ALCF Newsbytes

Argonne Leadership Computing Facility

Argonne National Laboratory

Volume 1, Issue 2 | July 2011

Another Year of Satisfied Users at the Argonne Leadership Computing Facility: 2010 User Satisfaction Survey Results

Each year, the Argonne Leadership Computing Facility conducts a survey to gauge user satisfaction with ALCF programs and service. The 2010 survey, comprised of 15 multiple-choice and open-ended questions, was distributed electronically to all users in late December. Nearly 250 users participated in the survey, resulting in a response rate of over twenty-nine percent. "We're very pleased with both the number of users who completed the survey and with the scores overall. Next year, we're hoping even more users will participate in the survey," said David Martin, manager of user services and outreach.

The survey polled users on project status, any performance issues or resource needs, and gave them an opportunity to share next steps for their project—including if and when there are plans to apply for an INCITE allocation or other award. Users were also asked to comment on their experiences with ALCF support staff and services over the year, and with specialized services offered in certain programs, such as the Catalyst staff aligned with INCITE projects. "The ALCF consistently has seen scores that are well above average on all our user surveys, and the 2010 survey was no exception," said Martin.

2010 ALCF User Satisfaction Survey Highlights

- 249 responses, response rate of over 29%
- 92% of users rated their overall satisfaction with the ALCF as above average or excellent.
- 90% of user were able to achieve their science goals.
- 94% of users agreed or strongly agreed that the ALCF staff provides accurate, complete assistance and/or answers to their questions.

Survey scores and comments are tabulated and mined for use in improvement initiatives, including potential user workshop topics and requests for new libraries. "User feedback like this is a primary driver of future programs and offerings at the ALCF," Martin continued.

And while every year the ALCF gets a myriad of user requests for a larger system and shorter queue times, the 2010 survey had a new one for the books. "Maybe it was the holiday timing of this particular survey," remarked Martin, "but we had one user actually ask for unlimited allocations for all with no wait time…and a pony. We're proud of our record of highly satisfied users, but out of those three requests, the pony is the most likely."

Table of Contents

Like Us On Facebook

Events of Interest

Grace Hopper Celebration of Women in Computing
November 9-12, 2011
Portland, Oregon
The world's largest gathering of women in computing in industry, academia, and government, the Grace Hopper Celebration of Women in Computing is a four-day technical conference designed to bring the research and career interests of women in computing to the forefront. For details, visit: www.gracehopper.org.

SC11
November 12-18, 2011
Seattle, Washington
The SC11 conference continues a long and successful tradition of engaging the international community in high-performance computing, networking, storage and analysis. For details, visit: http://sc11.supercomputing.org.

Argonne NATIONAL LABORATORY

Advanced Climate Modeling of the Earth Under Way at the ALCF

The earth's changing climate means an increase in the occurrence of droughts, heat waves, wildfires and floods. The Climate Science Computational End Station (CCES) advances climate science through improved model development by the Department of Energy laboratories and the National Center for Atmospheric Research. Climate simulations currently underway include the global carbon cycle and its feedback to the climate system, and its variability and modulation by ocean and land ecosystems. The research team is testing a new, highly scalable method for solving the fluid dynamics of the atmosphere for use in future climate simulations. This model, called the High Order Method Modeling Environment (HOMME), has run with a resolution as high as 1/8th of a degree of latitude on more than 80,000 computational cores. The team uses HOMME to perform standard climate model benchmark simulations for comparisons with other models. The vital information provided by these improved models will increase the understanding of present and future climates and will help guide U.S. environmental policy.

Using an INCITE allocation, researchers are also testing a new version of the Community Earth System Model (CESM), on the Blue Gene/P system at the Argonne Leadership Computing Facility. This state-of-the-art model is used to make national assessments of global climate change and to investigate other important climate problems. A configuration of the Community Atmosphere Model (CAM) component of CESM (known as CAM-HOMME) represents the best path towards scalable, high-resolution global atmosphere simulations. With CAM-HOMME, scientists achieved the first-ever integration of a climate model that included atmosphere, land, ocean and sea ice, on more than 100,000 cores.

