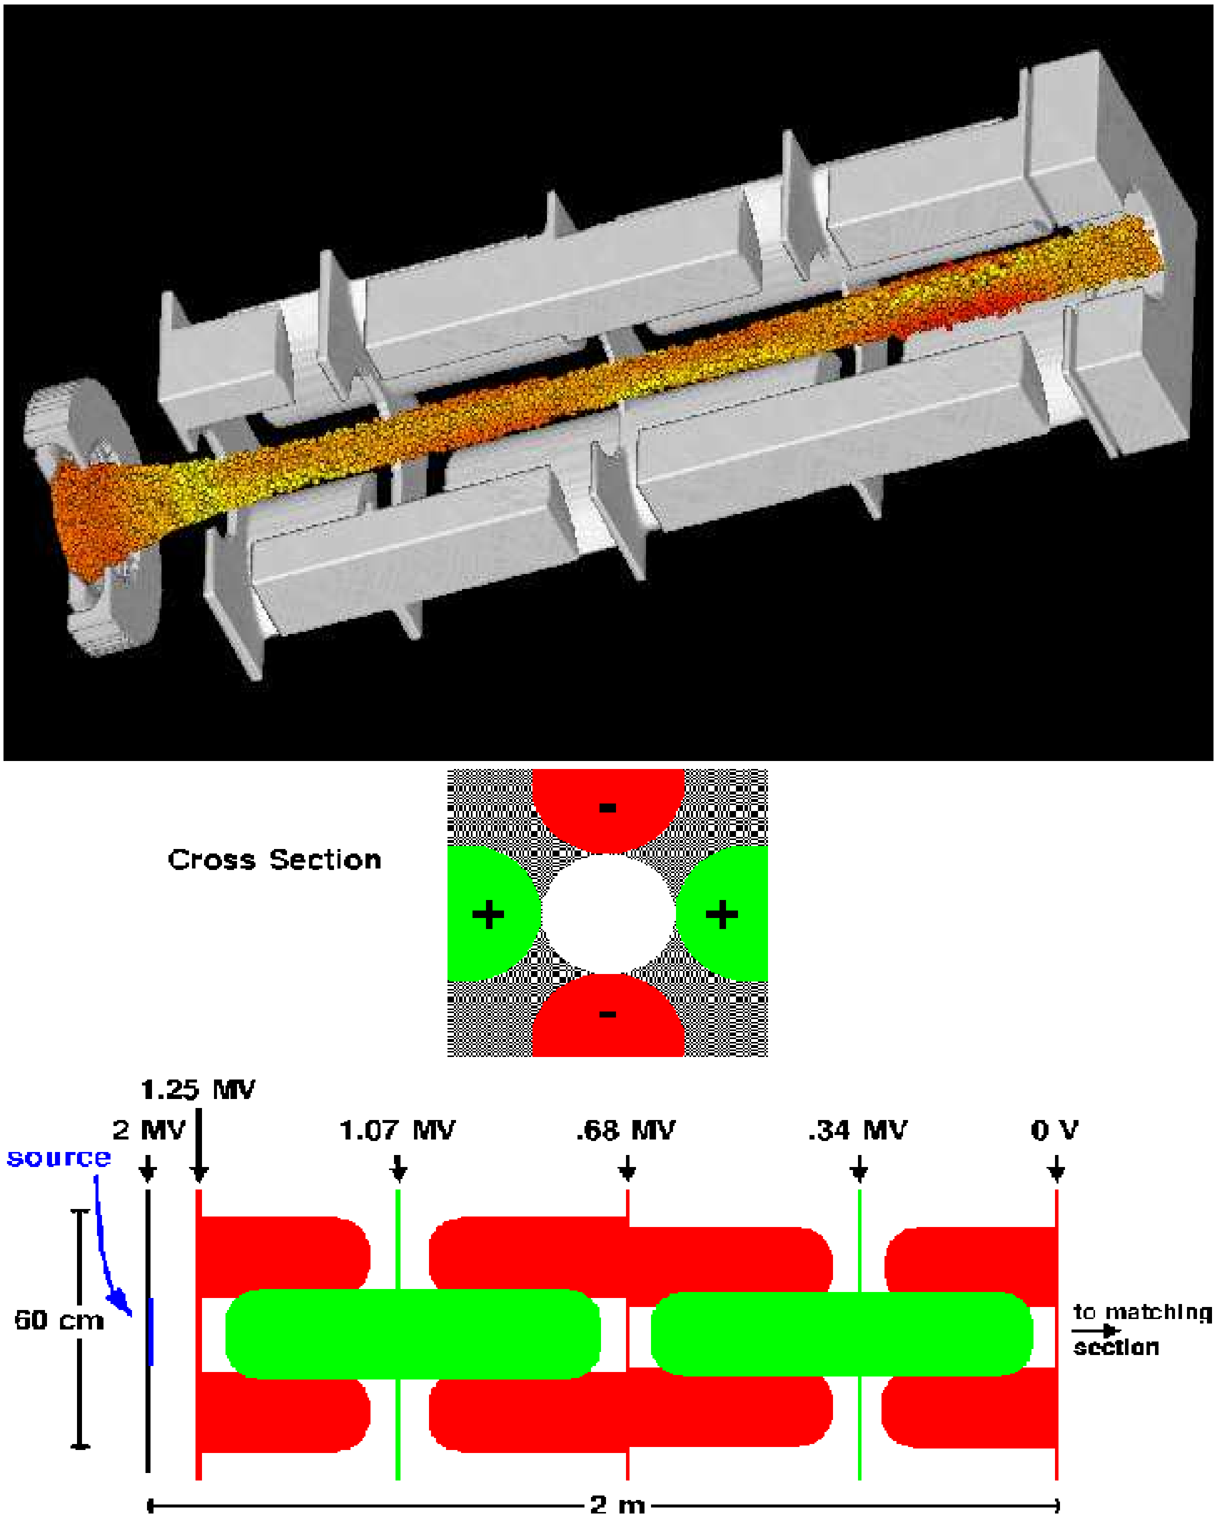

Figure 28: The layout of the ESQ injector. The beam is emitted from the source and accelerated to the right through the four quadrupole lenses. The voltages are arranged so that the beam is both strongly focused and accelerated.

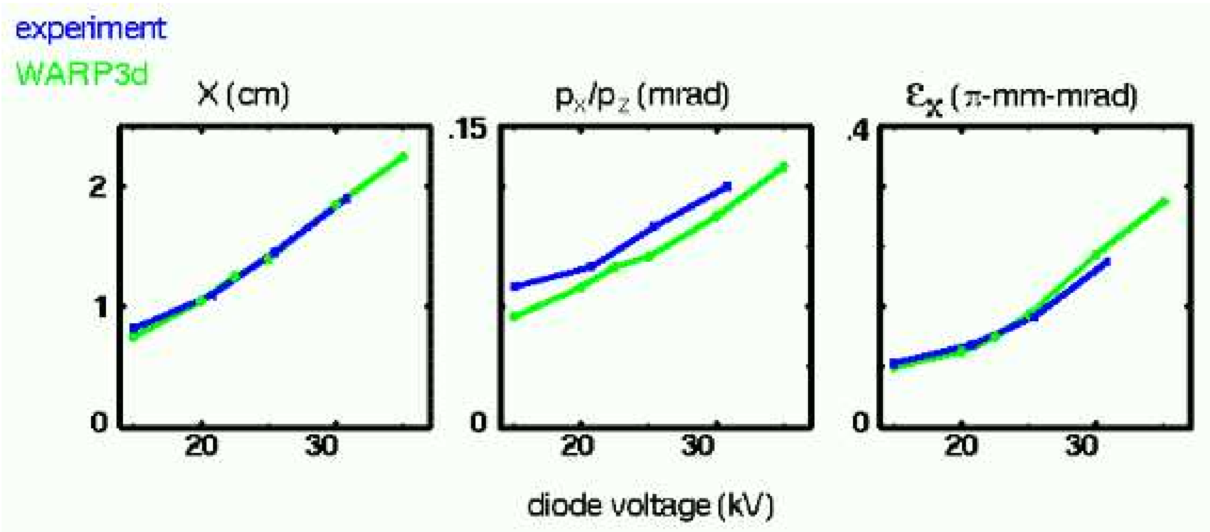

Figure 29: Comparison of experimental and simulated transverse beam parameters at the end of the scaled ESQ for a varying diode voltage. The emittance is a measure of beam quality; high quality beams have low emittance.

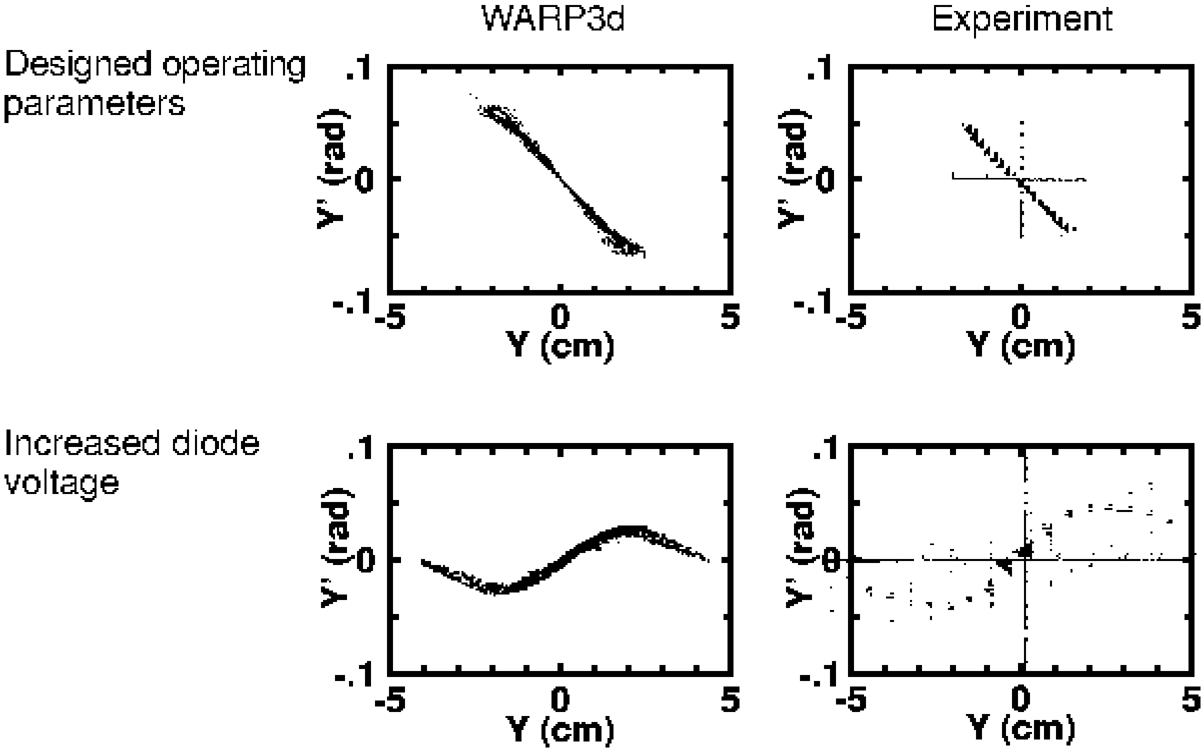

Figure 30: Comparison of experimental and simulation transverse phase space representations at the end of the full-size ESQ injector. The top compares the results at the design parameters. The bottom compares the results with an increased diode voltage; in this non-optimal case, the representations are significantly distorted from a straight line.

Recirculating Induction Accelerator

A recirculating induction accelerator, or "recirculator," offers the prospect of reduced cost relative to a conventional linear accelerator because the accelerating and focusing elements are re-used many times per shot. [26] The circumference of the recirculator can be much less than the length of the equivalent linear accelerator, and the accelerating cores can be smaller because there is no need to accelerate the beam at the maximum gradient. A small recirculator is being developed at LLNL to explore the beam dynamics of a full-sized fusion driver in a scaled manner. This research is being carried out as a sequence of experiments that will lead to the construction of a complete prototype.

The key dimensionless parameters that characterize the beam are similar to those of a driver-scale ring. The planned ring will have a circumference of 14.4 m, made up of 40 "half lattice periods," each with a permanent magnetic quadrupole for transverse confinement, an electric dipole for beam bending, and an induction acceleration gap (except in the insertion/ extraction region). A beam of singly charged potassium ions will be accelerated from 80 to 320 keV over 15 laps, at a current increasing from 2 to 8 mA. With its low temperature, the beam is strongly influenced by the self-electric fields arising from its own space charge.

Although a full-scale driver will use magnetic dipole elements for bending the beam, electric dipole plates were chosen for the prototype to minimize costs. However, the size of the plates is constrained by space considerations and voltage-holding requirements. As a result, they introduce significant nonlinear fields. Using WARP3d, the plate shape was adjusted to minimize 3D field nonlinearities and their influence on beam quality.

WARP3d simulations were used to follow the beam for all 15 laps. Typically, the 2-ms-long beam is made up of 100,000 simulation particles. The field grid is a moving window covering four half lattice periods. It has a 32x16 transverse grid (taking advantage of vertical symmetry), and a 128-cell axial grid. Figure 31 shows a beam of length two half lattice periods in the lattice.

The simulations show an initial transverse heating of the beam on the order of 50 percent as the beam "matches" itself from the straight insertion line to the bent recirculator lattice. This heating is due to radial separation of particles with differing axial velocities in the bends, which leads to nonlinearities in the space charge forces. The nonlinearities thermalize, producing an increase in temperature. The observed growth agrees well with theoretical analysis.[27]

After the initial heating, little further degradation in beam quality is seen over the 15 laps. The overall degradation is acceptable.

---

[26] J. J. Barnard, F. Deadrick, A. Friedman, D. P. Grote, L. V. Griffith, H. C. Kirbie, V. K. Neil, M. A. Newton, A. C. Paul, W. M. Sharp, H. D. Shay, R. O. Bangerter, A. Faltens, C. G. Fong, D. L. Judd, E. P. Lee, L. L. Reginato, S. S. Yu, and T. F Godlove, "Recirculating Induction Accelerators as Drivers for Heavy-Ion Fusion," *Physics of Fluids B* **5**, no. 7, 2698 (1993).

[27] J. J. Barnard, H. D. Shay, S. S. Yu, A. Friedman, and D. P. Grote, "Emittance Growth in Heavy Ion Recirculators," *1992 Linear Accelerator Conference Proceedings, AECL-10728*, 229 (1992).

Conclusions and Future Work

WARP3d simulations played an important role in the design and successful operation of the ESQ injector for the Heavy Ion Fusion Accelerator program at LBNL. For the small-recirculator experiments at LLNL, WARP3d simulations showed that a beam can be accelerated through 15 laps using electric-dipole bending, which gave us sufficient confidence to adopt this reduced-cost approach. The goal of the heavy-ion-fusion modeling effort is effective simulation of experiments and, ultimately, of a full-scale fusion driver.

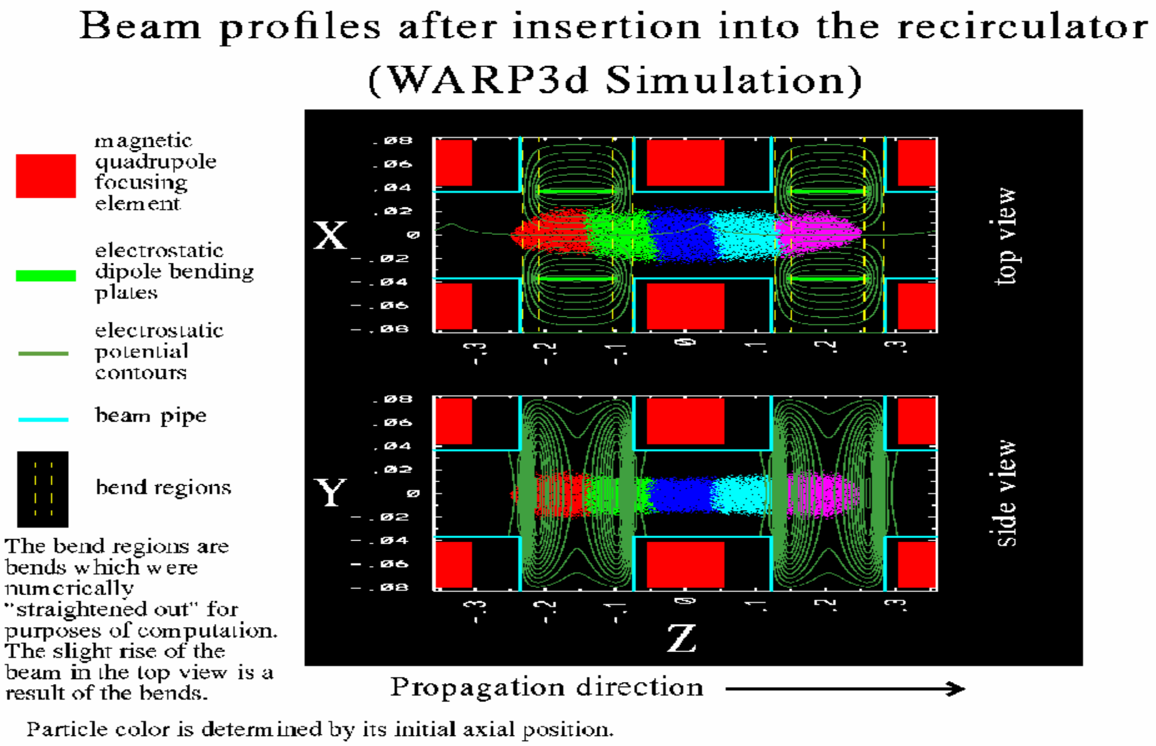

Figure 31: A beam in the lattice as represented in WARP3d. The figure shows a sampling of the particles that make up the beam. The locations of the focusing and bending elements are identified. The bend regions are sections where polar coordinates are used to model the bent accelerator. For the purposes of the plot, the bent regions were straightened out.

