BNL-72286-2001

The Center for Data Intensive Computing

in affiliation with the Departments of

Computer Science and
Applied Mathematics and Statistics
State University of New York at Stony Brook

www.bnl.gov/cdic

James Glimm, Director

November 2001

Brookhaven National Laboratory
P.O. Box 5000
Upton, NY 11973-5000
www.bnl.gov

Managed by
Brookhaven Science Associates, LLC
for the United States Department of Energy under
Contract No. DE-AC02-98CH10886

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 any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use 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 or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

FOR UNCLASSIFIED, UNLIMITED STI PRODUCTS

Available electronically at:

OSTI:
http://www.osti.gov/bridge

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

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

National Technical Information Service (NTIS):

Available for sale to the public from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22131
Phone: (800) 553-6847
Facsimile: (703) 605-6900
Online ordering: http://www.ntis.gov/ordering.htm

CENTER FOR DATA INTENSIVE COMPUTING

Page

1. Mission....3

2. Overview....3

3. People....4

4. Research....7

4.1 SciDAC: TSTT at CDIC....7

4.2 High Energy and Nuclear Physics....7
- 4.2.1 Muon Collider Target....8
- 4.2.2 Accelerator Design....10
- 4.2.3 Computing with the QCDOC....13
- 4.2.4 RHIC Data Management....14

4.3 Biological and Environmental Research....14
- 4.3.1 Evolution of Atmospheric Aerosols....15
- 4.3.2 Parallel Implementation of an Atmospheric Aerosol Model....19
- 4.3.3 Ocean Modeling with Automatic Mesh Refinement....20
- 4.3.4 Analysis and Visualization of Human Brain Function....23
- 4.3.5 Neuron Imaging....28
- 4.3.6 Cardiac Electrophysiology....30

4.4 Basic Energy Sciences
- 4.4.1 Magnetohydrodynamics of Multifluid Systems....31
- 4.4.2 Atomization and Spray....34
- 4.4.3 Quantum Mechanics of Magnetic Systems....35
- 4.4.4 First Principles Calculations of Melting Phenomena....38
- 4.4.5 Numerical Simulation of Photonic Crystals....38

4.5 Advanced Scientific Computing....40
- 4.5.1 Frontier-Lite....41
- 4.5.2 Uncertainty Quantification....42
- 4.5.3 Visualization....44
- 4.5.4 Brain Imaging....46
- 4.5.5 Cluster Computing....46

1. MISSION

CDIC will provide state-of-the-art computational and computer science for the Laboratory and for the broader DOE and scientific community.

We achieve this goal by performing advanced scientific computing research in the Laboratory's mission areas of High Energy and Nuclear Physics, Biological and Environmental Research, and Basic Energy Sciences. We also assist other groups at the Laboratory to reach new levels of achievement in computing.

We are "data intensive" because the production and manipulation of large quantities of data are hallmarks of scientific research in the 21st century and are intrinsic features of major programs at Brookhaven.

An integral part of our activity to accomplish this mission will be a close collaboration with the University at Stony Brook.

2. OVERVIEW

Modern science is critically dependent on computers and computing. A mix of capabilities is required, which range from advanced hardware to high level and focused software efforts to high speed networking - all aimed at the solution of forefront scientific problems which are consistent with our missions in High Energy and Nuclear Physics, Biology and Medicine, Materials and Chemical Sciences, Energy, Environmental Science, and National Security.

National plans for Information Technology in the early 21st century call for computing at the level of 100 Teraflops and data storage in the Petabytes. However, because this computational power will be obtained from massively parallel architectures there is no guarantee that scientific and engineering codes will actually achieve these speeds. This issue is central to the Department of Energy's five year program "Scientific Discovery Through Advanced Scientific Computing" (SciDAC). CDIC plays a prominent role in this program through TSTT - Terascale Simulation Tools and Technologies. TSTT is a multi institutional effort to develop interoperable adaptive grid and meshing tools for a wide variety of scientific problems. Targeted applications include climate modeling, fusion, and fluid dynamics.

Requirements for storage of experimental and observational data are higher and will reach hundreds of Petabytes. Examples at BNL include RHIC, the Relativistic Heavy Ion Collider; US-ATLAS, one of the detectors planned for the Large Hadron Collider at CERN; climate modeling and simulation; brain imaging; protein structure and function; and computational fluid dynamics. In addition, large-scale computational physics codes aimed at studying the fundamental properties of quarks and gluons (lattice quantum chromodynamics) are planned to run at tens of teraflops and will easily produce terabytes of data.

To be useful, this data must be processed, visualized, analyzed, and understood, as well as stored. Ultimately, data achieves value through its influence on some human decision. Increasingly,

decisions will be made on a computer-assisted basis. For this reason, the modern theory of data management goes beyond sheer volume to data intelligence.

CDIC will mount a research program in computational science to develop fundamentally new approaches to the problem of management of large-scale, highly complex data sets and to assess uncertainties inherent in their use.

CDIC is closely linked with the Department of Applied Mathematics and Statistics and with the Computer Science Department at SUNY Stony Brook, as well as the Information Technology Division at BNL. The Center Director and one of the staff members hold joint appointments, nine other faculty are CDIC Affiliates, and 12 graduate students contribute to CDIC projects. This arrangement adds critical strength to the Center, and has allowed a rapid initiation of its research programs.

Research in these areas provides a close link to Brookhaven's scientific programs. The Center will promote excellence, Laboratory wide, in scientific computing. CDIC will contribute to progress and productivity at BNL in a broad sense with special emphasis on:

- **High Energy and Nuclear Physics**
- **Biological and Environmental Science**
- **Basic Energy Sciences**
- **Advanced Scientific Computing**

3. PEOPLE

CDIC Staff

The Director of the Center for Data Intensive Computing (CDIC) is James Glimm; James Davenport is the Associate Director, and the Staff Assistant is Claire Lamberti. CDIC is located in Building 463, Room 255, at Brookhaven National Laboratory, Upton NY 11973-5000.

CDIC, established in 1999, presently has a staff of seven, which is expected to approximately double over the next four years. Recruitment of additional new faculty at SUNY Stony Brook, who will be associated with the Center and may hold joint appointments in the Center, will take place over the same time frame. As part of an active, ongoing collaboration with graduate students from Stony Brook, eleven Ph.D. students and several undergraduates have held scientific positions in CDIC.

TABLE 1. CDIC STAFF

Name	**Position**	**Phone/email**
James Glimm	Director	631-344-8155 (voice) 631-344-5751 (fax) glimm@bnl.gov
James Davenport	Associate Director	631-344-3789 (voice) 631-344-5751 (fax) jdaven@bnl.gov

Wonho Oh	Assistant Scientist	631-344-8156 (voice) 631-344-5751 (fax) woh@bnl.gov
Roman Samulyak	Assistant Scientist	631-344-3304 (voice) 631-344-5751 (fax) rosamu@bnl.gov
Isabel Campos	Research Associate	631-344-8603 (voice) 631-344-5751 (fax) campos@bnl.gov
Myoung-Nyoun Kim	Research Associate	631-344-3051 (voice) 631-344-5751 (fax) email
Nicholas D'Imperio	Physics Associate	631-344-8589 (voice) 631-344-5751 (fax) dimperio@bnl.gov
Claire Lamberti	Staff Assistant	631-344-3051 (voice) 631-344-5751 (fax) lamberti@bnl.gov

TABLE 2. CDIC AFFILIATES

Name	**Affiliation**	**Expertise**
Carl Anderson	BNL/Biology	Protein Structures
Carmen Benkovitz	BNL/Environmental Sciences	Atmospheric Transport
Michael Creutz	BNL/Physics	Lattice Gauge Theory
Efstratios Efstathiadis	BNL/ITD	Cluster Computing
Steven Evans	Beth Israel Medical Center	Cardiac Physiology
Flavio Fenton	Hofstra University	Cardiac Simulation
Joanna Fowler	BNL/Chemistry	Medical Imaging
Gene Gindi	SB/Radiology	Medical Imaging
John Grove	Los Alamos National Laboratory	Computational Science
Harold Hastings	Hofstra University	Dynamical Systems
Jiansheng Jiang	BNL/Biology	Protein Structures
Arie Kaufman	SB/Computer Science	Visualization
Xiaolin Li	SB/Applied Math. & Statistics	Computational Science
Brent Lindquist	SB/Applied Math. & Statistics	Porous Media; Image Analysis
Alfredo Luccio	BNL/AGS	Accelerator Design
Edward McFadden	BNL/ITD	Parallel Hardware
Michael McGuigan	BNL/ITD	Visualization & BNCT Simulations
James Muckerman	BNL/Chemistry	Combustion Chemistry
Klaus Mueller	SB/Computer Science	Visualization
Raymond Mugno	SB/Applied Math. & Statistics	Statistics

Hong Qin	SB/Computer Science	Parallel Computing
Arnold Peskin	BNL/CDIC & ITD	Visualization
John Pinezich	TST Corporation	Computational Fluid Dynamics
John Reinitz	SB/Applied Math. & Statistics	Computational Biology
Thomas Robertazzi	SB/Electrical & Computer Engineering	Grid Computing
Eunice Santos	Virginia Tech	Grid Computing
Stephen Schwartz	BNL/Environmental Sciences	Aerosol and Climate Modeling
David Sharp	Los Alamos National Laboratory	Computational Science
Leonard Slatest	BNL/ITD	Parallel Computing
William Studier	BNL/Biology	Protein Structures
Marcelo Vazquez	BNL/Medical	Neuron Analysis
Nora Volkow	BNL/Directorate	Medical Imaging
Michael Weinert	BNL/Physics	Condensed Matter Physics
Torre Wenaus	BNL/Physics	High Energy and Nuclear Physics
Wei Zhu	SB/Applied Math. & Statistics	Statistics; Image Analysis

TABLE 3. STUDENT AFFILIATES

Name	**Affiliation**
Mihail Busu	SB/Applied Mathematics & Statistics mibusu@amirani.ams.sunysb.edu
Seung Yeon Cho	SB/Applied Mathematics & Statistics sycho@ams.sunysb.edu
Erwin George	SB/Applied Mathematics & Statistics egeorge@ams.sunysb.edu
Wei Guo	SB/Applied Mathematics & Statistics wguo@ams.sunysb.edu
John Leita	SB/Electrical & Computer Engineering j.leita@ic.sunysb.edu
Jing Luo	SB/Computer Science jingluo@cs.sunysb.edu
Andrea Marchese	SB/Applied Mathematics & Statistics marchese@ams.sunysb.edu
Natasa Stojic	SB/Physics nstojic@grad.physics.sunysb.edu
Haibin Su	SB/Material Science and Engineering Hsbsu@sun2.bnl.gov
Tom Welsh	SB/Computer Science tfwelsh@cs.sunysb.edu
Christina Weaver	SB/Applied Mathematics & Statistics cweaver@ams.sunysb.edu

4. RESEARCH

4.1 SciDAC: TSTT at CDIC

Scientific Discovery through Advanced Computing (SciDAC) is a five-year program to develop the Scientific Computing Software and Hardware Infrastructure needed to use terascale computers to advance DOE's research programs.

At Brookhaven we have formed an alliance with researchers at Argonne, Livermore, Oak Ridge, Pacific Northwest, Rensselaer, and Stony Brook, to develop technologies that enable scientists to use complex mesh and discretization strategies easily, within a single simulation on terascale computers.

The Terascale Simulation Tools and Technologies Center (TSTT) will provide interoperable tools to facilitate use of advanced mesh and discretization technologies. The Center will develop standardized interfaces to local mesh refinement and adaptive higher order discretization codes. Existing codes of the Center partners and new codes will be supported to create a plug and play capability, whereby an application user can experiment easily with alternative technologies. Insertion of these tools into targeted applications is part of the Center work plan.

CDIC work within TSTT will focus on climate model and accelerator applications, and on finite element and front tracking contributions to TSTT technology.

4.2 High Energy and Nuclear Physics

See Table 4 for a list of projects.

Modern accelerators generate vast quantities of data. These machines are operated for worldwide collaborations involving hundreds of institutions and thousands of researchers. New techniques for retrieval, visualization, and distribution of these data sets are required.

For BNL, our own Relativistic Heavy Ion Collider as well as our participation in the Atlas detector at CERN, dictate that these issues be addressed. We will also address simulation issues in the design, operation, and control of accelerators. See Table 4.

TABLE 4. HIGH ENERGY AND NUCLEAR PHYSICS PROJECTS

Project Title	Key Personnel	Comments
Muon Collider Target	R. Samulyak, W. Oh, J. Glimm,	Simulation of the Muon Collider Target
Accelerator Design	A. Luccio, N. D'Imperio R. Samulyak	Parallelization of Design, Control, and Optimization Codes
Computing with QCDOC	J. Davenport, N. Stojic, I. Campos, N. D'Imperio	Development and implementation of simulation algorithms
RHIC Data Management	T. Robertazzi, J. Leita	Globus for Data Management

4.2.1 Muon Collider Target

Introduction: Muon Collider. In order to understand the fundamental structure of matter and energy, an advance in the energy frontier of particle accelerators is required. Advances in high energy particle physics are paced by advances in accelerator facilities. The majority of contemporary high-energy physics experiments utilize colliders. A study group was organized at Brookhaven National Laboratory to explore the feasibility of a high energy, high luminosity Muon-Muon Collider. For more information visit the home pages of the Muon Collider Group (http://www.cap.bnl.gov/mumu/mu_home_page.html) and the Center For Accelerator Physics (http://www.nsls.bnl.gov/AccTest/capfiles/CAP.html). Such a collider offers the advantages of greatly increased particle energies over traditional electron-positron machines (linear colliders). However, several challenging technological problems remain to be solved. One of the most important is to create an effective target able to generate the high-flux muon beam. The need to operate high atomic number material targets, in particle accelerators that will be able to withstand intense thermal shock, has led to the exploration of free liquid jets as potential target candidates for the proposed Muon Collider. The target will be designed as a pulsed jet of mercury (high Z-liquid) interacting in a strong magnetic field with a high energy proton beam. A global numerical study and an analytical estimation of the behavior of the mercury jet under such energy deposition is the main goal of the present research. The numerical simulation was performed using the computer code FronTier, developed to model the propagation of three-dimensional fluid jets.

Muon Collider Target. The Muon Collider target is shown schematically below (Fig. 1). It will contain a series of mercury jet pulses of about 1 cm in diameter and 30 cm in length. Each pulse will be shot at a velocity of 30-35 m/s into a 20 Tesla magnetic field at a small angle to the axis of the magnetic field. When the jet reaches the center of the magnet it will be hit with a 2 ns proton pulse. The proton pulse will deposit about 100 J/g of energy in the mercury.

