National Center for Computational Sciences Snapshot
The Week of May 11, 2009

The Dynamic Duo

Team simulates core/edge interaction in fusion plasma

Recent fusion simulations on Oak Ridge National Laboratory's (ORNL's) Jaguar supercomputer have verified what has long been speculated: the temperature and turbulence at the edge of a fusion plasma affect the temperature and turbulence of the plasma core.

A plasma is an ionized gas that serves as the fuel in a fusion reactor, so this finding has enormous implications in the quest for fusion energy—the process by which stars power themselves and, theoretically, one capable of providing mankind with a clean, virtually unlimited energy source.

In a fusion reactor, the plasma is confined with a magnetic field that helps to heat the plasma, a key metric in the success of a fusion reaction. Unfortunately, plasmas are susceptible to all sorts of temperature and density fluctuations, known as turbulence. These fluctuations, if too strong, can easily degrade the plasma confinement, thus weakening the fusion reaction and with it the hopes of economically feasible energy production.

Further underlying this dynamic is the relationship between the edge and core of the plasma. In the most popular design for a future fusion reactor, plasmas are housed inside a donut-shaped device known as a tokamak. The core is the center of the plasma, while what is known as the edge is generally the area in proximity to the tokamak's outer wall. Understanding the dynamics between the temperature and turbulence in the edge and their influence on the core will be key in eventually generating commercially viable fusion power.

In the latest simulations, researchers used the XGC1 code on Jaguar, the fastest system in the world for open science, to verify that turbulence in the well-confined edge can penetrate the core and boost its temperature, a phenomenon long postulated experimentally but now verified via simulation. In a well-confined edge—achieved spontaneously by the plasma upon adequate heating—the plasma temperature in contact with the wall is cool but rises sharply within a thin layer in a pedestal form. Given that a major problem in fusion is maintaining the core's temperature (ten times hotter than the surface of the sun) while keeping the edge plasma in contact with the wall cool, the simulation delivered good news. Generally the hotter the core, the better, and the fact that edge turbulence carries the high-temperature property to the core is a plus for maintaining the fusion reaction.

Researchers now have more accurate profiles for both the core and edge pedestal, strengthening their ability to optimize a plasma for energy production. A better understanding of the profile of both the edge and core is necessary if the upcoming prototype fusion reactor known as ITER is to function optimally.

These self-consistent (multiple-phenomena), multiscale simulations represent the first time the XGC1 particle-in-cell code, or one that tracks individual particles in a given state, has been used to simultaneously simulate both the edge and core. The data have yielded a treasure trove of information on not only the edge and core individually, but also the complex relationship that exists in their interaction.

The team, led by New York University's C.S. Chang, consumed more than 1 million CPU hours and gathered more than 1 terabyte of data. These most recent simulations used 20,000 of Jaguar's more than 30,000 cores. In the future, said Chang, the goal is to simulate the entire ITER device using all Jaguar's cores.

This will be a major achievement in ITER's success—to visualize the workings of the whole device, revealing the numerous relationships at play in the complexities of a working fusion reaction. "We expect to find what pedestal temperature is necessary for the successful operation of ITER," said Chang, alluding to the important role of the edge temperature of the plasma in achieving high core temperature.

Because of the complexity of these and future fusion simulations, only the world's premier supercomputing systems are capable of handling them in a reasonable timeframe. Chang is enthusiastic not only about the relationship between the edge and core, but also about that between his particle code, XGC1, and Jaguar.

"The purpose of the code and the purpose of the machine fit perfectly together," said Chang, adding that for these types of simulations, "Jaguar is number one."

From Photosynthesis to Fuel: The Next Generation of Ethanol
Cellulose holds great potential as source of biofuel energy

The push for alternative energy sources worldwide is leading to more advanced research in biofuels. Searching for new materials from which to produce such fuels is keeping researchers at ORNL and the University of Tennessee (UT) busy.

