Argonne Laboratory Computing Resource Center

FY2004 Report

Argonne Laboratory Computing Resource Center

FY2004 Report

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

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

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

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

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

Contents

Executive Summary

Argonne National Laboratory founded the Laboratory Computing Resource Center in the spring of 2002 to help meet pressing program needs for computational modeling, simulation, and analysis. The LCRC established a new Argonne supercomputer facility with capability and flexibility to address a wide range of computing needs across the Laboratory. The LCRC also is positioned to support a wide range of computer users, with expert assistance for both scientists and engineers who are new to parallel computing as well as those who are seasoned parallel computer users. The guiding mission is to provide critical computing resources that accelerate the development of high-performance computing expertise, applications, and computations to meet the Laboratory's challenging science and engineering missions.

The first goal of the LCRC was to deploy a mid-range supercomputing facility to support the unmet computational needs of the Laboratory. Based on a survey of Argonne's computational scientists and engineers, processor benchmarking, and applications analysis, the LCRC selected a Linux-based commodity cluster strategy as the most cost-effective option for a production computing system for the Laboratory. In September 2002, the Laboratory purchased a 350-node computing cluster from Linux NetworX.

This cluster, named "Jazz," achieved over a teraflop of computing power (i.e., 10^{12} floating-point calculations per second) on standard tests, making it the Laboratory's first terascale computing system and one of the fifty fastest computers in the world at the time. Jazz was made available to early users in November 2002 while the system was undergoing development and configuration. In April 2003, Dr. Hermann Grunder inaugurated the Laboratory's new user facility at a ceremony marking the transition to production operation.

The Jazz user community grew quickly and steadily. By the end of fiscal year 2004, there were 66 active projects on Jazz involving over 250 scientists and engineers. These projects represent a wide cross-section of Laboratory expertise, including work in biosciences, chemistry, climate, computer science, engineering applications, environmental science, geoscience, information science, materials science, mathematics, nanoscience, nuclear engineering, and physics. Overall use of the system has grown to the maximum capacity that can be sustained on such a shared resource. Most important, many projects have achieved results that would have been unobtainable without such a computing resource. In addition, the LCRC has increased the number of staff, postdoctoral associates, and their students who are versed in applying high-performance computing approaches.

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

The Laboratory Computing Resource Center

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

Inaugural System

The LCRC's first computer cluster was installed in September 2002. The cluster was named "Jazz" because we expected that the mix of disciplines and applications operating on the system would be diverse, wide ranging, and continually changing. Moreover, the name is short and easy to remember. Jazz is a "Beowulf" cluster, built largely from commodity components.[1] Beowulf clusters, introduced in the late 1990s, have become the predominant platform for high-performance computing. They are relatively inexpensive compared with specialized high-performance computing hardware, because they are built with components drawn from common workstation-class computers. They also take advantage of an extensive assortment of open-source software, ranging from the popular Linux operating system to the Argonne-produced MPICH tools for interprocess message passing. A large number of science and engineering applications have been ported to these architectures by universities, federal laboratories, and independent software vendors. Further, while Beowulf clusters do not perform optimally for all possible algorithms, they perform quite reasonably on most.

Application Services

Pervasive application of computation in science and engineering requires more than the availability of a supercomputer. In many of Argonne's disciplines, researchers are just beginning to make the transition to high-performance computing to address their computational needs. In other areas, scientists have been using parallel computers for some time, but new research avenues require new computational approaches. Therefore, in addition to supporting Jazz, the LCRC provides consulting services to the Argonne computational research community. These services include training in computation techniques from the fundamentals to advanced topics, assistance with code performance analysis, guidance with algorithm development, and general help and advice. Since the LCRC staff are familiar with many projects across the Laboratory, they often help identify potential collaborators with specialized expertise. The long-term goal is to develop a vigorous computational science and engineering community at Argonne.

---

[1] T. Sterling, D. Savarese, D. J. Becker, J. E. Dorband, U. A. Ranawake, and C. V. Packer, "BEOWULF: A parallel workstation for scientific computation," in *Proceedings of the International Conference on Parallel Processing, Vol. 1: Architecture*, pp. 11–14, CRC Press, 1995.

Resources

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

Goals

The LCRC has set the following goals for fiscal year 2005.

- Continue to operate Jazz as a highly effective production supercomputing resource.
- Expand the use of Jazz to new projects and disciplines, with an emphasis on the Laboratory's initiatives.
- Help Argonne personnel improve the scientific reach and performance of their computational applications.
- Help large-scale users of Jazz identify opportunities to obtain allocations at national computing resources such as NERSC, NLCF, MHPCC, PACI, and TeraGrid.
- Develop a plan and specifications for expanding LCRC computing resources in FY2006.
- With Argonne's information infrastructure team, examine the requirements and opportunities to provide comprehensive data management across the Laboratory.

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

The Jazz Computing Cluster

After the installation of Jazz by Linux NetworX in late September 2002, the LCRC staff immediately began configuring and tuning the system for Argonne-specific use. This process involved significant modification to various system services, as well as installation of the tools needed for parallel computing, project management, documentation, and user support. The first users began to test the system in October 2002, using code that was known to work on similar systems. Gradually the initial user base was enlarged in order to expand the range of tests being carried out on the system and to exercise the various components. In December Jazz was opened for exploratory use by the Argonne community; and after the system stabilized, Jazz went into formal production mode in April 2003. Dr. Hermann Grunder presided at a ribbon-cutting ceremony that marked the occasion (photo: from left to right: Remy Evard, founding director of the LCRC; Herman Grunder, Laboratory director; Susan Coghlan, systems administrator; and Rick Stevens, chair of the Computational Science Advisory Committee).

Cluster Configuration

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

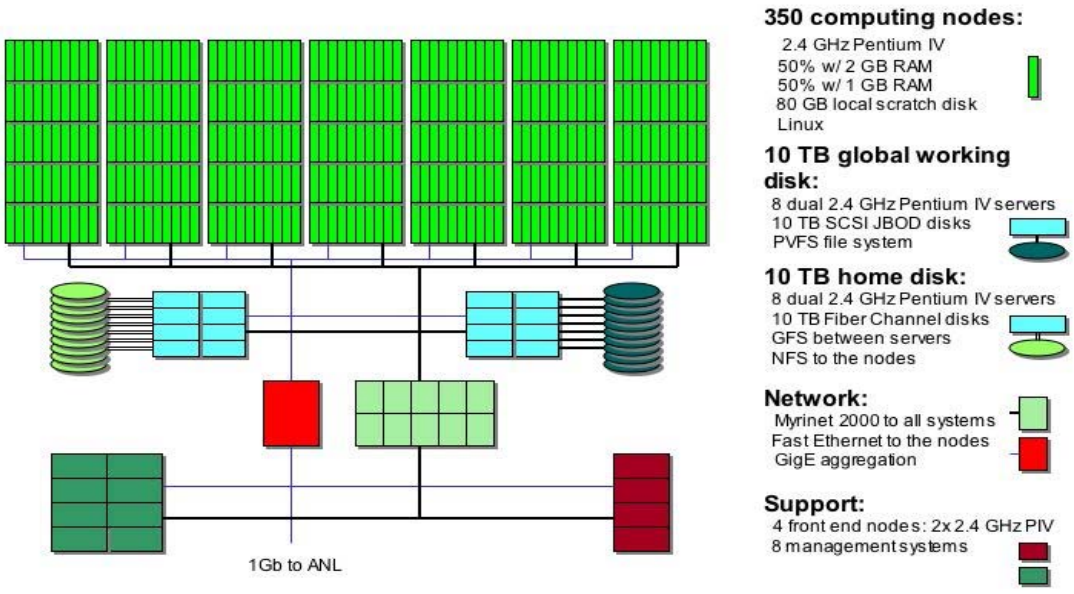

Benchmarks of the processors and memory system in Jazz compare very favorably with other Linux clusters installed in the same time period.

Jazz Usage

As can be seen in the chart below, the usage of Jazz grew rapidly, averaging 80% in fiscal year 2004 – in spite of the occasional outage. On this chart 100% would mean that all 350 nodes were used continuously for that day. The trend line rises to 85% by the end of the year. The maximum sustainable usage is about 85% over long periods, given that some nodes go idle when the smallest computation waiting to run is requesting more nodes than are available. Reservations for very large computations (e.g., over 200 nodes), create additional holes as nodes are held open for a while until the large computation has sufficient resources to begin.

Consider a scenario in which you were a travel agent chartering flights but only for tour groups consisting of 25 or more inseparable people. You could easily end up in a situation in which you would have to leave 24 or fewer seats open on a plane, even if you had other tour groups wishing to fly on that plane. The situation is analogous but more complicated on a cluster, when user jobs of varying sizes run for different, and often unpredictable, lengths of time. Thus, utilization is an indicator of load, capacity, and scheduling efficiency, but not demand. Demand is more a function of the number of computations waiting to run, and our observation is that there are nearly always jobs waiting in the Jazz queues.

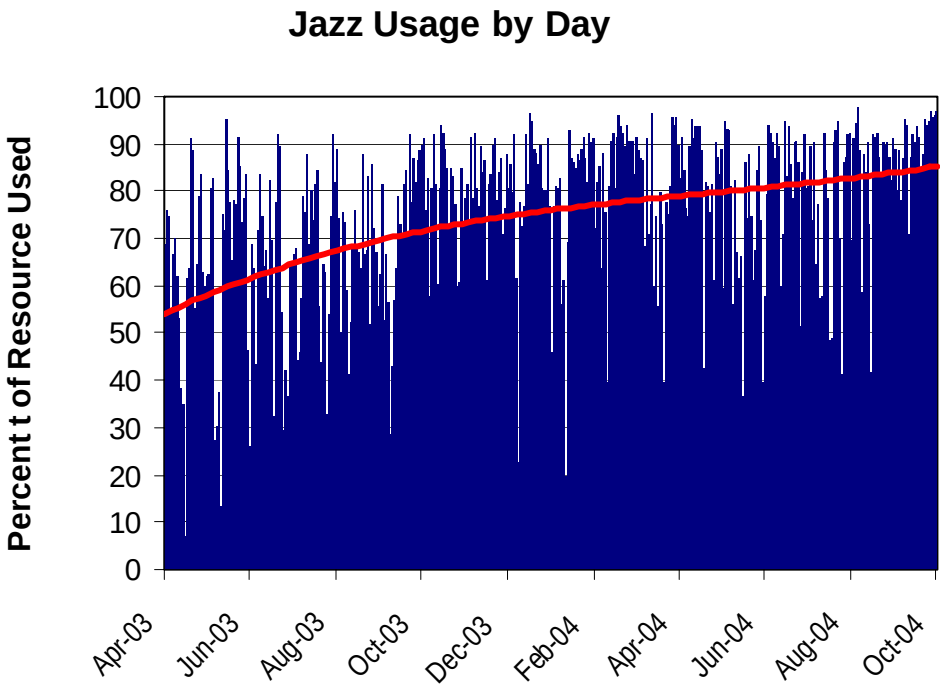

Jazz supports a wide mix of computations, from long runs with a few nodes to short runs with many nodes, with many variations between. On a project-by-project basis usage varies considerably over time. Projects often start slow, then ramp up, then increase and decrease in cycles as the project progresses, with interruptions for analysis, manuscript preparation, conferences, and vacations. The last section of this report provides more detailed statistics on usage patterns over time. If one considers the whole spectrum of projects using Jazz, a significant fraction is still in the ramping-up phase, indicating that although the Jazz system is saturated, needs for Laboratory computing resources are still growing.

Scientific and Engineering Applications

In a relatively short time, the LCRC has become an essential resource for a wide range of projects in many Argonne divisions, from production runs to exploratory modeling and simulation research to high-performance algorithm and application development. In the LCRC's first 18 months of production computing, 70 projects from 15 divisions used Jazz for research, employing its full range of capabilities on cutting-edge science and engineering problems.

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

Modeling, simulation, and informatics are integral components of essentially all of the Laboratory's strategic initiatives. Here the LCRC is playing a growing role in developing the next generation of research programs and capabilities, by providing both essential computational resources and expertise about high performance computing. Jazz is critical to strategic initiative LDRD projects as well as completion of programmatic work in strategic areas. In other cases Jazz is the testing ground for the concepts that will make up future LDRD and programmatic proposals. For many of these strategic projects Jazz is the only large computing resource available.

Another important component of LCRC's mission is to provide a fast and reliable resource for the Laboratory's production modeling and simulation activities. Some of these production codes are still under development; others are fully mature. Often Jazz is the largest resource available to a research group. For other production applications, Jazz fills the need for mid-range runs, where larger runs are done on systems bigger than Jazz at NERSC or other national centers. For example, the quicker turnaround on Jazz increases research productivity, while NERSC has resources for much larger computations. Also, parallel versions of a number of commercial science and engineering codes are available on Jazz, for example, for production computational fluid dynamics computations.

Exploratory projects are investigating approaches and algorithms for next-generation computations, developing prototype applications with new capabilities, and evaluating software performance, accuracy or programming models. While large parallel computing resources offer tremendous power, frequent access to them is essential to develop effective algorithms and codes. Challenges include strategies for dividing the problem into many pieces, optimization of the kernels of computation, and management of the massive flow of results to storage. Some exploratory projects are drawing on expertise in the Argonne community to explore new designs or models of complex phenomena; other projects are forging new paths to extend the range of existing applications. Some researchers are starting a parallel programming project from scratch and are drawing on the experience of the LCRC staff.