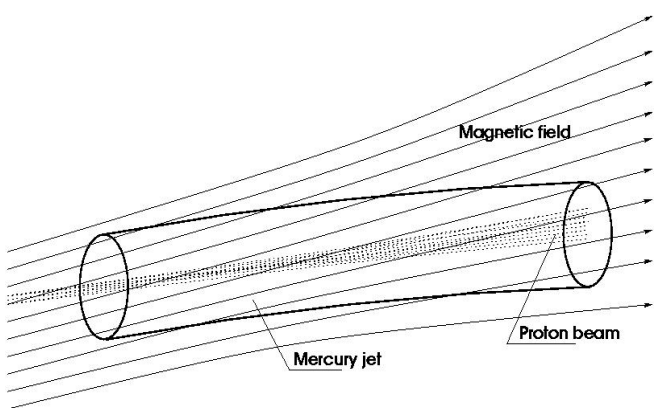

Figure 1

Numerical Simulations. Numerical simulations of the target evolution were performed using the computer code FronTier. The FronTier code is based on front tracking, a numerical method which allows the computation of dynamically evolving interfaces. In general, interfaces can model many types of discontinuities in a medium such as shock waves in gas dynamics, boundaries between fluid-gas states, different fluids or their different phases in fluid dynamics, component boundaries in solid dynamics, etc. The unifying feature of the models solvable by the method of front tracking

is their ability to represent the governing system of equations in the form of conservation laws, and to determine the evolution of the discontinuity through the solution of the Riemann problems. This method often does not require highly refined grids and it has no numerical diffusion across a tracked interface. For more information on the method of front tracking and FronTier code, visit the FronTier homepage (http://www.ams.sunysb.edu/~shock/FTdoc/FTmain.html).

To model numerically the behavior of mercury jets in magnetic fields, an MHD extension of the FronTier code has been developed (see Section 4.4.1, Magnetohydrodynamics of Multifluid Systems, for more information). We have shown that strong nonuniform magnetic fields cause significant distortion of conducting liquid metal jets. To avoid this phenomenon, the nozzle of the mercury supply system will be placed inside the target solenoid system (see Fig. 2). In addition, a magnetic system design which includes the superconducting, resistive and matching solenoids provides a nearly uniform magnetic field distribution along the path of the mercury target. We have shown [1] that the mercury target will not be significantly distorted during its motion in such a solenoid system. However, the magnetic forces speed up the natural instability and the pinchoff of the mercury jet.

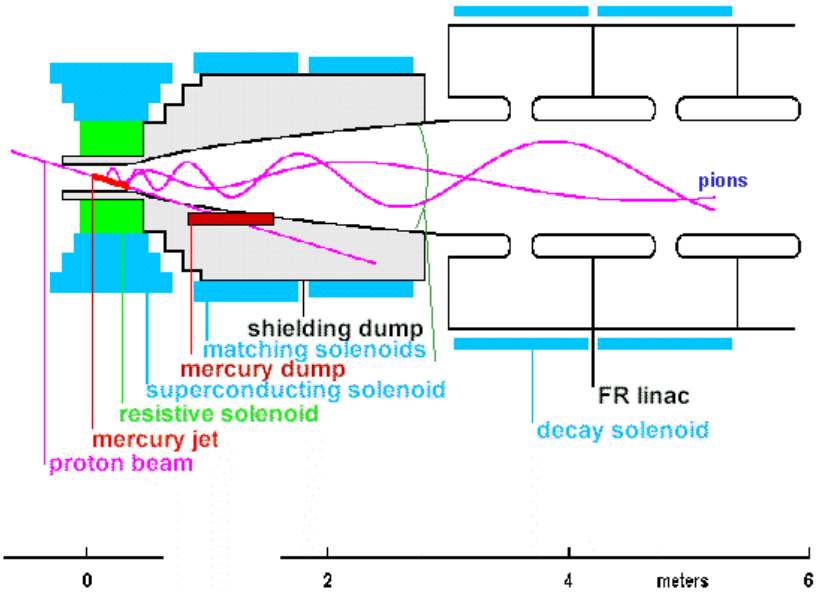

Figure 2. Target solenoid system design.

Numerical simulation also aids in understanding the behavior of the target under the influence of the proton beam, and in estimating the evolution of the pressure waves, surface instabilities, and the state of the target after the interaction with protons. The evolution of the mercury target after the interaction with a proton pulse is shown below (Fig. 3). According to the simulation results, the mercury target will be broken into a set of droplets due to the proton energy deposition and the radial velocity of the jet surface before the droplet formation is in the range 20 - 60 m/sec.

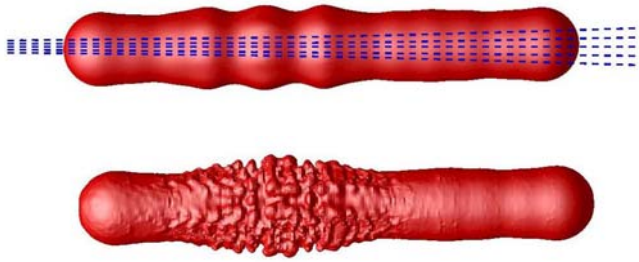

Figure 3. The first image depicts the initial state of the target. The second image corresponds to the state of the target after the interaction with the proton pulses at the time 35 microseconds. The energy deposition profile is Gaussian.

To model accurately the interaction of the mercury target with proton pulses, a tabulated equation of state for mercury was created in a wide temperature - pressure domain which includes the liquid-vapor phase transition and the critical point. The necessary data describing thermodynamic properties of mercury in such a domain were obtained courtesy of T. Trucano of Sandia National Laboratory. A simplified, two-phase, analytic equation of state containing liquid-vapor phase transition was also developed, assuming an isentropic regime of the flow. In the future, we shall improve our numerical simulations using the FronTier-MHD code and both the tabulated and analytic two-phase equations of state.

References

1. Ozaki, S., Palmer, R., Zisman, M., and Gallardo, J., Editors. Feasibility Study-II of a Muon-Based Neutrino Source.. BNL-52623, June 2001; available at http://www.cap.bnl.gov/mumu/studyii/FS2-report.html.
2. Kirk, H., et al. Target studies with BNL E951 at the AGS. In *Proc. Particles and Accelerators 2001*, June 18-22, 2001, Chicago, IL.
3. Glimm, J., Kirk, H., Li, X. L., Pinezich, J., Samulyak, R., and Simos, N. Simulation of 3D fluid jets with application to the Muon Collider target design. In *Advances of Fluid Mechanics III* (Editors: M. Rahman, C. Brebbia), pp. 191-200, WIT Press, 2000.

4.2.2 Accelerator Design

Advanced Computing for 21st Century Accelerator Science and Technology within the DOE initiative **Scientific Discovery through Advanced Computing.**

This project is a multi-institutional research and development effort involving four national laboratories and four universities. It will develop a comprehensive terascale Accelerator Simulation Environment (ASE) whose components will enable accelerator physicists and engineers to address a broad range of important issues in next generation accelerators and their design. These issues, which could not be addressed without terascale resources, range from the high accuracy requirement for modeling complex accelerating structures with tight tolerances, to the high resolution requirement needed for simulation beam halos in high intensity linear and circular accelerators.

The following target applications are relevant to BNL accelerator simulation research:

- 3D strong-strong beam-beam simulation for improving performance of RHIC.
- Simulation of beam dynamics and collective beam instabilities in collider injectors such as the AGS and BNL Booster accelerator.
- Simulation of beam dynamics in the Next Linear Collider (NLC) beamline.

The CDIC/BNL research activity is in the area of modeling the beam dynamics and collective beam instabilities in high intensity accelerators by means of Particle-in-Cell (PIC) codes. In particular, we have been developing computational methods and software for the particle beam – wake field simulations in PIC codes for both linear and circular accelerators.

The electromagnetic interaction of an intense charged particle beam with its vacuum chamber surroundings in an accelerator plays an important role for the beam dynamics and collective beam instabilities. Wake fields, generated by a moving particle in the accelerator's pipe and objects such as RF cavities, bellows, stripline monitors etc., affect the motion of all particles in the tail part of the beam. This effect is usually of the same order of magnitude or larger than the influence of the space charge. While the space charge forces approach zero in the ultrarelativistic limit, the wake field effect remains finite for an ultrarelativistic beam due to resistivity of the accelerator walls and non-smoothness of the chamber (existence of RF cavities, bellows etc.). The effect of wake fields is an increasingly important issue since operating regimes are continually moving towards higher currents and smaller bunches. If not properly dealt with, wake fields can lead to collective beam instabilities that limit a machine's performance. Therefore, an accurate numerical modeling of wake fields and their interaction with the beam is important for the development of high performance accelerator codes.

The main aim of this research is modeling of wake fields/impedances of resistive accelerator walls, RF cavities, bellows etc. and their influence on the beam dynamics in circular and linear accelerator. Since the impedance is the Fourier transformation of the corresponding wake field, the description of the wake force in terms of wake functions in the time domain is identical to its description in terms of impedances in the frequency domain.

Modeling of the particle beam – electromagnetic wake field interaction has been implemented in several PIC codes. In the current codes, however, the kick resulting from the total impedance budget for the whole ring is applied to all particles at once. In the present project, we make the first attempt to distribute the wake field influence along the ring by calculating the resistive wall impedance and impedances of non-smooth elements of the ring.

The wake function/impedance of resistive cylindrical accelerator walls can be calculated analytically [1] using a representation of the beam as a ring possessing multipole moments with $\cos m\theta$ distribution. The resulting longitudinal and transverse force components on a test charge in the *j*th macroparticle is

$$\left(\int ds \quad F_{\parallel}\right)_j = -\frac{1}{2} N_j e^2 W_0^{'}(0-) - \sum_{k>j} N_k e^2 \left[W_0^{'}(z_j - z_k) + y_j y_k W_1^{'}(z_j - z_k) \right] ,$$

$$\left(\int ds \quad F_{\perp}\right)_j = -\sum_{k>j} N_k e^2 y_k W_1(z_j - z_k) ,$$

where $N_j e$ is the macroparticle charge, y_j and z_j are the transverse displacement and the longitudinal position of the macroparticle, respectively, and $W_m, W_m^{'}$ are the transverse and longitudinal wake functions of *m*th momentum. We considered here simultaneously only $m = 0$ and $m = 1$ wake fields and ignored the higher order terms.

The wake functions and impedances for non-smooth accelerator elements will be computed numerically using the codes MAFIA and Omega3P. We will combine the wake field – particle beam interaction with the space charge forces using the operator splitting technique. The space charge forces will be obtained by solving the Poisson problem with open boundary conditions. This approach excludes the influence of image charges on the accelerator walls taken into account in the wake field calculations. We will implement modules for the beam - wake field interaction in both IMPACT, a linear accelerator code, and a code for circular accelerators. The code for circular accelerators will be developed based on MaryLie, a magnetic optics code with the Lie algebraic treatment of external forces.

Collaboration with BNL's Collider-Accelerator Department Computational Project

We have been working in collaboration with computational physicists of the BNL Collider Accelerator Department. The primary objective is to develop Orbit, a PIC code for circular accelerators, which models collective beam effects through a "two-and-half" dimensional space charge problem [2,3]. The code is being applied for modeling of the beam dynamics and beam instabilities in AGS and SNS storage ring.

In the Orbit Code, a "herd" of randomly generated macro particles is pushed through a lattice represented by a sequence of maps. Once all particles have reached a certain location, the space charge nodes, their charge density is calculated by binning to a grid, and the potential Φ is found by solving the Poisson equation with the perfectly conducting wall boundary conditions:

$$\Delta\Phi = -\frac{1}{\varepsilon}\rho, \quad \rho_{wall}(x, y, z) = 0.$$

Space charge force components (with coefficients to account for both the electrostatic and magnetic action) are calculated as derivatives of the potential and applied to each macro particle in the transverse direction. In a ring with long longitudinal bunches, the transverse motion can be uncoupled from the longitudinal one and the Poisson problem can be solved in parallel in many longitudinal beam slices. Figure 4 shows a "frozen" beam at a given time in a simple FODO channel in the moment of the space charge calculation.

Several Poisson solvers have been implemented in the Orbit code. A direct integration of particle forces (brute force) gives a very transparent solution. However, this method has limitations: (i) arbitrary treatment of poles arising from the field points accidentally coincident with source points,

(ii) difficulties, including images on walls in the space charge calculation, (iii) length of execution. The Orbit code has the capability to solve the space charge problem by the FFT method. The method needs a grid twice the size of the beam to avoid aliases. Parallel differential solvers for the Poisson problem were implemented recently in the code based on the LU decomposition and an iterative (SOR) technique. The x- component of the space charge force calculated using different approaches is depicted in Fig. 5.

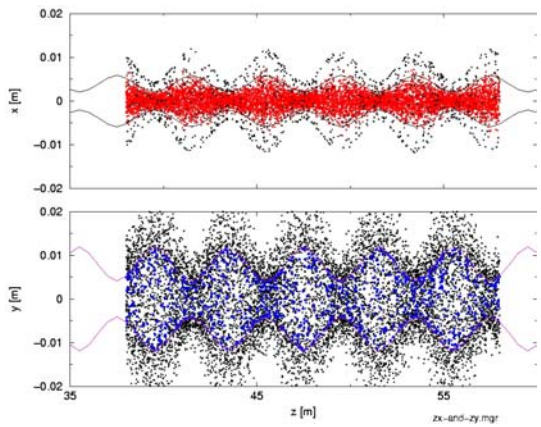

Figure 4. Frozen matched beam.
and the beam Envelope is also shown.

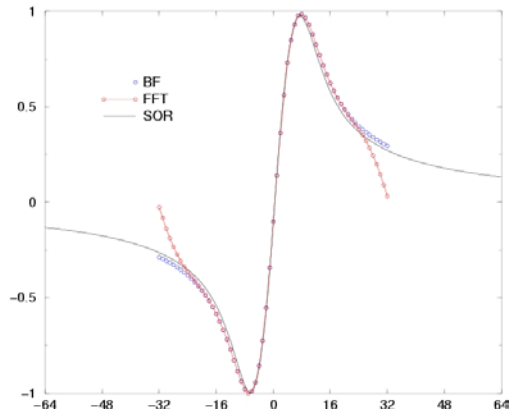

Figure 5. Space charge force component
calculated using different methods.

The performance of different parallel Poisson solvers was tested using the Brookhaven Galaxy Cluster. Linear performance scaling was observed with 32 processors.

References

1. Chao, A. *Physics of Collective Beam Instabilities in High Energy Accelerators*, John Wiley & Sons, Inc., New York, 1993.
2. Luccio, A. and D' Imperio, N. Tracking in accelerators with parallel computing. 6th International Accelerator Physics Conference, Darmstadt, Germany, Sept. 11-14, 2000.
3. Luccio, A., D'Imperio, N., and Samulyak, R. Solving by parallel computing the Poisson problem for high intensity beams in circular accelerators. In Proc. of the Conf. *Particles and Accelerators 2001*, June 18-22, 2001, Chicago, IL

4.2.3 Computing with the QCDOC

QCDOC is a special purpose machine being designed by the high energy physics community for lattice quantum chromodynamics calculations. We propose to investigate its potential and cost effectiveness for other types of calculations. The machine is under design by a team from Columbia University, IBM, and the Riken BNL Research Center. Funding for a portion of the design work has been obtained as part of a SciDAC project of DOE's Office of High Energy and

Nuclear Physics. This group is using experience gained with QCDSP - a massively parallel computer made up of 12,288 digital signal processors which is located at Brookhaven and achieves a peak speed of 0.6 teraflops. The new machine (QCDOC, quantum chromodynamics on chip) will utilize an IBM PowerPC RISC processor with a 1 GHz floating point unit, 256 Mb of memory, and communication hardware all on a single chip. This chip will consume only one to two watts of power, an important consideration for a machine which will be made up of roughly 10,000 nodes. In collaboration with Brookhaven's Information Technology Division, we plan to acquire a motherboard and several current generation PowerPC embedded processors. We will test the configuration on several of our software packages, including the Maxwell equation solver we have developed for the study of photonic crystals. The main purpose of this study is to determine the feasibility and cost effectiveness of this architecture for use in a variety of problems unrelated to QCD.

