ALCF Newsbytes

Argonne Leadership Computing Facility

Argonne National Laboratory

Volume 1, Issue 3 | October 2011

Michael Papka Appointed Director of Argonne Leadership Computing Facility

Michael Papka, deputy associate laboratory director (ALD) for Argonne's CELS Directorate, has assumed a new role as director of the Argonne Leadership Computing Facility (ALCF). Mike has served as interim director of the ALCF since August 2010. He will continue to serve as deputy ALD.

Mike has expertise in scientific visualization and data analysis, and over two decades of experience working with advanced display technology and high performance computing to promote greater insight into scientific simulation data. Mike earned his doctoral degree in computer science from the University of Chicago, his master's degree in computer science from the University of Illinois at Chicago, and his bachelor's degree in physics from Northern Illinois University.

Mike joined Argonne in 1992 and first served as the co-director of the Mathematics and Computer Science Division's Futures Laboratory, a role he continues to fulfill today as team lead of its visualization and data analysis group. Mike is also a senior fellow of the Argonne National Laboratory/University of Chicago Computation Institute, where he leads multiple visualization and scientific data analysis efforts.

Program Awards 185 Million Hours of ALCF Supercomputer Time

Seven research projects have been awarded 185 million processor hours of computing time at the Argonne Leadership Computing Facility (ALCF) through the U.S. Department of Energy's (DOE) initiative—the ASCR Leadership Computing Challenge (ALCC). Chosen through a peer-review process, the projects selected reflect areas of special interest to the DOE: energy, national emergencies, and broadening the community of researchers capable of using leadership-class computing resources.

"By providing millions of hours of computing time on the ALCF's powerful system, these awards facilitate the work of researchers around the country who are addressing some of the nation's most challenging scientific problems,"

ALCF Director, Michael Papka

continued ▸▸

Table of Contents

Events of Interest

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. See page 8 for more information about the ALCF at SC11.

DOE Issues ALCC Call for Allocations

The Department of Energy (DOE) has issued a call for researchers seeking allocations of compute hours through the ASCR Leadership Computing Challenge (ALCC). Peer-reviewed ALCC awards are available to researchers in national laboratories, academia and industry. Submissions will be accepted through February 14, 2012. For submission details, visit: http://science.energy.gov/ascr/facilities/alcc/. For information about ALCC projects underway at the ALCF, visit: http://www.alcf.anl.gov/collaborations/alcc/alcc10_index.php.

Argonne NATIONAL LABORATORY

Program Awards 185 Million Hours of ALCF Supercomputer Time (continued)

ALCC awards of compute time on Intrepid (the ALCF's Blue Gene/P) became available July 1 for the following recipients:

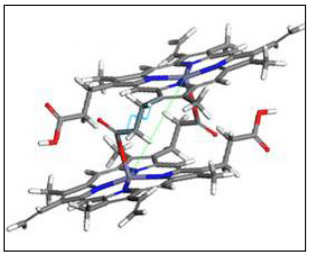

James R. Chelikowsky and ***Noa Marom*** in the Computational Materials Center at the Institute for Computational Engineering & Science at the University of Texas at Austin were awarded 12 million hours to work on crystal engineering from first principles for applications such as hydrogen storage, drug design, and non-linear optics.

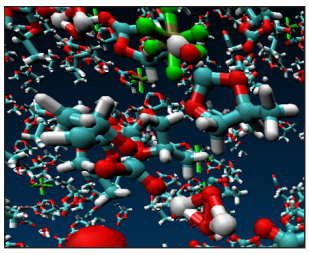

Larry Curtiss of Argonne's Materials Science Division and the Center for Nanoscale Materials at Argonne was awarded 30 million hours for the design and discovery of new materials critical to the nation's energy future, including electrical energy storage and catalytic materials.

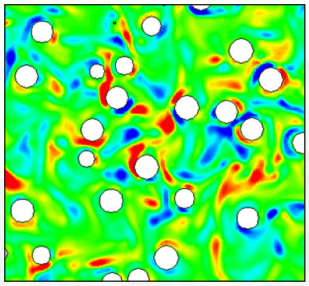

Said Elghobashi from the University of California, Irvine, received 10 million hours to study the two-way interactions between vaporizing droplets and turbulent flows using fully resolved, direct numerical simulation (DNS). The study is of particular relevance to liquid fuel spray combustion that is the predominant energy source in transportation systems.

