ALCF Newsbytes

Argonne Leadership Computing Facility

Argonne National Laboratory

Volume 2, Issue 3 | Summer 2012

Program Awards 247 Million Hours of ALCF Supercomputer Time

Nine research projects have been awarded 247 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.

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

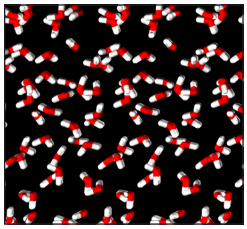

Jeffrey Greeley with Argonne National Laboratory was awarded 20 million hours for a groundbreaking study using electronic structure methods that combine modeling of solid/liquid interfaces with detailed mechanistic analysis of electrochemical deNOx chemistry to allow researchers to both understand and, ultimately, propose new catalysts with specific applications in environmental chemistry and broader applications to other electrocatalytic problems of interest.

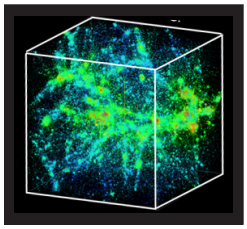

Katrin Heitmann with Argonne National Laboratory was awarded four million hours to create unprecedented simulations that will lead to a better understanding of neutrinos, the lightest known massive particle in the universe, and the role they may play in cosmic evolution and structure formation.

continued ▸▸

Table of Contents

Program Awards 247 Million Hours 1

Mira Ranks Third on TOP500 2

Rapid Procedure Unites Quantum Chemistry with Artificial Intelligence 4

Lithium-Air: The "Mount Everest" of Batteries for Electric Vehicles 5

Spotlight: Steve Crusan 6

ALCF Systems Shine on Green500 7

Mira Shares Graph 500 Top Spot 8

Leap to Petascale Workshop in Review 8

Events of Interest

Argonne's Energy Showcase
Saturday, September 15, 2012
9 a.m. – 4 p.m.

Argonne National Laboratory will open its gates to the community for a day of discovery and fun for the whole family. The event is free and open to the public. Advance registration is required as attendance will be limited. Visit http://www.regonline.com/builder/site/Default.aspx?EventID=1097223 for details.

Blue Gene Consortium – Blue Gene/Q Summit
Tuesday, October 2, 2012
Argonne National Laboratory
Hosted by the ALCF and IBM

Join current and future Blue Gene users for this daylong event of BG topics, including:

- Science on BG (INCITE and Early User science)
- BG/Q Programming Primer – lessons learned
- "Ask the Expert" Session
- Building Community through the Blue Gene Consortium

In addition, the ALCF is offering access to their BG/Q Test & Development system for preliminary investigations prior to the event. For Summit details, or to request system access, visit http://www.bgconsortium.org/BGQ-Summit.

Argonne NATIONAL LABORATORY

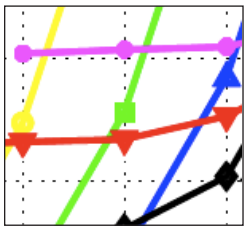

Terry Jones, Oak Ridge National Laboratory, received three million hours to test newly developed system software essential for harnessing the power of next-generation exascale supercomputers and for making these leading-edge resources more accessible to researchers throughout the scientific community.

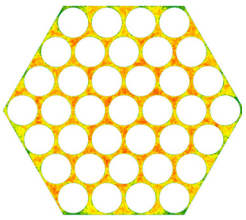

Elia Merzari, with Argonne National Laboratory, received 30 million hours to create extreme-fidelity physics models that will enable new reactor design analysis in existing and next-generation reactors on exascale computing platforms.

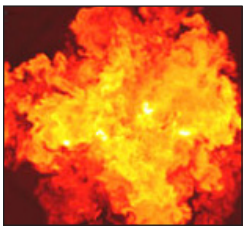

Parviz Moin, Stanford University, was awarded 60 million hours for large-eddy simulation of jet noise and validation of turbulence models in complex geometries.

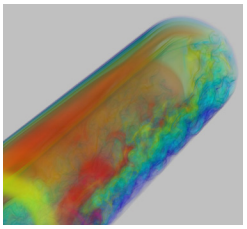

In collaboration with the Moscow Institute of Nuclear Energy Safety (IBRAE), ***Aleksandr Obabko*** and the SHARP team with Argonne National Laboratory received 30 million hours to advance high-fidelity computational capabilities for research into complex thermo-fluid phenomena that will help make safe nuclear energy a reliable, carbon-free energy source.