4.2.4 RHIC Data Management

Computing Grids

A "grid" is a network of distributed computers that provides computing power in a simple and seamless manner upon demand in much the same way electric power grids routinely provide electric power. At BNL we are using the Globus grid software package in support of BNL science. Globus (www.globus.org) is a forward looking software package for distributed computing. It provides a middleware infrastructure for distributed computing and resource sharing. It includes security features. The concepts behind Globus are described in the book *The Grid* by Ian Forster and Carl Kesselman.

At BNL, Globus was installed on a number of machines and its performance is being tested. We are examining ways in which Globus can be used to enhance collaborative science at BNL and other sites.

4.3 BIOLOGICAL AND ENVIRONMENTAL RESEARCH

See Table 5 for a list of projects.

Brookhaven has strong experimental programs in atmospheric chemistry and in brain physiology. Both programs are data intensive. The data challenge is not only data quantity, but especially data understanding. We are developing new tools for the visual and statistical analysis of data and we are actively using these tools in the search for scientific understanding. The tools themselves are based on advanced and novel concepts in computer science and statistics.

TABLE 5. BIOLOGICAL AND ENVIRONMENTAL RESEARCH PROJECTS		
Project Title	**Key Personnel**	**Comments**
Evolution of Atmospheric Aerosols	M. Busu, J. Davenport, D. Imre, K. Mueller, W. Zhu	Atmospheric Chemistry
Parallel Implementation of an Atmospheric Aerosol Model	C. Benkovitz, J. Luo, H. Qin, S. Schwartz, L. Slatest	Aerosols in Climate
Ocean Modeling with Automatic Mesh Refinement	M.-N. Kim, W. Oh	Western Boundary Currents
Analysis and Visualization of Human Brain Function	K. Mueller, N. Volkow, W. Zhu	Drug Abuse Alzheimer's Disease
Neuron Imaging	B. Lindquist, M. Vazquez, C. Weaver	Edge Detection Image Analysis
Cardiac Electrophysiology	S. Evans, F. Fenton, J. Glimm, H. Hastings, W. Oh,	Fibrillation, Sudden Cardiac Death

4.3.1 Evolution of Atmospheric Aerosols

Introduction

The purpose of this project is to develop data processing tools for the analysis and classification of atmospheric aerosol particles. Atmospheric aerosols play an important role both in climate and public health. Current research is focused on the geographical distribution, particle size, and chemical composition of aerosols. Recently, researchers in the Atmospheric Sciences Division at Brookhaven have developed a field deployable mass spectrometer capable of measuring these quantities with unprecedented precision. This instrument has already produced almost a terabyte of field data and future missions will yield much more. The analysis task is daunting, and efficient methods are urgently needed to harvest the information that is lurking in the vast amount of data.

Large-scale BNL data collections have reached sizes for which straightforward visualization techniques are starting to fail. In atmospheric science, automated classification methods, incorporating domain expert knowledge, have been developed to filter the particle data based on spectral features, mass, and other parameters. Clearly, the accuracy of the classification process is crucial, as potentially incomplete classification rules can lead to both false positives and false negatives. In that spirit, our proposed approach will couple a powerful classification engine with an intuitive and responsive interface to fine-tune the underlying classification model.

Classification Engine

We will classify aerosols based on their mass spectra through an iteration of expert-machine interaction. We will also implement automatic procedures to elucidate the structure of the compounds. To study the evolution of aerosols and the changes in atmospheric composition versus time, we will perform univariate and multivariate time series analyses to unravel trends and patterns. There are four major types of classification methods— clustering analysis, discriminant

analysis, neural network, and classification tree. Each has been adopted in the classification of mass spectra previously, often with success.

We propose to build a flexible human-machine interface such that large numbers of mass spectra can be classified automatically at the initial stage. The "mean spectra" of the classes will be produced and inspected by the field experts. The classifier will be modified if necessary and a second round of classification will ensue. The iteration continues until satisfactory results are achieved. Our experience suggests that a multivariate binary classifier based on expert input can be highly successful.

Different classifiers tend to focus on different aspects/features of the data, and thus what is optimal for one problem may be inefficient for another. Our literature search revealed that although a wide range of classification methods had been adopted in mass spectra analysis, no systematic comparison of the approaches exists. To fill this void and provide guidance for future work, we will integrate typical classifiers of each type into our study and compare their performance. These classifiers are widely available in professional mathematics/statistics software. For example, the classification tree is available through S-Plus (tree, rpart), SAS (enterprise miner), and CART. Various neural network procedures are also available through MATLAB Neural Network Toolbox.

A novel idea in modern classification, and that pursued here, is to combine different classifiers to achieve better discrimination. We will try stacking [8] and related methods. Committees can help when different learners have complementary strength for a given task. Stacking works like a committee where the membership (class) of each data point is determined by a majority vote; that is, the data point will go to the class with the most votes from the classifiers. Other methods differ from stacking in that only one classifier is of concern, which is then improved through bootstrap sampling.

For efficient classification and automatic structure elucidation, we have been constructing a molecule library where the signature profile and class membership of each molecule is established. The molecules are classified along a natural chemical classification tree with two categories—organic and inorganic, at the initial node. Subsequently, the organics are further divided into classes of carboxylic acids, aldehydes, ketones, alkenes, alkanes, aromatics, etc. We begin with the spectra of known molecules (NIST library or lab-generated).

We want to determine the composition of a given aerosol based on its spectrum. We allow for measurement errors by solving a linear system including unknown errors through the method of least squares.

Visualization

Visualization is a powerful and very effective paradigm to integrate and convey a massive amount of information on a single display. By using volume visualization techniques [1] – [6], scientists can quickly grasp the results of their computational simulations, which often give rise to data files of many GB.

Our system will offer a number of visualization modules, distinguished by their level of data abstraction. All modules will employ a unified interface, a flight simulator model, to navigate the data in 3D virtual space.

Traditional Plotting with Extended Functionality. The first module provides traditional three-dimensional plotting, but with added functionality. Figure 6 shows some screen shots of preliminary work that we have performed. The current software retrieves a portion of the particle data from disk and bins the data into a 2D grid formed by mass and acquisition time. A navigable plot is then constructed that shows the data as a 3D landscape shaped by the particle distribution. Tools are already available that allow the user to refine or reduce the time/mass resolution of the binning and in this way reveal more or less detail in the landscape.

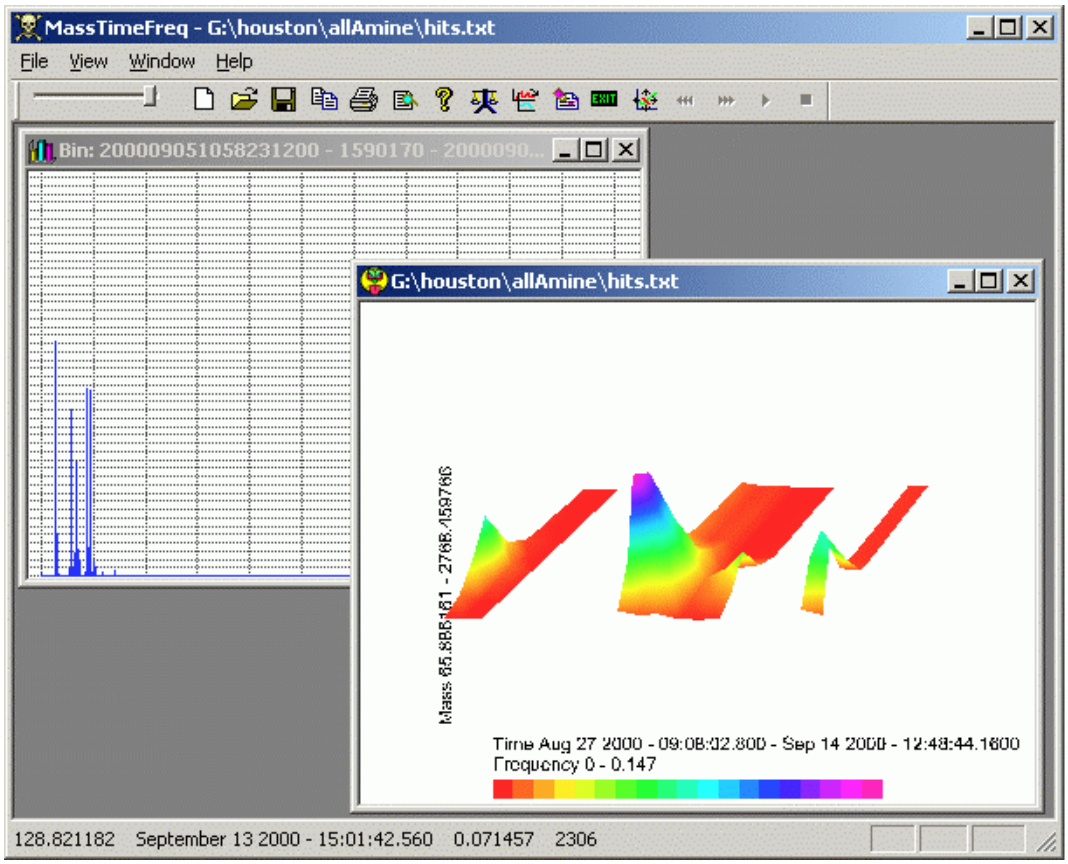

Figure 6. Screenshot of our plotting tool. The frame on the right shows the mass-time frequency plot of a class of particles. The user is free to rotate and zoom the landscape into any orientation and resolution. The frame on the left shows the spectrum of a selected particle. The user can click on one of the times and the tool will loop through all particles (or their cluster representatives) in that bin and display their spectra. The user is free to interrupt the animation at any time and remove a particle from the bin should the spectrum not fit the scientist's class criteria. This information is then used to train and refine the classification engine.

Abstracted Visualization of Classification Results with Classification Steering. The second module expands on the first in less traditional ways. It seeks to capture and convey the composition of the set of particles that were assigned to certain classes. Presently, scientists cross-validate the outcome of a classification by browsing through the spectra files to look for misclassified particles. Considering the huge number of particles that are potentially assigned into a class, this can be an almost insurmountable task. Furthermore, the particles that were *not* assigned into the class are false negatives which also have to be identified in a potentially even larger set.

The idea of the proposed approach is to cluster the data in each class (and its complement) and visualize these clusters in the form of a 3D terrain landscape of mountains and valleys. The clustering algorithm would use subspectra or a multi-resolution spectra representation, augmented by information on mass, location, acquisition time to cluster the particles in a higher-dimensional (n-) space. Based on the relative distance between these clusters, a 2D graph can be constructed where nearby nodes also have small distances in the hyperspace. We propose to use a force-directed algorithm [7] to efficiently find this graph embedding. In this type of algorithm each node pair is connected by a spring whose length encodes the desired distance.

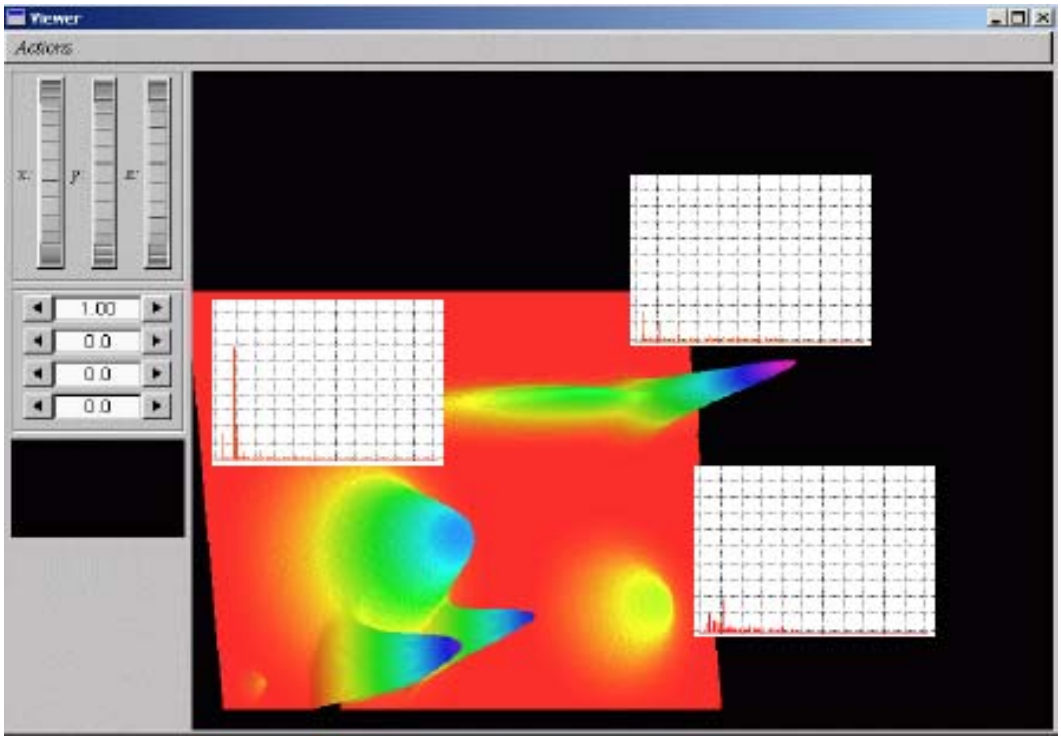

Figure 7. A preliminary tool for plotting the clusters in a selected bin or class. The proximity of the elevations code according to a similarity measure that can be any weighted function of mass, time of acquisition and particle spectrum. The spread of an elevation encodes the variation of the particles with respect to the similarity measure. The height encodes the number of particles fitting that measure. Billboards with a representative (for example average) spectrum will be erected next to the corresponding elevations (we have manually pasted these here for now).

Once the landscape has been constructed, representative spectra of each mountain can be erected on its top. The scientist will also be able to zoom into a certain mountain range, reveal more detail and finally drill down to the individual particles. Clicking on one of these "nuggets" will show its spectra as well as other recorded parameters. Then, when the scientist decides that this particle should belong to a different cluster or even class, a simple removal or dragging of the particle or set of particles into the more appropriate position (in another elevation) will trigger an automatic refinement of the classification rules and a retraining of the classification engine. This system will reflect this change in the classification rules by periodic updates of the landscape, possibly highlighting the modifications from the last configuration. In Fig. 7 a preliminary tool is shown for the visualization of the particle clusters within one bin or an entire class.