LCRC Research Highlights

In this section, we present 15 examples of research performed with Jazz through September 2004. These reports highlight the science and engineering advances being made with Jazz, tackling problems from subatomic scale to the size of cities. They span both the basic and applied missions of the Laboratory, including most of our strategic initiatives. This research encompasses a wide range of computational techniques, including agent-based simulations, specialized forms of computational fluid dynamics, Monte Carlo solvers and "first principles" solutions. These, and many other techniques, have been successfully carried out with parallel computing.

Catalysis in Nanoporous Materials (NANOCATALYSIS)

PI: Peter Zapol, Chemistry and Materials Sciences Divisions, and Center for Nanoscale Materials

Scientists have expressed considerable interest in understanding and controlling reactions catalyzed in nanoporous materials at the *molecular* level, spurred further by the recent discovery of how to construct metal oxide nanoporous membranes with atomic-level precision over a range of pore sizes (1–200 nm) and composition. Argonne has developed the capability to synthesize these membranes and has initiated new programs to develop them as catalysts, ion transport membranes, and biomolecule composites. Theoretical studies of reactivity in these types of nanoporous materials provide information on adsorption energies and activation barriers and insight into the reaction mechanisms for selective catalytic oxidation of hydrocarbons. The purpose of this project is to develop computational capabilities to carry out studies of reactivities in nanoporous materials by combining methods capable of studying systems of thousands of atoms with existing high-level quantum chemical methods that provide key information on reaction mechanisms.

These studies are focusing on three topics related to catalytic reaction pathways in nanostructured membranes: the catalyst-support architecture, the catalytic sites, and the reaction mechanisms. They feature high-level quantum chemical methods for studies of chemical reactions involved in selective activation of C-H bonds such as in propane and cyclohexane as well as the development of accurate models for the surface reaction sites. Access to Jazz allowed us to investigate the structure of an amorphous alumina support as a function of pore size using molecular dynamics studies with many-body potentials. At the same time, studies of initial steps in the propane dehydrogenation on vanadia

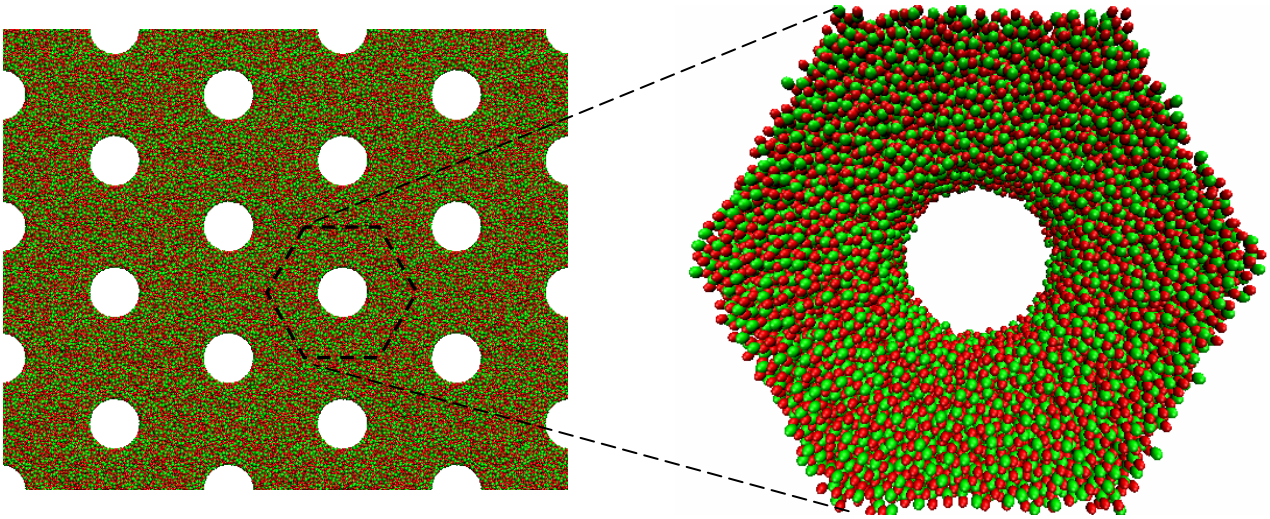

surface and supported vanadia clusters were performed at the quantum mechanical level. Whereas the adsorption reaction on a vanadyl oxygen on a flat surface was found to be highly unfavorable, we identified configurations of supported clusters leading to exothermic reactions. In the future, we will explore further steps in the oxidative dehydrogenation reactions and develop improved models for the substrate/catalyst interface that our results have indicated is important in determining the reaction rates.

Spatio-Temporal Chaos in Rayleigh Benard Convection (CHAOS)

PI: Paul Fischer, Mathematics and Computer Science Division

Nonequilibrium pattern formation and spatiotemporal chaos is an important field of study, both for its contributions to chaos and bifurcation theory and for its applications to a variety of problems such as chemical reaction-diffusion equations and ventricular fibrillation in a human heart. Rayleigh-Benard convection is one of the most productive systems for studying pattern formation, based on its accessibility to experimenters and its ability to be modeled theoretically.

Although spatiotemporal chaos has been an active area of study in experiment, theory, and numerics for a number of years, many of the basic characteristics remain poorly understood. Our calculations of the sensitivity to initial conditions and dimensions of the attractor via the Lyapunov exponents allow us to quantify the complexity of the dynamics and how it increases as system size and driving strength are increased. Our simulations of Rayleigh-Benard convection, for which the basic underlying equations are well understood and on which precise experiments are ongoing in various labs, provides a well controlled laboratory for developing understanding of these basic issues.

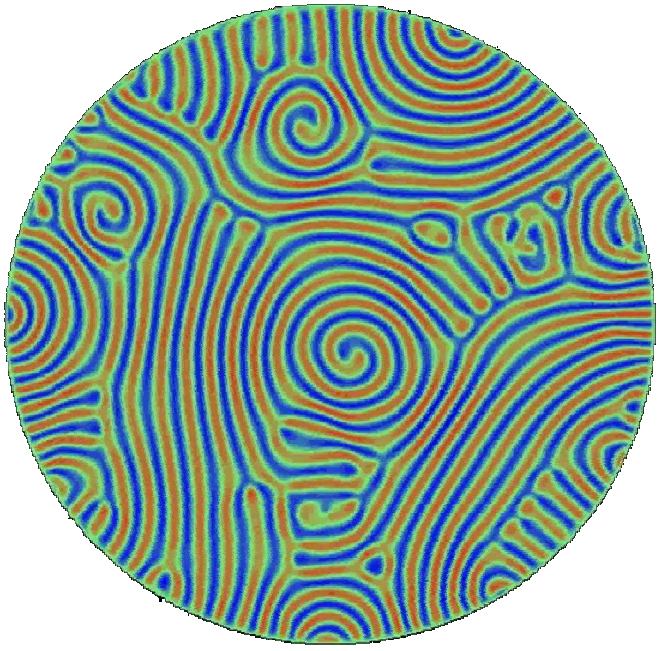

Simulations on Jazz showed that spiral and dislocation defects like these can result in non-Fickian transport of passive scalars in chaotic Rayleigh-Benard convection.

We modified the Nek5000 code to evolve multiple copies of the linearized perturbation equations simultaneously with the full solution. We performed numerous convergence and consistency tests in order to verify results. The largest Lyapunov exponents and their corresponding Lyapunov eigenvectors have been computed for various parameters. This work yielded two publications in 2004.

For aspect ratio 5, the largest Lyapunov exponent is indeed positive, confirming that the turbulent state discovered by Ahlers is truly chaotic. The Lyapunov eigenvector seems to be most localized in regions where defects are present. A more systematic exploration of the correlation between Lyapunov eigenvectors and defect nucleation/annihilation/motion is planned for the future. We would also like to compute the complete spectrum of Lyapunov eigenvalues for these systems.

Lyapunov exponents and eigenvectors have also been computed for aspect ratio 40 domain chaotic states in rotating cylinders. In preliminary runs, we found that the largest Lyapunov exponent increases

linearly with the control parameter. These results agree with Anandhan Jarayaman's aspect ratio 20 cases. We find, however, that for exactly the same parameters, the Lyapunov exponent increases as the aspect ratio increases from 10–40. We are also investigating the behavior of the Lyapunov eigenvector for domain chaos and its correlation with defect motion and domain wall motion. We plan to move to aspect ratio 60 and eventually compute all the positive Lyapunov exponents for our systems.

Multilevel Agent-Based Simulation of Bacterial Chemotaxis (REPAST_CHEMOTAXIS)

PI: Charles Macal, Decision and Information Sciences Division

A fundamental biology question is how specific cellular behaviors are the direct result of the biochemical processes that occur in the cell at the molecular level. Understanding these connections will provide fundamental insights about the nature of the basic mechanisms of cellular behavior, growth, adaptability, scalability, and diversity, as well as the factors responsible when these processes go awry. Recent laboratory findings have provided many details of these processes, in some cases forming a complete picture of the underlying biochemical processes.

Breakthroughs in computational modeling have demonstrated the feasibility of simulating the time-dynamics of cellular signal transduction processes, based on the approach of modeling individual molecules and reactions. The continuing rapid development of distributed computation architectures and multiscale simulation techniques allows the possibility of simulating intra-cellular processes for complexes of interacting cells existing in a dynamic environment. We are developing a computational process that characterizes some general principles underlying the functioning of biochemical networks.

A key challenge is to demonstrate an understanding of the complete process by which cells find and move toward chemical attractants. We have developed a multilevel simulation model to study important properties of the *E. coli* chemotaxis signal transduction network, a model system that has been well characterized experimentally. The chemotaxis network mediates the environment and movement of the cell, transforming molecular concentrations in the environmental medium into cell motion. This simulation framework models the dynamics of cellular processes operating simultaneously throughout various biological scales. The multilevel simulation models the dynamics of the signal transduction networks within cells at a fine-grained level and simulates the movement of the cells through a medium at a coarse-grained level of time and space. Hence, the macroscopic cell behavior is built from the ground up, based on the individual molecular reactions in the cell. The goal is to relate intracellular molecular events to the phenotype of a single cell and finally to the behavior of the entire population.

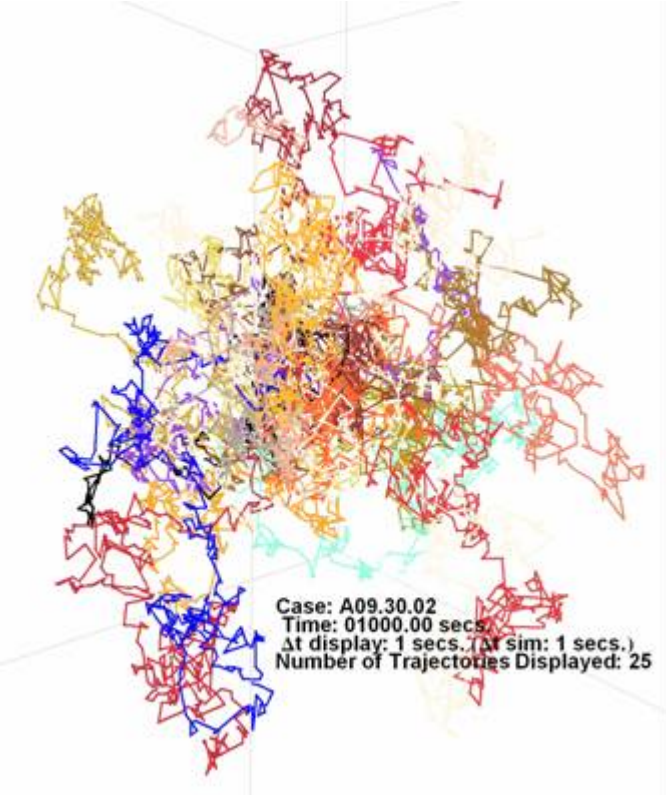

Trajectories of a population of 25 E. coli cells, based on laboratory experimental data.

The resulting simulation platform is AGENTCELL, a multiscale, agent-based model of the relationship between intracellular processes in individual cells and the behavior of a cellular population. Agent-based computational approaches have a natural modular architecture, which reflects the modular organization of biological systems. In the model of the chemotaxis system of swimming *E. coli* bacteria, each cell is an independent agent equipped with its own chemotaxis network, motors and flagella. Internal molecular fluctuations produce temporal variability in the behavior of individual cells. AGENTCELL is being further developed to incorporate time- and context-dependency of signaling pathways. These features will be used to study agent-interaction phenomena, such as quorum sensing, where the dynamic coupling between intracellular events and population dynamics is important.

Quantum Monte Carlo Calculations of Light Nuclei (QMC_FOR_NUCLEI)

PIs: Steven Pieper and Robert Wiringa, Physics Division

We use Quantum Monte Carlo (Green's function [GFMC] and variational [VMC]) methods to compute ground-state and low-lying excited-state expectation values of energies, densities, structure functions, astrophysical reaction rates, and so forth for light nuclei and low-energy scattering reactions involving these nuclei. Realistic two- and three-nucleon potentials are used. The goal is to accurately describe all of these systems by using a Hamiltonian that also provides an excellent description of nucleon-nucleon scattering and nucleonic matter. Such a "standard nuclear model" can then be used, for example, to compute low-energy astrophysical reactions, which cannot be experimentally measured. During FY2004 Jazz was used for many aspects of this project. The fast turnaround of Jazz makes it the machine of choice for all but the biggest calculations, which are done at NERSC or Los Alamos.