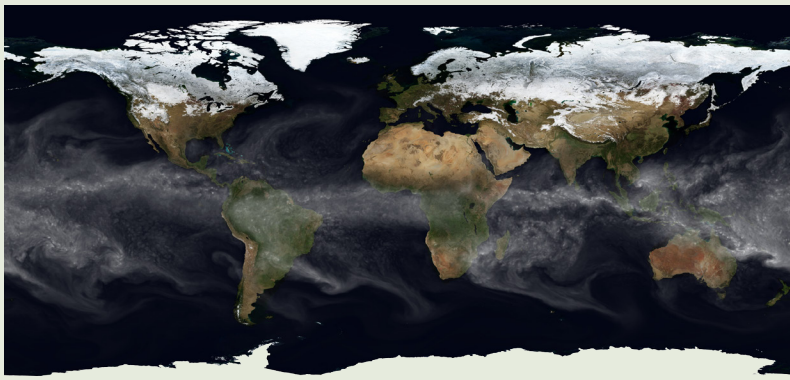

Total precipitable water, a measure of how much moisture is in the air from a single moment in time in the global simulation of the atmosphere at a resolution of half a degree of latitude. (Figure provided by Mark Taylor, Sandia National Laboratories.)

Contact: Warren Washington (PI) or Tom Bettge (wmw@ucar.edu, bettge@ucar.edu)

Predicting Nuclear Structure and Nuclear Reactions

Developing a robust and precise nuclear theory based on the underlying theory of strong interactions, quantum chromodynamics, represents a "holy grail" for physics with many practical applications. Predictions of the structure and reactions of nuclei are important for basic science and for efforts addressing the nation's energy and security needs.

A comprehensive description of all nuclei (stable and unstable) and their reactions requires investigating rare and exotic isotopes with unusual proton-to-neutron ratios that are difficult to produce and study experimentally because of their short lifetimes. Led by James Vary at Iowa State University, researchers are performing state-of-the-art simulations at Argonne and elsewhere to provide needed predictions of the structure and reactions of nuclei, where direct experiment is not possible or is subject to large uncertainties.

The researchers are deploying the best available, theoretical many-body physics tools, coupled with current theory of the strong interactions, including three-nucleon potentials. Among other goals, the researchers aim to provide an *ab initio* understanding of triple-alpha burning essential to life on Earth. So far, *ab initio* investigations of the role of three-nucleon forces have been limited to light nuclei. However, this project is exploring the role of the three-nucleon force in substantially heavier nuclei and an expanded range of nuclear reactions. Researchers are performing calculations for the entire nuclear mass table to improve understanding of the nuclear energy density functional. They are also calculating nuclear properties relevant for the description of nuclear reactions. Studies include various scattering processes in light nuclei and bulk properties for nuclei across the entire mass table.

Predicting Nuclear Structure and Nuclear Reactions (continued)

A major focus area of the project is a high-quality microscopic description of the nuclear fission process. To achieve this goal, researchers have developed a new nuclear energy density functional, "UNEDF1," that greatly improves results for fission over previous functionals. They expanded the experimental data set in the optimization of UNEDF1 to include excitation energies of superdeformed fission isomers in the actinide region. Further improvement of the nuclear energy density functional requires *ab initio* simulations of nuclear systems beyond experimental reach. For this purpose, the team performed *ab initio* simulations of "neutron drops" with 4 through 42 neutrons in external fields. Researchers will incorporate these results in the data set to further optimize the energy density functional.

Contact: James Vary (jvary@iastate.edu)

Gordon Bell Prep Support at the ALCF Yields Significant Accomplishments

Each year, Gordon Bell Prizes are awarded at the annual Supercomputing conference in recognition of outstanding achievement in high-performance computing. To assist users in the pursuit of winning this prize, the Argonne Leadership Computing Facility (ALCF) collaborated with two teams of researchers in March and April. The accomplishments are detailed below:

Project 1: Multiscale Simulation in the Domain of Patient-specific Intracranial Arterial Tree Blood Flow

PI: George Karniadakis, *Brown University*
Participants: Leopold Grinberg, *Brown University;* and Vitali Morozov, Joseph Insley, Kalyan Kumaran and Michael E. Papka, *Argonne National Laboratory*

Overview: The seamless integration of heterogeneous computer codes that implement discretizations of partial differential equations with codes that implement atomistic-level descriptions is key to the successful realization of parallel multiscale modeling of realistic physical and biological systems. Such multiscale modeling is crucial in many disciplines, e.g., to tune the properties of smart materials, to probe the function of living cells and organisms, and to predict the dynamics and interactions of rarefied plasmas with dense plasmas. In addition to advances in computational power of computer systems, fundamental new advances in algorithms are required to provide the proper mathematical interface conditions between atomistic and continuum systems that respect the physics of the problem and provide numerical accuracy, stability and efficiency. Moreover, post- or co-processing the results of multiscale simulations requires new quantitative visualization tools that can simultaneously handle particle-based and continuum-based parallel simulation data.

