

National Center for Computational Sciences Snapshot

The Week of February 16, 2009

Supercomputer Illuminates Supernovas in 3D

Petascale runs go where no supernova simulation has gone before

A team of astrophysicists led by Oak Ridge National Laboratory's (ORNL's) Anthony Mezzacappa has launched a new era in our understanding of the universe's primary element factories.

Using the world's most powerful open scientific supercomputer—ORNL's Cray XT Jaguar system—Mezzacappa and his colleagues are the first team ever to run realistic simulations of the exploding stars known as core-collapse supernovas in three dimensions. The simulations, which include nearly all the factors likely to be important to the explosion, will provide the feedstock for a new era of scientific understanding and could validate earlier discoveries made by the team. The results should also give researchers new insights into related consequences of the supernova explosion, including the creation of new elements and the kick with which the explosion sends the remnant neutron star on its way.

The simulations will also be the first to include various types of neutrinos—tiny particles that are nearly undetectable on earth but may play a major role in blowing massive stars into space.

“Given that the neutrino transport in particular is treated in multiple frequencies and given all of the physics included, the level of realism in the model is significant,” Mezzacappa explained. “This will be the first three-dimensional simulation that could be considered realistic.”

The increased complexity and realism of the simulations were made possible by recent upgrades to Jaguar. At a peak speed of 1.64 quadrillion calculations each second, Jaguar is the most powerful supercomputer ever built for open scientific calculations. Because of this power, the team is able to conduct simulations in three dimensions without oversimplifying the supernova process.

“The ability to go to three dimensions, I cannot overstate the importance of that,” Mezzacappa noted. “This is the first model of its kind. Even as a single model, this will advance supernova theory in a significant way for all involved in supernova theory.”

The death of a massive star is built on the details of its evolution. Over millions of years, the star burns through its nuclear fuel, fusing atoms to create increasingly heavy elements.

These elements form layers, from hydrogen at the surface through helium, carbon, oxygen, silicon, and iron at the core. The process hits a roadblock at iron, which does not create energy when atoms are fused. Eventually, the iron core becomes so heavy that it collapses under its own weight.

When the core reaches the limit of its collapse into a mass of protons and neutrons, it bounces on itself, creating a shockwave that eventually blows the star into space.

Recreating this process computationally is like reverse-engineering your favorite entrée; you know the dish but you have no recipe. The team's simulations include nearly every factor likely to be important in a core-collapse supernova. Fluid-dynamics calculations illuminate phenomena such as convection currents and the evolution of the shockwave. Nuclear-physics equations of state show how atomic nuclei respond to changes in pressure and temperature. The simulation handles gravity in three dimensions, with corrections conforming to Einstein's theory of relativity. And state-of-the-art physics equations handle the production of neutrinos at various energies.

“To use a cooking analogy, you're adding key ingredients in a recipe,” Mezzacappa explained. “If a recipe works, you have a soufflé. But if you go in and remove a key ingredient, either you have no soufflé or you have something unfit to eat.

“What we're doing here is finding the recipe for a core-collapse supernova. You add the ingredients—you add the physics—and you then carry out the simulations and see what the outcomes are. At the end of the day, if your model makes predictions that match observables, then your model is good.”

In fact, the only major ingredient not included in the model is a star's magnetic field, a gap the team plans to fill in the coming year.

The team recreates the supernova using a three-part software application known as Chimera, named after the three-sectioned monster of Greek mythology. For this application, the three components are MVH3/VH1, an astrophysical hydrodynamics code that computes the motion of material within the star; MGFLDTRANS, a neutrino radiation transport code that follows the evolution of neutrinos and their interaction with material in the star; and XNET, a nuclear kinetics code that traces the different elements created in the explosion and computes nuclear reactions that take place as the explosion evolves.

Three-month simulations

Core-collapse supernovas start as stars 10 to 25 times the mass of the sun, and Mezzacappa's team is beginning the project by simulating a star 15 times the mass of the sun. Each simulation takes about three months, with the outcome being an immensely detailed look at just under a second at the beginning of a supernova.

“Three-quarters of a second is not a magic number,” Mezzacappa explained. “You have to run the simulation sufficiently long so you initiate the explosion and you can determine characteristics of the explosion such as the explosion energy. At about three-quarters of a second, things are fairly well determined; if not, we would run longer.”

The simulations should help verify discoveries already made by the team. One is known as the standing accretion shock instability (SASI), in which the shockwave stalls early on in the process before eventually reviving. Another is the most plausible explanation to date for the ultimate spin of the leftover neutron star—that it is created by the shockwave as it spins in place before being revived.

“One has to run three-dimensional, multiphysics models in the end to really make definitive predictions for kicks and spins of neutron stars,” Mezzacappa noted, “and that’s what we’re doing. So one of the outcomes of this ongoing model and other 3D models beginning with different progenitors will be an analysis and determination of whether or not the SASI-induced spin continues to describe well-observed spins of young pulsars.”

