




Argonne's Laboratory Computing Resource Center

2006 Annual Report



About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne, see www.anl.gov.

Availability of This Report

This report is available, at no cost, at <http://www.osti.gov/bridge>. It is also available on paper to the U.S. Department of Energy and its contractors, for a processing fee, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone (865) 576-8401
fax (865) 576-5728
reports@adonis.osti.gov

Disclaimer

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

Argonne Laboratory Computing Resource Center

FY2006 Report

In the spring of 2002, Argonne National Laboratory founded the Laboratory Computing Resource Center, and in April 2003 LCRC began full operations with Argonne's first teraflops computing cluster. The LCRC's driving mission is to enable and promote computational science and engineering across the Laboratory, primarily by operating computing facilities and supporting application use and development. This report describes the scientific activities, computing facilities, and usage of LCRC operation in 2006 and the broad impact on programs across the Laboratory.

The LCRC computing facility, Jazz, is available to the entire Laboratory community. In addition, the LCRC staff provides training in high-performance computing and guidance on application usage, code porting, and algorithm development. All Argonne personnel and collaborators are encouraged to take advantage of this computing resource and to provide input into the vision and plans for computing and computational analysis at Argonne.

Steering for LCRC comes from the Computational Science Advisory Committee, composed of computing experts from many Laboratory divisions. The CSAC Allocations Committee makes decisions on individual project allocations for Jazz.

For further information about the LCRC and Jazz, please see the LCRC Web site at <http://www.lcrc.anl.gov/>, or send e-mail to consult@lcrc.anl.gov.

Ray Bair

Director, Laboratory Computing Resource Center
bair@lcrc.anl.gov

Contents

Executive Summary	1
The Laboratory Computing Resource Center	3
The Jazz Computing Cluster	4
Cluster Configuration	4
Cluster Usage	4
Users and Projects	6
Usage by Discipline	6
Job Size Characteristics	7
Scientific and Engineering Applications	8
LCRC Research Highlights	8
Integrated Modeling System for Evaluating the Impact of Aerosols on Regional- and Urban-Scale Climates (AERO-MODEL)	10
Process-Scale Modeling of Atmospheric Chemistry and Transport (ATMOS CHEM)	11
Development and Benchmarking of Eulerian 2-Phase Boiling Models for Computational Fluid Dynamics Simulation	13
Studies of the Water-Gas Shift Reaction over Mono- and Bimetallic Catalyst Surfaces (BiMETALLIC-WGS)	15
Wave Packet Dynamics with Cartesian Coordinates (CARTESIAN)	16
Development and Performance of the Community Climate System Model on Linux Clusters (CCSM2)	17
Spatiotemporal Chaos in Rayleigh-Benard Convection	19
Geometry Optimization for Molecular Electronic Structure Using Multireference CI Methods	21
Computational Nanoscience (COMPnano)	23
Electronic Excitation Calculations on Organic Semiconductor Copper Phthalocyanine (CuPc)	25
Kinetics of Enzymatic DNA Repair	27
Meson and Baryon Observables in a Dyson-Schwinger Approach (DSE_EXC_PION)	29
Earth System Modeling Framework (ESMF)	32
Fast Code Development for Large Computed Tomography Data Sets (FASTCTDEV)	34
Global Climate Modeling with the Fast Ocean Atmosphere Model (FOAM)	35
High-Throughput Genome Analysis	38

Evaluation of Commercial CFD Code Capabilities for Prediction of Aerodynamic Drag for Heavy Vehicles	40
Advanced 3-D Laser-Produced Plasma Hydrodynamics and Radiation Transport for EUV Lithography Applications	42
Physics of the High-Temperature Gas-Cooled Reactor.....	44
<i>Ab Initio</i> Investigation of CO Hydrogenation Mechanisms with Co Carbonyl Catalysts	45
Dynamics of Voltage-Gated K ⁺ Channel	48
Lattice Quantum-Chromodynamics (LATTICE-QCD).....	49
High-Fidelity Model Simulations in Urban Boundary Layer Applications: Air Pollution Meteorology and Hazard Cloud Dynamics (MANHATTAN).....	51
Monte Carlo Analysis of Accelerator-Driven Subcritical Systems	53
Molecular Dynamics Simulations for Solution X-Ray Diffraction (MD_FOR_SXD).....	56
Multimethod Linear Solvers in Terascale PDE-based Simulations (MULTIMETHOD_SOLVERS)	58
Multi-Scale Analysis of Protein Folding	61
Computational Nanocatalysis (COMPnano)	62
Computational Nanophotonics (NANOPHOTONICS)	64
Simulation of Water in Carbon Nanotubes.....	66
Electronically Nonadiabatic Direct Dynamics for Barrierless Reactions.....	68
Neocortical Seizure Simulation (NEOCORTEX_SIM).....	69
Heat Transfer in Advanced Reactor Cores	71
Integrated 3-D Simulation of Neutronic, Thermal-Hydraulic, and Thermo-Mechanical Phenomena (NUMERICALREACTOR).....	73
Parallel Tools Performance Testing (PARALLEL-TOOLS)	75
Parallelization of the Accelerator Modeling Code <i>elegant</i>	76
Estimation of Protein: Oligonucleotides Binding Energies	78
Quantum Monte Carlo Calculations of Light Nuclei (QMC_FOR_NUCLEI)	79
Multilevel Agent-Based Simulation of Bacterial Chemotaxis (REPAST_CHEMOTAXIS).....	83
Simulations of Spin Wave Excitations in Magnetic Nanoparticles	85

Executive Summary

Argonne National Laboratory founded the Laboratory Computing Resource Center (LCRC) in the spring of 2002 to help meet pressing program needs for computational modeling, simulation, and analysis. The guiding mission is to provide critical computing resources that accelerate the development of high-performance computing expertise, applications, and computations to meet the Laboratory's challenging science and engineering missions.

In September 2002 the LCRC deployed a 350-node computing cluster from Linux NetworX to address Laboratory needs for mid-range supercomputing. This cluster, named "Jazz," achieved over a teraflop of computing power (10^{12} floating-point calculations per second) on standard tests, making it the Laboratory's first terascale computing system and one of the 50 fastest computers in the world at the time. Jazz was made available to early users in November 2002 while the system was undergoing development and configuration. In April 2003, Jazz was officially made available for production operation.

Since then, the Jazz user community has grown steadily. By the end of fiscal year 2006, there were 76 active projects on Jazz involving over 380 scientists and engineers. These projects represent a wide cross-section of Laboratory expertise, including work in biosciences, chemistry, climate, computer science, engineering applications, environmental science, geoscience, information science, materials science, mathematics, nanoscience, nuclear engineering, and physics. Most important, many projects have achieved results that would have been unobtainable without such a computing resource.

The LCRC continues to foster growth in the computational science and engineering capability and quality at the Laboratory. Specific goals include expansion of the use of Jazz to new disciplines and Laboratory initiatives, teaming with Laboratory infrastructure providers to offer more scientific data management capabilities, expanding Argonne staff use of national computing facilities, and improving the scientific reach and performance of Argonne's computational applications. Furthermore, recognizing that Jazz is fully subscribed, with considerable unmet demand, the LCRC has framed a "path forward" for additional computing resources.

The Laboratory Computing Resource Center

The Laboratory Computing Resource Center was established in 2002, based largely on the recommendation of Argonne's Computational Science Advisory Committee. The driving mission of the LCRC is to enable and promote computational science and engineering across the Laboratory, primarily by operating computing facilities and supporting application use and development. The Mathematics and Computer Science Division operates the LCRC on behalf of the Laboratory.

The LCRC's first computer cluster was installed in September 2002. The cluster was named "Jazz" because we expected that the mix of disciplines and applications operating on the system would be diverse, wide ranging, and continually changing. Moreover, the name is easy to remember.

The long-term goal of the LCRC is to develop a vigorous computational science and engineering community at Argonne. Therefore, in addition to supporting Jazz, the LCRC provides consulting services to the Argonne computational research community. These services include training in computation techniques from the fundamentals to advanced topics, assistance with code performance analysis, guidance with algorithm development, and general help and advice.



LCRC resources are available to all Argonne personnel. Non-Argonne collaborators working with an Argonne principal investigator are also welcome. All such personnel who would like to use the Jazz system may sign up for an account on the LCRC Web site. New users are granted an initial allocation of 1,000 computing hours on the system to get started. An investigator wishing to use Jazz for a longer period of time may apply for additional computing allocation via the Web site. All project requests are judged on the basis of scientific merit by the LCRC Allocation Committee, composed of scientists from across the Laboratory.

During 2006, the LCRC put emphasis on the following objectives:

- Continue to operate Jazz as a highly effective production supercomputing resource.
- Encourage the use of Jazz to new projects and disciplines, with an emphasis on the Laboratory's initiatives.
- Help Argonne personnel improve the scientific reach and performance of their computational applications.
- Help large-scale users of Jazz identify opportunities to obtain allocations at national computing resources such as NERSC, NLCF, MHPCC, PACI, and TeraGrid.
- Continue to develop a plan for next-generation LCRC computing resources.

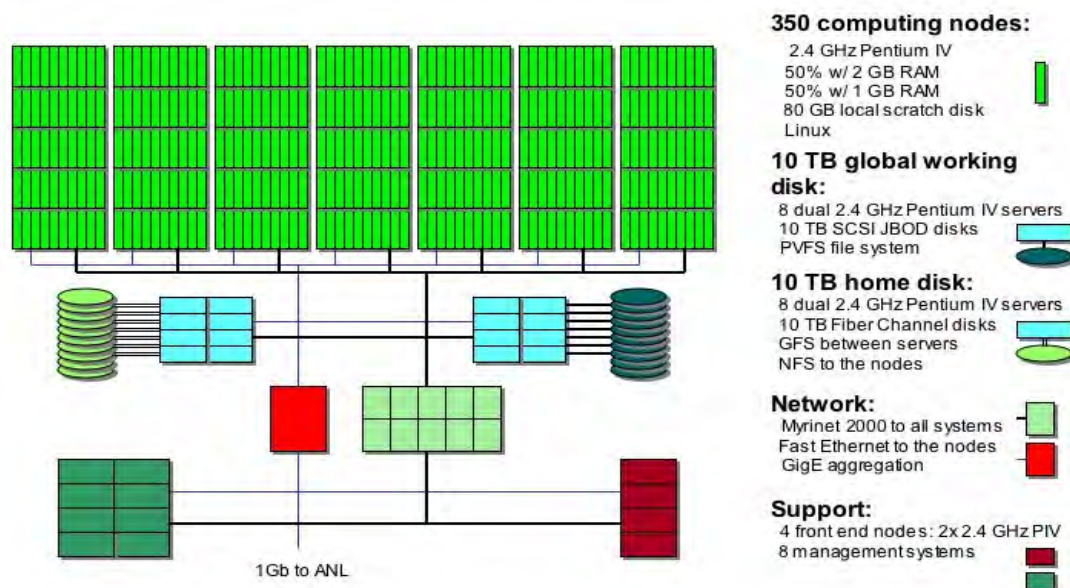
For further information about the LCRC and Jazz, please see the LCRC Web site at <http://www.lcrc.anl.gov/>, or send email to consult@lcrc.anl.gov.

The Jazz Computing Cluster

Jazz is a “Beowulf” cluster, built largely from commodity components. The cluster has, however, been tuned for Argonne-specific use, with diverse tools installed to support parallel computing and project management. Documentation has also been added for user support.

Cluster Configuration

Jazz comprises 350 computing nodes, each with a Pentium Xeon processor and a connection to both Myrinet and Ethernet communication networks. The system has 20 TB of disk storage, half in a global shared file system and half in a high-performance parallel file system.




Cluster Usage

As can be seen in the following chart, the usage of Jazz was high throughout most of fiscal year 2006, with usage often over 80% and exceeding 90% at times. On this chart, 100% would mean that all 350 nodes were used continuously for that day. The maximum sustainable usage on Jazz is about 85% over long periods, given that some nodes go idle when the smallest computation waiting to run is requesting more nodes than are available. Reservations for very large computations (e.g., over 200 nodes) create additional holes as nodes are held open for a while until the large computation has sufficient resources to begin.

With the exception of part of October and a week in January, demand for compute resources on Jazz was sufficient to maintain its maximum sustainable level. Over its production lifetime, Jazz averages 76% utilization compared to the wall clock hours available (including down time) and continued at this rate in 2006. However, utilization of Jazz was down during certain periods

owing to a few events. Most notably, compute node power supplies, which had begun failing in FY2005, continued failing in FY2006, reducing the availability of the cluster. The power supply failures caused intermittent communications disruptions on the high-speed interconnect network (Myrinet), so until affected compute nodes could be repaired, they were removed from the default job queue so that they wouldn't impact multinode jobs. While LCRC staff began to proactively replace power supplies at the start of the fiscal year, the manufacturer of the supplies was unable to provide enough replacements in advance to enable the staff to replace all the supplies before they failed. Consequently, there were periods of time when many nodes were down or not in the default queue while awaiting new power supplies. One such occurrence was during most of April and May when about 40 compute nodes (just over 10% of the cluster) were removed from the default queue due to failed power supplies. Another significant occurrence was in August when the remainder of the nodes (over 70) had their supplies replaced, completing the replacement of the power supplies in all of the compute nodes. Other notable reductions in availability occurred in December, when a Myrinet switch failed, and in March, May, and July when problems with the global, shared file system prevented jobs from accessing files stored in users' home directories.

Jazz Usage by Day



Jazz supports a wide mix of computations, from long runs with a few nodes to short runs with many nodes, with many variations between. On a project-by-project basis, usage varies considerably over time. Projects often start slow, then ramp up, then increase and decrease in cycles as the project progresses, with interruptions for analysis, manuscript preparation,

conferences, and vacations. This usage pattern also causes fluctuations in the overall utilization of Jazz.

Users and Projects


Argonne staff may obtain a startup account for 1,000 hours to become familiar with the Jazz system and evaluate its applicability to their work. Aside from these startup efforts, all computer time is allocated to projects, rather than individuals, where a project is focused on a specific technical activity or avenue of investigation. A project usually has several scientists or engineers working on it, and a person may be active on more than one project.

In FY2006, 13 new LCRC projects were started, and 116 new user accounts were created. By the end of FY2006, there were over 380 active user accounts and 76 active LCRC projects.

Usage by Discipline

Jazz supports research in a wide range of disciplines, as indicated by the chart below. Here the size of the slice is proportional to the fraction of the Jazz resource (number of node hours) used by all projects in a particular area during FY2006. These fractions change over time as new projects are added. In FY2006, nanoscience, physics, and biosciences dominated the Jazz usage.

FY2006 LCRC Usage by Discipline




Job Size Characteristics

Jazz is configured to support parallel computing work of many types, with jobs using a handful of nodes to hundreds of nodes. The number of nodes that a particular application can use effectively depends on many factors, including the computational approach taken, the problem size, and the ratio of computational effort to interprocessor communications required. Jazz is also a multiuser system, and it is almost entirely space-shared, rather than time-shared. Each large job gets a set of computational nodes dedicated to it for the duration of the computation. At any time, several jobs will be running, each using a part of the system.

Data about the number of computing nodes that jobs used provides an interesting view of how the system is being utilized and the degree of scalability or parallelism commonly utilized on Jazz. The chart below shows the amount of computing time used by jobs of various sizes in FY2006. Over 50% of the machine was used for jobs that employed 16–63 processors. Just under 10% of the jobs used 64–350 processors. The jobs that needed only one node accounted for about 12% (the same as FY2005) of the total computing time used on Jazz. We encourage users to submit large parallel jobs so that the interconnect on the machine gets properly utilized.

FY2006 LCRC Usage by Job Size



On a system such as Jazz that is heavily used, job turnaround is related to the number of nodes requested. That puts a practical bound on the turnaround for very large jobs. To compensate, users may reserve a portion of the system for a specific period during a night or weekend to carry out large computations or to meet an urgent computing need. All users are notified in advance of these reservations.

Scientific and Engineering Applications

The LCRC has become an essential resource for a wide range of projects in many Argonne divisions, from production runs to exploratory modeling and simulation research to high-performance algorithm and application development. Currently, 76 projects from 15 divisions use Jazz for research in science and engineering.

In addition, creation of the LCRC has provided an important catalyst for establishing a high-performance computing community in the Laboratory, building relationships that enable cross-disciplinary efforts and developing computational science skills in our mission areas. LCRC staff continue to help build this community by assisting researchers who are getting started with new parallel codes or expanding the range of their codes to address larger, more complex problems. Sharing experience in parallelizing codes can save huge amounts of time and effort. The LCRC also provides a launching point for researchers who gain experience on Jazz and then move up to larger systems at national centers.

Most of the projects using the LCRC fall into three broad, overlapping categories.

- **Strategic.** Jazz is critical to most all of the Laboratory directed research and development (LDRD) strategic initiatives as well as completion of programmatic work in strategic areas. In other cases, Jazz is the testing ground for the concepts that will make up future LDRD and programmatic proposals. For many of these strategic projects, Jazz is the only large computing resource available.
- **Production.** An important component of LCRC's mission is to provide a fast and reliable resource for the Laboratory's production modeling and simulation activities. Some of these production codes are still under development; others are fully mature. Often Jazz is the largest resource available to a research group; for other production applications, Jazz fills the need for quick turnaround on mid-range runs, where larger runs are done on systems bigger than Jazz at NERSC, NCCS, or other national centers. Also, parallel versions of a number of commercial science and engineering codes are available on Jazz, for example, for production computational fluid dynamics computations.
- **Exploratory.** Some Jazz projects are investigating approaches and algorithms for next-generation computations, developing prototype applications with new capabilities, and evaluating software performance, accuracy, or programming models. While large parallel computing resources offer tremendous power, frequent access to them is essential to develop effective algorithms and codes. Challenges include strategies for dividing the problem into many pieces, optimization of the kernels of computation, and management of the massive flow of results to storage. Other projects on Jazz are exploring new designs or models of complex phenomena, forging new paths to extend the range of existing applications, or starting a parallel programming project from scratch and drawing on the experience of the LCRC staff.

LCRC Research Highlights

In the following section, we present examples of research performed with Jazz through 2006. The examples span both the basic and applied missions of the Laboratory and include most of our strategic initiatives. These examples highlight the science and engineering advances being made with Jazz, tackling important problems in nanosciences, physics, biosciences, nuclear engineering, climate, and many other disciplines. The research also encompasses a wide range of computational techniques, including agent-based simulations, specialized forms of computational fluid dynamics, Monte Carlo solvers, and “first principles” solutions.

Integrated Modeling System for Evaluating the Impact of Aerosols on Regional- and Urban-Scale Climates (AERO-MODEL)

PI: V. R. Kotamarthi, Environmental Research

We have completed replacing the I/O interface in the CMAQ model by the parallel NetCDF software developed by the Mathematics and Computer Science Division at Argonne National Laboratory. A number of simulations to test the new interface were performed. We have made additional progress in developing the coupling between CMAQ and WRF/MM5 (Figure 1). The meteorological preprocessor in CMAQ, known as MCIP, has now been moved into CMAQ. This configuration makes it possible for us to perform the next step in coupling the CMAQ and MM5/WRF models using the Model Coupling Toolkit (MCT).

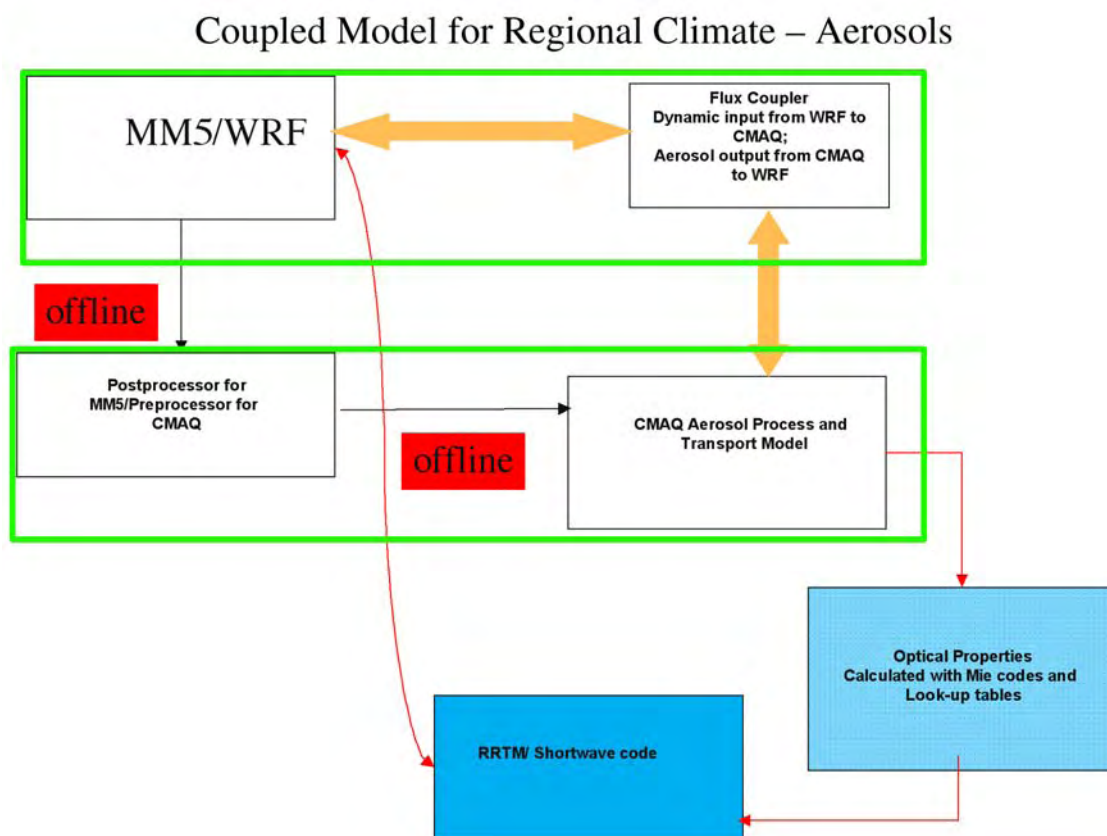


Figure 1. Development of a regional climate-aerosols model. The green boxes show the tasks that have been completed to date. Developing a coupling between MM5/WRF CMAQ using MCT remains unfinished.