3.4 Biological and Environmental Research

The Office of Biological and Environmental Research manages the Department's Biological and Environmental Research (BER) program. BER's extraordinary legacy spans 50 years and includes these four areas:

- Life Sciences: The Health Effects and Life Sciences Research Division administers research programs to develop fundamental biological information and to advance technologies for use in research on the health effects of energy-related agents and processes.
- Medical Applications and Biophysical Research: The Biological and Environmental Research Medical Applications and Biophysical Research Division supports research in several distinct areas of science and technology that address a broad range of fundamental and applied research applications in medicine and biophysics.
- Environmental Processes: Global Change Research: The Global Change activities include the process research and modeling efforts needed to
  - improve understanding of factors affecting the Earth's radiant-energy balance;
  - predict accurately any global and regional climate change induced by increasing atmospheric concentrations of greenhouse gases;
  - quantify sources and sinks of energy-related greenhouse gases, especially carbon dioxide; and
  - 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.
- Environmental Remediation: The Environmental Remediation research portfolio is focused on developing an understanding of the fundamental physical, chemical, geological, and biological processes that must be marshaled for the development and advancement of new, effective, and efficient processes for the remediation and restoration of the Nation's nuclear weapons production sites.

3.4.1 Molecular Structures of Carcinogen-DNA Adducts

Brian E. Hingerty, Oak Ridge National Laboratory
Suse Broyde, Department of Biology, New York University

The nucleic acids molecular mechanics program DUPLEX is specifically designed to compute atomic resolution structures of DNA that is damaged by linkage with carcinogenic environmental chemicals. These chemicals, polycyclic aromatic hydrocarbons and amines, are present in automobile exhaust, tobacco smoke, dyes, and barbecued meat and fish.

According to our current understanding, cancer is initiated by these substances, once they gain entry into the cell, by their biochemical activation to reactive moieties which then attack the DNA bases (A, G, C, T). Once bound to the bases, they can alter the DNA shape (conformation). This change can profoundly impact on the normal functioning of DNA during replication. In this process the replication machinery unwinds the double helix and synthesizes a daughter strand with bases that are the complement of the parent template–A and T are paired, and G and C are paired. Thus, for example, the machinery places a G in the daughter where the parent has a C. When the DNA is damaged by a carcinogen, this process can fail in accuracy such that a mismatch occurs at the modified base, or the modified base may be skipped altogether, or an extra base may be added at the lesion site. These errors constitute mutations, which can cause the synthesis of aberrant proteins that lead to cancer.

Obtaining atomic resolution views of carcinogen modified DNA has long been deemed essential for understanding the initiation of cancer by these substances. However, the task is formidable. Studies of recent years have revealed that a given substance may be mutagenic or not depending on what the bases adjacent to the modified one are (a sequence that is mutagenic is termed a "hot-spot" while a non-mutagenic one is a "cold-spot"). Both hot-spot sequences, and cold-spot sequences as controls, must be examined. Furthermore, it is now understood that the damaged DNA must be examined in both normal double helices (the shape of the DNA when it is not replicating) and at single strand-double strand junctions that more accurately mimic the DNA shape as it is unwinding during the replication process. Double strands must be studied because lesion repair processes may succeed or fail, depending on the damaged DNA's shape; single strand-double strand replication intermediates must be studied because they model the situation during replication when the mutagenic event actually takes place. Furthermore, a given carcinogen may damage the DNA at more than one locus, and all the damaged loci must be examined. Moreover, certain polycyclic aromatic hydrocarbons are carcinogenic, while chemically very similar ones, even mirror images, are harmless; it is necessary to study both harmful substances and harmless ones, as controls.

Thus, it is necessary to obtain molecular views of DNAs modified by many different chemicals, linked to the DNA bases at different sites, in many different base sequence neighbor contexts, and in double strands and single strand-double strand junctions. With a large atomic resolution data base, structural hallmarks that are related to known mutagenic outcome from biological studies will emerge. Indeed, our investigations have already uncovered some such hallmarks. Once the structural data base is achieved, it will be possible to screen chemicals by computing their conformational impact on the DNA. If mutagenic hallmark structural damage is computed for a given substance, it would likely be harmful. This computational prediction of the mutagenic and carcinogenic potential of environmental chemicals will be extremely valuable. Current screening procedures require laborious, expensive and controversial animal testing to be relevant to the human condition with each test costing upwards of half a million dollars. Computational modeling would both lower the cost and controversy of these mandated screening procedures.

However, obtaining atomic resolution views experimentally is a daunting undertaking. For unambiguous structural information at atomic resolution, x-ray crystallography is the tool. However, to date, there is not a single crystal structure of any DNA modified by a polycyclic aromatic carcinogen; it has, so far, been impossible to obtain crystals of suitable diffracting quality. Average structures in solution are obtained by high resolution nuclear magnetic resonance (NMR) methods. However, fewer than a dozen structures have been obtained by this method to date. Synthesis has been a great problem. In many cases, the synthetic problems have not yet been overcome. In others, very laborious synthetic procedures (one preparation can take several man-years) are required. Furthermore, even when synthesis succeeds, the preparation may fail to yield interpretable NMR data if the structure is very mobile. Moreover, the NMR data, in the form of inter-proton distances, requires molecular mechanics for translation of the distances into a structure.

Our group has been involved in the molecular mechanics aspect of the NMR studies in most of the structures published so far. Thus, the experimental approach cannot possibly examine the large number of compounds, in the multiple base sequence contexts, and in both double strands and single-double strand junctions, that must be investigated to identify structural hallmarks of mutagenicity and carcinogenicity which will permit a predictive computational approach. Molecular mechanics calculations, on the other hand, have the capacity to model them, including especially the many whose synthesis is extremely difficult and/or for which high resolution NMR data is uninterpretable. In addition, for full understanding, both static and dynamic views are needed, while the experimental methods offer only static ones. Molecular dynamics simulations can, however, offer the animation that will yield the full picture.

DUPLEX has been tailored with special features that permit reliable computation of carcinogen-modified DNA structures, by surveying the potential energy surface with energy minimization. These features include (1) employing the reduced variable domain of internal coordinates (torsion angle space), (2) studying small modified subunits extensively and employing special building strategies to generate larger structures, and (3) including experimental data in the form of penalty functions (restraints) when these are available. Currently, DUPLEX can compute up to 24 residues of double stranded DNA in principle. However, the practical limit is about a dozen, since obtaining a single such structure can take 50 - 100 hours on a Cray YMP-C90 (such as the a-machine at NERSC), due to the computational intensity of this work. More powerful computing resources would permit analysis of the huge array of structures that must be examined in a much more reasonable time frame. Furthermore, much longer DNA strands could be studied.

The 140 residues of the DNA nucleosome core that is the first order building block of mammalian chromosomes could be treated instead of the small model duplexes that are now being studied of necessity. Even larger biologically important structures should become treatable; one key example is supercoiled DNA comprised of hundreds of residues. Supercoiling, the higher order winding of the double helix in topologically constrained situations like closed circles, to form figure-8 and other inter-wound forms, plays a leading role in DNA replication; furthermore, carcinogens are known to alter this essential topology, yet nothing is known about it at atomic resolution. DUPLEX will be able to investigate this morphology at atomic resolution with more powerful resources. The only way to accomplish such calculations is from advances in massively parallel computer architectures in addition to algorithm improvements.

Recently, we have carried out an extensive study of a substance, present in barbecued meat, known as PHIP (2-amino-1-methyl-6-phenylimidazo [4,5-b]pyridine). PHIP (see Fig. 32) is mutagenic and carcinogenic in rodents. When biochemically activated it binds to DNA, primarily to the base guanine at a position known as carbon-8. PHIP causes the damaged guanine to pair with adenine instead of the normal partner cytosine during DNA replication, causing a mutation that is probably an initiating event in the carcinogenic process.

We have carried out extensive computations to achieve atomic resolution views of the PHIP-bound DNA in a double helix in which the modified guanine is mispaired with adenine. In this work we used the nucleic acids molecular mechanics program DUPLEX. The computations revealed that binding of PHIP caused the damaged guanine to rotate 180 degrees from its normal "anti" position in the double helix to an abnormal position termed "syn". In this orientation it can readily accommodate an adenine in opposition, placing the carcinogenic PHIP-moiety so that it is comfortably snuggled into the helix "minor groove." Thus, the computed structure offers a molecular view of an initiating mutagenic event in PHIP induced carcinogenesis: the mismatching of the modified guanine with adenine (see Figure 32). This is the only available atomic resolution structure of DNA bound by PHIP, as experimental studies have not yet been successful.

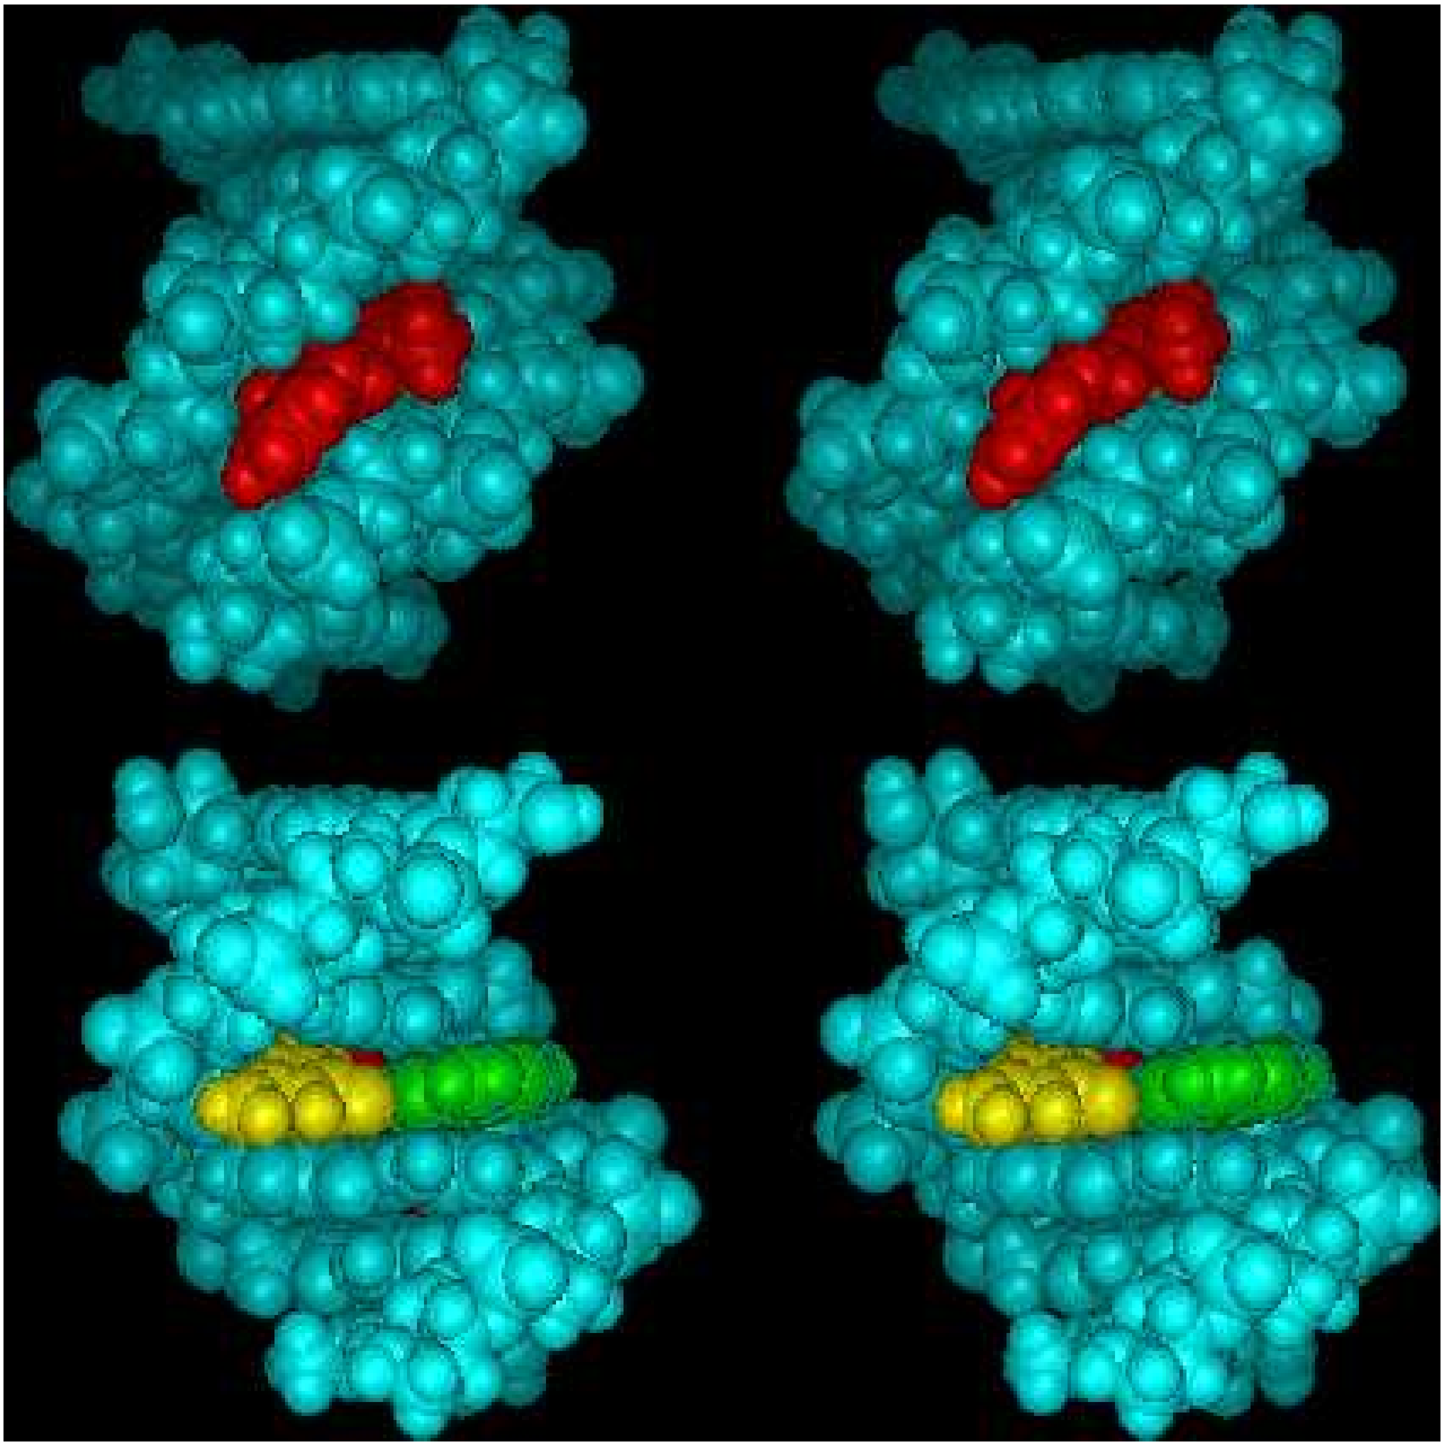

Figure 32: PHIP (2-amino-1-methyl-6-phenylimidazo [4,5-b]pyridine)

3.4.2 North-West Computational Chemistry Suite

Edoardo Aprà, David E. Bernholdt, David M. Elwood, Miles J. Deegan, Michel Dupuis,
George I. Fann, Martyn F. Guest, Maciej Gutowski, Herbert A. Früchtl,
Robert J. Harrison, Rick A. Kendall, Rik J. Littlefield, Jeffrey A. Nichols,
Jaroslaw Nieplocha, Hugh L. Taylor, Tjerk P. Straatsma, John B. Nicholas,
Ramzi A. Kutteh, Gregory S. Thomas and Adrian T. Wong
High Performance Computational Chemistry Group
Environmental and Molecular Sciences Laboratory
Pacific Northwest National Laboratory

NWChem is computational chemistry package for the study of large chemical systems that is under development and is designed to run on high-performance parallel supercomputers as well as workstation clusters. It aims to be scalable both in the ability to treat large problems efficiently, and in its usage of available parallel computing resources. The suite utilizes parallel programming tools (TCGMSG and the Global Array library) developed by PNNL staff. NWChem has been optimized to perform calculations on large molecules using large parallel computers; the code has been ported to the following MPPs: KSR-1/2, Intel Delta/Paragon, CRAY T3D and IBM SP1/2.