A team led by Jeremy Smith, director of the ORNL Center for Molecular Biophysics and the UT-ORNL Governor's Chair, will use the ORNL Jaguar and the UT Kraken supercomputers to run simulations that will help reveal the detailed workings of cellulose, a potential biofuel material. Cellulose is a complex carbohydrate that forms the cell walls of plants and gives leaves, stalks, stems, and trunks their rigidity. Figuring out how to unlock its sugar subunits, which can be fermented to produce ethanol, is a grand challenge of engineering. Tackling that challenge could enable full use of plants for cellulosic ethanol.

"The simulations we are performing are designed to provide a picture of biomass that will help experimentalists design plants with new, less resistant cell walls and enzymes that break the cellulose down more efficiently," Smith said. "This is basic research designed to help underpin the current major, worldwide effort in renewable energy research."

Cellulose strongly resists being broken down into glucose. Researchers call this natural resistance to decomposition “biomass recalcitrance,” and it is the major problem Smith and his team are addressing in their INCITE (for Innovative and Novel Computational Impact on Theory and Experiment) project. Hydrolysis, a reaction that uses water to break down chemical bonds, must take place for cellulose to become glucose. In lignocellulosic biomass, noncellulosic components slow the hydrolysis of cellulose and must be removed before hydrolysis can effectively occur.

Smith plans to create models of lignocellulosic biomass and cellulose that will show the structure, motion, and mechanics of the materials on a level never seen before. Supercomputing and lignocellulosic biomass are both relatively new areas of research, and using one to study the other is an even newer analytical approach.

The simulations will show the team what needs to happen to get trees and weeds into the gas tank. This work will lead to a greater understanding of cellulose composition and breakdown and serve as a reference for future research in biomass fuels. And if a process can be understood, it can be engineered.

“We hope to design plants that are less resistant and microbes that can overcome recalcitrance,” Smith said. “Putting this designed plant into fields would mean cheaper ethanol production, getting more per acre. It’s a big effort, worldwide, to understand ethanol processes, and we will likely see these efforts take effect in the next 5 years or so.”

Workshop and User Meeting Offer New Insights

Users introduced to Cray XT5 system

The U.S. Department of Energy–funded Oak Ridge Leadership Computing Facility (OLCF) and National Science Foundation–funded National Institute for Computational Sciences held the 2009 Cray XT5 workshop, “Climbing to Petaflop on Cray XT” at ORNL April 13–16.

“We were pleased here at the OLCF at the National Center for Computational Sciences (NCCS) to have such a good turnout of key users on the Jaguar leadership computing platform,” said Doug Kothe, NCCS director of science. “We presented and discussed the center’s status and plans and got very candid, constructive feedback on how we can be even more effective in supporting and enabling users’ breakthrough science goals.”

ORNL staff introduced the Cray XT5 system to principal investigators and their research teams. During this 4-day workshop, users participated in hands-on sessions with ORNL staff to become familiar with the supercomputer’s new features. The staff also led seminars and made presentations covering the architecture, issues, and effective programming of the Cray system.

“There’s been a lot of time to try things this week. It was good to get the face time with your liaison and to have a chance to really work with them. I think it will help going forward to have had those real personal face-to-face contacts,” said Laura Frink, research scientist at Sandia National Laboratories. “The first important part was to become fully up to speed on what the state of the art really is. The very detailed seminars about the machine architecture and compilers were not the flashiest to look at all the time, but it's important to get at the heart of what the state of the art is and to learn how to aim the highest you possibly can with the science that you do."

The workshop was followed by the OLCF Users' Meeting on Friday, April 16. Principal investigators and their research teams were able to discuss with ORNL staff the challenges and solutions of applications on the XT system.

Users, many of them attending their first users' meeting, had the opportunity to discuss their projects and find out what kinds of issues other users were encountering in their research.

"It's important once in a while to get out of the office and get some new ideas. There are some developments in computing and research that we are not aware of, and it's nice to know about some options," Wolfgang Balzer, research assistant at the University of Arizona, said. "We know that this is a very well-maintained site. They have an expert team here that is paid to run this computer. As a user, I just sit in my office and run my project. I really don't see what is going on behind the curtain and how much work is involved. And whenever I have a problem, they solve it very quickly."