Process-Scale Modeling of Atmospheric Chemistry and Transport (ATMOS CHEM)

PI: V. R. Kotamarthi, Environmental Research

The allocated resources were used in principle to perform the calculation discussed below. A significant amount of allocated time was spent on performing mesoscale simulations to complete the analysis of airport-related emissions on Chicago regional air quality. In 2006, a report was submitted to the National Aeronautics and Space Administration, detailing the results.¹

We also have started a project to evaluate the impact on air quality from the development of Oil Shale and Tar Sands (TSOS) in the Colorado, Utah, and Wyoming region under a contract from the Bureau of Land Management, National Science and Technology Center, Denver, Colo. An air-quality modeling analysis was initiated to assess the potential impact on ambient air quality from the TSOS project in the modeling domain (Figure 2).




Figure 2. Modeling domain for the TSOS assessment region.

The CALPUFF modeling system was used as the basis of the modeling analysis for both near- and far-field impact assessments. CALPUFF is a Lagrangian puff model with the capability to simulate regional-scale, long-range dispersion as well as local-scale, short-range dispersion. CALPUFF is recommended for a refined modeling analysis (as opposed to a screening-type analysis) to address the air-quality impacts of pollution transported over relatively long distances. It has three main components: CALMET (a diagnostic three-dimensional meteorological model), CALPUFF (the transport and dispersion model), and CALPOST (a post-processing package). During the past year, we ported and installed the CALPUFF model and its supporting software on Jazz.

We have used Jazz computing resources to complete two full years of MM5 model runs for this domain at 4-km resolution for the inner domain and a 12-km grid resolution outer domain for the years 2004 and 2005, respectively. CALMET, one of four main components of the CALPUFF modeling system (which also includes CALPOST and MM5), has a diagnostic wind model that combines surface and upper-air meteorological data with the diagnostic effects of terrain and other factors to generate three-dimensional wind fields. It also includes other interpolation algorithms that generate three-dimensional temperature, pressure, and other meteorological variables, as well as two-dimensional precipitation fields.

The data flow diagram in Figure 3 shows model input and observational data input used in generating the CALMET diagnostic wind fields used by CALPUFF. The 12-km MM5 hourly

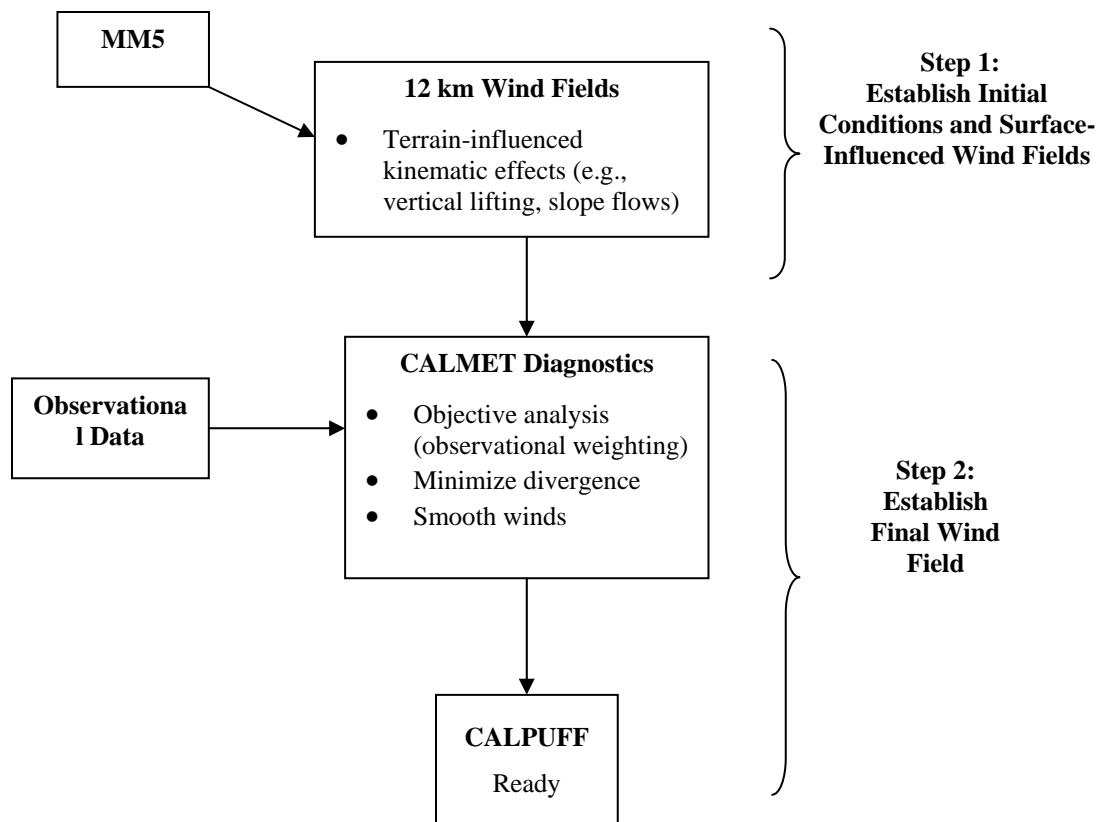


Figure 3. Establish wind fields for modeling domain.

meteorological fields are used for both establishing initial conditions and generating vertically and tangentially forced complex terrain-induced wind fields. Surface and upper air observations are used to produce interpolated and smoothed winds down to finer-scale resolution (between 4–12 km). We are completing the calculations using CALPUFF.

Publications/Presentations

1. M. Dutta, V. R. Kotamarthi, and D. J. Wuebbles, Impact of airport emissions on urban air quality in the Chicago area, Appendix A, March 2006.

Development and Benchmarking of Eulerian 2-Phase Boiling Models for Computational Fluid Dynamics Simulation

PI: W.D. Pointer, Nuclear Engineering

In this project, a combination of first principles and semi-empirical models are being implemented within the framework of the commercial computational fluid dynamics code Star-CD. This implementation provides a practical multidimensional simulation capability for application to the development of advanced fuel assembly designs for conventional and future light water reactors (LWRs). Existing Eulerian bubble/droplet models within the code have been extended to allow for mass and energy exchange in addition to momentum exchange between the two fluids. A flow regime map scheme is being developed and implemented to enable a smooth transition between the bubbly flow model and the droplet/wall film flow model as the vapor fraction increases.

Preliminary beta versions of the Eulerian two-phase boiling model were installed on the Jazz system with the LCRC's assistance in early 2006. Argonne selected the OECD-NEA/US-NRC Benchmark based on NUPEC BWR Full-size Fine-mesh Bundle Tests (BFBTs) as a means of validating existing capability and directing the future development effort. Argonne joined the benchmark program as a contributing participant. An initial model of the selected boiling water reactor (BWR) fuel assembly configuration was developed, as shown in Figure 4, based on the parameters provided as part of the benchmark specification. Initial simulations were completed for three cases selected by the benchmark committee. Although detailed data describing the vapor distribution in the assembly will not be released for comparison until final results are submitted later this year, limited sub-channel averaged data is available for comparison with predictions.




Figure 4. Computational model of the BFBT test assembly. Shown are the computational mesh distribution, the components of the complete model, and a detailed view of the simplified representation of a single spacer grid.

While these models are currently quite coarse and a manageable size, inefficiencies in the current scheme result in extremely long run times. The availability of large platforms such as Jazz allow these issues to be partially overcome by brute-force application of more processors to accomplish near-term applications goals, while longer-term code improvement efforts are under way. It is expected that the validation phase will be completed in early FY2007, and more significant effort will be focused on run-time improvement later that year.

To date, the results are the only multidimensional computational fluid dynamic results submitted to the benchmark. All other participants are focusing on empirical correlation-based one-dimensional subchannel modeling codes. The large-scale simulations completed on Jazz have provided direction for enhancements to the code. A significant improvement in the accuracy of the detailed vapor distribution well beyond the current state of the art in multidimensional boiling models is expected when these improvements are employed in early FY2007.

An Eulerian-Eulerian boiling model was implemented within the Star-CD CFD code by participants from the code vendor. Initial testing, model development, and validations were completed by Argonne participants.

Studies of the Water-Gas Shift Reaction over Mono- and Bimetallic Catalyst Surfaces (BIMETALLIC-WGS)

PI: John Krebs, Chemical Engineering

Our group was involved in the development of catalytic materials for the reforming, or conversion, of hydrocarbon fuels to a hydrogen-rich gas mixture that can be fed to hydrogen (H_2) fuel cells. In addition to H_2 , the breakdown of hydrocarbon fuels also produces carbon monoxide (CO) and carbon dioxide (CO_2). Carbon monoxide has been shown to inhibit the activity of H_2 fuel cells at levels of 10–100 ppm. Levels of CO in the product stream are as high as 10%; thus, methods of reducing CO levels have been implemented. The water-gas shift (WGS) reaction, $CO + H_2O = CO_2 + H_2$, is being used to remove the CO contaminant and produce additional H_2 . Currently, the most active WGS catalysts suitable for use in fuel processors are Pt-based catalysts. Although very effective, the Pt-based catalysts are expensive, and effort is being directed to develop a less expensive, nonprecious-metal-based WGS catalyst.

Ruthenium (Ru) and cobalt (Co) are two potential alternatives to Pt. Both Ru and Co have been reported to be highly active for the WGS reaction and are less expensive than Pt. However, a major issue with Ru or Co as WGS catalysts is that they also promote the formation of methane (CH_4) via the CO methanation reaction, $CO + 3H_2 = CH_4 + H_2O$. Formation of CH_4 and other short-chain hydrocarbons via other Fischer-Tropsch reactions is undesirable because it consumes the desired H_2 into a form that cannot be converted to electrical power by the H_2 -fuel cell. It is believed that the selectivity of Ru or Co toward the WGS or methanation reaction pathways is based on the nature of CO adsorption on Ru and Co surfaces. During WGS, CO is believed to adsorb associatively (no C-O bond breaking), while during methanation it adsorbs dissociatively.

Our goal was to develop a modified Ru or Co-based catalyst that can selectively operate as a WGS catalyst under the H_2 -rich environment of post-reformer product. We were pursuing the modification of the adsorptive properties of CO on Ru or Co via the addition of a second metal. Understanding how the secondary metal influences reactions over Ru- or Co-based catalysts is essential to developing bimetallic WGS catalysts that have improved selectivity to H_2 . Theoretical modeling calculations are useful in determining both electronic and energetic trends that may not otherwise be possible using physical techniques. Theoretical calculations examining the energetics of the WGS and methanation reactions were conducted over mono-metallic surfaces first to establish general trends over the periods of interest in the periodic table. Calculations of bimetallic systems are now being conducted. Experimental results coupled with theoretical studies of the WGS reaction over monometallic surfaces will be employed to design bimetallic, WGS-selective Ru- or Co-based catalysts.

Wave Packet Dynamics with Cartesian Coordinates (CARTESIAN)

PI: Stephen Gray, Chemistry

Understanding the behavior of atoms and molecules at the fundamental quantum mechanical level is important for understanding a wide range of problems in chemistry, physics, and nanoscience. This project's aim is to develop a Cartesian-coordinates-based program for obtaining such an understanding. The code will solve the Schrödinger equation governing the motion of atoms and molecules, leading to wavefunctions or wave packets. Such functions tell us the nature of the quantum mechanically allowed motions in the system. Most numerical methods for such problems involve internal coordinates and complicated angular momentum considerations, which are cost-effective for few-atom systems only. A conceptually simpler approach is to represent the wavefunction in Cartesian coordinates. While the wavefunction is much less compact, the required operations to generate it, though perhaps more numerous, are simpler and more readily implemented on parallel computers. As one begins to consider larger systems, or even atoms or molecules confined to move in nanostructured regions (Figure 5), a purely Cartesian representation, implemented in a parallel computing environment, should be a good way of obtaining quantum information.

Work on hydrogen molecules confined to carbon nanotubes was completed, which involved a mixed Cartesian and angular representation. We identified an interesting new problem to test the full Cartesian program—the endohedral complex Li@C_{60} . This system consists of a lithium atom encapsulated in the soccer-ball-shaped C_{60} molecule (“buckyball”). Preliminary calculations yielded vibrational energy levels that were in good accord with recent experimental results. A more thorough analysis is now in progress.

The calculations will allow experimental spectra to be better understood. They will also provide information that cannot be found easily by experiment, for example, the regions in the C_{60} molecule where the atom is most likely to be (not necessarily the center).

We developed an iterative, matrix eigenvalue program (Lanczos method) for the determination of the energy levels. Its scalability, performance, and so forth on Jazz remain to be determined in FY2007.

Publications/Presentations

1. T. Lu, E. M. Goldfield, and S. K. Gray, Quantum states of hydrogen and its isotopes confined in single-walled carbon nanotubes: dependence on interaction potential and extreme two-dimensional confinement, *J. Phys. Chem. B* 110, 1742–1751, 2006.




Figure 5. An endohedral complex consisting of an atom encapsulated by a C_{60} molecule.

Development and Performance of the Community Climate System Model on Linux Clusters (CCSM2)

PIs: Jay Larson, Robert Jacob, and Jong Kim, Mathematics and Computer Science

The Community Climate System Model (CCSM) is a fully coupled, global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climate states. CCSM is the primary model for global climate research in the United States and is supported on a variety of computer systems. It also is being used to assess climate change impacts for the Intergovernmental Panel on Climate Changes (IPCC) Fourth Assessment. The CCSM is the result of a community modeling effort sponsored by the National Science Foundation and the U.S. Department of Energy. More information is at www.cesm.ucar.edu. The Model Coupling Toolkit is the coupling software used in CCSM to join the atmosphere, ocean, sea-ice, and land models into a complete system.

In an effort to develop a computationally objective parameter tuning scheme of the CCSM, we performed a broad range of sensitivity analysis and parameter tuning experiments of the CCSM model's sea-ice modeling code (CICE) with the automatic differentiation (AD) technique. Two different grid systems were used to compute the sensitivities of global sea ice conditions with regard to various thermodynamic and dynamic modeling parameters. The problem size of the coarse (3-degree) grid system is 100 by 101 global grid points with five different ice categories and four ice layers in a single ice category. The multiyear spin-up result of ice thickness, concentration, and velocity fields on January 1 was used as the initial state for the experiments. A significant parallel performance of the AD-based CICE code was observed on Jazz (scaled parallel efficiency of 69% on up to 60 processors with a fine resolution of 0.4 degree problem set with 900 by 601 global grid points).

Sensitivity experiments were carried out with respect to two main dependent variables, ice thickness and hemispheric ice volume. We chose average ice thickness (the ice volume per unit area) as the dependent variable because it represents the integrated evolution of the ice pack through time and is critical for climate simulations, particularly for the fresh water balance of the climate system. Sea-ice concentration also plays a crucial role in Earth's energy budget because of the insulating and reflective properties of ice on the ocean surface; likewise, sea-ice velocity and associated deformation fields direct the pack-ice evolution. Overall, density, radiative, and conductive thermodynamic parameters strongly affect the ice thickness simulation. The ice-ocean drag coefficient and ridging parameters also play an important role locally. The details of the sensitivity experiments are found in a paper by Kim et al. (2006).¹ Figure 6 shows the computed sensitivities of ice thickness with regard to ice density on the fine-resolution grid system.

The development of CCSM involves various parameterization schemes that employ a number of model parameter values with different scales of uncertainty. Many model parameters can be tuned to improve various aspects of the simulated climate. The major impact of this work would be the enhancement of the CCSM modeling components with objective parametric tuning schemes based on AD techniques. Through this effort, a broadly useful computational method

has been implemented to make optimal use of climatological observation data to tune the CCSM modeling components. This has been demonstrated with the CICE model.

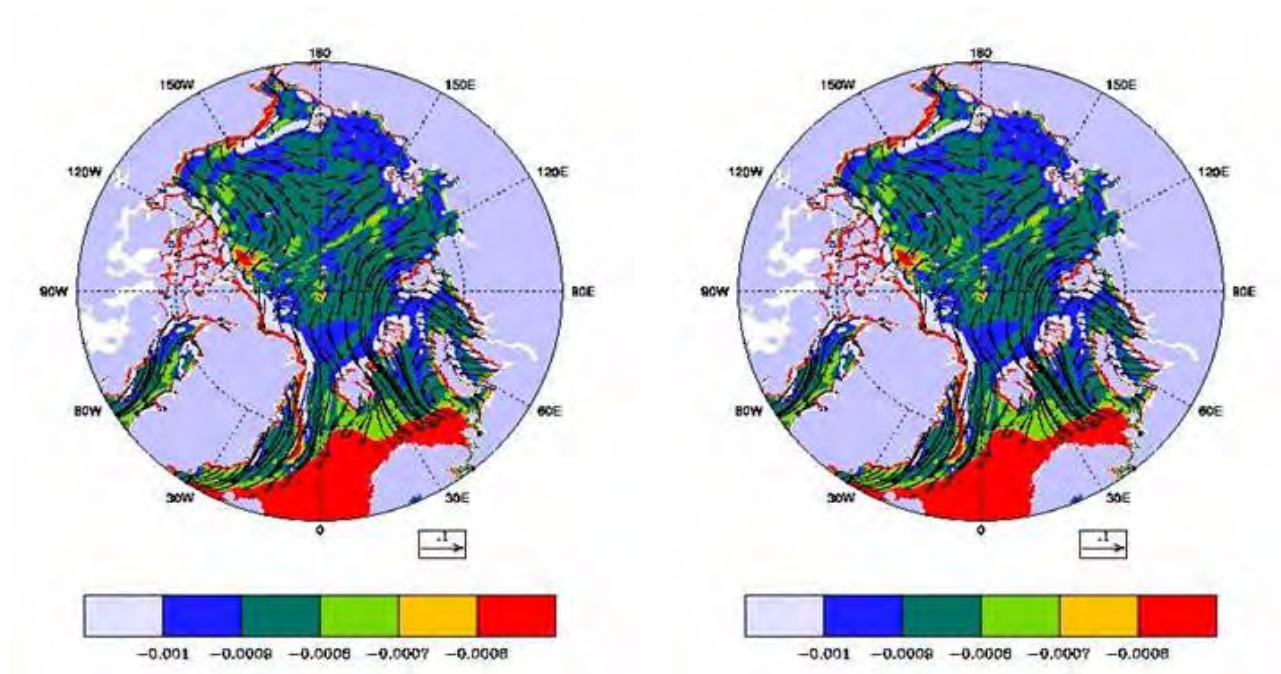


Figure 6. Color map of computed sensitivities of sea-ice thickness with regard to ice density modeling parameter.

We developed a derivative of the CICE model using automatic differentiation called CICE-AD. We also developed a new MPI replacement library for MCT called mpi-serial. This allows models coupled with MCT to function without a MPI library.

Publications/Presentations

1. J. G. Kim, E. C. Hunke, and W. H. Lipscomb, Sensitivity analysis and parameter tuning scheme for global sea-ice modeling, *Ocean Modeling Journal*, 14 (1–2), 61–80, 2006.

Spatiotemporal Chaos in Rayleigh-Benard Convection

PI: Paul Fischer, Mathematics and Computer Science

Many frontiers of science and numerous technological problems involve nonequilibrium systems that create intricate spatiotemporal patterns. These systems are difficult to analyze because of the nonlinear way in which spatiotemporal patterns affect transport of energy and matter and, in turn, changes in transport modify the spatiotemporal patterns. A particular challenge is to understand spatiotemporal chaos, a commonly observed behavior of nonequilibrium systems in which the properties of the system evolve nonperiodically in time (chaos) and space. This work is a computational effort to push as deeply as is currently possible toward a quantitative understanding of how complex dynamics arise in a spatially extended nonequilibrium system.