GFMC scattering: In the last few months of FY2004, we extended QMC to compute scattering states, obtaining fairly good agreement with experimental data. Previous calculations of the ^{5}He states ignored the fact that they are particle-unstable and treated them as true bound states. The GFMC could "see" this

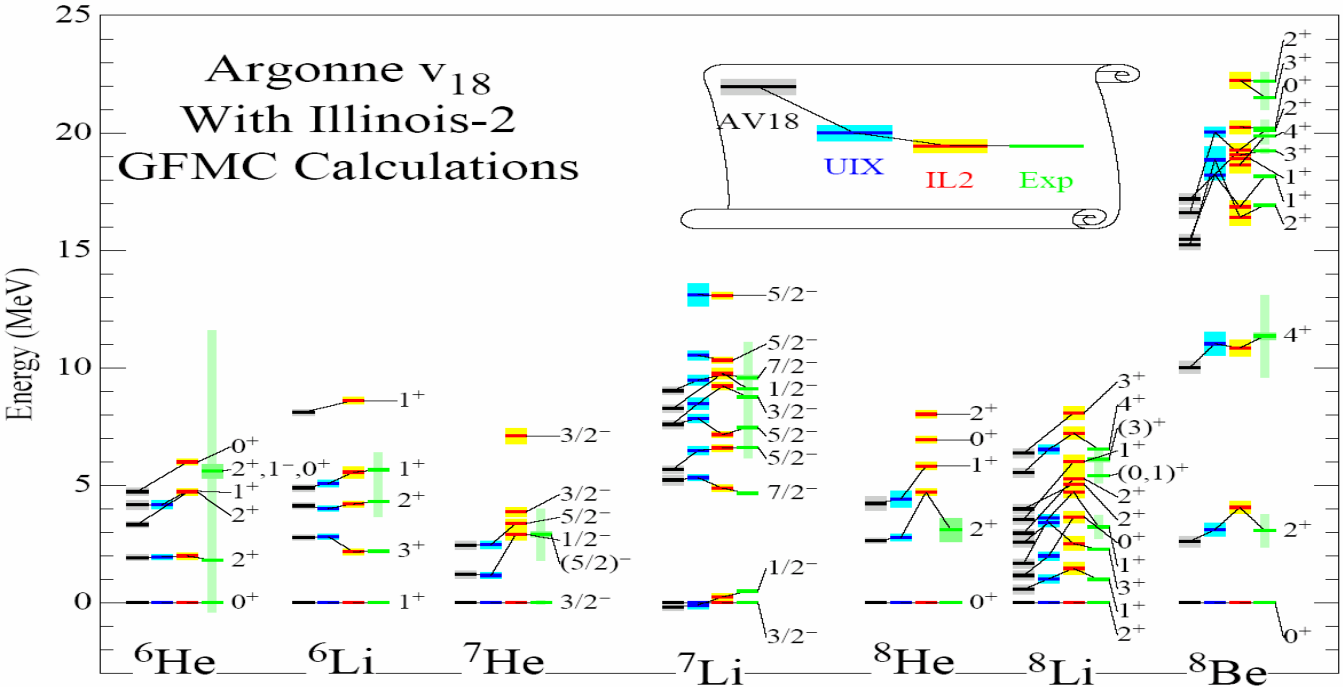

and produced slowly decaying energies, rather than a stable result. The scattering approach results in reliable calculations of such unbound states. Once validated, QMC can be compared against experiments for much larger nuclei. In addition to other single-nucleon scattering states, we will also study some astrophysically interesting capture reactions and the validity of approximate nuclear reaction methods; this work will require further development of the program.

Multiple excited states with the same quantum numbers: In FY2004 we completed an extensive set of calculations of the excited states of A=6,7,8 nuclei. This work was made possible by the realization that GFMC could produce stable results for multiple excited states with the same quantum numbers. Previously it had been assumed that GFMC would always go to the lowest state of given quantum number, even when the GFMC propagation is started from orthogonal trial wave functions. It was demonstrated that the orthogonality is preserved, allowing the computation of many more nuclear states for A=6,7,8 nuclei. We obtained good agreement with the experimental energies for the 36 known states in this region and made a number of predictions for other states. Most of this work was done on Jazz, with some of the A=8 calculations being done at NERSC. This work has been submitted to Physical Review C.

Spectroscopic factors: We also made an extensive set of calculations of overlap wave functions and spectroscopic factors for nuclei with A from 5 to 10. These overlap wave functions, or quasi-hole and quasi-particle amplitudes, can be used in reaction programs (such as the DWBA program Ptolemy) to compute differential cross sections for transfer reactions. Their first use was for the radioactive beam experiment d(^{8}Li,p)^{9}Li, which was done at Argonne's ATLAS facility; this experiment with QMC analysis has been submitted to *Physical Review Letters*. During these investigations, we found that a theorem that has been used since the 1970s to convert fixed-center shell-model spectroscopic factors to the desired translationally invariant ones is not really appropriate for light p-shell nuclei (where the effects of the fixed-center approximation are largest). Most of this work was done using Jazz, except for smaller cases done on desk-side systems.

RMS radii of weakly bound nuclei: A recent ATLAS experiment measured the RMS radius of the unstable nucleus ^{6}He with an accuracy of better than 1%. This led the attempt of a similarly precise calculation with Jazz, but they found that the GFMC produced a slowly fluctuating value of the RMS radius as the propagation proceeded, even though the computed energy was extremely stable. The obtainable accuracy is limited to several percent for ^{6}He; a similar limit probably applies to other weakly bound nuclei.

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

PI: W. David Pointer, Nuclear Engineering Division

The issue of energy economy in transportation has grown beyond traditional concerns over environment, safety, and health to include new emphasis on energy self-sufficiency. At the same time, new emissions standards for tractor-trailer vehicles introduce challenges in fuel economy for the manufacturers. Computer simulations offer a promising path to facilitating these improvements, but validation of results is imperative. With the U.S. Department of Energy Office of FreedomCAR and Vehicle Technologies' Heavy Vehicle Aerodynamic Drag Consortium, we are developing guidelines for the use of STAR-CD in vehicle design.

These guidelines are being developed based on experience gained with the generic conventional model, or GCM, and the STAR-CD suite of commercial computational fluid dynamics software tools. The GCM, developed by NASA Ames Research Center, is a one-eighth-scale generalized representation of a conventional U.S. tractor-trailer truck. The model is instrumented for the measurement of body forces and moments, surface pressure distributions, and near vehicle flow field distributions.

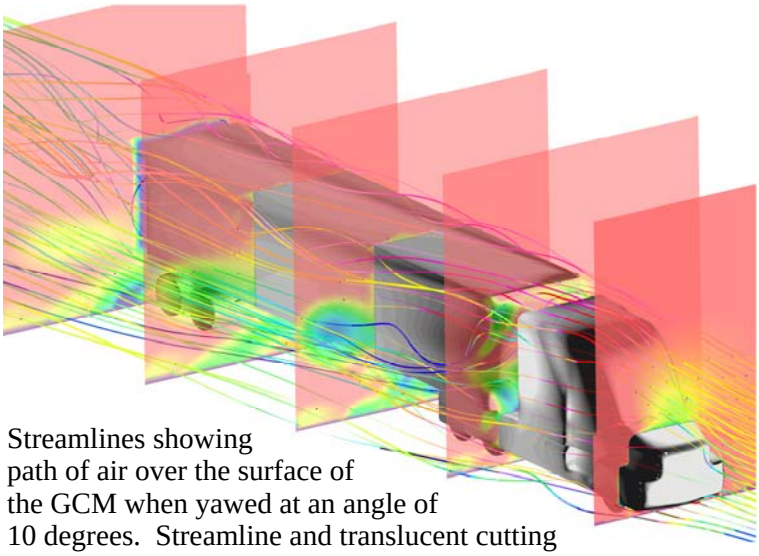

Streamlines showing path of air over the surface of the GCM when yawed at an angle of 10 degrees. Streamline and translucent cutting plane shading indicates local velocity magnitude and grayscale shading on vehicle surface indicates pressure contours.

In FY2004 we used the Jazz cluster to investigate the sensitivity of drag coefficient predictions and surface pressure distributions to the size and structure of the computational mesh, the selection of a turbulence model, and preliminary assessments of the prediction of the aerodynamic characteristics when the vehicle is positioned at non-zero yaw angles. These investigations have shown that simulations of wind tunnel experiments for a generalized tractor-trailer geometry using the STAR-CD can predict the detailed surface pressure distributions on vehicle surfaces with reasonable accuracy and resolve the drag coefficient within 1% of the experimentally measured value. These results are accomplished with only a small fraction of the computational power being employed by other approaches and also provide better accuracy. The Jazz high-performance computing cluster clearly demonstrated the excellent scalability of this aerodynamics problem.

As a result, we plan to extend the current approach to include much more complicated surface geometries and to further resolve regions of the computational domain where complex flow structures appear. These investigations will be carried out under a Collaborative Research and Development Agreement effort with PACCAR Inc., a U.S. manufacturer of heavy-duty Class 8 trucks.

Scalability of Numerical Algorithms (SCALABILITYOFALGS)

PI: William Gropp, Mathematics and Computer Science Division

We are evaluating the feasibility of porting scientific applications to next-generation hardware and software scenarios. The efforts are a component of the larger Argonne Petaflops project. Without an understanding of the hardware and algorithm demands of an application, researchers cannot decide the most effective computer architectures to use. Additionally, it is best to understand the utility of new computer architectures before time is spent porting codes. This effort seeks to answer a series of questions for scientific applications, including the following:

- What are the primarily performance bottlenecks for a given application?
- How can simple models be built and tested to predict performance on new architectures?
- How can scientific applications be classified by performance needs?
- What existing tools can be leveraged, and what tools need to be built?

The objective is to understand the needs of a broad spectrum of real scientific applications and to devise a methodology for performing similar analyses in the future.

Several stages of this process were performed on five Argonne applications spanning several different disciplines: QMC in physics, COLUMBUS in chemistry, Nek5000 for biofluid flows, FLASH in astrophysics, and pNeo in neurology. For these applications, we explored the maturity of the software process for each application. Then, we studied the performance of each application, defining benchmark simulations if none already existed and identifying the performance bottlenecks. We also analyzed the primary parallel and computational algorithms of the applications in order to understand theoretical peak performance of the codes and to further understand bottlenecks. Given these descriptions of the applications, we created analytical performance models to predict how these applications would perform on new systems.

Large fractions of the performance analysis and prediction calculations for these applications were completed on Jazz. The analysis and model of FLASH helped predict potential bottlenecks on the new IBM Blue Gene/L architecture and prompted some algorithm redesigns to avoid potentially huge problems scaling to thousands of processors on BG/L. As a partial result from these efforts, FLASH was one of the first applications to run on 16,384 processors and scale very well. Additionally, pNeo was able to explore limitations in its current implementation, prompting some redesign of the software. Overall, good progress has been made in developing a process for describing high-performance scientific codes.

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

PI: Tanju Sofu, Nuclear Engineering Division

We are leading the way to the first-ever three-dimensional simulations of the full core of a nuclear reactor. Specifically, we are coupling advanced numerical models of computational fluid dynamics (CFD) for thermal-hydraulic calculations and whole-core discrete integral transport for neutronics (DeCART) calculations into an integrated simulation capability. This application, based on rigorous treatment of multiphysics phenomena influencing the core design, involves computationally intensive studies of the operational and safety characteristics of various designs. Even an increase of 1% efficiency in current nuclear reactors would be a huge cost benefit, and an integrated approach like this will be essential for such analyses.

Initial development of this evolving production application was done on the Nuclear Engineering Division's 80-node Linux cluster, but the memory-intensive simulations outgrew that cluster. Jazz enabled us to conduct full-system runs of unprecedented size and resolution for core simulations. The simulations are some of the largest parallel runs using STAR-CD, the commercial computational fluid dynamics component. Future plans include expanding the model to incorporate other physics. We also plan to demonstrate that our code design is flexible enough to incorporate other CFD applications.

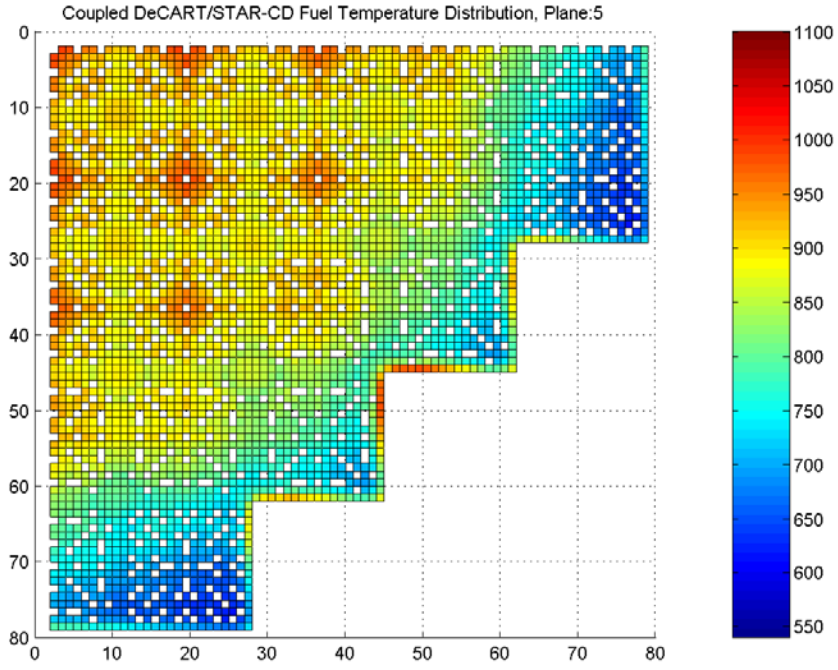

Efforts in 2004 focused mainly on the assessment of computational performance as well as quality assurance of the interface. We conducted a parametric study to establish proper convergence criteria and to optimize the data exchange frequency between the CFD and neutronics modules. Continuing work will focus on extending the integrated analysis capabilities by addressing two specific limitations that currently prevent application to full-scale light water reactor (LWR) models: (1) development of domain decomposition methods for the neutronics module to enable scaling LWR applications for solution on an arbitrary number of processors, and (2) development and validation of a transient analysis capability for the time-dependent method of characteristics solutions of the integral transport equation in DeCART.