Goal: To perform a first-of-its-kind, multiscale simulation in the domain of patient-specific intracranial arterial tree blood flow.

To resolve the large-scale dynamics of the blood flow, the team used Nektar, a high-order spectral element code. To resolve the mesoscale features of the flow, they used a modified version of LAMMPS software called LAMMPS-DPD. These two codes are coupled into a multiscale solver called NektarG, developed at Brown University. The team used several Nektar modules coupled to a LAMMPS-DPD module and scaled the entire run to 32 racks (131,072 cores) of the IBM Blue Gene/P, Intrepid. Communications between and within the modules are performed via MPI.

Demonstration Problem: As Figure 1 shows, the computational domain consists of tens of major brain arteries and includes a relatively large aneurysm–a blood-filled, balloon-like bulge in the wall of an artery. The team used Nektar software to resolve the large-scale features in this domain. The aneurysm region was resolved using a sub-domain of about two cubic millimeters, which calculated the blood flow using the dissipative particle dynamics solver, LAMMPS-DPD. Researchers scaled the size of the sub-domain by a factor of 200 (non-dimensional units), which led to O (800M) DPD particles (equivalent of 8 billion MD particles) simulation to assure a high-resolution and physiologically correct parameter. The large number of particles involved makes this work particularly suited for the unique capability of leadership-class resources like those available at the ALCF.

Gordon Bell Prep Support at the ALCF Yields Significant Accomplishments (continued)

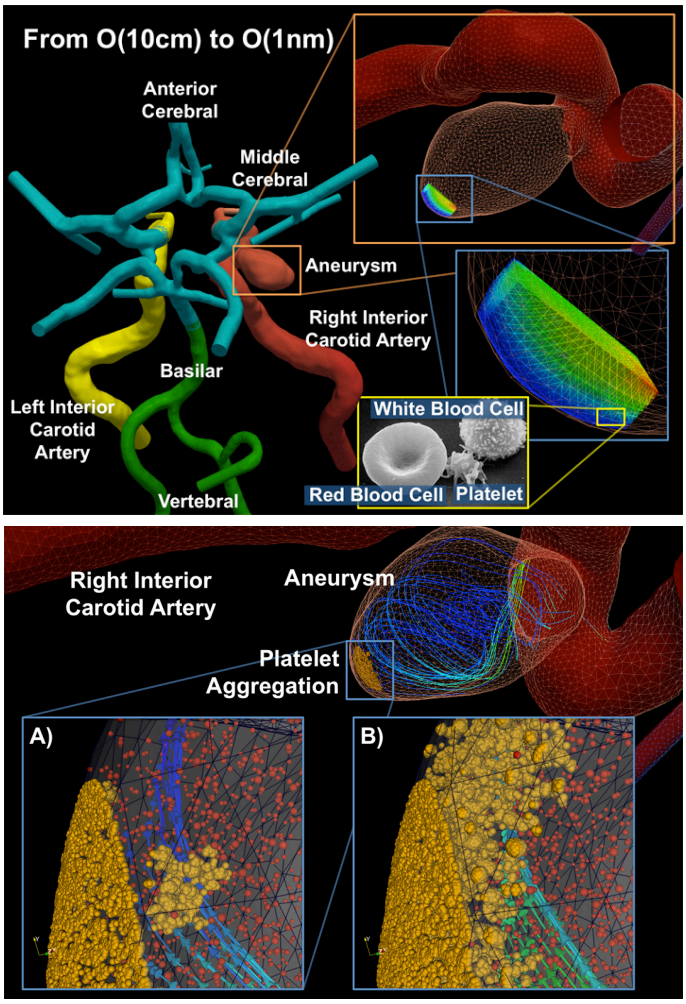

Figure 1: The top figure shows an overview of the entire effort. The leftmost piece is the macrodomain where the Navier-Stokes equations are solved; different colors correspond to different computational patches. The upper right inset shows the microdomain where dissipative particle dynamics is applied. Of interest in the present paper is the deposition of platelets to the aneurysmal wall. The bottom figure shows the overall flow through the artery and the aneurysm as calculated by Nektar, as well as that within the subdomain calculated by LAMMPS-DPD, shown in detail in insets A and B, along with platelet aggregation along the aneurysmal wall.

