NERSC NEWS

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Berkeley Lab

Office of Science
U.S. DEPARTMENT OF ENERGY

SEPTEMBER 2007

Goodbye Seaborg

IBM SUPERCOMPUTER WORKHORSE THAT HAS PRODUCED GREAT SCIENTIFIC RESULTS WILL RETIRE

Seaborg, the IBM supercomputer that has tackled some of the most challenging problems in astrophysics, climate research, fusion energy, chemistry and other scientific areas, is set to retire next January after seven years of serving NERSC users.

Since August 2001 Seaborg has provided a little over 250 million CPU hours to some 3,000 scientific users. Included in these numbers are 26.5 million CPU hours for 22 projects from the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, which was created by the DOE Office of Science to support large-scale, high-impact projects.

At the time of its purchase, the IBM SP RS/6000 represented the largest single procurement in Berkeley Lab's history, costing $33 million for the system and a five-year contract. Named after the Nobel prize winner and Berkeley Lab chemist Glenn Seaborg, the supercomputer began its service in January 2001 and was ranked No. 2 on the June 2001 TOP500 list, a semi-annual ranking of the world's most powerful supercomputers.

After all these years, Seaborg remains on the TOP500. The ranking released in June of this year put Seaborg at No. 158.

"From the day it went online, Seaborg has been the scientific workhorse for many of the most important computational science projects undertaken within the Office of Science," said Michael Strayer, Associate Director of Advanced Scientific Computing Research within the Office of Science. "Now the system is ready to retire at a ripe old age of seven, many researchers will remember Seaborg with appreciation."

Seaborg has been called on to run jobs that no other supercomputers could handle. For example, the scientists at the Max Planck Institute for Gravitational Physics relied on Seaborg to simulate colliding black holes and the resulting gravitational waves because Seaborg provided the memory required for such a large-scale project. The research, led by Ed Seidel in 2002, used 1.5 terabytes of memory (a third of what Seaborg could offer) and 2 terabytes of disk space for each run.

SEABORG'S SUCCESSES

These runs typically consumed 64 of the large memory nodes for 48 wall-clock hours at a time. In a two-month time period, Seidel's research used up 700,000 CPU hours.

Seaborg also ran the first three-dimensional simulations of the collapsar model, developed by Stan Woosley at UC Santa Cruz, that sought to explain the origin of gamma ray bursts.

In the collapsar model, the iron core of an aging star runs out of fuel for nuclear fusion and collapses from its own weight, creating a black hole or a dense neutron star. Material trying to fall onto this object forms a hot swirling disk and a narrow jet, which shoots out of the star in less than ten seconds at nearly the speed of light. When this jet erupts into interstellar space, it creates a "fireball" of gamma rays — the highest energy, shortest wavelength form of electromagnetic radiation. The rest of the star explodes as a supernova,

continued on page 5

Kudos for NERSC

A new operational review by DOE gives high marks to the supercomputing center

NERSC received praises from its first-ever Operational Assessment Review by a committee of scientists, national research facility leaders and managers at the DOE Office of Science.

The review, conducted on August 28, fulfills a requirement by the Office of Management and Budget to evaluate national research facilities for capital planning purposes. NERSC, Oak Ridge National Laboratory and the Molecular Science Computing Facility (MSCF), part of the Environmental Molecular Sciences Laboratory at the Pacific Northwest National Laboratory, underwent reviews this year. Argonne National Laboratory will be reviewed next year.

"The committee feels that we are doing a good job and that our operations are mature and efficient," said Bill Kramer, NERSC's General Manager. "We exceeded

continued on page 3

Water Effect

SCIENTISTS EXPLORE HOW WATER MOLECULES AFFECT FUEL CELL MEMBRANE PRODUCTIVITY

Three NERSC users have published a paper detailing the effects of temperature and hydration during a critical step in fuel cell energy production.

The work, published in the Journal of Physical Chemistry B in June, examined the physics of a type of DuPont-made polymer membrane called Nafion. In the polymer electrolyte membrane fuel cell (PEMFC), a well-hydrated membrane is needed to allow protons to easily pass through after a chemical reaction splits the hydrogen fuel into protons and electrons. The electrons, which can't go through the membrane, would instead move along an external circuit and pro-

continued on page 2

Open for Service

MORE NERSC RESEARCHERS CAN TAKE ADVANTAGE OF NEW GRID TOOLS

NERSC has deployed new services that make the Open Science Grid (OSG) more widely available to its researchers, who could use the grid to easily manage their work at multiple computing sites.