References

1. Meissner, M., Huang, J., Bartz, D., Mueller, K., and Crawfis, R. A practical comparison of popular volume rendering algorithms. *2000 Symposium on Volume Rendering*, pp. 81-90, Salt Lake City, October 2000.
2. Mueller, K., Shareef, N., Huang, J., and Crawfis, R. High-quality splatting on rectilinear grids with efficient culling of occluded voxels. IEEE Trans. Vis. and Comp. Graph. **5(2)**, 116-134 (1999).
3. Mueller, K., M`ller, T., Swan, J.E., Crawfis, R., Shareef, N., and Yagel, R. Splatting errors and antialiasing. IEEE Trans. Vis. and Comp. Graph. **4(2)**, 178-191 (1998).
4. Levoy, M. Efficient ray tracing of volume data. Acm Trans. Comp. Graph. **9(3)**, 245-261 (1990).
5. Drebin, R., Carpenter, L., and Hanaran, P. Volume rendering. *Proc. SIGGRAPH '88*, pp. 65-74, 1988.
6. Kaufman, A. *Volume Visualization*. IEEE Computer Society Press, Los Alamitos, CA, 1991.
7. Fruchterman, T. and Reingold, E. Graph drawing by force-directed placement. Software-Practice and Experience **21(11)**, 1129-1164 (1991).
8. Wolpert, D. Stacked generalization. Neural Networks **5**, 241-259 (1992).

4.3.2 Parallel Implementation of an Atmospheric Aerosol Model

We have constructed a parallel version of an atmospheric aerosol model to describe sulfate and SO_2 concentrations over North America and Europe. See [1].

The model is a 3D Eulerian transport code which also includes chemical reactions among 25 different species. Pollutants are transported by (externally supplied) velocity fields and allowed to react within each cell according to the local temperature (also externally supplied). The hemispherical model has one degree resolution, so the grid has 360 points for longitude and 81 for latitude, with 27 vertical levels. It was parallelized using MPI on the Cray T3E at the San Diego Supercomputer Center.

The relatively large size of the grid leads to a large array of data points, e.g., 360x81x27x25 = 19,683,000 numbers or approximately 150 megabytes in double precision. This is for a single time step. This problem is an ideal candidate for parallelization since much of the processor time is

spent on the chemistry within a single cell, and relatively little time is spent on communication among cells. In addition, the communication which does occur is between a cell and its nearest neighbors. Our domain-decomposition parallel algorithm has demonstrated excellent scalability (near 100%) when the problem size is increased. Therefore, we can carry out simulations efficiently with much higher resolution (>>360x81x27), which are not feasible without parallelization, on hundreds or thousands of processors.

References

1. Benkovitz, C.M. and Schwartz, S.E. Evaluation of modeled sulfate and SO_2 over North America and Europe for four seasonal months in 1986-87. J. Geophys. Res. **102,** 25305-25338 (1997).

4.3.3 Ocean Modeling with Automatic Mesh Refinement

Global climate modeling is one of the grand challenges of computational science. The ocean influences the global climate system, because the ocean and atmosphere are strongly coupled by fluxes of moisture, energy and momentum. The atmosphere is conditioned by the ocean at all time scales ranging through daily, seasonal, and climatological periods. The ocean is in turn conditioned by the atmosphere, but its response time is longer, and this inertia tends to smooth out high frequency oscillations in the atmosphere.

Modeling ocean dynamics is less complicated physically but more demanding computationally than modeling atmosphere dynamics. In the atmosphere there are many physical processes and variables that must be taken into account. The ocean, in contrast, is relatively simple. However, ocean modeling is more expensive, requiring longer iterations at finer resolution because of the broad range of temporal and spatial scales which encompass the relevant dynamics. The spatial scales in the ocean range from the diameter of mesoscale eddies, a few tens of kilometers, to the dimensions of ocean basins. In contrast, the scale of atmospheric mesoscale eddies is on the order of 1000 kilometers. The time scales in the ocean range from the two- to three-week time scale of mesoscale eddies to the 1000-year time scale for ventilation of the deep Pacific, which is much larger than the longest atmospheric time scale. Therefore, the use of modern numerical tools and technologies such as adaptive mesh refinement and load balancing in the global ocean modeling codes will help to achieve long time simulations at desirable spatial resolution (see [3] for some simulation results).

The Parallel Ocean Program (POP) was developed at Los Alamos National Laboratory under the sponsorship of the Department of Energy's CHAMMP program, which brought massively parallel computers to the realm of climate modeling. POP has been used to perform high resolution ocean simulations, using parallel supercomputers at the Los Alamos Advanced Computing Laboratory (ACL).

The MICOM group at Los Alamos National Laboratory collaborates with the RSMAS Division of Meteorology and Physical Oceanography at the University of Miami. It is engaged in research and graduate instruction using the Miami Isopycnic Coordinate Ocean Model as the principal tool. Simplified versions of the Miami model have been used for 15 years in a number of process studies. Large-scale applications of the fully configured model now include simulations of individual basin circulations and the global ocean. Support for this work is provided by the U.S. Department of Energy, the National Science Foundation, and the Office of Naval Research.

The theoretical background of the POP is the Bryan-Cox model [1]. The POP code uses a polar longitude-latitude type grid on a sphere with the North pole shifted to the center of either North America or Asia [2]. This configuration makes it possible to avoid the so-called "North pole problem" associated with the convergence of meridians at the North pole (the shifted South pole remains in Antarctica and therefore does not cause singularities).

The equations of the Bryan-Cox model are split into two parts: 2D barotropic equations (momentum equations) describing the vertically-averaged flow and a set of 3D baroclinic equations describing the evolution of temperature, salinity and deviations of the horizontal velocities from the vertically-averaged flow. The baroclinic equations are solved explicitly using finite differences on a polar longitude-latitude grid. The barotropic equations are elliptic (the pressure equation) and require solvers based on iterative techniques.

To improve the resolution of calculations in domains of interest (Western boundary currents, the Gulf of Mexico, the Gulf Stream, etc.) without increasing the total computational time, an adaptive mesh refinement technique is necessary. We have been working on the development of the adaptive mesh refinement technology for the longitude-latitude polar grid used in the POP code (see Fig. 8). The work is based on the algorithmic technologies and the AMR library of the Overture package, namely, on the elliptic solutions algorithms for adaptively refined meshes such as the asynchronous fast adaptive composite method and the corresponding technologies for hyperbolic solvers. Overture numerical algorithms [4] are based on high resolution finite difference and finite volume methods applied on overlapping grids. Overlapping grids are sets of structured curvilinear grids that overlap where they meet, and together can accurately represent the computational geometry for complex geometric configurations. The Overture grid generation research has focused on algorithms for generating structured component grids from a variety of geometrical representations, on automatic algorithms for overlapping these grids to form the computational grid-structure for complex simulations, and on adaptive mesh strategies that can be used on overlapping grids [5]. This software infrastructure is designed to effectively exploit architectures of modern distributed memory supercomputers.

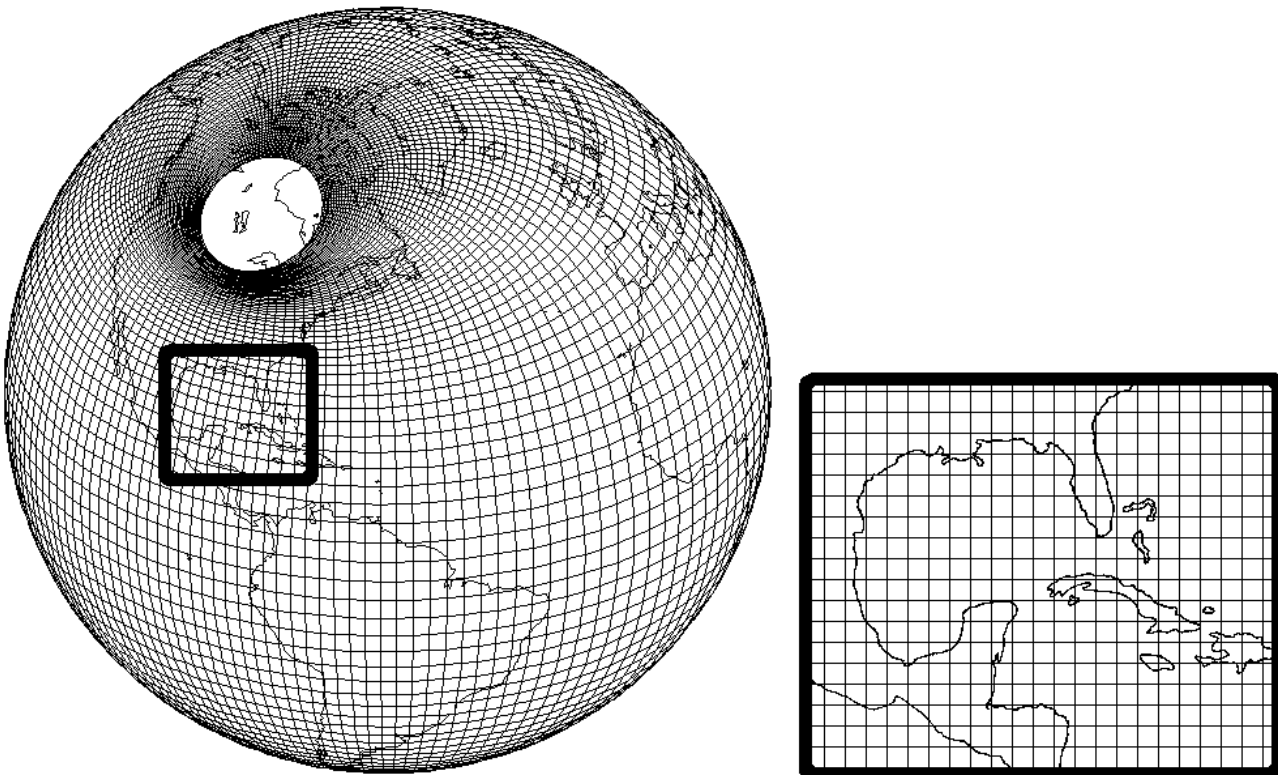

Figure 8. Left: Typical POP grid for ocean modeling simulation. Right: Schematic illustration of a refined grid path.

References

1. Smith, R.D., Dukowicz, J.K., and Malone, R. C. Parallel ocean general circulation modeling. Physica D **60**, 38 – 61 (1992).
2. Smith, R.D., Kortas, S., and Meltz, B. Curvilinear coordinates for global ocean models. Los Alamos National Laboratory Report, 1995, LAUR-95-1146.
3. Smith, R.D., Maltrud, M.E., Bryan, F.O., and Hecht, M.W. Numerical simulation of the North Atlantic Ocean at $1/10^0$. J. Physical Oceanography **30**, 1532 – 1561 (2000).
4. Henshaw, W.D. A primer for writing PDE solvers with Overture. Lawrence Livermore National Laboratory Report, UCRL-MA-132233 (available at http://www.llnl.gov/casc/Overture/henshaw/documentation/primerHTML/primerHTML.html.
5. Brown, D.L., and Henshaw, W.D. Adaptive mesh refinement routines for Overture. Lawrence Livermore National Lab. Report, UCRL-MA-140918 (available at http://www.llnl.gov/casc/Overture/henshaw/documentation/amrHTML/amrHTML.html.

4.3.4 Analysis and Visualization of Human Brain Function

Overview

A major goal of statistical analysis in brain research is to extract from a vast amount of brain image data (PET and fMRI) the essential features that characterize brain functions. We are developing two novel approaches for such analysis: (1) Correlational and Path Analysis to uncover brain functional networks, and (2) coefficient of variation (CV) -- a measure of brain regional variability or heterogeneity. These analyses are being implemented in BrainMiner, an interactive brain functional analysis and visualization tool [2, 3, 6].

Correlational and Path Analysis. The correlations in the brain activity are calculated on the basis of predefined anatomical regions-of-interest (ROIs). The correlation coefficient is then employed to quantify similarity in response, for various regions during an experimental setting. To account for anatomical variability, each test subject's volumetric brain data is first transformed into a common anatomical coordinate system (the Talairach-Tournoux space). Statistical analyses of brain functional relationships include:

1. The correlational analyses including the Pearson product-moment correlations, the partial correlations and the canonical correlations.
2. Cluster analysis.
3. Principal component analysis (PCA) and the factor analysis (FA).
4. Path analysis and time series analysis.

Measuring Brain Functional Variability. It is well known that the two most essential statistics are mean and variance. In ROI-based brain functional studies, the analysis of the regional mean is a fully developed practice; however, the same is not true for the variability. We have developed the analysis for brain functional homogeneity/heterogeneity based on the coefficient of variation (CV). The CV is the ratio of the standard deviation to the mean. It is a measure of regional variability independent of the underlying measurement unit. We have applied such analysis to the comparison of brains affected by Alzheimer disease (AD) with normal ones [3]. Our results indicate that such analysis is providing new insight disease etiology and additional information on diagnosis.

The amount of statistical data can be enormous, and effective tools are essential if the brain researcher is to grasp and discover functional relationships quickly from the statistical data. BrainMiner has a powerful visualization tool that facilitates this task [1] – [6].

Viewing in 2D

We visualize statistical brain relationships in 2D by overlaying the statistically significant voxels on top of a high resolution MRI. The data is presented as 2D slices, either in flip mode or side-by side:

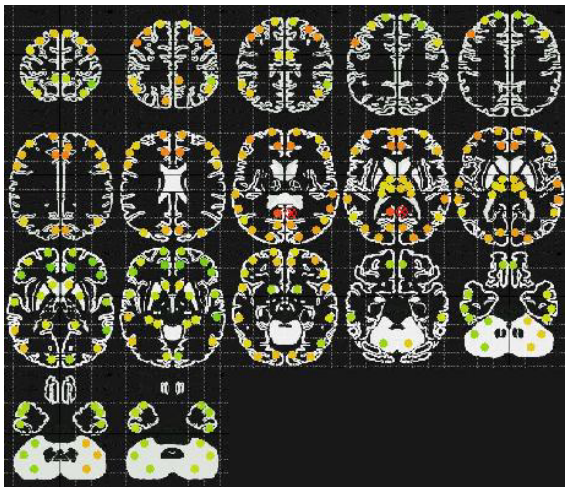

Figure 9. The circular ROIs are colored according to their correlation with respect to a root-ROI, marked by a red cross. The rainbow color scheme is used, where blue stands for highly negative correlation and red stands for highly positive correlation. Green and yellow stand for mildly negative and positive correlations, respectively. There are apparently no extremely strong correlations in this configuration.

The 2D approach works well as long as the axial dimension is not important. However, the decomposition of the dataset into 2D slices for visualizing 3D relationships becomes limiting when relationships are widely spread over the brain.

Viewing in 3D

To account for the problems with the 2D approach, we have developed (in addition to the 2D viewer) a 3D visualization interface that displays correlational data for each ROI along with an MRI volume and a digitized version of the Talairach atlas. Both can be sliced in 3 orthogonal directions and can be overlaid on each other. A basic view with a few ROIs is shown in Fig. 10, which shows the Graphical User Interface (GUI) of our newly developed 3D brain visualization software, along with a basic view of a small number of ROIs embedded into a cut-out area of a normalized/standardized MRI brain. Similar to the 2D viewer, the colors of the ROIs denote the strength of the correlational relationship, on a rainbow scale. The root ROI is colored in yellow. The GUI allows the user to slide the cutting planes up and down and back and forth, to rotate the volume, and to select certain brain surfaces, such as white matter, gray matter, and skull to be semi-transparently superimposed. The correlation thresholds can also be selected, and many more features are available.