We have calculated a defining characteristic of chaotic dynamics, the extreme sensitivity of the dynamics on initial conditions. This is illustrated by Edward Lorenz's now famous question, "Does the flap of a butterfly's wings in Brazil set off a tornado in Texas?" This sensitivity is quantified by the Lyapunov exponents, which give the exponential expansion of nearby solutions. At least one positive Lyapunov exponent is a common definition of chaos. However, the Lyapunov exponents are hard to access in real experiments, and, in practice, experimental systems are often said to be chaotic-based on weaker criteria, such as a broad-band component to the power spectrum of a time-dependent measurement. From knowledge of the Lyapunov exponents, we have calculated the spatial variation of the largest growing perturbations (Lyapunov vectors) and the fractal dimension of the chaotic attractor describing the dynamics as a function of system size. Jazz has played an important role in the development and exploration of these capabilities.

Knowledge of the Lyapunov exponents, Lyapunov vectors, and fractal dimension has allowed us to probe fundamental concepts of spatiotemporal chaos. We have shown that fluid convection is indeed chaotic for experimentally accessible conditions. We have also shown that the chaos is extensive in that the fractal dimension scales linearly with system size. From calculations of the Lyapunov vectors, we have determined that localized defect structures contribute significantly to the disorder of the patterns (Figure 7). The information gained from this fluid system yields insight that may find use in the study of the dynamics of more complex real-world systems.




Figure 7. Spectral-element numerical simulation of chaotic Rayleigh-Benard convection with aspect ratio 10 (layer depth to height). Shown is a horizontal two-dimensional slice at mid-depth (left). In the fluid-flow field, red indicates warm rising fluid, and blue indicates cool falling fluid (right). In the leading-order Lyapunov vector, blue indicates large negative growing perturbations, and red indicates large positive growing perturbations. The results suggest that highly localized defects significantly contribute to regions where perturbations are the largest.

We have continued to develop the Nek5000 code to calculate the spectrum of the largest and most dominant Lyapunov exponents. The Lyapunov exponents are computed by measuring trajectory separation in several independent linearized computations that are executed in parallel with a driving Rayleigh-Benard convection problem. The Gram-Schmidt orthogonalized vectors ordered by the exponents are the Lyapunov vectors. Knowledge of the Lyapunov spectra permits a calculation of the fractal dimension that describes the ball of initial conditions that neither exponentially grows nor shrinks on average under the dynamics.

Publications/Presentations

1. N. Becker, J. D. Scheel, M. C. Cross, and G. A. Ahlers, Effect of the centrifugal force on domain chaos in Rayleigh-Benard convection, *Physical Review E*, 73, 066309, 2006.
2. M. R. Paul, M. Einarsson, P. F. Fischer, and M. C. Cross, Extensive chaos in Rayleigh-Benard convection, *Physical Review Letters*, submitted in August 2006.
3. J. D. Scheel and M. C. Cross, Lyapunov exponents for small aspect ratio in Rayleigh Benard convection, *Physical Review E*, 74 (6). 0663012, ISSN 1539-3755, 2006.
4. J. D. Scheel, Rotating Rayleigh-Benard convection, Ph.D. thesis, physics, Caltech, published electronically, 2006.
5. M. Einarsson, Quantifying spatiotemporal chaos, master's thesis, Virginia Tech, adviser M. R. Paul, July 2006.
6. M.R. Paul, New insights from numerics: Spatiotemporal chaos in fluid convection and the stochastic dynamics of nanoscale cantilevers, University of California at Los Angeles, April 2006.
7. M. Einarsson, P. F. Fischer, and M. R. Paul, Quantifying spatiotemporal chaos in Rayleigh-Benard convection, American Physical Society Annual March Meeting, Baltimore, Md., March 2006.

Geometry Optimization for Molecular Electronic Structure Using Multireference CI Methods

PI: Thomas Mueller, Mike Minkoff, and Ron Shepard, Chemistry

We are using the COLUMBUS System for molecular electronic structure, jointly developed by the University of Vienna and Argonne National Laboratory. We specifically use the MR-CISD method that is generally applicable to accurate calculations of ground and excited state properties. The central issue is the evaluation of a matrix-vector product for which the matrix dimension easily exceeds 10^9 . This matrix is never directly calculated, and the computation of this matrix-vector product constitutes nearly the entire computation. Coarse-grained parallelism is obtained by using a segmented matrix-vector product approach. Our work specifically involves geometry optimizations using multireference CI methods. The basis sets involve a sequence of correlation-consistent Gaussian orbital basis sets (cc-pVDZ, cc-pVTZ, cc-pVQZ), which are increasingly more accurate and computationally intensive. The software is parallelized via MPI by using a SPMD programming model. The program uses the GA library, which is included in the program source distribution. Our benchmarking work involves ethylene (C_2H_4) as an example of the utility of our approach. The underlying electronic structure calculation involves an eigenvalue problem with dimensions that range from 25 million to nearly 1 billion and is solved via iterative techniques.

During this year, we progressed on the Ethernet aspect of our benchmarking project and related production studies. The benchmarking experiments involve increasing larger basis sets (cc-pVDZ, cc-pVTZ, and cc-pVQZ), using both Ethernet and Myrinet communication on the Jazz cluster. Our work on Myrinet has been limited because of errors in the GA library on that network on the Jazz system. We are collaborating with our colleagues at Pacific Northwest National Laboratory and Jülich in addressing this issue. A major advance in conducting this study is testing the previously developed software tool for modeling the performance of COLUMBUS and using it to determine optimal segmentation of the matrix-vector product over processors. The current COLUMBUS parallel code requires user intuition to conduct this optimization. We have found that the estimates of optimal segmentation obtained with our model were inaccurate for large numbers of nodes. We have enhanced the model to more accurately




Figure 8. The color patches encode the locality resolved number of task-type specific (n-ext. segtype/segtype) contributions to the CI matrix. The left and right segtype specifications refer to the vertical and horizontal range of the depicted subblock of the CI matrix. The color scale denotes a logarithmically increasing number of contributions; white denotes no contributions. The actual computational effort requires a task-type-specific scaling.

deal with communication costs. We are now completing these benchmarking studies using this enhanced segmentation prediction tool for the Ethernet network. Our collaborator has tested this model on other networks as well.

Late last year, we modified COLUMBUS to improve its communication requirements. The communication cost varies with the square of the number of segments. By reducing the number of segments, we were able to decrease the memory required per processor and the bandwidth requirement. This modification is useful when memory is limited (as for the TZ and QZ basis sets). In experiments with the segmentation reduction approach, we have been able to reduce computation time by up to one-half of that previously required.

Figure 8 illustrates the details of integral type in the CI matrix for obtaining the segmentation reduction for decreasing communication costs.

This work is a component of three related projects in the SciDAC/BES project: Advanced Software for the Calculation of Thermochemistry, Kinetics, and Dynamics (R. Shepard, PI; S. Gray, M. Minkoff, and A. Wagner, Argonne collaborating staff).

Publications/Presentations

1. T. Mueller, Parallel MR-CISD: Basic concepts, talk presented at the COLUMBUS Programmer's Workshop, a Theory Institute Workshop held by Argonne National Laboratory's Chemistry Division, August 15–19, 2005.
2. T. Mueller, Modeling the performance of parallel MR-CISD, talk presented at the COLUMBUS Programmer's Workshop, a Theory Institute Workshop held by Argonne National Laboratory's Chemistry Division, August 15–19, 2005.
3. M. Minkoff, Parallel CIUDG benchmarks, talk presented at the COLUMBUS Programmer's Workshop, a Theory Institute Workshop held by Argonne National Laboratory's Chemistry Division, August 15–19, 2005.

Computational Nanoscience (COMPNANO)

PI: Peter Zapol, Chemistry/Materials Science/Center for Nanoscale Materials

Reactivity Control by Shape and Size of Nanoparticles (H. Iddir, P. Zapol, L. Curtiss)

Computational nanoscience projects involving density functional theory, molecular orbital theory, and tight-binding methods were performed at the LCRC facility this year.

Metal nanoparticles are currently receiving a lot of attention because of their suitability for advanced fuel cell applications. This work is focused on noble metal nanoparticles in oxygen atmosphere to examine the effects of various adsorption configurations on the equilibrium shape. We also included the influence of oxide support in the study. Size, phase, and morphology of the nanoparticles are the critical parameters determining their performance in particular applications.

A thermodynamic model devised to describe the shape of nanoparticles as a function of size provides means to predict the stability of metal nanoparticles. Parameters of the model are determined by density functional calculations. Such calculations performed on Jazz are used to accurately determine surface and interface energies. We also investigated the effect on the surface energies of oxygen adsorption (at different coverage) on the three low-index surfaces (Figure 9). These calculations are very demanding, as larger surface cells and rather thick slabs are needed to get meaningful results. The calculations are still ongoing.

The effects of adsorption on nanocrystal morphology have been addressed and comparisons drawn with previously reported studies. Also, optimized structures and calculated electronic properties were compared with experimental results. These studies help experimentalists to understand the properties of metal nanoparticles and tune synthesis conditions for particular applications.




Figure 9. Top (left) and side (right) view of 1 ML of oxygen adsorbed on a (110) “missing-row” metal surface.

Reactivity of Patterned Carbon Nanotubes (M. Sternberg, D. Horner, P. Redfern, L. Curtiss, P. Zapol)

Carbon nanotubes (CNTs) have held promise for use in molecular electronic devices since the demonstration of a nanotube as a working transistor. More recently, progress toward using CNTs in quantum computing has been made by using a CNT with defects. Defects in the tube wall change the electronic properties of CNTs and can lead to interesting new applications.

The DFTB and the Gaussian03 quantum chemistry package were used for studying Stone-Wales and ad-defects in carbon nanotubes as well as reactions of different molecules with them. Jazz was essential for large-memory and CPU-intensive parallel jobs to overcome limitations of other locally available machines. It was found that formation energies of the ad-dimer defects are sufficiently low to observe these defects experimentally. The reactivity of ozone and other small molecules changes dramatically when the defects are present.

Extensive testing and comparison of ScaLAPACK and PETSc/SIPs solvers for use in DFTB for various dimensionalities of physical systems have shown that the SIPs is more efficient than ScaLAPACK for sparse matrices arising from physical systems of low dimensionality, such as nanotubes and nanowires. For 2-D systems, SIPs is on par with ScaLAPACK, but for extended 2-D and 3-D systems, ScaLAPACK does perform somewhat better. Progress also has been made on large-scale parallel replica dynamics to extend the time scale for nanoscale materials modeling.

Publications/Presentations

1. M. Sternberg, L. A. Curtiss, D. M. Gruen, G. Kedziora, D. Horner, P. C. Redfern, and P. Zapol, Carbon ad-dimer defects in carbon nanotubes, *Physical Review Letters* 96, 075506, 2006.
2. H. Zhang, B. Smith, M. Sternberg, and P. Zapol, SIPs: Shift-and-invert parallel spectral transformations, *ACM Transactions on Mathematical Software (TOMS)*, accepted May 15, 2006.

Electronic Excitation Calculations on Organic Semiconductor Copper Phthalocyanine (CuPc)

PIs: Chaminda Nalaka Kodituwakku, X-Ray Science Division, Argonne National Laboratory, and Western Michigan University, and Clement Burns, Western Michigan University

Organic semiconductors are expected to be widely used in future applications such as displays, light-emitting sources, and solar cells because of their low cost and high efficiency. Several applications such as displays in digital cameras, organic light-emitting diodes, and sensors are already important. The possibility of creating efficient, low-cost solar cells is an area of great interest. Currently, copper phthalocyanine (CuPc) is one of the main materials being tested for solar cell applications, and initial work shows great promise for developing low-cost solar cells. However, relevant fundamental properties of CuPc for device applications (such as the excited state energies and properties) are not well understood. We wish to carry out first principles calculations of the excitations in this material to better understand its properties.

We carried out energy optimizations on single-molecule CuPc. CuPc is an almost flat molecule that has almost D_{4h} symmetry with a small distortion perpendicular to the molecular plan (Figure 10). We used both symmetrized and actual coordinates to do the geometrical optimization. These calculations were done under B3LYP and Hartree-Fock theories. In B3LYP, the actual data optimization was carried out by using 8 nodes and took about 72 hours. The symmetric case with the 8 nodes took about 2.5 hours.




Figure 10. CuPc molecule.

Both optimizations were checked for stability and were confirmed to be stable. Time-dependent density functional theory (TD-DFT) calculations were carried out on these optimized molecules.

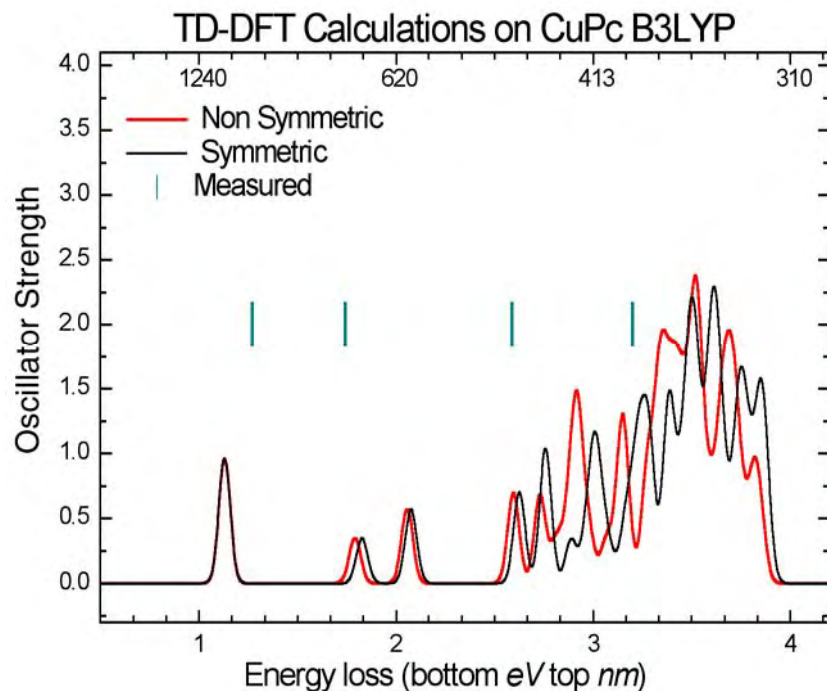


Figure 11. Electron excitation probability vs. energy loss. The Y scales were not normalized to each other.

Currently, we are carrying out extensive electronic excitation calculations of single molecule CuPc for higher energies. These calculations need a large amount of computer power (e.g., 18 nodes with 2 GB memory). Also, we are attempting to match our measured data, including intensities, to the calculated data (Figure 11). All these calculations were done by using Gaussian03. The Jazz cluster had Gaussian98 and upgraded to Gaussian03. This upgrade helped to do these calculations much faster with more functionality.

Kinetics of Enzymatic DNA Repair

PIs: Aaron Dinner and Stuart Rice, Chemistry, The University of Chicago, and Argonne Office of the Director

DNA repair proteins are essential for maintaining the integrity of genetic information. Our studies seek to understand how these proteins bind to damaged DNA, their substrate, by examining two paradigmatic systems: the human proteins uracil-DNA glycosylase (UDG) and O6-alkylguanine-DNA alkyltransferase (AGT). UDG is the prototypical base excision repair protein; it removes the normal RNA base uracil from DNA by cutting the C1'-N1 bond that links the DNA side chain to the DNA backbone. In contrast to UDG and most other repair proteins, AGT directly reverses damage. In addition to shedding light on the mechanisms by which these proteins interact with DNA, the studies have driven significant advances in general methods for treating dynamics in complex systems as detailed below.

We have applied a novel path-based method described below to the study of nucleotide flipping by AGT. In contrast to a simple “push-pull” mechanism, the simulations reveal a two-step mechanism involving a major intermediate in which the O6-methylguanine lesion (mG) is out of the base stack but not yet in the active site. Each step involves a minor intermediate as well—one in which the “arginine finger” (Arg128) is partially intercalated in the base stack. However, one side chain dihedral angle is in the *cis* rather than the *trans* rotamer, and another intermediate, in which the base is in the active site but not yet at its maximum depth, in the pocket. A complete determination of the mechanism of nucleotide flipping requires identification of transition states that connect the stable states and their characterization by a small number of physically meaningful coordinates. To this end, we define a transition state as a configuration that has an equal likelihood of leading to the product and reactant basins in additional molecular dynamics simulations initiated in that configuration with random velocities drawn from a Maxwell-Boltzmann distribution. Preliminary analysis of the first step of the flip shows that the transition state involves a triad between the guanidinium group and the bases of the separating Watson-Crick pair, which corresponds to a “hand-off” of the orphaned cytosine from the mG to Arg128 (Figure 12).

The path-based methods that we have developed enable, for the first time, unbiased treatment of DNA binding dynamics and can be applied to other molecular processes of comparable complexity. Most path-based methods minimize a functional of the potential energy of structures along the path. In contrast, transition path sampling (TPS) weights paths that connect two stable states just as they would be in straightforward molecular dynamics simulations. This aspect of the method is important for treating processes like DNA binding, in which variations in entropy are expected to play a significant role. TPS can be viewed as a Monte Carlo (MC) procedure in the space of whole trajectories. One significant problem with TPS, however, is that it is often very hard to obtain an initial reactive trajectory from which to harvest additional ones. We recently solved this problem by introducing a systematic procedure for relaxing from steered molecular dynamics (SMD) trajectories to unbiased ones. In SMD simulations, minima of harmonic restraints are advanced from one basin to another to drive the system along the speculated path of the reaction; a transition is always made because the force constant for the

restraints can be arbitrarily large. In the new procedure, we “anneal” the paths by shooting steered trajectories in TPS with progressively weaker bias.

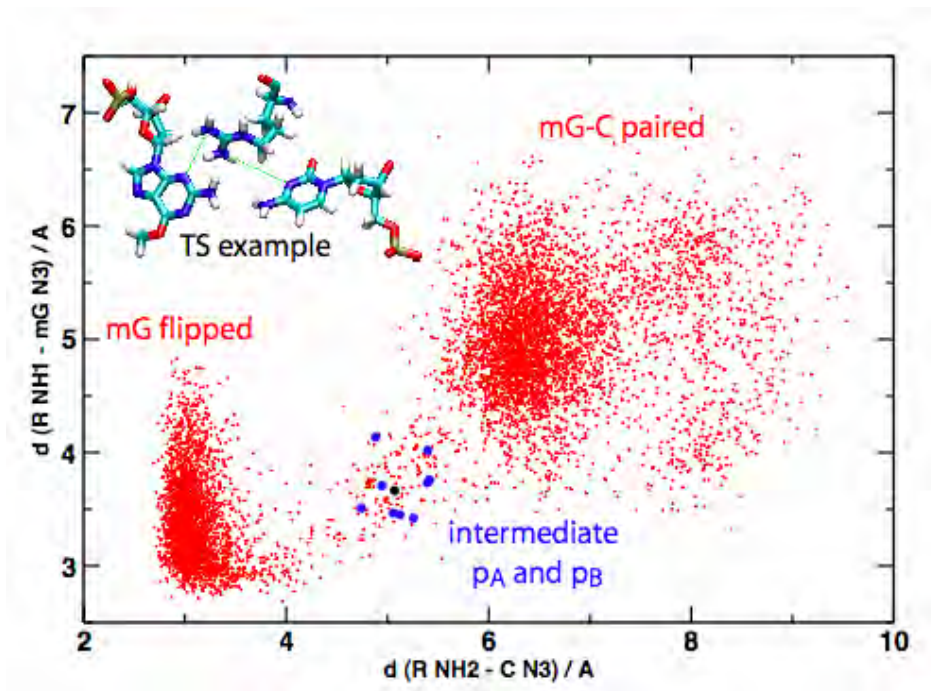


Figure 12. Red points show structures from trajectories in which the AGT arginine finger inserts and the methylguanine and cytosine separate. Blue/black points are dynamically defined transition states; the black point corresponds to the structure shown with the distances used to project the data indicated in green. The separation between the stable states and the transition states indicates that these coordinates provide a good description of the reaction progress.

We implemented the bias annealing method in the CHARMM program, which is widely used for biomolecular simulations. In addition, bugs associated with running TPS in parallel were fixed. These advances in method have been incorporated into the publicly distributed version of the program.

Publications/Presentations

1. A. Ma, J. Hu, M. Karplus, and A. R. Dinner, Implications of alternative substrate binding modes for catalysis by uracil-DNA glycosylase: An apparent discrepancy resolved, *Biochemistry*, accepted pending minor revision.
2. J. Hu, A. Ma, and A. R. Dinner, Bias annealing: A method for obtaining transition paths de novo. *J. Chem. Phys.*, in press.

Meson and Baryon Observables in a Dyson-Schwinger Approach (DSE_EXC_PION)

PIs: Andreas Krassnigg, Craig Roberts, and Stewart Wright, Physics

Our project deals with a nonperturbative continuum approach to QCD, the fundamental theory of the strong interaction between quarks and gluons. This approach uses coupled integral equations to calculate basic building blocks, the Green functions, of the theory. A typical solution in our approach yields functions that describe quarks and hadrons on a common basis (in practice with the same quark-gluon interaction) such that the symmetries of the underlying theory are respected and their manifestations visible in the numerical results. By describing quarks (the constituents) and hadrons as quark-antiquark or three-quark bound states on the same footing, we can reliably address one of the key questions of modern physics, namely, the mechanism behind confinement. This property of quarks means that no single quark has ever reached a detector.