Argonne's Mira Ranks Among Nation's

Mira, Argonne National Laboratory's new IBM Blue Gene/Q system, is the third fastest supercomputer in the world according to the TOP500 list announced June 18. Mira has 20 times the processing power of Intrepid, it's predecessor at the ALCF. At peak performance, Mira will be capable of 10 quadrillion floating-point operations per second.

Now in its 39th year, the TOP500 is the semiannual ranking of the world's most powerful supercomputers. The list began as a way to produce meaningful statistics for supercomputing systems and is a widely observed benchmark that sets performance targets for vendors to deliver increased capabilities to the most challenging computational science applications.

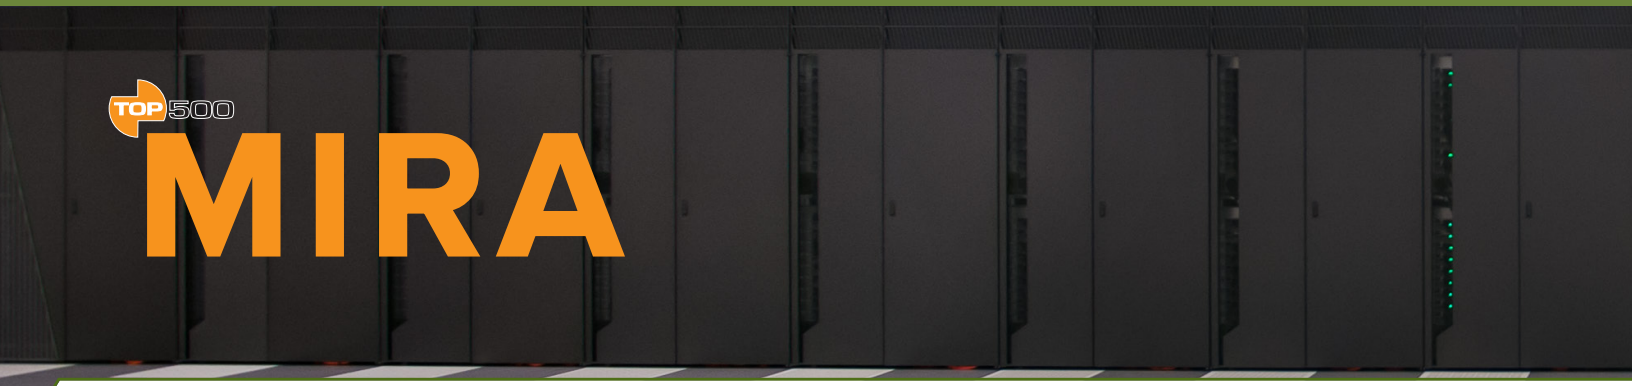

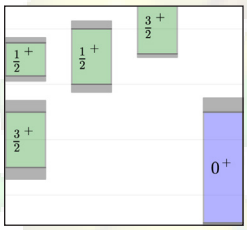

Kostas Orginos, College of William and Mary and Thomas Jefferson National Accelerator Facility (Jefferson Lab), was awarded 20 million hours to study the spectrum of light nuclei and hypernuclei directly from the underlying quantum field theory of the strong interactions, a central goal of nuclear physics research, where findings will impact the worldwide hypernuclear experimental program and the realm of nuclear astrophysics by refining our understanding of the lifecycle of stars.

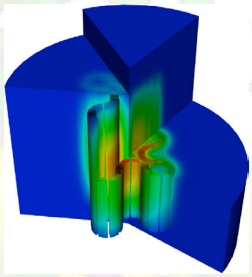

Micheal Smith, Argonne National Laboratory, was awarded 30 million hours to conduct validation and verification of heterogeneous nuclear code reactor calculations that will improve the ability of nuclear reactor physics codes to solve real-world design problems.

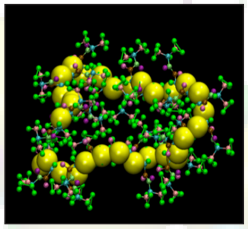

Subramanian Sankaranarayanan with Argonne's Center for Nanoscale Materials was awarded 50 million hours to study conformational transformations in thermo-sensitive oligomers and their macroscopic architectures such as polymer brushers and polymer gels.

Fastest Science Supercomputers

The TOP500 ranks supercomputing systems based on their Linpack benchmark score—a special-purpose computer code that scores application runs in quadrillions of floating-point operations per second, or petaflops. For the June list, Mira achieved 8.1 petaflops per second on the Linpack benchmark, using 786,432 processing cores on the machine's full 48 racks.

Expected to be fully operational later this year, Mira gives the ALCF the compute muscle to continue its mission to provide world-class resources for groundbreaking science and engineering, "We've entered the next great revolution in computational science, which will involve the development of predictive models of complex, interconnected processes," said ALCF Director Michael Papka. "Mira represents our continued efforts to design HPC systems that meet the unique priorities of science codes."

Rapid Procedure Unites Quantum Chemistry with Artificial Intelligence

By combining quantum chemistry with artificial intelligence (machine learning), researchers at the Institute for Pure and Applied Mathematics at the University of California, Los Angeles, achieved a scientific breakthrough expected to aid in exploring chemical compounds. Led by Anatole von Lilienfeld, with the Argonne Leadership Computing Facility, the interdisciplinary team from the Technical University Berlin (Germany), the Fritz-Haber Institute of the Max-Planck Society (Germany), and the Argonne Leadership Computing Facility (United States) dramatically increased the speed of calculating energies of small molecules with quantum chemical accuracy.

Quantum chemical methods permit scientists to calculate molecular properties on a computer from first principles (i.e., without having to conduct any experiments). They are necessary for many chemical applications such as catalysis, or the discovery of novel materials. Previously, such calculations demanded intensive computational resources. Machine learning, however, generates predictive models based on examples. An algorithm is fed examples, and the computer uses them to predict an outcome. When applied to quantum chemistry, thousands of quantum chemical reference energies have been calculated in order to "learn" a molecular model. The resulting machine permits the prediction of molecular properties with comparable accuracy within milliseconds, instead of hours. This speed-up paves the way for highly accurate calculations of an unprecedented number of molecules.

The team focused on a basic quantum chemical property: the energy tied up in all the bonds holding a molecule together—the atomisation energy. The researchers calculated atomisation energies for a database of 7,165 molecules, and used it to design a machine-learning model capable of predicting atomisation energies for new molecules out of sample. Intrepid, the Blue Gene/P supercomputer at the ALCF, will be used to extend this approach to millions of other compounds or materials. This will permit researchers to tackle more complex problems, such as finding novel catalysts for chemical reactions, designing bio-mimetic materials, or other purpose-built materials.

Articles featuring the research have been published recently in *Physical Review Letters, New Scientist,* and *Chemistry World.*

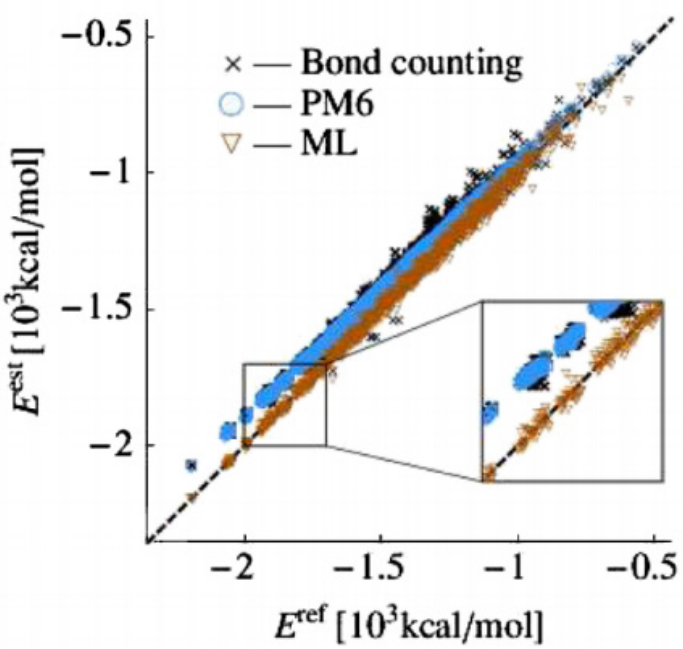

Machine learning models estimate molecular atomisation energies with errors smaller than 10 kcal/mol, outperforming semi-empirical methods, or the counting of covalent bonds.

Image courtesy of Anatole von Lilienfeld, Argonne Leadership Computing Facility

Lithium-Air: The "Mount Everest" of Batteries for Electric Vehicles

Is it possible to produce a novel lithium-air battery that could drive a car 500 miles per charge? Researchers from Argonne and Oak Ridge leadership computing facilities and IBM Research think so.

Weight-for-weight, a rechargeable lithium-air battery can store five to ten times the energy of a conventional lithium-ion battery. It has the highest energy density of any battery yet devised. Currently, the batteries in electric cars can't compare with gasoline in the amount of energy derived from a given weight of fuel. Cars with lighter, more powerful lithium-air batteries could overcome this challenge. The battery's enhanced energy storage capability would enable the widespread use of electric vehicles.

Realistic computational chemistry models are of crucial importance to developing a viable lithium-air battery. Led by Jack Wells, director of science at the Oak Ridge Leadership Computing Facility, a research team modeled how the lithium-air cell would work. Using supercomputing resources allocated by a Department of Energy INCITE award, the researchers studied the cell's basic chemistry, how catalysts affect its performance, and how to make the device more stable. They investigated the use of propylene carbonate (PC) as a potential electrolyte for a rechargeable lithium-air battery to establish the practicality of PC. By employing sophisticated models for both the surface and the electrolyte, they achieved sufficient detail to study the full complexity of propylene carbonate degradation by lithium-air discharge products. The team evaluated this problem using three application codes—CP2K, CPMD, and VASP—on Intrepid, the IBM Blue Gene/P system, at the ALCF.

For the first time, the scientists demonstrated that PC is not a suitable electrolyte for lithium-air-based batteries. Formerly, PC was believed to be stable, based on the behavior of similar systems. However, the simulations in this project revealed that during discharge, oxygen reduces to peroxide, forming layers of Li_2O_2 that immediately degrade PC. Moreover, special properties of the surface of Li_2O_2 act to enhance its reactivity. The findings also explained the recent failure to identify Li_2O_2 in the discharge products on the cathode surface during experimental research into this technology and determined the major chemical mechanisms behind the charge/discharge processes. The calculations that resulted in these insights used 26.5 million core-hours on Intrepid.

Groundbreaking research conducted with a U.S. Department of Energy INCITE award was featured on the March 2012 cover of Chemistry—A European Journal, accompanied by an article detailing the research findings. (Laino, T. and Curioni, A., "A New Piece in the Puzzle of Lithium/Air Batteries: Computational Study on the Chemical Stability of Propylene Carbonate in the Presence of Lithium Peroxide," Chem. Eur. J., Vol. 18, Issue 12, p. 3421, March 2012.)

Spotlight on ALCF Staff: Steve Crusan

This month, the Newsbyte spotlight shines on Storage Team member, Steve Crusan.

NEWSBYTES: What's your role here at the ALCF?

STEVE CRUSAN: I'm a member of the Operations Storage Team. I'm basically an HPC sys admin for our GPFS. It's my job to maintain, fix and upgrade development tools for GPFS—to make the system available, fast and reliable.

NB: You joined the Storage Team in December of last year, which makes you one of its newest members. What's your background and what brought you to the ALCF?

SC: I have a degree in Networking Security and System Administration from Rochester Institute of Technology. Before here, I worked at the University of Rochester on their single rack of Blue Gene/P and their Linux clusters. The scale of the ALCF systems intrigued me. Luckily, working on GPFS on a bigger machine is conceptually pretty much the same as working on a small machine. Mira, our new Blue Gene/Q, will have a 30-petabyte file system. Most people won't ever touch systems of this size, so it's pretty exciting.

NB: What's challenging about your work here?

SC: All parallel file systems have issues, and they're inherently difficult to test at the scale we use them on, especially for developers. If there's any weakness in the chain, the supercomputer will find it and the file systems will grind to a halt. So, there's always the stress of potentially losing data. Plus, if users have storage issues, I have to be able to respond immediately, which can make getting project work done more of a challenge.

NB: Besides you, what's new with the Storage Team?

SC: The "team" used to be just one person. Now it has five dedicated members. We also have stronger ties to the IBM GPFS development team so the collaboration is better.

NB: What does that mean for our users?

SC: It means system uptime has improved. For years, GPFS issues topped the list of reasons for system interrupts. We worked with IBM and identified overloaded cluster managers as a likely root cause. So, we added two new cluster managers and placed them on X86 boxes. So far, we haven't seen any more GPFS-timeout issues.

NB: What's on the storage horizon for you with the Mira installation underway?

SC: I'll be working with the Storage Team and the IBM GPFS development team on Mira's filesystem layout. The filesystem will be accessed via RDMA, versus over 10 gigabit ethernet on Intrepid. We're also addressing some of the storage capacity issues and filesystem subscription problems we've had on Intrepid by instituting quotas; per-user home quotas, and per-project quotas on the scratch and work filesystems.

NB: Newsbytes had a chance to ask ALCF Director of Operations, Bill Allcock, for his thoughts on you as one of the Storage Team newbies. He said, "Steve came in and addressed storage issues that have been bedeviling us for years." Your reaction?

SC: Sometimes I think it just takes a fresh set of eyes to look at a persistent problem to identify new possible solutions.

Argonne's Mighty Mira Joins Ranks of World's Most Energy-Efficient Computers

The 20 most energy-efficient supercomputers in the world are IBM Blue Gene/Q systems according to the latest Green500 list announced by Green500.org. Technology originating from U.S. national laboratories now leads the trend toward sustainable and environmentally responsible high-performance computing.

Mira, Argonne National Laboratory's Blue Gene/Q and a U.S. Energy Department petascale resource in support of scientific research, and its two testing and development racks, are among those leading the pack. Green500 ranks systems by energy efficiency, or performance per watt, using the Linpack performance runs that determines the TOP500 list of the world's fastest supercomputers.

The Blue Gene series of supercomputers is the result of a collaborative project between Argonne, Lawrence Livermore National Laboratory, and IBM. By densely packaging hundreds of thousands of low-power processors with highly efficient water-cooling, the Q design reduces floor space needs and energy use. Mira's 48 racks, containing over 768,000 processors, operate five times more efficiently than its predecessor Blue Gene/P, and roughly within the same footprint.

The Green500 list was a response to "performance-at-any-cost" computer operations that consume vast amounts of electrical power and produce so much heat that large cooling facilities must be constructed to ensure proper performance. Green500 promotes awareness to performance metrics other than speed, and offers lists to encourage supercomputing stakeholders to drive energy efficiency as a first-order design constraint.

Green Computing for Green Research

Mira will enable a broad range of research efforts aimed at finding or advancing solutions to the world's energy challenges, while itself being environmentally sustainable. As just one example, Argonne computational scientist Mihai Anitescu is developing advanced optimization methods for the U.S. electrical power grid in order to increase its reliability and reduce its carbon footprint. The operation and planning of the grid with these requirements involve an unprecedented amount of uncertainty in supply and demand brought on by the high variability of wind, solar, and other renewable power sources. The project aims to solve optimization problems of unprecedented size (tens of billions of variables and constraints), which will also help to advance the field of mathematical optimization.

"Mira will allow research teams to run larger-scale simulations and tackle larger problems in energy-critical areas,"

Michael Papka, director,
Argonne Leadership Computing Facility

www.green500.org/about.php

New Argonne Supercomputer Ties for First Place on the Graph 500 Benchmark for Data-Intensive Computing

Two U.S. Department of Energy (DOE) IBM Blue Gene/Q systems received top billing on the annual Graph 500 list, announced June 19 at the International Supercomputing Conference, outscoring the closest competitor by a performance factor of seven.

Mira, located at the Argonne Leadership Computing Facility (ALCF), and Sequoia, located at Lawrence Livermore National Laboratory, each achieved a score of over 3,500 GTEPS (giga traversed edges per second). The Graph 500 benchmark evaluates machine performance while running data-intensive analytic applications and is a measure of the machine's communications capabilities and computational power.

Graph 500 began in 2010 as a complementary benchmark to the widely observed TOP500 benchmark that evaluates a machine's ability to deliver very high performance to dense linear algebra (compute-intensive floating-point) calculations. Argonne's Mira was designed principally to deliver very high performance to such calculations since they are highly correlated to science and engineering applications.

Computers are routinely used to solve small graph problems (mapping optimal routes for a fleet of delivery trucks, for instance). Graph problems can become intractable, however, as they are scaled up to ever larger datasets (the difference between estimating the effect of climate on the economy of cities in a small region versus that of the entire country).

"Having machines that do well on these sorts of calculations will allow this very useful category of computing techniques to be applied to ever more areas," said Kalyan Kumaran, who manages ALCF's applications performance engineering team. "The range of problems that Mira can tackle is much wider than large-scale scientific simulation."

Vesta, Mira's testing and development rack, placed 6th on the Graph 500 list. Mira's predecessor at Argonne, an IBM Blue Gene/P, placed 16th.

Researchers Test Drive Blue Gene/Q at "Leap to Petascale" Workshop

Researchers eager to get their code ported to Blue Gene/Q converged on Argonne for the Leap to Petascale workshop held May 22-25, 2012, and got early access to ALCF's test and development rack of Q.

Geared especially for projects that are already scaled to multiple racks of Blue Gene/P, this year's Leap to Petascale focused on issues of scaling to Mira, the ALCF's 10-petaflops Blue Gene/Q system. A large portion of the workshop was devoted to hands-on tuning of applications with one-on-one assistance from ALCF's team of experts, plus access to the facility's Blue Gene/Q test and development rack. Special system reservations provided researchers with the opportunity for full-scale runs.

View all the presentations from the Leap to Petascale workshop, including talks on hardware and architecture, compilers, communication libraries, debuggers and more.

Photo credit: Neal Conrad

Like Us On Facebook

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