National Center for Computational Sciences Snapshot
The Week of February 18, 2008

NSF Grant Will Aid Supernova Research

Team will get better at blowing up stars

A team led by National Center for Computational Sciences (NCCS) astrophysicist Bronson Messer will use a National Science Foundation (NSF) grant to prepare its supernova simulation tool for the era of petascale computing.

The team, a collaboration among researchers from the University of Tennessee (UT), North Carolina State University, and Florida Atlantic University, has received an $800,000 NSF grant given through the agency's PetaApps program. The researchers will use the funding to upgrade the CHIMERA application for next-generation supercomputers that are orders-of-magnitude faster than today's most powerful systems. In particular, the program seeks important applications to take advantage of petascale supercomputers—machines capable of performing more than 1,000 trillion calculations each second (1 petaflop).

According to NSF materials, "Proposals are sought from researchers aiming to capitalize on emerging petascale computing architectures, catalyzing progress in science and engineering beyond the current state of the art."

"The scientific problems need to be there," Messer explained of the program, "but just as important, the codes need to be able to exploit effectively the platforms that NSF was going to field, and so they realized they needed to put some funding into this."

Messer is a principal investigator on the project. His collaborators include Anthony Mezzacappa and Raphael Hix, Oak Ridge National Laboratory (ORNL) researchers who are, like Messer, adjunct faculty in UT's Department of Physics and Astronomy. Other collaborators in the project include John Blondin of North Carolina State University and Stephen Bruenn and Pedro Marronetti, both of Florida Atlantic University. The Tennessee collaborators will receive $385,000 of the $800,000 total grant, which will allow them to bring postdoctoral researchers and UT students into the project.

The researchers are using CHIMERA to work out details of the core-collapse supernova, the death of a star at least 10 times as massive as the sun. In particular, they are investigating how the shockwave created when the star's core collapses eventually blows most of the star into space. They have made important contributions to the field using terascale supercomputers (capable of trillions of calculations a second), discovering, among other things, the first explanation for a pulsar's spin that matches observations from actual pulsars. The group was able to publish that discovery in the leading science journal *Nature*.

The team is working to tailor CHIMERA to take advantage of the dramatic advances in supercomputing power expected over the next 5 to 10 years. In particular, it is working on methods to effectively use upcoming computing chips as they load more and more processing cores into a single processor.

"We have lots of local microphysics that happen in CHIMERA," Messer explained, "neutrino interactions, nuclear burning, things like this that are not parallelized now but can be. We can use one socket ,which has four cores for instance on the new [Cray XT4] Jaguar upgrade, and be able to marshal the power of all four cores. That's a lot of what the students and the postdocs are going to be doing. And they'll be using the code that they work on to produce simulations and do science."

Researchers at ORNL Contribute to HPC Community

Tools under development to maximize HPC potential

The NCCS at ORNL is developing a suite of tools aimed at benefiting the larger high-performance computing (HPC) community. These new utilities will enable researchers everywhere to more efficiently tackle today's greatest scientific challenges as HPC moves to the petascale and beyond.

NCCS researchers Richard Graham and Galen Shipman recently ported the Open MPI implementation of the Message Passing Interface (MPI) 2.0 standard to run on Cray's XT Compute Node Linux (CNL) operating system, which interfaces with Cray's Alps run-time system. The Open MPI implementation can now take advantage of the shared-memory functionality supported on CNL, greatly improving the on-host communications performance and reducing the application's network demands. Researchers are currently implementing shared-memory collective communications optimizations, which will improve application scalability because these communications are a major factor in limiting it.

Richard Graham is also chairing the reconvened MPI Forum. This group of production and research staff from national laboratories, universities, the computer industry, and independent software vendors from around the world aims to update the decade-old standard for communications and process management for parallel computing. Among the improvements being explored are better support for petascale-class computers, including support for fault tolerance, high-performance one-sided communications, and nonblocking collective communications operations.