Current functionality includes:

- Self Consistent Field or Hartree Fock
  - RHF, ROHF, UHF
  - Resolution of the Identity integral approximation (RI-SCF)
- SCF (RHF, ROHF, UHF) gradients
- Geometry Optimization (Minimization and Transition State)
- MP2, RI-MP2
- MCSCF
- Gaussian Density Functional Theory (DFT) energy with many local and non-local exchange-correlation potentials
- DFT gradients
- Frequency Analysis by Finite Difference.
- Molecular Dynamics
- QM/MM with semi-empirical QM

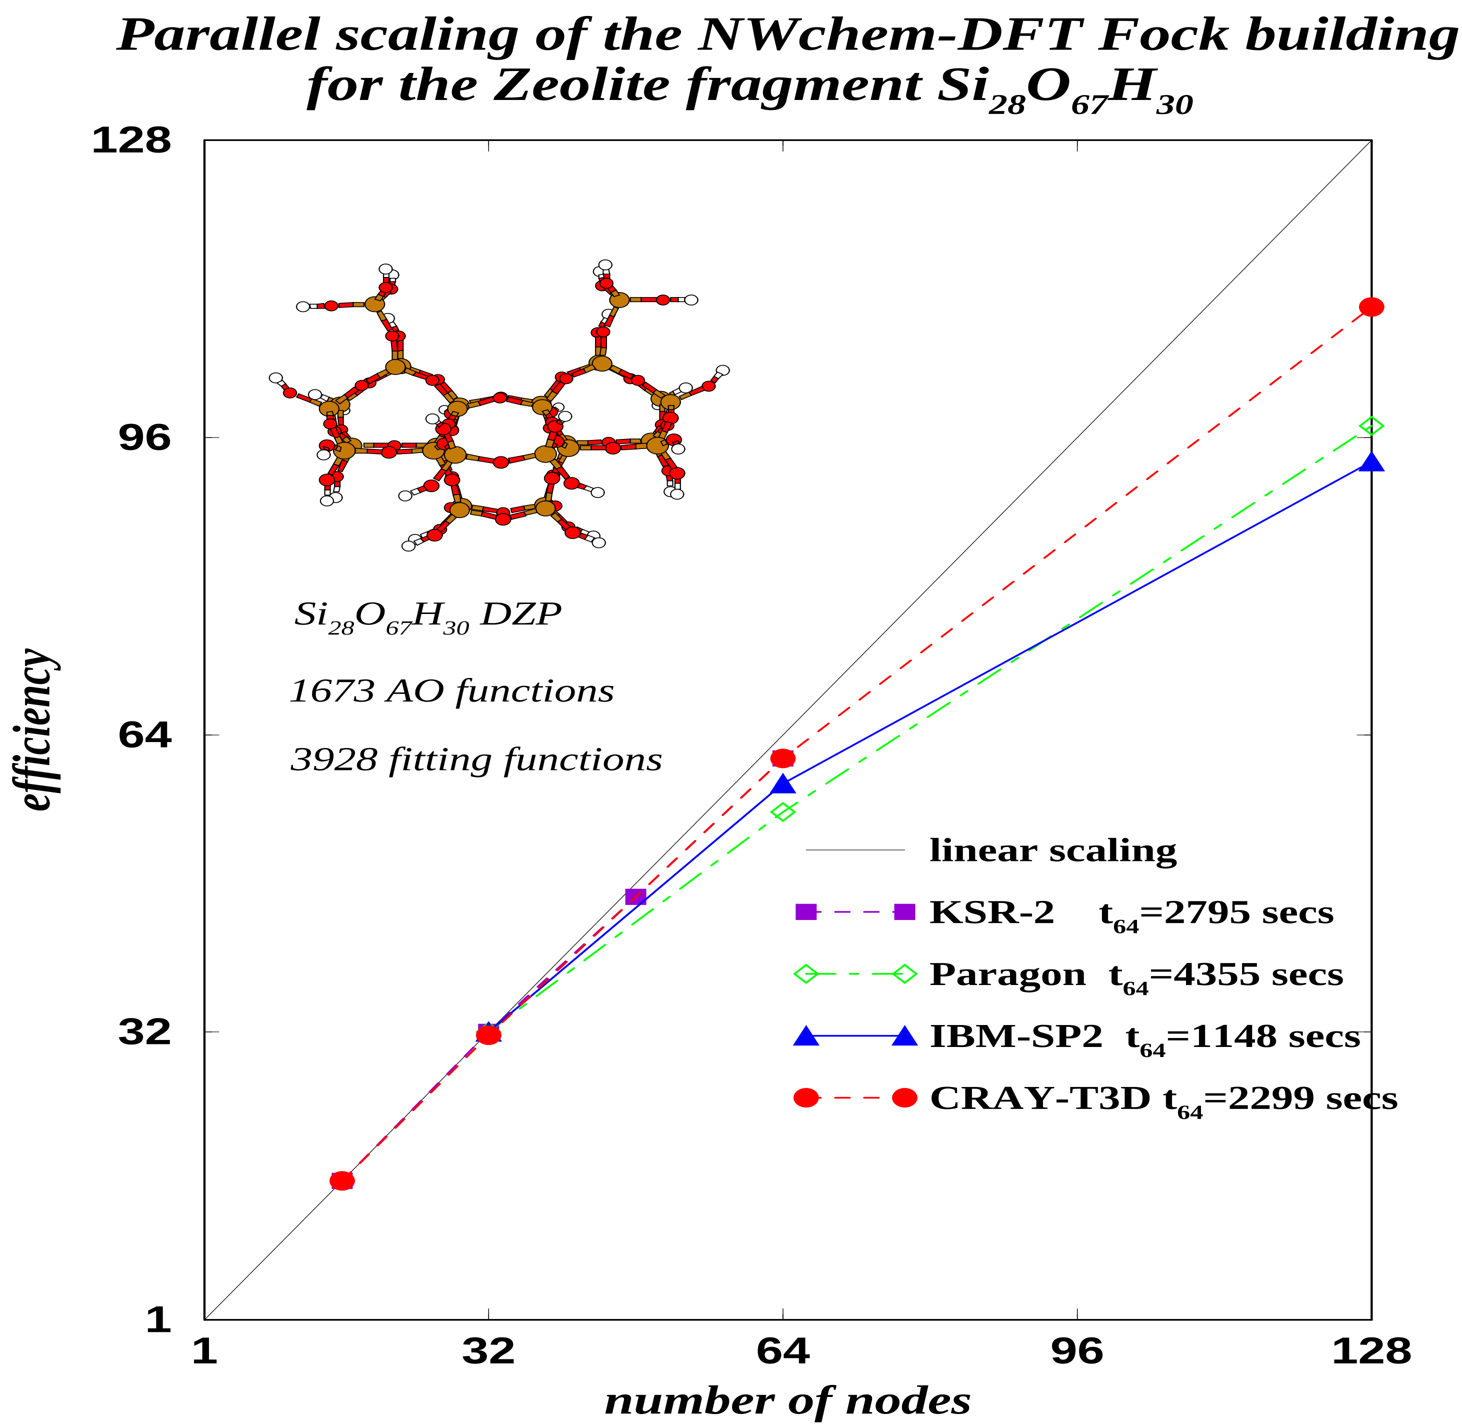

Figure 33: NWChem: DFT Scalability on Various MPP Supercomputers

Quantum Chemical study of the binding of alkali metal cations with crown-ethers: effects of computational parameters and methods.

Crown ethers have drawn much experimental and theoretical interest since they were first described by Pederson in 1967. These cyclic polyethers show a remarkable range of specificity for metal cations and are being considered as separation agents for radioactive cations in waste steams. As a prototypical interaction of these systems we have studied the binding of various alkali metal cations with multiple crown ether derivatives using Local and Non-Local DFT and MP2 methods. Binding energies for these system have been determined from moieties whose structures have been optimized using a variety of approximations within the DFT method, including: various basis sets (orbital and charge density fitting), numerical grids (various levels of accuracy, and types), and basis set superposition errors (orbital and grid).

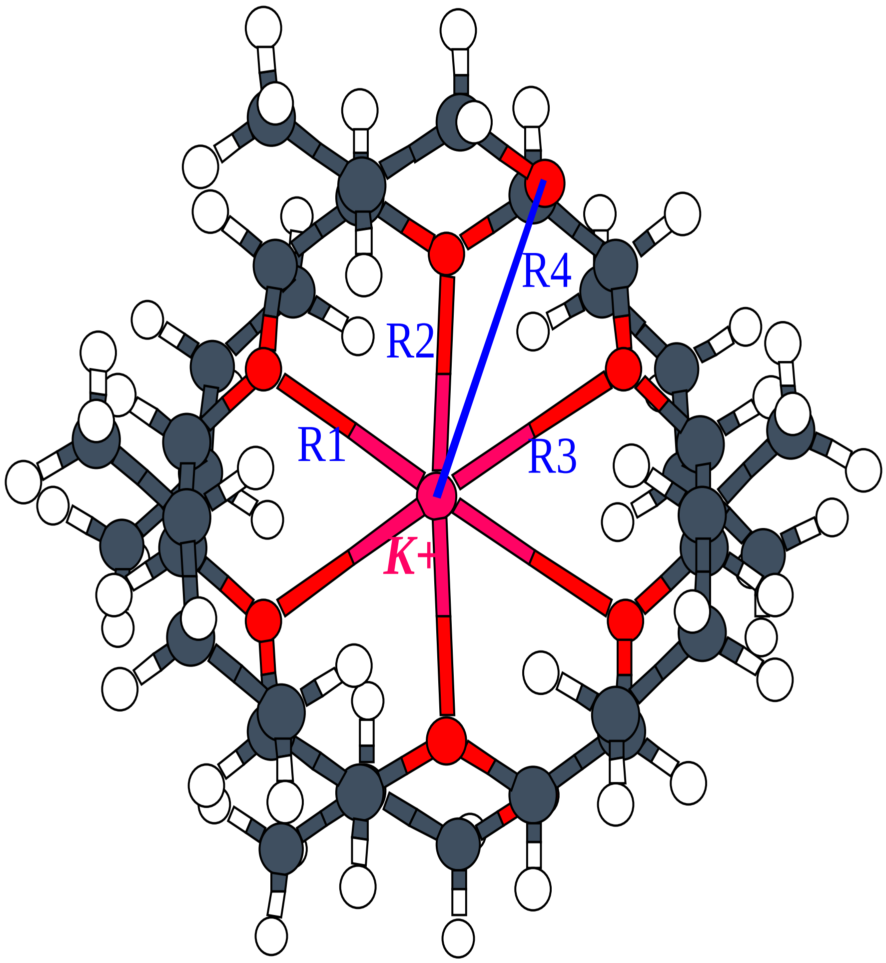

	HF/631+G* (hybrid)	DFT BLYP/6-31G*
R1	2.807	2.915
R2	2.806	2.807
R3	2.763	2.807
R4	3.924	3.894
BE	-77.1 (-75.8)	-75.8 (-70.6)

Figure 34: Binding energies (in kcal/mol) and bond lengths (in Å) for the modified 18-crown-6 molecule. The number in parentheses are obtained after estimating the BSSE error using the CP method.

3.4.3 Parallel Climate Modeling on High-Performance Computing Systems

Gerald L. Potter Lawrence Livermore National Laboratory

Parallel Climate Model

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 need for rational formulation of energy policy. Developing this knowledge requires a combination of field observations and numerical modeling, as well as theoretical advances. A comprehensive climate system model must describe the behavior of the atmospheric, oceanic, cryospheric and biospheric systems and the interactions between them.

The principal tools used to simulate the behavior of the global atmosphere and oceans are known as general circulation models. Large numbers of cells or degrees of freedom are needed to adequately resolve the circulation on the global scale.

The NCAR Community Climate Model version3 has been coupled with the Parallel Ocean Program (POP) and a sea ice model from the Naval Postgraduate School in a massively parallel computer environment. The target machines are the CRAY T3D/T3E and SGI Origin 2000. With the cooperation and assistance of NCAR's Scientific Computing Division, the Parallel Climate Model (PCM) model code has been ported to the Hewlett-Packard and SGI Origin 2000.

Based on the experience with the NCAR Climate System Model, in order to minimize the initial drift of the coupled system, the ocean/ice can be spun-up with forcing from previous CCM3 runs with prescribed sea surface temperature. This has also been useful in demonstrating *and improving* the kind of adjustments that occur in the ocean and ice due to coupling the CCM3, without having to run the more expensive coupled system. Lastly, the full system has been run and tested.

Atmospheric Model

The atmospheric component is the massively parallel version of the NCAR Community Climate Model version 3 (CCM3). This model includes the latest versions of radiation, boundary physics, and precipitation physics and is a state-of-the-art atmospheric component. The CCM3 also includes the land surface model (LSM) which takes into account soil physics and vegetation. This model has been coded to run on the T3D/E and SGI Origin 2000.

Ocean Model

The POP model's grid is $384\times288\times32$, with an average resolution of $\frac{2}{3}$ degree latitude and longitude with increased latitudinal resolution near the equator of approximately $\frac{1}{2}$ degree. Because of the displaced pole, there is relatively higher horizontal resolution in the eastern North Pacific, in the Arctic Straits near northern Canada and Greenland, and in the Gulf Stream area. In addition, the continents and bottom topography were carefully modified to obtain realistic flow in many regions throughout the globe. This model is being spun up with observed surface and subsurface forcing in preparation for coupling. The model is presently running efficiently on the NERSC CRAY T3E and LANL SGI Origin 2000. Plans are being made to add more realistic upper ocean parameterizations to the model. The model in its present form yields an extraordinary simulation of the Arctic Ocean, tropical Pacific, and boundary currents, such as the Gulf Stream, with eddies resolved in most basins (see figures 35-37). Tools have been developed to interpolate the model output to regular grids. More documentation on the POP model can be found in Los Alamos report LA-UR-95-1146 by Richard D. Smith, Samuel Kortas, Bertrand Meitz, *Curvilinear Coordinates for Global Ocean Models*, pp 38.

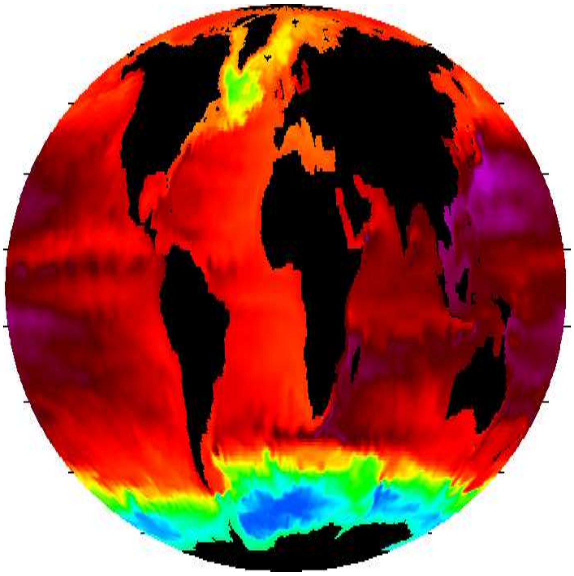

Figure 35: Global POP simulation

Sea Ice Model

The sea ice component was adapted for PCM from the Arctic model implemented by Yuxia Zhang (NPS) in an eddy-resolving ocean model. The Zhang and Hibler viscous-plastic ice rheology is used for dynamical processes, with thermodynamics from the Semtner and Parkinson-Washington models. The grid is transformed such that the horizontal resolution is constant (27 km), thus avoiding the problem of convergence near the pole as on a latitude-longitude grid. The model is configured for efficient computation on the CRAY T3X architectures, and is being converted for use on the SGI Origin and HP/Convex platforms by the SCD Computational Science Section.

Computing

The PCM has several configurations that can be run on the T3E - mainly, they involve activating/deactivating the various component models. The T3E requirements are then dictated by the specific PCM configuration that has been set-up. For example, the fully coupled PCM is restricted to 64 PEs because CCM3 is restricted to 64 PEs. If the inactive

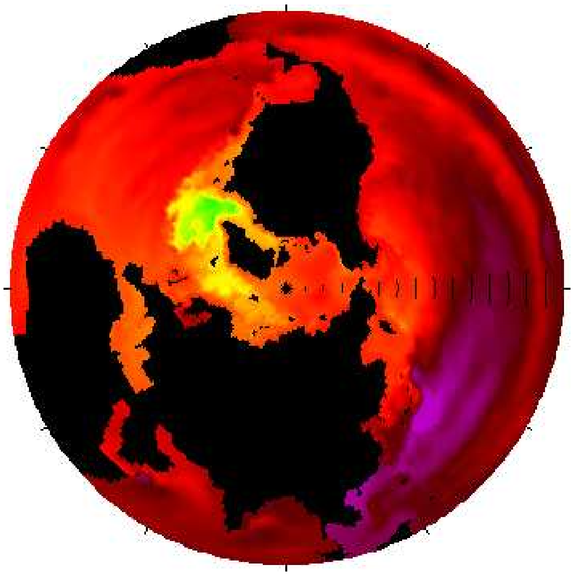

Figure 36: Northern Hemisphere POP simulation

CCM3 version is used to spin-up the ocean and ice systems, is possible to run with 512 PEs with very good scalability.

The Parallel Ocean Program (POP) developed at LANL is one of many efforts in the Department of Energy's CHAMMP[28] program aimed at increased understanding of the earth's climate through the use of modeling and high performance computing. POP has been used to perform the highest resolution ($\frac{1}{5}$ degree on average) global ocean simulation ever undertaken using the Thinking Machines CM5 computer located at LANL's Advanced Computing Laboratory (ACL[29]).

POP is a descendant of the Bryan-Cox model that is used frequently in ocean simulations. This earlier model has been substantially improved and adapted for use with massively parallel computers. Some of these improvements are:

- a surface pressure formulation that allows for an unlimited number of land masses (with no additional computational cost) and unsmoothed bottom topography.
- an implicit free surface technique that allows the air-sea interface to evolve freely.
- the equations of motion are formulated and discretized to allow the use of any locally orthogonal horizontal grid (instead of just latitude/longitude) which, for example, can let the Arctic ocean be resolved without the standard problems of convergence of meridians at the North Pole.

[28] c.f., http://www.epm.ornl.gov/chammp
[29] c.f., http://www.acl.lanl.gov

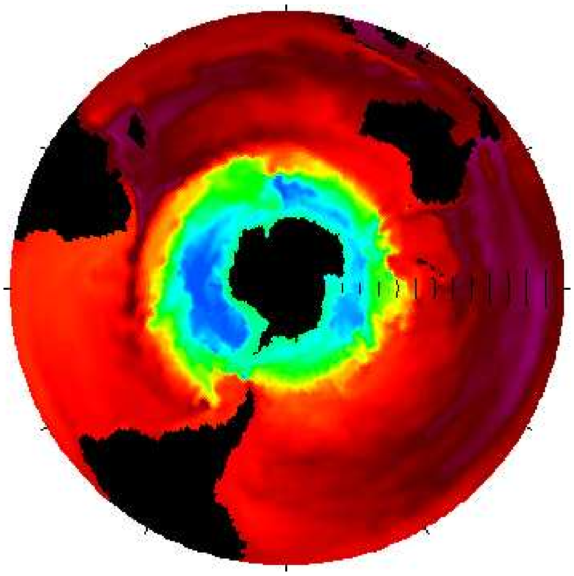

Figure 37: Southern Hemisphere POP simulation

- the code is written in Fortran90 and is capable of running on a variety of parallel and serial computer architectures.

A primary motivation for performing such high-resolution simulations is to resolve eddy motions that can play an important role in the dynamics of the ocean.

Project Participants

The following are the scientists/programmers involved in the coupled climate model effort in alphabetical order: J. Arblaster (NCAR), T. Bettge (NCAR), R. Chervin (NCAR), T. Craig (NCAR), J. Dennis (NCAR), J. Dukowicz (LANL), J. Hack (NCAR), S. Hammond (NCAR), E. Hunke (LANL), R. James (NCAR), P. Jones (LANL), R. Loft (NCAR), R. Malone (LANL), M. Maltrud (LANL), W. Maslowski (NPS), J. Meehl (NCAR), A. Semtner (NPS), R. Smith (LANL), G. Strand (NCAR), W. Washington (NCAR), V. Wayland, (NCAR), J. Weatherly (NCAR), D. Williamson (NCAR), and Y. Zhang (NPS).