The project's main directions are twofold. On the one hand, we investigate mesons and their radial excitations in order to study the long-range part of the strong interaction; this part of the interaction is not accessible by perturbative methods and requires a nonperturbative framework. On the other hand, we investigate the structure of nucleons by calculating properties such as their electroweak and strong form factors. These are key elements in our understanding of the nucleon and its structure. A manifestly Poincaré covariant approach such as ours is ideal to compute form factors, in particular, on the domain of high momentum transfer that is probed with modern accelerators.

In 2006, we used Jazz as follows:

- 1) We investigated masses and decay constants of pseudoscalar and scalar mesons from the chiral limit up to and beyond the charm-quark mass. For excited-state masses, we use ratios of quantities to reliably extract predictions from our results, since they depend on the model parameters. This is justified because these ratios remain constant over the domain of parameters under consideration. We have found that ground- and excited-pseudoscalar-meson leptonic decay constants reach a maximum magnitude at about the charm-quark mass. For higher quark masses, the decay constant decreases, which is analogous to the well-established heavy-quark limit behavior for a meson composed of one heavy and one light quark (Figure 13).
- 2) Our study of the first radial excitation of the kaon and its leptonic decay constant has been extended to pseudoscalar and scalar meson ground and excited states involving heavier quarks, such as D and D_s states.
- 3) We studied the possibility and reliability of methods to extract information about meson masses and decay constants from information obtained solely in the region of Euclidean momenta; such methods are used, for example, in lattice computations. The data for these studies were generated from solutions of the inhomogeneous Bethe-Salpeter equation for the pseudoscalar and scalar meson channels.

- 4) Additionally, Jazz gives us the ability to investigate possible extensions of our ongoing studies without the immediate danger of not being able to actually do the calculations required.

In FY2006, the meson code was reorganized completely. Various versions of the code have been combined into a single set of routines that can now handle any kind of meson and calculate (in principle, limited by the memory available) any number of radial excitations in a given channel. Some algorithmic improvements have increased the performance of the code. For the nucleon, code has been written to calculate electroweak and strong nucleon form factors using parallel routines. This enables more realistic calculations and the investigation of more demanding problems in the study of the nucleon.




Figure 13. Results for decay constants of pseudoscalar ground and excited mesons, f_{gr} and f_{ex} , plotted from the chiral limit to a quark mass beyond the charm-quark mass.

Publications/Presentations

1. V. V. Flambaum, A. Hoell, P. Jaikumar, C. D. Roberts, and S. V. Wright, Sigma terms of light-quark hadrons, *Few Body Syst.* 38:31-51, [nucl-th/0510075], 2006.
2. C. D. Roberts, M. S. Bhagwat, A. Hoell, and S. V. Wright, Aspects of hadron physics, to appear in *Hadron Structure and Nonperturbative QCD*, Proceedings of the 44th Winter School on Theoretical Physics (IUTP 44), Schladming, Austria: 11–18 March 2006, Lecture Notes in Physics (Springer-Verlag, Vienna).
3. C. D. Roberts, From quarks to reality—an issue of substance, colloquium given at UNAM, Mexico City, Mexico, November 17, 2005.
4. C. D. Roberts, Aspects of hadron physics, lectures given at the 44th Universitaetswochen fuer Theoretische Physik, Schladming, Austria, March 11-18, 2006.
5. C. D. Roberts, Dyson-Schwinger equations: An update, talk given at the Light-Cone QCD and Nonperturbative Hadron Physics (LC2006), University of Minnesota, Minneapolis, Minn., May 18, 2006.
6. A. Krassnigg, Meson spectroscopy and properties in a Dyson-Schwinger approach, talk given at the 9th International Workshop on Meson Production, Properties and Interaction (MESON2006), Krakow, Poland, June 13, 2006.

Earth System Modeling Framework (ESMF)

PIs: Cecelia Deluca and Robert Jacob, Mathematics and Computer Science

The Earth System Modeling Framework (ESMF) consists of software infrastructure for constructing and combining high-performance model components in the Earth science domain. It is in production use at the National Aeronautics and Space Administration's Global Modeling and Assimilation Office and the National Oceanic and Atmospheric Administration's National Centers for Environmental Prediction (NCEP) Global Forecasting System (GFS). In addition, it is being evaluated and adopted by many other groups across the country, including the Navy, Air Force, and Army, the Community Climate System Model (CCSM), GFDL, MIT, UCLA, and the Weather Research and Forecast (WRF) model. The ESMF hierarchical, component-based architecture facilitates the systematic construction of complex climate and weather applications and the interoperability of model components. ESMF also offers application developers an extensive toolkit for standard modeling functions such as grid transformations, data communication, logging, and calendar management.




Figure 14. As a key partner in the Battlespace Environments Institute, the ESMF team has redesigned its data structures to accommodate the multi-patch, curvilinear grids, and unstructured grids found in ocean, weather, space weather, and hydrological models.

In FY06, the ESMF project focused on transitioning prototype components and applications into production. ESMF codes transitioned to operations at NCEP and progressed toward operations at several U.S. Department of Defense sites.

The ESMF team used Jazz to run exhaustive nightly unit and system tests and examples, with a variety of the Jazz compilers. These tests were instrumental in ensuring that the ESMF software was robust and portable. The ESMF team also used Jazz as a development platform.

Widespread use of ESMF represents a paradigm shift in the way weather and climate models are constructed. Through increased code interoperability, community building, and standard tools, ESMF is beginning to make model development easier and facilitate new, multi-agency science collaborations (Figure 14). The end result is an Earth science community better equipped to explore basic research issues and answer questions about the impacts of Earth science on society.

We gratefully acknowledge the grants that were provided for the NOPP Open Source Community Sediment Transport Model and Extensions to Coupled COAMPS/NCOM, NRL PMW-180.

Publications/Presentations

1. T. Killeen, C. DeLuca, T. Gombosi, G. Toth, Q. Stout, C. Goodrich, A. Sussman, and M. Hesse, Integrated frameworks for earth and space weather simulation, preprint presented at the American Meteorological Society Meeting, Atlanta, Ga., January 29–February 2, 2006.

Fast Code Development for Large Computed Tomography Data Sets (FASTCTDEV)

PIs: Chris Deemer and Eugene Koehl, Nuclear Engineering

We have used JAZZ resources to provide 3-D tomographic images of projection x-ray images. The technique was adapted from the medical field and applied to engineering materials, such as fiber-reinforced ceramics, carbon and ceramic composite materials, and metal and ceramic components for use in heavy-vehicle propulsion systems and structural aerospace components. In many cases, the Jazz resources provide “near-real-time” visualization of these components. At other times, Jazz reduces to hours or days what would have required weeks to months to compute on stand-alone, single, and multiprocessor workstations.

The FastCTdev project’s goal is to render and display the details of a 10-inch, 3-D volume data set with 150–175 micron resolution from a series of 2-D x-ray projections within 30 minutes of acquisition. The project supports the development of nondestructive and analytical techniques employed in the examination of rare archeological and hard-to-obtain geological samples and one-of-a-kind prototype engineering designs.

The Jazz system was used to display data from a prototype aluminum casting process, Eurasian archeological and American geological samples, National Aeronautics and Space Administration shuttle components for the “Back to Flight” program, military body and vehicle armor, turbine blades and rotors and combustor liners for stationary and mobile power systems, and prototype fuel cells.

Because of the prototypic nature of the samples tested, the images and details resulting from the analysis are proprietary and cannot be released for publication at this time.

Global Climate Modeling with the Fast Ocean Atmosphere Model (FOAM)

PIs: Robert Jacob and Yannick Donnadieu, Mathematics and Computer Science

We are studying how the atmosphere, ocean, cryosphere, and land mass distribution interact to produce past and present climates. We use a coupled climate model called the Fast Ocean Atmosphere Model (FOAM). Its main features are a low-resolution model of the global atmosphere and an efficient model of the global ocean. We are exploring a wide range of climate problems with FOAM on Jazz.

This year, GEOCLIM, a coupled numerical model of the climate and global biogeochemical cycles, has been developed to investigate the consequences of the Pangea break-up on the atmospheric CO₂ content between 255 and 65 million years ago (Myr). The climate module of the GEOCLIM model is the FOAM atmospheric general circulation model, allowing the calculation of the consumption of atmospheric CO₂ through continental silicate weathering with a spatial resolution of $7.5^\circ \text{ long} \times 4.5^\circ \text{ lat}$. Thanks to the Jazz computing facilities, a large number of climate simulations has been performed. Eight time slices have been simulated from the Late Permian (260 Myr) to the Late Cretaceous (68 Myr). We have shown that the break-up of the Pangea supercontinent triggers an increase in continental runoff, resulting in enhanced atmospheric CO₂ consumption through silicate weathering. As a result, atmospheric CO₂ falls from values above 3,000 ppmv during the Triassic, down to rather low levels during the Cretaceous (around 400 ppmv), resulting in a decrease in global mean annual continental temperatures from about 20°C to 10°C (Figure 15).

We have also studied the behavior of the climate system during the Cretaceous (140-65 Myr). Using climate simulations performed with Jazz, we find that changes in geography from the Early to Mid-to-Late Cretaceous cause a large decrease of the seasonal cycle. First-order identified processes are the decreased continentality of the mid-to-high latitudes from the Mid-Cretaceous and the increase of the latent heat transport into the winter hemisphere, which induce a wetter and a cloudier atmosphere capable of diminishing the winter cooling of the continents.

The last, and maybe the most outstanding result we obtained with the climate-carbon model GEOCLIM, concerns the nature of the links between environmental changes and evolutionary innovation through time. In a paper to be submitted to *Nature*, we focus on the geological processes driving the atmospheric CO₂ levels and raise the question of causal links between the appearance of coccolithophores and the northward drift of Pangea. These calculations provide new insight into the relationship between continent configuration, weathering rates, atmospheric CO₂, and biodiversity.

New developments in FY06 include the ability of the atmosphere component in FOAM to read initial data in NetCDF format and write history files in NetCDF format. This eliminates the need for netcdf conversion tools when running FOAM. We removed all use of “system” calls, which enhances portability. We added the ability to do an arbitrary initial condition start in the land and sea ice model, which eliminates the need for a separate file during initial runs. Refactor land and river model initialization for clarity. Further, we began development of FOAM 2.0, including

switching from Fortran77 to Fortran90 compilers on all platforms and compiling with double precision.




Figure 15. The tectonic forcing on the evolution of atmospheric CO₂ level and biodiversity is shown in this figure. The thick black line represents the coccolithophore diversity. The red line represents total nannofossil diversity. Data represents species richness for a three million-year interval, plotted at the mid-point. The dashed grey line represents atmospheric CO₂ from a simpler GEOCARB III model. The four black points represent the steady-state atmospheric CO₂ levels achieved with the GEOCLIM model. Vertical bars on the four data points denote the upper and lower range of atmospheric CO₂ levels calculated, using a 20% increase and a 20% decrease in degassing flux, respectively. Horizontal bars represent duration of the geological stages from which paleogeographies have been derived. The light shadings represent the distribution of possible glacially derived sediments and glendonites through the Mesozoic.

Publications/Presentations

1. Y. Donnadieu, R. Pierrehumbert, R. Jacob, and F. Fluteau, Modelling the primary control of paleogeography on Cretaceous climate, *Earth and Planetary Science Letters*, 248, 426-437, 2006.
2. Y. Donnadieu, Y. Godd ris, R. Pierrehumbert, G. Dromart, R. Jacob, and F. Fluteau, Pangea break up and Mesozoic climatic evolution simulated by the GEOCLIM model, *G-cubed*, in press, 2006.
3. Y. Godd ris, Y. Donnadieu, and R. Pierrehumbert, Evolution of marine biocalcifiers linked to atmospheric CO₂ decline in the Late Triassic, to be submitted to *Nature*, 2006.
4. Y. Donnadieu, R. Pierrehumbert, and R. Jacob, Modelling the primary control of paleogeography on Cretaceous climate, International Symposium on the Cretaceous, Neuch tel, Switzerland, 2005.
5. Y. Donnadieu, R. Pierrehumbert, and R. Jacob, Modelling the primary control of paleogeography on Cretaceous climate, European Geosciences Union, Vienna, Austria, 2006.

6. Y. Donnadieu, Y. Godd ris, R. Pierrehumbert, G. Dromart, R. Jacob, and F. Fluteau, Pangea break up and Mesozoic climatic evolution simulated by the GEOCLIM model, European Geosciences Union, Vienna, Austria, 2006.

High-Throughput Genome Analysis

PIs: Natalia Maltsev, Dinanath Sulakhe, and Alexis Rodriguez, Mathematics and Computer Science

The GADU project is part of a system (PUMA2) that addresses the problem of annotating protein sequences automatically via protein comparisons with existing, well-annotated protein sequences from a list of several organisms from different kingdoms of life. GADU is responsible for making sequence homology computations using several tools. Millions of protein sequences need to be compared and analyzed through homology comparison tools. Computing sequence comparison homology will give scientists in the field a fast and educated method of annotating the function of a novel protein.

GADU system implementation has helped groups such as Gryphin and Open Science Grid (OSG) in developing new methods of automatic control of job submission to different Grid clusters (i.e., OSG, TeraGrid). Currently, the GADU project is one of a few systems that can submit effortlessly to different Grid clusters (Figure 16). Jazz has been extremely useful in job submission control, since it gives us access to a test bed to control what is going on at the node levels.

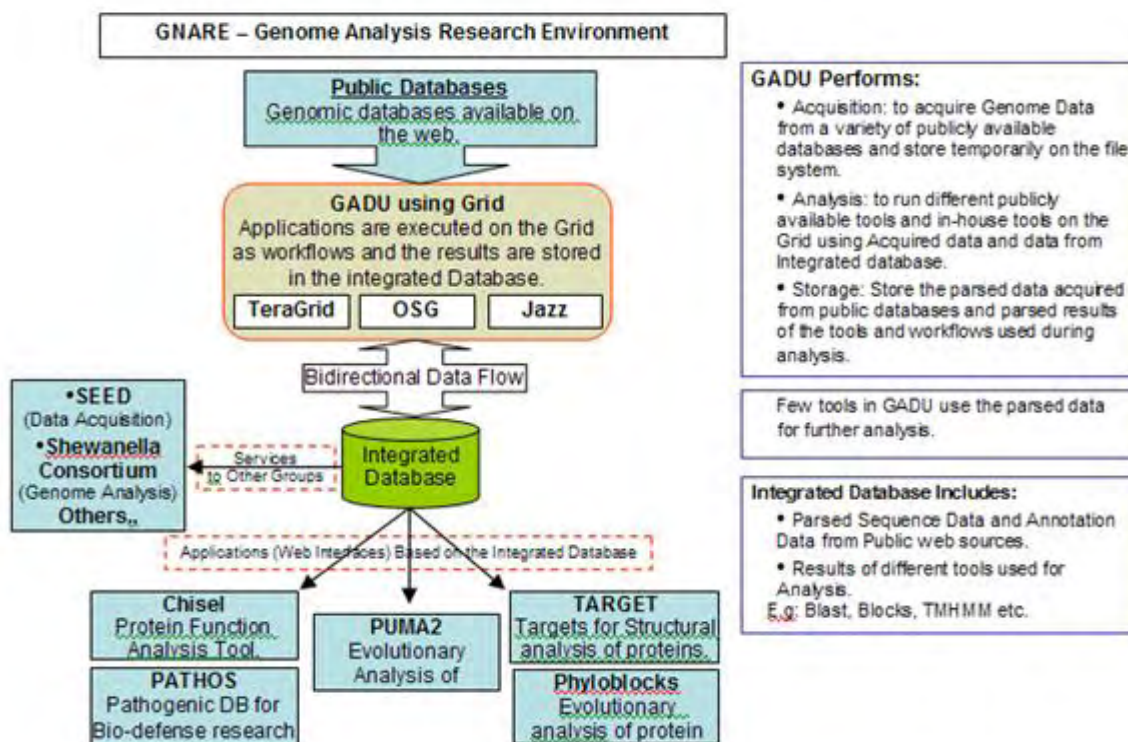


Figure 16. The GADU project is one of a few systems that can submit effortlessly to different Grid clusters.

Besides the Grid technology being addressed, we are also obtaining important results regarding the functions of new protein sequences. The results can be viewed in PUMA2. Computing time on Jazz helped us in analyzing rapidly growing genomic sequence data through various

bioinformatics tools such as BLAST, BLOCKS, and TMHMM. The amount of time it takes to run these tools has been significantly reduced because of the use of Jazz resources and other Grid resources (Table 1). We have also developed the CHISEL environment, which aids in evolutionary analysis of enzymatic proteins by creating clusters of proteins with similar taxonomic and functional characteristics and domain features.

New software capabilities accomplished in FY2006 are the addition of an automated job submission control. The scalability factors have also helped in analyzing more protein sequences and, at the same time, created protein clusters based on their function, taxonomy, and protein domains. We converted the execution of tools in the form workflows and automatically execute the workflows on high-performance clusters using Condor-G.

Table 1. Performance of Blast Using Jazz Resources

<p>One CPU (time in wall time): 100 sequences took 66 min on one CPU. One genome (~4000 seqs would take 2,640 min or 44 hr). 2314886 sequences takes ~1527825 min (i.e., 25463.746 hr, or 1,061 days).</p>
<p>200 CPUs: 2314886 sequences took 12,480 min (i.e., 208 hr or 8 days, 16 hr). So one genome (~4000 seqs) should take 22 min using 200 CPUs.</p>

Publications/Presentations

1. D. Sulakhe, A. Rodriguez, M. D'Souza, M. Wilde, V. Nefedova, I. Foster, and N. Maltsev, Gnare: Automated system for high-throughput genome analysis with grid computational backend, *Journal of Clinical Monitoring and Computing*, 19, 361–369, 2005.
2. D. Sulakhe, A. Rodriguez, M. D'Souza, M. Wilde, V. Nefedova, I. Foster, and N. Maltsev, Gnare: An environment for grid-based high-throughput genome analysis, *Proceedings of CCGRID*, IEEE, 2005.
3. N. Maltsev and D. Sulakhe, Gnare: An environment for grid-based high-throughput genome analysis, BioGrid, 2005.
4. M. D'Souza and D. Sulakhe, PUMA2: An interactive integrated environment for high-throughput genetic sequence analysis and metabolic reconstructions with grid-based computational backend, ISMB, 2005.

Evaluation of Commercial CFD Code Capabilities for Prediction of Aerodynamic Drag for Heavy Vehicles

PI: W. D. Pointer, Nuclear Engineering

This project is a component of a multilaboratory effort to develop a computational prediction capability for the heavy-vehicle industry to aid in improving fuel economy in Class 3-8 trucks. The Argonne component will provide best practice guidelines for the use of current-generation computational tools and a baseline for comparison with more advanced methodologies. The study includes evaluations of scalability issues, mesh construction methodologies, and turbulence modeling options. The prediction of drag coefficients for heavy vehicles is difficult as a result of the multiple, large-flow separation regions inherent in most truck designs. These separation regions lead to the development of multiple regions in the flow field populated by multimode vortex shedding phenomena.

This year's effort extended the previous year's studies of the effects of cross-wind consideration on the accuracy of simulations, quantification of expected accuracy when evaluating changes in vehicle configuration, the capability to create more realistic vehicle designs from manufacturers' CAD data, and a first application of the guidelines that have been developed. Most studies utilize the simplified geometry of the Generalized Conventional Model (GCM), shown in Figure 17. Figure 18 shows a comparison of the predicated and measured value of the drag coefficient as a function of the vehicle yaw angle for the GCM vehicle. Studies completed in FY2006 have shown that error in the predictions at higher yaw angles can be reduced to less than 5% with only a relatively modest increase in computational investment. The predictions indicate that the device reduces the drag coefficient by approximately 7%, which is consistent with both wind tunnel and road testing of a prototype device by project partner Aerovolution, Inc. A drag coefficient reduction of 7% would result in fuel savings of approximately 3.5% at highway speeds.

These calculations have provided an independent assessment supporting the applicability of commercial CFD software to heavy-vehicle aerodynamic evaluations. The study has provided a baseline simulation using a large number of computational elements, which requires a large computational facility with hundreds of available processors to complete and investigate the impact of reducing the resolution and size of the computational investment on the accuracy of the results and the problem turnaround time. As a result of these studies, in tandem with the assessments completed by the software vendors, the manufacturers have begun investing in small clusters and commercial CFD as part of their design process. Ongoing studies are supporting the development of deployable drag-reduction devices for use in the existing fleet.

Publications/Presentations

1. W. D. Pointer, T. Sofu, and D. Weber, Development of guidelines for the use of commercial CFD in tractor-trailer aerodynamic design, SAE Commercial Vehicle Exposition, 05CV-120, November 2005.
2. W. D. Pointer and S. Singh, Heavy-vehicle drag estimation using commercial CFD tools, Heavy-Vehicle Systems Optimization Program Review, April 18, 2006.