Time-dependent solutions of the integral transport equation have never been implemented in an application before. A major obstacle is the daunting data storage requirements stemming from the need to keep the history of the angular neutron fluxes from all the previous time steps at each mesh location. Even with proper decomposition, this implementation will require creative data management schemes consistent with the challenges of petascale computing.

Lattice Quantum-Chromodynamics in Extreme Environments (LATTICE-QCD)

PI: Donald Sinclair, High Energy Physics Division

We are studying nuclear matter at high temperatures and densities using numerical simulations of the dynamics of quantum chromodynamics (QCD). Hadronic/nuclear matter at finite temperatures and/or densities existed in the early universe, exists today in neutron stars, and is produced in relativistic heavy-ion colliders. Nuclear matter should also provide a good description of the cores of heavy nuclei, away from their surfaces.

Our goal is to study the features of the phase diagram of QCD relevant to the physics mentioned above. We use the hybrid molecular-dynamics method of lattice QCD. Because studies at finite baryon/quark-number density have sign problems that preclude the naive application of standard (importance sampling) methods, we study the related problem of finite isospin density, which does not have such difficulties.

We are performing simulations on Jazz to study the behavior of the finite temperature transition from hadronic/nuclear matter to a quark-gluon plasma for small isospin density, which is closely related to the behavior at small quark-number density. The two most interesting questions to address are how the transition temperature depends on density and whether at large enough density there exists a critical endpoint beyond which the finite temperature transition ceases to be a mere crossover and becomes a first order transition.

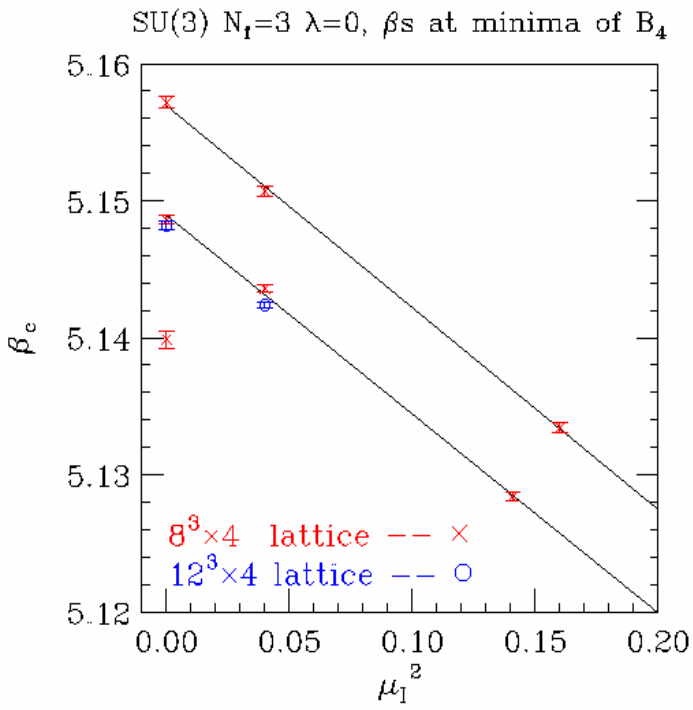

$\beta_c=6/g_c^2$ as functions of μ_I^2 from minima of B. The bottom point is for m=0.03. The middle set of points and line are for m=0.035. The top set of points and line are for m=0.04. Here μ_I is the isospin chemical potential.

Our simulations of the 2-flavor theory, started at supercomputer centers and finished on Jazz, showed that the transition temperature decreases slowly with increasing density. However, no sign of a critical endpoint was observed in the regime where finite isospin and finite quark-number density behaviors are expected to be related. (We did, however find such a change to a first-order transition at higher isospin densities.)

We are now studying the 3-flavor theory, which is expected to be closer to the physical world. Here it is believed that, by carefully choosing quark masses, the critical endpoint can be moved to arbitrarily small densities, where the finite quark-number and finite isospin densities should exhibit the same physics. We have observed the slow decrease of the transition temperature with mass for this case, as shown in figure. We have discovered that those quantities that determine the nature of the transition depend strongly on the updating increment δt, and we are now performing a systematic study of this δt dependence. When completed, this work should enable us to determine the position of the critical endpoint.

Geometry Optimization for Molecular Electronic Structure Using Multireference CI Methods (COLUMBUS)

PIs: Ron Shepard, Chemistry Division; Mike Minkoff, Mathematics and Computer Science Division; Thomas Mueller, Juelich Research Institute

This project involves the development, implementation, and application of theoretical methods for the calculation and characterization of potential energy surfaces involving molecular species that occur in hydrocarbon combustion. These surfaces require an accurate and balanced treatment of reactants, intermediates, and products. This challenge is met with general multiconfiguration self-consistent-field and multireference single- and double-excitation configuration interaction (MR-SDCI) methods. In contrast to the more common single-reference electronic structure methods, this approach can describe accurately molecular systems that are highly distorted away from their equilibrium geometries, including reactant, fragment, and transition-state geometries, and can describe regions of the potential surface that are associated with electronic wave functions of widely varying nature. The molecular geometry conformation optimization for the multireference expansions used in this study represents the largest and most flexible wave functions ever used in an analytic energy gradient optimization procedure. Such calculations are simply not practical without access to large-scale parallel machines such as the Jazz cluster.

Earlier we optimized the molecular geometry for ozone using two smaller wave functions. We immediately found deficiencies in the flexibility of the MR-SDCI and MR-AQCC calculations based on this reference space. In the MR-SDCI case, we found relatively large wave function amplitudes associated with movement between the b_1 lone-pair electrons and the (a_2,b_1) electron-pair orbitals of the reference space. Furthermore, we encountered convergence difficulties during the MR-AQCC optimizations based on this reference space, a feature often observed with inflexible reference spaces when using this method. We expanded the reference space in order to include these important wave function contributions directly. This expanded reference space consists of 81 reference terms, providing a reasonable description of the ozone wave function near the equilibrium geometry and in fairly accurate predictions of the equilibrium geometry. The wave function also eliminated any convergence problems. However, this wave function cannot describe smoothly the pseudo-rotation motion of the ozone molecule.

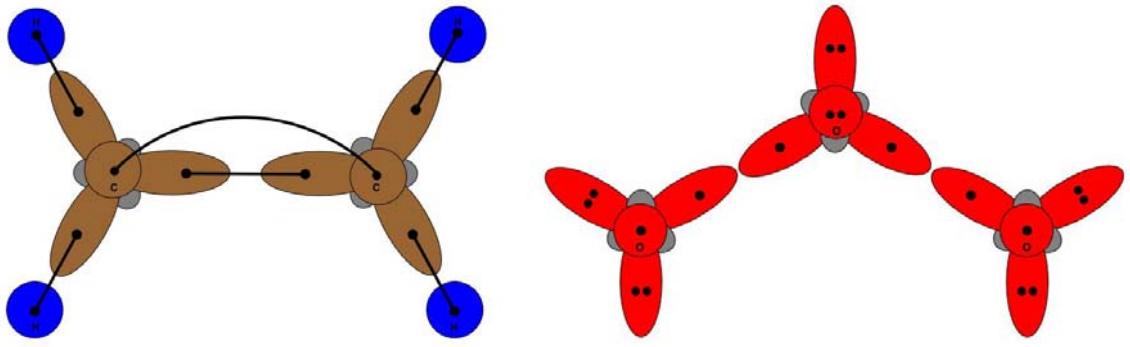

Ethylene (left) and ozone (right) geometry and bonding configuration.

To accurately describe the wavefunction of ozone, we needed to include wave function contributions important to interpair correlation with the b_1 lone pair, and configurations for the pseudo-rotation motion

associated with the isomerization motion that describes the three permutationally equivalent minima (O1-O2-O3, O2-O3-O1, and O3-O1-O2). The MR-SDCI calculations required the solution of an eigenproblem with rank of up to 520 million for the largest basis set used in this study. These wave function expansions are about a factor of 10 larger than those of the previous calculations, and they required access to the Jazz cluster in order to complete the optimization step. The computed bond lengths with this more flexible (12^{18}) wave function are a significant improvement over previous results. However, the ozone molecule still exhibits the largest errors, when compared to experimental values, of our 20-molecule sample. Further work is planned in order to determine the source of these errors.

Computational Nanophotonics (NANOPHOTONICS)

PI: Stephen Gray, Chemistry Division

Researchers have expressed much interest in focusing and manipulating light on the nanoscale, with the potential for developing new kinds of optoelectronics and chemical or biological sensing devices. Experimental groups at Argonne and throughout the world are actively engaged in such efforts. Metallic nanoparticles and related objects such as nanoholes in thin metal films are particularly promising device components because collective excitations of the electrons near the surface of the metal can resonantly interact with applied light, leading to intense, localized electric fields. These excitations are called surface plasmons. However, realistic electrodynamics simulations of arrays of metal nanoparticles and other potential nanophotonics systems can require very large computational efforts because nanometer to micrometer dimensions need to be considered in a given simulation. The purpose of this project is to develop and apply software tools for carrying out such simulations.

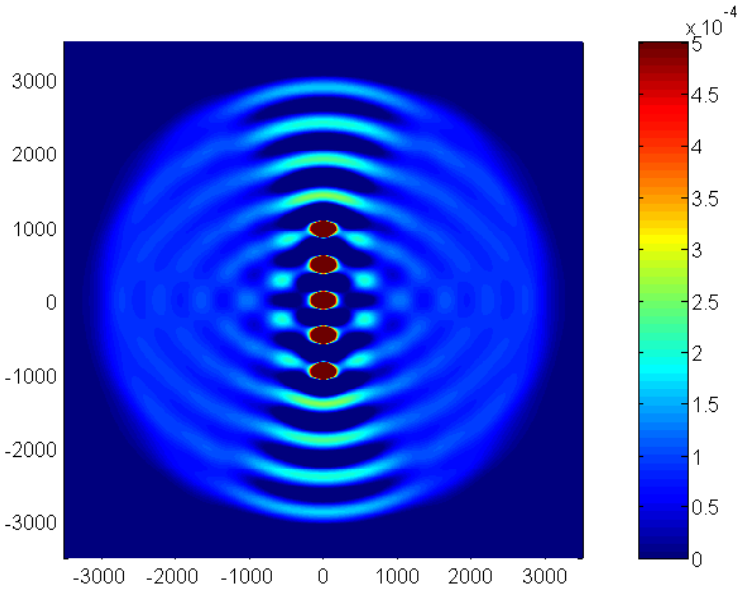

Simulation of the near-field intensity just above a 100 nm thick metal film that is irradiated with laser light from below. The alternating light and dark fringes are due to SPPs. The axis units are nm.

To this end we have developed a three-dimensional and parallel finite-difference time-domain code for determining the electric and magnetic fields of complex systems containing metallic components interacting with light. Motivated by experiments being carried out at Argonne's Materials Science Division, we carried out simulations of isolated nanoholes in thin gold films. Our joint experimental-theoretical paper showed how these nanoholes act as point sources of traveling surface plasmon waves called surface plasmon polaritons (SPPs), with implications for building devices that use SPPs as information carriers. The parallel code and access to Jazz allowed us to address important fundamental questions such as the role of experimental probe/sample interactions. A future direction is to look at

arrays of nanoholes and more complicated configurations. The image above shows the SPP waves that can come off a small linear array of nanoholes. In collaboration with MCS researchers we are also exploring new, more efficient routes to such simulations.

Wave Packet Dynamics with Cartesian Coordinates (CARTESIAN)

PI: Stephen Gray, Chemistry Division

The goal of this project is to develop and apply a wavepacket code for determining quantum dynamics based on Cartesian coordinates. We are concerned with the solution of the Schrödinger equation, which is normally solved with the use of specialized internal coordinates for a given problem at hand. The difficulty with this approach is that it is hard to extend to systems with more than a few atoms or to even few-atom systems in nonsymmetric environments, for example, atoms or molecules confined on or within complex nanostructures. Cartesian coordinates lead to a very simple structure that also can be more easily parallelized. We anticipate the simpler, Cartesian coordinates based approach to be superior. In addition to developing a general, parallel code for quantum dynamics simulations in Cartesian coordinates, another goal is the application of the code to benchmark problems including atom/diatom scattering and atoms or molecules confined within fullerenes such as buckyballs.

We have made significant progress on the development of a general, parallel Cartesian wave packet code. The parallelization makes use of the dispersion-fitted finite-difference (DFFD) approximation for the evaluation of the action of the Hamiltonian matrix on a vector (see the figure below). The code will allow, in principle, an arbitrary number of degrees of freedom. With the emergence of more powerful parallel computers, this will open the way to rigorous quantum simulations of systems with many more atoms than currently possible.

The code is written in C++ and uses MPI for parallelization. The object-oriented techniques provided by C++ will enable treating different quantum dynamics problems (e.g., requiring different initial wavepackets and/or different wavepacket analysis techniques) without modifications to the main part of the code. MPI parallelization ensures portability for various computing environments, such as widely used Linux clusters, clusters of multiprocessors, or massively parallel shared-memory computers (such as SGI Altix). We have fully designed the general structure of the code and written the majority of the code. Currently we are testing and debugging the code on Jazz, using the Totalview parallel debugger. We anticipate that the code will be completed by the end of January.