Accomplishments: The team's initial goal was to perform a coupled simulation for the duration of 20,000 time steps—representing approximately 15% of a cardiac cycle. Once accomplished, the team continued to work on a series of optimizations in NektarG, which resulted in a significant speed-up in function used to compute derivatives in tetrahedral elements and in writing data to disk by the DPD-LAMMPS module. In addition, they performed several optimizations in memory management. As a result, the team was able to scale NektarG to run on 32 racks of the Blue Gene/P (131,072 cores) with satisfactory performance. The team used 13.6 million core hours to accomplish the runs.

"We successfully integrated a solution of over 132,000 steps in a single, non-stop run on 32 compute racks of Blue Gene/P," said Leopold Grinberg. "This clearly shows the stability of our code; moreover, writing frequently about 20GB of data on disk did not adversely impact our simulation."

This simulation is the first of its kind to include the largest 3D cerebrovasculature (essentially all the arteries that can be reconstructed from an MRI), physiologically accurate blood cells, and at the same time model microheological mechanisms inside the aneurism.

Future work: Future work will focus on tuning parameters related to the DPD simulations. Due to substantial secondary flows, researchers may reduce the size of the time step and rescale the parameters associated with the sub-domain where the solution is computed with DPD. As a result of current successes, the team will endeavor next to simulate an entire cardiac cycle simulation—representing about 300,000 time steps in the setup—and will incorporate an investigation via visualization.

Project 2: Co-visualization of Full Data and *in situ* Data Extracts from Unstructured Grid CFD at 160 Thousand Cores

PI: Ken Jansen, *University of Colorado at Boulder*

Participants: Ken Jansen and Michel Rasquin, *University of Colorado at Boulder;* Pat Marion, *Kitware Inc.;* Venkat Vishwanath, Mark Hereld, Kalyan Kumaran, Michael Papka and Ray Loy, *Argonne National Laboratory*

Overview: PHASTA is a parallel, hierarchic (2nd-5th order accurate), adaptive, stabilized (finite element) transient, incompressible and compressible flow solver. Many of its application cases have been sufficiently complex that grid-independent results could only be obtained through the efficient use of anisotropically adapted unstructured grids or meshes capable of maintaining high-quality boundary layer elements, and scalable performance on massively parallel computers. PHASTA has been shown to scale to 288 thousand cores.

GLEAN is a MCS/ALCF-developed tool providing a flexible and extensible framework for simulation-time data analysis and I/O acceleration. GLEAN moves data out of the simulation application to dedicated staging nodes with as little overhead as possible. Previous work by the team had connected PHASTA to GLEAN. However, those runs were done at modest scale with all visualization filters executed on the staging nodes.

Gordon Bell Prep Support at the ALCF Yields Significant Accomplishments (continued)

Goal: The goal of the Gordon Bell effort was to integrate an improved version of GLEAN and to collect data at large scale for PHASTA+GLEAN for three real-time visualization scenarios to determine frame rate and solver impact. In the first two, visualization filters were run directly on the Intrepid compute nodes and the polygon data was transferred to Eureka; this was accomplished alternatively by sockets in VTK, and by GLEAN. The third method did not perform visualization filtering on Intrepid, but rather forwarded the raw dataset to Eureka via GLEAN.

Demonstration Problem: The demonstration problem simulates flow control over a full 3D swept wing. Synthetic jets on the wing pulse at 1750Hz, producing unsteady cross flow that can increase or decrease the lift, or even reattach a separated flow. Runs used meshes of size 22M, 416M, and 3.3B elements. Approximately 6.1 million core hours were required for the data collection.