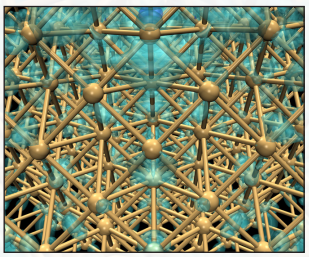

Jeffrey Greeley in the Theory and Modeling group at the Center for Nanoscale Materials at Argonne was awarded 20 million hours for research into clean energy solutions, including work to extend the life of catalysts in fuel cells and investigations of titania as a durable, efficient electrode in lithium ion batteries.

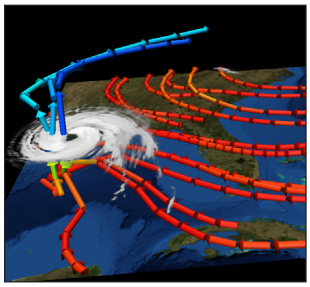

Led by ***Greg Holland***, a team of researchers with the National Center for Atmospheric Research received 13 million hours to develop high-resolution climate models to advance our understanding of Earth's climate. Efforts will increase the predictive ability of climate models for better national emergency preparedness and will serve to broaden the community of researchers capable of using leadership computing resources.

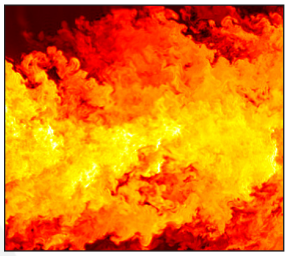

Parviz Moin, founding director of the Center for Turbulence Research, was awarded 60 million hours at the ALCF to further critical research into how jet noise is produced. The information will be used in noise-mitigation strategies in the design of next-generation supersonic transport vehicles.

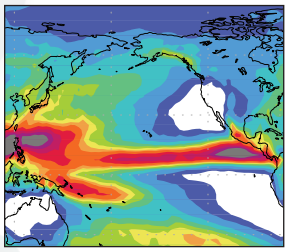

Laura Zamboni from Argonne National Laboratory received 40 million hours to investigate the sensitivity and uncertainty of precipitation in climate models. Using high-resolution models, this research will further risk assessments and water management worldwide, and deepen our understanding of processes in cloud formation and precipitation.

For more about ALCC, visit http://www.alcf.anl.gov/collaborations/alcc/alcc_index.php.

First User Advisory Council Give Users Greater Voice at the ALCF

To give users more input into the policies and services that shape their experiences at the ALCF, the first User Advisory Council was formed in the spring of this year. Comprised of seven users from an array of scientific and technical backgrounds, the UAC meets monthly to provide advice and feedback to facility leadership on the current state of ALCF resources and services from the user perspective.

UAC members are appointed by the ALCF director in January of each year. Collectively, they represent ongoing projects from all programs (INCITE, ALCC, Early Science, etc.) using the facility's computational resources. To qualify for a one-year UAC term, a user must be a member of an active project, but need not be a PI.

By sharing information and experiences, suggesting improvements, and voicing user concerns, the UAC furthers the effective use of leadership-class resources at the ALCF.

The UAC also provides input on policies that affect users and may also ultimately act as the oversight body of a larger ALCF user group.

To date, the UAC has provided feedback on a variety of topics, including proposed changes to job scheduling policies, suggestions on ways to increase user participation in the monthly ALCF User Call, and reactions to structure, content and navigation for the ALCF website.

Users who have suggestions or comments for the UAC are encouraged to contact Chair, Dean Townsley or UAC/ALCF liaison, Chel Heinzel. For more information about the UAC members, visit: http://www.alcf.anl.gov/about/uac.php.

UAC members for 2011-12		
Dean Townsley (Chair)	Astrophysics	University of Alabama
Andrew Binkowski	Biosciences	Argonne National Laboratory
Hank Childs	Visualization	Lawrence Berkely National Laboratory, UC Davis
Curtis Hamman	Turbulence/CFD	Center for Turbulence Research, Stanford
Philip Maechling	Environmental Sciences	Southern California Earthquake Center
Ron Minnich	Computer Sciences	Sandia National Laboratories
Steven Pieper	Nuclear Physics	Argonne National Laboratory

Chel Heinzel, UAC/ALCF staff liaison.

"The UAC is a great addition to the way the ALCF supports cutting-edge science by keeping in touch with user needs. As a standing venue for discussion, the UAC enables ALCF users to be more involved in decisions that shape their ability to accomplish science."

UAC Chair, Dean Townsley

New Algorithms for Numerical Linear Algebra Reduce Communication to Improve Efficiency on the ALCF's Blue Gene/P

Parallel dense matrix computations are computational bottlenecks in many scientific simulations. Researchers at the University of California-Berkeley, led by Edgar Solomonik, have derived novel algorithms for these kernels that map closely to modern supercomputer architectures such as Intrepid, the IBM Blue Gene/P at the Argonne Leadership Computing Facility. These algorithms exploit the full memory capacity of the machine and take advantage of the interconnect network topology to reduce communication. The methods are dubbed "2.5D algorithms" because they generalize known 2-D and 3-D algorithms and adapt to the underlying architecture.

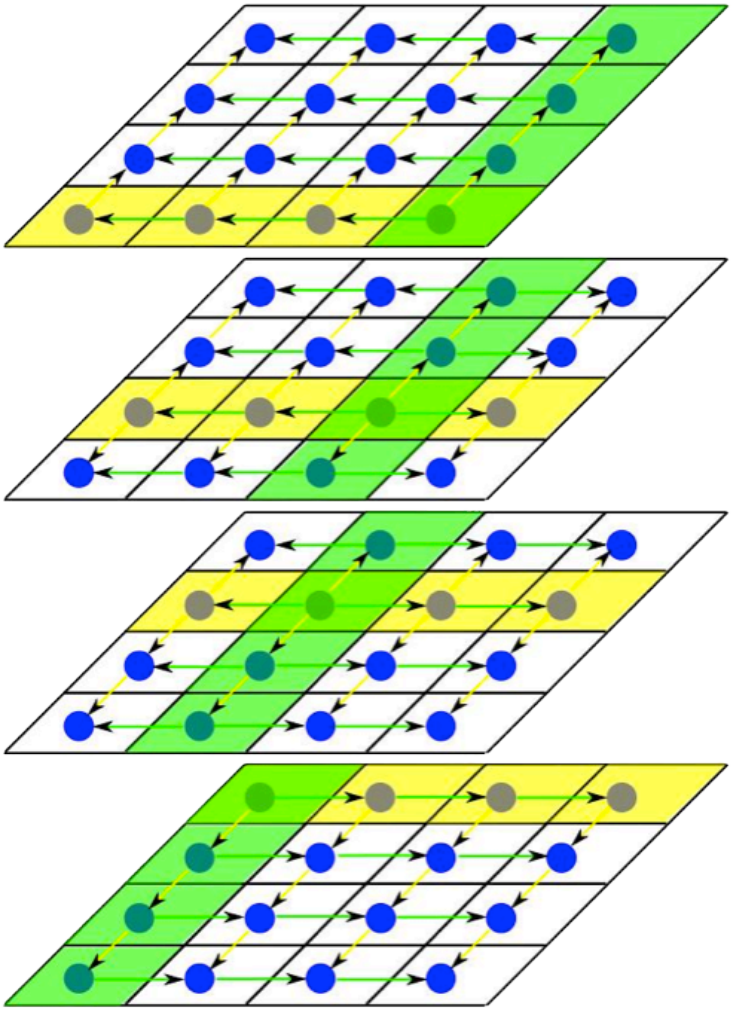

Matrix mapping to 3D partition of Blue Gene/P.

On the theoretical side, 2.5D algorithms are probably faster due to a lower communication complexity. The new algorithms also have practical advantages: On the Blue Gene/P, their adaptivity allows deployment of its unique communication infrastructure. This infrastructure performs collective communication with much more efficient, topology-aware algorithms. By improving communication performance, 2.5D algorithms improve parallel scalability and drastically increase efficiency for large-scale problems that are distributed over machines the size of Intrepid.

During a Computational Science Graduate Fellowship practicum at Argonne, Solomonik worked with ALCF Assistant Computational Scientist Jeff Hammond to extend and apply these algorithmic ideas to accelerate tensor contractions performed in quantum chemistry electronic structure calculations.

The published work, based on the research conducted at UC-Berkeley, has been recognized as a Distinguished Paper by the Euro-Par 2011 Conference, Bordeaux, France. Another paper on this work was accepted for publication and presentation at the ACM/IEEE Supercomputing conference, to be held in Seattle in November.

Contact: Edgar Solomonik (solomonik@berkeley.edu)

Argonne-developed Software Enables World-record Calculation

Leveraging software developed by the Argonne National Laboratory Leadership Computing Facility (ALCF), scientists at the Juelich Supercomputing Centre (JSC) have performed a world-record simulation of the N-body problem—calculating over three trillion particles in eleven minutes.

Ivo Kabadshow and Holger Dachsel of JSC developed a state-of-the-art algorithm and code for applying the fast multipole method (FMM) to the N-body problem. However, the researcher were unable to fully utilize the world's largest Blue Gene/P system (JUGENE) due to the code's dependence on one-sided communication software developed for traditional clusters, not supercomputers, rendering it unable to scale to more than 1000 nodes of this architecture.

Led by ALCF Assistant Computational Scientist Jeff Hammond, a team at Argonne National Laboratory solved this problem by creating a new communication library for Blue Gene/P that is faster and more scalable than existing libraries. Known as A1 ("Argonne 1-sided"), Hammond and Pavan Balaji designed the software from scratch for leadership supercomputer architectures. Hammond and Sreeram Potluri, a visiting student from Ohio State University, implemented it in the summer of 2010. Over the next six months, Kabadshow and Hammond worked to integrate A1 into the JSC FMM code and to further optimize its memory use and scalability, since the library needed to scale to all 294,912 cores of JUGENE. Because of the extreme scale of this machine, any operation that required time or memory proportional to the number of cores would become a bottleneck.

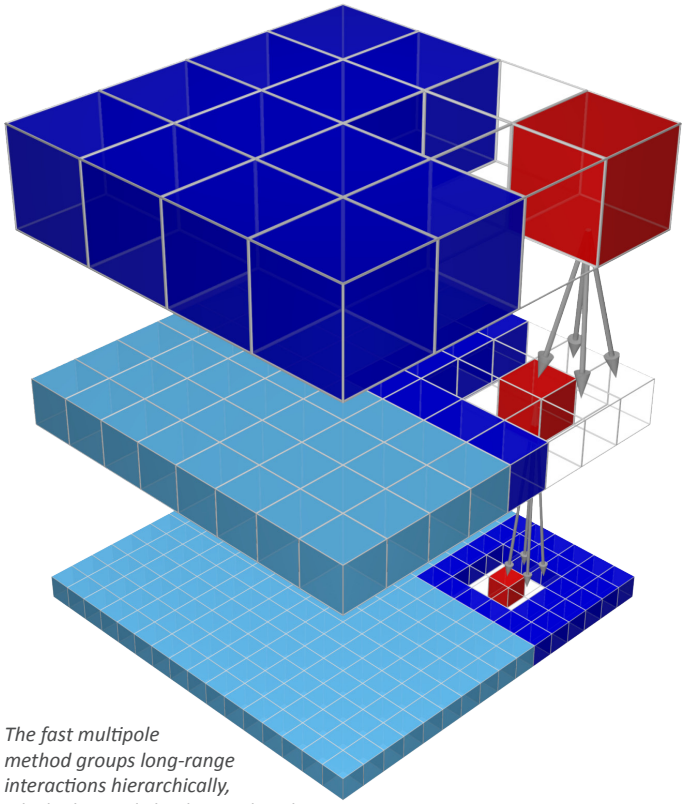

The fast multipole method groups long-range interactions hierarchically, which ultimately leads to reduced computational cost. Image courtesy of Ivo Kabadshow. Produced using systems at the Argonne Leadership Computing Facility and the Juelich Supercomputing Centre.

The collaboration between JSC and Argonne yielded an amazing outcome—the N-body problem for more than three trillion particles can be solved in only eleven minutes! For comparison, this simulation running on a single processor would take more than half a year, if it were even possible, since the memory required to perform such a simulation requires almost the entire memory of JUGENE (144 terabytes).

The results of these efforts are two-fold: more powerful simulation tools for the N-body problem, which is ubiquitous in fields such as astrophysics, biology and chemistry, and more powerful communication software for architectures like Blue Gene/P. With the arrival of the ALCF's Blue Gene/Q system in 2012—a machine that will have almost three times the number of cores as JUGENE—the scaling optimizations in the JSC FMM code and A1 pave the way for unprecedented N-body simulation in the near future. These researchers will continue to collaborate and tune each of their respective codes for next-generation supercomputers and intend to branch out beyond Blue Gene/P and Q to state-of-the-art Cray machines and other architectures.

For more information, contact Jeff Hammond, or visit the JSC website.

Blue Gene/Q Prototype Tops Green500 List of Energy-efficient Supercomputers

This past spring, the Green500 released its rankings of the most energy-efficient supercomputers in the world. Topping the list is the IBM Blue Gene®/Q—a prototype of the production machine the Argonne Leadership Computing Facility is scheduled to deploy in 2012, continuing the ALCF's long-standing tradition of cutting-edge efficiency in supercomputing.

Six of the top ten most energy-efficient supercomputers in the world are built on IBM high-performance computing technology. The list includes supercomputers from the United States, China, and Germany that are being used for a variety of applications such as astronomy, climate prediction, and life sciences.

Energy efficiency, including performance-per-watt for the most computationally demanding workloads, has long been a core design principle in developing IBM systems.

Energy-efficient supercomputers offer critical cost savings by lowering power consumption, reducing expenses associated with cooling, and scaling to larger systems while maintaining acceptable power consumption. For example, for every $1 spent on electricity with the largest petascale system on the Green500 list, less than $0.40 cents would be spent on a system based on the IBM Blue Gene/Q and would be 2.5 times more energy efficient.

THE GREEN 500

Blue Gene/Q Prototype 2 ranked #1 June 2011

More information about the Green500 list is available at http://www.green500.org.

ALCF's Intrepid Ranked #1 on Graph 500 List for Two Years Running

GRAPH 500

For the second consecutive year, Intrepid, the Argonne Leadership Computing Facility's (ALCF) IBM Blue Gene/P was ranked #1 on the Graph 500 list. The list ranks supercomputers based on their performance on data-intensive applications and thus complements the Top 500 list that is based on the LINPACK benchmark (http://www.top500.org/project/linpack). The latest rankings were announced at ISC'11 in Hamburg, Germany this spring.

Traditional benchmarks and performance metrics fail to provide useful information on the suitability of supercomputing systems for data-intensive applications. Backed by a steering committee of more than 30 international HPC experts from academia, industry, and national laboratories, Graph 500 established a new set of large-scale benchmarks for these applications. The new benchmarks will guide the design of hardware architectures and software systems intended to support such applications and help procurements.

An ALCF team, led by Kalyan Kumaran, ALCF Manager, Performance Engineer, worked with an Indiana University team led by Jeremiah Willcock and an IBM team coordinated by Fabrizio Petrini, to submit the #1 result on 131,072 cores of Blue Gene/P—four times larger than the previous result on 32,768 cores submitted at SC10 in November 2010—using a graph that likewise was four times larger. To date, the graph is the largest on record ever analyzed on a parallel machine of any kind. In addition, IBM submitted a result on the Blue Gene/Q on 8,192 cores with 11.2 Billion Traversed Edges Per Second (TEPS) with only 512 nodes. This result is #13 on the Graph 500 list and clearly indicates the potential of the new generation of Blue Gene machines.

The Graph 500 benchmark was first announced at ISC2010, and the premier list appeared at SC10. The next Graph 500 list will be announced at SC11.

For more information, contact Kalyan Kumaran, or visit the Graph 500 website.

Spotlight on ALCF Staff: Nick Romero

This month, Nick Romero, catalyst and assistant computational scientist at the ALCF gets caught in the Newsbytes spotlight.

NEWSBYTES: What's your educational background and areas of interest?

NICK ROMERO: I have a B.S. in physics from MIT and a Ph.D. in physics from the University of Illinois at Urbana-Champaign. My primary scientific interest is the electronic structure of condensed matter. I have mathematical and computational interests as well (numerical methods for optimization, programming models, etc.)

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

ROMERO: I was very interested in science at an early age. I picked up my first physics book in first grade and learned about the structure of atoms. I figured I'd end up being a mathematician, a physicist, a chemist, or a computer programmer. I even shuffled among those careers through most of college. In my present role, I am a little bit of each.

NB: You're a Catalyst here at the ALCF. Can you describe the role of a Catalyst?

ROMERO: The Catalysts' role is to help INCITE and ALCC projects achieve their scientific goals. A robust set of skills is required to reach this end, which includes providing performance analysis, debugging, I/O, visualization and even recommending software packages. Since it is unlikely that any single Catalyst is an expert in all these areas, we leverage the expertise of the entire Catalyst team and other resources within the ALCF to assist our users. We also have the ability to collaborate with other Argonne divisions when necessary, especially MCS (Mathematics and Computer Science Division).

NB: Describe your work here and how it has evolved over the years.

ROMERO: As a post-doc, I was tasked initially with improving the performance of the GPAW code on Blue Gene/P. When I was made staff, it became my responsibility to support material science projects with INCITE awards. Support can span the gamut of logistics (reservations, collecting reports) to code development. I still develop the GPAW code in support of a couple of INCITE projects. Recently, my attention has focused on working with our collaborators to prepare a strategic plan for electronic structure on the next-generation Blue Gene/Q, which will be delivered to the ALCF in 2012.

NB: How has the ALCF evolved since you've been here?

ROMERO: I started at the ALCF as a post-doc in 2008, and was re-classified to staff in 2010. We were a very small applications group then—two catalysts and three performance engineers. Now there are more than a dozen of us. It's great for the users to have the additional support, but it can be hard to fit us all in one conference room for meetings.

NB: What goes on here that our readers might not have considered?

ROMERO: Every new computer architecture requires us to re-think our electronic structure algorithms and, unfortunately, to a large extent, re-write our software. Someone more senior that I may say that this has always been the case; I personally think that the next 10 years of computing will introduce radical changes—the likes of which we haven't seen in a while.

NB: What's an exciting development that stands out in your mind?

ROMERO: Electronic structure software and supercomputers have come a long way, even since my first post-doc. Calculations that were impossible when I graduated in 2005 are now straight-forward with present codes on our Blue Gene/P. As part of Jeff Greeley's INCITE project, we have performed electronic structure calculations on nanoscale particles containing over 10,000 valence electrons.

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

ROMERO: I get to work with top-notch scientists on high-impact projects. I also enjoy being a part of the bleeding-edge of computing technology.

To read more about Jeff Greeley's INCITE project, see page 39 of the 2010 ALCF Science Brochure, available in our online media kit: http://www.alcf.anl.gov/news/media_kit.php.

For more information about the ALCF's Catalyst program, visit: http://www.alcf.anl.gov/resources/catalyst.php.

Join the ALCF at SC11 in November

For the twentieth time since the inception of the conference in 1989, Argonne is heading to SC. This year, two 48-foot semis packed with booth parts, audio/visual equipment, posters and supplies will set out in advance of a team of Argonne representatives making the annual pilgrimage. SC11 is November 12-18, at the Washington State Convention Center in Seattle, Washington.

The ALCF finds itself at home at SC, which brings together the international community in high performance computing, networking, storage and analysis. The conference attracts researchers, scientists, computing center staff members, IT and data center management, application developers, computer manufacturing personnel, program managers, journalists and congressional staffers making it a "must attend" for stakeholders throughout the technical computing community—including the ALCF.

As always, ALCF technical and scientific experts will be on hand at the Argonne booth, offering visitors the opportunity to ask questions and to find out more about the behind-the-scences efforts that go into running one of the nation's premier supercomputing facilities. The booth will also feature electronic posters of the exciting research underway at the ALCF, and a variety of speakers will cover a wide range of current topics of interest.

ALCF guest speakers include:

James Sexton, Program Director, Computational Science Center, T.J. Watson Research Center, will speak about the latest IBM HPC technology, including the Blue Gene/Q.

Gordon Bell finalist, ***Leopold Grinberg***, senior research associate at the Division of Applied Mathematics, Brown University, will discuss his work, "A New Computational Paradigm in Multiscale Simulations: Application to Brain Blood Flow."

William George, computational scientist with the National Institute of Standards and Technology will present, "Advancing the Materials Science of Concrete with Supercomputers."

David Lecomber, chief technical officer with Allinea Software, will present on scalable debuggers for the Blue Gene/Q.

Several Argonne divisions will be represented in the booth this year. For details, visit: sc11.anl.gov.

"Getting Started" Workshop Primes New Users and Staff

This year's second "Getting Started" workshop was held October 4-5 at the ALCF. Now in its fourth year, the workshop provides users with information they need to fully utilize the resources of the facility. Talks included an overview of ALCF services and resources, technical details of system architecture, debugging and performance tools, and visualization capabilities. The workshop brought ALCF experts and users together, providing time for collaboration and hands-on assistance in porting and tuning applications on the Blue Gene/P.

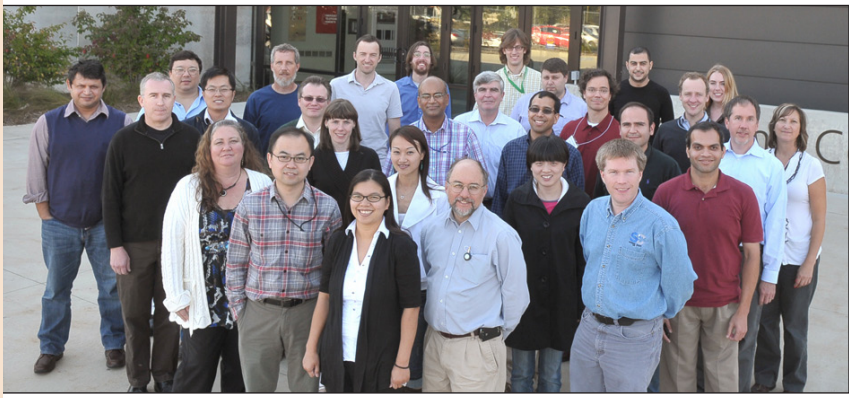

Representatives from Dow Chemical, Sandia National Laboratories, the University of Chicago and several divisions from within Argonne attended the event. The workshop also provided a unique training opportunity for a team of postdocs brought on to work with the Early Science Program (ESP) project teams, who, in collaboration with ALCF staff and IBM, will undertake intensive efforts to adapt their software to maximize the Blue Gene/Q architecture.

Presentation slides from the October Getting Started workshop are available at http://workshops.alcf.anl.gov/gs11fall/agenda.

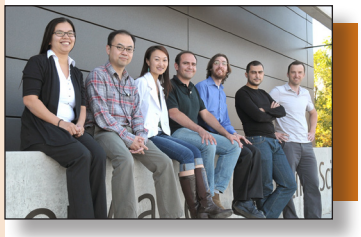

The Getting Started workshop was attended by postdocs brought on to work with the ALCF Early Science Program project teams, in collaboration with ALCF staff and IBM, to undertake intensive efforts to adapt their software to take advantage of the Blue Gene/Q architecture.

Left to right: Maricris Mayes, Heechang Na, Yun (Lyna) Luo, Alvaro Vazquez-Mayagoitia, Hal Finkel, Anouar Benali, and Geoff Ely.

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