Recommendations

The major issue is the 64 PE queues. Higher priority (i.e., more time) given to that queue each day would make the simulations much more efficient. To utilize the allocation, it would be necessary to have two 4-hour jobs make it through that queue through the allocation period. Another ocean spin-up phase will be initiated that will use the 128 PE queue.

The PCM requires no special compilers (we use the standard F90), and no special math libraries.

Our memory use on the T3E with 64 PEs is 9 Mwords per PE. We require 4 GB of temporary disk space during each job submission. Disk space is not currently a problem.

Each year we simulate we save, on average (again, depends upon active components), 2 GBytes. We will run in increments of 100-year experiments, so each experiment will be 200 GBytes of storage.

The NERSC T3E is regarded as the workhorse and NCAR is very satisfied with the performance. Long runs have been initiated in the 64 PE and 128 PE queues. NCAR is concentrating on long control and climate change experiments. Extensive simulations are being planned at NERSC over the next 6 months. Further down on the list of things to do is that we expect to bring up version of PCM that can run on IBM SP systems. Finally, Rodney James of NCAR and John Drake of Oak Ridge are updating the parallel version of CCM so that two-dimensional decomposition is used in CCM3. This will allow CCM3 to run efficiently at T42 on more than 64 PEs on a T3E. This is ongoing work.

Finally, concerning our performance on the T3E, the full PCM achieves about 40 Mflops/PE on 64 PEs - that's 2.5 Gflops. Without CCM3, on 512 PEs, NCAR achieves nearly 40x512= 20 Gflops in tests without I/O.

3.5 Computational and Technology Research

The Office of Computational and Technology Research (OCTR) provides a focal point in the Office of Energy Research (ER) of the Department of Energy for long-term computational and technology research. It manages research in forefront and diverse applied mathematical sciences, high performance computing, communications, and information infrastructure which spans the spectrum of activities from strategic, longer-term, fundamental research to technology research, development, and demonstration. It links ER's science programs and laboratories to national economic competitiveness by conducting long-term, high-risk industry relevant research and development projects in critical technology areas. The Office provides the technical, analytical, and management direction for development, implementation, and evaluation of long-term research programs which include major involvement in activities such as High Performance Computing and Communications (HPCC), the National Information Infrastructure (NII), the American Textiles Consortium (AMTEX), the Advanced Computation Technology Initiative (ACTI), and the Environmental Technology Partnerships (ETP). The Office leads and manages the Technical Information Management Program for the Department of Energy which provides direction and coordination for management and dissemination of scientific and technical information resulting from the Department's research and development programs as well as environmental programs. A key current initiative of this Office is the DOE2000 program.

This Office also explores advanced energy concepts not currently included within departmental energy technology organizational components and manages the DOE-wide Small Business Innovation Research and Small Business Technology Transfer (SBIR/STTR) programs. The staff works closely with other ER offices, other departmental programs, and other agencies. It performs outreach activities directed to ER stakeholders and current and potential partners in its assigned areas of responsibility.

The Office of Computational and Technology Research is composed of three Divisions:

- Laboratory Advanced Energy Projects and Technology Research Division, Walt M Polansky, Director.
- Office of Scientific and Technical Information, Walt Warnick, Director, R. Charles (Chuck) Morgan, Manager.
- Mathematical, Information, and Computational Sciences Division, Dan Hitchcock, Acting Director.

3.5.1 Development of DOE-Wide Computational Grid

Ian T. Foster, Argonne National Laboratory

Advances in networking technology and computational infrastructure make it possible to construct large-scale high-performance distributed computing environments, or computational grids, that provide dependable, consistent, and pervasive access to high-end computational resources. These environments have the potential to change fundamentally the way we think about computing, as our ability to compute will no longer be limited to the resources we currently have on hand. For example, the ability to integrate TFLOP/s computing resources on demand will allow us to integrate sophisticated analysis, image processing, and real-time control into scientific instruments such as microscopes, telescopes, and MRI machines. Or, we can call upon multiple supercomputers to address ultra-large scientific simulations, perform computationally intensive analyses of distributed PetaByte datasets, or simply steer a remote simulation from our desktop.

Computational grid applications place new demands on both networks and computational infrastructure. For networks, these applications introduce a need for high peak performance, performance guarantees, and high-speed interfaces to supercomputers. For supercomputers, grid applications introduce a need for on-demand scheduling, rather than the batch scheduling traditionally used in supercomputer systems. In addition, sophisticated software is required to provide users with a uniform view of geographically distributed resources.

Scientists at Argonne National Laboratory are working in collaboration with researchers at other laboratories and universities to create a large-scale prototype of such future computational grids. This prototype, called the Globus Ubiquitous Supercomputing Testbed, or GUSTO, comprises supercomputers and other computer systems at national laboratories and other institutions across the United States and in Europe. These systems are connected by a variety of high-speed networks, including ESnet and vBNS. Advanced management software developed in the Globus research project is being deployed on these systems, with the goal of providing a uniform set of services for authentication, scheduling, information, and so forth. (For more information on Globus, see http://www.globus.org.) The GUSTO testbed was unveiled at the SC'97 conference, held in San Jose in November 1997. At the conference, ten different application groups used computers at 15 different sites in the U.S., Germany, and Sweden, totaling 3600 processors and 2 TFLOP/s of compute power. Computers at Argonne, NERSC, and LANL were included in the testbed.

3.5.2 Center for Computational Science and Engineering

Chuck Rendleman and Phil Colella
Lawrence Berkeley National Laboratory

The Center for Computational Science and Engineering is conducting research on the computational fluid dynamics of combustion. The application under development will help improve the understanding of the dynamics of such systems as diesel engines and industrial burners. Currently, we can model 3D systems using relatively simple reaction kinetics models. Our codes run effectively on the T3E using a library of software written to support Adaptive Mesh Refinement (AMR) algorithms for fluid systems on distributed memory parallel systems. The present algorithm is well suited to the existing T3E architecture in terms of CPU use, but we under-utilize the available system memory. Presently, each production simulation results in a considerable amount of output, corresponding to roughly one tenth of the available RAM per simulation phase. We find that we can easily use 25GB per production job. Each production job is one part of an overall series of jobs that model one physical experiment. For that reason we make heavy use of mass storage facilities. Currently our I/O Bandwidths scale well for processor to disk storage up to 256 processors. Some system configuration or AMR library software changes will be required to use 512 processors. Our software environment is heavily dependent on C++ and Fortran compilation systems. We are very disappointed with the current debugging and performance analysis tools available on the Cray T3E.

The next generation version of our software is likely to require upwards of 10 times the floating performance of the current code, mainly because of the use of a more realistic reaction kinetics model. In addition, the memory requirements will increase by roughly a factor of two. One implication is that our Disk requirements will likely increase by a factor corresponding to the CPU requirement, a factor of 10. Disk usage per production run will increase to 250GB, with corresponding increases need in the requirements for mass storage.

3.5.3 Computational Chemistry for Nuclear Waste Characterization and Processing: Relativistic Quantum Chemistry of Actinides

Robert J. Harrison and Ricky A. Kendall, Pacific Northwest National Laboratory

Research Objectives

We aim to develop and apply the methods of relativistic quantum chemistry to assist in the understanding and prediction of the chemistry of actinide and lanthanide compounds.

Computational Approach

The work involves determination of the electronic structure of molecules including relativistic effects necessary for heavy elements. Most calculations are very challenging and well suited to the CRAY-T3E. There are five major categories of activities:

- **Benchmarking of methods**: Detailed and systematic comparison of various theoretical approaches with each other and with experiment. Few such studies are available for rigorous relativistic methods and fewer still for systems containing actinides.
- **Application work**: We are studying the speciation of aqueous uranium (VI) carbonates. A detailed understanding of the actinide-carbonate-water system is essential to modeling the fate and transport of actinides in the environment.
- **Method and computer program development**: Existing programs are being parallelized and extended to enable calculations on larger molecules at higher levels of accuracy.
- **Computer Science**: Extensions of Global Arrays, parallel I/O, new linear algebra, meta-computing, and prototyping of new parallel programming tools.
- **Collaboration tools**: We use Internet video conferencing to augment weekly telephone voice conferences and are evaluating CORE-2000 and other collaboration tools.

Accomplishments

All components of the project are well underway. For benchmark purposes we have parallelized existing all-electron Dirac-Hartree-Fock and four-component second-order perturbation theory which treat relativistic effects very rigorously. Calculations at these levels of theory can take several hours on 64 nodes of the T3E. These results are being used to validate calculations using relativistic effective core potentials on larger molecules. Determination of the vibrational frequencies of the molecule in the figure took about 30 hours on 256 nodes of the T3E using the NWChem DFT module. The data from calculations on a sequence of related molecules is being correlated with experiment data. A spin-orbit configuration interaction code has also been parallelized and has just commenced production use on the T3E.

Significance

Most radioactive waste involves actinides, and their large atomic number implies that relativistic effects have important chemical consequences. Our implementation of relativistic quantum chemical methods on massively parallel computers will provide capabilities for modeling heavy-element compounds similar to those currently available for light-element compounds. The theoretical and computational methodology so developed will supplement current, very expensive experimental studies of the actinides and lanthanides. This will allow limited experimental data to be extrapolated to many other regimes of interest.

The program objectives will be attained through a multi-laboratory, multi-university and multi-disciplinary collaboration. These techniques will be applied to important molecular systems and processes, including the interaction of actinides with: 1) organic complexing agents present in tank wastes; 2) natural aqueous systems (carbonates), in order to better understand fate and transport in the environment; and 3) new materials, such as phosphates and amides, for the design of in situ remediation technologies and separation systems.

DFT/RECP calculations have been performed on this polymeric uranium carbonate (c.f., Figure 38), and results agree well with experimental observations[30].

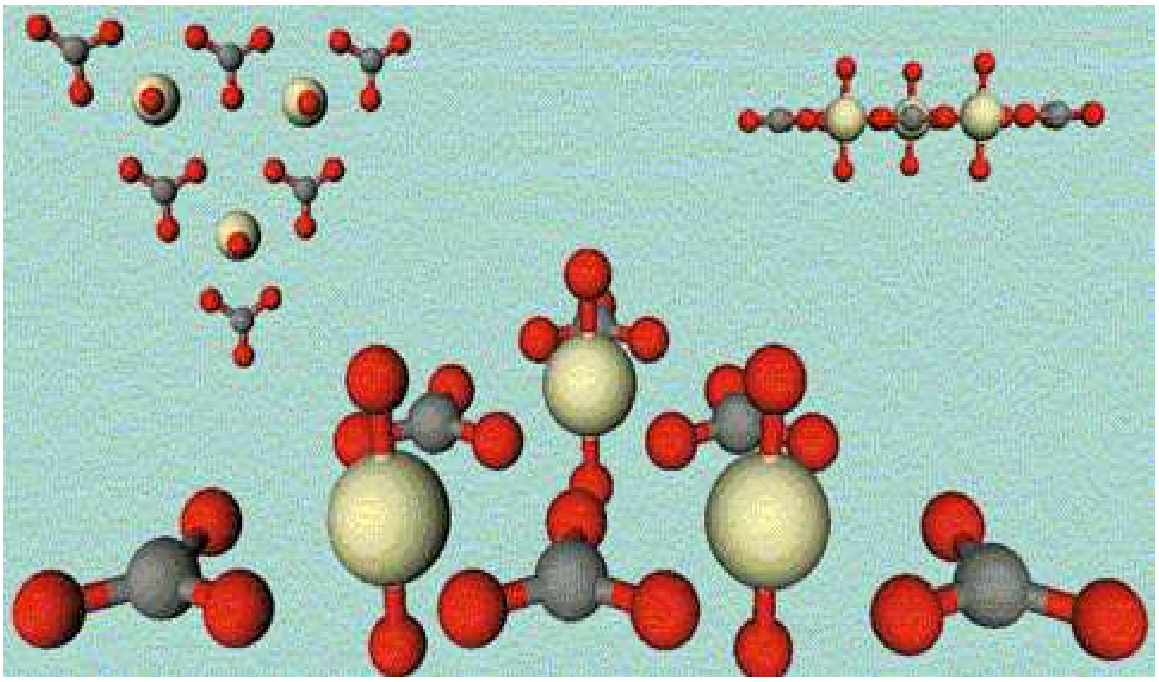

Figure 38: Polymeric uranium carbonate system

[30]URL: http://www.emsl.pnl.gov/pub/proj/tms/hpcc_actinides

Future Requirements

Within the scope of this project and the accuracy requirements of future scientific simulations the default minimum specifications for an MPP system are:

- 512 Mbytes memory per node,
- per-processor disk I/O bandwidth of 8 Mbytes/sec,
- 700 Mflops peak performance per-processor,
- 1 micro-sec communication latency,
- 500 Mbytes/sec communication bandwidth.

To reach all the simulation goals the project will need between 1.5 and 2 cpu months of this MPP system.

4 The Energy Research Computational Resources

The computational resources of the High Performance Computing Access Center (HPCAC) are the focal point of production high performance computing for the DOE Office of Energy Research (DOE-ER) scientific programs. These resources include not only the obvious assortment of hardware but also the application software, development software and the network connectivity to the national laboratories and academic institutions that utilize the facility. The HPCAC is a unique facility shared by essentially every research area supported by DOE-ER. Moreover, the NERSC facility aims to be a showpiece within the distributed computing environment for DOE-ER, where components of the computing environment are first prototyped and quickly put into large scale production use.

4.1 Hardware

The hardware resources housed in the HPCAC serve the needs of both capacity and capability computing. Capacity computing engines provide high computational speed and rapid throughput for all but the largest jobs in the ER spectrum. Capability computing engines have special features (e.g. very large memory, very large local disk, high speed archive and data transfer engines, etc.) for jobs that require some combination of these features. In the following sections we discuss briefly the different classes of supercomputers in use for computational science and the supporting hardware infrastructure to make effective use of these supercomputers. It should be noted here that the definition of a "supercomputer" is a very fluid and time dependent definition.

The HPCAC user base has it's primary experience with parallel vector supercomputers, but usage of RISC-based symmetric multiprocessors (SMP) systems and of massively parallel processing (MPP) systems is increasing. NERSC has offered its own MPP (A CRI/SGI T3E) since 1996, and now offers a mix of MPP and PVP hardware; a RISC-based SMP is scheduled to come on-line at NERSC in the Fall of 1998. The current and expected utilization and the experience gained will determine the hardware and software requirements for the next major expansion of the facility.

4.1.1 Parallel Vector Supercomputers

Parallel Vector Processor (PVP) supercomputers such as the CRI C-90 and J-90s are shared memory multiprocessors (SMPs) that are the computational workhorses of the current HPCAC facility. The C90 has a peak performance of approximately 1 Gflops per processor with 16 processors and 1 GB of memory. The J-90 Complex has a per processor peak performance of approximately 250 Mflops per processor, with 96 processors and 32 GB of memory. Thus the J-90 complex exceeds the C90 performance, at a fraction of the cost. However, any single-threaded simulation will perform at most a factor of four slower.[31]

The current top-of-the-line PVP equipment available from the vendor community is represented by the Cray T-90 and the NEC SX-5. For instance, the Cray T-90 has a peak performance of close to 2 GFlops per processor for up to 32 processors. The future evolution of PVP systems planned at NERSC is to phase out the C-90 at the end of CY 1998, and to upgrade the J-90 line with improved processors as they become available, and depending on budget and demand. An upgrade of at least 25% of the J90 cluster is currently budgeted. Without improved PVP capability in NERSC's next procurement, this marks a transition for the facility in that they will no longer provide the highest-end vector processing hardware available, but instead will try to meet the demand with larger aggregate capability using slower but more cost-effective solutions.

4.1.2 Massively Parallel Processing Supercomputers

Distributed memory Massively Parallel Processors (MPP) systems based on commodity chips have become increasingly important to supercomputing because of the low price performance ratios and because systems of this kind can offer extremely large memory. For suitable problems, and in particular for very memory intensive problems, these systems are capable of much higher computational performance and they are able to address certain Grand Challenge and Capability computing problem that cannot be done with PVP systems in any reasonable time frame.

The HPCAC T3E MPP is one of the largest configurations available anywhere. The system installed in July of 1997 is a 512 processor T3E-900, with 1.5 Terabytes of disk and 132 GB of memory. At the time of delivery, it was the largest I/O configuration ever built. This system has a peak performance of 900 Mflops per processor for a theoretical peak of slightly less than 0.5 Teraflops. The system has a measured performance of more than 0.25 Teraflops (265 Gflops on 512 processors). It is allocated to Grand Challenge and very large capability problems. In FY 98, the HPCAC will provide roughly 5 times the computing resource on the MPP systems[32] as exists on the PVP systems. Currently the MPP user base is smaller[33], and many applications will require restructuring to effeciently use these MPP systems.

There also exists Shared Memory Processors that are not Parallel Vector, such as the IBM SP and SGI Origin, and Sun HPC series. NERSC currently has prototype systems using this architecture (a pair of Sun Enterprise 4000 systems) and in FY 98 will receive a 64-node Origin 2000 to evaluate the state of a hybrid architecture for clustering SMPs.

[31] This is dependent upon the level of vectorization and the memory access patterns used in the application. This also demonstrates the need for education in the user community in using parallel mechanisms for critical applications on the PVP systems; so that users can access the aggregate performance of the larger J-90 system.

[32] based only on theoretical peak performance

[33] The active user base is roughly 2 PVP users to 1 MPP user

Indeed, the J-90 complex is already an early implementation of this style of computing. It is believed these systems will provide both fine and coarse grain parallelism in a single application.

Future directions of this system include the incorporation of clustered SMPs. This could be clustering using custom switches within an MPP to clustering using commodity network components. The number of processors per SMP is also a technical challenge, whether there will be a small number (2, 4 or 8) or a large number (32 or 64) of processors sharing local memory within an SMP node. One path, small SMPs clustered with specialized connections could be viewed as an evolution of today's MPP architecture. The other path, large SMPs clustered with commodity or custom networks, could be an evolution of today's SMP architecture. Adding the complexity of providing distributed memory access in a way that enables vector computing creates a large range of possible paths.