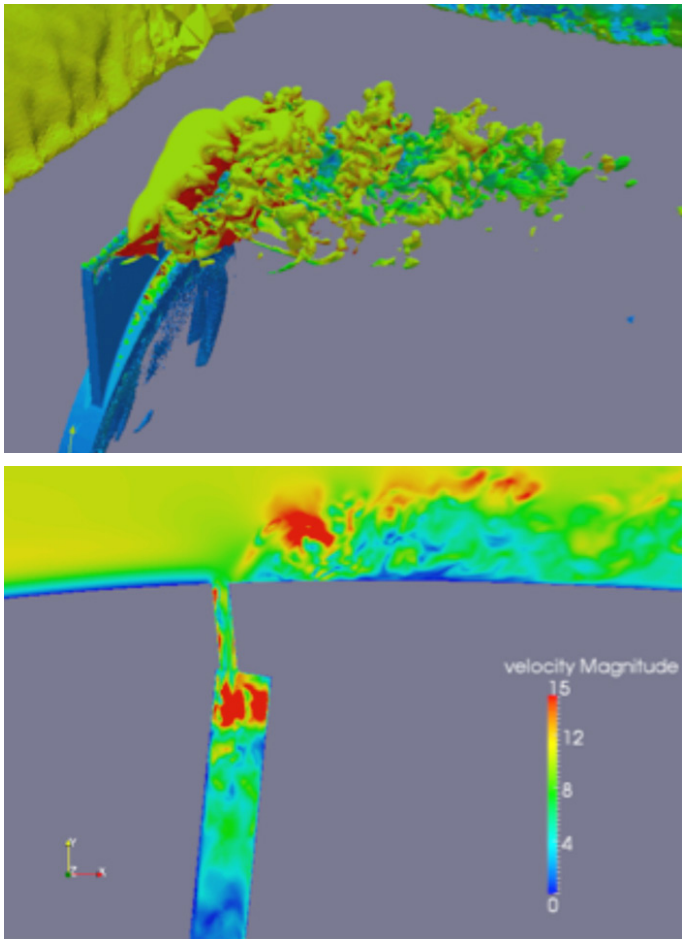

Figure 2: On the top is an isosurface of vertical velocity colored by magnitude of velocity and on the bottom is a cut plane through the synthetic jet (both on a 3.3 billion-element mesh). These are single frames taken from the real-time rendering of a live simulation.

Accomplishments: When the newly assembled code stack was first run at scale, it was clear that disk I/O during the program startup was not only slow but also having a major impact on the entire Intrepid system. This was traced to the initialization of Python, which unexpectedly caused 10-20 extension modules to be loaded on each rank. To overcome the issue, a tool called Python Freeze was used to embed the contents of the modules into C source files statically linked with the program. Only a single Python file is still loaded at runtime, but it is read by a single rank and broadcast to the other ranks via MPI. This reduced the Python startup from 50 minutes at full scale to less than one second.

Additional improvements include

- Preventing Python from traversing directories (PYTHONPATH) searching for extension modules.
- Enabling ParaView on Eureka to use both GPUs per node.
- Bug fixes and optimizations to the ParaView coprocessor library.
- Optimizing the socket polling frequency on the GLEAN server side.
- Improved handshaking for data transfer on Eureka between the GLEAN server and the ParaView client.

Future work:

- *In situ* data extraction (ISDE) to generate contours currently uses uniform binning for point merge in ParaView. An octree-based method is 10 times faster for classic covisualization (CCV) so we expect a significant improvement for ISDE when it is implemented.
- CCV transfers the complete dataset to the staging nodes (Eureka) for each time step. The mesh accounts for 90% of the data volume, but does not change frequently. This communication may be optimized by caching the mesh data on Eureka.
- Various other problems need to be investigated, including the desynchronization observed when the pvserver is not able to process a time step before PHASTA sends the data of the next time step; a memory leak; and certain cases of *in situ* data extraction with GLEAN fail.

Spotlight on ALCF Staff: William Scullin

This issue's spotlight shines on William Scullin, senior HPC systems administrator and member of the ALCF's Operations Team.

NEWSBYTES: What is your background? What interests you professionally?