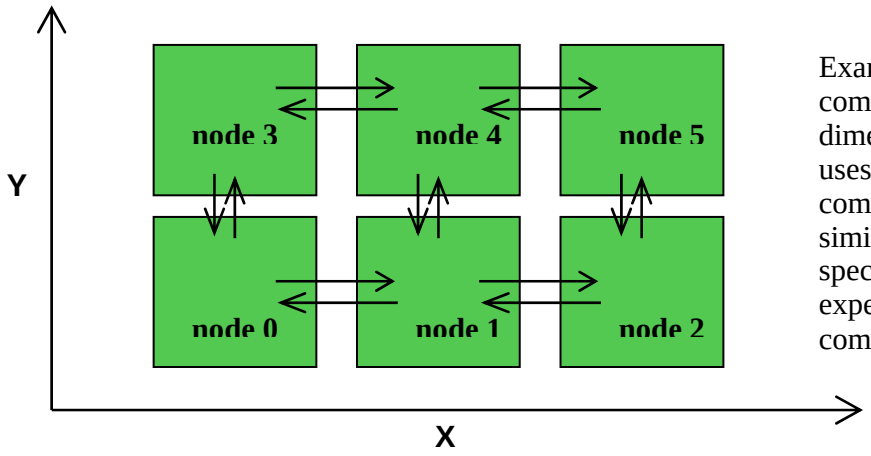

Example parallelization and communication scheme for two dimensions. The DFFD method uses only nearest-neighbor communication with accuracy similar to widely used pseudo-spectral methods, which require expensive collective communication.

First-Principles Modeling of Lithium-Battery Electrode Materials (LIBATTMODELING)

PI: Roy Benedek, Chemical Engineering Division

Lithium batteries, which power a wide range of consumer electronics applications, represent a multibillion-dollar industry. Since the lithiated cobalt-oxide-based electrode materials currently in use in lithium batteries pose cost and safety problems, a global research effort is under way to design and develop improved electrode materials. Among the most promising potential cathode materials to replace the cobalt-based materials are lithium manganese oxides in the layered and spinel structures. A technical obstacle, however, is that lithium insertion-extraction properties of manganese-based systems during electrochemical cycling are degraded by the Jahn-Teller distortion that occurs when the density of trivalent Mn ions in such insulating oxides is high. Dopants are commonly introduced into these materials to increase the Mn oxidation state. Both experimental and modeling efforts to characterize the structure and assess the electrochemical performance of such materials are being performed within the Battery Department in the Chemical Engineering Division.

The first-principles density-functional-theory calculations in this project assess the relative effectiveness of first transition-row and other dopants that substitute for Mn in stabilizing lithium manganese dioxide against the Jahn-Teller distortion. We performed calculations for materials with composition $LiMn_{0.75}M_{0.25}O_2$ in both the monoclinic and rhombohedral structures, where M is a selected dopant species. The calculations used the VASP code (developed by a group at the University of Vienna), a numerically efficient parallel implementation of density functional theory. The treated dopants include first row transition series elements as well as Al, Mg, Zn, and N (substituting for O). The monoclinic structure incorporates a cooperative Jahn-Teller distortion, and may be thought of as a deformed variant of the rhombohedral structure. The latter structure is unstable in the undoped system but can be stabilized by doping.

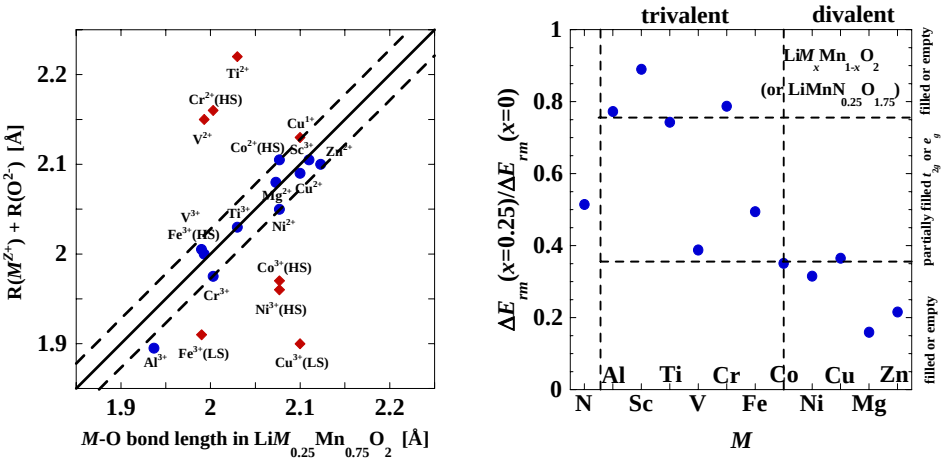

Left: comparison of M-O bond length from standard tabulated ionic radii with first principles calculations. Right: energy difference between rhombohedral and monoclinic structures at doping level x=0.25. Divalent dopants have the greatest stabilizing tendency.

The left-hand panel of the figure above reveals the oxidation state of transition element dopants in the rhombohedral structure, in which each dopant lies at the center of a symmetrical octahedron of oxygen

atoms. The plot compares the M-O bond length predicted based on standard tabulated ionic radii for different hypothetical valence states (abscissa) with the results obtained from first-principles calculations with the VASP code (ordinate). All dopants earlier than Co are trivalent, whereas Co and later dopants are divalent. The right-hand panel illustrates the stabilizing effect of the dopants on the rhombohedral relative to the monoclinic structure, where lower ordinate values indicate greater stabilization.

We observe that the divalent dopants, including late transition elements as well as Mg and Zn, exhibit the largest stabilizing effect. These results confirm the hypothesis that divalent dopants, which oxidize a Mn neighbor in the $LiMnO_2$ host material from 3+ to 4+, stabilize the rhombohedral structure by diminishing the concentration of trivalent Mn ions, thereby suppressing the Jahn-Teller distortion. The present analysis indicates that the oxidation state of the dopant is the primary factor that determines the stabilizing effect, and other factors such as solute size and local lattice relaxation effects are less significant. Up to now, our simulations have focused primarily on the layered lithiated manganese oxides, because of their simpler magnetic states than the spinel structure; however, one expects similar stabilizing effects of divalent dopants for the latter systems.

RIA Beam Dynamics Simulations (RIA-RND)

PI: Brahim Mustapha, Physics Division

The Rare Isotope Accelerator (RIA) is a proposed DOE scientific facility that will produce beams of radioactive and rare isotopes in amounts 10 to 1,000 times more than existing facilities. This ability to produce beams of short-lived isotopes will enable scientists to study the structure and interactions of nuclear matter with nuclei far from the stable regions. RIA will also have applications in nuclear medicine and in the materials and biological sciences.

RIA will consist of a driver accelerator to accelerate the primary beam at a target that produces radioactive isotopes and a postaccelerator to accelerate the selected isotopes for experimentation and analysis. Argonne's approach combines "fast gas catcher" technologies developed for the ATLAS facility with an advanced superconducting linear accelerator design (linac). This project is performing simulations to study the performance and limitations of the top two options for the driver linac – a baseline design with elliptical-cell cavities and a newer design with triple-spoke cavities. Of particular interest are the losses from various beam halo effects in the accelerator.

The TRACK beam dynamics code solves the equation of motion of beam particles through every element of the accelerator system: accelerating cavities, quadrupoles, solenoids, bending magnets, strippers, and so forth – about 1400 elements for the baseline design and 1200 for the triple-spoke

Baseline Design **Triple-Spoke Design**

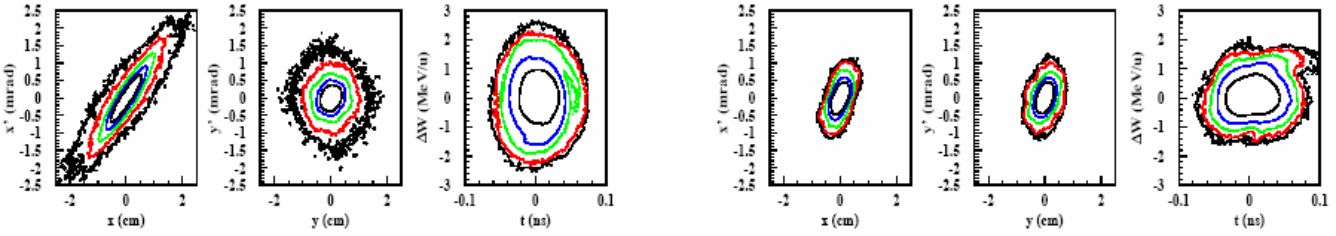

design. The linac is subdivided into a low-energy, a medium-energy, and a high-energy section separated by two stripping areas. The stripper simulation is based on the code SRIM including thickness fluctuations.

The performance of the whole accelerator is a function of the positioning of over a thousand elements and the accuracy of the phasing of the accelerating potentials. Hence very large parametric studies are needed to gauge the overall performance of each design. Deviations within expected tolerances of cavity and magnet alignment, target thickness, and so forth are generated according to Gaussian distribution. Such studies are possible only on a large cluster like Jazz, where a typical run uses 200 processors for 7 hours, and a minimal parameter study requires 10 runs (14,000 processor-hours) for one design option. The simulations with ^{238}U ions show that the transverse emittances are affected mainly by solenoid displacements and quadrupole rotation. Longitudinal emittance is affected largely by the radio frequency (RF) field and phase errors and by the fluctuation in the stripper thickness (see the figure above for one example). The baseline design shows a much larger sensitivity to RF errors and stripper thickness fluctuation compared to the triple-spoke design. Whereas no uncontrolled beam losses were observed for the triple-spoke design, the losses in the high-energy section of the baseline design exceed the 1 W/m limit set for hands-on maintenance. More careful tuning would be needed to make the baseline design competitive.

Impact of Aerosols on Regional and Urban Scale Climates (ATMOS_CHEM)

PI: V. R. Kotamarthi, Environmental Research Division

Reducing the uncertainty in evaluating the impacts of tropospheric aerosols on global- and regional-scale climate is critical for assessing the impacts of human activities on climate. This uncertainty can be attributed largely to the poor representation of sources, sinks, transport processes, and microphysical and chemical processes affecting aerosols as treated in the current generation of climate models. This project is pursuing a range of questions, from the effect of particulate pollution to the sensitivity of ozone concentrations to adjustments in fuel formulations.

For example, estimating the effect of wind-blown dust on aerosol concentrations in semiarid areas of the southwestern United States is crucial to defining the extent of suspended particulates and to distinguishing natural versus anthropogenic contributions to health exposure concerns. With data from a number of experiments in and around Las Vegas, Nevada, we used a dust-modified Community Multiscale Air Quality model (CMAQ) to calculate wind-blown dust and its transport characteristics in the Las Vegas Valley and the surrounding region. The objective was to understand the impacts of current and projected land use changes in Las Vegas and surrounding areas resulting from Bureau of Land Management land sales and the effects on particulate and gas phase pollution. Jazz was essential for these calculations. Simulations spanned 90 days, and each day took 75 minutes on 16 processors for each of the outer (regional) and inner (local) grids. These computations also generated very large datasets.

In another study the decrease in California emissions from 2000 to 2020 was shown to lead to significant changes in the contribution of boundary conditions to ozone in Las Vegas Valley. These sensitivity studies looked at several sets of parameters over 25 day spans, each day running in 75 minutes on Jazz, as above. The figure shows the difference between ozone generated for 2020 with emission reductions for Las Vegas included. The changes can be considered to be the contribution of California to the improvement of ozone levels in Las Vegas – approximately a 50% reduction. We note, however, that ozone production is quite complex and changes in the boundary conditions of ozone or its chemical precursors are not linearly related to the ozone reductions in the Las Vegas valley. These factors should be considered as an indicator of the potential impact of changes in California to air quality in Las Vegas.

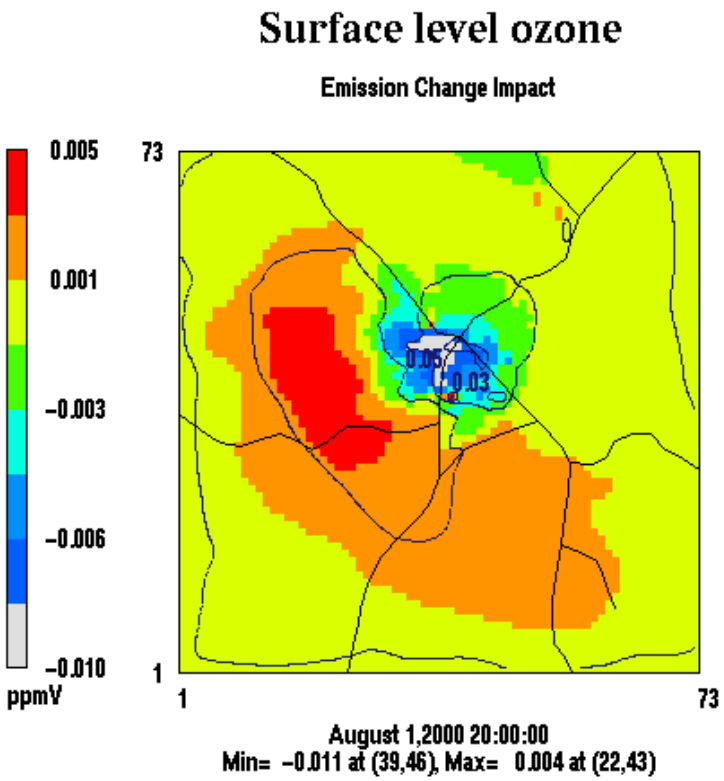

Neocortical Epileptiform Activity (NEOCORTEX_SIM)

PIs: Mark Hereld, Mathematics and Computer Science Division; W. van Drongelen, Department of Pediatrics at University of Chicago; H.C. Lee, Mathematics and Computer Science Division; and R.L. Stevens, Mathematics and Computer Science Division

This collaboration with the University of Chicago seeks to further the understanding of childhood epilepsy. At the level of the microcircuitry of the cortex, scientists hypothesize that epileptiform discharges are caused by malfunction of the neurons themselves (intrinsic properties), pathology in the neural networks (synaptic), or both. These seizures manifest themselves as discharges on an EEG, but researchers have only limited understanding in how the discharges on EEGs map to specific neural activity. This project is mapping specific brain activity during seizures to the observed EEGs and thereby helps explain exactly what is happening in the brain during these childhood seizures.