NERSC has begun exploring these directions in preparation for the next major acquisition in 1999, and the one in 2001/2002. These acquisitions could lead to a system with hundreds of processors that has Teraflops of measured performance in 1999, and 10s of Teraflops in 2002.

In all these scenarios, the programming methods are different than for the PVP or current SMP systems. In order to do increasingly more demanding problems that track the computational price performance curve, HPCAC will have to work with clients and providers to develop system and programming environment software which allows applications to capture a significant part of the aggregated computational power these new systems provide. New models of computing services are also needed to fully realize the potential of the new technology.

4.1.3 Symmetric Multiprocessors

Symmetric multiprocessors (SMPs) are available from a wide range of workstation vendors in various configurations, with up to 64 nodes in current systems from Sun and SGI. These cost-effective systems are comprised of similar or the same components (processor, memory, and disk subsystems) that make up vendor MPP offerings, but the typical programming model for applications on SMPs and most MPPs is different in that explicit message passing and management of memory locality is not required on the SMPs. However, SMPs can be programmed in the same fashion (e.g., with MPI) as MPPs so one component of their capability is providing an ideal development environment for MPP application software. Medium-size SMP systems are the capacity and capability engines for many departmental and working groups within the ER and university research community, although the largest systems of this kind are already out of the range of all but a few of the ER research projects.

4.1.4 Workstations

Workstations are a necessary local resource for the computational scientist. They serve a myriad of functions from desktop publishing, software development, to production scientific computing or capacity computing. Workstation components are essentially the commodity parts that make up both MPPs and SMPs available from the vendor community. Entry level Unix workstations cost less than $10K and high-end models seldom cost more than $50K. They are quite cost effective since the allocation of scientist-to-workstation is usually a few (less than 5) to 1 or better.

Clusters of workstations using tools like P4, MPI, Parasoft Express, PVM, or TCGMSG can enable software development activities for MPP applications. However, at present these systems still fall short in memory bandwidth and in scalability for truly large scale simulations. This trend may change with the advent of low level programmers interfaces for clustering systems being developed by the vendor and research communities. This still requires applications that are latency tolerant both in the memory and communication subsystems.

4.1.5 Database Servers

There are two primary functions that database servers fulfill; one is at the user archival level and the other is as information servers based on query languages. Examples of the archival type are the common file system (CFS) and the high performance storage system (HPSS) that are used at NERSC for data storage and data migration. The second function takes the form of integrated hardware and software to perform specific data base aggregation and search functions. These include database servers that are on the world wide web and those available throughout the ER user community. As the complexity and size of data mining problems increase, fully integrated data management systems will need to be developed.

4.1.6 Scientific Visualization Services

A general scientific visualization resource will contain high performance visualization station(s), multimedia station(s), and the associated software for applications or analysis codes to drive these tools. Visualization stations are distinguished from other desktop workstations by their high graphics performance, large memory, and large direct attached auxiliary storage.

The visualization and multimedia equipment must be located at the user sites. The scientific data represented as images, will need to be stored on either videotape or disk, and redisplayed with a high-resolution video system. In addition, basic editing capability should be available to construct presentations of scientific material. The system must be available for use by the scientific teams that need these resources and for presenting information to visitors. A mechanism should also be explored to facilitate at least scaled down versions of the visualizations to the WWW.

The current state-of-the-art graphics accelerator can render about two million polygons per second. To allow ER scientists to begin viewing complex scenes at interactive rates, a $3x$ or six million polygons per second is probably required. Depending on the visualization rendering paradigm, pixel fill rates could become the bottleneck to rendering performance. A high performance graphics system (HPGS) must provide a minimum of 160 million textured, anti-aliased pixels per second. This fill rate is available on a few systems today but will become more common in future generations.

4.1.7 Special Computers

The role of special or experimental computers is still important to the ER scientific community. Some special computer systems may be ideal for specific grand challenge problems but with the current fiscal constraints, dedicating resources to a single grand challenge problem are unlikely. The real role for experimental computers is to lay the ground work for the next generation of production computer systems. Experimental computers allow the evaluation of new hardware and software architectures. The evaluation must be performed in an application environment that stresses every component of the computer (hardware and software). The high performance computing research centers (HPCRC) and access to such systems provided by programs such as NERSC's MPP program have filled this role well. Support for these activities and centers should be continued. The success of an MPP system in the NERSC facility will be built upon the lessons learned at the HPCRC supercomputer centers.

4.1.8 Network

The Network situation has changed significantly since the last requirements update. Clearly the growth of the global internet and the World Wide Web (WWW) has increased the overall utilization of the internet; the demand is higher than the evolving supply. ESnet has taken significant steps to make the connectivity and required bandwidth available to the ESnet community. ESnet has done a stellar job in this regard. The advent of distributed computing tools such as desktop video conferencing tools, multimedia electronic mail, and experimental control systems demand continued expansion of network bandwidth. The connections to the internet fabric for specific sites will likely need to be enhanced for ER scientific efforts since the global "free" internet is heavily burdened by WWW traffic. This will require T3 or better links to each ESnet site and certain other critical facilities. In addition, international lines will need to be upgraded as the research community becomes more global in character.

As the scale of scientific collaborations and the volumes of data increase for the ER scientific programs the importance of the local and wide area network increases. The role of the local area network in coupling the massive data storage systems with the high-performance computational resources has become as important as the inter-processor communication subsystems of the parallel supercomputers. The wide area network (WAN) and the coupling between the HPCAC and the WAN has a critical role in enabling the collaborative science that spans many laboratories and institutions. As the size of datasets have increased it becomes impractical to duplicate this data at several sites and an application running at any site needs effective access to both local and remote data storage systems.

4.2 Software

The computing axiom "Nothing happens in a computer without good software" is still true today and even more so with the advent of MPP technology. The software infrastructure is what a user truly interacts with and is now becoming a measure that distinguishes MPP vendor offerings. The infrastructure includes distributed computing tools, operating systems, network software, compilers debuggers, performance tools, standards based libraries (e.g., BLAS, LAPACK, etc.), job control and queuing systems, as well as third party applications that are of great importance to many end users in the ER community. Below we discuss aspects of the software infrastructure that are required by the ER scientific computing user community.

4.2.1 Operating Systems

The operating system (OS) and the base level software infrastructure (BLSI) must be a robust software suite or the computational resource will only be utilized by the most sophisticated users. The OS and BLSI must allow users to (in a secure fashion) manage files, run jobs (both interactively and in a batch facility), export data, interact with other computational resources and significant third party applications, etc.. The BLSI often acts as an extension of the OS. For example, automatic file migration to the hierarchial storage system may not be a component of the base OS but is an important component of the BLSI.

4.2.2 Libraries, Compilers, Debuggers, and Performance Analysis Tools

Standard language compilers, software libraries, debugging and performance tools are essential components of any high performance computing resource. The languages that must be available are FORTRAN-90, C, C++ and Assembler. These should have fully integrated debug and performance analysis tools. The computational science community as a whole will build applications that use one or more of these languages, therefore, language interoperability is also important.

For MPP supercomputers it is important to support third party applications and libraries for data management and performance measurement. For example, the UPSHOT performance tools (from Argonne National Laboratory) should be available since they are in use at many DOE sites. Parallel software development and environment tools such as HPF, MPI, PVM, and the Global Array tools are used by a variety of ER computational scientists and access to these tools on the central NERSC MPP resources are essential. Furthermore, software libraries such as ScaLAPack and vendor specific parallel linear algebra tools are required for the development of many scientific applications.

4.2.3 Distributed Computing and Unified Production Environments

With the advent of software infrastructure tools supported by the vendor community such as the Distributed Computing Environment (DCE), one new challenge to, and opportunity for, the ER user community is to harness the power of physically distributed computational resources. The development of tools for distributed computing by the research community is also well underway as evidenced by the I-WAY demonstrations performed at Supercomputing 95. The Distributed Computing Coordinating Committee (DCCC) is providing various infrastructure tools for the DOE-ER community. The challenge for the ERSUG community is to help focus DCCC efforts on tools that will enable ER scientific efforts and application development.

NERSC is planning a Unified Production environment to facilitate the interaction of the various DOE computational communities. The basic goals for the Unified Production environment are:

- The centralized NERSC environment must serve all the high-end needs. No local user environment, no matter how rich, can hope to offer peak computational speeds of on the order of 200-500 GFlops or storage capabilities of on the order of 200 TByte. In short, a centrally located capability engine (MPP, cluster of SMPs or PVPs, or heterogeneous system) within a Unified Production Environment, with adequate internal balance in support infrastructure (disk, tertiary storage and complementary high-end capability shared memory platforms) remains a major goal.

- The NERSC environment must continue to offer all the services required by the high-end user with limited local resources: this includes code development, computing and assimilation capabilities, as well as information services and consulting. Information services, consulting and collaborative ventures with NERSC staff are becoming more important because of the complexity in the new programming and assimilation environments. Without expertise, the capability latent in the hardware will not be realized.

- Ancillary services must be easily accessible, and must be offered on the appropriate platform. The MPP is not currently a viable platform for code development; NERSC must facilitate the user's migration through all the phases of computational research projects by providing a system administration environment that makes this easy, and that allows the user to select the platform most appropriate to the computational task at hand. That is what is meant by the Unified Production Environment.

- For users with a rich local environment, the Center must strive to act as the high-end complement to the local environment. Many of the aspects of unification practiced within the centralized NERSC environment can be practiced within the NERSC - user (that is, centralized - remote) environment. Work in this area could begin with one or two pilot projects involving remote sites.

- As scientific collaborations have increased in size the role of distributed computing and a unified or seamless environment has grown in importance. The centralized environment at NERSC needs to be both compatible with resources at other large (centers) and small (desktop) sites in terms of application software and security services. Additionally effective network connections for data access between these sites is crucial to the scientific missions.

4.2.4 Scientific Applications: Third Party

NERSC currently supports a wide third party variety of software applications from the independent software vendor community, (such as Objectivity/DB, Iris Explorer, Gaussian, NASTRAN etc.). The continued support of these applications and the staff to maintain and provide consulting services is required to allow the ER scientific community to do their science. The obvious fiscal restraints mean that NERSC should continue their monitoring of third party application utilization and make appropriate additions and deletions of the third party applications that NERSC supports. This should include user-site fiscal responsibilities for applications that are not widely used. In other words, if an application is required by a very limited number of users then the NERSC facility should not have to bear the full burden of the cost and maintenance of this application software.

4.2.5 Scientific Applications: User Developed

Many ER scientists utilize in-house developed applications to do their science. The recent down-size of all federally funded research efforts has strengthened the push to do better science at a reduced cost. Currently redundant efforts are funded, as they should be, because this allows different scientists and disciplines to investigate the same or similar problems thus expanding the scientific knowledge base. The cost reduction efforts include the reduction of some of this redundancy and minimizing reinvention of the same tools and methods. The NERSC facility should provide resources to share these in-house developed applications in addition to software repositories that research disciplines want to establish. Sharing the application software to others helps meet the cost reduction goals outlined here and also enables more robust application software. New users to any software suite uncover bugs and design limitations of that software.

4.3 Basic Research and Human Resources

With respect to basic research and human resources the user community needs to be both involved in and have access to the results of extensive basic research efforts. The High Performance Computing and Communications Initiative Grand Challenge programs and the research efforts in the High Performance Computing Research Centers have provided much valuable insight. The NERSC user community now has the opportunity to put these developments into a true production environment. Cross-disciplinary exchanges and collaborations in addition to the "Users Helping Users" program should be encouraged to implement these new capabilities and identify additional opportunities in computational science.

4.4 The High Performance Computing Access Center Services

In addition to the standard support services that the NERSC consulting staff there are additional services that are required. The standard services include consulting on and facilities for:

- User Account initialization and security.
- Basic system utilization (batch and interactive computing).
- Third party or ISV application access and utilization.
- Tools for Principle Investigators and account managers to allocate their available resources and plan their own resource utilization.
- Tools that report the status of NERSC systems to users.
- Mechanisms to transition applications to MPP architectures (e.g., MPP program and applications group support).

Additional standard services that should be added are:

- Mechanisms for users to share their technology (c.f., Section 4.2.5.
- Mechanisms for users to report their results on-line (e.g., space for WWW pages, and/or registration of links at user sites).
- Support for the "users helping users" efforts.

One component of the decision to move the NERSC facility to Lawrence Berkeley National Laboratory was the ability to strengthen the ties of the facility to ER research programs. NERSC has planned to build strategic partnerships with various research efforts of the user community. We strongly encourage and support these efforts and hope that this will include not only traditional supercomputing efforts but those targeted at the utilization of the MPP resource.

5 Specific Recommendations

Based on the discussion presented, we offer these specific recommendations:

- Expand the high performance computing resources available at NERSC.
- Focus efforts to transition appropriate applications to parallel systems.
- Strengthen the networks.
- Encourage collaborations and integration among ER physical, computational, and computer science programs.
- Encourage software development.
- Encourage science of scale by enhancing the effective use of the capability computing systems.
- Encourage expansion of data intensive computing capabilities.

5.1 Expand the high performance computing resources available at NERSC.

The high performance computing resources must be expanded in response to the ever increasing, complex, and diverse programmatic mission requirements of the ER research community. The procurement of the next major system at NERSC is already underway, aiming for delivery in 1999. At present the computational resources at NERSC are saturated and large scale computational science efforts in ER are starved for resources. We recommend that the procurement specifications emphasize the importance of system performance on large scale computational problems. At the same time, the procurement must ensure an appropriate balance of computational equipment to support the wide range of large-scale applications that use the current facilities.

Not only must the aggregate capability be enhanced, but special attention must given to improving the component subsystems of each class of capability resource (PVPs, MPPs, SMPs); e.g., fast CPUs, increased memory capacity and memory bandwidth, increased interprocess communication (mechanisms for both shared and distributed memory programming models) and disk I/O performance. It is clear that improving these component subsystems will require next generation technology in both hardware and software.

The next system at NERSC will co-exist with the current T3E MPP and the J-90 PVP cluster (with upgrades). Access to all classes of production capability resources (PVPs, MPPs, SMP systems and clusters, c.f. section 4) is essential. Access to the many classes of computers has the benefit of allowing the user to appropriately match the resource capability to the scientific task or job. The current resources meet one HPCCI goal of exporting parallel computing technology to a broad user base. An additional benefit of the co-located resources at NERSC is that it allows to set up test-beds for any distributed computing environment in a production facility.

5.2 Focus efforts to transition appropriate applications to parallel systems.

There are many large scale user applications that are not presently operational on a parallel platform, but that would benefit from being converted to use one. Efforts to transition applications to the NERSC MPP and future parallel systems should be focussed on these applications. The DOE High Performance Computing Research Centers have done an excellent job in providing resources to the grand challenge community to build their experience base. NERSC is in a unique position to provide both development and production systems to the majority of the ER scientific computing user community.

The choice of a portable programming model and the optimization of software for a particular parallel platform are often in conflict. There is a need to get applications working quickly on the new resource and a need to be portable since the high performance computing systems available for the next NERSC procurement have yet to be completely designed. To meet these multifaceted requirements, including user training, NERSC must have a highly qualified and diverse parallel computational science and applications group that should build strategic partnerships and collaborations with the core ER physical and computational science research efforts.

5.3 Strengthen the networks.

The network is the glue that holds the computational resources and scientific collaborations together. The use of the World Wide Web (WWW), email services, data exchange, remote visualization, etc., are daily operations of professional scientists. Many sites operate distributed file systems (AFS, DFS, NFS, etc.) and when the local network is down work is not possible. The operation of networks between sites are handled by ESnet. We recommend that specific site-to-site connections and increased bandwidth be added on a continuing basis to meet the rapidly increasing communication requirements, but defer any specific recommendations to the various ESnet governing committees.

5.4 Encourage collaborations and integration among ER physical, computational, and computer science programs.

New architectures and the constant change in the computational technology area require large investments of time for development efforts to maintain and augment scientific application codes and for understanding the implications of these changes with respect to algorithmic choices. It is quite clear that the advances of computational science over the last two decades are due to the coupling of advances in both computational technology and algorithm design. The goal or "holy grail" in this area is the single source that encompasses the entire software suite's functionality which is portable from the workstation to the MPP system and is also extensible or adaptable to architectural changes. There are two conflicting requirements here in that in short-term the application must run efficiently on the resources at hand and in the long-term the application must migrate to new technology.

Although, progress can be made by individual physical and computer science efforts, many research efforts have made significant progress by the strong integration between physical scientists and computer scientists. This idea can and should be extended to interaction among physical science efforts as well. NERSC is now in a position to not only be more closely involved in ER research programs but also to facilitate collaborations by acting as the "Collaboratory meeting place" for ER scientific efforts. NERSC has taken steps via strategic partnerships with specific areas and our recommendation is to continue this and associated efforts. We encourage collaborations and integration among ER physical, computational, and computer science programs, and also encourage that the funding and performance reviews include a measure of the value of such integrated collaborations.

5.5 Encourage software development.

Encourage software development for the advancement of ER missions. Extensive software development must be encouraged to take advantage of current and new opportunities for all of ER computational sciences. The functionality and performance of application codes can be directly impacted by software development efforts within other areas of ER. Large applications, as discussed above, have both long-term and short-term requirements that pose significant managerial problems. These applications represent many man-years of development and cannot be easily re-engineered for MPP systems, unless good software development practices are followed. It is nearly impossible to simply port a PVP application to an MPP and have it work efficiently.

New or continued software development efforts should include mechanisms for the redesign of applications to allow for architectural changes. These efforts should also be done in cooperation with computer scientists. Without the proper computer science in the design and implementation of software applications, they become difficult to maintain and change. NERSC has neither the funding or the mandate to take on these tasks for all development efforts, but the facility and staff should serve as a focal point or facilitator in these activities and where appropriate as a partner.

Finally, the utilities (compilers, debuggers, performance analysis tools, visualization tools, mathematical and scientific libraries, data management systems, etc.) and operating systems on the computational resources need to be robust and full featured. The user interface to these utilities must also be intuitive and easy to use to minimize the learning curve and get the most from their utilization. This is especially important with respect to these tools on the MPP system.