The number of ROIs to be displayed, however, can become quite large (about 120-140), which poses challenging problems in the visualization task: in a space too crowded with statistically significant ROIs, it becomes very hard, if not impossible, for the user to tell the 3D positions of the individual ROIs. To overcome these difficulties, a number of techniques were investigated:

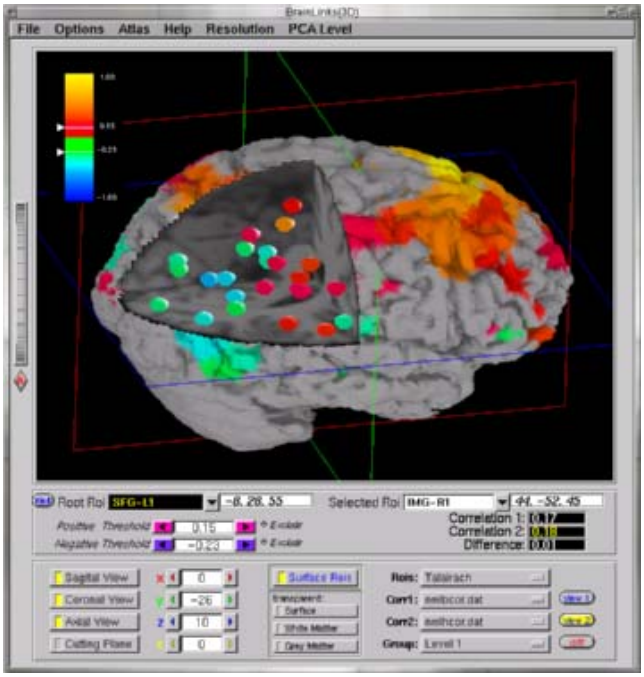

Figure 10. GUI for 3D brain visualization software.

- Superimposing a Talairach atlas slice that can be slid up and down the volume.
- A single light source placed above the volume in a fixed position, providing specular lighting cues for the height and depth of each ROI sphere (this can be seen in Fig. 10).
- Enhancing the ROIs by colored halos, where the colors code their height and depth on a rainbow color scheme. The ROIs are connected by iso-lines to the MRI volume cuts which suggests their position in 3D space (see Fig. 11).

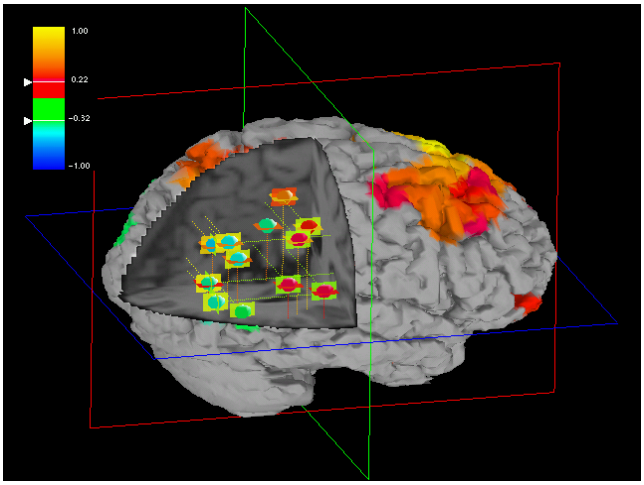

Figure 11. ROI halos, painted in colors corresponding to ROI height and depth (the rainbow color scheme is used). Dashed iso-height and iso-depth lines emanate from the ROIs and pierce the MRI volume slices at the ROI depth and height.

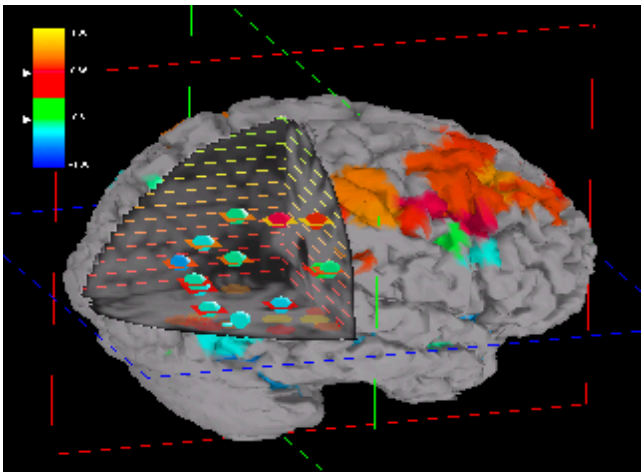

Figure 12. Iso-lines, with height coded into rainbow colors, are drawn onto two of the three orthogonal MRI volume cuts. The ROI halos are coded in height using the same color scheme. The ROI position with respect to the third MRI cut is suggested by shadows cast by the ROI spheres onto that plane.

- Projecting a colored grid onto the volume cuts, again encoding height and depth on a rainbow map. Colored shadows cast onto the exposed volume slices provide additional cues.
- Projecting the ROIs onto the brain iso-surface, such as white or gray matter, or skin.

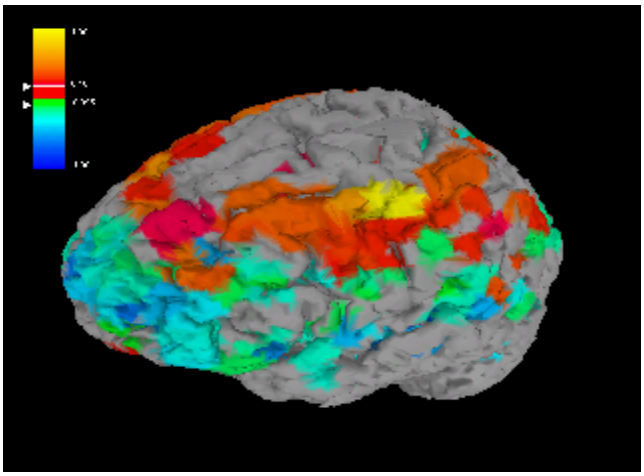

Figure 13. Since the ROIs are mostly located close to the brain surface, i.e. on the brain cortex, one can generate a comprehensive, EEG-like, view by projecting the ROIs onto the cortex surface and painting the projection in the correlation color.

- Grouping ROI networks into composite polygonal objects, which reduces the object complexity of the scene:

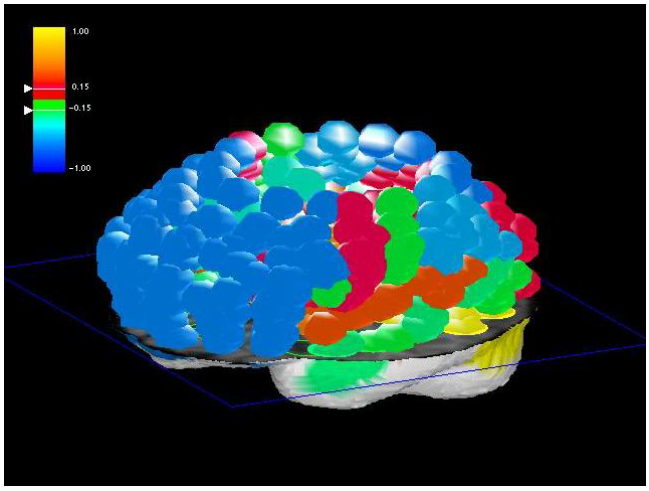

Figure 14. For now we simply increase the radius of the ROI spheres until they just touch. This approach is rather effective. For the future we plan to estimate the actual hull of a set of ROIs of similar brain function and display this hull as a polymesh.

Results

We are continuing our effort towards the discovery of brain functional networks. Intriguing results on brain functional homogeneity/heterogeneity have already been obtained. In analyzing the PET images of 35 subjects with AD and 35 age and gender matched healthy controls, we found that AD subjects showed a significantly larger CV for the whole cortex but had a significantly greater number of regions with increased homogeneity (decreased CV) than the controls. At the regional level the reductions in CV in the AD subjects were associated with reductions in metabolism that were most accentuated in the parieto-temporal cortex; except for the precuneus, which had a marked reduction in metabolism but increased CV. Discriminant analyses showed that the combination of variables that best discriminated between AD and controls were metabolism and CV in left parietal, CV in right precuneus and global CV (94% sensitivity and 100% specificity) [3].

Conclusions

BrainMiner allows many types of interactive explorations of brain functional image data. It is currently in use by a number of brain researchers and is being refined continuously.

References

1. Meade, J., Zhu, W., Mueller, K., Welsh, T., and Volkow, N. Brain Tableau – the 2D Brain Functional Visualization & Analysis Software. Medical Imaging 2001, 17-23 February, San Diego, CA.
2. Mueller, K., Welsh, T., Zhu, W., Meade, J., and Volkow, M. Brain Miner: A Visualization Tool for ROI-Based Discovery of Functional Relationships in the Human Brain, *New Paradigms in Information Visualization and Manipulation (NPIVM) 2000*, Washington DC, November, 2000.

3. Volkow, N., Zhu, W., Felder, C., Mueller, K., Welsh, t., Wang, G.-J., and De Leon, M. Changes in Brain Functional Homogeneity in Subjects with Alzheimer's Disease. Submitted to *Psychiatry Research* (2001).
4. Welsh, T., Mueller, K., Zhu, W., Meade, J., and Volkow, N. Brain Explorer– the 3D Brain Functional Visualization & Analysis Software. Medical Imaging 2001, 17-23 February, San Diego, CA.
5. Welsh, T., Mueller, K., Zhu, W., Volkow, N., and Meade, J. Graphical Strategies to Convey Functional Relationships in the Human Brain: A Case Study, Visualization '01, San Diego, October 2001.
6. Zhu, W., Volkow, N., Mueller, K., Welsh, T., and Meade, J. BrainMiner – An Interactive Tool for Brain Functional Connectivity Study. Human Brain Mapping 2001, Brighton, England, June 2001.

4.3.5 Neuron Imaging

We have developed a numerical method to analyze digitized microscope images of retinal explants and quantify neurite outgrowth [1]. Few parameters are required as input and limited user interaction are necessary to process an entire experiment of images. This eliminates fatigue-related errors, user-related bias, and labor-related data analysis limitations common to manual analysis, while providing results that are nearly identical. The method does not require stained images and it handles images of variable quality. The crux of our algorithm lies in tracing the ridges of a polar function of greyscale intensity, which represent the neurite outgrowth. See Fig. 15 to compare automatic and manual analysis of a retinal explant culture.

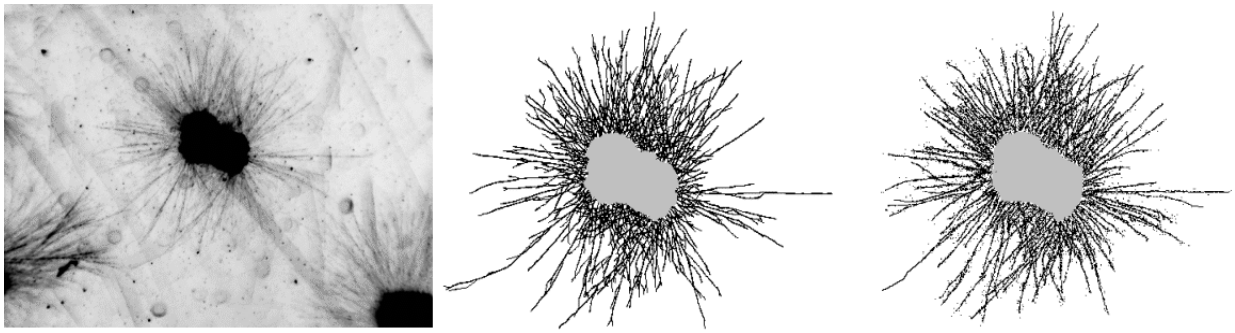

(a) (b) (c)

Figure 15. (a) A digitized microscope image, (b) automatic reconstruction of the explant region and neurite outgrowth, as compared with (c) manual analysis.

The method has been used to assess the neurotoxicity in retinal explants of high-energy iron particles found in galactic rays. This assessment will help to determine whether the heavy ion component of these rays might inactivate critical cells in the brain that could not be repaired or replaced by the body. As shown in Fig. 16, no significant neurotoxic effect is detected prior to 72 hours after explantation; at 72 hours significant growth reduction compared to the control is observed at even the smallest radiation dose.

Future work includes determining the effects of shielding on the neurotoxicity, as well as determining neurotoxicity of various ions as a function of linear energy transfer (LET). Shielding the explant cultures from the ion beam may reduce the neurotoxic response, may show no significant reduction, or may even further inhibit neurite growth. In the LET study, it is expected that a high-LET ion will have a greater neurotoxic effect on the explant cultures than a lower-LET ion at the same radiation dose. Additionally, the algorithm is being modified to quantify neurite outgrowth from individual neuron cells (see Fig. 17 for preliminary results), in order to study the cellular morphological responses induced by methamphetamine (METH). We expect that increasing doses of METH will result in reduced neurite outgrowth. This research was supported in part by a grant from the Swartz Foundation.

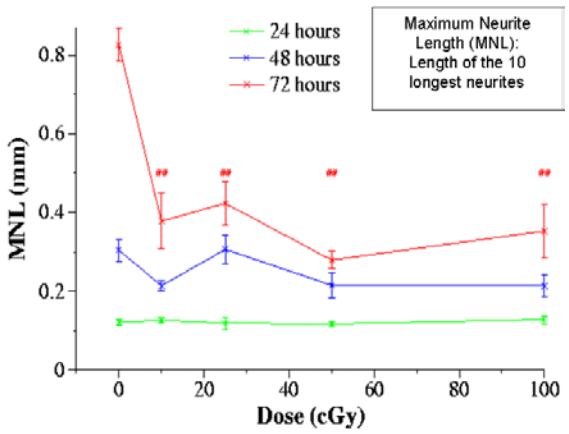

Figure 16. Automated analysis of the time-dependent effects of the dosage of 1 GeV/nucleon Fe ions on maximum neurite length. No significant effect is detected prior to 72 hours after explantation; at 72 hours significant growth reduction compared to the control is observed at all radiation doses.

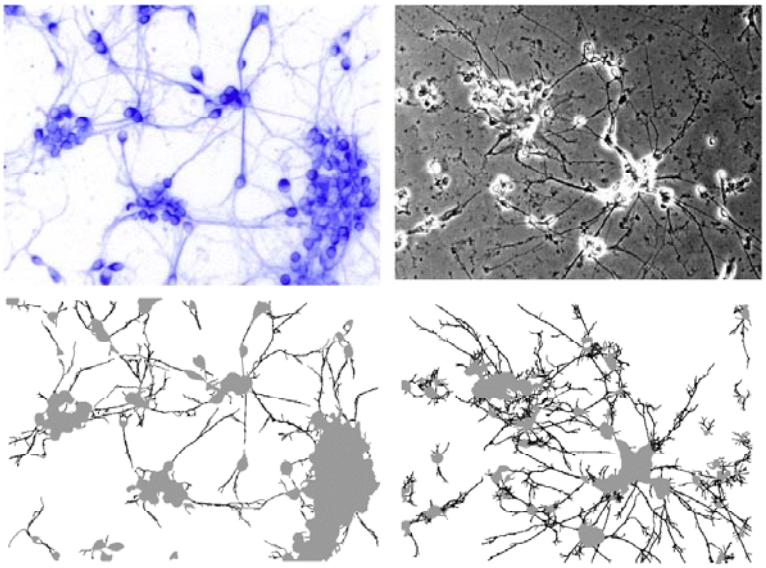

Figure 17. Digitized microscope images of both stained and non-stained neuron cell cultures treated with METH. Below the raw images are preliminary results of the modified algorithm as applied to each image.