Simulations of large neural networks have the potential to provide unique insights in the study of epilepsy, from the effects of local changes in neuron environment, to the effects of large scale wiring anomalies. However, simulations with sufficient detail have typically been limited to cell counts that are far smaller than scales measured by typical experimental probes. The resources needed for problem solving in this domain call for very large scale computing, from teraflops to petaflops.

We are developing a simulation of a large network of neurons (current working size is of order of 100,000 cells). The cells and network wiring are configured to simulate electrical behavior of neocortical tissue samples, and in particular activity that corresponds to epileptic episodes in children. Parameters that define the model come from the literature, from collaborators' laboratory data, and from detailed studies of the behavior of simulation runs on Jazz. On 256 processors of Jazz, a 75,000-cell

simulation for 0.1 seconds with 10 microsecond time steps typically takes 18 hours to complete. The model allows us to test a very wide range of hypothetical causes and modulating effects, from very local properties of individual cells to general parametric sensitivity of cell populations and anomalies in network wiring. Such experiments are difficult, or even impossible, to perform either in laboratory samples or in the surgical setting.

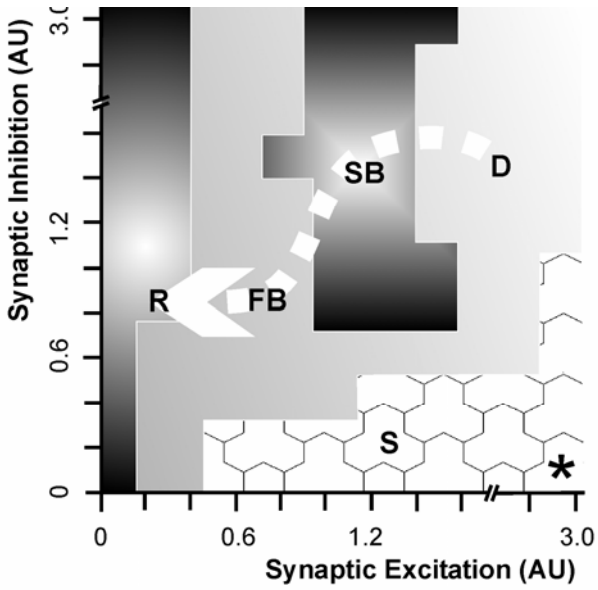

In 2004 we performed a fairly detailed study of the firing behavior of a patch of simulated cells as a function of two global parameters: inhibition and excitation sensitivities of the synaptic connections. The interesting conclusion from this study is that epileptiform activity corresponds to lower intensity of excitatory connections, a result that was not expected and in fact seems to be contrary to general opinion. The figure above shows a two-dimensional parameter scan presented as a map of activity patterns. A potential path to seizure onset (white arrow) is indicated, as well as the area of strong excitatory and weak inhibitory coupling (marked with *) commonly viewed as the area where seizures could occur. These results were presented at a recent meeting of medical and biological engineering researchers, and a description of the model has been published in *Neurocomputing.*

Further LCRC Research

Here we provide brief descriptions of many other active LCRC projects in FY2004. These 35 vignettes show the great breadth of fundamental and applied research being carried out with Jazz. Clearly, large-scale computing has become an integral element of advanced research at Argonne.

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

PI: V.R. Kotamarthi, Environmental Research Division

As part of our work in evaluating the impacts of tropospheric aerosols on climate, we have developed high-resolution models for aerosols and are assessing their effect on regional-scale climate. In 2004, we replaced the I/O interface with the parallel NetCDF software. We have also started using Weather Research and Forecasting IWRF) models and single columns of the Community Atmospheric Model. Work in 2005 will involve coupling the new model and WRF/MM5.

Illinois BioGrid (BIOGRID)

PI: Gregor von Laszewski and David Angulo, Mathematics and Computer Science Division

Much of genomics involves the study of proteins based on their primary structure (DNA sequence or amino acid sequence). Working with phylogenetics collaborators at the Field Museum of Natural History in Chicago, we are discovering feasible evolutionary relationships of given taxa by looking at differences in DNA sequences and determining the evolutionary tree. The IBG Workbench is a genomics toolkit being developed to facilitate high-throughput genomics studies. It includes FASTA, BLAST, and Smith-Waterman algorithms, all transformed to handle large batches of input sequences and also to take advantage of a distributed Grid computing environment. We are also developing Grid-enabled, parallel software algorithms that take raw proteome data from a mass spectrometer and derive the amino acid sequence of the proteins in the sample. Future work will target a collection tools for prediction of protein structure from sequence.

Development and Performance of the Community Climate System Model on Linux Clusters (CCSM2_1)

PI: Jay Larson, Mathematics and Computer Science Division

The Community Climate System Model (CCSM) is a fully coupled global climate model that provides state-of-the-art computer simulations of the Earth's past, present, and future climate states. Primarily developed at the National Center for Atmospheric Research under NSF funding, CCSM has recently received development support from DOE under the SciDAC program. Our first major task was to modify the latest version of CCSM's complex build system to support Jazz specifically. This effort was successful, and Jazz will be one of the recognized platforms for running CCSM in the next public release, joining machines at NCAR, NERSC, ORNL, LANL, and the Pittsburgh Supercomputing Center. The modifications can be easily extended to any Linux cluster using MPICH and thus greatly increase the number of people who can run their own experiments with CCSM. Finding the optimum processor configuration for CCSM on Jazz is also an ongoing effort. Because long (1000-year) climate model integrations often cannot be completed in a single submission to a machine's queue, the ability to checkpoint and restart is crucial for a climate model.

Large Scale Cumulative Reaction Probability Calculations (CRP)

PI: Mike Minkoff, Mathematics and Computer Science Division

The calculation of Cumulative Reaction Probabilities (CRPs) provides an exact measure of reaction rates. We are developing massively parallel methods for problems with high degrees of freedom - on the order of 10 degrees of freedom. These results can be used to evaluate statistical approaches for such large problems.

Kinetics of Enzymatic DNA Repair (DNA_REPAIR)

PIs: Aaron Dinner, University of Chicago; Stuart Rice, Argonne Directorate

The DNA repair enzyme uracil DNA-glycosylase (UDG) removes bases that are normally restricted to RNA (uracil) from DNA. Misincoporation of deoxyuridine and deamination of cytosine cause several hundred uracil bases to arise in the DNA of a human cell every day. Except in cases where it replaces a thymine, such uracil is premutagenic and must be eliminated to maintain genomic integrity. UDG initiates this process by cleaving the glycosylic bond of deoxyuridine in DNA to yield an apyridiminic site. We are performing dual-topology free-energy perturbation simulations in order to resolve apparent discrepancies between simulations and experiments concerning the role played by specific substrate phosphate groups.

Excited States of the Pion in a Dyson-Schwinger Approach (DSE_EXC_PION)

PI: Andreas Krassnigg, Physics Division

The formalism of Dyson-Schwinger equations provides a nonperturbative way to study QCD at low energies, for example, hadron spectra and other properties. Compared with quantum mechanical models, the Dyson-Schwinger approach makes a covariant and also field theoretical treatment possible, once a truncation scheme is found, which preserves the symmetries of the theory. At least one such scheme exists, which has been studied over the past few years and was found to be successful in describing pseudoscalar and vector mesons. Present and future studies also deal with other mesons and with baryons. In addition, we are studying steps beyond the present truncation.

Earth System Modeling Framework (ESMF)

PI: Jay Larson and Robert Jacob, Mathematics and Computer Science Division

The scientific challenge of developing advanced Earth system applications is daunting. Independently developed components typically have incompatible interfaces or may be written in different computer languages. The high-performance computer platforms required by numerically intensive Earth system applications are complex, varied, rapidly evolving and multipart systems themselves. The goal of the Earth System Modeling Framework project is to build high-performance, flexible software infrastructure to increase ease of use, performance portability, interoperability, and reuse in climate, numerical weather prediction, data assimilation, and other Earth science applications. Our aim is to create a framework usable by individual researchers as well as major operational and research centers, engaging the user community in its development.

ASCI Flash Project (FLASH)

PI: Robert Rosner, Physical, Biological and Computing Sciences; Katherine Riley, Mathematics and Computer Science Division

The ASC/Alliances Center for Astrophysical Thermonuclear Flashes aims to solve the long-standing problem of thermonuclear flashes on the surfaces of compact stars such as neutron stars and white dwarf stars, and in the interior of white dwarfs (i.e., Type Ia supernovae). The Center researchers are developing a parallel code called FLASH, which shows excellent scalability up to thousands of processors. (Typically, active science research is done on hundreds to thousands of processors.) In this project we are using Jazz primarily for development, performance testing, and smaller science runs.

Grid Physics Network (GRIPHYN)

PI: Michael Wilde, Mathematics and Computer Science Division

We have been testing the feasibility of integrating the Jazz cluster into the large Data Grids required by the ATLAS high-energy physics experiment and the Sloan Digital Sky Survey, both part of the GriPhyN (Grid Physics Network) project. A significant aspect of GriPhyN progress in the past year has been the use of a common “virtual data framework” developed at Argonne, the University of Chicago, and the University of Southern California Information Sciences Institute to express the scientific workflow of multiple projects in an uniform manner and to capture the provenance of results across a Grid using a common schema. This framework, named “Chimera,” has been accepted by most GriPhyN experiments and is being used beyond GriPhyN as well, initially in MCS computational biology projects. The International Virtual Data Grid Laboratory project, a sister project of GriPhyN, is currently constructing a new Grid, named “Grid3.” Having Jazz as a resource on Grid3 will enable authorized Argonne Jazz users to extend their computational resources with the other resources of Grid3, currently targeted at about 500 CPUs.

Investigation of Nano-Particulate Production for Homogenous Charge Compression Ignition Combustion (HCCI)

PI: Adrian Tentner, Nuclear Engineering Division

Considerable effort has been devoted to reducing harmful emissions generated by internal combustion engines. Homogenous Charge Compression Ignition (HCCI) is a combustion concept that has the potential to provide high thermal efficiency while reducing harmful emissions. We are conducting a modeling and simulation effort in conjunction with an experimental effort to analyze the factors that control the HCCI combustion. Our simulation work uses computational fluid dynamics (CFD) and chemistry kinetics codes to describe fuel/air mixture zones, complementing experimental work. The results of the CFD model also provide insight into placing the sampling probe and the x-ray beam for analysis of HCCI combustion experiments. Hence CFD analysis will allow the experimenters to fine-tune the experiments for maximum effectiveness.

Discharge Produced Plasma Hydrodynamics and Radiation Transport for EUV Lithography (HEIGHTS-EUV)

PI: Vitali Morozov, Energy Technology Division

Discharge produced plasma (DPP) devices are a potential light source for extreme ultraviolet (EUV) lithography. To simulate the environment of the EUV source and to optimize its output, we have developed an integrated model to describe the hydrodynamic and optical processes that occur in the DPP. The model addresses three main features: plasma evolution and magnetohydrodynamics (MHD) processes, detailed photon radiation transport, and interaction between plasma/radiation and materials. It uses the total variation diminishing scheme in the Lax-Friedrich formulation to describe magnetic compression and diffusion in a cylindrical geometry. Several models of radiation transport, based on both direct methods and Monte Carlo implementations, are coupled with the MHD. The calculation of radiation transport depends on the accuracy and the level of details of the optical coefficients of absorption/emission, which are obtained from a collisional-radiation equilibrium (CRE) model. The calculation of absorption for free-free, bound-free, and bound-bound radiation transitions is based on detailed atomic physics calculations.

Physics of the High Temperature Gas Cooled Reactor (HTGR)

PI: Taek Kim, Nuclear Engineering Division

The very high-temperature gas-cooled reactor (VHTR) is an important element in the U.S. Department of Energy program to revitalize nuclear power for the U.S. suite of energy generation options. Construction of the first VHTR, called the Next Generation Nuclear Plant (NGNP), is planned for the next ten years. In order to have confidence in the NGNP design, computational tools used for the design studies must be validated. As a first step, an assessment of lattice physics capabilities has been performed for prismatic fuel assembly designs. To support this validation, we used the MCNP4C Monte Carlo code for generating reference solutions. Since the MCNP4C code requires huge computation time, however, the Jazz cluster was necessary for this work. During FY2004, we considered two Monte Carlo models: one with regular lattice distribution and one with random distribution of the coated fuel particles in a graphite matrix. The results for these models were found to be self-consistent. Comparison of the models, however, revealed differences larger than those attributable to statistical uncertainties. Specifically, we determined that the MCNP4C eigenvalues obtained with the regular lattice distributions are consistently smaller than those obtained with random distributions, with differences between the models as high as ~500 pcm. These differences increase with the amount of fuel and the kernel size, and the magnitude is also dependent on the fuel isotopic vector.

LCRC Application Engineer Support, Benchmarking and Tutorial Work (LCRC-APP-ENGR, LCRC-BENCHMARK, LCRC-TUTORIALS)

PIs: Katherine M Riley and Michael Dvorak, LCRC

We are helping the Jazz community with a variety of activities, for example, design and development of utilities for applications. We also regularly help users with performance evaluations on Jazz to guide applications development and expectations of the system. In addition, we devote a small amount of time to hands-on help during LCRC tutorials.