5.6 Encourage science of scale by enhancing the effective use of the capability computing systems

. Supercomputers should be used primarily for the applications that require the specific and special capabilities of the resource. Although local computing has expanded enormously since the last requirements document, little has been done to eliminate the capacity vs. capability computing problem. It is quite true that a capability supercomputer without significant capacity is not a useful resource. The use of supercomputers as general-purpose capacity engines is driven by the fact that these resources are pre-paid ("free or allocated") systems that are very stable and well maintained. Unfortunately the primary mission of supercomputers has been overloaded by filling the capacity requests. This situation must be reversed but there are numerous and obvious political and sociological obstacles that will be encountered.

There are many ways to accomplish this. One is by expanding the local and distributed computing resources further, a place for many mid-scale and small-scale projects/tasks (c.f., section 2.1) will exist. These services need to be integrated with the supercomputer systems to allow some omnipotent scheduling process (currently users that understand available resources and their applications) to move appropriate tasks to the local and distributed computing class. A compatible applications software environment across the range of computational systems (local, distributed, high performance) is necessary to achieve this goal.

The numerous political and sociological obstacles are compounded by the infancy of true distributed computing environments and the fact that there is little experience in the user community in using coupled but distributed resources and tools. The NERSC facility is again in a unique position to prototype and foster the development of such tools and environments in a production facility. This may require the installation of capacity resources at the NERSC facility and will definitely require the cooperation of the user community.

5.7 Encourage Expansion of Data Intensive Computing Capabilities

For some classes of scientific endeavors in Energy Research which are very data intensive, the capability to select both large and small samples from very large data sets is as equally critical to the science as is applying compute cycles to those samples once they are selected. This will involve the development both of software and necessary networking hardware for this capability to be effective. Selecting and staging files from shell scripts via ftp is simply inadequate for large complex data sets. A true object database capability is required for the scientist to be able to select the appropriate data sets or "data objects" which have meaning for the scientific study. Two issues involved in this capability are optimizing data layout so that selecting small samples is efficient as well as providing adequate network bandwidth so that large samples are retrieved in a timely way.

6 ERSUG Action Items

Based on the requirements analysis in this document the following action items have been delineated for action by ERSUG with guidance by ExERSUG.

- Expand the "Users helping User" working group to facilitate the transition to MPP computing.
- Determine if the user community should establish "application benchmarks" or template benchmarks that can be used to characterize the performance of each class of computer outlined in this document, e.g., from workstation to MPP (c.f., section 4).
- Resolve to strengthen our participation in the design and implementation of distributed computing applications and tools in conjunction with the DOE2000 and similar groups.
- Work with NERSC to improve the mechanisms to match jobs/tasks to appropriate computational resources.

ExERSUG Membership

The members of the executive energy research scientific computing users group (ExERSUG) represent each of the five offices in the office of Energy Research of the US Department of Energy. ExERSUG composes the primary authors of this requirements document.

Office of Basic Energy Sciences

Rajiv Kalia
Concurrent Computing Laboratory for Materials Simulations
Louisiana State University
Baton Rouge, LA 79803-4001
Office no: (504) 388-1112
Fax no: (504) 388-5855
Email: kalia@bit.csc.lsu.edu

Ricky A. Kendall, Past Chair
Staff Scientist, Theoretical and Computational Chemist
MSIN K1-83, PO Box 999
High Performance Computational Chemistry Group
Environmental Molecular Sciences Laboratory
Pacific Northwest National Laboratory
Richland, WA 99352-0999

Office no: (509) 375-2602
Fax no: (509) 375-6631
Email: ricky.kendall@pnl.gov

Office of Computational and Technology Research

Phil Colella
Staff Scientist Center for Computational Sciences and Engineering Lawrence Berkeley National Laboratory Mail Stop 50D-106 Berkeley, CA 94720 Office no: (510) 486-5412 Email: pcolella@euler.me.berkeley.edu

Mike Minkoff
Computer Scientist
Math and Computer Science Division
Argonne National Laboratory
Office no: (630) 252-7234 Fax no: (630) 252-5986 Email: minkoff@mcs.anl.gov

Office of Fusion Energy

Bastiaan J. Braams, Chair
Research Assoc. Professor
Magneto-Fluid Dynamics Division
Courant Institute
251 Mercer Street
New York University
New York, NY 10012
Office No. (212) 998-3276
Fax no: (212) 995-4121
Email: Braams@mfdd2.cims.nuy.edu

Viktor K. Decyk
Research Physicist
University of California
405 Hilgard Ave.
Los Angeles, CA 90024-1547
Office no: (310) 206-0371
Email: ecj9vkd@mvs.oac.ucla.edu

Steve Jardin
Deputy Head, Theoretical Department
Princeton Plasma Physics Laboratory
Princeton University
Princeton NJ 08543
Office no: (609) 243-2635
Fax no: (609) 243-2662
Email: jardin@pppl.gov

Alan D. Turnbull
Staff Scientist
General Atomics
PO Box 85608
San Diego, CA 92186-9784
Office no: (619) 455-4042
Fax no: (619) 455-3586
Email: turnbull@gav.gat.com

Office of Biological and Environmental Research

Brian Hingerty, Past Chair, current Vice-Chair
PO Box 2009, MS 8077
Biology Division -Bear Creek Road
Health Sciences Research Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831-8077
Office no: (615) 574-0844
Fax no: (615) 574-1274
Email: beh@ornl.gov or u16127@nersc.gov

Jerry Potter
Deputy Director, Program for Climate Model Diagnosis
PO Box 808
Lawrence Livermore National Laboratory
Livermore, CA 94550
Office no: (510) 422-1832
Fax no: (510) 422-7675
Email: potter@oliver.llnl.gov

Office of High Energy and Nuclear Physics

Kwok Ko
Physicist
Stanford Linear Accelerator Center
PO Box 43492 (2575 Sand Hill Road)
Stanford CA 94309 (Menlo Park CA 94025)
Office no: (650) 926-2349
Email: kwok@slac.stanford.edu

Steve Pieper
Physicist, Nuclear Physics
Physics, Bldg. 203
Argonne National Laboratory
Argonne, IL 60439
Office no: (630) 252-4232
Fax no: (630) 252-3903
Email: pieper@theory.phy.anl.gov

Robert Ryne
LANSCE-1, Mail Stop H817
Los Alamos National Laboratory
Los Alamos, NM 87545
Office no: (505) 667-2839
Fax no: (505) 665-2904
Email: ryne@lanl.gov

Doug Toussaint
Professor
Physics Department
University of Arizona
Tucson, AZ 85721
Office no: (602) 621-2554
Fax no: (602) 621-4721
Email: doug@klingon.physics.arizona.edu

Emeritus Members of ExERSUG

Jack Byers, Past Chair (OFE)
Staff Scientist
Magnetic Fusion Energy
PO Box 808
LLNL
Livermore, CA 94550
Office no: (510) 422-9861
Fax no: (510) 423-3484
Email: byers@kaje.llnl.gov

Thom Dunning, (BES)
Batelle Fellow,
PO Box 999
Environmental Molecular Sciences Laboratory
Pacific Northwest National Laboratory
Richland, WA 99352
Office no: (509) 375-6863
Fax no: (509) 375-6631
Email: th_dunning@pnl.gov

Bill Herrmannsfeldt (OHNP)
Physicist, High Energy Physics (accelerator theory)
Stanford Linear Accelerator Center
PO Box 43492 (2575 Sand Hill Road)
Stanford CA 94309 (Menlo Park CA 94025)
Office no: (650) 926-3342
Fax no: (650) 926-5368
Email: wbhap@slac.stanford.edu

Rich Nebel, (OFE)
Group Leader for Plasma Theory, T-15
MS B217
Los Alamos National Laboratory
Los Alamos, NM 87545
Office no: (505) 667-7721
Fax no: (505) 655-5757
Email: rnebel@lanl.gov

7 Acknowledgments

The ExERSUG Chair and the docuement editor would like to thank all the ExERSUG members for a thorough reworking of the original text. ExERSUG would like to thank all those who supplied research highlights for our requirements analysis efforts. We would also like to express our appreciation to those, too numerous to mention, that supplied supplemental information for this report and also acknowledge the support of many staff at NERSC who supplied countless bits of technical information and critical reviews of the draft documents as they were generated.

In the following Appendices we summarize relevant supporting documentation and reports from the overall computational science community. We direct the reader to the complete document where appropriate.

A Requirements for Supercomputing in Energy Research: The Transition to Massively Parallel Computing

Only the executive summary of this document is enclosed here. This document is known as the Blue Book and is a Technical Report (DOE/ER-0587) available from DOE/ER.

A.1 Executive Summary

Key Recommendations:

- The procurement of a massively parallel computer in the Access Program capable of 200-300 billion operations per second for delivery in Third quarter FY94.
- Integration of the massively parallel computer with mass storage facilities and an appropriate distributed computing environment.
- Research and development of the remaining components of the software environment required for production use of parallel computing.
- Integration of the software achievements of the High Performance Computing Research Centers (HPCRCs) in the initial production environment and the use of HPCRC resources to aid in the transition to general use of parallel computing in the Access Program.

The mission of the Department of Energy's Office of Energy Research is accomplished by programs in Fusion Energy, Basic Energy Sciences, Health and Environmental Research, High Energy and Nuclear Physics, Scientific Computing, and the Superconducting Super Collider. All of these programs depend critically on computational modeling. The application of large-scale computation is necessary for the understanding of fundamental scientific problems in the areas covered by Energy Research programs, and the design of multi-million dollar experiments cannot be done without computer simulation of the experimental apparatus and its components.

This study analyzes the needs of the range of disciplines in Energy Research (ER) in which large-scale computation is a crucial component. The goals are clear in each of these areas, and in most of them good estimates can be made of the computational capability needed to accomplish both near- and long-term milestones. In the cases of climate studies and quantum chromodynamics of elementary particles, for example, the problems are studied by methods which make use of numerical grids. To achieve the needed accuracies, the grids must be made finer than those it is possible to treat with today's computers. In other cases, the interest is focused more on moving from two to three dimensions or additional physics rather than on finer resolution in the models. The ways in which the demands of these calculations scale are understood and provide estimates of the speed and capabilities of the computational hardware necessary to perform them. The scaling of quantum chemical and molecular dynamics calculations, computational fluid dynamics, and plasma physics, among many other disciplines, are all understood.

From these observations this study concludes that across the spectrum of research and design in Energy Research, many (but not all) critical programmatic goals can be accomplished with a computer capable of approximately 200-300 billion operations per second. It is expected that this machine will be based on a massively parallel computer architecture. Reasonable estimates based on current trends in computational technology place the availability of such a machine sometime in calendar 1994.

The development of massively parallel computers by several manufacturers has progressed from first generation machines to second or third generation machines with progressively better hardware and systems software. The efforts of vendors and those of the DOE High Performance Computing Research Centers (HPCRCs) have, in combination, brought us to a point where it is practical to employ a massively parallel machine capable of delivering 200-300 GFlops performance as a production platform. As with most of the early-model supercomputers acquired by the supercomputer Access Program in the past, some work will be required to produce a robust software environment for production. But with new generation machines, it is clear that a massively parallel machine will be an invaluable production platform for Energy research users, especially in conjunction with current supercomputer resources in the Access Program. Making the first massively parallel production platform available alongside conventional supercomputers will allow users to make the transition from the older to the newer technology in an orderly manner.

This study recommends that such a system be acquired in early 1994. An appropriate timetable to do so is to finalize the requirements for the procurement in the first quarter of calendar 1993, and begin the procurement immediately thereafter. Evaluation of vendors' offerings should be completed by the end of the third quarter of calendar 1993, and the machine and its configuration chosen by that time. Budget authority is required for a third quarter FY94 installation.

Once the vendor is chosen, the Access Program should work intensively with the HPCRCs to move as much as possible of the software environment they have developed to the coming production platform.

This study recommends that a small percentage of the resources of the HPCRCs be set aside to help in the conversion of applications codes to massively parallel architectures. For the larger community of users of the Access Program, the transition to massively parallel machines is impeded by a combination of lack of access and expert help. Prior to the arrival of the massively parallel machine recommended here, the Access Program should provide and via consultation and programming assistance to a subset of it users who are granted access to the HPCRCs or other DOE parallel computing facilities.

In a number of ER research disciplines the requirements for the storage of computed data also increase dramatically as research moves to the problems which must be solved over the next three to five years. The entire configuration of a critically needed computational resource for Energy Research is therefore addressed in this study, especially including data storage requirements. By the time of the recommended procurement, it will be possible to satisfy the requirements of the programs by integrating such a machine into the appropriate data storage and distributed computing facilities.

Further conclusions of this report specify the necessary computational software which must be available on such a computer and the services which must be provided to enable workers in Energy Research programs to make the transition to the use of such a massively parallel computer for the majority of their large-scale calculations. The initial software environment should have batch processing and some time-sharing capabilities. Tools for debugging and libraries of application software must be present. Primary data storage on disk must be integrated with mass storage facilities for data migration and archival requirements. All of these features are present in some form today on massively parallel machines from various vendors. By the time of delivery of a system to the Access Program, these software components will have improved further so that, building on the collective experience of the HPCRCs, a satisfactory production software environment can be assembled quickly and continually improved.

Thus, the final recommendation of this report is that research and development for software libraries, programming tools, operating system features, etc. for a parallel computing production environment be undertaken with clear near-term and long-term goals. This effort would be a part of the Software Components and Tools area of the Advanced Software Technology and Algorithms (ASTA) component of the DOE Program Component of the Federal High Performance Computing and Communications Program.

B ExERSUG and User Input to the OSC

What follows is a letter and document regarding cost-effective, computational resources not available to the ER user community at the National Energy Research Supercomputer Center.

B.1 Letter to Office of Scientific Computing, December 1995

John S. Cavallini December 5, 1994
(Acting) Associate Director,
Office of Scientific Computing, ER-30
Department of Energy
Washington, D.C. 20585

Dear John,

ExERSUG for some time has realized we need to initiate planning documents and to target our concerns to the Office of Scientific Computing (OSC), the Supercomputer Access Committee (SAC), and Program Offices, collectively, and to the management of the National Energy Research Supercomputer Center (NERSC). Washington has made it clear that they want ExERSUG to be an independent body, not being directed by OSC or others in DOE, thus having our stated positions available to DOE as a truly independent view from the users. We believe we can strengthen this process by contributing our views and the views of the users we represent on a frequent and timely basis.

The enclosed document, "View from the Energy Research Users" is ExERSUG's presentation of the users' collective views and concerns to OSC on issues we consider of paramount importance. These views are strongly held by a large fraction of our users. We make specific, carefully weighted, constructive policy recommendations, aiming at a more efficient use of existing financial resources. We urge you to give our document your serious attention.

The document makes a strong statement of support for keeping the Massively Parallel Processor (MPP) acquisition by NERSC on track. The NERSC MPP team has done a stellar job in structuring their procurement to get the latest technology that will do the job. But with delays being evident, especially in these budget cutting times, ExERSUG perceives that we need to make a strong statement in support of MPP procurement.

The document additionally argues for NERSC adding additional capacity, by way of cost-effective Symmetric Multi-Processors (SMP). The SMPs would act both as development platforms for the MPP, which have been relatively poor in software, and as cost-effective capacity engines that would allow offloading the C90, allowing better access by our established base of supercomputer vector codes, and additionally allowing expansion of the new-capability Special Parallel Processing (SPP) use of the C90. Furthermore SMPs with larger memories would allow new capability, true supercomputing large-memory computations not feasible on the C90.

Sincerely yours,
Jack A. Byers, Chairman, ExERSUG

Co-signees, members of ExERSUG:

Bastiaan J. Braams, Viktor K. Decyck, William B. Herrmannsfeldt, Brian E. Hingerty, Stephen C. Jardin, Rick A. Kendall, Steven C. Pieper, Gerald L. Potter, Alan D. Turnbull

CC:

OER Program Directors:

Iran Thomas, (Acting) Associate Director, Office of Basic Energy Services, ER-10

Wilmot N. Hess, Associate Director, Office of High Energy and Nuclear Physics, ER-20

N. Anne Davies, Associate Director, Office of Fusion Energy, ER-50

Aristides A. Patrinos, (Acting) Associate Director, Office of Health and Environmental Research, ER-70

NERSC Managers:

C. William McCurdy, Director, NERSC

Michel G. McCoy, Deputy Director, NERSC

John M. Fitzgerald, Assistant Director, NERSC

B.2 View from the Energy Research Users

The National Energy Research Supercomputer Center (NERSC) is the primary supercomputing resource available to the Energy Research programs. The user base of NERSC is large and diverse with many differing needs. We believe the following describes the most significant concerns and needs of the NERSC user community.

B.2.1 Summary of users' concerns

NERSC represents the primary supercomputing resource available to the Energy Research scientists, providing them with an unparalleled environment of high-end computing equipment. One of the primary virtues of NERSC is that it provides a unified integrated computing environment that can accommodate the spectrum of supercomputing needs of the ER community, from the Grand-Challenge-scale single jobs to the tremendous volume of medium-to-large size jobs. We feel strongly that this dual role is essential for NERSC, and thus we urge in the strongest terms that the unified integrated computing environment of NERSC be maintained.

NERSC's vision of the Unified Production Environment is an appropriate model. NERSC should be there to provide those things (both capacity and capability) that we do not have available to us locally and which we require in order to do our jobs as an ER researcher. In addition to hardware, this also means software, technical leadership, and consulting. The users urge Washington to support this model. We think this is so important that we urge you to consider changing your present policies, charters, etc., if that is what is necessary for you to support it.

The users' overall concern is for NERSC to provide them with the best computing infrastructure for their scientific needs. This requires that the NERSC budget be maintained or, failing that, protected to the highest possible degree. It also requires that NERSC be given the freedom to spend its funding wisely in support of the users' spectrum of computing needs. Our detailed concerns can be summarized in the following three points.

First, the area that we perceive to be most ignored, additional cost-effective resources should be made available in order to allow use of the C90, and eventually of the MPP, to be dedicated to the most appropriate tasks for these machines. The workstation resources available at the users home sites are insufficient to meet this need.