Beyond this line of research, he said, the team will move on to the collapse of even larger stars, known as hypernovas. Whereas the stars being simulated by the team now yield ultramassive neutron stars after they explode, stars more than 30 times the mass of the sun are associated with even more bizarre leftovers, namely black holes and gamma ray bursts.

“Our intent is to continue into a higher mass range down the road. We won’t stop at 25 solar masses. We intend to look at higher masses and to look at these other scenarios and to look at these other subclasses of core-collapse supernova explosions. But right now, we’re going after the standard canonical ones.”

Supercomputing Seeks Energy Savings

ORNL facility takes all-angles approach

As high-performance computing (HPC) enters the petascale age, the scientific challenges facing researchers have never been greater. Nor has the might of today’s production petascale machines.

The recent exponential growth in the power of modern supercomputers has gone hand-in-hand with an increased demand on resources—as machines have gotten bigger and faster, the amount of resources required for their operation has likewise increased.

As a result, HPC centers now face unprecedented power demands from the very machines they rely on to tackle today’s most daunting scientific challenges, from climate change to the modeling of biological processes. However, recent energy-saving innovations at ORNL are setting a new standard for resource-responsible HPC research. The laboratory has taken an all-angles approach, seeking energy savings from a suite of different areas.

ORNL’s flagship system, a Cray XT known as Jaguar, is now the fastest computer in the world for open science. With this great power comes great responsibility, especially when it comes to energy consumption. “We take energy utilization very seriously,” said ORNL’s National Center for Computational Sciences (NCCS) Project Director Buddy Bland.

Jaguar is capable of a maximum speed of 1.6 petaflops, or approximately 1.6 quadrillion calculations a second. “The scale of this machine is just phenomenal. There are very few places in the world where this computer could have been built,” said Bland.

Needless to say, feeding this animal is no small task: simulation at the petascale requires robust power and cooling networks to ensure maximum production from these machines. But now those necessary support networks, and the system itself, have been designed with unprecedented efficiency, responsibly satisfying Jaguar’s energy appetite. These advances make ORNL among the most energy-efficient locations for HPC, enabling groundbreaking research with minimal resource impact.

It all starts with the building. ORNL’s Computational Sciences Building (CSB) was among the first Leadership in Energy and Environmental Design (LEED)–certified computing facilities in the country, meaning that its design satisfies criteria used by the U.S. Green Building Council to measure the efficiency and sustainability of a building.

Take the computer room for example: it’s sealed off from the rest of the building by a vapor barrier to reduce the infiltration of humidity. The air pressure inside the computer room is slightly higher than the surrounding area so air will flow out of the computer room without the air outside flowing in. Because ORNL is located in an area of the country with high humidity, keeping moisture out of the air is a high priority said Bland, one that the building was designed to tackle as efficiently as possible. Too much moisture in the air can lead to water condensation on equipment, while too little moisture can cause static electricity to build up—both of which can be problematic for a room filled with expensive electronics. Both removing moisture from or adding it to the air uses a lot of power, so keeping the humidity stable is a great tool for reducing energy consumption.

Another computing building on the ORNL campus adjacent to the CSB was recently certified LEED Gold, the second-highest ranking, and Bland points out that the laboratory plans on an equal rating for any future HPC facilities. But the innovation doesn’t stop with the building—there is plenty more under the roof.

Jaguar requires huge amounts of chilled water to keep the machine cool. To accomplish this as efficiently as possible, the laboratory uses high-efficiency chillers, which are the first step in a multifaceted, efficient cooling design.

A newly introduced Cray cooling system for Jaguar, dubbed ECOphlex, complements the chillers and the CSB’s efficiency. Using a common refrigerant and a series of heat exchanges, ECOphlex efficiently removes the heat generated by Jaguar to keep the computer room cool. The combination of air- and refrigerant-based cooling is much more efficient than traditional systems, which rely almost solely on air for temperature control. Without ECOphlex, the number of air-based units would not fit into the CSB’s computer room. This high-efficiency cooling system makes Jaguar possible.

ECOphlex also allows the NCCS to reduce the amount of chilled water used to cool Jaguar by accommodating a broader inlet temperature range for the cooling water. Considering the fact

that thousands of gallons of water per minute are necessary to keep Jaguar cool, a reduction in the volume of necessary chilled water means a proportionate reduction in the energy used to cool it. Simply put, warmer water can mean big energy savings for the NCCS and the taxpayer. Whereas most centers use 0.8 watts of power for cooling per every watt of power used for computing, ORNL's NCCS enjoys a far more efficient ratio of 0.3 to 1, one of the lowest of all data centers measured.

Another important innovation is one that ORNL has been working on with Cray for several years. Instead of using the more common 208-volt power supply that Jaguar used in the past, the system now runs directly on 480-volt power. This seemingly "minor" change is saving the laboratory \$1 million in the cost of copper used in the power cords for the cabinets. Furthermore, keeping the voltage high allows a lower current, which means lower resistance and less power turned into heat as it travels down the wires. The reduction in electrical resistance will reduce energy costs by as much as half a million dollars.