Figure 17. Particle traces and translucent contour planes showing distribution of velocity magnitude in the flow field surrounding the Generalized Conventional Model (GCM).

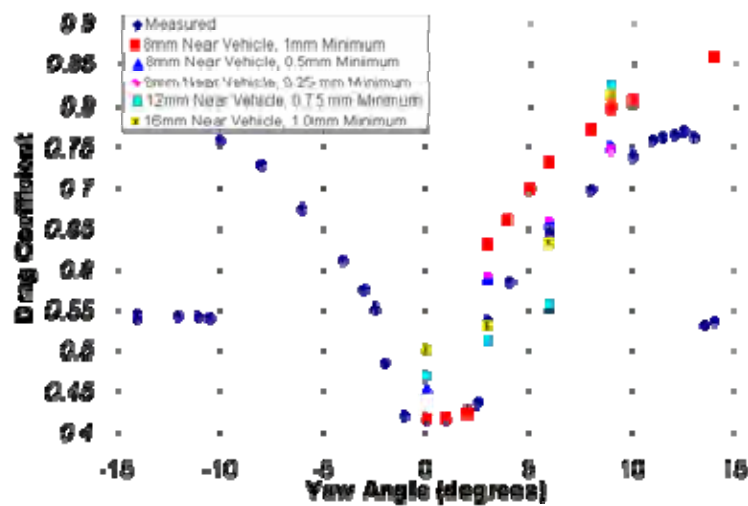


Figure 18. Comparison of predicted and measured drag coefficients versus yaw angle, which is the apparent angle perceived by the vehicle as a consequence of cross-winds.

Advanced 3-D Laser-Produced Plasma Hydrodynamics and Radiation Transport for EUV Lithography Applications

PIs: Tatyana Sizyuk and Vitali Morozov, Mathematics and Computer Science

Extreme ultraviolet (EUV) lithography at wavelength near 13.5 nm is considered by many as the strongest candidate to succeed the perspective lithography for large-scale manufacturing of integrated circuits. However, the EUV radiation source has been identified as the largest obstacle to its successful implementation, because no source technology currently meets all the requirements for a production tool. Several concepts are now under intense development, which include electrically driven plasmas (DPP) as well as laser-produced plasma (LPP) concepts. Despite noticeable progress in development during recent years, the main limitation of the current designs is insufficient total in-band ($13.5 \pm 1\%$ nm) power, and correspondingly, the low conversion efficiency of the sources. At the moment, theoretical investigation and experimental work are directed in a way to find the most optimal parameters for the proposed devices, which would allow us to maximize the efficiency of the EUV source.

The HEIGHTS team has been working intensively on the development of a comprehensive 2-D and 3-D computer simulation package for modeling the properties and radiation output of both DPP and LPP EUV lithography sources. The purpose of the development is to provide leading source manufacturers with a powerful, flexible, and extensive research tool to optimize operational regimes for their devices, and/or suggest alternatives with higher stability and efficiency. The developed models address the following subjects: plasma evolution and magnetohydrodynamic (MHD) processes of compressible gas/plasma, heat conduction, atomic data and plasma properties, detailed photon radiation transport, interaction between plasma/radiation and material, laser absorption in LPP, and external electric circuit in DPP devices. Regions with differing propagation speeds of perturbation require accurate numerical solutions of the MHD equations. The total variation diminishing scheme in the Lax-Friedrich formulation for the description of magnetic compression and diffusion is worked out and used in our model. Parallel implementation of the 3-D HEIGHTS computer simulation package is based on the message-passing interface (MPI) and each block of the package is implemented independently.

This year, the HEIGHTS team mainly modeled various LPP devices. The highest importance of the device is its conversion efficiency (CE) of laser energy to extreme ultraviolet radiation. We have made a parametric study of a typical LPP device and found the dependence of the CE on a host of parameters (such as laser intensity, wavelength, target geometry, duration of laser pulse, etc.).

The results of calculated CE from laser wavelength and laser beam intensity (tin droplet target with diameter 100 μm) are shown in Figure 19.

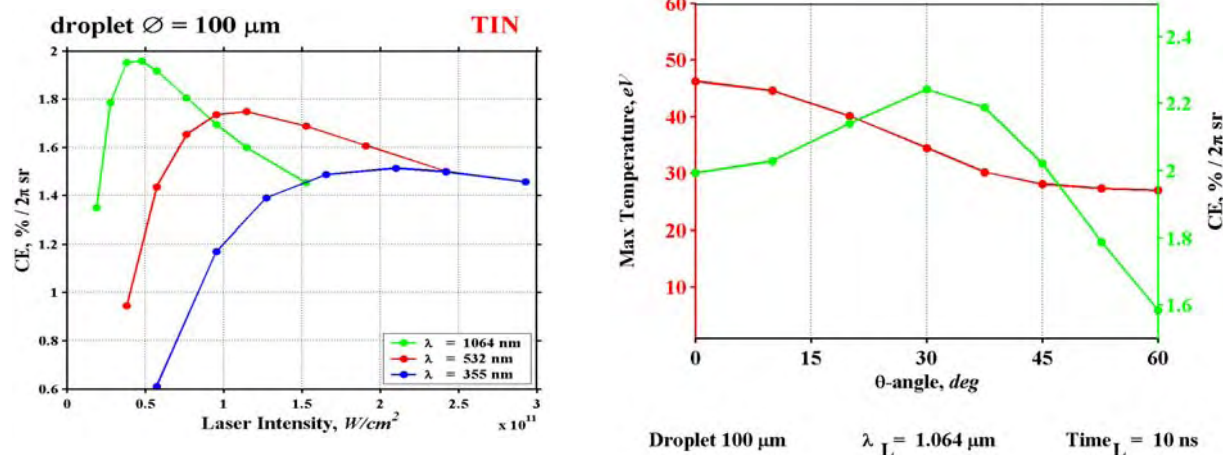


Figure 19. Conversion efficiency vs. laser intensity for tin droplet target (left) and temperature dependence vs. position of the laser (right).

The use of different laser pulses, multiple incident lasers with various directions, and focusing points allows noticeably improved EUV radiation efficiency. The constructions of multiple incident lasers may help to warm up the target more uniformly, which enables enlargement of the EUV emitting plasma area. The results of the HEIGHTS computer simulation of a tin droplet warming up by three lasers are presented in Figure 19 (right). The positions of the lasers were changed, and, as a result, the total laser energy intensity on droplet surface and heating area was varied.

Publications/Presentations

1. V. Sizyuk, A. Hassanein, V. Morozov, V. Tolkach, and T. Sizyuk, Numerical simulation of laser-produced plasma devices for EUV lithography using the HEIGHTS integrated model, *Numerical Heat Transfer, Part A*, 49, N3, 215-236, 2006.
2. G. V. Miloshevsky, V. Sizyuk, M. B. Partenskii, A. Hassanein, and P. C. Jordan, Application of finite-difference methods to membrane-mediated protein interactions and to heat and magnetic field diffusion in plasmas, *Journal of Computational Physics*, 212, 25-51, 2006.
3. A. Hassanein, V. Sizyuk, V. Morozov, and T. Sizyuk, Three-dimensional simulation and optimization of laser-produced plasma for EUV lithography devices, Emerging Lithographic Technologies X Conference, San Jose, Calif., February 21-23, 2006.

Physics of the High-Temperature Gas-Cooled Reactor

PI: Taek K. Kim, Nuclear Engineering

The Global Nuclear Energy Partnership (GNEP) has been launched to develop and demonstrate new proliferation-resistant technologies to recycle nuclear spent fuel and reduce waste. The burning of recovered transuranium from light water reactor (LWR) spent nuclear fuel in the advanced burner reactor (ABR) has been considered to minimize nuclear waste. The utilization of the deep-burn concept, based on the graphite-moderated gas-cooled reactor (DB-MHR) in the GNEP fuel cycle, has been evaluated during FY2006.

To provide confidence in the DB-MHR core design, validation and verification (V&V) was required for the computational tools used at Argonne. As one of the V&V efforts, a benchmark calculation was performed using the High-Temperature Test Reactor (HTTR). To obtain the reference solution of the HTTR benchmark problem, the MCNP model (Figure 20) was developed, and the results were obtained by using the Jazz machine. In addition, in order to assess the Argonne suite of fast reactor analysis codes to design ABR core, a numerical benchmark problem was developed, based on the reference Advanced Burner Test Reactor (ABTR) metal core design with weapons-grade plutonium fuels. Like the DB-MHR case, the reference solution was obtained from the MCNP calculations by using the Jazz machine.




Figure 20. MCNP model of HTTR benchmark problem.

***Ab Initio* Investigation of CO Hydrogenation Mechanisms with Co Carbonyl Catalysts**

PIs: Randall Meyer, Department of Chemical Engineering, University of Illinois at Chicago, and Jerry Rathke and Robert Klingler, Chemical Engineering, Argonne National Laboratory

The rapid depletion of oil necessitates the discovery and development of new ways to make other potential energy sources viable. One such process is the Fischer-Tropsch catalytic process to take syn-gas, a mixture of carbon monoxide (CO) and hydrogen produced from coal, to make higher molecular weight hydrocarbons of greater value such as those used for diesel fuels.

Unfortunately, Fischer-Tropsch catalysis is a complex process whose mechanisms and surface intermediates are not well understood. Therefore, as a first step to developing improved Fischer-Tropsch catalysts with a high selectivity to particular products, we have chosen to examine a homogeneous analog, $\text{HCo}(\text{CO})_4$. If we can precisely identify the reaction pathways for the homogeneous catalyst and understand what factors influence the catalyst's selectivity, then we can ultimately use this knowledge to aid in the design of new heterogeneous catalysts.

We have used the VASP (Vienna *Ab Initio* Simulation Package) density functional theory code to investigate the hydrogenation of CO with Co-based homogeneous catalysts. Previously, Argonne researchers had generated a vast amount of experimental data (primarily using NMR), allowing for identification of the structures of stable intermediate species, which leads the way toward an improved understanding of the reaction mechanism. However, many unanswered questions remain, including the precise location of the hydrogens in the $\text{Co}(\text{CO})_4\text{H}_3$ transition-state and the influence of ligands on the catalysis. The primary aim of our work is to test various hypotheses for the mechanisms in an effort to understand how to improve selectivity of these catalysts. Calculations utilizing Jazz allow for rapid and efficient computation of both the thermodynamics and kinetics of all relevant reaction steps in this system.

We have calculated the energies of all stable intermediates in the reaction cycle, allowing for the creation of a potential energy surface as depicted in Figure 21. Two different product pathways are possible, based on the insertion of CO, leading to either ethylene glycol and methanol or methyl formate and methanol. Previous calculations in the literature have indicated much higher energy paths for these reactions than have been observed experimentally. However, our own examination of this system (based on an improved approach through the application of superior computational power) has revealed significantly better agreement with the experiment, as shown in Table 2 for several key reaction steps.




Figure 21. A view of the potential energy surface for CO hydrogenation. The black line indicates the energy of intermediates leading to methyl formate and methanol production. The red line indicates an alternative path leading to the formation of methanol and ethylene glycol.

Table 2. Comparison of Several Key Reaction Steps

Reaction (energies expressed in kcal/mol)	Ziegler et al.	Bursten et al.	Current Work (VASP)	Experimental
Co-Co dissociation	35.4	29.9	21.9	19.7
Decarbonylation	45	26	32.9	26
Co ₂ (CO) ₈ hydrogenation	6.5		0.1	4.7
CO hydrogenation			37	40.7

Since our project has been active for only five months, we have only begun to understand the reaction cycle for CO hydrogenation for the production of methanol. However, our results achieve a previously unknown level of agreement with experiment for this system. Furthermore, we have identified the transition state in the reaction cycle of methanol production: (CO)₃H₂CoCOH. The results also indicate that (CO)₄CoCOH should be an observable intermediate, giving guidance to new NMR experiments to be performed by Klingler and Rathke. We can also use our calculations to explain product distributions based on the relative reaction enthalpies. New results will more thoroughly examine the influence of ligands such as triethylphosphine (replacing carbonyl groups) upon the product selectivity. The next phase of the project will focus on the production of ethanol by the homologation of methanol, another homogenous catalytic process of great importance given the recent increase in demand for ethanol-based fuels.

Presentations/Publications

1. R. Meyer, B. Klingler, and J. Rathke, A combined experimental and theoretical study of the mechanism of CO Hydrogenation with Co carbonyl catalysts, being submitted to the *Journal of the American Chemical Society*. The intent is to present this work for the first time at the American Chemical Society Meeting in Chicago, Ill. in March 2007.

Dynamics of Voltage-Gated K⁺ Channel

PIs: Benoit Roux and Yuqing Deng, Biosciences

Our project addresses questions about voltage-gated potassium channels (kv). These channels open to conduct potassium ions under the influence of the transmembrane voltage. So far, we have used our time with Jazz to carry a 40 ns molecular dynamics simulation of a detailed model of the Kv1.2 channel with explicit solvent and membrane lipid (around 75,000 atoms). With this simulation, we hope to obtain the spatial variations of the voltage in three-dimensional space. Figure 22 shows the potential obtained by using a continuum electrostatic (Poisson-Boltzmann) approximation. The spatial variation of the potential is key for understanding the mechanism of gating of these kv channels.

The entire simulation was carried out on Jazz using a MPI run of CHARMM in parallel. The code was ready to use without extensive modifications. Without ready access to Jazz, this project would have been delayed significantly.

These computations should clarify the solvated environment of the atomic charges that play a key role in channel gating and the transmembrane voltage.




Figure 22. The Kv1.2 channel embedded in a lipid bilayer is represented. The levels of the transmembrane voltage, calculated using a Poisson-Boltzmann equation, are shown. Each color represents about 15% of the total transmembrane potential.

Lattice Quantum-Chromodynamics (Lattice-QCD)

PI: Donald Sinclair, High Energy Physics

We are studying hadronic/nuclear matter at finite temperature and density. In particular, we are interested in the finite temperature transition to a quark-gluon plasma. Our studies use quantum-chromodynamics (QCD), the theory that describes the strong interactions. These studies are relevant to the physics of the early universe, neutron stars, and relativistic heavy-ion collisions at CERN-LHC, BNL-RHIC, GSI, CERN-SPS and BNL-AGS.

For these studies, QCD is defined on a discrete space-time lattice. The behavior of this quantum theory is mapped onto that of a classical theory evolving in a fictitious “time.” Periodic interactions with a heat-bath assure ergodicity.

Finite density is achieved by using a chemical potential. We use a finite isospin chemical potential, which avoids the sign problem associated with a baryon number chemical potential. At the small chemical potentials at which we work, the physics associated with these two different chemical potentials is very similar. The most striking feature of the phase diagram in the region accessible to heavy-ion colliders is the critical endpoint at which the transition becomes a first-order phase transition rather than a smooth crossover. It is this feature that we are seeking.

Last year, we had determined that the binder cumulants that we use to determine the nature of the transition from hadronic/nuclear matter to a quark-gluon plasma were very sensitive to discretization errors in our numerical integration of the classical equations of motion (Figure 23). This year, we replaced our simulation algorithm with the new RHMC algorithm, in which a global accept/reject step compensates for these discretization errors, making it an exact algorithm. These simulations were run on Jazz and other high-performance computers. Although intrinsically slower than these other machines, Jazz has better turnaround, which helped in maintaining steady progress toward our goals. The conclusion from these runs with the RHMC algorithm for 3-flavour lattice QCD is that there is no sign of a critical endpoint in the small-chemical-potential region in which we performed our simulations. This shows that the association of this critical endpoint with the critical mass at zero chemical potential, which some have suggested, is most probably incorrect. Use of the RHMC algorithm enabled us to show that this critical mass is 0.0264(5), rather than the published 0.033(1). Others have confirmed this.

As indicated above, these computations have established the RHMC as the algorithm of choice for such simulations. They have led us to implement the RHMC for another of our projects. Our




Figure 23. Binder cumulant at the transition from hadronic matter to a quark-gluon plasma as a function of isospin chemical potential at quark mass $m=0.03$ in lattice units.

calculations have caused other groups to repeat their simulations, also using the RHMC algorithm. All claims to have observed a critical endpoint in lattice QCD calculations have been called into doubt, and new directions will need to be taken to answer the question as to whether this endpoint exists and, if so, where.

The main new software capability developed in FY2006 involved the coding of the RHMC algorithm to replace the older HMD(R) algorithm. Even though the RHMC algorithm is *a priori* a little more expensive than the HMD(R) method, the fact that one no longer needs to vary the updating increment to remove discretization errors from observables actually makes it less expensive in the long run.

We gratefully acknowledge the ERCAP grant that was provided for the allocation at NERSC and the NRAC grant that was provided for allocations at the NSF supercomputing centers.

Publications/Presentations

1. J. B. Kogut and D. K. Sinclair, The RHMC algorithm for theories with unknown spectral bounds, [arXiv:hep-lat/0608017].
2. J. B. Kogut and D. K. Sinclair, 3-flavour lattice QCD at finite density and temperature: QCD at finite isospin density revisited, [arXiv:hep-lat/0509095].
3. J. B. Kogut and D. K. Sinclair, Finite dt dependence of the binder cumulants for 3-flavor QCD at finite temperature and isospin density, [arXiv:hep-lat/0504003].
4. J. B. Kogut and D. K. Sinclair, The finite temperature transition for 3-flavour lattice QCD at finite isospin density, *Nucl. Phys. Proc. Suppl.* 140, 526 [arXiv:hep-lat/0407041], 2005.
5. D. K. Sinclair, presentation at breakout session, HEP Division DOE Review, March 2006.
6. J. B. Kogut and D. K. Sinclair, Searching for the elusive critical endpoint at finite isospin density and temperature, presentation by D.K.Sinclair at Lattice 2006, Tucson, AZ, July 23-28, 2006.
7. J. B. Kogut and D. K. Sinclair, Lattice QCD at finite temperature and/or densities, talk presented by D.K.Sinclair at QCD in Extreme Conditions, Brookhaven National Laboratory, July 31-August 2, 2006.

High-Fidelity Model Simulations in Urban Boundary Layer Applications: Air Pollution Meteorology and Hazard Cloud Dynamics (MANHATTAN)

PI: Michael Lazaro, Environmental Science

In collaboration with EPA scientists, we are exploring the possibility of developing or adapting the products from computational fluid dynamics (CFD) simulations to support rapid exposure and risk models to guide urban emergency response and emergency management for chemical, biological, or radiological attacks and accidents. The first phase of our study involved application of Fluent code to lower Manhattan in New York City to evaluate the transport of aerosols generated from the World Trade Center collapse and aid in the study of toxicological effects from exposures to fine aerosols. Experimental data from a scale model (from the USEPA wind tunnel) of lower Manhattan was used to support the Fluent simulations. The second phase focused Manhattan simulations to support the New York City portion of the Urban Dispersion Program led by the U.S. Department of Homeland Security.



Figure 24. Simulation of tracer releases from several locations around Madison Square Garden and Penn Plaza.

Our objective is to explore, extend, and establish best CFD practices to open-scale outdoor urban environment applications. Our ultimate goal is to produce and test an advanced, computationally efficient, high-fidelity modeling system with integrated data assimilation that will have applications in reducing the inherent large uncertainties in existing models used to predict hazard cloud movement in complex urban settings. This system, when fully implemented, will provide analysts with enhanced capability to advise incident commanders and postresponse officials on critical decisions involving population evacuation/sheltering, contamination mapping/cleanup, and reentry/reuse. It will provide the needed special capability in highly vulnerable and densely populated urban centers where critical infrastructure protection and counterterrorism are of importance to our national homeland security goals. In addition, this project, with its focus on advanced simulation capabilities, is consistent with Argonne's strategic commitment to high-performance computing.

In FY2006, we used Fluent on the Jazz cluster to simulate tracer releases from several locations around Madison Square Garden and Penn Plaza. These simulations provided a unique understanding of complex urban canyon flow influences. Small changes in relative location of releases with respect to building groups illustrated that release positions can have as important an influence on downwind dispersion patterns as changes in boundary layer or free stream winds. The release locations in Figure 25 are indicated in Figure 24 by the green and blue circles. The relative significance of this influence is shown in Figure 25. These and other lower Manhattan and Midtown simulations were also important in developing and testing CFD data preparation and setup methods and simulation techniques unique to open-scale outdoor environment

applications. During 2007, the methods that have been developed will be applied to support production runs needed for comparison with field program measurements.




Figure 25. Plume 1: Release at south side-center of MSG Building (green rel. pt). Plume 2: Release at north side-center of Penn 1 building (blue rel. pt.).

Although this project is using a commercial code, best practice and special application methods that are being developed can be extended to future developments using open-source CFD codes. New capabilities that will provide the ability to piece together solutions to format a larger domain will be tested and applied in nested mega-domains during FY2007. Methods for interface CFD simulations with larger grid-scale routine meteorological models will be developed and tested this year.