The computer industry in recent years has turned away from the high-end maximum performance supercomputer. This change stems in part from the decline of support of this arena by Defense Programs, which all ER benefited from and now is feeling the lack of, and in part from a host of industry and economic changes. The biggest advances in recent years have been in lower-than-highest-end, cost-effective machines of increasingly powerful performance. NERSC's plan, to add newly available SMPs as part of its Unified Production Environment seems to us a very viable solution. Replacing the expensive-to-maintain Cray-2s with SMPs should be part of this plan. We urge Washington management to push this plan forward with all possible speed as the best possible way to ease the access to supercomputing at NERSC.

Second, it is essential that NERSC remain on the forefront of high-end, new-capability supercomputing. We believe that this should be achieved by keeping the acquisition of a general purpose Massively Parallel Processor (MPP) acquisition at as high a priority as possible. The NERSC MPP team has done a stellar job in structuring their procurement to get the latest technology that will do the job. We urge that the MPP acquisition by NERSC be kept on track. If the MPP procurement fails there are concerns from some of us that this could lead to the disappearance of the NERSC facility all together.

Third, it is essential that the C90 supercomputing resource be maintained. Pushing the C90 into more new-capability modes such as the Special Parallel Processing (SPP) program is to be encouraged, but this plan *requires* that the broader base of users at the same time be accommodated by other platforms at NERSC.

Special Parallel Processing (SPP) vs. normal vector use of C90 this last year became a big issue. On the one hand SPP has provided some truly new-capability scientific computing for a small but significant subset of our users. The use of the C90 in SPP mode is competitive with present day MPP machines. On the other hand, the SPP program forces other users off for significant fractions of the total time available. The success of last year's SPP was such that there were active plans by our Washington managers to ask for a significant expansion this year; these plans had to be largely shelved or reduced in scope when NERSC's analysis reported too large a negative effect on the rest of our users. So we are faced with some truly new capability solutions with the existing supercomputer platform, but nowhere to push the rest of traditional vector supercomputer codes for that fraction of time used by the new capability. The SPP use represents a suddenly appearing new jump in demand. A clearer case for the need for additional supercomputing resources is hard to imagine.

Given that SPP has become a successful mode of operation, it is our view that the SPP allocation process should be absorbed into the normal procedures. We urge that this change be expedited.

B.2.2 On the Need for SMPs at NERSC

It has often been suggested by managers and others that NERSC cannot be a place for "midlevel" computing, or "routine" computing, and that such computing should be moved elsewhere. We believe the distinctions of "midlevel" or "routine" vs. "true" supercomputing to be largely improper and artificial. Something that a manager might consider "midlevel" may be viewed by the user as a necessary part of his entire supercomputing endeavor, and one that requires substantial computing capacity. When there are no new resources, some hard judgments have to made about allocating existing resources. But the computer industry has been making big strides in the advancement of lower price/performance platforms, particularly with the now available SMPs. Acquisition of such platforms is evidently cost-effective. The stumbling block appears to be the policies in Washington, over two particular issues:

- the unavailability of "lease-to-own", and
- the OSC charter which we believe has been too rigidly interpreted to mean that NERSC should have the high-end machine and nothing else.

This view seems to have taken hold, to the exclusion of any serious discussion of other options, even when cost-effective solutions are increasingly available.

Big SMPs with big memories not only provide a cost-effective route to badly needed increased capacity, but can also provide high capability compared to the kind of workstation resources that are typically available at local sites. Another extremely favorable point for SMPs: for many of our codes there is no large new programming effort needed. As development platforms, the SMPs would directly contribute to the best use of the MPP, for which it is more difficult to develop software. We expect development on an SMP followed by production on the MPP to be the most successful transition path to effective use of the MPP. We urge that high priority be given to the acquisition of some of these cost-effective and versatile platforms.

B.2.3 Conclusion

We close with three points and a final statement.

One, the problems the community is experiencing are a consequence of our success. Overcrowding occurs because NERSC is successful at providing a unified computing environment to a broad base of users. And the decrease in general access caused by the sudden jump in SPP only occurred because of the success of this new way to use the C90: computations heretofore impossible are now routine in the SPP program, extending the scope of research to new regimes.

Two, the most important product of NERSC is the new scientific research that comes from the use of the computing resources. If this product is not continually pushed towards excellence, we shouldn't be in business at all.

Three, we believe that Washington has an opportunity to renew its support for the model of broad based supercomputer access. This includes extending the high end at the pace available from the computer industry (via MPP) and supporting its very broad base of users at the same time (via increasingly cost-effective SMP platforms now available). For us to continue to adhere to policies that actively advance the cause of only the highest-end users will only ensure the slow decay of NERSC; computational scientific research throughout ER will be the loser.

A diverse group of OER researchers uses NERSC to produce high-quality scientific results. It is now time to take advantage of the enormous advances made by the American computer industry in increasing the compute cycle/dollar ratio. The near-future acquisition of an MPP is seen as the best way to enhance the NERSC computing resources to push science beyond its current scope. In the meantime the addition of SMP systems will improve the facility by: (1) providing a development platform for the labor intensive process of converting research applications to use MPPs, (2) allowing the expansion of the SPP program to allow an increased subset of users the ability to do new science, and (3) providing a capacity engine for appropriate applications and thus reducing the load on the current vector supercomputers.

C Findings of the Second Pasadena Workshop on System Software and Tools for High Performance Computing Environments

Thomas Sterling[34]
Paul Messina[35]
and Jim Pool[2]
November 30, 1995

Introduction

The Second Pasadena Workshop on System Software and Tools for High Performance Computing Environments was conducted in January, 1995 and involved more than a hundred invited participants from industry, academia, and government. The objective of the workshop was to identify key factors critical to expanding the effectiveness and use of HPC systems through software technology and to propose actions by the community that would advance the quality of HPC programming environments. The assembly considered these issues from a number of perspectives including applications, parallel system software development, HPC architecture, relationship between research and commercial software, and several programming paradigms. Among the principal findings was that the HPC community is unable to sustain necessary R&D independently and must leverage industry investment in the rapidly emerging workstation clusters and symmetric multiprocessors. Of equal importance is the need for a common user interface across tools of equivalent type and shared interfaces between tools for interoperability. This paper presents the issues addressed and approach taken by the workshop, as well as the findings and recommendations generated by the workshop working groups.

The purpose of this second workshop on HPC software technology was to focus in on the very real problem of getting useful tools in the hands of applications developers to make scalable parallel computing systems an effective means of solving large-scale real-world computational problems. The first workshop established a baseline of understanding in this domain, providing broad coverage of the many components comprising HPC software technology and identifying specific critical and immediate needs. This second workshop extended these findings to include the enabling infrastructure for migrating experimental results and tools from the research community to production grade environments in support of the HPC applications developers. In so doing, this workshop involved hardware systems vendors, independent software vendors, researchers in system software, end-users, and governmental representatives responsible for technology policy and planning. During the course of the workshop, this expert body identified issues in the development, deployment, and use of

[34]NASA Goddard Space Flight Center
[35]California Institute of Technology

HPC system software and tools and provided recommendations for resolving them.

C.1 Focus Questions

The basis for discussions at the workshop was five central issues that were to be addressed if the end goal of realizing effective production quality HPC programming and execution environments is to be achieved. These are:

1. Progress since First Pasadena Workshop

   What progress has been made in system software and tools since the First Pasadena Workshop? What are the major successes and failures? To what extent have the recommendations of the First Pasadena Workshop been adopted—why and why not?

2. Characteristics of HPC Scientific and Engineering Applications

   What are the characteristics of scientific and engineering applications and algorithms that require high-performance computer systems? What are the highest priority issues for system software and tools implied by these applications and algorithms? How will these applications and algorithms contribute to the long-term commercial viability of high-performance computing systems? What particular system software and tools will attract independent software vendors to participate in the commercialization of these applications and algorithms?

3. Use of System Software and Tools

   What types of systems software and tools are needed and can be provided to facilitate the creation, monitoring, and analysis of portable, scalable application software in high-performance computing environments? What would be required to motivate the user to use such system software and tools?

4. Interaction between System Software and Architecture

   What are the constraints imposed upon system software designers by current and anticipated architectures? Are key developments in system software hampered because architectural support is lacking? What aspects of high-performance systems software are likely to place significant new requirements on future architecture?

5. Transition from Research to Products

   What barriers prevent system software and tools developed by the research community from becoming products supported by system vendors and/or independent software vendors? What are the barriers to the use of these products by application developers? What should be the strategy for removing these barriers and explicitly what is the government's role in that strategy?

C.2 Major Findings

Many issues related to the prevailing conditions and their causes were revealed by the nine working groups comprising this workshop. Their details may be found in the individual working group reports. Presented here are those findings which dominate the current state of HPC system software and the future of HPC in general. In many cases, individual findings reflect the conclusions of a number of working groups and in certain cases findings are derived from a truly global consensus. While the focus issues described in the previous section provided a framework for discussions, the major findings below are organized for convenience in terms of cause and effect. The first set of findings describes the achievements over the first half of the HPCC program. The second set of findings identifies the critical challenges that are facing the HPC community. And the third set of findings uncovers some of the driving factors that contribute to the causes of the problems.

C.2.1 Progress and Accomplishment

Following the first related workshop some progress has been made in HPC system software and tools. Two important standards for programming languages for high performance computing systems have emerged through community led forums. MPI provides a unified message passing library that integrates ensembles of independent processing elements from clusters of heterogeneous systems to tightly coupled arrays of like processors. HPF provides a high level data parallel programming model suitable for many applications in the sciences and engineering while hiding much of the complexity of managing the underlying resources. Both programming systems are being evaluated by the HPC to determine their potential long term impact for portability and programmability.

Two successful multi-agency collaborations have been conducted in response to the findings of the first workshop. The Joint NSF-NASA Initiative in Evaluation or *JNNIE* engaged researchers from many parts of the country and across computational science disciplines to analyze the performance and behavior of HPC systems under diverse real-world workloads. Pablo is an important development in performance profiling of parallel computing which was co-sponsored by ARPA and the NSF.

Although not an accomplishment in the area of HPC system software, it is important to note the emergence of SMP (*symmetric multiprocessor*) systems as a major new platform type. These systems evolved from the workstation community and integrate a small number of processors in a tightly coupled shared memory configuration. The significance of this trend is that many ISVs are targeting for the first time a class of parallel systems on which to host their products. This introduces a rapidly growing parallel program development industry unprecedented in the history of computing, presenting an opportunity for parallel HPC systems to leverage this investment.

Another trend since the first workshop that will impact HPC system software is the increasing use of workstation clusters which are finding productive application by industry for real workloads and are being treated as a low cost entry point into very coarse grain parallel computation. Such systems have already been a major factor in the success of de facto standards in message passing such as MPI and PVM.

Heterogeneous computing is emerging as a way to utilize the NII and HPC. In practical terms, this emerging field is still in its infancy and considerable basic infrastructure remains to be developed.

C.2.2 Critical Problems and Challenges

In spite of the achievements of the intervening three years since the first workshop, critical problems and challenges face today's HPC community which will require innovative, coordinated, and concerted efforts to address. Some key issues confronting HPC are:

1. HPC application developers currently use almost none of the experimental research software that was designed to make their tasks easier. This software needs to be made far more reliable and portable before it can have the desired impact.

2. Market penetration of HPC systems remains low with limited impact on national productivity and competitiveness.

3. Software tools and applications exhibit poor portability. This is true not only between HPC systems but between SMP systems and across SMP to HPC systems as well.

4. Lack of interoperability among evolving system software and tools precludes mutual leveraging of capabilities and severely limits benefits to applications developers as well as hinders cross platform portability. For example parallel compilers, debuggers, and performance analyzers on the same HPC platform but from independent sources are rarely capable of cooperation. At least in one sense, an notable exception is PVM which permits cross platform cooperation at the application level.

5. Products of academic research in system software and tools are rarely commercialized, nor find wide use. Commercialization for many such tools may never be viable due to the small market but are nonetheless very useful and important to productivity and the advancement of HPCC.

6. Management of scalable mass storage by HPC systems is poorly understood but a critical requirement and is the focus of some recent major collaborative research initiatives.

C.2.3 Driving Factors

Given the challenges identified above, a number of critical factors have been determined that are contributing to them. These must be addressed by the HPC community in general and those involved in advancing the capability and utility of parallel systems software in particular. Among those issues of most concern are:

1. The HPC market has been too small for commercial viability to justify ISV R&D investment in sophisticated system software and tools or in large parallel applications packages. If anything, this condition has been aggravated by the functional disparities among HPC vendor offerings.

2. The multiplicity of architectural models for HPC has been a significant source of difficulty in parallel programming and resource management. System software has been unable to fully compensate for these shortcomings and provide the user with the necessary tools for efficient programming or the means of conveniently achieving efficient execution. One possible exception is SIMD parallel architectures, at least for a limited range of applications.

3. Tools generally provide inadequate functionality.

4. Experimental tools have low reliability.

5. Tools to aid parallel programming often work for small numbers of processors but can not effectively manage the context of a very large number.

6. Tools of the same type are frequently inconsistent across platforms which hinders acceptance by ISVs and other end-user application developers.

7. System software and tools for the same platform are often not interoperable.

8. The path from research experiment to commercial product varies dramatically depending on nature of the tool, community demand, and ease of integration into diverse contexts. No single model of prototyping and commercialization can support and encourage this process (e.g.,. Linux, Mosaic, HPF, Pablo).

9. HPC system market growth will be driven by commercial applications that can be enabled but not forced by availability of software tools. A successful example is the oil and gas industry who has applied HPC technology to exploratory geophysics through seismic analysis among other means.

10. HPC hardware and software systems must be positioned as high end of a continuum of parallel system capability from workstations, through SMPs, to scalable HPCs with common programming, environment, and application framework.

11. Rapid HPCC system software development, evolution, and deployment is critical to the health and aspirations of the HPCC community and the objectives it is addressing.

12. A key gap in the funding of technology transfer is that agencies sponsor research and prototypes but it is unclear what source of support is available to engage the next step towards broad dissemination of robust systems or even commercialization.

C.3 Recommendations

The workshop developed a broad set of recommendations in response to the above findings. In each case, the recommendations are proposed to enhance the utility of high performance computing systems through the advance of robust environments. The recommendations were strongly influenced by current trends and market forces in commercial computing. These recommendations are organized into two sets. The first establishes what needs to be achieved to make HPC systems technology competitive and capable of pushing the boundaries of performance. The second set of recommendations are a set of broad activities that together will bring about accomplishment of the first set. From the outset, it is clear that progress towards these goals will require the active cooperation of the hardware and software industry, academic research groups, and strong leadership from appropriate government entities.

C.3.1 What Needs To Be Achieved

The disparity among programming and execution models across HPC system architectures has inhibited the development and migration of third party applications and required substantial effort for ports as well as learning curves as users move between HPC systems. HPC systems must include a basic set of capabilities and tools with similar, if not identical, user interfaces. Systems of diverse components can be used as an ensemble for computing but critical infrastructure and resource management systems must be developed to make this practical.

To a large extent, new tools and software subsystems for HPC environments are developed as stand-alone products. Interoperability, especially with the diverse and rapidly changing HPC system software components, is rarely a near term requirement. Yet, the inability of a user's tool-sets to work together or the lack of means to integrate a new tool into an already established software context constrains the effectiveness, usability, and progress of advanced HPC software tools. Tools must, in general, be able to share common data, exchange products, and work together.

It is end-user applications of economic importance that will drive the demand and market for HPC systems, and provide a sustaining economic base for their future development. But the uncertainties involved in developing major codes for HPC systems impose too much risk to warrant investment by third party-software vendors. A base of economically important applications must emerge for HPC systems and, initially, these must be identified and sponsored through government support in cooperation with industry that would provide the initial sustaining customer base. The degree and rate of success will depend on the extent to which the base of applications can be broadened beyond the conventional science and engineering applications.

Because a major aspect of HPC system software is to simplify the programming interface to and manage the myriad resources of complex parallel systems, the difficulty in realizing effective system software is in large part due to the structure and low level logical interface which make up the architecture of such systems. Although processor architecture is driven by the workstation market, within that constraint processor/system architecture needs to evolve to be more supportive of scalable parallel processing applications and tools.

Symmetric Multi-Processors (SMP) are emerging as a major class of computing system, derived from and extending the range of capabilities of high end workstations. These are likely to be the first high volume parallel computers and the target for commercial third party parallel software. While not truly scalable beyond a small number of processors, these systems will be the dominant form of parallel architecture with sufficient market presence to justify industrial investment. The degree and likelihood of the future success of HPC systems will be significantly enhanced through their ability to leverage software investment for SMPs. This requires an almost seamless relationship between SMPs and HPC where end-user applications may migrate up and down the spectrum of parallel computer organizations from single user workstations to HPCs comprising thousands of processors.

C.4 Implementation Initiatives

This section provides action items that will go a long way toward implementing the recommendations. They involve the formation and sponsorship of a collection of committees with focus charters. Note the use of the term "circle of excellence." This is a notion that achieves the same thing as a center of excellence but has the talent distributed around the country. This allows people in other organizations to be funded to collaborate on specific work. It does not require the lead time to set up a building nor is it as likely to be perceived as in competition with all other research groups. Rather it leverages talent and capabilities already established and partly redirects it in a coordinated manner to specific objectives. Through such a structure, new skills can be easily added permitting new research directions to be exploited quickly and within a given budget. Its last advantage is that it is easier to scale up and down as funding permits/constrains. Implicit in this model is a key change in the process of engaging the research community to address government specified objectives. Technical decisions would be moved out to the field, providing rapid exploitation of new opportunities, while government oversight is maintained.

- Convene an Oversight Committee on HPC System Software and Tools to develop a national software technology plan and to monitor progress, engage the HPC community, and report status on an ongoing basis. The Oversight Committee should include technical representatives from academia, government, and industry.

- Convene a standing committee on HPC software tool set functionality to determine the critical capabilities that need to be available on all parallel platforms from SMPs to HPCs. Conduct a focused workshop on this specific issue to initiate the process.