Reference

1. Weaver, C. M., Pinezich, J. D., Lindquist, W. B. and Vazquez, M. E. An algorithm for neurite outgrowth reconstruction. Brain Research, submitted 2001.

4.3.6 Cardiac Electrophysiology

The purpose of this project, carried out in collaboration with researchers from Hofstra University and Beth Israel Medical Center, is to develop a family of robust, usable computer models of the cardiac electrical system, including realistic implementations of ventricular anatomy and cell physiology. The cardiac electrical system can be described by a system of partial-differential equations of reaction-diffusion type, whose solutions are various forms of traveling waves in three dimensions. Developing realistic models poses a significant computational challenge because the ventricle is a multi-scale, complex system: cardiac equations are sensitive to details, cell dynamics involves 20 or more types of ion channels and the ventricles have a complex, anisotropic geometry, requiring 10^6 - 10^{10} nodes, depending upon the degree of spatial resolution required.

Our approach begins with capturing cell dynamics at the intra-cellular and channel level in intermediate-scale models using a minimal set of phenomenological currents. (See Fig. 18.) This approach requires only 10^6 - 10^7 nodes and a 100-microsecond time step. In addition, the models will be parallelized so as to run rapidly (one minute simulation yielding one second of cardiac dynamics) on multiprocessor systems such as the Brookhaven Galaxy. These simulations will be validated both analytically and by comparison with animal and computer experimental data. They will then be used as "production code" in a large-scale study of wave stability and scroll wave dynamics in realistic 3D anatomy, using 3D stereoscopic visualization techniques and statistical analysis of the results of multiple simulations.

This project will yield well-tested production code for cardiac electrical dynamics. Large-scale simulations will provide critical insights into basic cardiac physics and dynamics. Moreover, the use of these models will allow experiments which cannot be performed with animal models because variability among individuals within a species makes it difficult to isolate causal factors.

Parallelization effectively exploits the multiload memory hierarchy of many modern supercomputers and clusters (multiple processors on a bus with a network to connect many such mother boards). Open MP and MPI are both used to this effect.

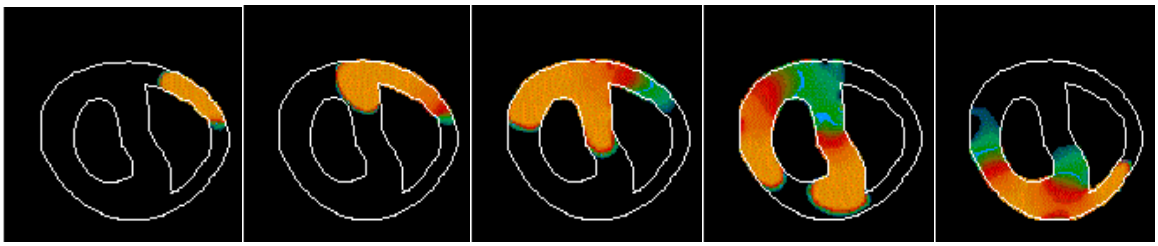

Figure 18. Simulation of a propagating wave (unexcited tissue is shown in black; excited tissue in yellow) on a 2D slice of ventricle taken from the anatomical rabbit model.

References

1. Fenton, F. and Karma, A. Vortex dynamics in 3D continuous myocardium with fiber rotation: fibrillation, CHAOS **8**, 20 (1998).
2. Fenton, F., Evans, S., and Hastings, H.M. Memory in an excitable medium: A mechanism for spiral wave breakup in the low excitability limit. Phys. Rev. Lett. **83**, 3964-3967 (1999).
3. Hastings, H.M., Fenton, F., Evans, S., Hotomaroglu, O., Geetha, J., Gittelson, K., Nilson, J., and Garfinkel, A. Alternans and the onset of ventricular fibrillation. Phys Rev E **62**, 4043 - 4048 (2000).

4.4 BASIC ENERGY SCIENCES

The Department of Energy's Office of Basic Energy Sciences supports research in materials and chemical sciences, geoscience, engineering, and energy biosciences. This broad-based program of fundamental research is also responsible for a number of national user facilities, including synchrotron x-ray and neutron sources and electron microscope facilities. CDIC has engaged in numerous projects in computational fluid dynamics, many body physics, and optics, which are described here. See Table 6.

TABLE 6. CDIC BASIC ENERGY SCIENCES PROJECTS

Project Title	Key Personnel	Comments
Magnetohydrodynamics of Multifluid Systems	R. Samulyak, W. Oh, J. Glimm, J. Davenport	Multiple Applications
Atomization and Spray	J. Glimm, A. Marchese, C. Tzanos, R. Samulyak, X. Li	Fuel Injection for Diesel Engine
Quantum Mechanics of Magnetic Systems	J. Davenport, I. Campos	Quantum Monte Carlo
First Principles Calculations of Melting Phenomena	J. Davenport, N. Stojic	Density Functional Theory and Classical Monte Carlo
Photonic Crystals	J. Glimm, Y, Deng, N. Stojic	Parallel FDTD Maxwell Solver

4.4.1 Magnetohydrodynamics of Multifluid Systems

Introduction. The main goal of this project is to develop numerical algorithms and computational software for the magnetohydrodynamics (MHD) of multi fluid (gas) systems with strong discontinuities in physical properties of the system components. This objective requires scalable parallelized algorithms for the effective solution of elliptic and parabolic equations in three dimensions with strong, geometrically complex discontinuities in the coefficient matrix defining the Laplacian operator. The proposed numerical algorithms are implemented as a magnetohydrodynamics extension of the FronTier code. FronTier, a hydrodynamics code with free interface support, is based on front tracking, a numerical method which allows the computation of

dynamically evolving interfaces. The method is able to detect and resolve changes in the topology of the moving front. It allows reduction in the level of grid refinement and has no interfacial numerical diffusion. The method of front tracking is ideal for problems with strong discontinuities. We implemented a grid moving algorithm in the FronTier-MHD code to bring distorted, but logically rectangular, grid elements into alignment with discontinuity surfaces, so that no discontinuity crosses a finite element boundary. See Fig. 19. This algorithm requires a high level of topological robustness in its interface handling, including intersection detection and interface topology determination, which are issues basic to computational geometry. The grid is used for finite element elliptic and parabolic solvers of the magnetic field and current density evolution equations based on Raviart-Thomas and vector Whitney elements. Development of the MHD extension of the FronTier code will enable us to model a variety of MHD phenomena in multicomponent (multiphase) systems. In particular, we shall apply the code to model the behavior of the Muon Collider target which is a mercury jet interacting with high energy proton pulses in the presence of a strong magnetic field, liquid lithium flows in the proposed Tokamak cooling system, and melted metal flows in experiments with exploding wires.

Mathematical formulation. The system of equations governing a compressible inviscid conducting fluid in a magnetic field is

$$
\begin{aligned}
\frac{\partial \rho}{\partial t} &= -\nabla \cdot (\rho u) \\
\rho\left(\frac{\partial}{\partial t} + \mathbf{u}\cdot\nabla\right)\mathbf{u} &= -\nabla P + \rho \mathbf{X} + \frac{1}{c}(\mathbf{J}\times\mathbf{B}) \\
\rho\left(\frac{\partial}{\partial t} + \mathbf{u}\cdot\nabla\right)U &= -P\nabla\cdot\mathbf{u} + \frac{1}{\sigma}\mathbf{J}^2 - \frac{1}{c}\mathbf{u}\cdot(\mathbf{J}\times\mathbf{B}) \\
\frac{\partial \mathbf{B}}{\partial t} &= \nabla\times(\mathbf{u}\times\mathbf{B}) - \nabla\times\left(\frac{c^2}{4\pi\sigma}\nabla\times\mathbf{B}\right) \\
\nabla\cdot\mathbf{B} &= 0,
\end{aligned}
$$

where ρ, $\mathbf{u}$, P, and U are the fluid density, velocity, pressure, and the total energy, correspondingly, X is the external force, B is the magnetic field induction and J is the current density distribution. This system must be completed with an appropriate equation of state. The boundary conditions:

$$
\begin{aligned}
\mathbf{n}\cdot(\mathbf{B}_2 - \mathbf{B}_1) &= 0, \\
\mathbf{n}\times(\mathbf{H}_2 - \mathbf{H}_1) &= \frac{4\pi}{c}\mathbf{K},
\end{aligned}
$$

must be satisfied at the free fluid interface, where $\mathbf{B} = \mu\mathbf{H}$, μ is the magnetic permeability and **K** is the surface current density. The MHD system is an example of a coupled hyperbolic – parabolic system, each of which can be solved separately at every time step using the operator splitting technique.

In many important physics applications, the influence of the eddy currents on the external magnetic field can be neglected and the current density distribution can be found by solving the Poisson equation

$$\mathbf{J} = \sigma\left(\nabla\phi + \frac{1}{c}\mathbf{u}\times\mathbf{B}\right)$$

$$\Delta\phi = -\frac{1}{c}\nabla\cdot\left(\mathbf{u}\times\mathbf{B}\right) \ ,$$

with

$$\left.\frac{\partial\phi}{\partial\mathbf{n}}\right|_{\Gamma} = -\frac{1}{c}(\mathbf{u}\times\mathbf{B})\cdot\mathbf{n}$$

The MHD equations represent in this case a hyperbolic system coupled to a scalar elliptic equation.

Numerical implementation. The MHD system consists of an elliptic/parabolic system coupled to a hyperbolic system. The coupling is solved via operator splitting. The equations are solved in each of two domains, separated by the free surface defined by FronTier's interface tracking numerical techniques. The evolution of the free fluid surface is obtained through the solution of the Riemann problem for compressible fluids. The elliptic equation for the current density distribution is solved using a finite element method based on Raviart-Thomas elements. We have been working on the parabolic solver for the magnetic field evolution equation using Whitney vector finite elements. The tetrahedral finite element grid is rebuilt at every time step and conformed to the new shape and position of the dynamically evolving interface. This is accomplished by using a point shifted grid which conforms to the tracked interface, and which has a rectangular index structure. Scalable parallelism will use direct methods to solve the elliptic/parabolic term on each processor and a wire basket outer iteration to ensure matching conditions across processor domain boundaries.

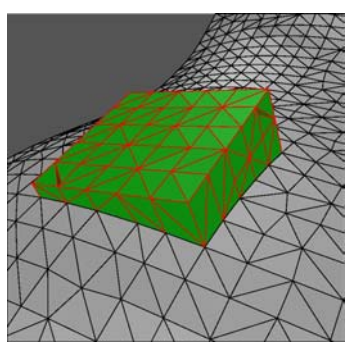

Figure 19. Triangulated tracked surface and tetrahedralized hexahedra conforming to the surface. For clarity, only a limited number of hexahedra have been displayed.

Results of numerical simulation. As a specific application problem, we have started modeling the motion of free mercury jets in strong (20 Tesla) solenoids. This problem arises as part of the design of the target for a proposed high energy accelerator, the Muon Collider [7]. Preliminary numerical simulation of the motion of a liquid jet in 20 Tesla magnetic field is shown in Fig. 20.

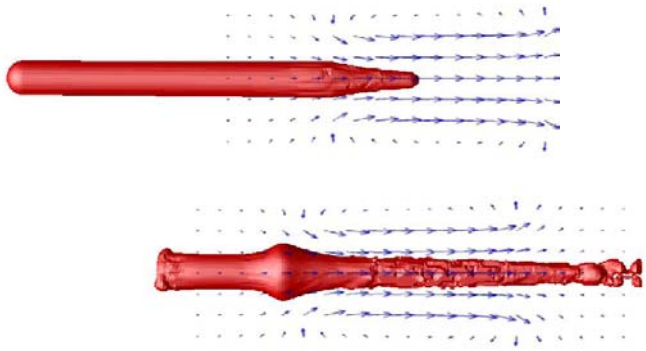

Figure 20. Motion of a conducting liquid jet in a 20T solenoid.

Other applications will include the modeling of free surface lithium flows in a proposed Tokamak cooling system and the dynamics of melted metal jets in experiments with exploding wires.

References

1. Landau, L.D. and Lifshitz, E.M. *Electrodynamics of Continuous Media*, Addison-Wesley Publishing Co., Reading, Mass., 1960.
2. Chandrasekhar, S. *Hydrodynamic and Hydrodynamic Stability*, Clarendon Press, Oxford, 1961.
3. Glimm, J., Grove, J.W., Li, X.L., Shyue, K.M., Zhang Q., and Zeng, Y. Three-dimensional front tracking. SIAM J. Sci. Comp. **19**,. 703-727 (1998).
4. Glowinski, R. and Wheeler, M. Domain decomposition and mixed finite element methods for elliptic problems. In R. Glowinski, G. Golub, G. Meurant, and J. Periaux, Eds., *Proceedings of the First International Symposium on Domain Decomposition Methods for Partial Differential Equations,* pp. 144-172. SIAM, Philadelphia, 1988.
5. Chan, T.F. and Matthew, T.P. Domain decomposition algorithms. Acta Numerica **3**, 61-143 (1994).
6. Cowsar, L.C., Mandel J., and Wheeler, M.F. Balancing domain decompositions for mixed finite elements. Math. Comput. **64**, 989-1015 (1995).
7. Ozaki, S., Palmer, R., Zisaman, M., and Gallardo, J., Editors. Feasibility Study-II of a Muon-Based Neutrino Source. BNL-52623, June 2001; available at http://www.cap.bnl.gov/mumu/studyii/FS2-report.html.

4.4.2 Atomization and Spray

The purpose of this research is to determine the size distribution for coherent mixing structures. Here we address the break up of a fuel jet in a diesel engine jet breakup [2] and spray flow [1,3,4]. We will build on our prior studies of fluid mixtures which result from acceleration driven mixing (See Figs. 25 and 26). There is substantial literature on the jet breakup problem. However, a predictive simulation capability for cavitation and particle size distribution does not appear to exist. We are developing a version of FronTier, with phase transition (fuel boiling) included.

Cavitation starts at the nozzle inlet, and a time dependent slug or laminar two-phase flow develops in the nozzle itself. FronTier will allow accurate simulation of the cavitation (liquid/vapor fuel)

interface. Mass transfer (evaporation) across the interface will require modification of the front propagation Riemann solvers in FronTier. We have already installed an analytic isentropic EOS for the vapor/liquid diesel fuel, which includes the fuel vaporization phase transition.

References

1. Amsden, A.A., Ramshaw, J.D., O'Rourke, P.J., and Dukowicz, J.K. KIVA: A computer program for two- and three-dimensional fluid flows with chemical reactions and fuel sprays. In *Computer Simulation for Fluid Flow, Heat and Mass Transport and Combustion for Reciprocating Engines*, pp. 239-344. Hemisphere Publishing Corp., New York, 1989.
2. Andrews, M.J. The large-scale fragmentation of the intact liquid core of a spray jet. Atomization and Sprays **3,** 29-54 (1993).
3. Lefebure, A. *Atomization and Sprays*. Hemisphere Publishing Corp., New York, 1989.
4. Watkins, A.P. Three dimensional modeling of gas flows and sprays in diesel engines. In N. C. Markatos, editor, *Computer Simulation for Fluid Flow, Heat and Mass Transport and Combustion for Reciprocating Engines*, pp. 193-238. Hemisphere Publishing Corp., New York, 1989.