WILLIAM SCULLIN: I have two undergraduate degrees—one in history, one in political science—from Louisiana State University in Baton Rouge. My formal education has always been a red herring as I've been using and coding on Unix systems since I was in 8th grade. Since then, I've worked in groups doing everything from gravitational astrophysics to storm surge simulation to web development. I'd like to think of myself as a generalist, though I aspire to the title of computational philosopher.

NEWSBYTES: Tell us about your role here at the ALCF.

SCULLIN: I am a Senior HPC Systems Administrator in the Operations group. I share responsibility for the care and feeding of the Blue Gene/P. I mostly spend my time working support requests, maintaining development tools, performing hardware maintenance, and helping ensure the Blue Gene/P remains the best possible platform for computational science and high-performance computing.

NEWSBYTES: What did you want to be when you grew up?

SCULLIN: When I was little, I wanted to design aircraft. My parents are both engineers. They took my sister and me to air shows, balloon festivals, and read us lots of bedtime stories about the Apollo program—and they couldn't stop themselves from explaining the science behind it all.

NEWSBYTES: When did you start at the ALCF and how has it changed since you've been here?

SCULLIN: I started at the ALCF in 2006. Since then, I've seen three buildings, three models of Blue Gene, three directors, and more meetings than I'd care to admit. It's been interesting watching the ALCF mature. I miss the energy and the community of the early ALCF, but I'm glad the ALCF has grown to a point of steady-state operations.

NEWSBYTES: How has your role evolved over the years?

SCULLIN: As we've added some very strong staff in Operations and built expertise, I've been able to spend less time firefighting and more time focused on making the system and tools for administration and application development better.

NEWSBYTES: What goes on here that would surprise our readers?

SCULLIN: We spend a lot of time on analysis and reporting. Every ticket, every job issue, every request, each gets looked at in depth—even if all the user sees in the end is a brief email response.

NEWSBYTES: What's an exciting accomplishment you recall that stands out in your mind?

SCULLIN: It may have been mostly invisible to our users, but figuring out how to enable multiple minor versions of the IBM XL compilers to co-exist still remains a major personal victory. IBM doesn't officially support doing this nor do they document exactly how the compiler obtains its configuration; the end solution came through piecing together documentation from older compiler versions, versions of the compiler on other platforms, tracing compiler system calls, and experimentation.

NEWSBYTES: What do you like most about working for the ALCF?

SCULLIN: I like the people. Just about everyone here has a great sense of humor—even though it sometimes gets lost in translation. Likewise, management is very good about making sure we have the support we need to get things done. It makes for a very positive working environment.

NEWSBYTES: We expect the usual cerebral hobbies from our computer guys, but we hear you have a more unusual pastime and a strong right hook.

SCULLIN: I actually have two unusual hobbies—boxing and freestyle wrestling—and despite their reputations, both require a good deal of brainpower. They are amazing forms of exercise and great outlets for stress. As a bonus, I've been both in the ring and on the mat with ALCF users, though I'm not sure it's the type of close collaboration and contact the DOE had in mind when it created the INCITE program!

Questions for William can be sent to wscullin@alcf.anl.gov

Scaling Success at ALCF's Leap to Petascale Workshop, June 7-9

Scaling scientific code to work effectively on a system the size of the ALCF's Blue Gene, Intrepid, is no small task. To support the efforts of the scientific community working to exploit the full power of leadership-class systems, the ALCF offers an annual scaling workshop, Leap to Petascale (L2P). Scaling success was especially sweet for two teams this year who ramped their codes to an impressive 32- and 40-racks of Blue Gene.

To take part or assist in L2P at the ALCF, researchers traveled from Rice University, the University of Colorado, Stony Brook University, the University of Wisconsin-Madison, the University of Chicago and several others institutions. Twenty users and as many ALCF experts rolled up their sleeves and hunkered down for three intensive days of hands-on scaling work—several with the goal of scaling their code to demonstrate computational readiness for a 2012 INCITE award.

Tulin Kaman, from Stony Brook University, and his team successfully ran a multiphysics simulation code, FronTier, on 32 racks. Jean Perez from the University of New Hampshire Space Science Center achieved the full 40-rack milestone with his MHD turbulence code used in the study of solar wind.