Use of Predictive Lithostratigraphy to Significantly Improve the Ability to Forecast Reservoir and Source Rocks (LITHOSTRATIGRAPHY)

PI: Thomas Moore, Energy Systems Division

We have been developing an implementation of a climate model for paleoclimate applications that detect the influence of orbital cycles. The FOAM model, although it has been used for paleoclimate applications, was not sufficient for our modeling goals. FOAM therefore was heavily modified to accommodate higher grid resolution. Subsequent climate model runs on Jazz showed that FOAM was primarily limited in equatorial regions because of the "cold tongue" problem. This problem arises when cold equatorial waters impact equatorial precipitation rates, which in turn can lead to incorrect interpretations. The model is now being improved to correct this problem.

Numerical Simulations of Magneto-Hydrodynamical Turbulence (MAGNETO)

PI: Robert Rosner, Physical, Biological and Computing Sciences; Alexandros Alexakis, Mathematics and Computer Science Division

We are using a pseudo-spectral code to simulate magnetohydrodynamical turbulence driven by a random forcing field in the presence of a mean magnetic field. The goal is to systematically determine how the presence of a mean magnetic field affects the properties of turbulence in a plasma. In particular we are interested in the anisotropy in the turbulence resulting from the imposed mean magnetic field that has been the focus of recent theoretical work in plasma turbulence. To date, we have performed parameter studies for a high-resolution three-dimensional grid studying the impact of different mean magnetic field values.

Computational Fluid Dynamics Modeling and Computational Science Studies Model Simulation, Comparison and Testing in Urban Boundary Layer Applications (MANHATTAN)

PI: Michael Lazaro, Environmental Assessment Division

In collaboration with EPA scientists, we are exploring the possibility of developing or adapting the products from CFD simulations to support rapid exposure and risk models to guide urban emergency response and emergency management for chemical, biological or radiological attacks and accidents. This work is expected to provide opportunities for collaboration with the Department of Homeland Security's New York City Urban Dispersion Program, which will be able to use our CFD applications to aid the ongoing air monitoring strategy and intensive field measurement studies. The CFD applications will also support development of simple reliable models to improve and provide practical tools for emergency first responders.

Multimethod Linear Solvers (MULTIMETHOD-SOLVERS)

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

The solution of large-scale, nonlinear PDE-based simulations by using implicit and semi-implicit methods typically depends on the performance of sparse linear solvers, which may be invoked at each nonlinear iteration. We are exploring polyalgorithmic multimethod solvers in the context of several large-scale applications involving computational fluid dynamics to potentially improve the execution

time and reliability of linear system solution. We are considering composite solvers, which provide reliable linear solution by using a sequence of preconditioned iterative methods on a given system until convergence is achieved. We are also exploring adaptive solvers where the solution method is selected dynamically to match the attributes of linear systems as they change during the course of a simulation.

Molecular Dynamics Simulations of Nuclear Receptor Proteins (NUCLEARRECEPTORS)

PIs: M. Cunningham and K. Nettles, Biology Division

The nuclear receptor (NR) superfamily of transcription factors responds to lipophilic signals including steroid hormones, bile acids, fatty acids, and oxysterols. Currently these therapeutic targets represent 10-20% of the $400 billion spent worldwide on prescription drugs. The NRs represent primary targets for cardiovascular, metabolic, and endocrine disorders, as well as cancer and the steroid anti-inflammatroy drugs. Our x-ray crystallography studies have identified the mechanism through which NRs activate transcription, which involves the ligand-mediated recruitment of transcriptional coactivator proteins. A molecular switch in the carboxy terminus of the NR is regulated by ligand, and forms one side of the coactivator binding site. Molecular dynamics simulations will be employed to understand the regulation of this switch.

Optimization of Many-Body Wavefunctions (NUCLEARSTRUCTURE)

PI: Richard Chasman, Physics Division

The purpose of this work is to obtain a good description of low-lying nuclear states. We need to go beyond the mean-field approximation when the nucleus of interest is transitional in some sense, for example, transitional between a spherical shape and a deformed shape.

Parallel Tools Performance Testing (PARALLEL-TOOLS)

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

The parallel tools project encompasses development of several software packages, including PVFS, PVFS2, Parallel-netCDF, MPICH, and ROMIO. The size and speed of Jazz makes it an excellent platform to exercise the scalability and performance of our parallel tools. Our early investigation into PVFS scalability yielded unexpected results, which motivated our making improvements to PVFS. Jazz has also been a good cluster to evaluate our new MPICH collective operations and to study data-type handling. Since Jazz uses PVFS for its high-performance file system and since many applications leverage MPI and MPI-IO, improvements to these packages have immediate and direct benefits to other Jazz users.

Petascale Data Quest (PDQ)

PIs: Michael Wilde, Jens-S. Voeckler, and Veronika Nefedova, Mathematics and Computer Science Division

The NCSA Alliance Expedition “A PACI Petascale Data Quest” aims to achieve the technical advances, technology deployment, and application integration required to enable data-intensive applications to harness Alliance and TeraGrid resources, thus transforming these systems from potential Grid resources into demonstrably effective Grid powerhouses. We are using the LCRC cluster to study the virtual data

automated workflow for two scientific problems: running the computational biology tools BLAST and PFAM on frequently updated genomes, and producing data to be made available as a pubic service.

PETSc Software Library Testing (PETSC)

PIs: Barry Smith, Satish Balay, and Richard Katz, Mathematics and Computer Science Division

PETSc (Portable, Extensible Toolkit for Scientific computation) provides tools for the parallel (as well as serial) numerical solution of PDEs that require solving large-scale, sparse nonlinear systems of equations. We have found the Jazz system to be cost effective and very stable for the PETSc software libraries. We have used Jazz both for basic benchmarking and for several ongoing algorithm and application development projects. For example, one project (conducted in collaboration with the Chemistry Division) involves eigenvalue calculations with the density-functional tight-binding method; the objective is to improve the parallel algorithms and implementation required by electronic structure calculations. A second project involves geodynamic simulations of subduction with a code written using PETSc and containing more physical features than have ever been incorporated before.

Family-Specific Substitution Matrices for Classification of Very Large Proteins (PROTEIN)

PI: Richard Vilim, Nuclear Engineering Division

The human BRCA1 protein contains 1,863 amino acid residues. Learning how BRCA1 interacts with other proteins is an important step in understanding the pathology of the familial forms of breast and ovarian cancer. Despite the relative importance of BRCA1, the structures of the inner regions of this protein are largely uncharacterized through chemical studies. Using sequence homology, however, researchers have determined replicate domains within the BRCA1 protein. We are developing and testing algorithms to find the substitution matrix and position weights that classify the protein training set sequences according to the values of their attribute tags. Substitution matrices provide a way to determine the likelihood that an amino acid replacement will not affect the structure and function of a protein. On Jazz we were able to exhaustively test our algorithm on immunoglobulin sequence data and confirm superior classification performance compared with generic substitution matrix methods. We ran a statistically meaningful number of cases using the cross-validation method and showed that our algorithm performs more accurate classification than do algorithms based on aggregate classification.

Application Performance Study for a New Program Execution Model (PXM_APPS)

PI: Mark Hereld, Mathematics and Computer Science Division

We are characterizing the computational and communication performance of application kernels to generate performance models for various program execution models. Our focus is on large networks of biological neurons, wired to model neocortical slices. Specifically, this project involves development and testing of pNeo, which is a simulation model based on our production code, pGENESIS.

Monte Carlo Burnup Calculations for Nuclear Reactors (REBUS-MCNP4)

PI: Sai-Chi Mo, Nuclear Engineering Division

The goal of this project is to perform parallel Monte Carlo burnup calculations for analyses related to conversions of research reactors from the use of highly enriched uranium to low enriched uranium. To

date, we have developed a computational system that combines a reactor burnup code, REBUS, with a Monte Carlo code, MCNP4C, to solve complicated time-dependent reactor burnup problems.

Regional Climate Simulation (REGCLIM)

PI: Jay Larson, Mathematics and Computer Science Division

We are using a mesoscale weather model to simulate climate at the regional scale over the continental United States. Our chief goals include decadal-scale simulations and model sensitivity experiments.

Use of Computational Fluid Dynamics for Analysis of Core Designs under the Reduced Enrichment for Research and Test Reactors Program (RERTR_CFD)

PI: Patrick Garner, Nuclear Engineering Division

The redesign of reactor cores to use low-enriched uranium fuel requires a series of nuclear physics and thermal-hydraulics calculations in order to assure that the fuel produces the maximum required power while still being adequately cooled. The spatial variation of heat generation requires the use of detailed three-dimensional calculations rather than just dealing with averages. In the present project we are expanding the level of spatial detail by using the STAR-CD code for computational fluid dynamics calculations.

Fundamental Forces in Sediment Transport (SEDIMENT)

PI: Gary Leaf, Mathematics and Computer Science Division

Particle transport is a process common to many physical systems. One aspect of such transport involves particles in suspension in oscillatory flow, for example, wave-generated sediment transport. The principal factors in this process are the lift and drag on the particles in the presence of such flows. Because the forces involved are so small, experimental approaches are severely limited. Computer simulations, on the other hand, provide a means for obtaining these measurements. For our initial study, we considered a fixed sphere near a bounding wall subject to a single frequency oscillatory flow, and measured the lift and drag experienced by this particle as a function of time. The time behavior of these net forces determines whether a particle in suspension will remain in suspension. The net lift on a particle, at any instant in time, is the result of the viscous contribution and the pressure acting on the particle. The numerical simulations of the 3D periodically forced were based on approximating a nondimensional Navier-Stokes equation with the spectral element code Nek5000, which achieves linear scaling to at least 64 processors for the problem sizes required for this project. For the lift in an oscillatory flow, in the large-gap short-period regime the lift is dominated by its viscous contribution, whereas in the small-gap long-period regime the lift is dominated by the pressure.

Simulated Maximum Likelihood Estimation of Returns to Schooling (SIMAXLIKE_SCHOOLING)

PI: Gayle Boyd, Decision and Information Sciences Division

We are using simulated maximum likelihood methods to estimate structural dynamic programming lifecycle schooling and consumption decision models for individuals in the United States. Our goal is to obtain a model that describes how individuals make their schooling decisions, separating the roles

played by uncertainty, costs, expectations, and credit constraints in the decision. By linking this decision to the expected consumption decision, we can model the decisions under different assumptions about how credit markets operate. The estimated model can then be used to simulate, for example, the effect that different policies (such as tuition subsidies) would have on schooling and consumption decisions.

Large-Scale Parallel Simulation of Subduction Zone Geodynamics (SUBDUCTION)

PIs: Matthew Knepley and Richard Katz, Mathematics and Computer Science Division

Our goal in this project is to better understand the role of subduction zones as a link between three of the Earth's chemically distinct planetary reservoirs: the hydrosphere, the crust, and the mantle. Long-term evolution of these reservoirs depends critically on the flux of volatiles, sediments, and oceanic lithosphere through the 'fractionation machine' that is subduction. Volatiles liberated from the down-going slab enter the overlying mantle, lower the melting temperature of the rock, and lead to magma genesis. Surface observables such as magma composition, spatial distribution, volume, and temperature are determined here at depth. We are seeking to run a very large 2D, non-Newtonian thermal convection simulation of a subduction zone. We will compare the relative efficiency of multigrid and sparse direct preconditioners on medium sized problems of immediate interest, with an eye to extremely large three-dimensional simulations in the near future.

Systems Administration Tasks on the LCRC Jazz Cluster (SUPPORT)

PI: Susan Couglan, Mathematics and Computer Science Division

We are carrying out various systems administration tasks on Jazz. These include debugging problems, testing software upgrades and modifications, and running scaling and performance tests.

Numerical Optimization on Jazz (TAO)

PI: Steven Benson, Mathematics and Computer Science Division

Semidefinite programming has been an actively studied area of numerical optimization. One reason for the high level of interest is that applications involving this class of optimization problem have been found in fields as diverse as structural design, control theory, and combinatorial optimization. A second reason is that although interior-point methods adopted from linear programming have proven reliable on small and medium-sized semidefinite programs, the computational and storage demands of these methods easily exhaust the resources of most computers and limit the size of problems that can be solved. Much of the research in this field focuses on solving medium-sized problems more quickly and on solving large-scale problems by any means possible. We are using Jazz to develop a parallel dual-scaling algorithm for semidefinite programming.

A Framework to Support Comparative Analysis of Genomes (THE_SEED)

PIs: Robert Olson and Ross Overbeek, Mathematics and Computer Science Division

We are implementing a system to support comparative analysis of genomes. In the process we find it important to compute similarities between protein and DNA sequences, detection of motifs, and so forth. These are computationally intensive, but straightforwardly parallelized, computations. We will attempt to maintain a more-or-less complete collection of similarities between sequences until we can accurately replace the collection with protein families (which should dramatically simplify the situation). We

expect the system we are constructing to become widely used by genome projects as a basic framework for investigating genomes.