Monte Carlo Analysis of Accelerator-Driven Subcritical Systems

PI: Yousry Gohar, Nuclear Engineering

The governing physics and characteristics of accelerator-driven subcritical assembly facilities have been studied. Target materials with high atomic numbers are used for generating neutrons. Several studies have been carried out to define the physics of the subcritical assembly facilities. The neutron source intensity, the neutron spectrum, the spatial neutron generation, and the spatial energy deposition in the target material have been examined as a function of the beam parameters and target material selection. The neutron flux of the subcritical assembly has been examined as a function of the uranium fuel enrichment, the uranium density, the reflector material and thickness, and the target material selection for k_{eff} of ~ 0.98 .

The spatial energy deposition density per incident electron is defined as a function of the target length for different electron energies. The peak value occurs a few millimeters away from the electron beam window. This peak decreases and shifts further from the electron beam window as the electron energy increases. The results show that a beam power density of $\sim 2 \text{ KW/cm}^2$ and electron energy $\geq 100 \text{ MeV}$ satisfy the different requirements. Reducing the beam power density increases the target cross-section area, which reduces the neutron flux in the subcritical assembly. Decreasing the electron energy reduces the required target length and increases the peak energy deposition, which require very thin layers of the target material and reduces the neutron yield.

The spatial distribution of the generated neutrons was examined. The results show that the use of high-electron energy reduces the neutron fraction that leaves from the beam tube in the beam direction. It should be noted that the neutron current from the beam window is larger than the neutron current from the other end of the target, which calls for locating the target midpoint under the subcritical assembly mid-plane. This placement increases the chance for the neutrons leaving the beam window to reach the subcritical assembly. The use of high-energy electrons in the range of 100 to 200 MeV is beneficial, and increasing the electron energy in this range enhances the achieved performance. The neutron spectrum from the target was analyzed to quantify the high-energy component as a function of the electron energy. The calculated neutron spectra show that the peak value of the neutron spectrum is about 1 MeV. The high-energy component of the spectrum is not sensitive to the electron energy, and its magnitude is very small. Also, the neutron yield from the target was analyzed as a function of the target length. A 0.08 m target length is adequate for the uranium target with 200 MeV electrons.

The target material has a cylindrical geometry, and its axis coincides with the beam tube axis. The water coolant flow direction is perpendicular to the target axis. This arrangement results in a stack of disks forming the target configuration. A clad material is utilized for the uranium disks to avoid water coolant contamination with fission products. The power density distribution in the uranium target assembly is shown in Figure 26, which is calculated by using low-enriched uranium for the subcritical assembly with uranium density of 3 g/cm^3 and beryllium reflector. The neutron source strength is $3.3 \times 10^{14} \text{ n/s}$ from the 100 KW electron beam with 200 MeV electron energy.

The uranium enrichment and the reflector material selection are varied to study the impact on the neutron flux field. The targets disks are placed in an aluminum tube surrounded by hexagonal canisters that match the fuel geometry. The beam tube is located at the center of the subcritical assembly, and its axis coincides with the centerline of the subcritical assembly. Three-dimensional models of the subcritical assembly, including the detailed target configuration, were developed to define the subcriticality level and the neutron flux distribution. The target and fuel assemblies are modeled explicitly without any geometrical approximation or material homogenization to get an accurate prediction of the subcritical performance. MCNPX computer code with continuous energy data libraries and $S(\alpha, \beta)$ thermal data was used for the analyses. Analyses were performed to define the neutron flux field for K_{eff} of ~ 0.98 with different combinations of the reflector and target materials, fuel density, and fuel enrichment.

The calculated neutron spectra show that about half of the neutrons in the irradiation channels have energies below 100 keV, and the other half are in the energy range of 100 keV to 20 MeV. The neutron fraction with energy above 20 MeV is very small. The use of beryllium reflector produces a higher flux level relative to the water. Figure 27 shows the neutron flux map averaged over the active fuel length of the subcritical assembly and the corresponding power deposition distribution from the 100 KW beam.






Figure 26. Uranium target power density distribution.



Neutron Flux Map



Power Density Map

Figure 27. Neutron flux and power density maps perpendicular to the beam tube averaged over the active fuel length of the subcritical assembly.

Molecular Dynamics Simulations for Solution X-Ray Diffraction (MD_for_SXD)

PIs: David Tiede and Xiaobing Zuo, Chemistry

In situ characterizations of the structure and dynamics of supramolecules in solution and solute-solvent interactions are critical for achieving a fundamental understanding of their functions at the molecular level. They are also crucial for achieving the predictive design of materials with applications ranging from solar energy conversion to next-generation pharmaceuticals.

Molecular dynamics (MD) simulations are widely used as the primary tool for characterizing the conformational landscapes of supramolecules in solution and for predicting DNA conformation, DNA conformational transitions, and DNA-drug interactions, as well as for investigating solvation interaction. Further, MD simulations and the associated force fields underlie structural determination using solution nuclear magnetic resonance (NMR) techniques. However, until now, there has been no direct, independent method for testing the accuracy of supramolecular structure and dynamics simulations in solution and other condensed phase noncrystalline media, or for testing the appropriateness of the underlying force fields, initial conditions, and simulation parameters.

In FY2006, we first finished a project started in FY2005 in which we developed methods for quantitatively testing the accuracy of MD simulations to represent conformational landscapes of DNA in solution. Fourier transforms are used to produce reciprocal space “fingerprints” of atomic pair distance correlations from MD simulation and to make direct, quantitative comparison to experimental solution X-ray diffraction (SXD) measurements. More simulations were carried out in FY2006 using AMBER 8 running on Jazz for a series of DNA sequences. The simulations tested the effects of DNA nucleotide base sequence, length, and range of simulation parameters on the resulting MD ensemble. In contrast to the experiment, MD simulations were found to be consistently skewed toward a single DNA conformer (the so-called B-form conformer with 10.0 base pairs per helix turn), although other conformers were found to be present to various extents as minority species within MD ensembles, more due to the defect of

current AMBER force fields than to simulation condition parameters. Reciprocal space fingerprints were found to be a useful means to efficiently compare simulation to experiment and for searching conformational ensembles to identify conformational subsets that provide a closer match to the experiment than the ensemble average. The results are summarized in Table 3.

Table 3. Molecular Dynamics Simulations						
Trial	Time (ns)	Initial Structure	Solvent Box (Å)	Cutoff (Å)	B'-form (%)	B'-form RMSD
d(A)₁₀						
A10.0	2.0	B	Implicit solvent		14.3	0.096
A10.1	10.8	B	8	8	21.8	0.076
A10.2	8.6	B	10	8	23.4	0.071
A10.3	8.2	B	10	12	46.2	0.069
A10.4	1.7	B	10	15	41.8	0.067
A10.5	6.0	B'	8	8	13.2	0.078
A10.6	4.8	B'	10	12	47.0	0.072
A10.7	4.0	B	30	12	34.8	0.067
d(A)₂₀						
A20.1	6.7	B	8	8	19.5	0.092
A20.2	4.8	B	10	12	49.6	0.071
d(AT)₁₀						
AT10.1	6.7	B	8	8	5.2	---
AT10.2	2.8	B	10	12	9.1	---

We also started a second project in FY2006 aiming to investigate the solvent-solute interaction by combining MD simulations and SXD measurements. Solvent-solute interactions are critical in many important chemistry processes and applications, for example, intra-/intermolecular electron transfer in solar energy conversion systems. Unfortunately, not many techniques can measure such solvation interaction. Solution X-ray diffraction does directly measure the solvent-solute interaction, but the data is often difficult to interpret. In this project, MD simulations provide a theoretical prediction on how solvent molecules pack around a solute molecule at atomic level, which will be used as structural models to understand SXD data and develop a solvent model in SXD theory. One simulation shown in Figure 28 is the relative water molecule density around a β -cyclodextrin obtained from AMBER-based MD simulations. More simulations and data analysis will be continued in FY2007.




Figure 28. Relative molecular density (in orange color) of water around β -cyclodextrin (in white-red-cyan) obtained from MD simulations.

Multimethod Linear Solvers in Terascale PDE-based Simulations (Multimethod_Solvers)

PIs: Boyana Norris and Lois McInnes Curfman, Mathematics and Computer Science

Many large-scale scientific simulations involve the parallel solution of time-dependent and/or nonlinear partial differential equations. Overall simulation time is often dominated by the parallel solution of large-scale, sparse linear systems. Typically, application developers select a particular algorithm to solve a given linear system and keep this algorithm fixed throughout the simulation. However, it is difficult to select *a priori* the most effective algorithm for a given application.

Thus, we are exploring *polyalgorithmic multimethod linear solvers* in the context of several parallel applications, including flow in a driven cavity, compressible Euler flow, and most recently, radiation transport, to potentially improve the execution time and reliability of the linear system solution. We are developing composite solvers, which provide linear solution by using a sequence of preconditioned iterative methods on a given system until convergence is achieved. This approach enhances reliability by overcoming the potential failure of a particular method. Moreover, for long-running applications in which the numerical properties of the linear systems change as the simulation progresses, a single algorithm may not be best throughout the entire simulation. This situation has motivated us also to develop an adaptive, polyalgorithmic approach, where the solution method is selected dynamically to match the attributes of the linear systems as they change during the course of a simulation.

Our work on Jazz during FY2006 focused on developing an adaptive, polyalgorithmic approach for parallel linear solvers in the context of radiation transport simulations, which govern the evolution of photon radiation in an optically thick medium. This application employs Newton-Krylov methods for a fully implicit solution at each time step. We used Jazz to help develop a technique for adaptively selecting the parallel linear solver methods to match the numeric properties of the linearized Newton systems as they evolve during the course of the nonlinear iterations. Our approach combines more robust, but more costly methods when needed in particularly challenging phases of the solution, with faster, though less powerful, methods in other phases.

We also used Jazz to explore machine-learning methods for solver selection. Machine-learning methods determine a set of suitable solvers by generating classifiers that map linear systems to solvers. The information on linear systems and solvers is stored in a database, where each entry consists of (1) the feature set of the matrix that defines the linear operator, (2) the parameters of the solver, and (3) labels representing the ranking of the solution of the matrix by the solver with respect to a performance parameter.

The learning algorithm builds the classifier by searching for a mapping that makes accurate predictions on the dataset. Once the classifier has been built, the feature set of a new matrix with similar characteristics to the matrices forming the database and the set of solvers, without labels, is provided to the classifier. The output then labels each solver in the set according to its

suitability for the new matrix. In our current implementation, we use alternating decision trees coupled with AdaBoost as the learning algorithm.

These computations extended our prior research on sequence-based adaptive heuristics in a single processor environment to solve a more complicated parallel application. We demonstrated by using Jazz that the adaptive polyalgorithmic approach can be easily parallelized, is scalable, and can lead to improvements in overall simulation time. This work provides the foundation for further investigations on parallel adaptive solvers in a variety of scientific applications, possibly including fusion and accelerator models.

Our research on multimethod solvers for large-scale, PDE-based simulations using Jazz has motivated us to develop software infrastructure for *computational quality-of-service*. In particular, we developed a component infrastructure that supports performance monitoring, analysis, and adaptation of important numerical kernels, such as nonlinear and linear system solvers (Figure 29). We defined a simple, flexible interface for the implementation of adaptive nonlinear and linear solver heuristics. We also provided components for monitoring, checkpointing, and gathering of performance data, which are managed through two types of databases. The first is created and destroyed during runtime and stores performance data for code segments of interest, as well as various application-specific performance events in the currently running application instance. The second database is persistent and contains performance data from various applications and different instances of the same application. This database can also contain performance information derived through offline analysis of raw data.

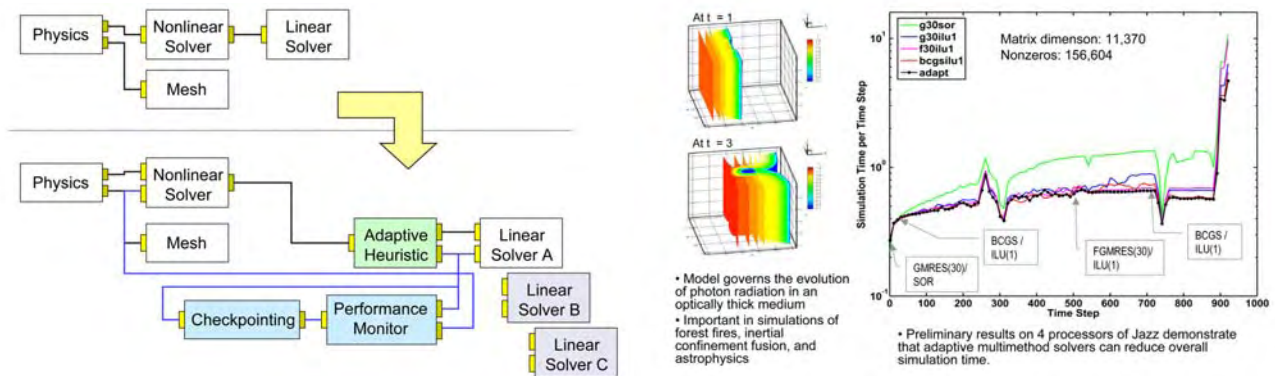


Figure 29. Multimethod linear solver components in nonlinear PDE solution (left).

Multimethod solvers: radiation transport (right).

Publications/Presentations

1. S. Bhowmick, D. Keyes, Y. Freund, V. Eijkhout, and E. Fuentes, Application of machine learning to selecting solvers for sparse linear systems, submitted to the *International Journal of High-Performance Computing Applications*, September 2006.

2. B. Norris, L. McInnes, and I. Veljkovic, Computational quality of service in parallel CFD, Proceedings of the 17th International Conference on Parallel CFD, Argonne National Laboratory preprint, ANL/MCS-P1283-0805, 2006.
3. B. Norris, S. Bhowmick, D. Kaushik, and L. McInnes, Middleware for dynamic adaptation of component applications, Argonne National Laboratory preprint, ANL/MCS-P1365-0806, to be submitted by September 30, 2006 to the Proceedings of WoCo9, Grid-based Problem Solving Environments: Implications for Development and Deployment of Numerical Software.
4. S. Bhowmick, D. Kaushik, L. McInnes, B. Norris, and P. Raghavan, Parallel adaptive solvers in compressible PETSc-FUN3D simulations, Proceedings of the 17th International Conference on Parallel CFD, Argonne National Laboratory preprint, ANL/MCS-P1279-0805, 2006.
5. B. Norris, L. McInnes, and I. Veljkovic, Computational quality of service in parallel CFD, Proceedings of the 17th International Conference on Parallel CFD, Argonne National Laboratory preprint, ANL/MCS-P1283-0805, 2006.
6. S. Bhowmick, D. Kaushik, L. McInnes, B. Norris, and P. Raghavan, Parallel adaptive solvers in compressible PETSc-FUN3D simulations, Proceedings of the 17th International Conference on Parallel CFD, Argonne National Laboratory preprint, ANL/MCS-P1279-0805, 2006.
7. L. McInnes, J. Ray, R. Armstrong, T. Dahlgren, A. Malony, B. Norris, S. Shende, J. Kenny, and J. Steensland, Computational quality of service for scientific CCA applications: composition, substitution, and reconfiguration, Argonne National Laboratory preprint, ANL/MCS-P1326-0206, February 2006.
8. B. Norris, S. Bhowmick, D. Kaushik, and L. McInnes, Middleware for dynamic adaptation of component applications, Argonne National Laboratory preprint, ANL/MCS-P1365-0806, to be submitted by September 30, 2006 to the Proceedings of WoCo9, Grid-based Problem Solving Environments: Implications for Development and Deployment of Numerical Software.
9. S. Bhowmick, D. Keyes, Y. Freund, V. Eijkhout and E. Fuentes, Application of machine learning to selecting solvers for sparse linear systems, submitted to the *International Journal of High-Performance Computing Applications*, September 2006.
10. S. Bhowmick, D. Keyes, Y. Freund, V. Eijkhout, and E. Fuentes, Application of machine learning to selecting solvers for sparse linear systems, SIAM Conference on Parallel Processing for Scientific Computing, San Francisco, Calif., February 22-24, 2006.
11. B. Norris, Computational quality of service for scientific components, IFIP Working Conference WoCo9, Grid-based Problem Solving Environments: Implications for Development and Deployment of Numerical Software, Prescott, Ariz.

Multi-Scale Analysis of Protein Folding

PI: R. Stephen Berry, The University of Chicago; Chemistry

We have performed a first comprehensive application of multiscale simulation to the process of proteins folding in water. Multi-scale simulation is the use of Monte Carlo simulation over long time scales and molecular dynamics over short time scales where the short-time dynamics are important. The complete result is a distributable computation methodology that can explore vastly longer time scales than a single molecular dynamics trajectory with more detail than Monte Carlo results.

We analyzed the folding behaviors of the SH3 protein using three separate force fields. This required simulation of 150 long-time trajectories on the Jazz cluster using the Tinker molecular dynamics package. From the resulting trajectories, we also used the cluster to analyze the folding properties of the system. This analysis included several energy calculations, radius of gyration, and root-mean squared deviation to the native fold of the protein. It also was necessary to visualize the resulting data set, so we generated viewable data files and graphics files for analysis. This complemented our work on Ubiquitin and Trp-Cage proteins done elsewhere.

We were able to demonstrate a preferred pathway to protein folding that develops tertiary and secondary structure progressively in distinct steps rather than either a strictly cooperative or all-or-none approach. Further, we demonstrated that computational folding times were consistent with first-order kinetics throughout their folding trajectory rather than reaching a fast folding region near the folded state. Moreover, we explored differences in how well different force fields predicted protein folds.

Most of the results have already been included in the Ph.D. thesis of David Kendall, who completed the requirements for his degree in June 2006. We are now preparing a manuscript based on this work for publication.

Computational Nanocatalysis (COMPnano)

PI: Peter Zapol, Chemistry, Materials Science, Center for Nanoscale Materials

Computational nanocatalysis projects involving density functional theory, molecular orbital theory, and tight-binding methods were performed at the LCRC facility this year. Theoretical studies of reactivity in these types of nanoporous materials are an important complement to experiment because they can provide information on adsorption energies and activation barriers for insight into the catalytic reaction mechanisms, such as for selective catalytic oxidation of hydrocarbons. The following paragraphs summarize our achievements in simulations of nanoporous catalysts.

Molecular Dynamics Simulations of Amorphous Alumina Nanopores (S. Adiga, P. Zapol)

Nanostructured membrane catalysts based on anodized aluminum oxide (AAO) hold promise for improving catalyst performance and stability. The ability to design AAO membranes with tailor-made compositions makes them an ideal support for ultra-uniform catalysts. AAO has tunable pore diameters as small as 30 nm, which can be further shrunk to only a few nanometers by atomic layer deposition of amorphous Al_2O_3 layers with precise control of stoichiometry and thickness. The surface structure of these alumina films is of great interest since they serve as a support for highly dispersed catalytically active species.

The surface structure of dry and hydroxylated amorphous Al_2O_3 has been studied by using atomistic molecular dynamics simulations. The density profiles indicate that oxygen is preferred at the surface, causing Al enrichment just below in the case of dry surfaces. In the case of hydroxylated alumina surfaces, hydroxyl ions are preferred at the surface with hydrogen ions occupying the outermost layer. The hydroxyl termination also leads to disappearing subsurface Al peak. The surface hydroxyl groups are predominantly singly coordinated with Al with a fraction of double coordinated. Hydroxylation also brings about an increase in surface roughness (Figure 30).



Figure 30. Surface contour (left); snapshot of hydroxylated amorphous alumina film (right).

We used MOLDY, a parallel molecular dynamics code that uses many-body potentials, to carry out AAO surface structure studies on Jazz with 32–48 PEs for about 50,000 atom calculations.

Catalytic Sites and Reaction Mechanisms (S. Adiga, P. Redfern, M. Sternberg, P. Zapol, and L. Curtiss)

The Mars-van Krevelen mechanism of oxidative dehydrogenation of propane on the (010) surface of V_2O_5 , as well as on supported vanadia species, was studied by using a hybrid density functional. This study was motivated by Argonne's experimental research of this reaction on supported membrane nanocatalysts. The surface was modeled by using a periodic representation of the V_2O_5 surface. The B3PW91 density functional method was used in most of the calculations. The calculations were performed with the CRYSTAL03 code. The propane adsorption energies were calculated and compared with cluster calculations.

The potential energy surface was found to involve a singlet/triplet state curve crossing with very high barriers. The barriers were found to be significantly lower for the supported vanadia catalysts as compared to crystalline vanadia surface. The periodic calculations were used to validate key energies obtained using the cluster models and confirm the main findings.

CRYSTAL calculations are performed on system sizes of ~50 atoms. The jobs run for 8–10 days on 6–10 processors.