4.4.3 Quantum Mechanics of Magnetic Systems

The quantum mechanics of multi electron systems is fundamental to an understanding of magnetism. Progress on this problem will have profound implications for chemistry and condensed matter physics. Correlated electron systems are also important technologically. Materials containing manganese which display colossal magnetoresistance are likely candidates for high density storage devices; and superconductors with high transition temperatures could lead to enormous energy savings.

Naively the correlated electron problem is exponentially hard because the number of states which must be considered grows exponentially with N, the number of electrons. Quantum Monte Carlo techniques have great promise because they reduce the multi electron problem to a sequence of single electron problems so that the overall effort scales roughly as the third power of N.

We have utilized the quantum Monte Carlo algorithm of Blancenbecler, Scalapino, and Sugar [1,2] to study the single orbital Hubbard model in three dimensions with, typically, N = 64, 216, or 512 sites. The model is defined by the hopping parameter, t and the Coulomb interaction, U. For U = 0, the bandwidth or energy range of the allowed states is determined by t (it is 12t in 3 dimensions). For t = 0, the electrons are completely localized at sites.

This model displays a phase transition at low temperature to an antiferromagnetic ground state for U > 0. Because many actual antiferromagnets are layered structures, we have also varied the hopping between layers on the lattice. For zero interlayer hopping the model becomes two-dimensional, for which it is known that there is no phase transition (except at zero temperature).

Since the bulk of the computer time is spent on matrix multiplication we have paid special attention to this aspect. Loop unrolling keeps data within the cache on our Linux cluster. Systematic cache optimization is provided by the Automatically Tuned Linear Algebra Software (ATLAS) library

[3]. We have developed two ways to parallelize the code. For small or large systems, multiple copies of the code can be run on different processors with different random number seeds, generating independent sequences of states. For larger systems, the matrix and its multiplication is distributed over a set of processors so that the same random numbers apply to each. Each processor can then perform the required matrix operations on a subset of the matrix elements. Clearly the two parallelization methods can be combined if desired.

The algorithm requires the inversion of matrices which become nearly singular for low temperatures. Following White et al. [2] we have implemented a scheme which uses singular value decomposition of the matrices and greatly expands the accessible temperature range.

To search for the phase transition we have studied the staggered magnetization as a function of temperature. The magnetization, or average value of the spin, is nearly zero since for antiferromagnets the spins point alternately up and down on adjacent lattice sites. The staggered magnetization is formed by taking the average spin but inverting it on every other site. It would be one for a classical antiferromagnet (and zero for a ferromagnet where all the spins are parallel or for an antiferromagnet at high temperatures where the order has been lost). The staggered magnetization is shown versus temperature (in units of t, the hopping parameter) in Fig. 21 for a 4x4x4 = 64 site cluster with periodic boundary conditions.

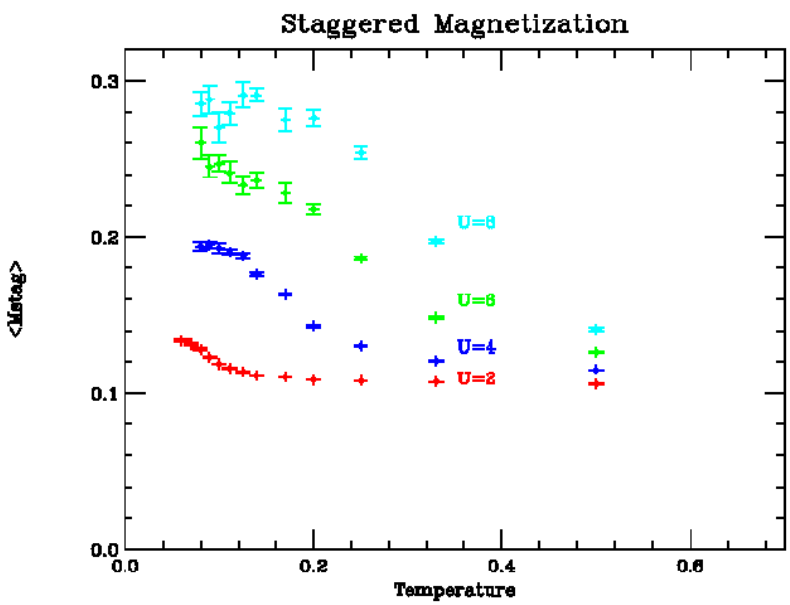

Figure 21. Staggered Magnetization versus temperature, T for a 4x4x4 lattice. T and U are in units of t, the hopping parameter.

As the Coulomb interaction U increases, the magnetization becomes larger because the electrons are more localized. The rise in magnetization as the temperature decreases is the result of ordering. The phase transition temperature is near the midpoint of the rise. However, the best way to determine the temperature is to study the magnetic susceptibility versus temperature. Results of such a study are shown in Fig. 22 and compared with previous work. As can be seen, the new results give a significantly lower transition temperature than found previously. We attribute this fact to the availability of longer runs, the use of periodic boundary conditions, and the utilization of the susceptibility to determine the phase transition. Further work on larger lattices is needed to confirm this result.

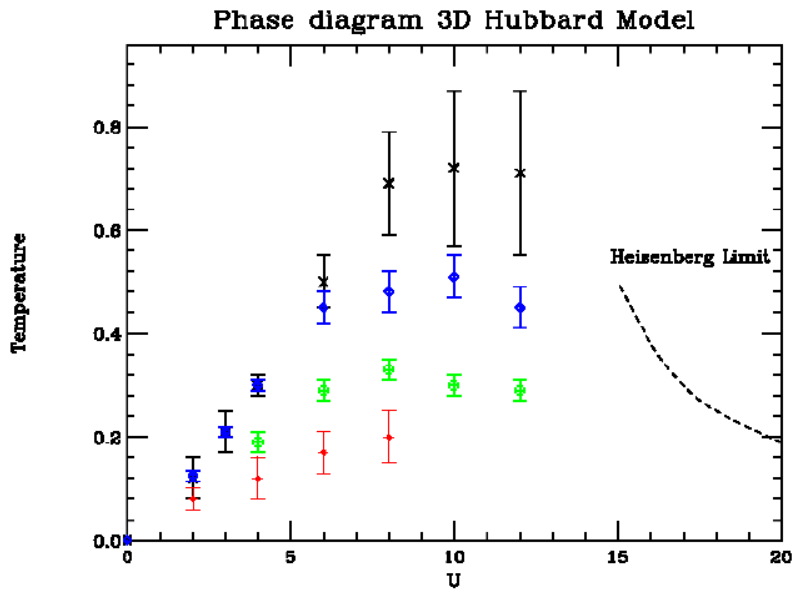

Figure 22. Neel Temperature versus U (both in units of t): red - current work, black [4], blue [5], and green [6] are previous quantum Monte Carlo results. The curve at layer L is the result for the Heisenberg model [7].

The single orbital Hubbard model has contributed greatly to our understanding of magnetic and superconducting systems but essentially all real magnets involve more than a single orbital. In order to compare with actual experiments, our future plans call for the study of the multi orbital Hubbard model with realistic values of the parameters.

We also plan to extend the message passing parallelization scheme to one involving a mixture of message passing and threads in keeping with recent trends is in supercomputing.

References

1. Blankenbecler, R., Scalapino, J.D., and Sugar, R.L. Monte Carlo calculations of coupled boson-fermion systems. I. Phys. Rev. D **24**, 2278 (1981).

2. White, S.R., Scalapino, D.J., Sugar, R.L., Loh, E.Y., Gubernatis, J.E., and Scalettar, R.T. Numerical study of the two-dimensional Hubbard model. Phys. Rev. B **40**, 506 (1989).
3. www.netlib.org/atlas/
4 Hirsch, J.E. Simulations of the three-dimensional Hubbard model: Half-filled band sector. Phys. Rev. B **35**, 1851 (1987).
5. Scalettar, R.T., Scalapino, D.J., Sugar, R.L., and Toussaint, D. Phase diagram of the half-filled 3D Hubbard model. Phys. Rev. B **39**, 4711 (1989).
6. Staudt, R., Dzierzawa, M., and Muramatsu, A. Phase diagram of the three-dimensional Hubbard model at half filling. Eur. Phys. J. B **17**, 411 (2000).
7. Sandvik, A.W. Critical temperature and the transition from quantum to classical order parameter fluctuations in the three-dimensional Heisenberg antiferromagnet. Phys. Rev. Lett. **80**, 5196 (1998).

4.4.4 First Principles Calculations of Melting Phenomena

Many fluid and plasticity simulations require accurate materials properties, including the equation of state at high pressures and temperatures. A first principles definition of the energy of an atomic configuration is given by density functional theory (DFT). Typically, this energy is obtained by solving a self consistent Schrodinger-like equation, usually by iterative diagonalization with a preconditioner.

We apply classical Monte Carlo simulations to determine melting points. These use DFT to compute the energy as a selection criterion in the Metropolis algorithm to accept or reject randomly chosen atom configurations. Trajectories can be easily generated both at constant volume and constant pressure. The free energy and other thermodynamic functions are determined by thermodynamic integration. The crossing of free energy curves for the liquid and the solid determines the melting temperature. Previous work with aluminum suggests 5% errors compared to experiment are possible. Further work on heavier metals is planned.

4.4.5 Numerical Simulation of Photonic Crystals

Photonic crystals are a novel class of optical materials fabricated with at least two different dielectric permittivities in a periodic arrangement. They have the ability to suppress, enhance, or otherwise control the emission of light in a selected frequency range by altering the density of states. A complete photonic band gap (PBG), i.e. a range of frequencies for which light cannot propagate through the crystal in any direction, is a spectral region where the density of states in an infinite crystal vanishes. Photonic structures have a number of potential applications: light emitting diodes, low loss optical devices with increased extracting capability (high Q, low volume and thin cavities, surface grating couplers), lasers, waveguides, low-loss waveguide bends, junctions, couplers and many more.

We have developed a simulation algorithm for photonics, capable of handling complex geometry, realistic initial and boundary conditions, finite size effects, dispersive and nonlinear media, and surface waves. Our primary focus is on the simulation of realistic experimental structures, such as waveguides, photonic band gaps, and totally reflecting boundaries. We used the finite-difference time-domain method to solve the Maxwell equations.

We use a mesh which is divergence-free with respect to its electric and magnetic fields, and suitable for specifying field boundary conditions and singularity [1]. We developed a parallel code which was run on the Brookhaven Galaxy cluster. The calculational domain consists of 19 photonic lattice cells in the periodic structure for the dielectric constant in the x and y directions and 1/5 of a photonic cell in the z direction. Periodic boundary conditions were used in all three directions. Each photonic lattice cell has been divided into 40x40 computational mesh cells, but due to duality of the calculational mesh (see Yee, [1]), we effectively calculated each field on only 20x20 points inside a photonic lattice cell. The period of the oscillation was divided into 90 time steps of the numerical integration. A typical 50-period (4500 time steps) run took 50 minutes.

We have applied this algorithm to study the localized transverse magnetic (TM) defect modes in a two-dimensional, triangular-lattice photonic crystal [2]. The sample consists of an array of circular air-cylinders in a dielectric medium (chosen to represent gallium arsenide). To validate our computations, the results for the transverse electric (TE) case were checked against experimental and numerical results using a different method. We study the spatial symmetry for the TM modes, obtained by changing the dipole excitation frequency. Also, we varied the defect-cylinder radius to tune the resonant frequency across the band gap. The TM mode is found to be highly localized at the defect in the photonic lattice. Figure 23 shows the spatial distribution of the magnetic field for the E1 and B2 modes. Figure 24 illustrates the eigen frequency dependence on the radius of the defect rod.

Currently, we are applying our results on defect modes and localization effects to laser structures. The novel aspects of this study are the choice of a different geometry, new absorbing boundary conditions, and resulting different modes.

(a) (b)

Figure 23. The spatial distribution of the magnetic field after 100 oscillation periods for two modes: (left) E1 mode [scaled angular frequency $a/\lambda = 0.411$ (a is the lattice constant and λ is the wavelength of the oscillating electric field)] and (right) B2 mode ($a/\lambda = 0.48$).

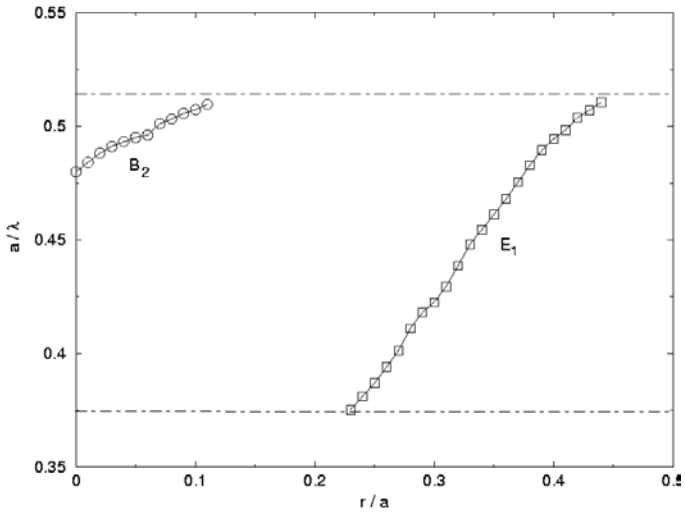

Figure 24. The eigen frequency of the localized defect modes as a function of the defect rod. The ordinate is the normalized frequency and horizontal lines represent the boundaries of the photonic band gap.

References

1. Yee, K.S. Numerical Solution of Initial Boundary Value Problems in Isotropic Media. IEEE Trans. Antennas and Propagation **14**, 302 (1996)
2. Stojic, N., Glimm, J., Deng, Y., and Haus, J.W.. Transverse magnetic defect modes in two-dimensional triangular-lattice photonic crystals. Physical Review E (in press).

4.5 ADVANCED SCIENTIFIC COMPUTING

See Table 7 for a list of projects.

TABLE 7. CDIC ADVANCED SCIENTIFIC COMPUTING PROJECTS		
Project Title	**Key Personnel**	**Comments**
Frontier-Lite	X. Li	Tech Transfer
Uncertainty Quantification	J. Glimm	Decision Making
Visualization	K. Mueller, A. Peskin, M. McGuigan	Viz Theater
Brain Imaging	K. Mueller, W. Zhu	BrainMiner
Cluster Computing	N. D'Imperio, E. McFadden	Cost Effective

4.5.1 Frontier-Lite

The Front Tracking code FronTier has achieved a notable agreement with laboratory measurements and with theoretical calculations for mixing rates for acceleration driven fluid mixing flows. See Figs. 25, 26 and Table 8 below.