NCCS researcher Ken Matney is likewise developing a number of tools to address users' day-to-day file management needs. These tools will leverage the parallel file systems by distributing these operations across many worker clients and will include parallel tools to copy, compress/decompress, and archive data. The first of these tools, the Staging Parallel Distributed Copy (SPDCP), allows an arbitrary tree or set of trees to be copied from one location to another, propagating Lustre striping information. Furthermore, says Matney, "it balances client resources for each copy against server resources, exploiting parallelism in input/output [I/O], at the file and data levels." In contrast, current tools are

limited in performance to a single server node of the file system. “This tool makes use of resources that are already there but aren’t being used at the time,” says Matney, adding that it essentially “enhances the productivity of the researcher.”

Matney’s other project, the Scalable Parallel Executor to Effect Distribution of Utility Programs (SPEEDUP), allows the user to fully leverage the capabilities of the parallel file system. This will provide the ability to execute an arbitrary Linux command in parallel across the Cray XT. Consequently, the time needed to run “sed,” “bzip2” compression/decompression, and various other Linux commands on large datasets can be greatly reduced. While so far the program has scaled to a limited number of processors, Matney says that “given the right data set, there is no reason why it won’t scale to 1,000 processors.” Everything is parallel (including I/O requests), with minimal synchronization requirements. “There are times when it’s best to allow disks to perform I/O,” says Matney, “but there are times when it is more efficient to use brute force to do the parallelization.”

An I/O componentization, dubbed ADIOS (ADaptable I/O System), is being developed by NCCS researcher Scott Klasky. ADIOS, says Klasky, is an easy-to-use application programming interface that can be inserted into Fortran and C/C++ codes along with external metadata (XML) files, which contain all of the extra information about the variables in a simulation. Furthermore, ADIOS has been integrated with some of the largest data-producing codes at ORNL (such as GTC, GTS, and CHIMERA). Among its advantages, says Klasky, are its ability to allow plug-ins for different I/O implementations, the fact that both synchronous and asynchronous transports are supported without code changes, its free hooks into visualization and workflow methods, and the optimized I/O implementations provided for each transport method. Klasky has also demonstrated that with some methods, the I/O rate can exceed 25 GB/second on Jaguar, the NCCS’s Cray XT4.

Finally, a collaboration between ORNL, Argonne National Laboratory, IBM, UT, Indiana University, The University of Wisconsin, and Renaissance Computing Institute is currently building a scalable tool infrastructure aimed at petascale machines and beyond. The group is close to completing the application programming interface definition and component architecture. ORNL’s main focus is the scalable startup in the first implementation phase, with Alps and ssh/rsh, a security utility, being the targeted resource “managers” for which an implementation will be developed.

As HPC advances to the petascale and beyond, researchers need the best tools available to effectively use the power of petascale machines. The tools being developed at ORNL will greatly improve the HPC efficiency, allowing researchers everywhere to more effectively tackle the country’s most pressing scientific problems.

NCCS Sponsors Multiple Workshops

Users Meeting, XT, Lustre all upcoming

This spring the NCCS will host a series of workshops aimed at assisting its users and the wider HPC community.

A total of three workshops, held at ORNL, will be hosted in mid-April, beginning with the NCCS Cray XT workshop on April 14, 15, and 16. Staff from the NCCS, ORNL's Joint Institute for Computational Sciences, Cray, and chipmaker AMD will discuss XT issues, and researchers attending the meeting will participate in hands-on sessions with the Cray XT system. Computational scientists will gather with vendors and staff experts at ORNL to discuss strategies for making the most of their time on Cray XT supercomputers.

On April 16 the NCCS will host a Lustre workshop, focused on helping application scientists get the most from the Lustre File System. The workshop will be presented by Oleg Drokin and Wang Di, two file system engineers with Sun Microsystems and both seasoned Lustre developers.

Closing out the workshop series will be the NCCS Users Meeting on April 17 and 18. Principal investigators and members of their research teams will gather with NCCS staff and vendors to discuss challenges and solutions in areas such as porting and scaling of applications on the XT system. Each project new to the Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program will be invited to give a 10-minute presentation on its upcoming work, while renewing projects will review their work at an evening poster session. NCCS staff will also give presentations highlighting the complete range of resources, capabilities, and services available to the center's users.

To register or for more information, visit www.nccs.gov .