Finally, ORNL gets a little help from history. The power grid for the city of Oak Ridge was designed when the research conducted during the Manhattan Project used one-seventh of all the electricity in the country. The grid was constructed with every protection possible out of the fear that any interruption in supply would drastically set back development. The result: an extremely resilient local power grid.

Because of this grid, said Bland, Oak Ridge doesn't need huge uninterruptable power supply (UPS) systems, which generally consume lots of electricity. However, the laboratory does have flywheel-based UPSs in case of an emergency. If there is a problem, the flywheel keeps generating power, which is a much more efficient process than conventional UPSs and therefore a greener method of supplying backup power. Because the flywheel-based UPS is mechanical as opposed to battery-operated, it also generates less waste in the long-term as battery replacement is not a concern.

While all of these steps are important, taken together they are greater than the sum of their parts. "There is no silver bullet," said Bland. By tackling energy efficiency from multiple angles, ORNL is helping to ensure that the groundbreaking research taking place on its petascale machines is conducted as responsibly as possible, setting new standards in both HPC and energy responsibility.

ORNL Hosts Lustre Workshop

Community gathers to address issues

The NCCS, Sun Microsystems, and Cray, Inc., recently held a two-day workshop on Lustre scalability at ORNL.

Participants from all the major Lustre sites, including Sandia National Laboratories, Lawrence Livermore National Laboratory, NASA, and Pacific Northwest National Laboratory, met to identify key scalability issues and develop a realistic roadmap for Lustre by 2012. Some of the issues addressed were the changes that need to be made to the architecture of the petaflop

file system and where Sun Microsystems, who developed Lustre, will need to allocate resources to meet some of those scalability challenges.

“The importance of the workshop is that all the simulation codes that we’re looking at at the petascale and beyond require high levels of I/O throughput,” said Galen Shipman, technology integration group leader at ORNL’s NCCS. “We’re checkpointing and writing up information from each processor so we can restart our jobs without having to roll back to the beginning. We’re gathering the entire Lustre community to agree upon the primary goals and the gaps that we have now in a parallel file system environment and how best to address those gaps.”

The workshop, held February 10–11, set the stage for a follow-up Lustre Scalability Summit to be held in April, which will focus on scalability challenges through 2014. The summit will be held in conjunction with the 2009 Lustre User Group meeting.

2009 HPC Workshops and Training Sessions Announced

NCCS offers a variety of training opportunities to users

The NCCS will conduct several workshops and teaching sessions in 2009 to help current and future users of HPC become familiar with the new petascale systems.

- NCCS and the National Institute for Computational Sciences (NICS) will jointly sponsor a four-day 2009 Cray XT5 Quad-Core Workshop entitled “Climbing to Petaflop on Cray XT” April 13–16, 2009, at ORNL. The workshop will cover the important issues in obtaining increased performance from the powerful new XT systems. Among the topics to be covered will be XT5 architecture, XT5 Nonuniform Access Memory Access issues, and programming effectively for the XT5.

The workshop will include lectures from NCCS, NICS, and Cray staff as well as hands-on sessions. Past workshops have been extremely successful and you are encouraged to attend to get the most out of the most powerful computers in the world. The registration site, with complete agenda, can be found at:
<http://www.nccs.gov/user-support/training-education/workshops/2008-cray-xt5-quad-core-workshop/>.

- The 2009 NCCS Users’ Meeting will be held at ORNL from 8:00 a.m. till 4:30 p.m. on Thursday, April 16. For more information, please visit <http://www.nccs.gov/user-support/training-education/workshops>.
- The NCCS and NICS will jointly host the annual High Performance Storage System (HPSS) User Forum, HUF09, March 11–13 in the Joint Institute for Computational Sciences auditorium. HPSS users are encouraged to share successes, lessons learned, and solutions to site issues. Developers and support representatives will talk about the latest HPSS directions and new features.
- Undergraduate, graduate, postdoctoral students, and faculty from southeastern universities are invited to attend HPCA Con ’09 April 3–4 at the Pollard Center in

downtown Oak Ridge. Computer science, computational science, engineering, and computational engineering students are encouraged to participate. The conference will give students an overview of parallel concepts and architecture, message passing interface (MPI), makefiles and batch scripts, debugging, and coding examples.

- HPCA Con '09 will begin with a workshop on April 3 on parallel programming models to give faculty and students an introduction to such models as MPI and OpenMP, two of the more popular paradigms seen in applications at the NCCS. The session will provide educators with material for their curricula and introduce students to parallel programming basics.
- The NCCS and ORNL will host CUG 2009, a meeting of the Cray Users Group, May 4–7 at the Omni Hotel in Atlanta. “Compute the Future” is the theme of the meeting, which is a major supercomputing event. Participants will learn how the Cray XT and future powerful systems will be used to answer the critical scientific questions of our age.
- The NCCS and NICS will sponsor a workshop called “Introduction to HPC on Cray XT” in May at Oak Ridge Associated Universities (ORAU) in downtown Oak Ridge. The date is to be determined.

For registration information on the workshops, go to <http://www.nccs.gov/user-support/training-education/workshops/>. For CUG 2009 meeting information, go to <http://www.cug.org/1-conferences/CUG2009>.