In Fig. 25, we display early and late time interfaces between heavy and light fluids in a gravitationally inverted situation, leading to acceleration driven fluid mixing. In Fig. 26, we plot a mixing height, h, vs. a scaled t^2 variable. The left frame is height vs. (scaled) t^2. Front tracking is compared to Total Variation Diminishing (TVD) simulations at two grid levels. The right plot has been rescaled to account for effects of diffusive mixing. The Atwood number, $A = (D_2 - D_1) / (D_2 + D_1)$ is a dimensionless measure of the density contrast. For the diffusive simulations, A = A(t) is a function of t. The left plot of Fig. 27 shows height h vs. 2ll A(s) gds ds. Note that this rescaling, which accounts for the diffusive mixing through a time dependent Atwood number, brings the diffusive simulations into agreement with experiment and theory, as well as with our nondiffusive simulations. In Table 8, we list the mixing rate for experiment, theory and these simulations. After renormalization of the diffusive simulation, all mixing rates are in approximate agreement. FronTier completely eliminates numerical mass diffusion across a tracked fluid interface, the apparent reason for its success in achieving consistency with theory and experiment.

For this reason, there is a growing interest in transfer of the Front Tracking technology to other simulation codes. The FronTier-Lite project is to produce a simplified, robust version of FronTier, and to transfer this technology for inclusion into other simulation hydrocodes. We will combine FronTier with the Overture code developed at Livermore and the code SAGE developed at Los Alamos. Both codes will add an automatic mesh refinement capability to FronTier-Lite. This project will also add a front tracking capability to these, and in principle, to other codes.

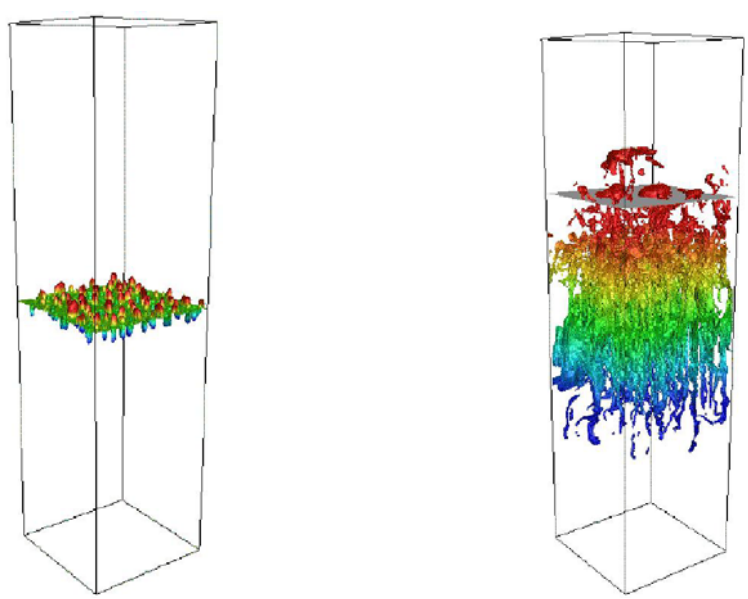

Figure 25. Simulation of acceleration driven fluid mixing. Simulation by Xiaolin Li, Zhiliang Xu, and Erwin George.

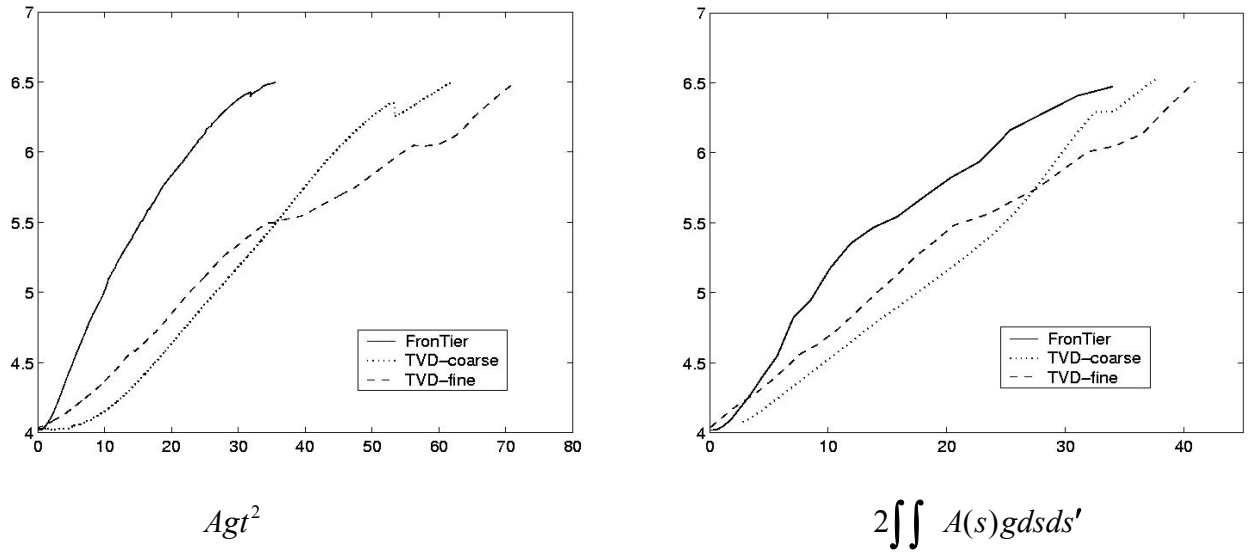

Figure 26. Mixing growth comparison of a *FronTier* (nondiffusive) with a TVD (diffusive) simulation.

TABLE 8	
Method	"
Experiment	0.05 – 0.077
Theory	0.05 – 0.06
FronTier simulation (unrenormalized)	0.07
FronTier simulation (renormalized)	0.07
TVD simulation (unrenormalized)	0.035
TVD simulation (renormalized)	0.06

Table 8. Values of " determined from experiment, theory, and simulation. All values are consistent except the unrenormalized TVD value (with " determined from a time independent $t = 0$ Atwood number).

4.5.2 Uncertainty Quantification

Quantification of Uncertainty and Computer Assisted Decision Making

The need for computer assisted decision making is driven by two related factors. The first is the importance of complex scientific/technical decisions, such as those related to global warming, for which controlled experiments are not feasible. The second is the need for rapid or timely decisions, using incomplete information, such as in shortening the time to market duration of a product design cycle, mandating a reduction of the role of the human in the loop.

The result is a strong increase in the requirement of fidelity in physical, mathematical, and computer based models of reality. Also required, and the focus of the work being conducted here, is a measurement of accuracy, or error inherent in model prediction. Assessment of errors has three uses. The first is to allow refinement of the model, in conjunction with observational data. The second is to quantify risk associated with model based predictions. The third is to refine the simulation adaptively to reduce the error.

We have several main results in the analysis of numerical simulation errors. We have developed a Gaussian model for simulation errors and used it in a Bayesian framework to (a) solve a stochastic version of the history matching, or parameter determination inverse problem with (b) confidence intervals for forward prediction, based on a verified, objective mathematical formalism, see [1, 2, 3, 4, 5]. This study was conducted in the context of petroleum reservoir simulation, in joint work with Chevron, BP, and Los Alamos National Laboratory. The Gaussian error model is defined by its mean and covariance. We have recently reformulated the model in terms of arrival time variables [4]. The result is simpler (only five parameters) and robust (independent of geology correlation length).

A second main result [5] was to show that simulation errors can be usefully represented by Gaussian statistics, provided the observed quantities have been averaged over at least a few mesh cells [a quantity larger than the numerical shock wave rise distance]. In a shock physics simulation study, such quantities as wave location and arrival times were shown to be approximately Gaussian up to about 1.5 sigma. Larger errors were distinctly subGaussian, and appear to be essentially bounded in an L_∞ norm.

Future work will emphasize simplicity of the error model, with physically meaningful model parameters. We expect thereby to achieve transferrability of the error models, so that they can be developed and parameterized "off line" in at least somewhat simplified contexts and then used in more complex situations.

References

1. Glimm, J., Hou, S., Kim, H., Sharp, D., and Ye, K.. A probability model for errors in the numerical solutions of a partial differential equation. CFD Journal 9, 2000; Proceedings of the 8th International Symposium on Computational Fluid Dynamics, Bremen, Germany, 1999; SUNY-SB Preprint Number SUNYSB-AMS-99-11.
2. Glimm, J., Hou, S., Kim, H., Lee, Y., Sharp, D., Ye, K., and Zou, Q. Risk management for petroleum reservoir production: A simulation-based study of prediction. Report No. SUNYSB-00-12, State University of New York at Stony Brook; J. Comp. Geosciences (to appear).
3. Glimm, J., Hou, S., Lee, Y., Sharp, D., and Ye, K. Prediction of oil production with confidence intervals. Presented at SPE Reservoir Simulation Symposium, Houston, Texas, 11-14 February 2001.
4. Glimm, J., Lee, Y., and Ye, K. A simple model for scale up error. Contemporary Mathematics, submitted, 2001.

5. DeVolder, B., Glimm, J., Grove, J.W., Kang, Y., Lee, Y., Pao, K., Sharp, D.H., and Ye, K. Uncertainty quantification for multiscale simulations. J. Fluids Engineering (submitted 2001); Los Alamos National Laboratory Preprint LA-UR 01-4011; SUNY Stony Brook Preprint SUNYSB-AMS-2001-08.

4.5.3 Visualization

The Center for Data Intensive Computing offers an array of advanced visualization facilities and tools associated with both Brookhaven and the University at Stony Brook. The Visualization Theater at BNL uses a stereographic projection facility that projects its three-dimensional images into a conference room setting, using polarized light for stereo separation. Viewers wear passive polarized glasses to see the stereo effect. The 3D illusion created by stereo is enhanced by one or more other visualization techniques including shading, perspective, and motion. High energy and nuclear physics simulations, biomedical imaging, and protein structure analysis are among the applications that have utilized this facility (http://www.ccd.bnl.gov/visualization/).

The Visualization Laboratory at Stony Brook, headed by Leading Professor Arie E. Kaufman (http://www.cs.sunysb.edu/~vislab), is developing volume rendering techniques for use in scientific visualization and virtual reality applications. Currently, there are projects in: architectures for volume rendering, parallel methods for volume rendering (regular and irregular grids), development of tools for visualization, and flow visualization.

Klaus Mueller, a faculty member of the Visualization Laboratory, is developing efficient volume rendering algorithms for interactive visualization of large datasets occurring in science, medicine, engineering, and entertainment (http://www.cs.sunysb.edu/~mueller).

A popular volume rendering algorithm is splatting, which pairs good image quality with an efficient volume projection scheme as well as a sparse object representation. The object is modeled as a collection of fuzzy points which are projected on the screen where they composite to the desired image. The originally proposed approach, however, leads to a number of inaccuracies. The effect of these is less noticeable in still images, but clearly revealed in animated viewing, where disturbing popping of object brightness occurs at certain view angle transitions, and a sparkling of object features diminishes the viewing experience under perspective distortion. A further drawback is the severe blurring incurred in magnified viewing. Our work in this arena has yielded a new generation of the splatting algorithm, termed image-aligned sheet-buffered splatting, that no longer suffers from these limiting artifacts. Popping is eliminated by adding the splats within sheets, aligned parallel to the image plane, and compositing these sheets in front-to-back order. Invisible, occluded splats are eliminated from the projection pipeline at an early stage via a process termed early splat elimination, which is reminiscent of early ray termination in raycasting. Depth-adaptive anti-aliased splats are splats that adjust their frequency bandwidth depending on distance from the image plane. They eliminate the sparkling artifacts that are so annoying under perspective viewing conditions, such as flights across terrains or endoscopic applications. Finally, blurring under magnified viewing is avoided by reordering the volume rendering pipeline and performing classification and shading after splat rasterization on the sheet-buffer plane. Examples of splatting are shown in Figure 27 (a)-(d).

But despite efficient volume rendering algorithms, such as splatting, volume rendering at interactive frame rates remains a challenge, especially with today's increasingly large datasets. We are currently developing a framework, using concepts from Image-Based Rendering (IBR), that decreases the required frame rate for the volume renderer significantly. All the volume renderer needs to supply is a set of volume renderings at 'key' view points, and the IBR renderer will interpolate the intermittent frames at good accuracy. The IBR provides methods to handle both opaque and transparent datasets, and is an inexpensive process that can be run on the user's desktop machine.

The Visualization Laboratory features a comprehensive visualization tool, the VolVis system. This volume visualization system unites numerous visualization methods within an advanced visualization environment, providing a flexible tool for the scientist and engineer as well as the developer and researcher. It provides a wide range of functionality with numerous methods incorporated in each functional component. The VolVis user interface is organized into functional components, providing an easy to use visualization system. The structure of the VolVis system is designed to allow a visualization programmer to add new representations and algorithms with ease. The system, written in C, is highly portable, running on most Unix workstations supporting X/Motif. VolVis is freely available as source code to not-for-profit organizations. A small fee is charged for-profit organizations.

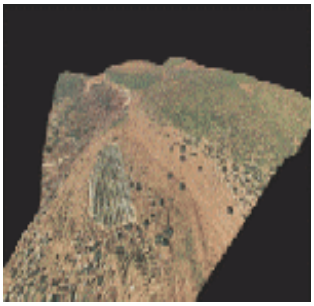

(a)

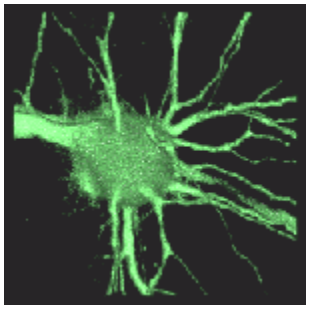

(b)

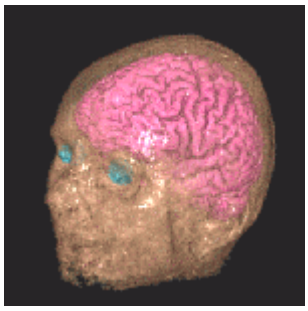

(c)

(d)

Figure 27. (a) Anti-aliased splatting; (b) image-aligned splatting (eliminates splatting's popping artifacts); (c) high-quality splatting on anisotropic grids with efficient occlusion culling; (d) splatting without the blur.

4.5.4 Brain Imaging

See Section 4.3.4.

4.5.5 Cluster Computing

Linux clusters have become important providers of computing cycles in the scientific world. Ours is modeled on similar clusters at Stony Brook (http://galaxy.ams.sunysb.edu/) and at BNL (http://www.ccd.bnl.gov/bcf/cluster/). It consists of 156 Intel Pentium III processors running at speeds from 500MHz to 1 GHz. Arranged in dual processor nodes, most have 1 Gbyte of memory per node. They are connected with a 100 Mbit/sec Cisco switch. A storage device with 1.4 terabytes of disk is attached to this machine. Parallel computations are performed using the message passing interface, MPI (www-unix.mcs.anl.gov/mpi) and through threads using Open MP. We plan to augment the system with an additional 200 processors.

Our goal is to continue to explore the frontier of the commodity component world, as processor, motherboard bus, and memory speeds increase, as fiber cable decreases in price, and switch speeds increase in performance and decrease in price. We will determine optimal configurations and architectures, within the overall framework of our design. Results from this research will be published and posted on the network, to benefit others seeking a similar price/performance advantage. Similar experiments with regard to disk storage will be conducted.

The Galaxy computer will provide computing cycles to CDIC affiliates and projects. Convenient and assured access to local parallel computing is an essential requirement for the development of parallel simulation codes, and it is thus a key component of the CDIC strategy to bring state-of-the-art computing technology to the benefit of the BNL scientific programs.