Among other L2P efforts were teams at work on:

- debugging the PHASTA-GLEAN connection and creating a new "common" build of the code to facilitate collaboration with ALCF;
- scalability and performance tuning of a non-equilibrium umbrella sampling application and its use in simulating the folding and unfolding of RNA molecules;
- addressing MPI library memory issues when scaling MOAB, the HDF5 scientific format library, higher than 16k cores;
- porting cosmology simulation code and testing components such as fast Fourier Transform on Blue Gene/P; and
- performance analysis of QMCPACK—an electronic structure (quantum mechanics) code used to calculate the properties of materials.

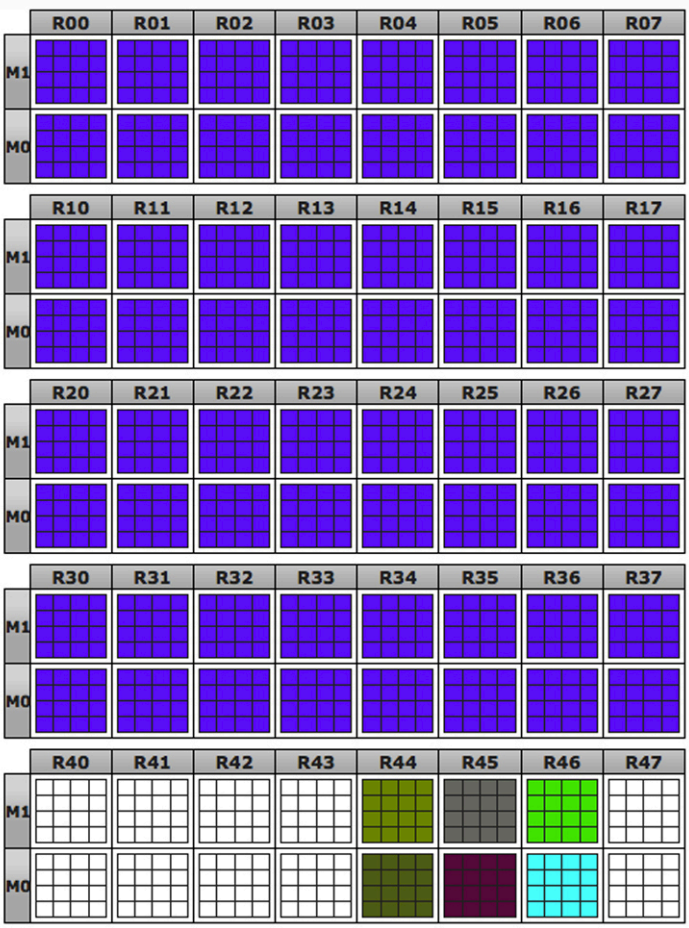

This snapshot of the ALCF System Status Display during the Leap to Petascale workshop June 7-9 reveals eight solid purple rows. Each purple row is equal to eight racks of code running to scale on Blue Gene/P, for a total of 32 racks. Scaling code at this level indicates readiness for production runs on the resources of a leadership-class facility, like the ALCF.

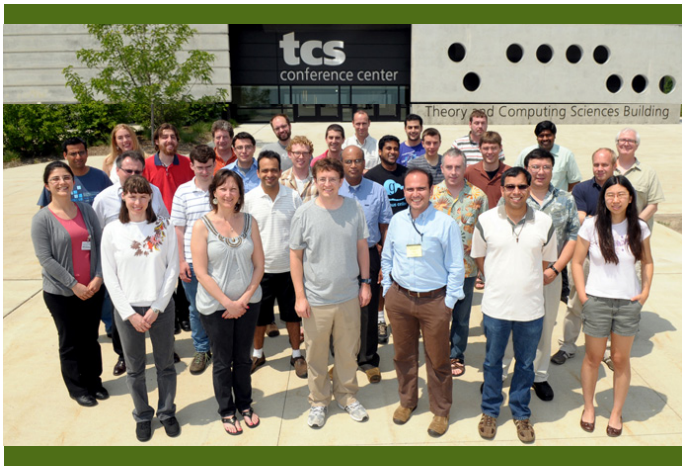

ALCF Newsbytes is a publication of the Argonne Leadership Computing Facility.
Visit us online at www.alcf.anl.gov Contact the editor at info@alcf.anl.gov

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC

U.S. DEPARTMENT OF ENERGY