Terascale Simulation Tools and Technologies (TSTT)

PI: Paul Fischer, Mathematics and Computer Science Division

The Terascale Simulation Tools and Technologies (TSTT) project is part of DOE's Scientific Discovery through Advanced Computing strategic initiative. The goal of TSTT is to provide cross-cutting technologies for large-scale simulation-based science projects such as climate modeling, fusion reactor design, accelerator design, and magnetohydrodynamics of stellar atmospheres. Our focus is on the development and deployment of high-order numerical discretizations that will facilitate multiscale and multiphysics simulations. Validation of these tools under realistic settings often requires large-scale simulations involving millions of grid points and millions of time steps or iterations. Most of the development work will be done using the simulation code Nek5000 (1999 Gordon Bell Prize for scaling to 4096 processors on ASCI-Red). Specific projects include the development and test of high-order discretizations for anisotropic diffusion in tokamaks, incompressible MHD for verification of liquid metal experiments in the United States and abroad, solution of Maxwell's equations for nanophotonics, turbulent multiphase flows, pebble bed reactor analysis, and biological flows

Vortex Matter in Layered Superconductors (VORTEX_MATTER)

PI: Alexei Koshelev, Material Sciences Division

We are exploring the phase diagram of vortex matter in layered superconductors using large-scale numerical simulations. The project consists of two components. First is development of a phase diagram of vortex chains. A small magnetic field applied at a finite angle with respect to crystal axes generates isolated vortex chains. The chains have quite rich internal structure depending on the magnitude and tilt angle of the magnetic field and the temperature. Typically, at small temperatures and large tilt angles, the vortex chain is composed of two perpendicular sublattices: a chain of Josephson vortices, oriented along the layers, and a chain of point (pancake) vortices, oriented perpendicular to the layers. We compute the ground state configurations at different parameters and construct the phase diagram of an isolated vortex chain. Second is thermodynamics of vortex lattice melting in very anisotropic layered high-temperature superconductors. Because of the large amplitude of thermal fluctuations, the vortex lattice in very anisotropic high-temperature superconductors melts at very small magnetic fields, and most of the field-temperature phase diagram is occupied by the vortex liquid state. In spite of considerable progress in recent years in theoretical understanding of vortex lattice melting, the quantitative theory is still absent. The major technical difficulty is the presence of two types of competing interlayer coupling: long-range magnetic coupling and neighboring-layer Josephson coupling. Our simulations are planned to include both these couplings.

Appendix: Additional Usage Statistics

The LCRC is a production supercomputing facility, with login accounts for each authorized user, each user assigned to one or more approved projects, and each project allocated a certain number of resource hours and a specific priority. Jazz system usage is tracked at the granularity of individual jobs, users and projects. For certain commercial applications and tools, the LCRC manages a pool of licenses that are "checked out" when the software is used. However, we do not track the usage of particular applications. In this appendix, we display and discuss usage statistics that illustrate utilization of the system.

One must keep in mind the challenge of effectively managing a large parallel computer used by a large and diverse science and engineering community. In the LCRC user community, some projects require hundreds of processors for extended periods, while other projects need to run a large number of jobs involving 50–100 processors. Some projects are making short turnaround test runs on small numbers of processors, while others are doing background work on single processors. Also, the resources needed by any project change over time, as new models are developed, tested, and run. For example, the period just prior to manuscript due dates for the major society meetings can be very intense for scientists in a particular domain, followed by a lull for a while. Consequently, no single component of the demand for LCRC resources is constant, and the challenge is to understand demand and usage and offer a combination of job and resource management tools that optimizes the value of the Jazz system to the Argonne community. The LCRC's Allocation Committee and Systems Administration staff have done an outstanding job of meeting user needs with the Jazz resources available.

Users and Projects

Since production began in April 2003, the number of user accounts and projects on Jazz has grown rapidly. Any staff member may obtain a startup account for 1000 hours, to become familiar with the Jazz system and evaluate its applicability to their work. Aside from these startup efforts, all computer time is allocated to projects, rather than individuals, where a project is focused on a specific technical activity or avenue of investigation. A project usually has several scientists or engineers working on it, and a person may be active on more than one project. By the end of FY2004, there were over 250 active user accounts, and the number of active LCRC projects has increased rapidly to 66 (see chart).

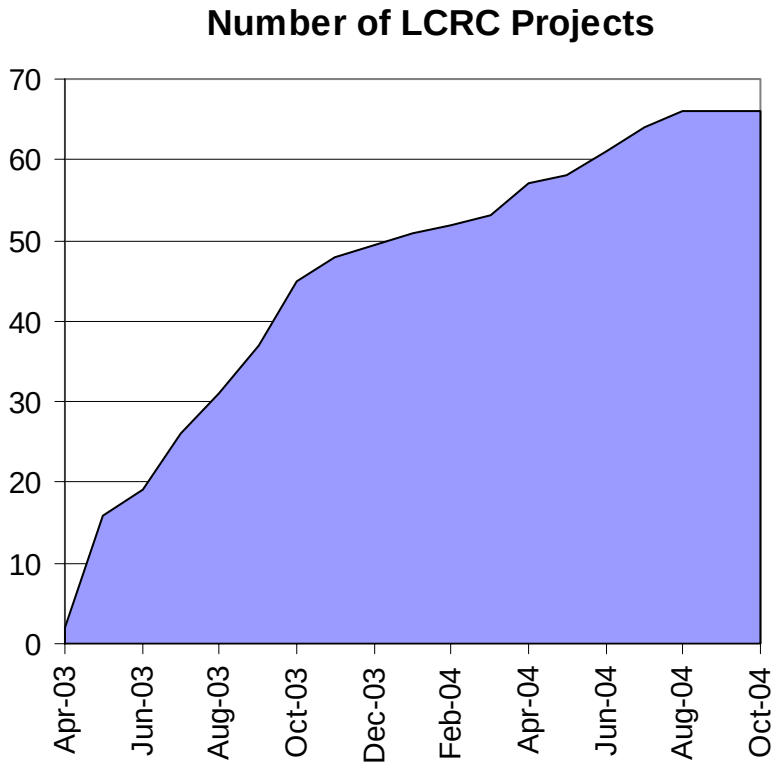

Usage by Discipline

Jazz supports research in a very wide range of disciplines, as indicated by the chart below. Here the size of the slice is proportional to the fraction of the Jazz resource used by all projects in a particular area, during FY2004. The sizes are proportional to the total number of node hours used. These fractions change over time as new projects are added. Initially more people in the physics and applied math areas had prior experience with parallel computers and made up a large fraction of the early LCRC user community. As more Argonne scientists and engineers have learned how to take advantage of Jazz in their research, however, the fraction of projects from other areas has increased.

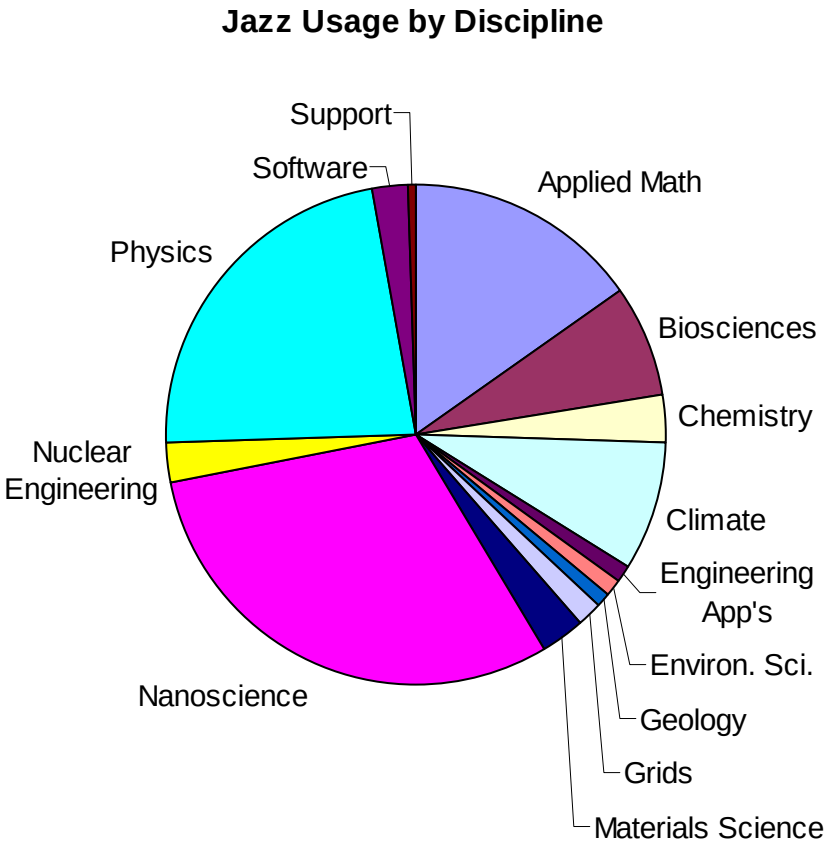

Another way to look at this information is to examine it over time. The following chart shows how the pattern of usage varied from month to month in FY2004 It demonstrates a number of interesting characteristics of large computer usage. In this chart specific discipline usage tends to both increase and decrease from month to month in ways unpredictable by those not involved in the specific projects, for example obligations to other projects, professional conference schedules, and the arrival and departure of postdocs.

Also projects often tend to have a lengthy development phase early in their lifecycle as researchers learn to use the system, project software and data are installed, and codes are ported. LCRC's tutorial and consulting efforts help shorten this cycle as much as possible. During the early phases projects use little

time. Once everything is place there is a ramp up in usage as problems of scale are addressed. This was the case with the Biosciences effort in early FY2004, and previously with Physics and Nanoscience efforts in FY2003.

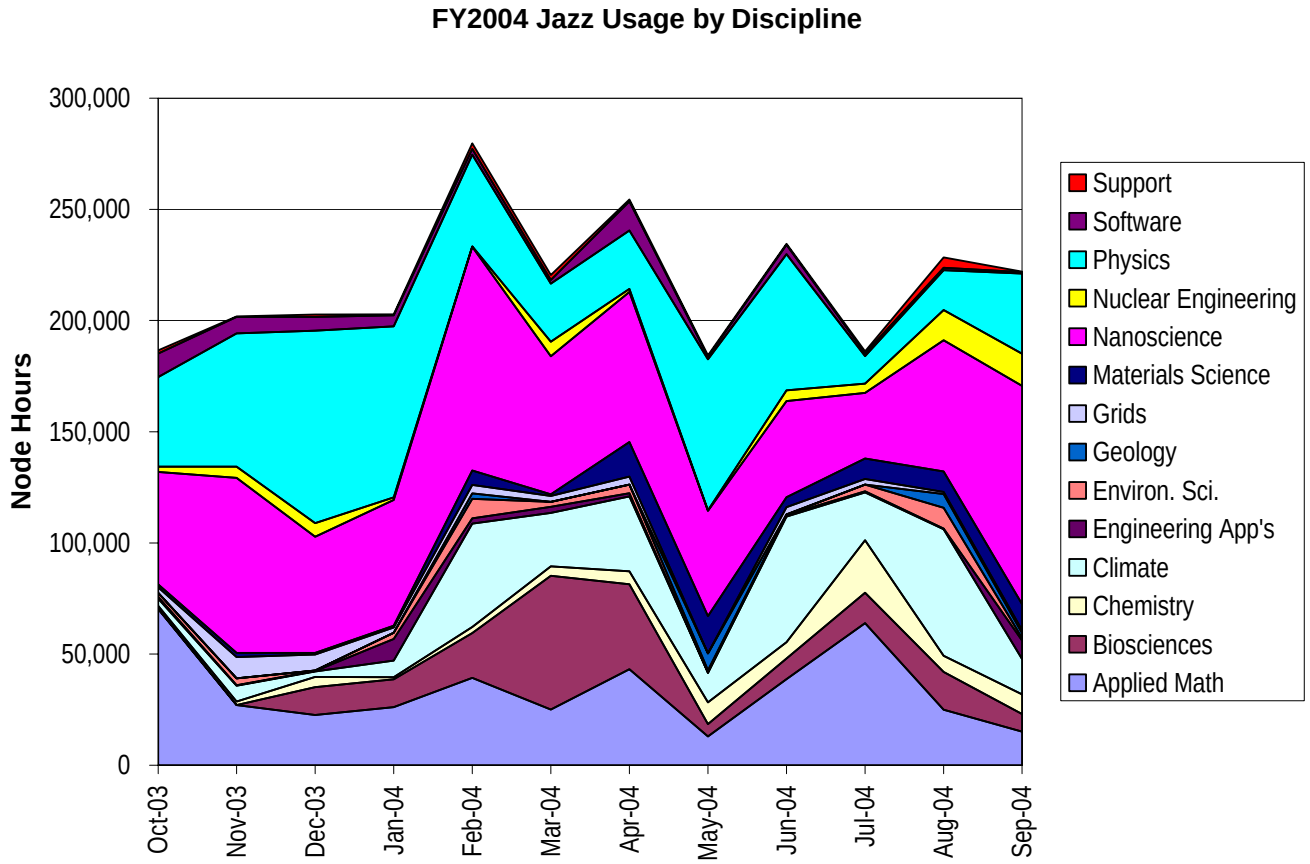

Job Size Characteristics

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

Data about the number of computing nodes that jobs used provides an interesting view of how the system is being utilized and the degree of scalability or parallelism commonly utilized on Jazz. The chart below shows the amount of computing time used by jobs of various sizes in FY2004. Jobs that ran on only 1 node accounted for only 6% of the total computing time used on Jazz. Most of the resource, 64%, was used by jobs that needed 32 nodes or more.

Clearly, the job size on Jazz is quite varied. Indeed, some users have used the entire machine with one job, while other users have run many, many 16-node jobs. Over time, the pattern of utilization is moving toward larger jobs. On a system like Jazz that is heavily used, job turnaround is related to the number of nodes requested. That puts a practical bound on the turnaround for very large jobs. To compensate, users may reserve a portion of the system for a specific period during a night of weekend to carry out large computations or to meet an urgent computing need. All users are notified in advance of these reservations.