"Engaging with computational and computer science users on a personal, face-to-face level afforded by our users' meeting is the best way to ensure that their needs, both now and in the future, are met," Kothe said. "It also helps to forge long-lasting research collaborations between center staff and the science teams with allocations on our systems."

In addition to one-on-one interactions, the research teams presented their computational work, and ORNL staff gave presentations on the range of resources available to users at the OLCF.

"There's no other place, as far as I can tell, where this particular problem in modeling biological membranes can be run," Frink said. "There are other resources that are big, but they typically are dedicated to other applications, and not as open as Jaguar. So the fact that OLCF has this mandate to do open science is why we're here."

ORNL hosts global audience at HUF '09

Researchers gather to discuss data storage needs

An international audience shared problems and solutions to data storage challenges at HUF 2009, the annual High Performance Storage System (HPSS) users' forum held March 11–13 in the Joint Institute for Computational Sciences auditorium at ORNL. Organizers expected 35 attendees, but more than 70 participants from national laboratories, universities, government, and private industry—from the U.S. and abroad—came to the meeting, which was hosted by the OLCF and the National Institute for Computational Sciences.

Stanley White, site administrator of HPSS operations at ORNL and Arthur (Buddy) Bland, director of the Leadership Computing Facility, gave the introductory talks. There were nine site presentations, in which managers described how they are configuring their storage systems to meet the needs of data generation. In a further 25 talks over 2 days, participants described the status of their individual HPSS systems as well as their plans and wish lists for new or enhanced features. Vendors that made presentations included DataDirect Networks, SUN, Spectra Logic, and IBM.

The meeting took place against the rapidly changing background of petascale computing and data generation, which represent a huge challenge for those who manage and write the code for the HPSS. The storage system is a complex piece of software at the heart of an intricate system of data transport, storage, and retrieval. Running on dozens of Dell servers, it works with an IBM-owned database software called DB2, which contains the metadata for the files and the instructions for what HPSS needs to do with them. DB2 tells HPSS to move the data coming in from runs on Jaguar and Kraken to a series of disk movers. From there, the data are migrated to tape for long-term storage. Requests for data will retrieve a file from tape and stage it to disk, to be passed to the requesting user.

“The people who are using the HPSS are here because they had extreme archival needs,” White explained in an interview. “We’re starting out at the baseline with people who have the top needs in the world for storage.” Speed is a key issue for the researchers. “Everything that is done on each separate computer is going on at the same time that I’m trying to ingest that information or retrieve information to give back to the user. So whatever is being done on the individual computers, in theory, I need to do twice as much,” White said. “I need to be able to communicate, store, retrieve, and migrate files for that system at the same time as other systems and at the same time as other groups. If you have a 200-petabyte file to store in HPSS, how long does it take to store that to a tape drive at 120 megabytes a second? You just can’t get it to tape that fast. That’s a big problem for the tape technology. ”

Another key issue for the system managers is the huge volume of coincident activity—from Jaguar, Kraken, Atmospheric Radiation Monitoring, the Earth Science Grid, and the Open Institutional Cluster (ORIC), which all have access. “I may be talking to all of these people at the same time. I’ve got tape drives that are slower than any file system I’ve got, and a lot slower than the file systems on the Cray, and I’ve got to get data to tape for archival. It’s an impossible challenge. You do the best you can with what you’ve got, and it’s never going to be good enough on the tape side. It won’t even come close; the technology just doesn’t exist to solve the long-term problem of massive storage because what we generate today will be so great.”

White along with Jason Hill and the OLCF user assistance group are trying to recover storage space by implementing new retention policies. They will give researchers some window of time to decide where to move their files. Then files will be purged from tape. Also, HPSS is a scalable system—it grows as tapes and tape drives are added. There are now more than 25,000 tapes for storage on up to 74 tape drives inside four silos. Current storage is 12.4 million files, or 4.6 petabytes. Over the past 10 years, the storage capacity has doubled each year, and there is potential for more than doubling annually, White said.