NERSC's IBM p575 POWER5 system named Bassi and the Opteron cluster named Jacquard are the latest additions to a list of systems joining the OSG, whose distributed computing network is made up of 53 sites at universities, national labs and computing centers in the North and South America, Europe and Asia. Another NERSC system, the SGI Altix 350 named DaVinci, joined the OSG in April this year.

The OSG will save valuable time and reduce headaches for scientists who carry out their research at several computing facilities. Instead of dealing with different authentication processes and software at each site, the scientists can go through the OSG to manage their computing jobs and file data at various computing sites.

"You can submit, manage and coordinate jobs across NERSC systems and other grid sites using a common interface directly from a client machine," said Shreyas Cholia, a member of NERSC's Open Software and Programming Group who has worked on developing the software for linking the NERSC systems to the OSG. "You can move data back and forth between various systems."

Cholia described the new services and how they work during a presentation at the annual meeting of the NERSC Users' Group this month.

The OSG marshals the software, computing, storage and other expertise of many of its 119 member sites which, in addition to the sites that provide computing resources, include research institutions that use the OSG and software developers responsible for building the

continued on page 5

Fuel Cell *continued from page 1*

duce electrical power, eventually meeting up with protons and oxygen at the other side of the membrane to form water.

The research is part of President Bush's Hydrogen Fuel Initiative, which aims to develop commercially viable hydrogen fuel cells. Using this clean and efficient technology would help to reduce the world's reliance on fossil fuel and lessen greenhouse gas emission.

"In order to advance PEMFC technology, it is important to understand the nanostructure of the membrane and the fundamental processes and physics governing protons and small molecule transport in these membranes," wrote the paper's three authors from the Pacific Northwest National Laboratory (PNNL), Arun Venkatnathan, Ram Devanathan and Michel Dupuis.

The paper, "Atomistic Simulations of Hydrated Nafion and Temperature Effects on Hydronium Ion Mobility," graced the cover of the journal's June 28 issue.

The researchers set out to create simulations that examined the impact of hydration and temperature on the positively charged hydrated protons and water molecules. Understanding these dynamics could lead to polymer membranes that are better engineered for transporting protons while controlling electrode flooding by the water molecules. One of the goals is develop PEM membranes that need little water.

Venkatnathan and his fellow researchers chose Nafion because it is the state-of-the-art membrane for building efficient fuel cells. Nafion has excellent proton conductivity and chemical and mechanical stability.

Even though many scientists have studied Nafion, they have yet to fully understand the membrane's structure and its behavior at a nanoscale, particularly at different levels of hydration.

"Experimental studies are inadequate to understand proton dynamics, because it occurs below nanoscale. This is where NERSC's computing power becomes indispensable. By using advanced computer models, we are getting a grasp of the complex processes at the molecular level in polymer membranes," said Devanathan.

Using classical molecular dynamics simulations, the research team investigated the impact of four levels of hydration and two different temperatures, 300 Kelvin and 350 Kelvin. The scientists calculated structural properties such as radial distribution functions, coordination numbers and dynamical properties such as diffusion coefficients of hydronium ions and water molecules.

The results of their calculations showed that protons and water molecules are bound to sulfonate groups in the membrane

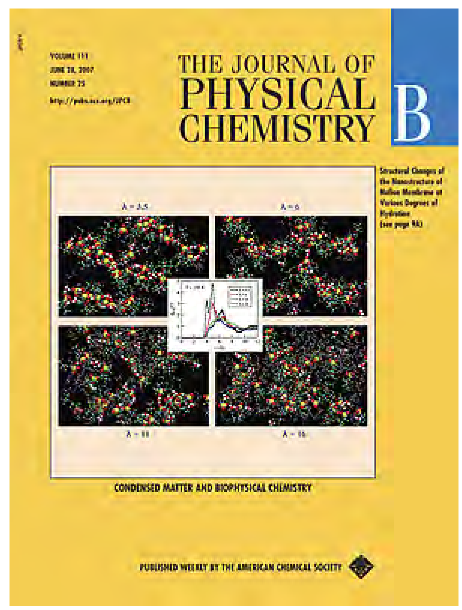

at low hydration levels. As the hydration level increases, the water molecules become free and form a network along which protons can hop. This leads to a dramatic increase in proton conductivity.