Publications/Presentations

1. S. P. Adiga, P. Zapol, and L. A. Curtiss, Atomistic simulations of amorphous alumina surfaces, *Phys. Rev. B*, 2006.
2. P. C. Redfern, P. Zapol, M. Sternberg, S. P. Adiga, S. A. Zygmunt, and L. A. Curtiss, Quantum chemical study of mechanisms for oxidative dehydrogenation of propane on vanadium oxide, *J. Phys. Chem. B* 110, 8363, 2006.
3. S. P. Adiga, P. Zapol, and L. A. Curtiss, Atomistic simulations of nanoporous anodic aluminum oxide, in *Nanoparticles and Nanostructures in Sensors and Catalysis*, edited by C-J. Zhong, N. A. Kotov, W. Daniell, and F. P. Zamborini, MRS Symposium Proceedings, 2005.

Computational Nanophotonics (NANOPHOTONICS)

PI: Stephen Gray, Chemistry

Nanophotonics is an exciting area of nanoscience that involves the study of light interacting with nanostructured systems. Nanophotonics is leading to new biological and chemical sensors, as well as the development of nanoscale devices that use photons to replace electrons as information carriers. Often, the special properties of surface plasmons (SPs) are exploited. SPs are collective excitations of electrons near the surfaces of nanostructured metallic systems. Computational modeling is essential both for interpretation of experiments and for suggesting new directions. This project involves the development and application of computational electrodynamics methods to study nanophotonics problems. The finite-difference time-domain (FDTD) method is one of the main approaches we use, which involves grid-based solution of the relevant Maxwell's equations to generate the electric and magnetic. SP excitations can involve strong field variations over 1–10 nm scales near the surfaces of much larger, 100–1000 nm scale metallic structures. Such problems require the use of large, dense grids and can consume significant amounts of computer time and memory. Parallel computing approaches are needed, with the simulation of many realistic problems being possible only in a parallel environment.

We showed how simple slits in metals such as silver can transport light across subwavelength bends (Figure 30). This is an encouraging result for the development of dense, nanoscale, all-optical circuits. Jazz supported our work by providing a parallel computing environment that enabled this prediction.

The slit modeling has resulted in an experimental project at Argonne aimed at verifying the light bending phenomenon.

We completed work on a two-dimensional version of *Shapes*, a parallel 2-D FDTD adaptive mesh refinement FDTD application. The code is now generally available from our website (<http://www.mcs.anl.gov/Nanophotonics/>) and described in detail in a paper appearing in *Computer Physics Reports*. We also developed working 2-D and 3-D versions of NEKCEM, an efficient, parallel, and spectrally accurate finite-element type code to complement our FDTD capabilities. We developed 2-D and 3-D sequential and parallel visualization capabilities for *Shapes* HDF5 data, with support for interactive and batch operation using the Rocketeer visualization suite. We also developed a prototype graphical tool for the creation of *Shapes* configuration files for 2-D simulations.




Figure 30. Time sequence of the electric field as light is input from the bottom of a slit waveguide in silver. Remarkably, the light is efficiently bent, i.e., relatively little light is reflected back.

Publications/Presentations

1. M. S. Min, T.-W. Lee, P. F. Fischer, and S. K. Gray, Fourier spectral simulations and Gegenbauer reconstructions for electromagnetic waves in the presence of a metal nanoparticle, *J. Comp. Phys.* 213, 730–747, 2006.
2. T. W. Lee and S. K. Gray, Subwavelength light bending by metal slits, *Optics Express* 13, 9652–9659, 2005.
3. K. L. Shuford, M. A. Ratner, S. K. Gray, and G. C. Schatz, Finite-difference time-domain studies of light transmission through nanohole structures, *Appl. Phys. B* 84, 11–19, 2006.
4. J. M. Oliva and S. K. Gray, A computational study of the interaction of light with silver nanowires of different Eccentricity, *Chem. Phys. Lett.* 427, 383–389, 2006.
5. Z. Meglicki, S. K. Gray, and B. Norris, Multigrid FDTD with Chombo, *Comp. Phys. Comm.*, in press, 2006.

Simulation of Water in Carbon Nanotubes

PIs: Alexander I. Kolesnikov, Intense Pulsed Neutron Source; Christian J. Burnham, Department of Physics, The University of Houston

Quasi-one-dimensional water encapsulated inside single-wall carbon nanotubes, here referred to as *nanotube-water*, was recently studied by neutron scattering and molecular dynamics (MD) simulations. The system is regarded as a model for the study of the transport of water in biological pores. Earlier results revealed an anomalously soft dynamics characterized by pliable hydrogen bonds and large-amplitude motions of the hydrogen (quasi-liquid-like behavior) at temperatures as low as 50 K. Using the Jazz cluster, we performed comprehensive simulations of water behavior confined in nanotubes of different radii, at different temperatures and pressures. Figure 31 shows the structure of nanotubes-water at 50 K.




Figure 31. Structure of water confined in a single-wall carbon nanotube (blue) from molecular dynamics simulation. The confined water molecules take the form of a soft “ice-shell plus water-chain” structure: an ice sheath (red and white) wrapped around a 1-D chain (green and white).

Using high-resolution quasielastic neutron scattering, we investigated the temperature dependence of single-particle dynamics of water confined in single-wall carbon nanotubes with the inner diameter of 14 ± 1 Å. The temperature dependence of the alpha relaxation time measured on cooling from 260 K to 190 K exhibits a crossover at 218 K from a Vogel-Fulcher-Tammann law behavior to an Arrhenius law behavior, indicating a fragile-to-strong dynamic transition in the confined water. This transition may be associated with a structural transition from a high-temperature, low-density (< 1.02 g/cm³) liquid to a low-temperature, high-density (> 1.14 g/cm³) liquid found in our “parallel tempering” MD simulation at about 200 K. At this temperature, the nano-ice melts into a liquid-like phase (Figure 32). These results appear in good agreement with our quasi-elastic and deep inelastic neutron scattering data.




Figure 32. Radial/thermal distribution function of waters inside carbon 10/10 single wall carbon nanotube of radius 14 Å. Black surface: oxygen distribution, blue surface: hydrogen distribution. The central peak in the H-distribution corresponds to the chain of waters running down the center of the nanotube. The two outer peaks in the O-distribution are from shell molecules.

Publications/Presentations

1. A.I. Kolesnikov et al., *Phys. Rev. Lett.* 93, 035503, 2004.
2. E. Mamontov et al., *J. Chem. Phys.* 124, 194703, 2006.
3. G. Reiter et al., *Phys. Rev. Lett.* 97, 247801, 2006.

Electronically Nonadiabatic Direct Dynamics for Barrierless Reactions


PI: Stephen Klippenstein, Chemistry

Transition state theory (TST) may be used to predict accurate rate coefficients for many chemical reactions. Some reactions, however, involve coupled electronic states, and these electronically nonadiabatic systems may not be accurately modeled by using conventional implementations of TST. We are developing computer codes and methods capable of performing semi-classical trajectory simulations, which may be used to estimate rate coefficients for nonadiabatic events. These results may be used in conjunction with TST calculations to more fully characterize a wide variety of chemical systems.

We have developed a coupled-states trajectory code capable of direct evaluations of potential energy surfaces and their gradients by using the electronic structure packages Gaussian and Molpro. Nonadiabatic dynamics may be modeled by using either the fewest-switches surface hopping or the decay-of-mixing semi-classical trajectory methods. Methods for treating spin-orbit coupled systems are currently being developed and implemented. Strategies for incorporating other kinds of nonadiabatic couplings will be developed for future applications.

Preliminary investigations have been made of the dynamics of the OCH_2CH_2 biradical, which is the initial complex formed by the reaction $\text{O}(3\text{P}) + \text{CH}_2\text{CH}_2$ (Figure 33). The biradical can either decay on the triplet surface or undergo intersystem crossing and decay on the singlet surface. There is experimental evidence for both decay mechanisms, but the rate for intersystem crossing remains uncertain. Calculations are under way to estimate the rate for intersystem crossing for this process.

Trajectory simulations show excellent scalability with the number of processors, and, because of the computational cost of evaluating the potential energy surfaces and gradients, parallel simulations on Jazz are necessary to obtain meaningful statistics.



The ability to accurately predict the importance of nonadiabatic effects in chemical systems using reliable potential energy surfaces and surface couplings will be useful in formulating reaction mechanisms and interpreting experimental results for a wide variety of chemical systems. The previously developed computer code ANT was extensively modified to include improved methods for selecting initial conditions and interfaces with the existing electronic structure programs Gaussian and Molpro.

Figure 33. Snapshot from an O atom insertion into CH_2CH_2 computed using ab initio semi-classical trajectory simulations. The incoming O atom first abstracts an H atom, and the resulting OH molecule then adds to CH_2CH .

Neocortical Seizure Simulation (Neocortex_Sim)

PIs: Mark Hereld, Mathematics and Computer Science; Rick Stevens, Computing and Life Sciences and The University of Chicago; Wim van Drongelen and Hyong Lee, The University of Chicago

Epilepsy, characterized by repeated seizures with no direct external cause, is among the most commonly diagnosed neurological diseases. In electroencephalogram (EEG) recordings of epileptic patients, seizures are often seen to initiate from a specific area of the brain (the focus) and spread, so the epileptic focus is generally suspected of possessing abnormal neurons, abnormal connectivity between neurons, or both. Furthermore, seizure foci often exhibit seizure-like behavior that does not spread, implying that conditions need to be right in order for full clinical seizures to occur.

Ideally, one would like to record from many cells simultaneously *in vivo* to create a picture of how seizures start and spread from the focus, much as current EEG records possess upwards of 128 channels. Unfortunately, current *in vitro* experimental recording techniques are time-consuming and generally limited to a handful of simultaneous recording sites or cells; they thereby provide an incomplete picture of the process at best. *In vivo* experiments are even more challenging.

We are attempting to overcome some of these limitations by investigating seizure generation in a scalable computational model of neocortex, a virtual focus in some sense. The cells and connectivity included in the model are, of course, simplifications of reality, but they preserve what we consider the essential elements—multiple inhibitory and excitatory cell types, Hodgkin-Huxley-based ion channels, and several types of interneuron connections that vary in their topology and signaling speed. The model helps address a number of open questions in the field: Are seizures more of a cellular (ion channels) or network (connectivity) phenomenon? How small can a focus be? What conditions govern whether seizure-like behavior propagates or dies out? The model can provide insight into these questions because it is much more easily manipulated than real neurons in an experimental setup.

In 2006, we continued to investigate the mechanisms that cause seizure-like discharges in EEG. Two interesting results stand out. A series of simulations manipulating the conductivity of a particular channel (persistent sodium) suggested that it could be responsible for network bursting observed in mouse brain preparations *in vitro*. Experimental work using a pharmacological agent to reduce the conductivity of this channel appears to corroborate this result.

In previous work, the simulation showed seizure-like behavior propagating between two laterally separated neocortical patches that were loosely coupled. One of these patches included a persistent sodium channel that drove the oscillations. The oscillations in that single patch now appear to result from an interplay between two populations of excitatory neurons in the patch (the superficial and deep pyramidal cells). The overall pattern cannot be discerned in the activity of the individual cells; it becomes apparent only when the populations as a whole are examined. This mechanism, if upheld, lends insight into why it has been difficult to find differences between epileptic and nonepileptic tissue at the level of single-cell measurements. Analyzing the

simulations in this way was straightforward, but not possible experimentally, owing to the impossibility of separating contributions to EEG from different, but overlapping, populations of cells.

As Figure 34 illustrates, this work has led to a new perspective on generation of EEGs similar to those seen in epileptic episodes. The discovered relationship to synaptic excitation is novel and unexpected. We hope these results will lead to a better understanding of possible mechanisms in real brains and lead to testable hypotheses as we improve the fidelity of our model and develop means to relate these computed results to measurements in both the laboratory and the live brain.

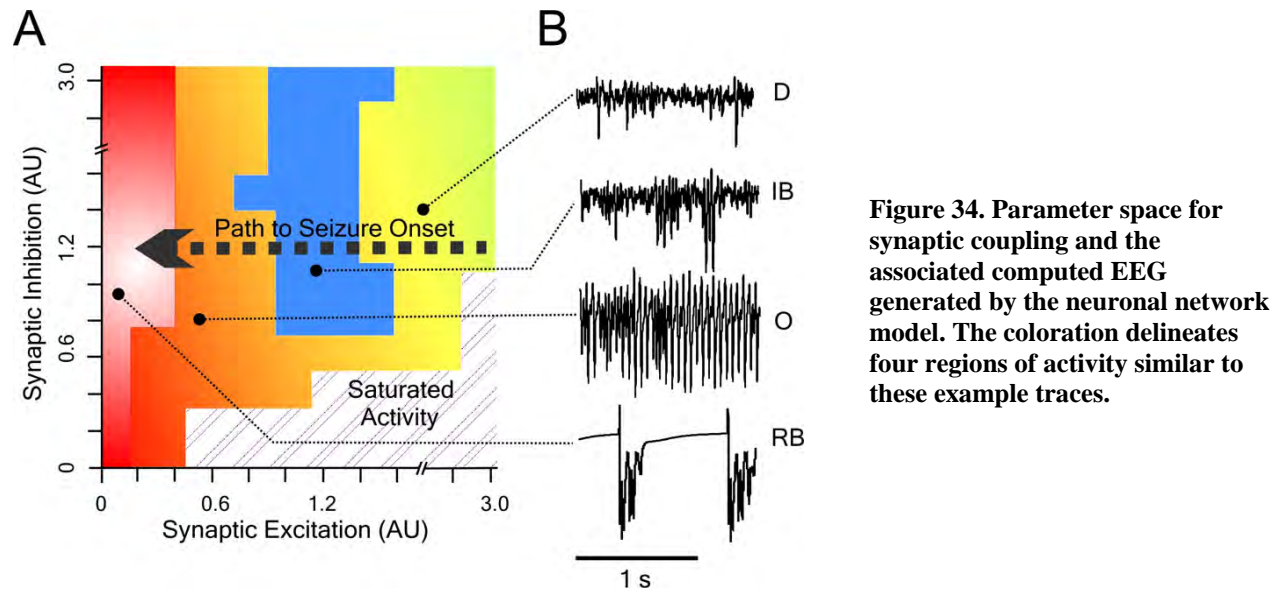


Figure 34. Parameter space for synaptic coupling and the associated computed EEG generated by the neuronal network model. The coloration delineates four regions of activity similar to these example traces.

Our model scripts have been expanded to incorporate an automated parameter search capability. This will allow better utilization of a large machine like Jazz to search for specific combinations of model parameters for which the model behavior matches pre-categorized experimental traces (e.g., seizure, postseizure) in a more objective manner.

Publications/Presentations

1. W. van Drongelen, H. Koch, F. Elsen, H. Lee, A. Mrejeru, E. Doren, C. Marcuccilli, M. Hereld, R. Stevens, and J.-M. Ramirez, The role of persistent sodium current in bursting activity of mouse neocortical networks in vitro, *Journal of Neurophysiology*, in press.
2. M. Hereld, R. Stevens, H. Lee, and W. van Drongelen, Framework for interactive million-neuron simulation, presented at the international conference, "An Overview of Epilepsy Research: What, Where, When and Why," May 2006, to appear in *Journal of Clinical Neurophysiology*.
3. W. van Drongelen, H. Lee, R. Stevens, and M. Hereld, Propagation of seizure-like activity in a model of neocortex, presented at the international conference, "An Overview of Epilepsy Research: What, Where, When and Why," May 2006, to appear in *Journal of Clinical Neurophysiology*.

Heat Transfer in Advanced Reactor Cores

PIs: Gary Leaf and Paul Fischer, Mathematics and Computer Science

This project's goal is to understand the fundamental thermal mixing phenomena within the core of advanced burner reactors. To leading order, the hydrodynamics can be decoupled from the thermal and neutronics transport and therefore studied independently. The thermal transport, however, is directly dependent on the hydrodynamics. To better understand the core thermal mixing, we have commenced a sequence of large-eddy simulations (LES), in which the transport-relevant scales are directly computed through numerical simulation while the smaller (sub-grid) scales are modeled.

The objectives are as follows:

- Provide a detailed hydrodynamics test data for validation of coarse-grained codes that are central to production simulations of reactor core thermal hydraulics
- Assess the degree of mixing within the subassembly
- Predict the coolant bypass flux at assembly walls
- Predict the pressure drop across the subassembly
- Predict the thermal and hydrodynamic entrance length in the reactor core
- Assess the influence of wire wrap parameters (e.g., pitch) on mixing efficiency and reduction of T_c for a given power load

The computations are being performed with the Argonne code Nek5000. The Jazz effort is directed at developing Nek5000 to provide thermal transport capabilities tailored specifically for coupling to the Argonne neutronics code Unic. This coupling will provide a high-fidelity reactor core simulation capability specifically targeted for advanced burner reactors that are a critical element in the Global Nuclear Energy Partnership (www.gnep.energy.gov) being led by the U.S. Department of Energy.

While most of the simulations have been done on Argonne's Blue Gene/L platform, Jazz has been instrumental for developing the computational models (i.e., mesh, boundary conditions, physical parameters) for this project. With these computations, we have been able to map out the mean and fluctuating velocity fields in an array of wire-wrapped fuel pins (Figure 35). While many groups have performed LES in rod-bundle arrays in the past, these are the first to examine the significant influence of the wire-wrap spacers that are fundamental to promoting interchannel mixing within the core.

We are still analyzing the current data from simulations of a single wire-wrap pin in a periodic array at three different Reynolds numbers. These, coupled with computations to be undertaken this year, will allow us to characterize pressure drop and degree of mixing as a function of wire-wrap pitch. In addition, we expect to derive coarse-grained models from these computations that can be used directly in coupled thermal hydraulics-neutronics computations of burner reactor behavior.



Figure 35. Spectral element computation using Nek5000 showing turbulent structures in the boundary-layer of a wire-wrapped fuel pin within a fast reactor. Understanding such turbulent flows will lead to improved thermal efficiency of advanced burner reactors.

We have made significant changes to Nek5000 this year, primarily to allow scaling to 32,000 processors on IBM's Blue Gene Watson platform. The changes included rewriting our domain partitioning software to handle large numbers of elements and processors, our interprocessor communication utility to eliminate a scaling bottleneck in the earlier version's setup phase, and our I/O to support parallel I/O. All of these were mandatory for satisfactory performance above $P=2048$ and obviously not a requisite for Jazz. Jazz's interactive capability and ease of use, however, made it a particularly productive platform on which to debug and develop scalable software that could not readily be tested on smaller platforms or on batch-only machines.

We have also developed a limited capability to couple the thermal hydraulics computations of Nek5000 to the neutronics power density computations of Unic and have performed a preliminary set of computations with this coupling. These codes will be more tightly integrated in the coming year.

Publications/Presentations

1. G. Palmiotti, M. Smith, C. Rabiti, E. Lewis, W. Yang, M. Leclerc, A. Siegel, P. Fischer, D. Kaushik, J. Ragusa, J. Lottes, and B. Smith, Status report on high fidelity reactor simulation, ANL-AFCI-175.
2. P.F. Fischer, Scaling to petaflops, presented at the Advanced Nuclear Energy Systems Workshop, Washington, DC, August 15–17, 2006.
3. P.F. Fischer, Preconditioning spectral element methods, presented at the Symposium on Advanced Numerical Methods in Engineering Applications, Brussels, July 3–4, 2006.

Integrated 3-D Simulation of Neutronic, Thermal-Hydraulic, and Thermo-Mechanical Phenomena (NUMERICAL REACTOR)

PI: Tanju Sofu, Nuclear Engineering

We are developing next-generation computation tools for design and analysis of light-water reactor (LWR) nuclear energy systems based on high-fidelity, first-principles-based simulations. The product software system, known as the Numerical Nuclear Reactor (NNR), is capable of performing integrated neutronic and thermal-hydraulic analysis for a reactor core at an unprecedented scale and level of fidelity. With NNR, a whole-core LWR simulation with pin-by-pin representation of the fuel assemblies is feasible with intrapin level thermal feedback and explicit representation of individual coolant channels based on CFD techniques (Figure 36).

Originally developed under the DOE I-NERI program for pressurized water reactor (PWR) systems, NNR's analysis capabilities have been extended for boiling water reactor (BWR) applications in FY2006. The methodology of the neutronics component, DeCART, also was extended to treat the nonperiodic structure of BWR fuel assemblies by removing the pin-cell symmetry assumption and introducing an assembly-level modular ray-tracing algorithm. Also, the new Eulerian two-phase boiling model of the thermo-fluid analysis module, STAR-CD, has been integrated with software system. In parallel, specific activities related to the verification and benchmarking of the BWR version of the neutronics module, experimental validation of the two-phase flow boiling model, and application of the NNR to address the operational and design issues for BWRs have been initiated. The availability of Jazz was critical for our demonstrations in addressing significant scale core design and operational issues.




Figure 36. Pin power distribution (W/cc) in core mid-plane (left) and exit void fraction (right) for a small fuel bundle using the integrated NNR simulations with intrapin level thermal feedback.

Our calculations to date have set the stage for radically changing the nuclear reactor core simulations. In part as a result of NNR's successes, several new advanced simulation initiatives recently have flared both nationwide and worldwide to take advantage of the high-performance computing infrastructure for high-fidelity, first-principles-based simulations to support the broad spectrum of nuclear engineering enterprise elements.