- Fund development of infrastructure and common interfaces for interoperability of HPC software tools. These efforts should reflect and respond to the findings from the tool set functionality committee. An early workshop should be conducted with strong vendor and SMP community participation to expose practical issues associated with achieving interoperability across systems and software types.

- Develop and deploy a basic tools set compliant with the functionality and interface requirements determined from the above studies.

- Identify, sponsor, and select key applications of commercial economic and national importance and sponsor ISV porting to HPC platforms. As an example, emerging area of synergy is support for massive data assimilation on very large archives through use of HPC resources. Providing access to HPC resources as services on the NII can enable this development.

- Engage the HPC community to determine minimum requirements for hardware support of system software and tools in order to provide a guideline to processor/system manufacturers as to what capabilities need to be included and to procurement agencies to aid in selection of appropriate systems.

- Create a software tools evaluation and test center.

- Support the establishment of a software tool development infrastructure.

- Develop and apply a *total cost to solution* metrics approach for evaluating progress in the HPC arena.

- Establish a common framework for ease of integrating and sharing computing products to enable HPC system interoperability. To this end develop transparent mechanisms for storing and handling all forms of data to be exchanged among systems and programs.

- Extend application/system evaluation and characterization studies such as JNNIE to provide detailed analysis of hardware impact on application efficiency and scalability as well as programmability and usability.

- Support rapid deployment consortia to produce prototypes/modules and related test suites that serve as an early implementation vehicle.

- Extend the charter of the NHSE (National HPCC Software Exchange) to include support for forming and administering a group representing a broad portion of the community. Empower and fund the NHSE to evaluate and robustify software.

- Create a National Distributed Heterogeneous Testbed as a vehicle for developing and evaluating system software for harnessing the combined capabilities of disparate systems types. Such a testbed should be distributed across the country and may include components at existing HPC centers.

C.5 Conclusions

The Second Pasadena Workshop on System Software and Tools for High Performance Computing Environments was an important step in crystallizing the issues governing the direction of HPC software technology. Its sharp focus on the realities dictating practical constraints was a valuable achievement in itself and produced a high confidence in the workshop findings. While the recommendations may be difficult to accomplish, even if the implementation plan is followed, there is reasonable expectation that its path is valid.

The timing of this workshop was important both in its relation to the HPCC program and with respect to industry directions. The reemergence of shared memory parallel systems such as the CRI T3D and cache coherent systems such as the SGI Power Challenge and Convex Exemplar has shifted the path anticipated for HPC at the beginning of the program. The impact of shared memory in its many forms on the evolution of system software is significant, although there are divergent opinions on the specific ramifications. The juxtaposition of successes in achieving high sustained performance for certain applications on large distributed memory systems on the one hand, and apparent difficulties in programming them on the other provides new understanding about the opportunities and challenges confronting the HPC community, not available at the beginning of the HPCC program. Finally, since the first workshop, the community has benefited from additional workshops in applications and petaflops scaled computing that has contributed to increased understanding of system scalability issues. With these advances since the first workshop, it was clearly time to engage in a major reassessment of the state and direction of HPC system software.

One of the realities dealt with by this ensemble of colleagues was the difficulty in harnessing resources. Simply put, this is because colleagues in all venues of the HPC community are over committed. With shrinking budgets and rapid shifts in priorities resulting in increased planning responsibilities, it is certain that the stress factors will only increase, at least in the near future. The community must find a better way of tasking itself towards common goals. Time constants are too short to permit long lead times typical of top-down driven projects. The "Circle of Excellence" model is one approach to gaining maximum benefit from existing capabilities with only small changes. Another component of the solution is to augment the capabilities of government leadership in guiding initiatives by moving the tactical direction of such initiatives closer to the technical contributors. Overall objectives, approach, and project evaluation must remain within the corridors of government, but, if left in the critical path, contention for those conduits of high level decision can become barriers to rather than channels for accomplishment. A means of delegating as much responsibility to the community as can be safely done should be found to expedite and facilitate successful accomplishment of designated initiatives.

If there was a single type of achievement of this workshop, it should be described as clarification. Few final answers were developed. Rather the major result of the workshop was an essential advance in understanding of the dynamics determining the progression in the HPC field and the sharing of that understanding to establish a shared consensus across the community. In a way, it set a new starting point from which to consider the next steps; steps that will lead not only in a somewhat different direction but on more than one course around the obstacles obscuring the final objective. Having pointed to the next mile posts, it is up to our community to make those first steps.

D DOE/OCTR/MICS White Paper

White Paper
on
The Energy Research High Performance Computing Access Program
till the End of the Century

By
T. A. Kitchens

D.1 Conclusion and Strategy from the discussion below:

The target heterogeneous distributed computing system for Energy Research would satisfy the capability and the capacity needs for Energy Research investigators. However, the target computing system may well not be attainable in the present outlook on the DOE budgets for the remainder of the century. A plan with less risk is to make the access center of highest priority because of the flexibility of the high performance computing and storage system and the strongest ability to accomplish the major computational tasks of Energy Research. Provided there is sufficient data communications technology, the centralized facility can tradeoff capability computing for capacity computing while the distributed system of workstations and symmetric multiprocessing units without the access center cannot provide capability computing. Moving away from the access center in the upcoming tight budget times will also place an additional load on the ER science and technology offices to provide local computing services when they can least afford them. This approach also depends less on technologies to be developed and progresses naturally toward the perceived nature of computing in the 21st century - heterogeneous distributed computing. Funds would be provided later for powerful workstations as their capability improves and their price comes down.

D.2 The Nature of ER Computing/Communications Today

The National Energy Research Supercomputer Center (NERSC) provides Energy Research investigators with about 22 gigaFLOPS of classical vector-processor supercomputer capacity with a 4-processor CRAY-2, an 8-processor CRAY-2, and a 16-processor 256 megaword C-90. The NERSC Central File System has about 18 terabytes of mass storage. The ESnet backbone has a peak bandwidth of 45 megabits/sec. Most local sites have workstations capable of 10-50 megaFLOPS on local area networks with bandwidths of 10-20 megabits/sec without resource sharing or scheduling. A few sites are more advanced and have resource sharing, such as FSU/SCRI, and a few with substantial compute servers as well, such as ANL and SNL, and even a discipline focus with substantial site/program support, such as PNNL and LANL. Today the only High Performance Computing Research Center with a cutting-edge machine, the 2000-node Intel Paragon, is at ORNL. Virtual reality research facilities exist at ANL and information infrastructural elements are under development at several sites.

D.3 The Nature of Computing/Communications in the Year 2000

D.3.1 Expecting some slowdown in the progress of technology:

Workstations capable of a 0.6 gigaFLOPS or more with 0.4 gigawords of RAM and about 1 terabytes of mass storage will be available for about $200,000. The major limitation will be the cost of memory. Optical memory, even for mass storage, will still be in the future.

Local area networks with bandwidths of gigabits/sec should be commercially available, albeit expensive. Research testbeds will be running at 10-50 gigabits/sec.

Compute servers, based on symmetric multiprocessing, capable of 20-40 gigaFLOPS with 20 gigawords of RAM will cost $2-3 million. Associated storage servers will have 10-25 terabytes of mass storage for $1-2 million. The major limitation in the compute servers, other than memory, will be interprocessor communications, and this will keep the class of machine from scaling to more than 16-32 processors. Mass storage devices will still primarily be magnetic.

Wide area networks with bandwidths of OC/12 will be numerous but driven by edutainment requirements. The technology and the politics/funding will be under different paradigms than today. For example, communications could be through stripped ports provided by a wireless communications vendor into the communications cloud, or use of Redundant Arrays of Inexpensive Networks (RAIN). The limitations on the capabilities are difficult to predict, but it is likely that there will still be substantial problems with security, authentication and verification.

Massively parallel processors, based on vector processing units, capable of a teraFLOPS or more with about 0.5-1.0 terawords of main memory and 100-200 terabytes of mass storage will be available for a cost of about $60 million. The major limitation will be the cost of memory and interprocessor communications. The latter has caused the favored architecture to go toward more efficient design involving vector processing and emphasis on very high bandwidth access to a (virtual) shared memory with 8000 or less processors and high I/O bandwidth to mass storage.

D.4 Target for Energy Research Computing in the Year 2000:

The DOE should set its sights on utilizing the elements of technology discussed above to form the efficient coordinated heterogeneous distributed system described below. The budget and other limitations such as our limited ability to predict the technology and arrangements for high speed digital communications will preclude our achieving this goal. The strategy, see above, outlines a path toward this goal that will reduce risk and provide suitable services in the transition.

D.4.1 Local sites:

Most small sites, DOE University contracts, might consist of 1-3 workstations as described above connected to the Internet of the year 2000. In a good moderately-sized local installation, the workstations described above will be networked together and with more numerous smaller workstations by a gigabit/s local area network. These sites will have one or more 20-40 gigaFLOPS local compute servers and a storage/information server with a dozen or so terabytes on the local gigabit network. By 2000, some local systems will have a dynamic scheduler which will allow sharing the local compute servers and the background capacity of the workstations on the local area network. These schedulers might be able to schedule jobs that exceed the capacity of the local system on other systems in the ER organization. As developed below, ER should have a single large massively parallel system at an "access site" with the capability of a teraFLOPS or more with a memory of a teraword and a mass storage system of 100 terabytes. Software development and debugging would be done locally in the object-oriented environment of the workstations where the volume of the workstation market insures robust systems software and tools - not at the access site. Most local compute servers will be top-end symmetric multiprocessor systems available from the workstation manufacturers and will provide the same robust software environment. Thus the local compute servers would provide the resources for prototyping the large jobs that require the access center capabilities as well as providing much of the production (capacity) computing.

More specifically, computational scientists could expect a peak of 3 or more gigaFLOPS at night and more than 3 terabytes of mass storage from a typical local installation of ˜4 large workstations and ˜10 smaller workstations. If a compute server is local, the gaggle of workstations would most likely be larger and the site would have peak compute power of about 30-50 gigaFLOPS and have a mass storage system with 30-50 gigabytes. Most code and data would be maintained locally and only transported to the access site when needed there. (The large mass storage at the access site would be used to provide flexibility in the overall system, temporary storage for the input and output of calculations on the massively parallel system and as a repository of data used more than once by the high performance system.) More than 90% of all cycles used would be satisfied by local capacities. For example, most discrete calculations, simulations of a million or less particles, and most commercial applications requiring less than a few gigaFLOPS would be prevalent at the local sites. Local sites would have good graphics/visualization capabilities, possibly even virtual reality CAVEs, and other information infrastructural resources that need to be in the proximity of the users.

D.4.2 Access site:

From the 1993 Report "*Requirements for Supercomputing in Energy Research: The Transition to Massive Parallel Computing*" there are at least six key areas of Energy Research that have computational problems that require 100-500 gigaFLOPS computing systems along with substantial needs for main memory and mass storage. These problems need "capability computing" and the needs are not met with the abilities of the clusters of workstations, or the compute servers, or the aggregate local site described above. At the turn of the century, local and wide area bandwidths and software technology will be lacking. The computing system at the access center would need to be at the cutting edge to fulfill the needs of ER. Such systems probably would not have robust software. However, the development of code for the teraFLOPS system would be done on the workstations/symmetric multiprocessors using the highly tuned workstation environment and cross-compilers. The access site exists to provide "capability computing" - optimized for the computationally intensive major projects of the caliber of the Grand and National Challenges like the DOE's Numerical Tokamak Project, Global Climate Change simulations, combustion and other reactive flow modeling, materials by design, environmental chemistry applications, Genome data base searches, and simulated flow of oil and ground water. Projects will qualify to use the access center if they require more than about 30 gigaFLOPS or have restrictions on wall clock production time or need a main memory of 100 gigawords to compute efficiently.

The access site may exist because of its leading edge mass storage capabilities as much as its housing the state of the art computing system; this issue will be especially important until there is a substantial change in the high performance data storage technology and the bandwidth to the storage system. Additionally, the access site could be a site for research and development of prototypes for mass storage technology, graphics/visualization, virtual reality apparatus, and other techniques for the information infrastructure – all in collaboration with the information industries.

D.4.3 The plan:

The early steps toward the ER computing scenario for the turn of the century are in the NERSC Implementation Plan. The CRAY-2s will be retired and the expense of maintenance and power will allow NERSC to prototype a likely "local compute server" system as envisaged in for the year 2000. The NERSC staff will be looking into interoperative systems with the prevalent workstations and the issues of cross-platform (standards driven) object-oriented applications development software. The OSC has also sponsored developmental work on distributed computing tools and applications as well as security and interoperability issues through the Distributed Computing Coordination Committee. With the completion of the Environmental Molecular Sciences Laboratory at PNNL, a well-funded discipline-focused local site will exist and provide an opportunity for OSC to prototype the kind of distributed computing appropriate for the year 2000. Meanwhile, the NERSC acquisition of a massively parallel machine in FY1996 will provide a prototype of the Year 2000 access site. Thus the components of the computing infrastructure for 2000 are presently under development.

The target ER computing system described above would also have its management and support distributed. The local site capabilities would be the responsibility of the site and programs: this would insure that each local system is responsive to the missions of the site and associated programs and that the local capabilities are properly scaled. The access site would be a center of high performance computing research and would be supported by the OSC. OSC would be responsible for providing means for coordination of the capabilities for the heterogeneous distributed computing (interoperating computing and communications software, and some services) to insure that the system can work. The local sites need to be coordinated with the access site through the Offices of Energy Research. This plan would return the emphasis on top-end computing in the services/capabilities provided by OSC funding.

D.4.4 Beyond 2000:

In five years a paradigm shift is likely in this very active field of research and commercial interests but it is not possible to predict. This paper is built on the assumption that technology will maintain the present course. With no shifts in the technology development the workstation and MPP technologies will be convergent and the needs of the access center for high-performance computing purposes would likely vanish provided network bandwidths increase substantially as expected. However, mass storage resources may continue to benefit from economy of scale and justify the existence of a centralized tertiary storage system.

E Opportunities and Challenges of High Performance Computing in Chemistry

Editors: M. F. Guest, R. A. Kendall, J. A. Nichols, T. H. Dunning, Jr., and M. S. Gordon

Authors: J. Almlöf, R. A. Bair, B. R. Brooks, C. L. Brooks, III, M. Colvin, D. A. Dixon, T. H. Dunning, Jr., M. F. Guest, R. J. Harrison, M. Kalos, A. Kuppermann, B. McCurdy, W. P. Reinhardt, A. P. Rendell, J. Saltz, K. Schulten, R. Stevens, P. R. Taylor, and A. F. Wagner

E.1 Executive Summary

The field of high performance computing is developing at an extremely rapid pace. Massively parallel computers offering orders of magnitude increase in performance are under development by all the major computer vendors. Many sites now have production facilities that include massively parallel hardware. Molecular modeling methodologies (both quantum and classical) are also advancing at a brisk pace. The transition of molecular modeling software to a massively parallel computing environment offers many exciting opportunities, such as the accurate treatment of larger more complex molecular systems in routine fashion, and a viable, cost effective route to study physical, biological, and chemical "grand challenge" problems that are impractical on traditional vector supercomputers. This will have a broad effect on all areas of basic chemical science at academic research institutions and chemical, petroleum and pharmaceutical industries in the U.S., as well as chemical waste and environmental remediation processes. But, this transition also poses significant challenges: architectural issues (SIMD, MIMD, local memory, global memory, etc.) remain poorly understood and software development tools (compilers, debuggers, performance monitors, etc.) are not well developed. In addition, researchers that understand and wish to pursue the benefits offered by massively parallel computing are often hindered by the lack of expertise, hardware, and/or information at their site.

A Conference and Workshop organized to focus on these issues was held at the National Institute of Health, Bethesda (February, 1993). The specific aim of the Conference was to provide a snapshot of current state-of-the-art and background material on the following:

1. current work in chemistry on parallel computers, with a special emphasis on massively parallel computers,
2. the chemical problems which are likely to benefit substantially from the use of massively parallel computers, and
3. the impediments, programming and otherwise, preventing the migration of existing codes to, and the development of new codes for, massively parallel computers.

This report is the culmination of the organized workshop and the efforts of the participants; it discusses issues (1) and (3) at some length, and provides an update of activities that have taken place since the workshop . The main conclusion of the report is that a drastic acceleration in the present rate of progress is required for the chemistry community to be positioned to exploit fully the emerging class of Teraflop computers, even allowing for the significant work to date by the community in developing software for parallel architectures.

E.2 Recommendations

The steps required to address this issue of increased rate of progress are encapsulated in eight specific recommendations, which are summarized below.

RECOMMENDATION-1: Recognizing that many of the scientifically driven needs of computational chemistry can only be met by future generations of MPP computers, software and algorithm development should take both problem and resource scalability as a basic tenant.

RECOMMENDATION-2: Scalable parallel computational chemistry applications demand both advanced algorithm development and vendor-based hardware and software capabilities, targeted to the requirements of the computational chemistry community. A well-defined list of requirements must be tabulated, including those of portability, performance, enhanced development environments, scalable MPP subsystems (communication, disk I/O, memory capacity), compact expression of chemical algorithms, and available parallelism, including non-uniform memory access and multiple levels of granularity.

RECOMMENDATION-3: The development of new methods and algorithms should be encouraged and supported. A key aspect of this support should involve the integration of the computer science, computational science, and computational chemistry fields at the undergraduate and postgraduate levels.

RECOMMENDATION-4: Close interaction of computational chemistry application developers across chemistry disciplines and with computational scientists, computer scientists, and vendors should be strongly encouraged to foster co-education of these varied communities.

RECOMMENDATION-5: Development and production use of MPP supercomputers are necessary to make full use of today's MPP systems, while preparing for the next generation of computing resources. MPP supercomputer centers must be configured to meet the needs of computational chemistry applications and users.

RECOMMENDATION-6: MPP supercomputer centers should act as a repository for algorithms and expertise, and provide a focal point for communication between the computational communities, vendors, and application users.

RECOMMENDATION-7: Technology transfer to industrial researchers must be streamlined to increase the productive use of MPP-based computational chemistry applications.

RECOMMENDATION-8: A suite of computational chemistry benchmarks (c.f., the NAS kernels) needs to be developed to demonstrate requirements for comparison and development of software environments and hardware, as well as to provide performance measurements of MPP and conventional supercomputers.