The findings have helped in interpreting experimental results that indicate a major structural change taking place in the membrane with increasing hydration. By combining these results with calculations from first principles, the PNNL team hopes to get insights that would lead to future PEMFC membranes that could operate with very little water.

NERSC Users Group Meeting

NERSC STAFF SHOWCASE NEW SERVICES AND TOOLS FOR AN ANNUAL GATHERING OF RESEARCHERS

Dozens of NERSC users convened at the supercomputer center this month for a four-day meeting and training, where they learned about new services, software and the Cray XT4 that is scheduled for deployment before the end of the year.

The annual NERSC Users' Group meeting provided an opportunity for researchers to meet with the NERSC staff and get a close-up look at NERSC resources that would support their scientific pursuits.

"It was good to meet other users and learn about services that will enable us to carry out our projects smoothly," said Stephane Ethier, vice-chair of the NERSC Users Group Executive Committee and a researcher from the Princeton Plasma Physics Laboratory. "The users also received critical training on the new NERSC-5 system, the 19,320-core Cray XT4 computer Franklin, which will allow them to make efficient use of this exciting new platform."

The event began with presentations and scientific talks by NERSC staff and users, followed by training on using Franklin, including compiling, debugging, profiling, and optimizing jobs for the best performance.

Bill Kramer, NERSC's General Manager, gave an overview of the center's operations and the large-scale projects being carried out at the center. He also discussed procurement plans for the next high-performance system. Dan Hitchcock, acting director of the Facilities Division in the Advanced Scientific Computing Research program office within the DOE Office of Science, outlined the federal agency's plans for funding and advancing research. He spoke via teleconferencing from Washington, D.C.

Jason Hick, head of the Mass Storage Group, gave an update of the High Performance Storage System (HPSS) while Brent Draney, head of the Networking, Security and Servers Group, discussed security risks and NERSC's efforts to minimize them.

Cecilia Aragon from the Analytics Team described the data management, visualization and analysis tools that her team provides to researchers. David Skinner, head of the Open Software and Programming group, talked about a proposed web portal for scientists to access the NERSC Global Filesystem, a new feature that would allow users to easily access data via the web.

Katy Antypas, a member of the User Services Group, introduced a new debugging tool called DDT (Distributed Debugging Tool). DDT, developed by Allinea Software, will be deployed on Franklin. The debugging tool allows scientists to search for and fix problems with their codes. Antypas compared DDT with the TotalView's software that has been deployed on other NERSC systems, and demonstrated how DDT would meet the researchers' needs in a cost-effective manner.

John Shalf, head of the Science-Driven System Architecture Team, gave presentations on characterizing the NERSC workload and on multi-core performance.

Jonathan Carter, head of the User Services Group, presented an update on Franklin's acceptance testing results. Helen He, also in the User Servces Group, presented an overview of Franklin, including its architecture, networking capabilities and software. She also was one of the NERSC staff who trained researchers on using Franklin.

Jim Craw, head of the Computational Systems Group, outlined the plan for decommissioning Seaborg. Shreyas Cholia, a member of the Open Software and Programming Group, explained the benefits of using the Open Science Grid.

Richard Gerber, a member of the User Services Group, provided instructions on how to submit and launch projects on Franklin. Zhengji Zhao, from the User Services Group, described NERSC's plans for installing software on Franklin.

Osni Marques, from the Computational Research Division at Berkeley Lab, presented the Advanced Computational Software Collection (ACTS), a software library that provides the tools for researchers to develop large-scale applications. Harvey Wasserman, on the Science-Driven System Architecture Team at NERSC, talked about benchmark performance. Cray consultants also held training sessions.

Three researchers, Julian Borrill from the Berkeley Lab, Adrianne Middleton from the National Center for Atmospheric Research and Ravi Samtaney from the Princeton Plasma Physics Laboratory, presented highlights of their research.

Check out http://www.nersc.gov/about/NUG/meeting_info/Sep07/ to learn more about the presentations.

Operational Review *continued from page 1*

the metrics for the review."

The committee evaluated each supercomputing center's business plan, financial management, innovation, scientific achievements and customer satisfaction. The inaugural review also serves as a baseline for future assessments.

Reviewers emphasized the importance for each center to seek feedback from its users. Each center conducts its own user survey; at NERSC the survey participation hovers around 10 percent of all authorized users. The survey yields a valuable assessment of the software, hardware and service offerings at NERSC.

In fact, researchers gave NERSC high marks in the 2006 survey, in which the respondents gave an average score of 6.3 — on a scale of 1 to 7 with 7 being "very satisfied" — for the question about their overall satisfaction with using NERSC resources.

The hard work by NERSC staff paid off when the committee gave the center kudos for its services and management.

"I would like to send along my thanks and congratulations to you and your team for such a fine job that you did on the Operational Assessment," wrote Vince Dattoria, a review committee member from the Advanced Scientific Computing Research Program Office within the Office of Science, in an email to Kramer.

"The professionalism and organizational maturity was evident in the presentations and responses to the impromptu questions from the reviewers. This is especially noteworthy since this was a baseline review with no precedents to draw from for guidance. Well done," Dattoria added.

Disturbance in Plasma Force

Most people do not think about turbulence very often, except when they are flying and the captain turns on the "Fasten Seat Belts" sign. The kind of turbulence that may cause problems for airplane passengers involves swirls and eddies that are a great deal larger than the aircraft. But in fusion plasmas, much smaller-scale turbulence, called microturbulence, can cause serious problems — specifically, instabilities and heat loss that could stop the fusion reaction.

Investigating the effect of such microturbulence is the mission of two SciDAC-funded projects: Center for Gyrokinetic Particle Simulations of Turbulent Transport in Burning Plasmas and Multiscale Gyrokinetics Particle Simulation of Magnetized Plasmas, both headed by Wei-li Lee of the Princeton Plasma Physics Laboratory. Lee and his fellow researchers have been using the Gyrokinetic Toroidal Code (GTC) they have developed to simulate instabilities in tokamak plasmas using the particle-in-cell (PIC) method.

"Particle-in-cell simulation, which began in the late sixties, uses finite size particles on a grid to dramatically reduce the numerical noise associated with close encounters between the particles, while leaving intact their long range interactions outside the grid," Lee explained. "This approximation reduced the number of calculations for particle interactions and greatly reduced the computational time."

In fusion research, all of the conditions necessary to keep a plasma dense and hot long enough to undergo fusion are referred to as confinement. The retention of heat, called energy confinement, can be threatened by microturbulence, which can make particles drift across, rather than along with, the plasma flow.

At the core of a fusion reactor such as a tokamak, the temperatures and densities are higher than at the outside edges. As with weather, when there are two regions with different temperatures and densities, the area between is subject to turbulence.

In a tokamak, turbulence can allow

continued on page 6

SPOTLIGHT

NEW STAFF

Prabhat

After working as a graphics system analyst at the Center for Computation and Visualization (CCV) at Brown University for six years, Prabhat has joined Computing Sciences at Berkeley Lab. He is now a member of the Analytics Team at NERSC and the Visualization Group in Computational Research Division.

Prabhat's research interests include computer graphics, scientific visualization, high-performance rendering and computing. He is also interested in human computer interaction, information visualization, general-purpose computation on graphics processors and machine learning.

At CCV, Prabhat supported a variety of inter-disciplinary scientific projects ranging from virtual archaeology, high-dimensional mathematics, developmental biology, computational neurosciences to planetary geosciences. For example, he worked with geoscientists on developing a virtual terrain exploration tool for missions to Mars.

Prabhat, who goes by a single name, earned a master's degree in computer science from Brown in 2001. For his graduate study, he worked on the interactive visualization of four-dimensional objects using the CAVE, or Cave Automatic Virtual Environment. Originally developed in the Electronic Visualization Laboratory at the University of Illinois in Chicago, the CAVE consists of a room with several large screens for projecting images on the walls and the floor, creating a virtual reality.

Prabhat earned a bachelor's degree in computing science from the Indian Institute of Technology in New Delhi, India in 1999.

VISIT FROM JAPAN

Six members of Japan's premier science and technology research institution, RIKEN, visited NERSC for a series of talks and a tour last month.

Ryutaro Himeno, director of RIKEN's Next Generation Supercomputer R&D Center, spoke about RIKEN's HPC activities and the status of its next-generation supercomputer development project. The supercomputer will start partial operation in April 2011, and its performance will be increased to 10 petaflop/s by March 2012. The basic design of the system was completed last March. The supercomputer will focus on carrying out nanoscience and life science projects.

Himeno also discussed RIKEN's current HPC activities in both hardware and software, such as the MD-GRAPE3 special purpose computer.

Himeno received his Ph.D. in engineering from the University of Tokyo in 1988. In 1979, he joined Nissan Motor, where he was engaged in the research of applying computational fluid dynamics to developing aerodynamic cars. He joined RIKEN in 1998. He is also a visiting professor at the University of Tokyo.

Horst Simon (left), head of Computing Sciences and NERSC, gave the Japanese delegation from RIKEN a tour of the supercomputing facility.

Seaborg *continued from page 1*

but that event is eclipsed by the brighter gamma ray bursts.

Woosley and Weiqun Zhang, a Ph.D. candidate, also at UC Santa Cruz, had been running simulations of the jet formation of the model at NERSC when observations of gamma ray bursts by NASA's High-Energy Transient Explorer (HETE) in March 2003 confirmed the collapsar model as the only theory supported by the observational data. Woolsey's simulation work was done under the Scientific Discovery through Advanced Computing (SciDAC) program, run by the Office of Science.

Seaborg underwent a significant upgrade in 2003, increasing the number of its processors from the initial 2,528 to 6,656. The upgrade, costing $30 million, readied the supercomputer for large-scale projects from the SciDAC and INCITE programs.

To prepare for Seaborg's decommissioning, NERSC has started archiving inactive accounts. Jim Craw, head of the Computational Systems Group at NERSC, presented a timeline and discussed data retention issues during the annual NERSC Users' Group meeting this month.

Seaborg is handing off its workload to a new and more powerful system, a Cray XT4 named Franklin. Franklin, which boasts 10 times more computing power than any of the current NERSC supercomputers, is scheduled to begin production this fall.

After retirement, Seaborg could be re-used by other DOE facilities or salvaged.

COMPUTING ON SEABORG

	Total CPU Hours[2]	CPU Hours Used by Scientists	CPU Hour Used by INCITE Projects	No. of INCITE Projects	No. of Total Active Users[3]	No. of Active Science Users
2007[1]	32,652,500	32,537,200	4,283,200	7	1,048	980
2006	51,904,400	51,554,100	3,992,500	3	1,991	1,868
2005	50,180,000	49,809,100	7,281,800	3	1,903	1,749
2004	56,821,900	56,086,500	6,670,800	3	1,458	1,354
2003	39,913,300	38,088,300	2,502,400	3	1,224	1,115
2002	22,649,400	22,279,200	1,785,700	3	897	831
Total	256,696,400	252,787,500	26,516,300	22	3,165	2,996

[1]Data reflect usage through mid-September.

[2]Total CPU hours includes time used by researchers, NERSC staff and HPC vendors.

[3]Active users include researchers, NERSC staff and HPC vendors.

Open Science Grid *continued from page 2*

OSG infrastructure.

Although the PDSF system at NERSC has been part of the OSG for several years, it's primarily for projects in high energy physics and nuclear science. Making other NERSC supercomputers available over the OSG has been a priority over the past year. NERSC managers tapped Cholia to work on integrating OSG services and hired Jeff Porter to be part of the OSG team that validate and test software tools. Bill Kramer, NERSC's General Manager, chairs the OSG Council.

Connecting NERSC systems to the OSG requires bridging different pieces of software and building a new infrastructure, particularly when some of the NERSC computer architectures are new to the OSG. Cholia is responsible for examining the OSG software stack and figuring out ways to support it, either through using solutions already available or by developing new ones.

Porter, on the other hand, works on making sure the software released by OSG several times a year can be deployed without glitches. He runs validation tests and solicits feedback from selected users who also get to run the new tools over a testbed, or a small grid.

Adding Jacquard and Bassi brings parallel computing resources to the OSG and makes such resources available for scientists who need to run MPI codes and large parallel jobs. Cholia and Porter will put the new Cray XT4, named Franklin, on the OSG in the future.

Having NERSC be part of the OSG helps the OSG to fulfill one of its goals — attracting a greater variety of users from different disciplines. NERSC hosts research from a wide range of scientific communities, including life sciences, material sciences, climate research and chemistry.

NERSC also has set aside computing time specifically for projects carried out over the OSG, as part of an effort to attract new research and users.

Learn more about the Open Science Grid at http://www.opensciencegrid.org. Check out Cholia's presentation at http://www.nersc.gov/about/NUG/meeting_info/Sep07/.

Plasma Force *continued from page 4*

charged particles in the plasma to move toward the outer edges of the reactor rather than fusing with other particles in the core. If enough particles drift away, the plasma loses temperature and the fusion reaction cannot be sustained.