Publications/Presentations

1. D. P. Weber et al., High-fidelity LWR analysis with the numerical nuclear reactor, *Nuclear Science and Engineering*, accepted for publication in May 2006.

Parallel Tools Performance Testing (PARALLEL-TOOLS)

PIs: Robert Latham and Robert Ross, Mathematics and Computer Science

The parallel tools project continues to develop and study highly scalable software tools and libraries for parallel applications. Our efforts include a parallel file system (PVFS), an MPI-2 implementation (MPICH2), an MPI-IO implementation (ROMIO), and work on higher-level libraries (Parallel-NetCDF).

PVFS provides a highly capable file system for today's large clusters as well as future clusters. Designed to scale to hundreds of servers and tens of thousands of clients, PVFS provides a set of features beneficial for scientific applications. As a parallel file system, it provides a global namespace for all client nodes. PVFS provides methods to efficiently express the noncontiguous I/O access patterns common in scientific and high-end computing applications. The PVFS design, motivated in no small part by MPI-IO requirements, integrates seamlessly into MPI-IO implementations.

Our MPI-2 implementation gives us a vehicle for exploring low-latency, high-performance MPI-2 routines. Our research in scalable collective routines yields higher performance on clusters of all sizes. ROMIO, our implementation of the I/O chapter of the MPI-2 specification (commonly known as MPI-IO), provides the basis for virtually all MPI-IO implementations.

While PVFS2 and MPI-2 offer powerful tools for achieving high performance, they do not always provide the most convenient interfaces for application writers. To that end, higher-level libraries such as Parallel-NetCDF can help scientists achieve high I/O rates, while allowing them to focus on their science.

As a whole, these parallel tools projects provide a way to expose the performance of underlying hardware while allowing application writers to write "performance portable" code that can work well on the wide variety of high-performance computing resources available today, as well as the new designs of the future.

As we do every year, we made extensive use of Jazz for evaluation of novel approaches for high-performance I/O. We explored ways that processes could cache data on the client side in ROMIO while still maintaining MPI-IO semantics. In PVFS2, we have started experiments with various data replication strategies. Further, all three fronts (PVFS, MPICH2, and pnetcdf) saw extensive testing and benchmarking. Thanks to the nightly tests we run on a portion of Jazz, we have caught several regressions in both PVFS2 and ROMIO's PVFS2 driver. Jazz's node count and Myrinet interconnect made for a useful testbed for our experiments.

Research in PVFS2 data replication and client-side MPI-IO caching is ongoing. Jazz results have helped us confirm the validity of our approaches or suggest refinements that we should make to maintain or improve performance.

Work done on Jazz has made its way into several product releases: PVFS2-1.5.1, MPICH2-1.0.4, and Parallel-NetCDF-1.0.1.

Parallelization of the Accelerator Modeling Code *elegant*

PI: Michael Borland, Accelerator Systems

Argonne has several electron accelerator facilities that provide unique and state-of-the-art capabilities—in particular, the Advanced Photon Source (APS) and the Argonne Wakefield Accelerator. Each of these has an important role to play in pushing the design of the next generation of accelerators, from synchrotron light sources to high-gradient colliders. However, the computational tools available for these fields have not kept pace with the requirements or state-of-the-art computing. The ability to provide powerful computational tools for simulating accelerators is critical for both existing accelerators and the design of next-generation machines.

The accelerator simulation code *elegant*, developed at the APS, has a large suite of capabilities and a worldwide user base in the accelerator community. Among other projects, it has had a prominent role in the evolution of the APS and the design of the Linac Coherent Light Source.

Jazz was used for developing, debugging, and tuning the performance of our parallel implementation of *elegant*. Some existing tools, such as totalview and Jumpshot, are useful for taking advantage of this high-performance facility efficiently. It also served as a medium step before porting our software to the more advanced supercomputer (i.e., the BlueGene/L).

Several APS physicists also have benefited from the current evolution of the parallel *elegant* (*Pelegant*). It achieved very good performance for some practical simulations, such as multiparticle tracking with synchrotron radiation and emittance blowup in the vertical rf kick scheme. This scheme promises a breakthrough in production of short, intense x-ray pulses from a high-energy light source. *Pelegant* has allowed much faster exploration of issues related to implementation of this scheme (Figure 37).




Figure 37. *Pelegant* simulation of the evolution of beam in the APS with 1 kHz pulsed crab cavities at various deflecting voltages, showing the increase in the vertical emittance. The speed of *Pelegant* allows performing such parameter explorations very quickly.

We parallelized many time-intensive elements in the accelerator simulation code *elegant*. This reduced the simulation time for some practical problems significantly. For example, for an APS simulation, including symplectic element-by-element tracking, accelerating cavities and crab cavities, the simulation time was reduced from 14.3 days to 42 minutes on 512 CPUs of the BlueGene/L supercomputer. The speedup for this particular run was 484 with efficiency near 95% (Figure 38).




Figure 38. Performance of parallel *elegant* for a crab cavity simulation as a function of the number of processors used on BlueGene/L.

Publications/Presentations

1. Y. Wang and M. Borland, *Pelegant*: A parallel accelerator simulation code for electron generation and tracking, to appear in the Proceedings of the 12th Advanced Accelerator Concepts Workshop.

Estimation of Protein: Oligonucleotides Binding Energies

PI: Matthew Wilce, Department of Biochemistry and Molecular Biology, Monash University, Victoria, Australia

Interaction of proteins and polynucleic acids, RNA, and DNA is key in all cells. Yet the molecular basis of protein-oligonucleotide affinity and specificity is poorly understood. Through the comparison of computer simulations and experimental measurement, we aim to get closer to understanding these interactions at an atomic level. This project utilizes molecular dynamics simulations as implemented with the computer program NAMD. In particular, it focuses on estimating the change in Gibb's free energy of protein oligonucleotide interactions using both free energy of perturbation (FEP) and steered molecular dynamics (SMD) protocols. We are concurrently experimentally determining the interaction energies of targeted molecular systems for comparison.

To date, we have focused on the interaction of polyC binding proteins (PCBP) and their cognate RNA. PCBPs bind to mRNA and affect translational efficiency (i.e., the production of proteins from RNA) as well as mRNA stability. We have determined the crystal structures of a number of polyC bind proteins in complex with RNA and have measured their affinities and binding kinetics experimentally. These results put us in an ideal position to use computational tools to see whether we can effectively model these interactions.

These calculations have been carried out on 16 to 96 nodes on the Jazz cluster. We anticipate publishing our results (some of which are summarized below) later in 2007.

Key results include the following:

- FEP protocols have been utilized where we have grown the extra O₂ oxygen of RNA on a DNA molecule binding to PCBP. This allows us to estimate the difference in binding energy of RNA vs. DNA. The *in silico* calculations suggesting that PCBP's preferential binding DNA over RNA are consistent with those determined experimentally using Surface Plasmon Resonance.
- SMD protocols have been utilized to pull the oligonucleotide away from an anchored domain of PCBP. We record the energy required to pull the oligonucleotide away from the protein against time throughout the simulation and then integrate to calculate the total energy. Since we know the distance traveled by the oligo, we can calculate the total work. The amount of work required to pull off the oligo can be converted to the change in Gibb's free energy using the Jarzinski Identity. This work is ongoing.

Quantum Monte Carlo Calculations of Light Nuclei (QMC_FOR_NUCLEI)

PIs: Kenneth M. Nollett, Muslema Pervin, Steven C. Pieper, and Robert B. Wiringa, Physics

This project uses quantum Monte Carlo [Green's function (GFMC) and variational (VMC)] methods to compute ground-state and low-lying excited-state expectation values of energies, densities, structure functions, transitions, astrophysical reaction rates, and so forth for light nuclei. Realistic two- and three-nucleon potentials are used. Our goal is a description of all these systems and processes using a Hamiltonian that also provides an excellent description of nucleon-nucleon scattering and nucleonic matter. Such a "standard nuclear model" can then be used, for example, to compute low-energy astrophysical reaction rates that cannot be experimentally measured.

During FY2006, Jazz was used for many aspects of this project. The fast turnaround of Jazz makes it the machine of choice (along with Blue Gene) for all but our biggest calculations, which are done at the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory or at Los Alamos National Laboratory. The specific projects this year fell into five categories: GFMC scattering, excited states in light nuclei, transition matrix elements, isospin mixing in ^8Be , and spectroscopic factors.

GFMC Scattering

In FY2006, we completed our first calculations of nucleus-nucleon scattering using GFMC. The system studied in this development work is scattering of a neutron from ^4He . We finished a long test series and found conditions on computational parameters that are in principle arbitrary but that affect the results at the desired high precision levels. These conditions allow us to achieve precision of better than 50 keV in scattering energy. The final result for our best potential, shown as points with error bars in Figure 39, agrees very well with laboratory data. All of these calculations, except those of Carlson, were performed on Jazz.




Figure 39. Cross sections for ^4He -neutron scattering in the $3/2^-$ (blue), $1/2^-$ (red), and $1/2^+$ (green) partial waves. The solid curves are from a fit to laboratory data. The points with Monte Carlo error bars are the results of GFMC calculations performed on Jazz, using our best potential: Argonne v18 (two-body), plus Illinois-2 (three-body).

The completion of this work clears the way for many more calculations in light nuclei, though the difficulty of achieving precise results in the ^4He -neutron system indicates a cautious choice of follow-on problems. The next step will most likely be to study the ^4He -proton system and address any difficulties that arise from the Coulomb interaction of the scattering particles. The next step up in complication (because the nuclei are more weakly bound than ^4He) will be neutrons on ^3H and protons on ^3He , which can be compared with the results of other theoretical calculations. Over the longer term, there are many interesting problems to be addressed at $A=7$ and above. All of the GFMC scattering calculations in the immediate future will be of sizes well-suited to Jazz.

Excited States in Light Nuclei

In FY2006, we continued work on higher excited states in the $A=9,10$ nuclei, including multiple states of the same spin and parity. We also worked on “unnatural-parity” excitations, which come from the promotion of one nucleon to the next-higher shell. GFMC calculations reproduce the observed order of $1/2^+$, $5/2^+$, $3/2^+$, and $9/2^+$ for the first four such states of ^9Be and get about the correct total energies. In FY2006, we completed the first VMC calculations of unnatural-parity states in ^{10}Be and ^{10}B , and obtained the observed ordering of states, which is 1^- , 2^- , 3^- , 4^- in ^{10}Be , but 2^- , 3^- , 4^- , 1^- in ^{10}B . GFMC calculations of these states are on the agenda for FY2007. All the initial VMC diagonalizations and many of the final GFMC calculations were made on Jazz.

Transition Matrix Elements

An important consequence of having a predictive theory of at least light nuclei will be the ability to compute reliably the matrix elements that determine the rates of nuclear decays. These include the $B(M1)$ and $B(E2)$ matrix elements for photon emission and the $B(GT)$ matrix elements for weak decay. In the past, we were able to compute these matrix elements only using the VMC method. These VMC calculations came within 5% to 10% of the experimental results at $A=6$ and 7 but did not fare as well in $A=8$ nuclei.

In the last few months of FY2006, we modified the VMC program to read in wave function samples from GFMC calculations and compute mixed matrix elements for these transitions. Preliminary results from the GFMC calculations for $A=6,7$ nuclei are generally closer to the experimental data than the VMC calculations, sometimes within 1% to 2%. A major effort in FY07 will be to study many additional transitions in $A=6-10$ nuclei and compare them with an extensive body of published experiments.

Isospin Mixing in ^8Be

Isospin (T) is an important quantum number in nuclear physics, specifying, for example, the difference between protons and neutrons. The bulk of nuclear forces conserve isospin, so specific nuclear states can often be characterized as having a definite isospin. Exceptions occur when two states of different isospin are close together in energy, allowing smaller parts of the nuclear Hamiltonian that do not conserve isospin, such as charge-symmetry-breaking and electromagnetic forces, to mix them.

A famous case in light nuclei is ${}^8\text{Be}$ in the 16–19 MeV excitation range, where two sets of three states ($2+$, $1+$, $3+$) appear with $T=0$ and $T=1$. In particular, the pair of $2+$ states is approximately 60% to 40% mixed. This fact has been used to determine an empirical mixing matrix element, $\langle 1|H|0\rangle$, of 149 keV, about half of which was expected to be due to the Coulomb part of electromagnetic interaction.

We have recently started adapting the GFMC method for transitions discussed previously to this problem, and preliminary results give about 120 keV for $\langle 1|H|0\rangle$. We expect to complete this work in FY2007 and extend it to other states, some of which are candidates for sensitive weak interaction experiments where it is crucial to know the admixture of different isospin components.

Spectroscopic Factors

We have been making an extensive set of calculations of overlap wave functions and spectroscopic factors for nuclei with A from 5 to 10. These overlap wave functions, or quasi-hole and quasi-particle amplitudes, can be used in reaction programs (such as the DWBA program Ptolemy) to compute differential cross-sections for transfer reactions. These results have been used successfully in publications describing radioactive beam experiments at the Argonne Tandem Linac System (ATLAS): ${}^2\text{H}({}_8\text{Li},p){}_9\text{Li}$ and ${}^2\text{H}({}_6\text{He},p){}_7\text{He}$, both cases in which the energies and nature of the low-lying states have been in question. In FY2006, we continued work on a general survey paper to cover many other cases in light nuclei, which have recently become experimentally accessible with the advent of radioactive beams. We made specific calculations for a new ATLAS experiment, ${}^2\text{H}({}_8\text{Li},{}_3\text{He}){}_7\text{He}$, which requires two overlaps.

The VMC code was expanded to read in configurations computed by the GFMC code and to use these configurations in calculations of $B(M1)$, $B(E2)$, and $B(GT)$ matrix elements. The codes had already been optimized for performance on Jazz in earlier years.

Publications/Presentations

1. R. B. Wiringa, Pair counting, pion-exchange forces, and the structure of light nuclei, *Physical Review C* 73, 034317, 2006.
2. K. Nollett, Quantum Monte Carlo: Not just for energy levels anymore, 3rd ANL/MSU/JINA/INT RIA Theory Workshop, Argonne National Laboratory, April 4-7, 2006.
3. R. B. Wiringa, Recent developments in nuclear quantum Monte Carlo, New Developments in Quantum Monte Carlo Workshop, Tempe, Ariz., May 15–18, 2006.
4. S. C. Pieper, Recent results from quantum Monte Carlo calculations of light nuclei, Elba IX Electron-Nucleus Scattering Workshop, Marciana Marina, Elba, Italy, June 19-23, 2006.

5. K. Nollett, Quantum Monte Carlo studies of bound and unbound nuclear states, Seventh International Conference on Radioactive Nuclear Beams Cortina d'Ampezzo, Italy, July 3–7, 2006.
6. K. Nollett, Quantum Monte Carlo studies of bound and unbound nuclear states, Nuclear Structure '06, Oak Ridge, Tenn., July 24–28, 2006.

Multilevel Agent-Based Simulation of Bacterial Chemotaxis (REPAST_CHEMOTAXIS)

PIs: Philippe Cluzel, Thierry Emonet, and Panos Oikonomou, Institute for Biophysical Dynamics, The University of Chicago; Charles Macal and Michael North, Center for Complex Adaptive Agent Systems Simulation, Decision and Information Sciences Division, Argonne National Laboratory

Cells process information and make decisions using biochemical networks. Some biological functions are carried out at the single-cell level, while other behaviors emerge at the population level mediated by cell-to-cell communication. The main focus of our project is to try unveiling design principles of biological networks using agent-based numerical modeling and mathematical methods borrowed from statistical physics.

We analyzed the stochastic dynamics of the chemotaxis pathway of *Escherichia coli* bacteria at the single-cell level. We found that the adaptation module of the system uses an ultra-sensitive architecture to control the timing of the chemotactic response, converting small changes in kinase activity into long adaptation times, and therefore, long runs. This mechanism provides an explanation for the observed large behavioral variability measured in single cells. The model also predicts a relationship between chemotaxis response and behavioral variability, in which the noisiest cells exhibit the largest chemotaxis response.

We confirmed these predictions using the Jazz cluster to run large-scale AgentCell simulations involving populations of 420 independent digital bacteria. AgentCell (www.agentcell.org) is an agent-based model of bacterial chemotaxis capable of simulating thousands of digital *Escherichia coli* bacteria in a 3-D medium containing gradients of chemo-attractants (see FY2005 report). The computing time on the Jazz cluster was instrumental for this research. We presented our results at the International Conference on Complex Systems (Boston, June 2006) and submitted a manuscript for publication in a premier journal on computational biology. Incidentally, these results led to the hiring of Thierry Emonet as assistant professor by the Molecular, Cellular, and Developmental Biology Department at Yale University.

We studied the role of network topology on evolution. Large biological systems can be viewed as complex networks of many interacting components. Such networks have been shown to follow a particular architecture called scale-free topology. We have used a model of Boolean neuron-like rules to simulate the dynamics and evolution of complex networks. Our results show that populations containing scale-free networks evolve rapidly and smoothly towards some target function (Figure 40). By contrast, equivalent random networks evolve slowly, through a succession of rare fortuitous random mutations. We used the Jazz cluster to simulate the dynamics of Boolean networks and their evolution. These results were recently published in *Nature Physics*.




Figure 40. Diagram of a scale-free network that contains components with a highly diverse level of connectivity. Some components form highly interconnected hubs, while other components have few connections. There are many levels of interconnectivity in between. Scale-free networks are pervasive in biology. Computer simulations on Argonne’s Jazz cluster show that scale-free networks are able to evolve to perform new functions more rapidly than an alternative network design.

Without the computing time on Jazz, these projects would be hard to complete, if not impossible.

AgentCell uses a legacy code written in C++ called StochSim. Because of the use of global variables in StochSim, we could use only one digital cell per Java VM. With the help of one of the original authors of StochSim (Tom Shimizu), we removed all global variables from StochSim and modified AgentCell to allow for many cells per Java VM. This opened the door for modeling with AgentCell of biological problems involving cell-to-cell communications.

Publications/Presentations

1. T. Emonet and P. Cluzel, Relationship between cellular response and behavioral variability in chemotaxis.
2. P. Oikonomou and P. Cluzel, Effects of topology on network evolution, *Nature Physics*, 2, 532–536, 2006.

Simulations of Spin Wave Excitations in Magnetic Nanoparticles

PIs: Ming Yan, Gary Leaf, and Barry Smith, Mathematics and Computer Science

The problem being addressed is that of magnetic excitations in confined geometries. It has no analytical solution and, with the advent of nanoscience and the technological interest in small magnetic particles, it is a problem of extreme current relevance. Over the past three years, the collaboration between Argonne's Mathematics and Computer Science and Materials Science divisions has resulted in a computational solution to this problem. By solving the dynamical equations of motion of a spin grid that represents the sample, we have been able to compute the magnetic normal modes of small magnetic particles (Figure 41). The frequency of these modes has proven to be intimately related to magnetic instabilities that are, in turn, crucial to developing an understanding of the stability and controllability of magnetic information bytes. At present, the codes that we have developed are unique to Argonne. These have enabled us to be at the forefront of this rapidly expanding research front.

The simulations require massive computational infrastructure. Jazz has allowed us to simulate particles that are realistic in terms of real applications. Without Jazz, it would not have been feasible to simulate problems of technological importance. In addition, we were able to conduct the necessary parameter sweeps in a timely fashion. To date, these simulations provide the only way to obtain information on the magnetization mode-dynamics in magnetic particles.

For our latest study, we needed to simulate the experimental setup in which a particle is being driven by an ac magnetic field. Moreover, the response measurements of the particle had to emulate the measurements taken in the laboratory. These requirements resulted in a substantial increase in the software demands.




Figure 41. Magnetic normal modes of a square Co particle in a vortex ground state.

Publications/Presentations (2005-2006)

1. G. Leaf, H. Kaper, M. Yan, V. Novosad, P. Vavassori, R. E. Camley, and M. Grimsditch, Dynamic origin of stripe domains, *Phys. Rev. Lett.* 96, 017201, 2006.
2. M. Yan, G. Leaf, H. Kaper, R. Camley, and M. Grimsditch, Spin wave modes in a cobalt square vortex, *Phys. Rev. B* 73, 014425, 2006.
3. M. Yan, G. Leaf, H. Kaper, V. Novosad, P. Vavassori, R. E. Camley, and M. Grimsditch, Stripe domains, dynamic modes and magnetic phase transitions, submitted to Proceedings of Intermag.
4. R. E. Camley, G. Leaf, H. Kaper, M. Yan, and M. Grimsditch, Phase transitions and reversal mechanisms in nanosized magnetic bars: the influence of bulk and localized spin wave modes, submitted to *Phys Rev B*.
5. M. Yan, P. Vavassori, G. Leaf, F. Fradin, and M. Grimsditch, Non-linear response of driven spin excitations, in preparation.



Argonne National Laboratory

9700 South Cass Avenue

Argonne, IL 60439

www.anl.gov

www.lcrc.anl.gov



UChicago ►
Argonne_{LLC}

A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC