Oak Ridge Leadership Computing Facility

SCIENCE at the PETASCALE 2009

Pioneering Applications Point the Way

ORNL 2009-G00700/jbs

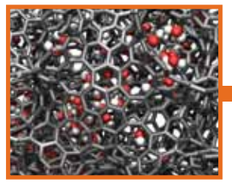

OLCF
Oak Ridge Leadership Computing Facility

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

CONTENTS

1 Introduction

DOMAINS

3 Biology
5 Chemistry
9 Climate
13 Combustion
15 Fusion
21 Geoscience
23 Materials Science
33 Nuclear Energy
37 Nuclear Physics
41 Turbulence

APPENDICES

43 Specifications
44 Applications Index
46 Acronyms

U.S. DEPARTMENT OF ENERGY | Office of Science

INTRODUCTION

Now is an especially good time to be engaged in computational science research. After years of planning, developing, and building, we have entered the petascale era with Jaguar, a Cray XT supercomputer located at the Oak Ridge Leadership Computing Facility (OLCF).

By petascale we mean a computer system able to churn out more than a thousand trillion calculations each second—or a petaflop. As impressive as this number is, the takeaway is that we are able to simulate physical systems—as large as galaxies and as small as the building blocks of subatomic particles—with previously unthinkable realism. This new fidelity brings with it new confidence as we model weather and climate, search for new or more efficient forms of energy, and seek to create new materials with custom-designed properties.

Opportunities for groundbreaking computational research will continue to expand as we move in the next decade toward the exascale—millions of trillions of calculations each second.

Jaguar kicked off the era of petascale scientific computing when it came on line in late 2008 with a peak performance of 1.64 petaflops. As the most powerful system ever created for open scientific research, Jaguar was the first supercomputer to host scientific applications that broke the petaflop barrier. (For more information on the OLCF and Jaguar, please visit our website at http://www.nccs.gov and see appendix.)

The OLCF recently completed a 6-month program of early scientific discovery on this new petascale system, inviting the world's most sophisticated scientific computing teams to take advantage of Jaguar's unprecedented power. The program allowed the OLCF to embrace its scientific user community, deliver early breakthrough science, and harden the system for production use when it is made available through avenues such as the Innovative and Novel Computational Impact on Theory and Experiment program (http://www.science.doe.gov/ascr/INCITE/).

Our early science efforts were a resounding success. As the program draws to a close, nearly 30 leading research teams in climate science, energy assurance, and fundamental science will have used more than 360 million processor-hours over these 6 months. Their applications were able to take advantage of the 300-terabyte, 150,000-processor Cray XT5 partition of the Jaguar system to conduct research at a scale and accuracy that would be literally impossible on any other system in existence.

The projects themselves were chosen from virtually every science domain; highlighted briefly here is research in weather and climate, nuclear energy, geosciences, combustion, bioenergy, fusion, and materials science. Almost two hundred scientists flocked to Jaguar during this period from dozens of institutions around the world.

- **Climate and Weather Modeling.** One team used century-long simulations of atmospheric and ocean circulations to illuminate their role in climate variability. Jaguar's petascale power enabled unprecedented resolution, while finer-scale coupled physics models enabled better depiction of tropical convection and climate. Another team modeled high-impact weather events for the Atmospheric Model Intercomparison Project, performing decadal simulations to compare with earlier, coarser models. This team's work will provide a newfound capability for predicting climate change on a regional scale. The increased resolution that Jaguar now offers is a game changer for climate models, demonstrating that the existing science model has a real propensity for true prediction.
- **Nuclear Energy.** Large simulations of nuclear reactor physics can help reduce the need for costly experiments and shed light on the intricacies of neutron transport within the hostile environment of operating reactors. A scalable neutron transport application allowed researchers to conduct the largest known simulation of a nuclear reactor, with 8 trillion degrees of freedom (or variables) in the reactor, 26 billion in the core, and 126 billion in the assembly. Before this project, simulations could typically accomodate only tens of millions of variables in each of these areas.

  The simulations allowed researchers to resolve what had been only approximations of the neutrons' individual energies and directions of movement, giving the team a more accurate picture of fission power distribution. This is critically important in designing next-generation reactors as well as improving existing operating reactors, the source

of the majority of the carbon-emission-free power in the United States. Another simulation modeled the neutron distribution (position, direction, energy) in an experimental zero-power reactor core with 120 billion degrees of freedom, helping researchers move from existing approximations toward a detailed resolution of the energy domain.

- **Geosciences.** Researchers were able for the first time to calculate the flux of contaminants into the Columbia River basin at the Hanford 300 waste disposal site in Washington State. The team's application was able to model the complex nonlinear chemistry and transient fluid flow caused by fluctuations in the Columbia River. In doing so this research greatly enhances the Department of Energy's (DOE's) ability to address a critical challenge: the remediation of America's aging underground contaminant storage facilities.
- **Combustion Science.** Jaguar's size and scalability allowed researchers to look at lifted flames from a direct-injection diesel-engine jet and, for the first time, perform a three-dimensional (3-D) simulation capable of fully resolving flame and ignition features such as chemical composition, temperature profile, and turbulence flow. In another project, researchers performed the largest simulation to date of turbulent dispersion, enabling a greater understanding of the geometry of turbulence and multiparticle statistics and helping to bring more efficient engines to market.
- **Biology.** A highly scalable molecular dynamics application elucidated the atomic-level structure, dynamics, and degradation pathways of lignocellulosic biomass, a waste plant material with rich potential as a biofuel. By providing insight into the interaction between lignin molecules and cellulose fiber, the simulations may help eliminate a bottleneck in the production of bioethanol. This work seeks to deepen our understanding of how enzymes perform their function in nature and, as a result, lead to new enzymes that work faster and more effectively in the production of biofuels.
- **Fusion Research.** Physicists suspect heat losses in the ITER fusion reactor will be dominated by electron transport. Electron physics in turbulence plays a critical role in ion heat transport and in generating plasma rotation, both of which are critical to achieving fusion power production. Simulations on the petascale Jaguar system confirmed electron-scale fluctuations in the flowing plasma, providing valuable insight that could one day make fusion power production a reality.
- **Materials.** Jaguar allowed researchers to study inhomogeneities and disorder in cuprate superconductors and the possible role played by this disorder in the materials' superconducting behavior. Other teams used first-principles approaches to study the dynamics of the solvent structure and electron transport properties of water. And still another group performed electron-scale calculations of zinc oxide nanosystems to see how the atomic and electronic structure of the surfaces reflects interior ground states. This research promises to play a critical role in nanostructure solar cell research.

These achievements would be unthinkable without parallel advances in applied mathematics and computer science. Our evolving understanding of physical systems in climate science, combustion science, and fusion energy is governed by partial differential equations whose complexity mandates a computational approach. Computational methods, in turn, are based on algorithms that solve more fundamental mathematical problems such as large systems of linear and nonlinear partial differential equations. Scientists ultimately rely on software that combines state-of-the-art applied mathematics and computer science to wrest new discovery from today's largest, fastest supercomputers such as the Jaguar leadership system at the OLCF.

No project of this magnitude would be possible without dedicated institutional support. In particular, we would like to thank DOE's Office of Science and its Advanced Scientific Computing Research (ASCR) program, as well as management at Oak Ridge National Laboratory (ORNL) and staff at the OLCF.

The projects featured in this document used pioneering applications to facilitate understanding in areas of research that are critically important to DOE's mission and to problems of national importance. With these applications and huge computational allocations on Jaguar, these researchers were able to punch through long-standing barriers and create new and useful knowledge. We are confident that their achievements will stand over time as substantial contributions to science and technology.

Doug Kothe

Director of Science,
Oak Ridge Leadership Computing Facility

Cellulosic Ethanol: A Simulation Model of Lignocellulosic Biomass Deconstruction

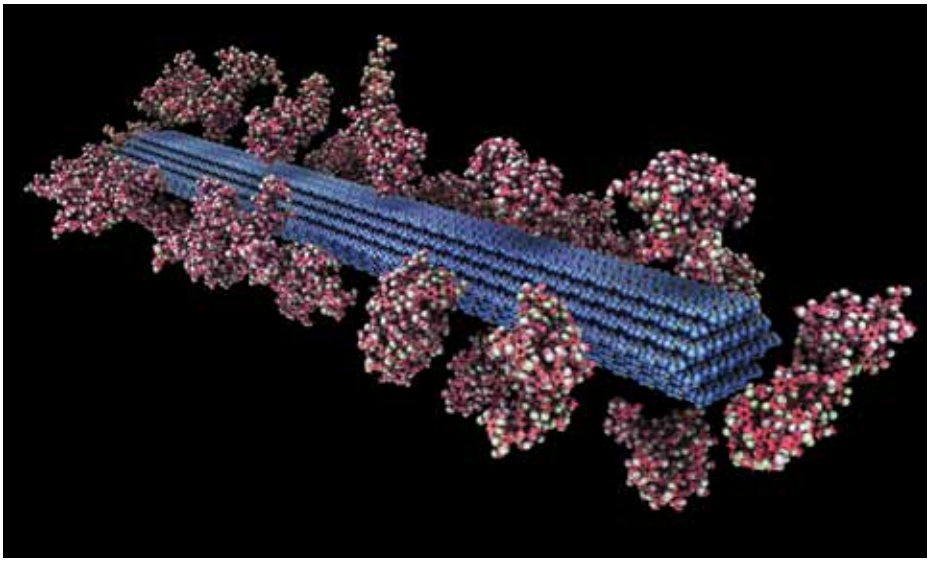

Abstract

Jeremy C. Smith and his research team are using more than 30 million hours to model plant cell wall lingocellulosic biomass and its interaction with enzymes that help turn this biomass into cellulosic ethanol.

Toward a fundamental understanding of lignin cellulose interaction: 3.3M atom lignocellulose system, cellulose fibril (blue), lignin molecules (pink), water (not shown). Visualization by Jamison Daniel, ORNL.

Key collaborators and affiliations

Jeremy C. Smith (Governor's Chair, University of Tennessee (UT) and Director, UT/ORNL Center for Molecular Biophysics); Loukas Petridis, Roland Shulz, Benjamin Lindner (ORNL).

Role of the Jaguar leadership system

Jeremy Smith's research team is consuming more than 30 million processor-hours on Jaguar to develop a molecular-level understanding of the structure, dynamics, and degradation pathways of cellulosic and lignincellulosic materials and their potential as sources for bioethanol. One of the most viable forms of renewable energy, bioethanol can be derived from the glucose in cellulosic and lignincellulosic materials. Hydrolysis, a chemical reaction in which water breaks down chemical bonds, must take place before cellulose can become glucose. Unfortunately, these materials strongly resist being broken down into usable components, a resistance known as biomass recalcitrance. A slow, costly, and labor-intensive pretreatment process is currently used to accelerate hydrolysis. The researchers are using a molecular dynamics application to simulate the structure, dynamics, and degradation pathways of these materials to provide a picture of the biomass that will help experimentalists design new, less resistant cell walls and enzymes that break the cellulose down more efficiently.

Investigating the biophysical properties of a large and complex system such as lignocellulose is a challenging task. However, the large scale of the simulation involved, approximately 5 million atoms, makes this work ideal for Jaguar. The researchers are using a molecular dynamics application called GROMACS to model the motions and behaviors of atoms that cannot be observed in real time. GROMACS is providing simulation insight into the action of enzymes on lignocelluloses at the molecular level. This approach requires massive computational resources and has hitherto been restricted to systems of approximately 100,000 particles and to the roughly 100-nanosecond (100 billionths of a second) timescale. However, the computing power of Jaguar is allowing this team access to microsecond (millionth of a second) timescales and micrometer (millionth of a meter) length systems, enabling it to observe reactions in unprecedented detail.

Importance of the research

Cleaner, more readily abundant alternatives to fossil fuels are crucial to the sustainability and renewability of global resources. Biofuels are currently being spotlighted in this search, and finding less labor- and material-intensive production methods is essential to the practicality of their use. Lignocellulose is a promising source of biofuel, but its resistance to hydrolysis, the process by which the material is broken down into usable components, is not well understood. Current methods of pretreatment and production are costly and time consuming. By modeling the interaction of this biomass with certain enzymes, the team hopes to find better, more efficient ways of overcoming the resistance to hydrolysis and to further contribute to the understanding of lingnocellulosic properties. The results are of particular interest to the computational chemistry community, DOE's BioEnergy Science Center (http://genomicsgtl.energy.gov/centers/center_ORNL.shtml), and the biofuel community in general.

Computational science applications employed

The molecular dynamics simulations performed are using GROMACS and Jaguar's petascale resources to study hydrated cellulose fibrils consisting of 300,000 atoms, requiring 1.5 million processor-hours. The researchers are also simulating four cellulose lignin molecules of 3.4 million atoms each, requiring 28.5 million processor-hours. To calculate the electrostatic interactions between atoms, they employ the reaction field method, which allows them to explore long-range molecular interactions. Using this method, 3-million-atom to 5-million-atom biological systems scale well up to 30,000 processor cores, modeling approximately 30 nanoseconds of movement per day. Atomistic simulations of biological systems on this grand scale are beginning to approach results in timescales of microseconds, a feat never before achieved. The team is also improving the parallel efficiency of the GROMACS code as a result of this work. The team's latest benchmarks indicate that GROMACS can scale to the entire 150,000 cores available on Jaguar's XT5 partition, producing 20 nanoseconds per day for a 100-million-atom system.

Institutional collaborations

The work is strengthening the partnership between UT and ORNL. Furthermore, the simulation models were built using experimental data from DOE's BioEnergy Science Center.

Acknowledgments

Support is made available by DOE through the Office of Biological and Environmental Research Bioenergy Science Center; the Office of Biological and Environmental Research Lab; an Experimental Program to Stimulate Competitive Research Implementation Award, "Neutron Scattering Research Network for EPSCoR States"; and a DOE Laboratory-Directed Research and Development grant, "Methodological Development of Computer Simulation in Molecular Biophysics."

Relevant publications

Schulz, R., B. Lindner, L. Petridis, and J.C. Smith. 2009. "Scaling of multimillion-atom biological molecular dynamics simulation on a petascale supercomputer," *J. Chem. Theory Comput.* **5**(10), 2798–2808.

Research results to date and/or path forward

These simulations of lignocelluloses include cellulose and lignin polymers. Research planned for 2010 includes a focus on hemicellulose and pectins, which will provide further insight into the atomistic interaction of lignocelluloses. The obtained atomic-detailed data will also serve as the basis for coarse-grained models.

Contact

Jeremy Smith

University of Tennessee-Knoxville

smithjc@ornl.gov

Quantum Monte Carlo Calculations of the Energetics, Thermodynamics, and Structure of Water

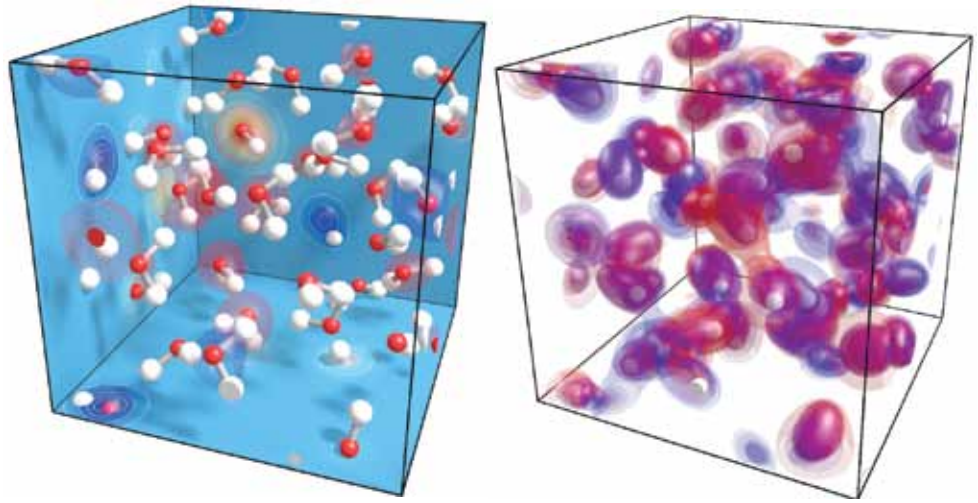

Abstract

The researchers are seeking a better description of the microscopic physics of water, including the bonding behavior of the hydrogen atoms. The team used more than 30 million hours to do advanced (quantum Monte Carlo) calculations and first-principles simulations of liquid water, its structure, and energetics—a benchmark for first-principles simulations of many-body systems.

A sample of liquid water, representative of what was used in the calculations. The red balls are the oxygen atoms, and the grey rods are the white hydrogen atoms with intramolecular bonds. The light shading and contours represent the electron density. Visualization by John Gergely,UIUC.

Key collaborators and affiliations

David Ceperley, Jeongnim Kim (National Center for Supercomputing Applications, University of Illinois–Urbana-Champaign [UIUC]); John Gergely (UIUC); Giulia Galli, Francois Gygi (University of California–Davis [UCD]); Eric Schwegler (Lawrence Livermore National Laboratory [LLNL]).

Role of the Jaguar leadership system

Researchers do not currently have a description of the microscopic physics of water, including the bonding behavior of the hydrogen atoms. The aim of this project is to calculate the interaction of water molecules in a small sample to improve the models used to describe the energetics, thermodynamics, and structure of water. The researchers prepared several hundred molecular configurations generated from a simulation based on an empirical model. They then used a much more computationally costly but accurate procedure, QMC, to assess the accuracy of the model for each of those configurations. QMC techniques are computer algorithms that simulate quantum systems, with the goal of accurately describing the electronic structure at the atomic scale. Once these simulations are complete, the team will go back and improve the models based on the results.

The water system addressed by this research, the largest such system calculated using QMC, is inherently disordered (the molecules do not lie on an ordered structure), requiring the team to consider many different configurations to ensure it had a representative sample. The team performed several hundred configurations, each requiring a substantial portion of Jaguar to reach the needed precision (1 part per million in total energy, meaning the errors on the energy are less than 10 degrees Celsius). With Jaguar the team for the first time had a large enough system and enough computer resources—more than 30 million hours—to run these configurations.

Importance of the research

Molecular biology researchers and the physics and quantum chemistry simulation communities need an accurate microscopic description of water's structure and dynamics. One outcome of this research is an increased understanding of how proteins surrounded by water in the human body are affected by the forces and movement of the water molecules. Also, understanding the nature of water in highly confined spaces, such as nanotubes,

could lead to the development of new technologies to supply clean water. To reliably simulate typical processes involving water, the total energy of a water configuration must be accurate to better than room temperature—or in electronic units, better than 1 million hartrees, or more than 1 million times the absolute value of the electric potential energy of a ground-state hydrogen atom. Achieving this level of accuracy will have breakthrough impacts on atomic-scale biological simulations and constitute a prominent benchmark for ab initio simulation of many-body systems (ten or more systems interacting at the atomic scale).

Computational science applications employed

QMCPACK, an application developed at the University of Illinois that can simulate correlated quantum materials on high-performance computers, is used to perform these simulations. The researchers were able to use the entire Jaguar system, although typical runs used a smaller fraction of the system (45,000 of the more than 150,000 cores of the Jaguar XT5 partition).

Institutional collaborations

This project was done within the framework of the QMC Endstation Project and the Scientific Discovery through Advanced Computing (SCIDAC) program. The Endstation Project brings together researchers from UIUC, Cornell University, The College of William and Mary, North Carolina State University, Florida State University, and ORNL. (This project was carried out only at UIUC.) The SciDAC program collaborators are from UIUC, UCD, and LLNL.

Acknowledgments

Support comes from the DOE Endstation Project (Ceperley principal investigator) and from a DOE SciDAC project administered at UCD (Galli principal investigator).

Relevant publications

Esler, K.P., J. Kim, D. Ceperley, W. Purwanto, E.J. Walter, H. Krakauer, S.W. Zhang, P.R.C. Kent, R. Hennig, C. Umrigar, M. Bajdich, J. Kolorenc, L. Mitas, and A. Srinivasan. 2008. "Quantum Monte Carlo algorithms for electronic structure at the petascale; the Endstation Project," *J. Phys.: Conf. Ser.* **125**(012057), 1–15.

Gergely, J. 2008. "Quantum Monte Carlo Methods for First Principles Simulation of Liquid Water." PhD thesis, Physics Department, University of Illinois Urbana–Champaign, August.

Morales, M.A., E. Schwegler, D. Ceperley, C. Pierleoni, S. Hamel, and K. Caspersen. 2009. "Phase separation in hydrogen–helium mixtures at Mbar pressures," *Proc. Natl. Acad. Sci. USA* **106**, 1324–1329.

Research results to date and/or path forward

The team is planning to simulate a host of microscopic systems at this level of accuracy. The goal is to go beyond pure water and do more complex studies. In the long term the researchers are striving to perform the electronic calculations for any type of system of ions and electrons.

Contact

David Ceperley
University of Illinois at Urbana-Champaign
ceperley@uiuc.edu

The Free Energy of Transfer of Hydronium from Bulk to Interface: A Comprehensive First-Principles Study

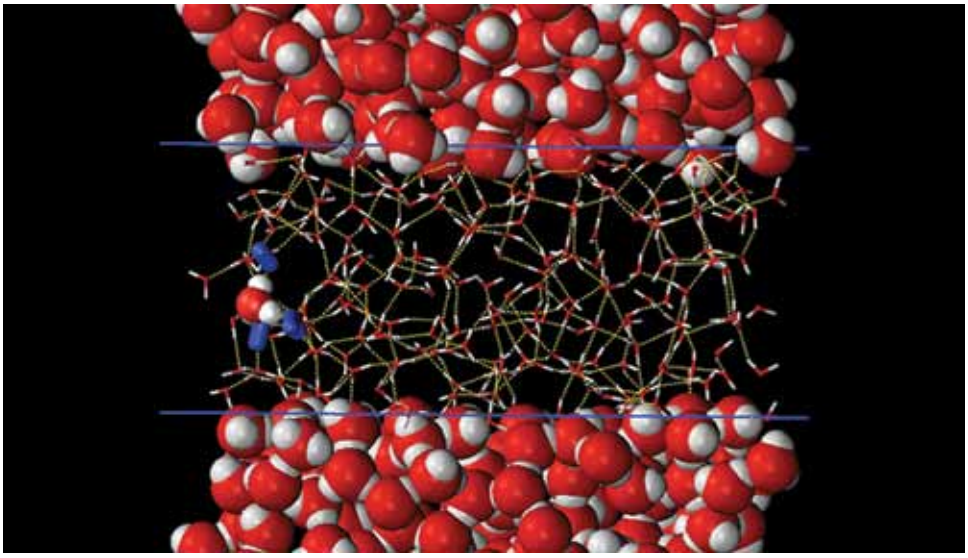

Abstract

The researchers are using 4 million hours to calculate the free energy transfer of hydronium from bulk water to the interface between water and air. They are using molecular dynamics simulation and novel sampling techniques that work to simultaneously umbrella sample a hydronium species in 215 water molecules.

H_3O^+ undergoing chemical transformations at the air-water interface. The H_3O^+ is depicted in the center using large red (oxygen) and three white (hydrogen) spheres. Blue bonds denote nearest hydrogen-bonded neighbors that are involved in interfacial chemistry. Image courtesy of Chris Mundy, PNNL.

Key collaborators and affiliations

Christopher Mundy, Roger Rousseau, Shawn Kathmann, Greg Schenter Pacific Northwest National Laboratory [PNNL]); Alessandro Curioni (IBM Research–Zurich, Switzerland).

Role of the Jaguar leadership system

The researchers have used more than 4 million computational hours on Jaguar to calculate transfer of the free energy of hydronium—an extremely reactive ion of hydrogen that is readily dissolved in water—from solution to interface (where the interface is a boundary of water with some other material such as air or another liquid). Ions are atoms or molecules that have acquired a positive or negative charge by gaining or losing one or more electrons. Water has the ability to autoionize, meaning molecules of water can spontaneously transition from an electrically neutral state to a lower-energy ionized state. This results in positively charged hydronium ions and negatively charged hydroxide ions. In pure bulk water, the concentration of these ions is low, but there is much disagreement among experiments about the presence (or lack of) hydronium and hydroxide ions at the interface. By studying both hydroxide and hydronium ion concentrations at the air–water interface, the researchers hope to provide an exact value of the pH of the surface of water and potentially resolve current experimental and theoretical discrepancies. They also hope to provide insight into the chemical reactions in condensed phases involving ionic/covalent bond making and breaking, excited states, hydrogen bonding, charge transfer, and radicals.

Direct observation of hydronium and hydroxide ions at water interfaces is experimentally difficult, as these species can intermittently change positions. The combination of free energy sampling methods and density functional theory interaction potentials in this work is the largest and most extensive set of electronic structure-based molecular dynamics simulations ever carried out on aqueous systems. Molecular dynamics simulations track the motion of atoms in a system over a given amount of time under the known laws of quantum physics. The researchers are simulating random sampling of the hydronium species in 215 water molecules, requiring the use of thousands of processors at once to complete the sampling in a reasonable amount of time.

Importance of the research

The abundance of water in nature, its role as a universal solvent, and its importance in many biological processes responsible for sustaining life are the driving forces behind the need for a fundamental understanding of its physical and chemical properties under different conditions and in various environments. The behavior of ions at the air–water interface has, until recently, been poorly understood.

Computational science applications employed

To study the air–water interface, the researchers used molecular dynamics simulations based on forces derived from density functional theory and novel sampling techniques, in particular the CP2K code, which performs atomistic and molecular simulations; CPMD, designed for molecular dynamics; and MCR, used to study liquid–vapor interfaces. The applications simultaneously umbrella sample a hydronium species in 215 water molecules.

Institutional collaborations

The researchers have actively used the Cray XT5 supercomputer at ORNL, the IBM Blue Gene/P supercomputer at Argonne National Laboratory (ANL), and the Cray XT4 supercomputer at the National Energy Research Scientific Computing Center (NERSC) to pursue studies of reactions in condensed-phase systems—matter that is in either its solid or liquid state—using unique calculations to determine interaction potentials.

Acknowledgments

Support for this project is provided by DOE's Office of Basic Energy Sciences, Chemical Sciences program.

Relevant publications

Kathmann, S.M., I-F.W. Kuo, and C.J. Mundy. 2008. "Electronic effects on the surface potential at the vapor–liquid interface of water," *J. Am. Chem. Soc*. **130**(49), 16556.

Mundy, C.J., R. Rousseau, A. Curioni, S.M. Kathmann, and G.K. Schenter. 2008. "A molecular approach to understanding complex systems: computational statistical mechanics using state-of-the-art algorithms on terascale computational platforms," *J. Phys.: Conf. Ser.* **125**, 012014.

Research results to date and/or path forward

The researchers plan to study ion pairing in aqueous solutions as well as reactions in heterogeneous environments.

Contact

Christopher Mundy

Pacific Northwest National Laboratory

chris.mundy@pnl.gov

Explorations of High-Impact Weather Events in an Ultra-High-Resolution Configuration of the NASA GEOS Cubed-Sphere Global Climate Model

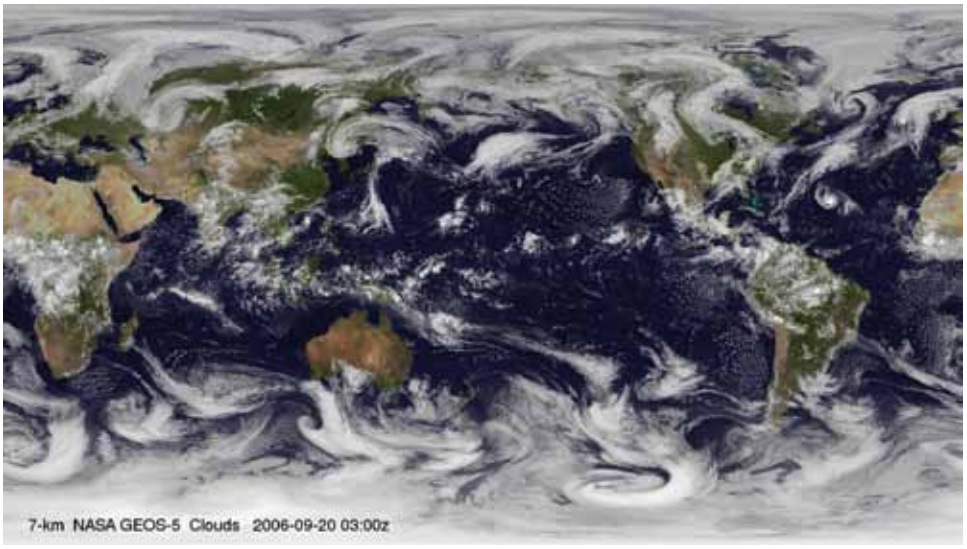

Abstract

Max J. Suarez and his research team have 20 million hours to determine the effects of individual weather events on global climate and vice versa. The team is creating atmospheric simulations of unprecedented resolution to compare with coarser resolutions, using the Goddard Earth Observing System Model (GEOS)-5 integrated models.

A global snapshot of forecasted cloud cover from the GEOS-5 model at 7 km global resolution, highlighting the existence of numerous cloud types represented at these resolutions. These include hurricane Helene in the tropical Atlantic Ocean, convective clusters throughout the tropical Pacific and Indian oceans, and the marine stratocumulus layer off the west coasts of North and South America, along with the characteristic structures of mid-latitude cyclones forming along fronts in both the Northern and Southern hemispheres. Image courtesy Max J. Suarez, NASA.

Key collaborators and affiliations

Max J. Suarez, William Putman (National Aeronautics and Space Administration [NASA]/ Goddard Space Flight Center); S.J. Lin (National Oceanic and Atmospheric Administration [NOAA]/Geophysical Fluid Dynamics Laboratory [GFDL]).

Role of the Jaguar leadership system

This team is examining the effects of individual weather events on global climate and the effects of global climate on individual weather events. The team has undertaken seasonal simulations using an unprecedented resolution of approximately 14 kilometers, or 1/8 degree (typical climate models have resolutions in the range of 55 to 110 kilometers, 1/2 to 1 degree). The researchers are also conducting a few very-high-resolution (3.5-and 7-kilometer) forecasts as contributions to the Year of Tropical Convection, a year-long program of coordinated observation, modeling, and forecasting that focuses on tropical convection prediction. They are conducting some 40 cases of 10-day forecasts at 7-kilometer resolution and a couple of 5-day forecasts at 3.5-kilometer resolution. Only a machine with Jaguar's computational power can efficiently compute the integrated physics models that the researchers are using to create these simulations, the highest-resolution atmospheric simulations to date.

Importance of the research

The research gives climate science its first look at the impact of weather events on climate variability and vice versa. The high-resolution simulations provide a methodology to scale climate-change predictions and projections downward to the regional level, using global models in preparation for contributions to the Intergovernmental Panel on Climate Change (IPCC) fifth assessment report. As such, these simulations have the potential to transform climate science in the area of climate change projection at regional scales.

Computational science applications employed

The researchers run their forecast applications using analyses from the GEOS-5 Data Analysis System (http://gmao.gsfc.nasa.gov/systems/geos5/). GEOS-5 is a system of models integrated using the Earth System Modeling Framework, a collection of flexible software intended to improve Earth-science applications. The researchers also use mathematical models developed by GFDL to run their simulations. The core component of

the atmospheric model, the dynamical core, has been scaled out to 50,000 cores on Jaguar. Extensive effort has been dedicated to improving the memory efficiency of the entire GEOS-5 system. At the core of this effort is the efficient use of parallel I/O to avoid costly and memory-consuming global communication operations that limit scalability. At the 7-kilometer resolution, the GEOS-5 system produces a global, regional weather/climate simulation, at a computational throughput of 50 days per day, using 30,000 cores.

Institutional collaborations

The researchers ran Atmospheric Model Intercomparison Project simulations using the GEOS-5 integrated systems (developed by the NASA Goddard Global Modeling Assimilation Office) and mathematical models developed at GFDL. The Atmospheric Model Intercomparison Project experiments also represent a joint science deliverable of DOE, "NOAA", and Year of Tropical Convection, a World Climate Research Programme-approved project.

Acknowledgments

Support for this project is provided by NASA's Modeling Analysis and Predictions Program, Earth Science Division, Science Mission Directorate, NASA Headquarters.

Relevant publications

Lee, M.-I., S. D. Schubert, M. J. Suarez, J.-K. E. Schemm, H.-L. Pan, J. Han, and S.-H. Yoo 2008. "Role of convection triggers in the simulation of the diurnal cycle of precipitation over the United States Great Plains in a general circulation model." *J. Geophys. Res. Atmos.* **113**, D02111, doi:10.1029/2007JD008984.

Ott, L.E., J. Bacmeister, S. Pawson, K. Pickering, G. Stenchikov, M. Suarez, H. Huntrieser, M. Lowenstein, J. Lopez, and I. Xueref-Remy 2009. "Analysis of convective transport and parameter sensitivity in a single column version of the Goddard Earth Observation System, Version 5, General Circulation Model." *J.Atmos. Sci.* **66**, 627.

Research results to date and/or path forward

The researchers are conducting a few very-high-resolution forecasts as contributions to the Year of Tropical Convection. These simulations are at resolutions of 7.5 and 3 kilometers, the highest-resolution climate models to date.

Contact

Max J. Suarez

National Aeronautics and Space Administration,
Goddard Space Flight Center

max.j.suarez@nasa.gov

Petascale CHiMES – Coupled High-Resolution Modeling of the Earth System

Abstract

Venkatramani Balaji and his research team have 20 million hours to model natural and forced climate variability at high resolution. They are using a coupled ocean–atmosphere model called MOMp25 that is capable of simulating 100 years in a single month at resolutions in the 25- to 50-kilometer range.

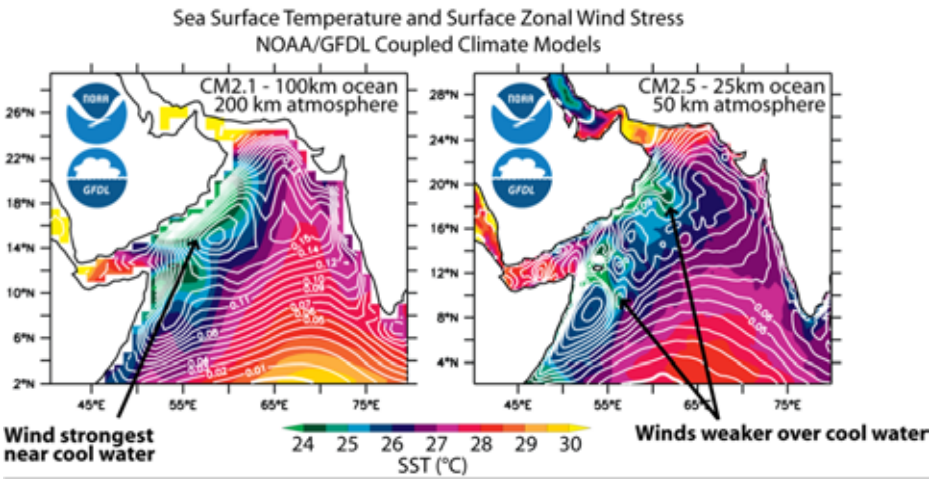

The figure shows Arabian Sea winds and sea surface temperature. Jaguar provides increased resolution for climate researchers, giving them the ability to capture a model atmosphere covariability on the mesoscale. Figure courtesy Whit Anderson, Rusty Benson, Tom Delworth, Tony Rosati and Gabe Vecchi, NOAA/GFDL.

Key collaborators and affiliations

Venkatramani Balaji, Tom Delworth, Tony Rosati, Isaac Held, S.-J. Lin (NOAA/GFDL); Christopher Kerr (University Corporation for Atmospheric Research).

Role of the Jaguar leadership system

This team is using 20 million computational hours to better understand and predict natural and forced variability of the coupled climate Earth System at high resolution. The simulations are being performed using Earth System models (ESMs) (http://www.gfdl.noaa.gov/earth-system-model), which are coupled mathematical codes that act as climate simulation tools. They couple an atmospheric circulation model with an oceanic one that includes representations of land and sea-ice dynamics. Current resolutions of IPCC–class models are mostly in the 100-kilometer range, but Balaji's research team seeks to create models in the 25- to 50-kilometer range, allowing regional climate prediction. Two scientific themes emphasized in this Coupled High-Resolution Modeling of the Earth System (CHiMES) project are decadal predictability of the Earth System and correlation of tropical cyclone frequencies and intensities with climate change.

Each of the models that make up the ESMs requires more than a thousand processing cores for computation, resulting in the need for thousands of cores for a given ESM simulation. Jaguar provides the massive computational power necessary to calculate coupled algorithms and accurately create Earth climate simulations at unprecedented resolutions.

Importance of the research

The cutting edge of climate research and policy questions is now moving to the issue of understanding and predicting climate variability and change on regional scales. Such regional scales are the ones of most direct relevance to society and decision makers. This research provides an early look at scientific issues associated with the ability to resolve mesoscale features in the atmospheric and oceanic circulations and their implications for understanding forced and natural variability of the climate system. Results from these simulations should provide insight into what to expect in the near future in terms of understanding regional climate change and therefore inform the design of international modeling campaigns aimed at addressing such questions.

Computational science applications employed

The study is based on runs of a coupled ocean–atmosphere model known as MOMp25. The atmosphere model is built on the cubed-sphere dynamical core (a modular component within the global modeling efforts at NASA, NOAA/GFDL, the National Center for Atmospheric Research, DOE, and other institutions) at 50-kilometer resolution, and the ocean model on the MOM4 (a computational tool used to study the global ocean climate system) model at a resolution of 25 kilometers. MOMp25 is capable of simulating 100 years per month using 2,400 processors on Jaguar.

Institutional collaborations

Components of the ESMs used by the team to create its simulations are developed at NOAA/GFDL. Preliminary results from the CHiMES project were obtained in 2008 in a series of simulations, including runs performed on Jaguar at the OLCF and the Cray XT4 at Lawrence Berkeley National Laboratory's (LBNL's) NERSC. The Cooperative Institute for Climate Science fosters research collaborations between Princeton University and the GFDL.

Acknowledgments

Support for this project comes from DOE's Office of Biological and Environmental Research and NOAA.

Relevant publications

Balaji, V., J. Anderson, I. Held, M. Winton, J. Durachta, S. Malyshev, and R.J. Stouffer. 2006. "The Exchange Grid: a mechanism for data exchange between Earth System components on independent grids," *Parallel Computational Fluid Dynamics: Theory and Applications, Proceedings of the 2005 International Conference on Parallel Computational Fluid Dynamics*, Elsevier Publishing Company.

Collins, N., G. Theurich, C. DeLuca, M. Suarez, A. Tranyanov, V. Balaji, P. Li, W. Yang, C. Hill, and A. da Silva. 2005. "Design and implementation of components in the Earth System Modeling Framework," *Int. J. High Perform. Comput. Appl.* **19**, 341–350.

Lin, S. –J. 2008. "Dynamical core advances in NOAA," presentation to the Climate Working Group's Climate Research and Modeling Program Review delivered March 25, 2008, http://www.joss.ucar.edu/cwg/mar08/presentations/3b_lunch_lin.pdf.

Research results to date and/or path forward

The researchers are planning to use 40 million computational hours on OLCF and GFDL resources to make contributions to the IPCC's Fifth Assessment Report (http://www.ipcc.ch/activities/activities.htm#1). The team anticipates using these allocations to perform experiments along two lines: time-slice experiments that focus on regional U.S. climate change and coupled model experiments that expose the role of the ocean in climate variability and change.

Contact

Venkatramani Balaji
Princeton University
balaji@princeton.edu

Direct Numerical Simulation of Lifted Diesel Jet Flame Stabilization

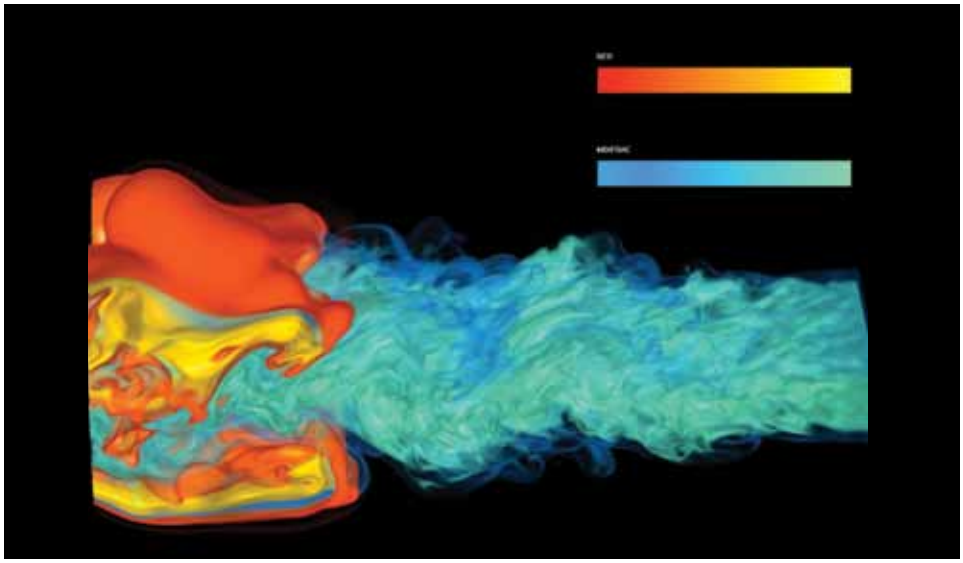

Abstract

The researchers received 50 million hours to perform 3-D direct numerical simulation (DNS) studies of lifted flames to determine how they are stabilized in diesel-engine thermochemical environments. The goal is to optimize the design and operation of evolving fuels in advanced engines.

Three-dimensional volume rendering of the mixture fraction and H_2O mass fraction showing a lifted n-heptane diesel jet flame. The data were generated through a 3-D coarse mesh DNS of a n-heptane jet flame at 50 atm, fuel-air equivalence ratio of 3.0, coflow air stream of 1,300K, fuel jet temperature of 9,00K, and Reynolds number of 4,000. Image courtesy Jacqueline H. Chen, SNL.

Key collaborators and affiliations

Jacqueline H. Chen, Edward Richardson, Ray Grout, Chun Sang Yoo (Sandia National Laboratories [SNL]); Ramanan Sankaran (ORNL); Tianfeng Lu (University of Connecticut).

Role of the Jaguar leadership system

The researchers performed DNS of the stabilization of lifted diesel jet flames to understand how a lifted turbulent jet flame is stabilized in an autoignitive, hot air coflow in diesel-engine thermochemical environments. They investigated the competitive roles of multistage autoignition, how flame is propagated in a partially premixed mixture, and how large eddies are carried along in the flow and mixed in the stabilization of the lifted flame. They also examined multistage ignition processes representative of diesel-engine combustion, using a diesel surrogate fuel, n-heptane, at engine pressures between 40 and 50 atmospheres. At these pressures and at low temperatures (less than 900 Kelvin), n-heptane exhibits "cool flame" behavior. A cool flame results from the transition between the first-stage (low-temperature) ignition to intermediate-temperature ignition, as the reactions are weakly exothermic and result in the formation of intermediate species. Cool flame behavior may facilitate transition to high-temperature ignition and the establishment of a flame—or the faster propagation of a flame—through a reactive mixture. DNS provides unprecedented detail unavailable through measurement.

Jaguar's petascale capability was used to investigate the feasibility of performing 3-D DNS of lifted diesel jet flame stabilization at high pressure (i.e., 50 atmospheres) with n-heptane and dimethyl ether (a biofuel surrogate) chemistry. The researchers began with "half-resolved" simulations of lifted diesel jet flames with an air coflow temperature of 1,300 Kelvin. On a half-resolved 1.85-micron (millionth of a meter) grid, the radical layers and even the heat release rate were only marginally resolved. Major species and temperature were adequately resolved in the temperature simulation, but further testing indicated a 0.4-micron grid would be required at these conditions to fully resolve all species; this is prohibitive even on a petaflop machine like Jaguar.

Importance of the research

In many modern combustion systems, fuel is injected into an environment of hot gases, and a flame may be stabilized through the recirculation of hot air and combustion products. This

leads to a turbulent lifted flame, and the hot environment makes it possible for autoignition to be used to control the stabilization of the flame base and determine the liftoff length (i.e., where the flame base resides downstream of the injector). Shorter ignition delays translate into shorter liftoff lengths and a smaller degree of fuel–air mixing. In addition to autoignition, flame propagation and large-scale mixing may contribute to the stabilization of the lifted flame base. In a diesel engine, the liftoff length affects the soot emission processes downstream. Experiments have shown that soot levels decrease as more premixing of fuel and air occurs upstream of the liftoff length.

Computational science applications employed

Preproduction DNS runs using S3D, a massively parallel DNS solver, are first required with reduced n-heptane and dimethyl ether mechanisms (58 and 30 transported species, respectively). Two production runs are then required next (also using S3D), one at an airstream coflow temperature of 1,200 Kelvin and the second at 1,300 Kelvin, each sized for 20 million processor-hours on Jaguar. The two targeted cases are designed to pinpoint the differences between single and multistage ignition characteristics on flame stabilization. Researchers found that a coflow airstream temperature of 1,300 Kelvin exhibits single-stage ignition, whereas a temperature of 1,200 Kelvin exhibits two-stage ignition. An extremely small grid size of 0.4 micron is required for full resolution of all species at 50 atmospheres—0.8 microns for marginal resolution of minor species.

Institutional collaborations

This work is a collaboration between ORNL and SNL.

Acknowledgments

Funding for this project was provided by the DOE Office of Science Basic Energy Sciences Chemical Sciences Program and ASCR.

Relevant publications

Liu, S. , J. C. Hewson, J. H. Chen, and H. Pitsch. 2004. “Effects of strain rate on high-pressure nonpremixed n-heptane autoignition in counterflow,” *Combust. Flame* **137**, 320-339.

Lu, T.F., C.K. Law, C.S. Yoo, and J.H. Chen. 2009. “Dynamic stiffness removal for direct numerical simulations,” *Combust. Flame* **156**(8), 1542–1551.

Yoo, C.S., R. Sankaran, and J. H. Chen. 2009. “Three-dimensional direct numerical simulation of a turbulent lifted hydrogen jet flame in heated coflow: flame stabilization and structure,” *J. Fluid Mech*., in press.

Research results to date and/or path forward

This research is motivating future efforts to step back and first perform two-dimensional (2-D) parametric studies of lifted diesel jet flame stabilization varying the coflow temperature, jet velocity, fuel/air equivalence ratio, and Reynolds number of the jet. Differentiation between ignition properties of n-heptane and oxygenated fuels like di-methyl ether and isobutanol will also be examined to help guide the eventual simulation of a small number of 3-D DNS cases to be used for model validation and development.

Contact

Jacqueline H. Chen

Sandia National Laboratories

jhchen@sandia.gov

Steady-State Gyrokinetic Transport Code (TGYRO)

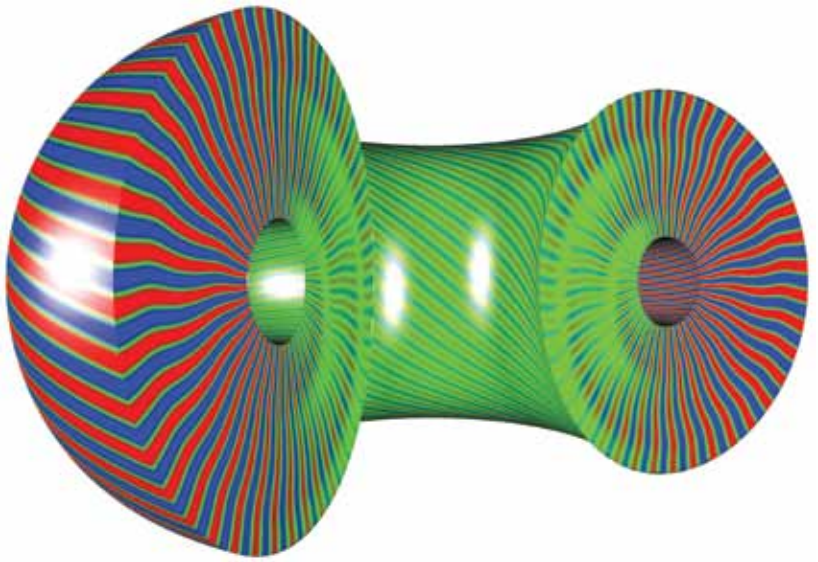

Abstract

For the first time, Jaguar is allowing researchers to make accurate plasma profile predictions based on the natural turbulences within the plasma. This predictive capability is critical in the eventual design and construction of fusion reactors.

Turbulent electron density fluctuations in a plasma with substantial cross-sectional elongation. Image courtesy Jeff Candy, GA.

Key collaborators and affiliations

Jeff Candy Ron Waltz, Emily Belli, Aaron Collier (General Atomics [GA]); Chris Holland (University of California–San Diego); Mark Fahey (ORNL)

Role of the Jaguar leadership system

This team is simulating the approach of turbulent plasma to a steady state by balancing power input (from external heating sources) with the power losses due to turbulence. This work requires the execution of hundreds to thousands of independent direct turbulence simulations to accurately converge on the solution. This approach to the problem is possible only as a result of significant algorithmic advances in the GA turbulence code, GYRO, combined with access to ever-expanding computational resources such as those offered by Jaguar. The final result is a set of predictions of the plasma profiles that arise given a specified input power.

The original TGYRO application results were obtained on Franklin, a Cray XT4 massively parallel processing system at NERSC (http://www.nersc.gov/). As TGYRO matured, the petascale Jaguar system became an ideal match for it because of the high availability of the system for jobs in the 3,000-to-6,000-core range. Although TGYRO has been shown to scale perfectly beyond 100,000 cores, scientific productivity is currently optimal at somewhat smaller sizes (about 5,000 cores).

Importance of the research

Tokamak confinement of magnetic plasmas, or containment of the necessary heat to sustain a fusion reaction, is limited in principle by the turbulent state of the plasma in the core and edge of the tokamak geometry. For roughly a decade now, so-called gyrokinetic codes have been able to simulate the evolution of this turbulence directly, but not in a predictive fashion. These codes evolve initial plasma profiles of a given temperature and density for only a very short time (typically less than a millisecond). In reality, however, the plasma profiles evolve to a self-organized state in which heat loss will balance the heat input from external heating sources (or from thermonuclear reactions, in the case of a reactor). Simulating this approach to the steady state, and determining the steady-state plasma profiles, has heretofore been done only with reduced models, which often rely upon gross approximations that neglect key features of the problem to maintain computational tractability. Now, for the first time, enabled by Jaguar and the TGYRO application, researchers can begin to make critical profile predictions.

Computational science applications employed

The researchers are using the GYRO application for the turbulent contribution to the plasma transport (heat and particle loss) and the NEO application for the neoclassical capability. Both of these applications represent the highest accuracy available for the solution of the fundamental kinetic equations. Moreover, GYRO is known to be a very fast and efficient gyrokinetic code. Processor scaling of TGYRO is also near optimal efficiency.

Institutional collaborations

ORNL (Fahey) and GA (Candy) have had a longstanding collaboration. However, most of the code work was done in-house at GA.

Acknowledgments

Funding for this research is provided by SciDAC through ASCR.

Relevant publications

Candy, J., C. Holland, R. Waltz, M. Fahey, and E. Belli. 2009. "Tokamak profile prediction using direct gyrokinetic and neoclassical simulation," *Phys. Plasmas* **16** (2009), 060704.

Candy, J. 2009. "Status of GYRO/NEO/TGYRO," lecture presented at Scientific Grand Challenges in Fusion Energy Sciences and the Role of Computing at the Extreme Scale, Washington, D.C., March 18–20, 2009.

Candy, J., H. Nordman, M. Fahey, T. Fülöp, and E. Belli. 2008. "Transport in ITER-like plasmas in neoclassical, fluid and gyrokinetic descriptions," poster presented at 50th annual APS DPP meeting, Dallas, Texas, November 17–21, 2008.

Candy, J., M.R. Fahey, and R.E. Waltz. 2007. "Progress on a fully gyrokinetic transport code," paper and poster presented at 34th European Physical Society Conference on Plasma Physics, Warsaw, Poland, July 2–6, 2007.

Research results to date and/or path forward

TGYRO is poised to become the transport-modeling tool of choice for many fusion science groups. It already has various users worldwide.

Contact

Jeff Candy

General Atomics

candy@fusion.gat.com

Global Gyrokinetic Turbulence Simulations of Toroidal Momentum and Electron Thermal Transport and Comparison with the National Spherical Torus Experiment

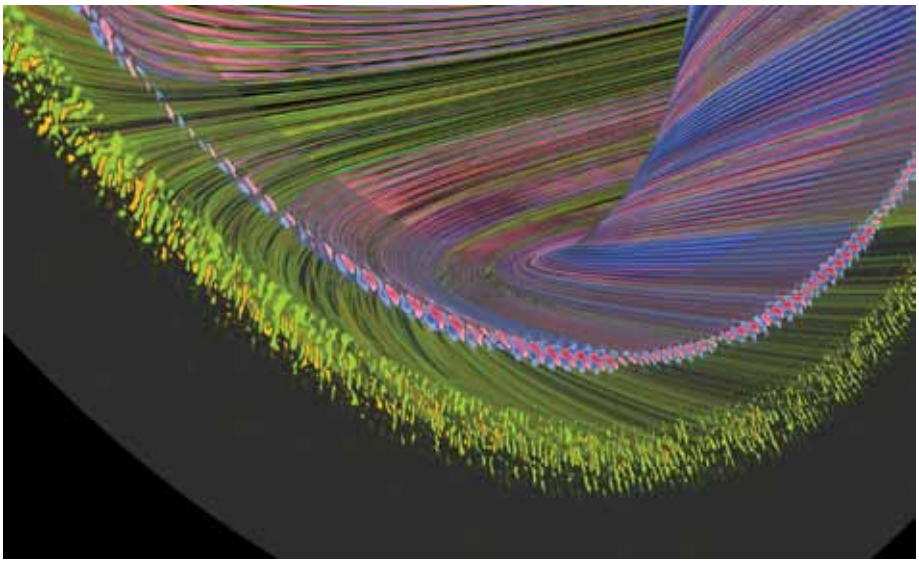

Abstract

Understanding the plasma energy transport and rotation behavior in toroidal, or doughnut-shaped, experimental fusion devices is critical in the design and implementation of experiments in the upcoming ITER fusion reactor.

Three-dimensional image showing electron-scale electrostatic potential fluctuations generated by ETG instabilities in a global gyrokinetic PIC simulation of the NSTX. Visualization by Kwan-Liu. Ma, UCD.

Key collaborators and affiliations

Weixing Wang, Stephane Ethier [PPPL]; Mark Adams (Columbia University); Francesca Poli (University of Warwick, United Kingdom); Kwan-Liu Ma (UCD); Scott Klasky (ORNL); visualization team at ORNL: Sean Ahern, Dave Pugmire, Jeremy Meredith.

Role of the Jaguar leadership system

This project addresses two important questions: (a) Do electron-scale fluctuations driven by electron temperature gradient (ETG) instabilities exist in plasmas of PPPL's National Spherical Torus Experiment (NSTX) fusion reactor, and if so, are they relevant to the high level of electron energy transport universally observed in fusion experiments? (b) How do plasmas rotate in toroidal systems? This is a critical issue for both the large-scale (macroscopic) stability and the confinement of hot plasma in the multinational ITER fusion reactor.

To answer the latter question, researchers must discover the underlying mechanisms for turbulence-driven momentum transport in toroidal systems. Moreover, this research tries to validate global gyrokinetic particle-in-cell (PIC) simulations against existing fusion experiments, a critical step to achieving predictive capability for the assessment of confinement performance in ITER.

Using the Gyrokinetic Tokamak Simulation (GTS) code developed at PPPL and the newly available petascale resource on Jaguar, this team carries out large-scale PIC simulations of modern laboratory fusion experiments with billions of "particles" (representing microscopic blobs of plasma) on tens of thousands of processors. These global turbulence simulations require petascale computing because the global ETG simulation for realistic experimental parameters must resolve electron gyroradius-scale fluctuations for a wide spatial range of the NSTX device. The number of grid points required for the current ETG simulations, for example, is approximately half a billion. These simulations also require a large number of particles—more than 20 billion—to ensure good sampling of the phase space and keep numerical fluctuations low. Furthermore, a global multiscale resolution is needed to simulate trapped-electron-driven turbulence of laboratory fusion plasmas, which involves widely disparate timescales between electrons and ions.

Importance of the research

The rapid growth of supercomputing capability in the past decade has enabled aggressive simulation efforts that have greatly advanced current understanding of energy transport in the ion channel in magnetic fusion experiments. The challenge remains to understand how energy is lost through electrons—perhaps fusion energy's greatest mystery and one that has puzzled scientists for decades. Anomalously large electron thermal losses are observed in all toroidal magnetic confinement devices. The majority of energy losses in ITER are also expected to be carried by electrons. Another poorly understood problem is the plasma rotation generation in fusion experiments, which exhibits a complex phenomenology. Understanding the toroidal momentum transport for driving plasma rotation is an urgent issue for ITER because a healthy rotation is essential for achieving high-performance plasma confinement.

Computational science applications employed

The GTS scientific application used in this project is a highly scalable PIC code that includes several levels of parallelism to efficiently use the resources on Jaguar. The current ETG simulations are carried out on 65,536 Jaguar cores and consume about 75 percent of the memory available on the nodes. A parallel solver implemented with the Portable, Extensible Toolkit for Scientific Computation (PETSc) library (http://www.mcs.anl.gov/petsc/petsc-as/) is used to evaluate the fields generated in the simulations.

Institutional collaborations

Researchers at Columbia University, the University of Warwick, the Scientific Computing Group at ORNL, the SciDAC Institute at UCD, and the OLCF visualization team assisted researchers at PPPL with the project.

Acknowledgments

Support for this work comes from DOE's Office of Science magnetic fusion energy program and the SciDAC project for gyrokinetic particle simulation of turbulent transport in burning plasmas.

Relevant publications

Mazzucato, E., R.E. Bell, S. Ethier, J.C. Hosea, S.M. Kaye, B.P. LeBlanc, W.W. Lee, P.M. Ryan, D.R. Smith, W. Wang, J.R. Wilson, and H. Yuh. 2009. "Study of turbulent fluctuations driven by the electron temperature gradient in the National Spherical Torus Experiment," *Nucl. Fusion* **49**, 055001.

Poli, F.M., W. Wang, S. Ethier, T.S. Hahm, and E. Mazzucato. 2009. "Synthetic diagnostics for validation of gyrokinetic nonlinear simulations of electron temperature gradient driven turbulence against experiments on NSTX," Transport Taskforce Workshop TTF 2009, San Diego, California, April 28–May 1, 2009.

Wang, W., T.S. Hahm, S. Ethier, P.H. Diamond, S.M. Kaye, and W. Solomon, "Gyrokinetic turbulence-driven non-diffusive toroidal momentum transport and effect of residual fluctuations in strong mean ExB shear," Transport Taskforce Workshop TTF 2009, San Diego, California, April 28–May 1, 2009.

Research results to date and/or path forward

The number of measurement channels will be increased to cover a lower wave-number range, where, according to the team's simulation results, fluctuations are peaked. Further, a high-wave -number diagnostic in the poloidal direction will be developed, which will enable experimentalists to obtain the 2-D spectrum of electron-scale density fluctuations.

Contact

Weixing Wang

Princeton Plasma Physics Laboratory

wwang@pppl.gov

Gyrokinetic Particle Simulation of Transport Barrier Dynamics in Fusion Plasmas

Abstract

An incomplete understanding of the physics of turbulent transport—or heat and particle loss—in a fusion plasma is one of the greatest obstacles facing the deployment of economically feasible fusion energy. Using more than 30 million hours, researchers are coming to understand this phenomenon in unprecedented detail.

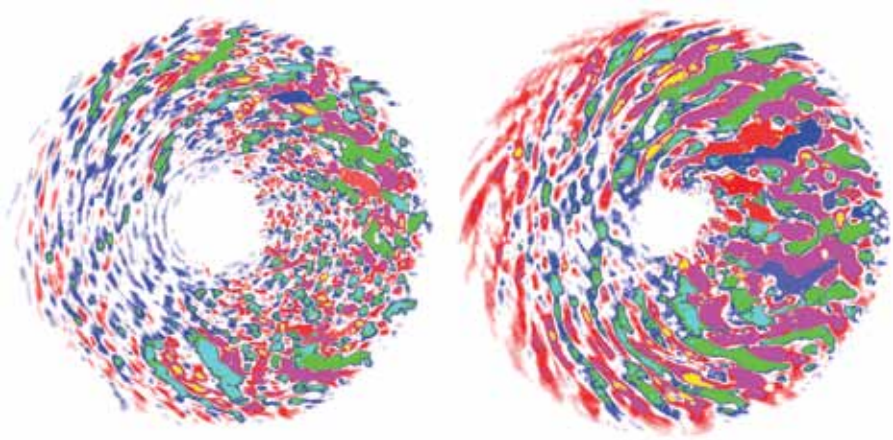

Electron transport in fusion plasmas is important in a burning plasma experiment like ITER. The simulation clarifies the self-regulatory dynamics of electron turbulence by the spontaneous generation of sheared flows. Image, right, shows the natural tendency of turbulence to form elongated eddies in the radial direction which induce large transport. Image, left, shows how self-generated sheared flows break up these eddies, leading to lower transport. Visualization in Phys. Rev. Lett.

Key collaborators and affiliations

Zhihong Lin (University of California–Irvine [UCI]); Scott Klasky (ORNL); Jay Lofstead (Georgia Tech); Chandrika Kamath (LLNL); Nathan Wichmann (Cray Inc.); Viktor Decyk (Unversity of California–Los Angeles [UCLA]).

Role of the Jaguar leadership system

Fusion ignition (attainment of a fusion reaction) in hot plasmas is sustained by a balance between the heat lost from turbulent transport and self-heating by fusion reaction products. As a result, understanding the physics of the turbulent transport is one of the most important scientific challenges facing magnetically confined, burning-plasma experiments. Current fusion reactor design relies on extrapolation of the transport level in present-day fusion experiments to much larger future devices, such as that being developed for the international ITER reactor. This team is studying electron heat transport in the largest-ever fusion simulations using the gyrokinetic toroidal code (GTC).

Numerical simulations of electron turbulence and its associated heat transport are more challenging than those of ion turbulence because electrons have much smaller mass (and thus large thermal velocity), which introduces smaller spatial scales and faster timescales. Energetic particles in ITER experiments are subjected to turbulence by thermal and energetic particles. Concurrent simulation of both scales is needed because of wave and particle dynamics. Particle dynamics of all species (thermal ions and electrons and energetic particles) with a large range of velocity must be treated simultaneously. Therefore, realistic simulations must cover multiple spatial-temporal scales.

Importance of the research

A fundamental understanding of electron transport in fusion plasmas is more elusive than that of ion transport. Nevertheless, electron transport is more important in a burning plasma experiment such as ITER because electrons tend to be heated first by the energetic fusion products. Furthermore, the confinement of these energetic particles is a critical issue in achieving ITER ignition because the plasma is heated by these energetic particles. Given that no existing fusion experiment can consistently reach significant ignition conditions, energetic particle confinement in ITER can currently be studied only by large-scale simulations. These simulations are providing the necessary physics insights for optimizing ITER experiments.

Computational science applications employed

The project uses GTC and XGC1, which calculate the interaction of billions of charged particles (e.g., electrons and protons) in the electromagnetic turbulence to predict the transport of mass, momentum, and heat in magnetic fusion plasmas. Turbulent transport is arguably the most important scientific challenge in burning plasma experiments for ITER, the crucial next step in the quest for fusion energy.

Institutional collaborations

These breakthrough simulations were carried out in close collaboration with computational scientists in the areas of I/O, parallelization and optimization, workflow, and data analysis. For the GTC simulation, collaborators came from UCI, ORNL, Georgia Tech, LLNL, Cray, UCLA, the University of Southern California, UCD, and The Ohio State University. For the XGC1, collaborators came from New York University, and for questions of theory from the University of California–San Diego.

Acknowledgments

This project acknowledges the following DOE Office of Science SciDAC fusion projects: Gyrokinetic Particle Simulation of Turbulent Transport in Burning Plasmas, Gyrokinetic Simulation of Energetic Particle Turbulence and Transport, and the Center for Plasma Edge Simulation.

Relevant publications

Xiao, Y., and Z. Lin. 2009. "Turbulent transport of trapped electron modes in collisionless plasmas," *Phys. Rev. Lett.* **103**, 085004.

Lin, Z., Y. Xiao, I. Holod, W. Zhang, W. Deng, S. Klasky, J. Lofstead, C. Kamath, and N. Wichmann. 2009. "Advanced simulation of electron heat transport in fusion plasmas," *J. Phys.: Conf. Ser*. **180**, 012059.

Chang, C.S., S. Ku, P.H. Diamond, Z. Lin, S. Parker, T.S. Hahm, and N. Samatova. 2009. "Compressed ITG turbulence in diverted tokamak edge," *Phys. Plasmas* **16**, 056108.

Research results to date and/or path forward

Future simulations will address multiscale dynamics intrinsic to the transport barrier formation and evolution, including realistic source and sink of particle, momentum, and energy to maintain the background plasma profiles and proper collision operators, which could have significant effects on the long-term behaviors of turbulence. The nonperturbative simulation method, which employs a full plasma distribution function, is also needed for the energetic particle simulation to handle the strong radial redistribution of energetic particles.

Contact

Zhihong Lin

University of California-Irvine

zhihongl@uci.edu

Modeling Reactive Flows in Porous Media

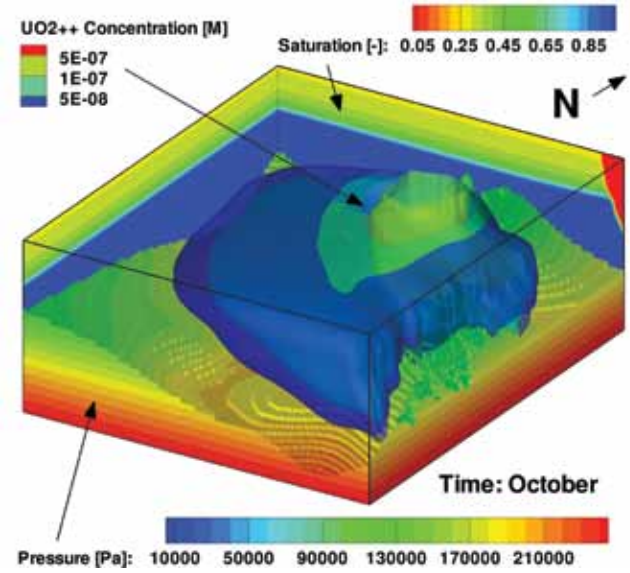

Abstract

Peter Lichtner and his team are using more than 9 million hours to model and calculate the flow rate of uranium pollution located underground in Hanford, Washington. This project will yield vital results for environmental research in subsurface contaminants.

Scaling of PFLOTRAN on the Cray XT5 and initial distribution of uranium for the Hanford 300 Area problem. Image courtesy Peter C. Lichtner, LANL.

Key collaborators and affiliations

Peter C. Lichtner (Los Alamos National Laboratory [LANL]); Glenn E. Hammond (PNNL); Richard T. Mills (ORNL).

Role of the Jaguar leadership system

The science challenge for this project is prediction of the natural attenuation rate, or loss of uranium, from an underground uranium plume located at the Hanford, Washington, 300 Area neighboring the Columbia River. The plume persists to this day, although estimates in the early 1990s based on a simple constant retardation model predicted it would dissipate in 10 to 15 years. Calculations are carried out on Jaguar using 4,096-core ensemble runs for a 28-million degree-of-freedom problem using multicomponent chemistry. Integration time steps are necessarily limited to 1 hour to capture the fluctuations in the Columbia River's stage. The team's work demonstrates that the release rate of uranium into the river is slow and can explain the long lifetime of the plume. Results obtained are in good agreement with field observations made at the site.

The large number of compute nodes available on Jaguar enables sets of ensemble simulations using 20,000 cores or more. Early access to the system gives the research team the opportunity to prepare for larger-scale and higher-fidelity runs by testing new at-scale developments in the solver algorithm.

Importance of the research

This research is important in helping DOE evaluate the appropriate remediation strategies for key sites that contain legacy waste generated over the last 50 years.

Computational science applications employed

Calculations are performed using the PFLOTRAN application, which employs key transport laws (e.g., Darcy's Law) to describe the flow of fluid through porous media. Lichtner and his team are able to demonstrate that PFLOTRAN can effectively use 131,072 cores for a 2-billion-degree-of-freedom, high-resolution problem. Parallel efficiency is excellent on up to 65,536 cores, and the team expects to increase efficiency at higher core counts by using new algorithms and faster inter-processor communication techniques to improve the overall speed and execution of the program.

Institutional collaborations

Collaborators came from LANL, PNNL, and ORNL.

Acknowledgments

This project is supported by the SciDAC-2 groundwater project.

Relevant publications

Mills, R.T., G.E. Hammond, P.C. Lichtner, V. Sripathi, G. Mahinthakumar, and B.F. Smith. 2009. "Modeling subsurface reactive flows using leadership-class computing," *J. Phys.: Conf. Ser.* **180**, 012062.

Mills, R.T., F.M. Hoffman, P.H. Worley, K.S. Perumalla, A. Mirin, G.E. Hammond, and B.F. Smith. 2009. "Coping at the user level with resource limitations in the Cray Message Passing Toolkit MPI at scale: how not to spend your summer vacation," *Proceedings of the 2009 Cray Users Group (CUG) Meeting, May 4–7, 2009, Atlanta, Georgia.*

Mills, R.T., V. Sripathi, G. Mahinthakumar, G.E. Hammond, P.C. Lichtner, and B.F. Smith. 2009. "Experiences and challenges scaling PFLOTRAN, a PETSc-based code for subsurface reactive flow simulations, towards the petascale on Cray XT systems," *Proceedings of the 2009 Cray Users Group (CUG) Meeting, May 4–7, 2009, Atlanta, Georgia.*

Research results to date and/or path forward

Further work in modeling the Hanford 300 Area will involve considerations of heterogeneous media, those comprising various types of materials, and at multiple scale lengths to further study the attenuation rate of uranium at the site.

Contact

Peter C. Lichtner

Los Alamos National Laboratory

lichtner@lanl.gov

Simulations of Disorder and Inhomogeneity Effects in High-Temperature Superconductors

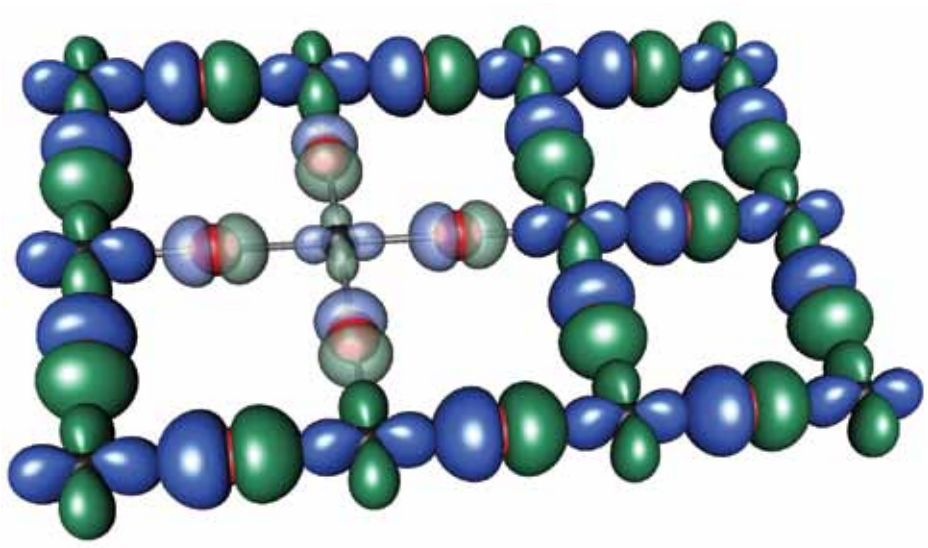

Abstract

Superconducting materials—which conduct electricity without resistance—promise enormous benefits, but they must be kept very cold (-190°F or colder). This project explores unevenness in the electronic structure of high-temperature superconductors, asking whether it causes superconducting behavior and whether optimal unevenness precipitates the highest superconducting temperature.

The basic electronic structure of the cuprates mapped onto the single-band Hubbard model, using the superconductor $La_{2x}Sr_xCuO_4$. Substituting Sr for La removes electrons from the CuO_2 planes, rendering the material superconducting, and distorts the electronic structure of the CuO_2 planes in its vicinity. A Hubbard model with random disorder in the orbital energies, where the impurity concentration is coupled to the hole doping, is used to simulate this system. Visualization by Jeremy Meridith, ORNL.

Key collaborators and affiliations

T. C. Schulthess (Swiss Federal Institute of Technology in Zurich and Swiss National Supercomputing Centre); T.A. Maier, G. Alvarez, M. Summers (ORNL).

Role of the Jaguar leadership system

Recent experiments have revealed nanoscale inhomogeneities in the electronic structure of a number of high-temperature superconducting cuprates (materials containing copper oxide). This observation raises several important questions: Do inhomogeneities cause high-temperature superconductivity, or do they merely play a peripheral role? Do they enhance or suppress the electron pairing mechanism that gives rise to superconductivity? And, perhaps most importantly, is there an optimum inhomogeneity that maximizes the critical temperature (i.e., allowing superconductivity at the highest possible temperature)? This effort seeks to answer these important questions by performing simulations for an inhomogeneous 2-D Hubbard model of the cuprates. This model is a simplified coarse-grained description of the electronic degrees of freedom of the cuprates, accounting for the dynamics of a single band of electrons in the copper oxide planes where superconductivity takes place.

The complexity and scale of the problem becomes clear when one considers, for example, a Hubbard model with random disorder and a minimal problem simulation to describe the random spatial modulations in the electronic structure. Quantum cluster simulations of the homogeneous Hubbard model without disorder use about 10 teraflops and run on the order of 1,000 processors. Disorder averaging adds another level of complexity, increasing the computational requirements by a factor of 100 to 10,000. This jump turns simulations of a disordered Hubbard model into a petascale to exascale computing problem.

Importance of the research

In the mid-1990s neutron scattering experiments led to the discovery of inhomogeneous electronic and magnetic states in cuprates, spurring many studies of the relationship between these inhomogeneities and superconductivity. This phenomenon appears to be crucial to understanding high-temperature superconductivity, and a comprehensive theory would be incomplete if it didn't account for the presence of these states.

In addition, the Yazdani Group at Princeton University (http://www.phy.princeton.edu/~yazdaniweb/Home.html) recently showed with variable-temperature scanning tunneling microscope studies that regions with paired charge carriers exist in the materials at temperature well above the superconducting transition. This opens the possibility that simulations focused on a detailed understanding of the role of inhomogeneities in cuprates may lead to new strategies in the search for materials with even higher transition temperatures.

Computational science applications employed

The project team is developing and evolving the DCA++ application, which employs dynamic cluster approximation models of high-temperature superconductors, such as the 2-D Hubbard model, and extensions that allow studies of disorder effects and nanoscale inhomogeneities. The DCA++ application is tuned to perform well at scale. As part of this early science project, a simulation of disorder effects in the Hubbard model was awarded the 2008 Gordon Bell Prize for first sustained petascale performance (1.3 petaflops) in a computation using 150,000 cores (the full Jaguar system).

Institutional collaborations

Collaborators in this project come from ORNL and the Swiss Federal Institute of Technology in Zurich and Swiss National Supercomputing Centre.

Acknowledgments

Support for this research is provided by two sources: ORNL Laboratory Directed Research and Development Funding, "Joining Ultrascale Computing and Neutron Scattering Studies to Enable Revolutionary New Materials Development," and funding from the Division for Scientific User Facilities of the Office of Basic Energy Sciences through the Nanotheory Institute of the Center for Nanophase Materials Sciences at ORNL.

Relevant publications

Alvarez, G., M.S. Summers, D.E. Maxwell, M. Eisenbach, J.S. Meredith, J.M. Larkin, J. Levesque, T.A. Maier, P.R.C. Kent, E.F. D'Azevedo, and T.C. Schulthess. 2008. "New algorithm to enable 400+ TFlop/s sustained performance in simulations of disorder effects in high-Tc superconductors," *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, IEEE Press, Piscataway, NJ.

Meredith, J.S., G. Alvarez, T.A. Maier, T.C. Schulthess, and J.S. Vetter. 2009. "Accuracy and performance of graphics processors: a Quantum Monte Carlo application case study," *Parallel Comput*. **35**, 151–163.

Summers, M.S., G. Alvarez, J. Meredith, T.A. Maier, and T.C. Schulthess. 2009. "DCA++: A case for science driven application development for leadership computing platforms," *J. Phys.: Conf. Ser.* **180**, 012077.

Research results to date and/or path forward

Building on the experience gained from these studies, researchers plan to seek answers to the important question of optimal inhomogeneity (i.e., a certain type of inhomogeneity for which the critical temperature is maximized). This work will entail simulations of other types of inhomogeneity, such as the charge- and spin-modulated striped state that has been observed in experiments. The team also plans to explore the related question of whether critical temperatures can be increased in composite systems such as artificial cuprate multilayers or heterostructures, in which inhomogeneity is artificially imposed. Research on these systems could prove especially fruitful. Based on the insights derived from the efforts, scientists may be able to predict systems that can be artificially designed with specific optimized properties as well as functionality.

Contact

Thomas Schulthess

Swiss Federal Institute of Technology in Zurich
and Swiss National Computing Centre

schulthess@cscs.ch

Charge Patching Method for Electronic Structures and Charge Transports of Organic and Organic-Inorganic Mixed Nanostructures

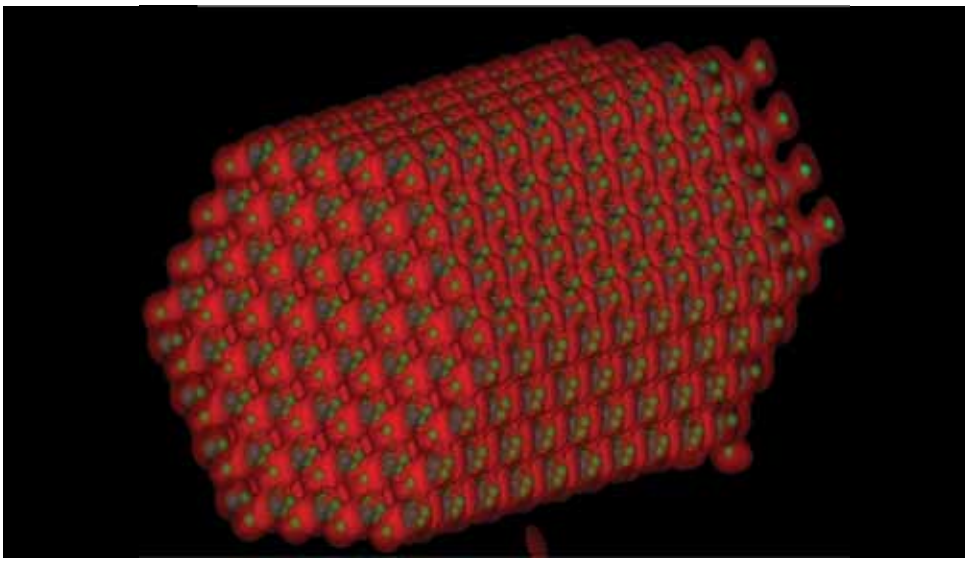

The ab initio, self-consistent charge density of a 2,010-atom ZnO nanorod with 17,520 electrons using the LS3DF code. Image courtesy Lin-Wang Wang, LBNL.

Abstract

There is no easy experimental measurement of a nanosystem's internal electronic states, and the complexity quickly overwhelms first-principles simulations. Using a new, linear-scaling, 3-D code, researchers are calculating the internal electric fields of large nanocrystal structures.

Key collaborators

Lin-Wang Wang, Sefa Dag, Shuzhi Wang (LBNL).

Role of the Jaguar leadership system

This team of researchers is working to better understand the intrinsic electronic properties of nanostructures such as the electrical dipole moment (the internal electronic state) of a zinc oxide nanorod. Jaguar's petascale capability is enabling unprecedented fidelity in modeling of the surface atomic structure of this nanorod and, in particular, how molecules are attached to that surface. LS3DF, a new, linearly scaling, 3-D electronic structure application based on density functional theory, is used to self-consistently predict the charge density of the nanosystem. Currently, computed nanosystems consist of a few thousand atoms, but the novel linear scaling method in LS3DF is algorithmically efficient enough for the task. The computed charge density and associated potential lead to the electronic structure of the system. It is the first time such a realistic surface model has been constructed using density functional theory to describe the polar attributes of this large nanostructure. Wang's research team, winners of the 2008 ACM Gordon Bell Prize for algorithm innovation in their LS3DF application, are also able to compute the change of electric fields as induced by elastic stress (the piezoelectric effect).

Importance of the research

This electronic structure application for first-principles petascale investigations will open the way for new nanostructure and material-science investigations. The systems under investigation can be as large as tens of thousands of atoms, the solutions of which can shed new light, for example, on how the internal electric field is formed and structured. These electric fields in a nanosystem can fundamentally change the electronic structure of the system (e.g., reduce the band gap, alter the carrier transport, and change the exciton wave functions). Researchers have also studied nonequilibrium semiconductor alloy systems, which have potential application for solar cells and solid-state lighting. Using the results of this team's calculations, indications are that the ZnTe:O alloy can have a theoretical solar cell efficiency of 60 percent, far exceeding the traditional limit.

Computational science applications employed

The researchers are using the LS3DF application, which employs a Gordon Bell–winning novel algorithm that results in linearly scaled density functional theory electronic structure solutions. A benchmark test run, for example, showed that LS3DF was able to sustain a 442-teraflop performance on 150,000 processing cores on Jaguar.

Institutional collaborations

The LBNL researchers used the early science applications allocation on Jaguar given them by the OLCF.

Acknowledgments

The research is supported through the DOE Basic Energy Sciences program.

Relevant publications

Wang, L.W. 2009. "Computational challenges for nanostructure solar cells," *Energy & Env. Sci.* **2**, 944–955.

Wang, L.W. 2008. "Large scale ab initio nanostructure electronic structure calculations for energy applications," *J. Phys.: Conf. Ser.* **125**, 012059.

Wang, L.W., B. Lee, H. Shan, Z. Zhao, J. Meza, E. Strohmaier, and D. Bailey. 2008. "Linear scaling 3D fragment method for large-scale electronic structure calculations," *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing,* IEEE Press, Piscataway, NJ.

Research results to date and/or path forward

The researchers found a surprisingly large dipole moment in the nanorod. The implications of this finding lead to the conclusion that in real nanosystems there could be compensation mechanisms that tend to reduce the total dipole moment. As a next step the team plans to investigate different compensation mechanisms in realistic nanosystems. The researchers are at an early stage in understanding the charging effects in such nanosystems, which replace the usual doping effects in bulk semiconductor systems. The current large-scale calculations allow them to characterize some important aspects of those charging effects in nanosystems.

Contact

Lin-Wang Wang

Lawrence Berkley National Laboratory

lwwang@lbl.gov

Petascale Simulations of Strongly Correlated Electron Systems Using the Multi-Scale Many-Body (MSMB) Formalism

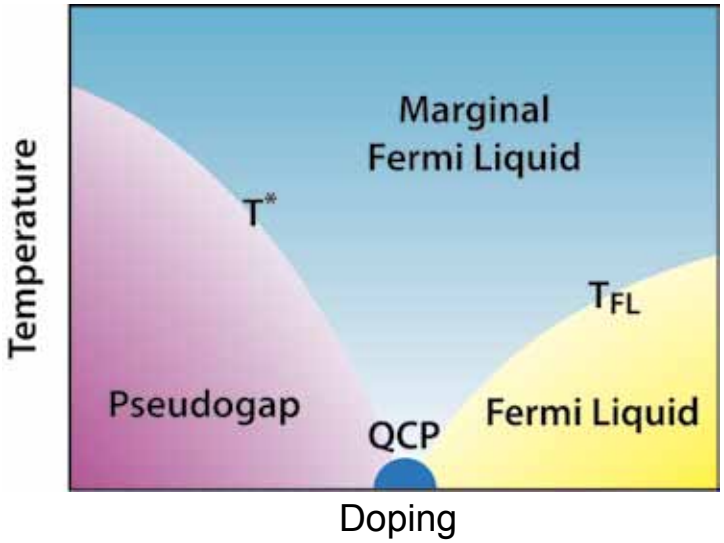

Abstract

Materials containing strongly correlated electrons may well open the door to new devices in the semiconductor industry, yet the demands of simulating such systems quickly overwhelm even the most powerful supercomputers. This project explores a promising method for introducing approximations into a simulation that make it manageable without significantly reducing the accuracy of the solution.

The generic phase diagram of the cuprates. Dynamical cluster approximation simulations of the Hubbard model now reproduce all of the essential features, including the superconductivity the quantum critical point, the pseudogap, strange metal, and Fermi liquid doping regions. Image courtesy Mark Jarrell, LSU.

Key collaborators and affiliations

Mark Jarrell (Louisiana State University); Karen Tomko (The Ohio Supercomputer Center); Ed D'Azevedo, Thomas Maier (ORNL); Zhaojun Bai, Richard Scarlettar (UCD); Juana Moreno (University of North Dakota).

Role of the Jaguar leadership system

The computational resources required to perform chemically specific simulations of the quantum physics of many-body correlated electronic systems currently grow exponentially with the number of particles being simulated. (For example, if it takes 10 seconds to simulate two particles, it will take 10^2, or 100 seconds, to simulate four particles and 10^4, or 10,000 seconds, to simulate eight particles.)

Because it is essentially impossible to take this approach with large systems, Mark Jarrell and his team use the MSMB formalism to judiciously introduce approximations that have a minimal impact on the final solution. This approach provides an exact solution to the interaction of particles when they are close together and quantum effects are, therefore, at their strongest, but it approximates the interaction as particles are farther apart. The approach is also far more manageable, as the required resources grow algebraically with the size of the system rather than exponentially.

Nevertheless these solution methods demand petascale—or greater—resources. Consider the simplest system researchers might want to solve—with 20 lattice sites and 100 time slices. The computational time scales to several thousand hours. Despite these significant requirements, the algebraic scaling displayed by the MSMB methodology represents a huge advantage over previous methods. However, if researchers want to treat a more realistic system with 100 lattice sites, multiple bands, and the same number of time slices, then the total memory and processor time required would surpass even the petascale.

Importance of the research

Strongly correlated materials have many potential applications. Transition metal oxides show great promise for novel applications in the semiconductor industry, going beyond complementary metal-oxide semiconductor transistor devices for future information processing technologies based, for example, on spin variables. Highly correlated electron

systems exhibit coupling between orbital, charge, and spin ordering and may enable new devices by greatly enhancing their sensitivity to different applied fields.

Computational science applications employed

The inner loops of the team's applications are LAPACK calls, which rely on the BLAS level-three engines (mostly GEMM); hence, the floating point rate per processor is extremely high. The main challenge the team encountered during this project was hiding the latency of the message passing interface all-to-all communication primitives. By improving the data layout and using other tricks, this team is able to effectively hide the latency and achieve nearly linear scaling up to more than 10,000 processors.

Institutional collaborations

The applications were developed by an interdisciplinary team involving researchers from The Ohio Supercomputer Center, Louisiana State University, ORNL, and UCD.

Acknowledgments

Support for this project is provided by DOE's SciDAC program (http://www.scidac.gov/).

Relevant publications

Khatami, E., C.R. Lee, Z.J. Bai, R.T. Scalettar, and M. Jarrell. 2009. "Dynamical mean field theory cluster solver with linear scaling in inverse temperature," arXiv:0904.1239.

Yang, S.X., H. Fotso, J. Liu, T.A. Maier, K. Tomko, E.F. D'Azevedo, R.T. Scalettar, T. Pruschke, and M. Jarrell. 2009. "Parquet approximation for the 4×4 Hubbard cluster," arXiv:0906.4736.

Research results to date and/or path forward

There are too many possible applications of the MSMB formalism to list. In the short time the team has had access to Jaguar, the method has matured very quickly, and the team's research promises new advances in high-temperature superconducting cuprates.

Contact

Mark Jarrell

Louisiana State University

jarrellphysics@gmail.com

Accurate Quantum Understanding of Adsorption on Surfaces: Water on Graphene (AQUA)

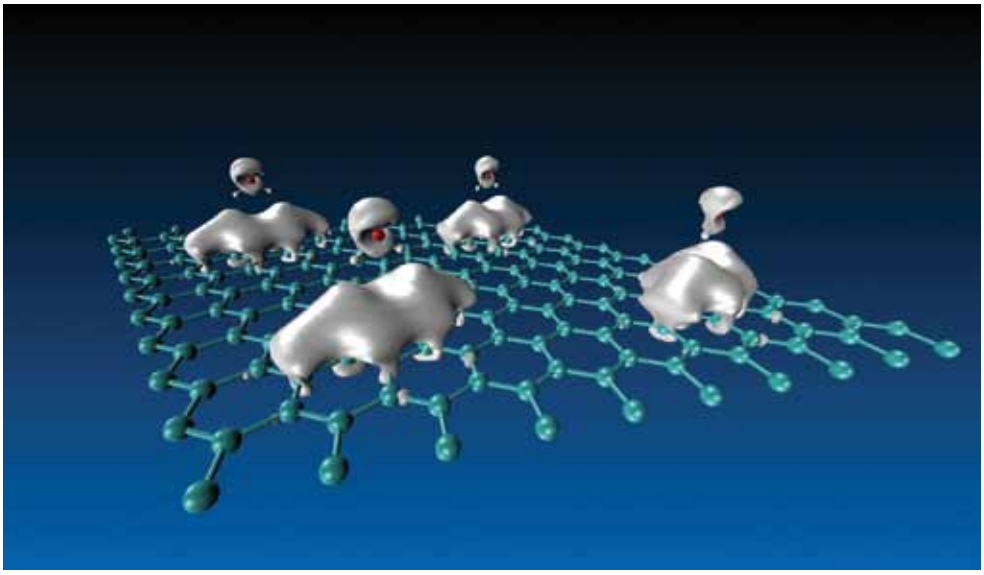

Abstract

Researchers are using 2 million hours to calculate the energy and geometry of water molecule adsorption on a graphene sheet. The Monte Carlo code CASINO, which they have scaled to 40,000 cores, will ultimately help explain how Earth forms materials at the extreme conditions of its deep interior.

The image shows how electrons rearrange when four water molecules bond to a sheet of graphene. Image courtesy Dario Alfe, UCL.

Key collaborators and affiliations

Dario Alfe, Angelos Michaelides (University College London).

Role of the Jaguar leadership system

The researchers are using 2 million processor-hours on Jaguar to calculate the adsorption energy and geometry of a water molecule on a graphene sheet (a single layer of graphite). Adsorption is the accumulation of atoms or molecules on the surface of a material. In a material, the atoms bond with other atoms, but surface atoms are not wholly surrounded by other atoms and therefore can attract adsorbates. Most surfaces in nature are covered by a film of water. The nature of the bonding depends on the species involved. It is difficult to describe such systems because of the large number of particles and the many degrees of freedom they possess. The researchers use a QMC technique, which has never before been used to the level of accuracy required by this project. In QMC methods, scientists attempt to solve the quantum mechanical Schrödinger equation by simulating a random process that is mathematically equivalent. A classic example of a random process is repeatedly flipping a coin to determine the chance of achieving heads, a technique that requires longer and longer simulation times to achieve higher precision.

The Monte Carlo method the researchers employ has never been used at this level of accuracy. It is highly computer intensive, requiring the generation of random numbers on tens of thousands of compute nodes. Such work cannot be performed without unprecedented computational power. The technique requires between 1,000 and 10,000 times more computer time than more conventional—and more approximate—quantum mechanical tools, such as density functional theory. It is only with the availability of computational power of the size of Jaguar that research at this level of accuracy is possible.

Importance of the research

To date, despite their great importance, the adsorption energies and structure of water molecules have not been well established on any carbon surface, either experimentally or theoretically. This research applies theoretical-computational first-principles techniques to problems that ultimately will explain how the Earth forms materials at the extreme conditions of its deep interior and to surface problems related to catalysis and hydrogen storage. Knowledge of the interaction between water and carbon surfaces is also crucial

to understanding lubrication to reduce friction between moving parts. It is also important to understanding how ice forms on plane wings. In space, molecules rarely form on their own because of the very low density of atoms, but carbon or silicates do bond with water to form interstellar dust grains. A detailed understanding of the structural and phase behavior of water at the nanoscale will contribute to all such research areas.

Computational science applications employed

The researchers use QMC techniques, and in particular the CASINO application, developed at Cambridge University with some important contributions from University College London. The researchers are using 40,000 cores (well beyond their previous record of 8,000 cores), and they envisage additional tunings and developments that will bring the scalability of CASINO to more than 100,000 cores.

Institutional collaborations

The CASINO code is under development at Cambridge University, and the researchers at University College London made adaptations to it. The current research is a collaboration between University College London and ORNL.

Acknowledgments

The research is supported by the Engineering and Physical Sciences Research Council of the United Kingdom.

Relevant publications

Ma, J., D. Alfe, A. Michaelides, and E. Wang. 2009. "The water-benzene interaction: insight from electronic structure theories," *J. Chem. Phys.* **130**, 154303.

Posse, M., and D. Alfe. 2008. "Structural properties and enthalpy of formation of magnesium hydride from quantum Monte Carlo calculations," *Phys. Rev. B* **77**, 04103.

Sola, E., and D. Alfe. 2009. "Melting of iron under earth's core conditions from diffusion Monte Carlo free energy calculations," *Phys. Rev. Lett.* **103**, 078501.

Research results to date and/or path forward

To date the researchers are able to calculate two out of the approximately 30 points needed to characterize the adsorption geometry and energy with the necessary accuracy in the calculated energies of approximately 5 million electric volts. The researchers are beginning to determine the binding energy curve between a water molecule and graphene. With initial partial results, they have indications that the water-graphene interaction is very weak.

Contact

Dario Alfe

University College London

d.alfe@ucl.ac.uk

Full-Band, Atomistic Simulation and Design of Next Generation Nanotransistors

Abstract

The number of transistors that can fit on a modern computer chip may soon reach its limit, and the materials have been pushed to their performance peak. A new generation of novel transistors will increase the performance of today's electronic devices. Engineers are simulating candidates at the atomic scale.

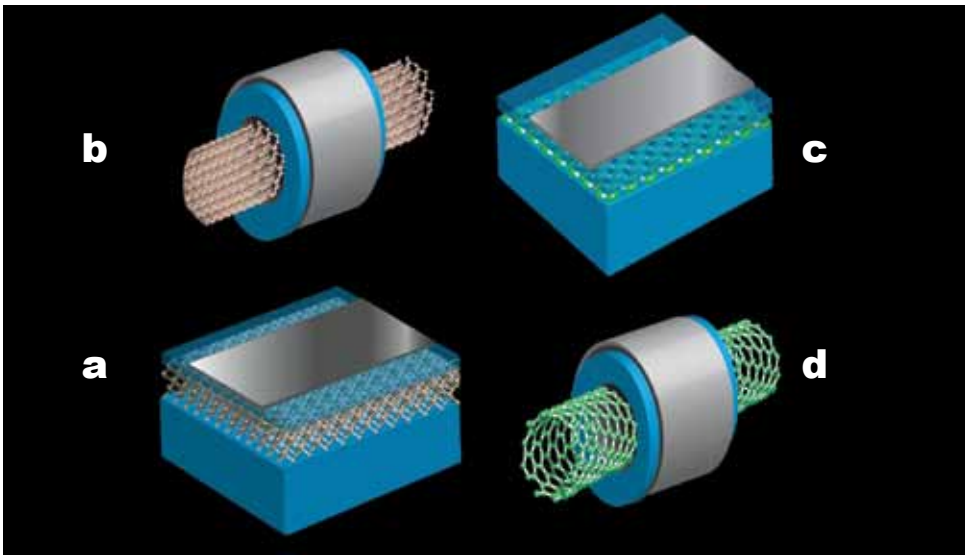

(a) Single-gate and double-gate ultra-thin-body FET made of Si, Ge, or III-V semiconductors; (b) gate-all-around nanowire FET; (c) graphene nanoribbon FET; and (d) coaxially gated carbon nanotube FET. Image courtesy Gerhard Klimeck, Purdue University.

Key collaborators and affiliations

Gerhard Klimeck, Mathieu Luisier (Purdue University).

Role of the Jaguar leadership system

As the size of transistors—the active components of electronic devices—approaches the nanometer scale, semiconductor manufacturers find it increasingly difficult to fully understand and control the quantum mechanical effects that arise. Conventional transistor structures such as the metal-oxide-semiconductor field-effect transistors (FETs) must be replaced by novel device concepts like nanowire or tunneling FETs. Technology computer-aided design tools will accelerate the development of future transistor architectures. However, no technology computer-aided design tools exist that can go beyond classical approximations and treat the flow of electrons and holes on a quantum mechanical level. An atomistic description and a full-band approach that offers an accurate representation of the electronic properties of any semiconductor material are required. More physics is not enough, though; rapid simulation turnaround is needed to impact engineering. Hence, researchers have developed OMEN, a massively parallel, atomistic, full-band, one-,2-D, and 3-D quantum transport solver. An analysis shows that most of the computational tasks of an atomistic quantum transport solver, especially the solution of several hundred thousand sparse linear systems of equations, can be distributed across multiple processing elements (cores). This means that with a high number of cores, large, realistic device structures with up to 200,000 atoms can be simulated in a relatively short time. Such experiments, which would take years on a single processor or weeks on a small cluster, can now be executed within hours using more than 100,000 cores on Jaguar.

Importance of the research

Moore's Law suggests that the number of transistors per chip doubles every 18 months. Forty years of aggressive scaling have pushed the size of devices down to the nanometer scale. Currently, the gate length of a transistor (i.e., the distance between the place carriers are injected and the place they are collected) measures less than 45 nanometers and will soon be scaled down to below 32 nanometers. At the same time, the power density per area in a chip has continued to increase and is now in the same order as that of a nuclear reactor. This has led the semiconductor industry to look for alternative transistor architectures that consume less power. Ideally, new designs should first be investigated using physics-

based computer simulators. The growing availability of large computational resources like Jaguar has boosted the development and use of complex device simulators so that a "cheap" design of next-generation nanotransistors is now possible.

Computational science applications employed

OMEN is based on the solution of the Schrödinger and Poisson equations with open-boundary conditions and on the nearest-neighbor tight-binding model. The team solves sparse systems of equations for different bias points, injection energies, and momentum vectors that form a three-level MPI parallelization. Hence, OMEN can simultaneously solve thousands of quantum transport problems with an almost perfect scaling of the simulation time. There is a fourth level of parallelism, spatial domain decomposition, that requires more communication between processors. To approach the petascale, the workload of each task is predicted in advance, resulting in a dynamic core allocation. OMEN used more than 4 million processor-hours and up to 147,456 cores on Jaguar's 150,000-core XT5 partition with four levels of MPI parallelization and reached a sustained performance of 504 teraflops.

Institutional collaborations

Several semiconductor companies have formed a consortium called Semiconductor Research Corporation (SRC) and a Nanoelectronics Research Initiative to leverage the development of novel transistor architectures. The research team is part of the Midwest Institute for Nanoelectronics Discovery, where it collaborates with the University of Notre Dame, Penn State University, the University of Illinois, the University of Michigan, the University of Texas, and Florida State University.

Acknowledgments

Support for this work comes from the National Science Foundation, SRC, the Nanoelectronics Research Initiative, and Intel.

Relevant publications

Luisier, M., and G. Klimeck. 2009. "Atomistic full-band simulations of Si nanowire transistors: with the effects of electron-phonon scattering," accepted for publication in *Phys. Rev. B.*

Luisier, M., and G. Klimeck. 2009. "Atomistic full-band design study of InAs band-to-band tunneling field effect transistors," *IEEE Elec. Dev. Lett.* **30**, 602–604.

Luisier, M., and G. Klimeck. 2008. "Full-band and atomistic simulation of realistic 40nm InAs HEMT," *Proceeding of the International Electron Device Meeting (IEDM) 2008*, 887–890.

Research results to date and/or path forward

In the next phase of the research, the simulator will be calibrated toward experimental data provided by the Midwest Institute for Nanoelectronics Discovery partners. Many simulations on 10,000 to 100,000 cores will be conducted to identify potential deficiencies of the physical models. The final goal is to propose a device structure that can be manufactured and will help reduce the energy consumption of electronics devices. In parallel, the numerical algorithms of OMEN will be improved to tackle even larger transistor structures.

Contact

Gerhard Klimeck
Purdue University
gekco@purdue.edu

Denovo, a Scalable HPC Transport Code for Multiscale Nuclear Energy Applications

Abstract

For nuclear reactor simulation, the size of the equations is tremendous—five orders of magnitude in space and ten in neutron energy. Researchers at ORNL are extending a parallel transport solver called Denovo to develop a first-of-its-kind, mathematically consistent, two-level approach to the multiscale challenge of simulating complex processes in novel nuclear reactors.

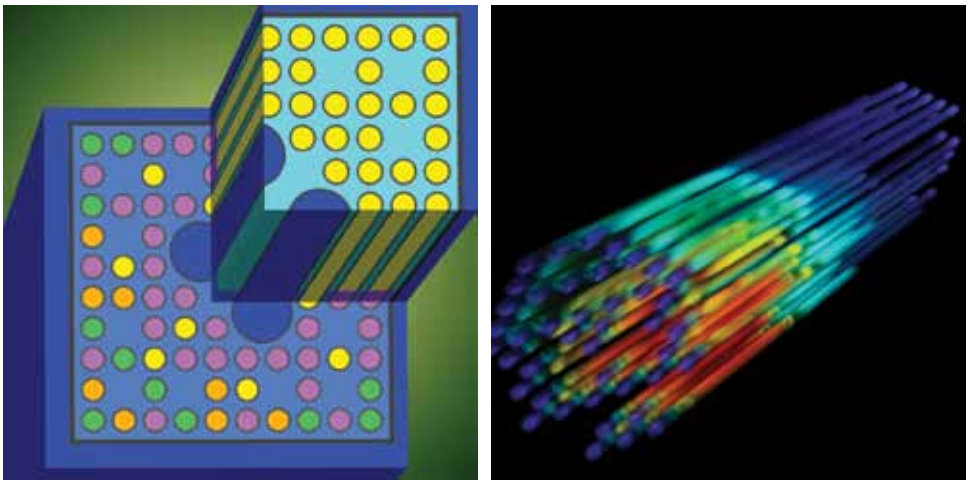

Illustration of boiling water reactor assembly simulations. Images courtesy Tom Evans, ORNL.

Key collaborators and affiliations

Tom Evans, Scott Mosher, Kevin Clarno (ORNL).

Role of the Jaguar leadership system

Using the Denovo discrete ordinates (S_N) radiation transport formulation, this team is using 5 million computational hours on Jaguar to solve a variety of large-scale problems in nuclear engineering and technology. For fuel storage and transport, calculations were performed for criticality safety, safeguards, and site-boundary radiation dosimetry (calculation of the absorbed dose in matter resulting from exposure to radiation). For nuclear facilities, they were performed for shielding, dosimetry, and material safeguards. For reactor analysis, they were performed for core design, fuel performance, material/component integrity, and dosimetry. And for ITER design analysis, they were performed for shielding, dosimetry, heating, and material integrity.

Nuclear technology applications require very accurate descriptions of high-energy radiation fields. Radiation transport calculations are among the most computationally intensive simulations in computational physics—in both memory capacity and operational count. The fundamental state vector in radiation transport, the radiation intensity, is a function of seven independent variables representing the radiation (e.g., neutrons, gamma rays, photons, x-rays) space, angle, energy, and time. Until very recently, even moderately sized (many millions of unknowns) 3-D calculations were beyond the reach of most computational architectures. However, with the petascale computational capability of Jaguar at the OLCF, the Denovo team can now perform simulations of unprecedented size, with greater than 88 trillion degrees of freedom (total solved unknowns). It is therefore now possible to apply high-fidelity transport simulations to a whole new class of problems whose solutions have traditionally resulted from inadequate and unreliable low-order approximations.

Importance of the research

New and existing nuclear reactors will have an ever-increasing role in U.S. energy production. The design and certification of new reactors will require extensive high-fidelity simulation tools, as many experimental facilities used to design and test existing reactors are no longer available. In addition to direct reactor design analysis, the expansion of nuclear

power will require a host of supporting engineering calculations, particularly in fuel performance, material and fuel performance integrity, and safety and dosimetry. Furthermore, high-risk/high-payoff, novel approaches to energy production such as the ITER fusion project have significant modeling and simulation needs. High-fidelity modeling and simulation will be immensely important in the development of U.S. energy production for the 21st century and beyond.

Computational science applications employed

The Denovo S_N transport code is a very recent development by ORNL's Nuclear Science and Technology Division, dating back only to August 2007. Its primary purpose is to provide a new, robust, 3-D radiation transport capability for a wide variety of nuclear technology applications investigated within DOE and Nuclear Regulatory Commission programs supported by the Nuclear Science and Technology Division. Denovo uses the latest state-of-the-art parallel solvers that have been developed within the applied math and radiation transport communities over the last decade. In this project Denovo was routinely used to run large simulations requiring between 20,000 and 60,000 cores on Jaguar. Denovo was released to the general user community in Standardized Computer Analyses for Licensing Evaluation (SCALE) 6. SCALE is available at http://rsicc.ornl.gov.

Acknowledgments

ORNL Laboratory Directed Research and Development provided funding for this project.

Relevant publications

Evans, T., K. Clarno, and J. Morel. "A transport acceleration scheme for multigroup discrete ordinates with upscattering," submitted to *Nuc. Sci. and Eng.*

Ibrahim, A., D. Peplow, T. Evans, P. Wilson, and J. Wagner. "Modifications of the hybrid Monte Carlo/SN techniques for high-fidelity/high-resolution modeling," submitted to ANS Radiation Protection and Shielding Topical Meeting.

Peplow, D., T. Evans, and J. Wagner. 2009. "Simultaneous optimization of tallies in difficult shielding problems," *Nuc. Tech.* **168**, 785–792.

Research results to date and/or path forward

These projects are geared toward performing full-core reactor analysis using a new hybrid radiation transport methodology, namely one that combines the strengths of the S_N discrete ordinates approach with a Monte Carlo method. Building on Denovo's capabilities, the researchers will use very accurate Monte Carlo calculations to perform full-core reactor analysis calculations. This is a new approach to reactor analysis, as Monte Carlo is generally thought to be too computationally intensive for full-core calculations.

Contact

Tom Evans

Oak Ridge National Laboratory

evanstm@ornl.gov

Scalable Simulation of Neutron Transport in Fast Reactor Cores

Abstract

Researchers are using UNIC, an unstructured, deterministic neutron transport code, to develop and test neutron transport algorithms for fast reactors. Through high-fidelity simulations, the team aspires to greatly reduce the uncertainties in reactor design calculations and improve reactor safety.

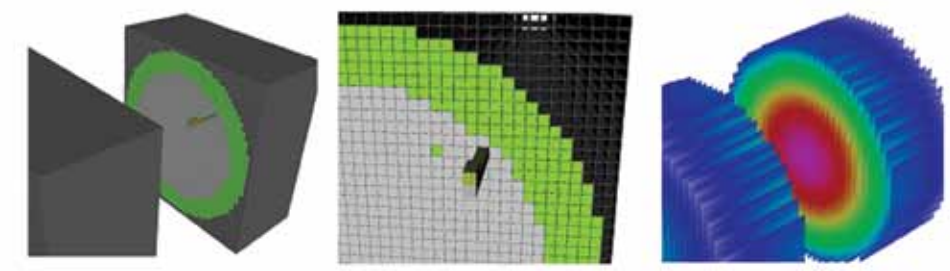

Two pictures (left and center) of ZPR 6/6A geometry and uranium-235 plate power distribution (with separated matrix halves). The gray indicates the matrix tube and drawer fronts that are loaded into each tube position. The solid green squares are 2 in. depleted uranium metal blocks directly loaded into the tubes surrounding the main core and acting as a neutron blanket. The plot at the right shows the enriched uranium plate power with the matrix halves separated. Images courtesy of Dinesh Kaushik, ANL.

Key collaborators and affiliations

Dinesh Kaushik, Satish Balay, Paul Fischer, Paul Hovland, Lois McInnes, Mike Minkoff, Boyana Norris, Andrew Siegel, Barry Smith, Michael Smith, Tim Tautges, Allan Wollaber, Won Sik Yang (ANL); Pablo Carrica, Fred Stern (University of Iowa); William Gropp (UIUC); David Keyes (Columbia University and King Abdullah University of Science and Technology); Giuseppe Palmiotti, Cristian Rabiti (Idaho National Laboratory); Jean Ragusa (Texas A&M University).

Role of the Jaguar leadership system

This team is performing detailed, high-resolution simulations of neutron transport in fast reactor cores using the UNIC code, which is being developed as part of the DOE Nuclear Energy Advanced Modeling and Simulation Program. The team is carrying out Zero Power Reactor Experiment (ZPR) 6/6a simulations with high fidelity in space, angle, and energy. Simulations are yielding solutions to this complex multiphysics problem with more than 120 billion degrees of freedom (variables) with excellent weak scalability on up to 122,800 processor cores. The goal of these simulations is to reduce the uncertainties and biases in reactor design calculations by progressively replacing existing multilevel averaging (homogenization) techniques with more direct solution methods based on highly accurate solutions of the governing equations.

UNIC is being developed to solve large-scale nuclear reactor core problems governed by the seven-dimensional Boltzmann equation for neutron transport. To model the complex geometry of a reactor core, billions of spatial elements, hundreds of neutron directions ("angles"), and thousands of neutron energy groups are necessary, which leads to problem sizes with petascale degrees of freedom. Such calculations are not possible without large-scale parallel platforms such as Jaguar.

Importance of the research

Fast reactors—which do not use a moderator to slow the neutrons that split fissile nuclei—show promise as both a source of electrical power and one producing minimal nuclear waste. There are two key obstacles to deployment: capital cost and demonstration of safety. UNIC can help in overcoming these barriers. Existing modeling and simulation tools are based on various approximations and rely on correlations derived from experiments.

However, a modern, integrated design tool is needed that will allow automated data transfers among different design activities and facilitate design optimizations. UNIC is striving to provide such a tool to reactor designers. Through these high-fidelity simulations, the team aspires to (1) free reactor designers to explore innovative design concepts with minimal need for new physical "calibration" experiments and (2) reduce uncertainties for existing designs to enable more economically efficient design/operation.

Computational science applications employed

UNIC allows a highly detailed description of a nuclear reactor core. It uses state-of-the-art algorithms to solve large, complex problems arising in nuclear reactor analysis. Written in FORTRAN 95, UNIC uses the PETSc library (http://www.mcs.anl.gov/petsc/petsc-as/) for solving the linear systems arising from the discretized form of the Boltzmann equation. With these early science resources, the researchers were able to carry out highly detailed simulations of the ZPR 6/6a, resolving the spatial heterogeneity for the first time.

Institutional collaborations

The UNIC team works in close collaboration with researchers from ORNL, LLNL, and PNNL. The industrial partners (General Electric and IBM) are also involved in code validation for different types of reactors. The meshing capability is provided by a CUBIT geometry and mesh generation toolkit (http://www.cubit.sandia.gov/) from SNL, and postprocessing is done using VisIT software (https://wci.llnl.gov/codes/visit/) from LLNL.

Acknowledgments

Support comes from DOE and the DOE Office of Nuclear Energy, Science and Technology through the Advanced Fuel Cycle Initiative.

Relevant publications

Kaushik, D., M. Smith, A. Wollaber, B. Smith, A. Siegel, and W.S. Yang. 2009. "Enabling high fidelity neutron transport simulations on petascale architectures," *Proc. of Supercomputing 2009* (Gordon Bell Award Finalist Paper), Portland, Oregon. Accepted.

Smith, M.A., D. Kaushik, A. Wollaber, W.S. Yang, B. Smith, C. Rabiti, and G. Palmiotti. 2009. "Recent research progress on UNIC at Argonne National Laboratory," International Conference on Mathematics, Computational Methods & Reactor Physics, Saratoga Springs, New York, May 3–7, 2009.

Fischer, P., D. Kaushik, D. Nowak, A. Siegel, W.S. Yang, and G.W. Pieper. 2008. "Advanced simulation for fast reactor analysis," *SciDAC Review* **9** (Fall), 12–21.

Research results to date and/or path forward

This team will continue validating the ZPR 6/6a simulation results with experimental data on the leadership systems at the OLCF. Over the coming months, members aim to improve the per-processor performance while maintaining high parallel efficiency by employing better algorithms (such as p-multigrid and parallelization of energy dimension).

Contact

Dinesh Kaushik

Argonne National Laboratory

kaushik@mcs.anl.gov

Lattice QCD

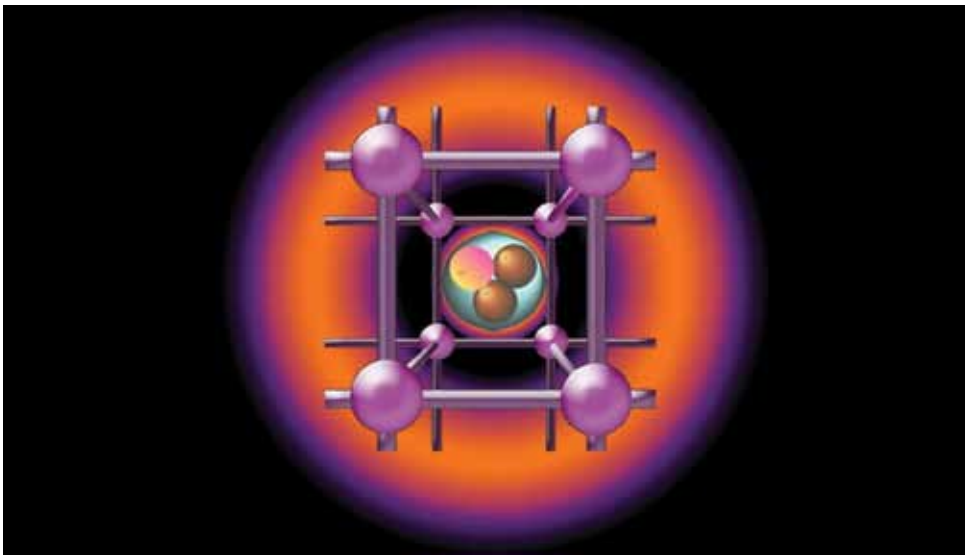

Abstract

The strong nuclear force by which gluons bind quarks—thereby creating protons and neutrons—is also responsible for holding the protons and neutrons together to form the nuclei of atoms and, therefore, us. This project takes an in-depth look at the energy states created by mesons (with two quarks) and baryons (with three).

An artist's conceptualization of a proton or a neutron and its quarks, sitting inside a lattice, surrounded by the strong force. Lattice QCD calculates how subatomic particles behave—a miniature, mathematical representation of the minuscule space between particles. Graphic by David Chopard, Jefferson Lab.

Key collaborators and affiliations

Paul Mackenzie (Fermi National Accelerator Laboratory); Jozef Dudek, Robert Edwards, David Richards, Bálint Joó, Christopher Thomas (Jefferson Lab); John Bulava, Justin Foley, Colin Morningstar (Carnegie Mellon University); Eric Engelson, Stephen Wallace (Univeristy of Maryland); Keisuke Juge (University of the Pacific); Michael J. Peardon, Sinead M. Ryan (Trinity College, Dublin, Ireland); Nilmani Mathur (Tata Institute, Bangalore, India); Huey-Wen Lin (University of Washington); Saul Cohen (Boston University).

Role of the Jaguar leadership system

This project targets better understanding of quantum chromodynamics (QCD), or the theory that matter is made up of particles called quarks, which are bound together by gluons. These gluons are massless particles, whereas the quarks have small masses. The gluons then bind the quarks into protons and neutrons, and residual QCD forces in turn bind protons and neutrons into nuclei, which are at the heart of atoms. It has been found that the masses of the quarks make up less than 10 percent of the mass of a proton, and the rest appears to come from the binding energy of the gluons. Mackenzie and his team are investigating the properties of the QCD force through exploration of the energy states of the simplest QCD systems: mesons (two quarks) and baryons (three quarks). In particular, the researchers are trying to understand whether the observed energy levels are just excitations of ordinary QCD states or are other kinds of states such as multiparticle states. QCD also predicts states in which gluons bind to themselves. These hybrid states are being sought experimentally, and their production forms the centerpiece of the GlueX experiment, which is a flagship component of the planned 12-gigaelectron-volt upgrade to the accelerator at Jefferson Lab's Continuous Electron Beam Accelerator Facility (http://www.gluex.org/). The team is attempting to compute properties of these states and estimate their production rates to guide the lab's experimental program.

Simulations of QCD are computationally extremely demanding. At their heart they need to solve the Dirac equation, which describes the interaction of quarks and gluons, many, many times over. The simulation technique is statistical, with one data sample made from the previous one via a molecular dynamics algorithm. And the quality of the final results improves as more data samples are generated. Combined, these factors make lattice QCD a voracious application. Generation of the data samples alone can

take teraflop-years worth of computer resources. To make timely headway, it is necessary to use terascale and petascale systems such as Jaguar.

Importance of the research

To understand the properties of QCD is to understand the matter of which everything is composed. More concretely, however, the team hopes to guide and inform current and future experiments at Jefferson Lab. The computations are also necessary to meet milestones and performance measures set out in the Nuclear Science Advisory Committee (NSAC) long-range plan (http://www.science.doe.gov/np/nsac/nsac.html). In particular, the calculations directly contribute to satisfying a 2015 NSAC performance measure that calls for "precision measurements of fundamental properties of the proton, neutron, and simple nuclei for comparison with theoretical calculations to provide a quantitative understanding of their quark substructure."

Computational science applications employed

The team developed the code base used in this project, namely Chroma, which is a C++ code designed for lattice QCD computations. Current simulations use 16,384-core partitions. Strong scaling prevents the researchers from using larger partitions with the current lattice size (which was considered large on the previous generation of machines). In the future, however, the team aims to calculate larger systems, which will require larger partitions.

Institutional collaborations

This project involves several diverse institutions in the United States and abroad. Currently the data samples are generated on Jaguar and then shipped to Jefferson Lab, where the researchers make them available to team members and the U.S. QCD collaboration for secondary analysis. The team is primarily interested in the spectrum of hadrons and the photoproduction of hybrid mesons. However, the data are also used by the Nuclear Physics with Lattice QCD collaboration, which is investigating forces between nuclei.

Acknowledgments

The project is primarily supported by DOE through programs at Jefferson Lab. Software development is funded through the SciDAC program as part of the lattice QCD application area. Some university colleagues are funded through the National Science Foundation in the United States and through their respective national funding agencies abroad.

Relevant publications

Bulava, J., R.G. Edwards, E. Engelson, J. Foley, B. Joó, A. Lichtl, H.-W. Lin, N. Mathur, C. Morningstar, D.G. Richards, and S.J. Wallace. 2009. "Excited state nucleon spectrum with two flavors of dynamical fermions," *Phys. Rev. D* **79**, 034505.

Lin, H.-W., S.D. Cohen, J. Dudek, R.G. Edwards, B. Joó, D.G. Richards, J. Bulava, J. Foley, C. Morningstar, E. Engelson, S. Wallace, K.J. Juge, N. Mathur, M.J. Peardon, and S.M. Ryan. 2009. "First results from 2+1 dynamical quark flavors on an anisotropic lattice: Light-hadron spectroscopy and setting the strange-quark mass," *Phys. Rev. D* **79**, 034502.

Joó, B. 2008. "Continuing progress on a lattice QCD software infrastructure," *J. Phys.: Conf. Ser.* **125**, 012066.

Research results to date and/or path forward

Currently the quarks in these simulations are a little heavier than in nature, and the blocks of space-time being simulated correspond to a box of about 4 femtometers. The simulation data generated as part of this project will be carried at least into the next year before it is sufficient to support further research. The ultimate goal is to simulate with quarks that have the same mass as those in nature, with a box of about 6 femtometers.

Contact

Paul Mackenzie

Fermi National Accelerator Laboratory

Mackenzie@fnal.gov

Ab Initio Structure of Carbon-14

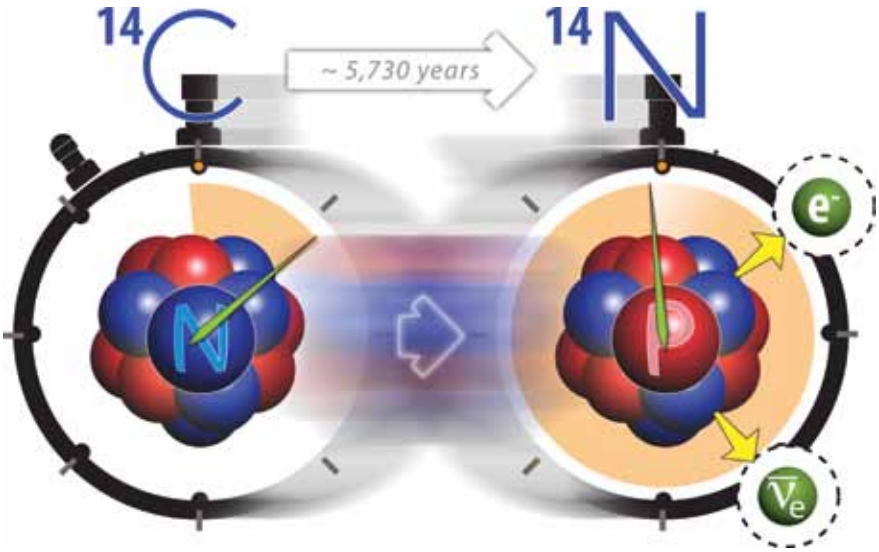

Abstract

Carbon-14's unexpectedly long half-life is a puzzle, enabling accurate dating of relics going back as far as 60,000 years but defying theoretical explanation. This project, the first of its kind to boost realism by including three-nucleon forces, takes a first-principles approach to calculating the low-energy nuclear structure and pertinent reaction rates of this important isotope.

Carbon-14 decays through beta decay, in which the nucleus emits an electron and an antineutrino and becomes a nitrogen-14 nucleus. Graphic by Andy Sproles, ORNL.

Key collaborators and affiliations

David J. Dean, Hai Ah Nam (ORNL); James Vary, Pieter Maris (Iowa State University); Erich Ormand, Petr Navratil (LLNL).

Role of the Jaguar leadership system

This study looks at the nuclear structure of carbon-14 in its low-energy states. Carbon-14 has an unexpectedly long half-life—5,730 years—which has made it useful for accurately dating organic materials as far back as 60,000 years. So far, however, its long half-life has eluded explanation and is, in fact, significantly underestimated by conventional two-body nuclear interactions. This project calculates the low-energy nuclear structure and pertinent reaction rates of carbon-14 using an ab initio (first-principles) approach. It also takes a step beyond earlier calculations by including three-nucleon forces, previously unattainable due to the lack of computational resources. This project would be impossible on any existing system other than Jaguar, as it is the first system with enough memory to solve the nuclear many-body problem in the no-core shell-model framework with a three-body interaction in a nucleus this size.

Importance of the research

One major challenge for nuclear science is to provide a consistent theoretical framework that accurately describes all nuclei and can be used predictively. Toward this aim, this project investigates ab initio two- and three-body interactions derived from effective field theory within the microscopic no-core shell-model approach. These interactions provide a long-sought-after bridge between low-energy nuclear physics and QCD, the fundamental theory describing the strong interaction. Previous calculations have successfully described light nuclei for elements such as lithium-7 (with three protons and four neutrons) and boron-10 (with five protons and five neutrons), but they have fallen short in describing the anomalous half-life of carbon-14. Accurate characterization of carbon-14 not only validates the model and improves our understanding of nuclear interactions, but also is potentially applicable to other elements.

Computational science applications employed

The MFDn application achieves near-perfect scaling and uses nearly the full capacity of Jaguar for carbon-14 calculations. Because it features intensive integer arithmetic operations, however, flop rates are not a realistic measure of true performance. On the other hand, MFDn's capabilities appear unmatched by those of other related applications. With support from SciDAC–UNEDF, the team is continuing to improve the I/O, per-processor calculation efficiency, efficiency of memory usage, and more.

Institutional collaborations

This project has collaborators at ORNL, Iowa State University, and LLNL.

Acknowledgments

Funding for this work comes from the DOE ASCR program office and the SciDAC–UNEDF project.

Relevant publications

Dean and his team completed the largest model space calculations at the end of the calendar year 2009, utilizing over 30 million core-hours, with results presented in the forthcoming paper in the first half of 2010.

Research results to date and/or path forward

The researchers will continue their systematic studies of light nuclei to achieve a better understanding of nuclear structure. They plan to continue pushing the limits of nucleus size and model space with both two- and three-body interactions.

Contact

David J. Dean

Oak Ridge National Laboratory

deandj@ornl.gov

Complex Transport Phenomena in Turbulent Flows

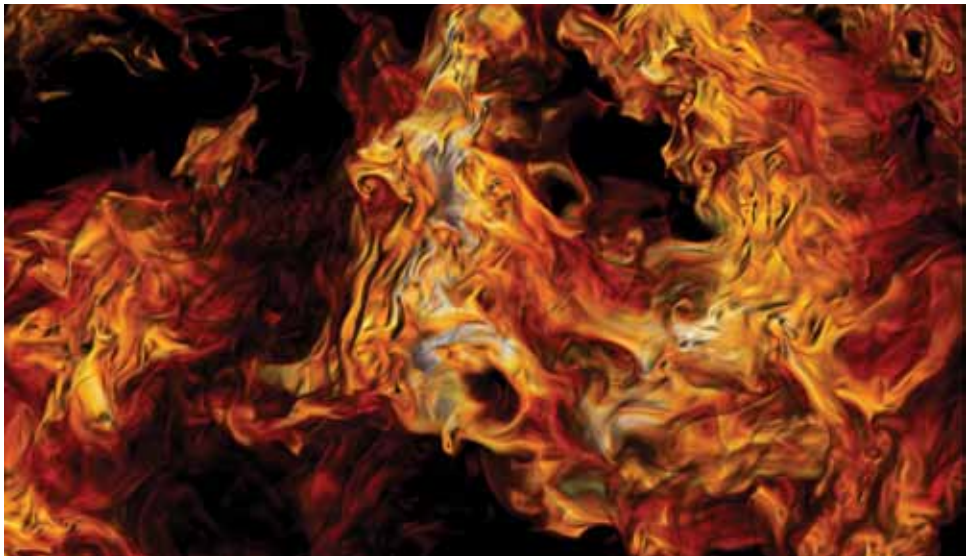

Abstract

Despite decades of study, turbulent flows have resisted theoretical description, which has limited the ability to predict natural phenomena and design better technologies in aeronautics, meteorology, and combustion. The researchers used nearly 10 million computational hours in what was the first simulation of turbulent mixing and dispersion at a resolution of 64 billion grid points.

These simulations track the motions of large numbers of small particles in fluids and closely examine their movement away from one another under the influence of differing lengths and timescales, providing unique insights into flow physics and aiding model development. Visualization by David Pugmire, ORNL.

Key collaborators and affiliations

P.K. Yeung (Georgia Institute of Technology); R.D. Moser (University of Texas); S.B. Pope (Cornell University); A. Pouquet (National Center for Atmospheric Research); J.J. Riley (University of Washington); B.L. Sawford (Monash University, Australia); K.R. Sreenivasan (International Center for Theoretical Physics, Italy, and University of Maryland); D.A. Donzis (University of Maryland); D. Pekurovsky (San Diego Supercomputer Center).

Role of the Jaguar leadership system

In the science of fluid dynamics, turbulent flow is characterized by continuous, chaotic changes in properties, changes that occur over a wide range of scales in both space and time. Turbulent fluids exhibit disorderly fluctuations ranging from the size of the system to the very small. The phenomenon of turbulence can be thought of as a collection of randomized eddies of different scales, with energy gradually being transferred from larger eddies to smaller ones. Despite decades of study, the complexity of turbulent flows has resisted theoretical description, which has limited researchers' ability to predict natural phenomena and design better technologies. This team has used close to 10 million computational hours on Jaguar to perform the first of a series of detailed computations of turbulent mixing, dispersion, stratification, and rotation in fluids.

The researchers have been able to focus on parameter ranges that have previously not been accessible. The main topic is the dispersion at a high Reynolds number (indicating a turbulent flow). In turbulent motion, transport of contaminants is primarily a result of disorderly motions over a range of scale sizes, and the fluid medium can be described as a collection of infinitesimal fluid particles. The growth in the linear size of a pollutant cloud, for example, is described by how quickly a pair of fluid particles may move apart. The range of lengths and timescales in a turbulent cascade increases strongly with the Reynolds number, a fluid dynamics parameter that indicates the spectrum of scales in a system. Many grid points are required to accurately capture the complete range of scale sizes of interest, and massive computational power is necessary. This work is the first to take simulations to the resolution of 64 billion grid points.

Importance of the research

Turbulence is perhaps the most important unsolved problem in classical physics, and an improved understanding will have great significance for both fundamental science and a wide range of applications in nature and engineering. Fluids in motion impact fields such as aeronautics, meteorology, and combustion as well as environmental problems. Turbulence in our atmosphere and oceans affects short- and long-term climate. Closer to home it helps produce engine thrust to keep airplanes in the sky and disperses pollutants in our environment. Coupled with chemistry, turbulence holds the key to the development of cleaner and more efficient combustion devices. All the proposed simulation tasks are set up to achieve results that will be qualitatively new and highly accurate. The science impacts will include new models of pollutant dispersion based on a rigorous understanding of turbulence-flow physics; the modeling of heat and mass transport in combustion, industrial processing, and the environment; and new insights into oceanic (and atmospheric) mixing and dispersion phenomena, which play an important role in global climate change.

Computational science applications employed

The team is using a Fortran MPI code, which is highly scalable and has been successfully benchmarked at a problem size involving half a trillion grid points distributed among the memory of up to 131,072 processing cores. Currently, production simulations are focused on ones involving billions of grid points, which, given the power of the machine, can be carried out quite efficiently.

Institutional collaborations

Several leading researchers based at other institutions provided collaborative input. Important contributions on the computing side were made by D.A. Donzis (Texas A&M University) and D. Pekurovsky (San Diego Supercomputer Center).

Acknowledgments

P.K. Yeung leads two collaborative grants from the National Science Foundation under Office of Cyberinfrastructure PetaApps and Fluid Dynamics programs.

Relevant publications

Donzis, D.A., P.K. Yeung, and K.R. Sreenivasan. 2008a. "Energy dissipation rate and enstrophy in isotropic turbulence: resolution effects and scaling in direct numerical simulations," *Phys. Fluids* **20**, 045108.

Donzis, D.A., P.K. Yeung, and D. Pekurovsky. 2008b. "Turbulence simulations on O(104) processors," Science Track paper at TeraGrid'08 Conference, June 2008.

Lamorgese, A.G., S.B. Pope, P.K. Yeung, and B.L. Sawford. 2007. "A conditionally cubic-Gaussian stochastic Lagrangian model for acceleration in isotropic turbulence," *J. Fluid Mech*. **582**, 423–448.

Research results to date and/or path forward

The team is striving to simulate with unprecedented precision the fluctuations in temperature and molecular concentrations of the infinitesimal fluid elements in turbulence using parameters typical in air and water. In addition to reaching a high Reynolds number, the team will also tackle the special challenges of simulating mixing in parameter regimes in which molecular diffusion is slow compared to the diffusion of momentum. Understanding mixing is especially important in stratified flow such as that found in the ocean.

Contact

Pui-kuen Yeung

Georgia Institute of Technology

pk.yeung@ae.gatech.edu

SPECIFICATIONS

Jaguar Specifications	Total	XT5	XT4
Peak Teraflops	1,645	1,382	263
Quad-Core AMD Opterons	45,376	37,544	7,832
AMD Opteron Cores	181,504	150,176	31,328
Compute Nodes	26,604	18,772	7,832
Memory (TB)	362	300	62
Disk Bandwidth (GB/s)	284	240	44
Disk Space (TB)	10,750	10,000	750
Interconnect Bandwidth (TB/s)	532	374	157
Floor Space (ft^2)	5,800	4,400	1,400
Cooling Technology		Liquid	Air

During the early science period (January–July 2009) at the OLCF, the Jaguar system consisted of an 84-cabinet quad-core Cray XT4 system and 200 new Cray XT5 cabinets, also using quad-core processors. Both the XT4 and XT5 partitions of the system have 2 gigabytes of memory per core, giving the users a total of 362 terabytes of high-speed memory in the combined system. The two systems are combined by linking each to the Scalable I/O Network (SION) and the Spider file system. The XT5 system has 256 service and I/O nodes providing up to 240 gigabytes per second of bandwidth to SION and 200 gigabits per second to external networks. The XT4 has 116 service and I/O nodes providing 44 gigabytes per second of bandwidth to SION and 100 gigabits per second to external networks.

There are four nodes on both the XT4 and XT5 boards. The XT4 nodes have a single AMD quad-core Opteron 1354 “Budapest” processor coupled with 8 gigabytes of DDR2-800 memory. The XT5 is a double-density version of the XT4. It has double the processing power, memory, and memory bandwidth on each node. The XT5 node has two Opteron 2356 “Barcelona” processors linked with dual HyperTransport connections. Each Opteron has directly attached 8 gigabytes of DDR2-800 memory. The result is a dual-socket, eight-core node with 16 gigabytes of shared memory and a peak processing performance of 73.6 gigaflops.

Each node runs Cray’s version of the SuSE Linux operating system that minimizes interruptions to the application codes running on the system, thus giving predictable, repeatable run times for applications. The SuSE Linux operating system on the nodes joins the system services, networking software, communications, I/O, and mathematical libraries, as well as compilers, debuggers, and performance tools to form the Cray Linux Environment. Jaguar supports MPI, OpenMP, SHMEM (shared memory), and PGAS (partitioned global address space) programming models. The OLCF supports compilers from PGI (The Portland Group), Pathscale, and the GNU Project on Jaguar.

Networking capacity at the OLCF is being expanded in parallel with its computing capability to ensure accurate, high-speed data transfer. High-throughput networks among its systems and upgraded connections to ESnet (Energy Sciences Network) and Internet2 have been installed to speed data transfers between the OLCF and its worldwide users. The OLCF core local area network consists of two Cisco 6500 series routers along with a Force 10 E1200 router. The core network provides more than 100 10-gigabit Ethernet ports for intraswitch connections, as well as directly connected hosts using 10-gigabit Ethernet. The OLCF provides more than 1,200 ports of gigabit Ethernet for machines with lesser data-transfer needs.

With the high power density of about 2,000 watts per square foot, Jaguar could not be built without using liquid cooling to remove the heat. At 4,400 square feet, the XT5 is as large as an NBA basketball court. It would be very difficult to provide evenly controlled temperature and air pressure to each of the 200 cabinets using traditional under-floor, forced-air cooling. Jaguar solves this problem by using Cray’s new ECOphlex cooling technology. This technology uses R-134a refrigerant, the same as in automobile air conditioners, to remove the heat as the air enters and exits each cabinet. The result is that the OLCF is saving 900 kilowatts of electricity that would be required just to power the fans in a traditional forced-air cooling system. Further savings are realized from the 480-volt power supplies in each cabinet.

CODE INDEX

Domain	Project	Applications
Biology	**Cellulosic Ethanol: A Simulation Model of Lignocellulosic Biomass Deconstruction** PI: Jeremy Smith, UT/ORNL	Gromacs LAMMPS
Chemistry	**Quantum Monte Carlo Calculation of the Energetics, Thermodynamics, and Structure of Water** PI: David Ceperly, UIUC	QMCPACK
	The Free Energy of Transfer of Hydronium from Bulk to Interface: A Comprehensive First-Principles Study PI: Christopher Mundy, PNNL	CP2K CPMD
Climate	**Explorations of High-Impact Weather Events in an Ultra-High-Resolution Configuration of the NASA GEOS Cubed-Sphere Global Climate Model** PI: Max Suarez, NASA/Goddard Space Flight Center	GEO S-5
	Petascale CHiMES—Coupled High-Resolution Modeling of the Earth System PI: Venkatramani Balaji, NOAA	CM
Combustion	**Direct Numerical Simulation of Lifted Diesel Jet Flame Stabilization** PI: Jacqueline Chen, SNL	S3D
Fusion	**Steady-State Gyrokinetic Transport Code (TYGRO)** PI: Jeff Candy, GA	TYGRO GYRO
	Global Gyrokinetic Turbulence Simulations of Toroidal Momentum and Electron Thermal Transport and Comparison with the National Spherical Torus Experiments PI: Weixing Wang, PPPL	GTS
	Gyrokinetic Particle Simulation of Transport Barrier Dynamics in Fusion Plasmas PI: Zhihong Lin, UCI	GTC XGC
Geoscience	**Modeling Reactive Flows in Porous Media** PI: Peter Lichtner, LANL	PFLOTRAN

Domain	Project	Applications
Materials Science	**Simulations of Disorder and Inhomogeneity Effects in High-Temperature Superconductors** PI: Thomas Schulthess, Swiss Federal Institute of Technology in Zurich/Swiss National Supercomputing Centre	DCA++ LSMS-WL
	Charge Patching Method for Electronic Structures and Charge Transports of Organic and Organic-Inorganic Mixed Nanostructures PI: Lin-Wang Wang, LBNL	LS3DF
	Petascale Simulation of Strongly Correlated Electron Systems Using the Multi-scale Many-Body Formalism PI: Mark Jarrell, Louisiana State University	MSMB
	Accurate Quantum Understanding of Adsorption on Surfaces: Water on Graphene (AQUA) PI: Dario Alfe, University College London	CASINO
	Full-Band, Atomistic Simulation and Design of Next Generation Nanotransistors PI: Gerhard Klimeck, Purdue University	NEMO
Nuclear Energy	**Denovo, a Scalable HPC Transport Code for Multiscale Nuclear Energy Applications** PI: Tom Evans, ORNL	Denovo
	Scalable Simulation of Neutron Transport in Fast Reactor Cores PI: Dinesh Kaushik, ANL	UNIC
Nuclear Physics	**Lattice QCD** PI: Paul Mackenzie, Fermi National Accelerator Laboratory	Chroma
	Ab Initio Structure of Carbon-14 PI: David Dean, ORNL	MFDn
Turbulence	**Complex Transport Phenomena in Turbulent Flows** PI: P.K. Yeung, Georgia Institute of Technology	PSDNS

ACRONYMS

2-D = two dimensional

3-D = three dimensional

ACM = Association for Computing Machinery

ANL = Argonne National Laboratory

ASCR = Advanced Scientific Computing Research program (DOE)

CHiMES = Coupled High-Resolution Modeling of the Earth System

DNS = direct numerical simulation

DOE = Department of Energy

ESM = Earth System model

ETG = electron temperature gradient

FET = field-effect transistor

GA = General Atomics

Georgia Tech = Georgia Institute of Technology

GEOS = Goddard Earth Observing System Model

GFDL = Geophysical Fluid Dynamics Laboratory

GTC = gyrokinetic toroidal code

GTS = Gyrokinetic Tokamak Simulation

I/O = input/output

IPCC = Intergovernmental Panel on Climate Change

LAN = local area network

LANL = Lawrence Berkeley National Laboratory

LBNL = Lawrence Berkeley National Laboratory

LLNL = Lawrence Livermore National Laboratory

MPI = message passing interface

MSMB = multi-scale many-body

NASA = National Aeronautics and Space Administration

NERSC = National Energy Research Scientific Computing Center

NOAA = National Oceanic and Atmospheric Administration

NSAC = Nuclear Science Advisory Committee

NSTX = National Spherical Torus Experiment

OLCF = Oak Ridge Leadership Computing Facility

ORNL = Oak Ridge National Laboratory

PETSc = Portable, Extensible Toolkit for Scientific Computation

PIC = particle in cell

PNNL = Pacific Northwest National Laboratory

PPPL = Princeton Plasma Physics Laboratory

QCD = quantum chromodynamics

QMC = quantum Monte Carlo

SCALE = Standardized Computer Analyses for Licensing Evaluation

SciDAC = Scientific Discovery through Advanced Computing (DOE program)

SION = Scalable I/O Network

SNL = Sandia National Laboratories

SRC = Semiconductor Research Corporation

UCD = University of California–Davis

UCI = University of California–Irvine

UCLA = University of California–Los Angeles

UIUC = University of Illinois–Urbana-Champaign

UNEDF = Universal Nuclear Energy Density Functional (SciDAC project)

UT = University of Tennessee

ZPR = Zero Power Reactor Experiement

Science at the Petascale: Pioneering Applications Point the Way

Production Team

Creative Director: Jayson Hines

Writers: Agatha Bardoel, Gregory Scott Jones, Dawn Levy, Caitlin Rockett, Wes Wade, and Leo Williams

Editors: Priscilla Henson and Leo Williams

Graphic Design and Layout: Jason B. Smith

Graphic Support: Tina Curry, LeJean Hardin, and Jane Parrott

Photography: Jason Richards

Contact

Doug Kothe
Director of Science
Oak Ridge Leadership Computing Facility

Oak Ridge National Laboratory
P.O. Box 2008, Oak Ridge, TN 37831-6008
PHONE 865-241-6536
FAX 865-241-2850
EMAIL help@nccs.gov
URL www.nccs.gov

The research and activities described in this report were performed using resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract no. DE-AC0500OR22725.