The growth of the microinstabilities that lead to turbulent transport has been extensively studied over the years. Understanding this process is an important practical problem, and it is also a true scientific grand challenge which is particularly well suited to be addressed by modern terascale computational resources.

With the availability of terascale computers in recent years, the GPSC team has been able to carry out simulations of experiment-sized plasmas with improved physics fidelity. Typical global PIC simulations of this type have used one billion particles with 125 million grid points over 7,000 time steps to produce significant physics results. Simulations of this size would not be feasible on smaller computers.

With nearly two orders of magnitude increase in particle numbers, the GPSC project has been able to resolve long-standing uncertainty about the effect of discrete particle noise on the long-term transport predictions of turbulent gyrokinetic PIC simulations.

The "noise" referred to here involves not just particle interactions that are not relevant to energy transport, but primarily numerical sampling noise, because PIC simulations involve Monte-Carlo sampling of a collection of "marker" particles. Recent work shows that this numerical noise has little effect on the resulting energy transport when a reasonable number of particles is used.

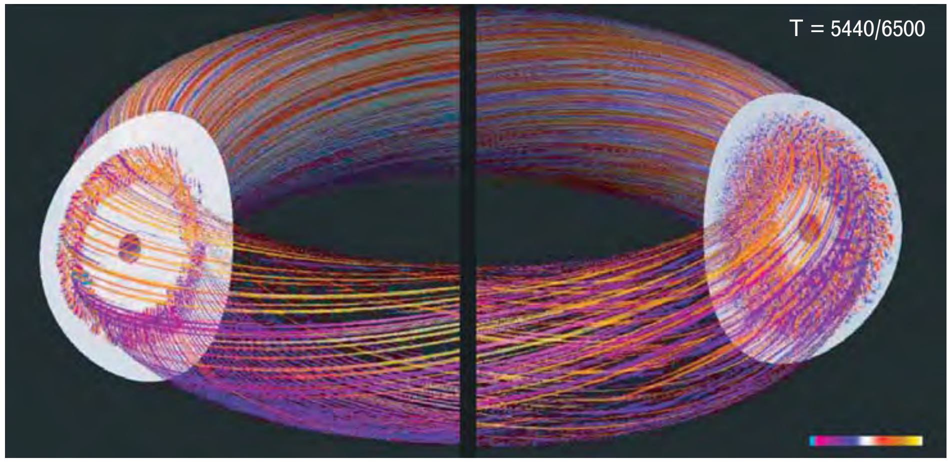

Turbulence spreading (left to right) as depicted by the perturbed potentials of ITG turbulence on the poloidal plane as they follow the magnetic field lines around the torus.

When the GTC code was applied to a geometry similar to the ITER experiment, an interesting new phenomenon was discovered: the turbulence spreads radially from a localized region to eventually cover most of the poloidal plane. This discovery was made possible by a simulation volume that is large enough to allow a clear scale separation between the turbulence eddy size and the device size.

"The simulation clearly shows that small-scale turbulence eddies are typically generated in the unstable region and flow along the streamers to the stable region," Lee said. "In addition, the streamers are found to break and reconnect, resulting in a very complex dynamical evolution. These new results have raised intense interest in the fusion theory community on the fundamental physics of turbulence spreading."

Clearly, there is a lot more work to be done in modeling tokamak plasmas, and with petascale computers coming online and the addition of more detailed physics to the GTC code, the GPSC team is eager to continue. With trillion particle simulations, they hope to find detailed solutions to problems such as electron thermal transport, the scaling of confinement with plasma size and the effects of different ionic isotope species such as tritium on plasma burning.

The success of these efforts will depend on close collaboration with other SciDAC centers, including the Terascale Optimal PDE Simulations (TOPS) Center, the Scientific Data Management (SDM) Center, the Ultrascale Visualization Center, and the Visualization and Analytics Center for Enabling Technologies (VACET).

Learn more about Lee's research at http://w3.pppl.gov/theory/GPSC.html.

WHAT IS NERSC NEWS?

NERSC News publishes every other month and highlights the cutting-edge research performed using the National Energy Research Scientific Computing Center, the flagship supercomputer facility for DOE's Office of Science. NERSC News editor Ucilia Wang can be reached at 510 945-2402 or Uwang@lbl.gov. Find previous NERSC News articles at http://www.nersc.gov/news/nerscnews.