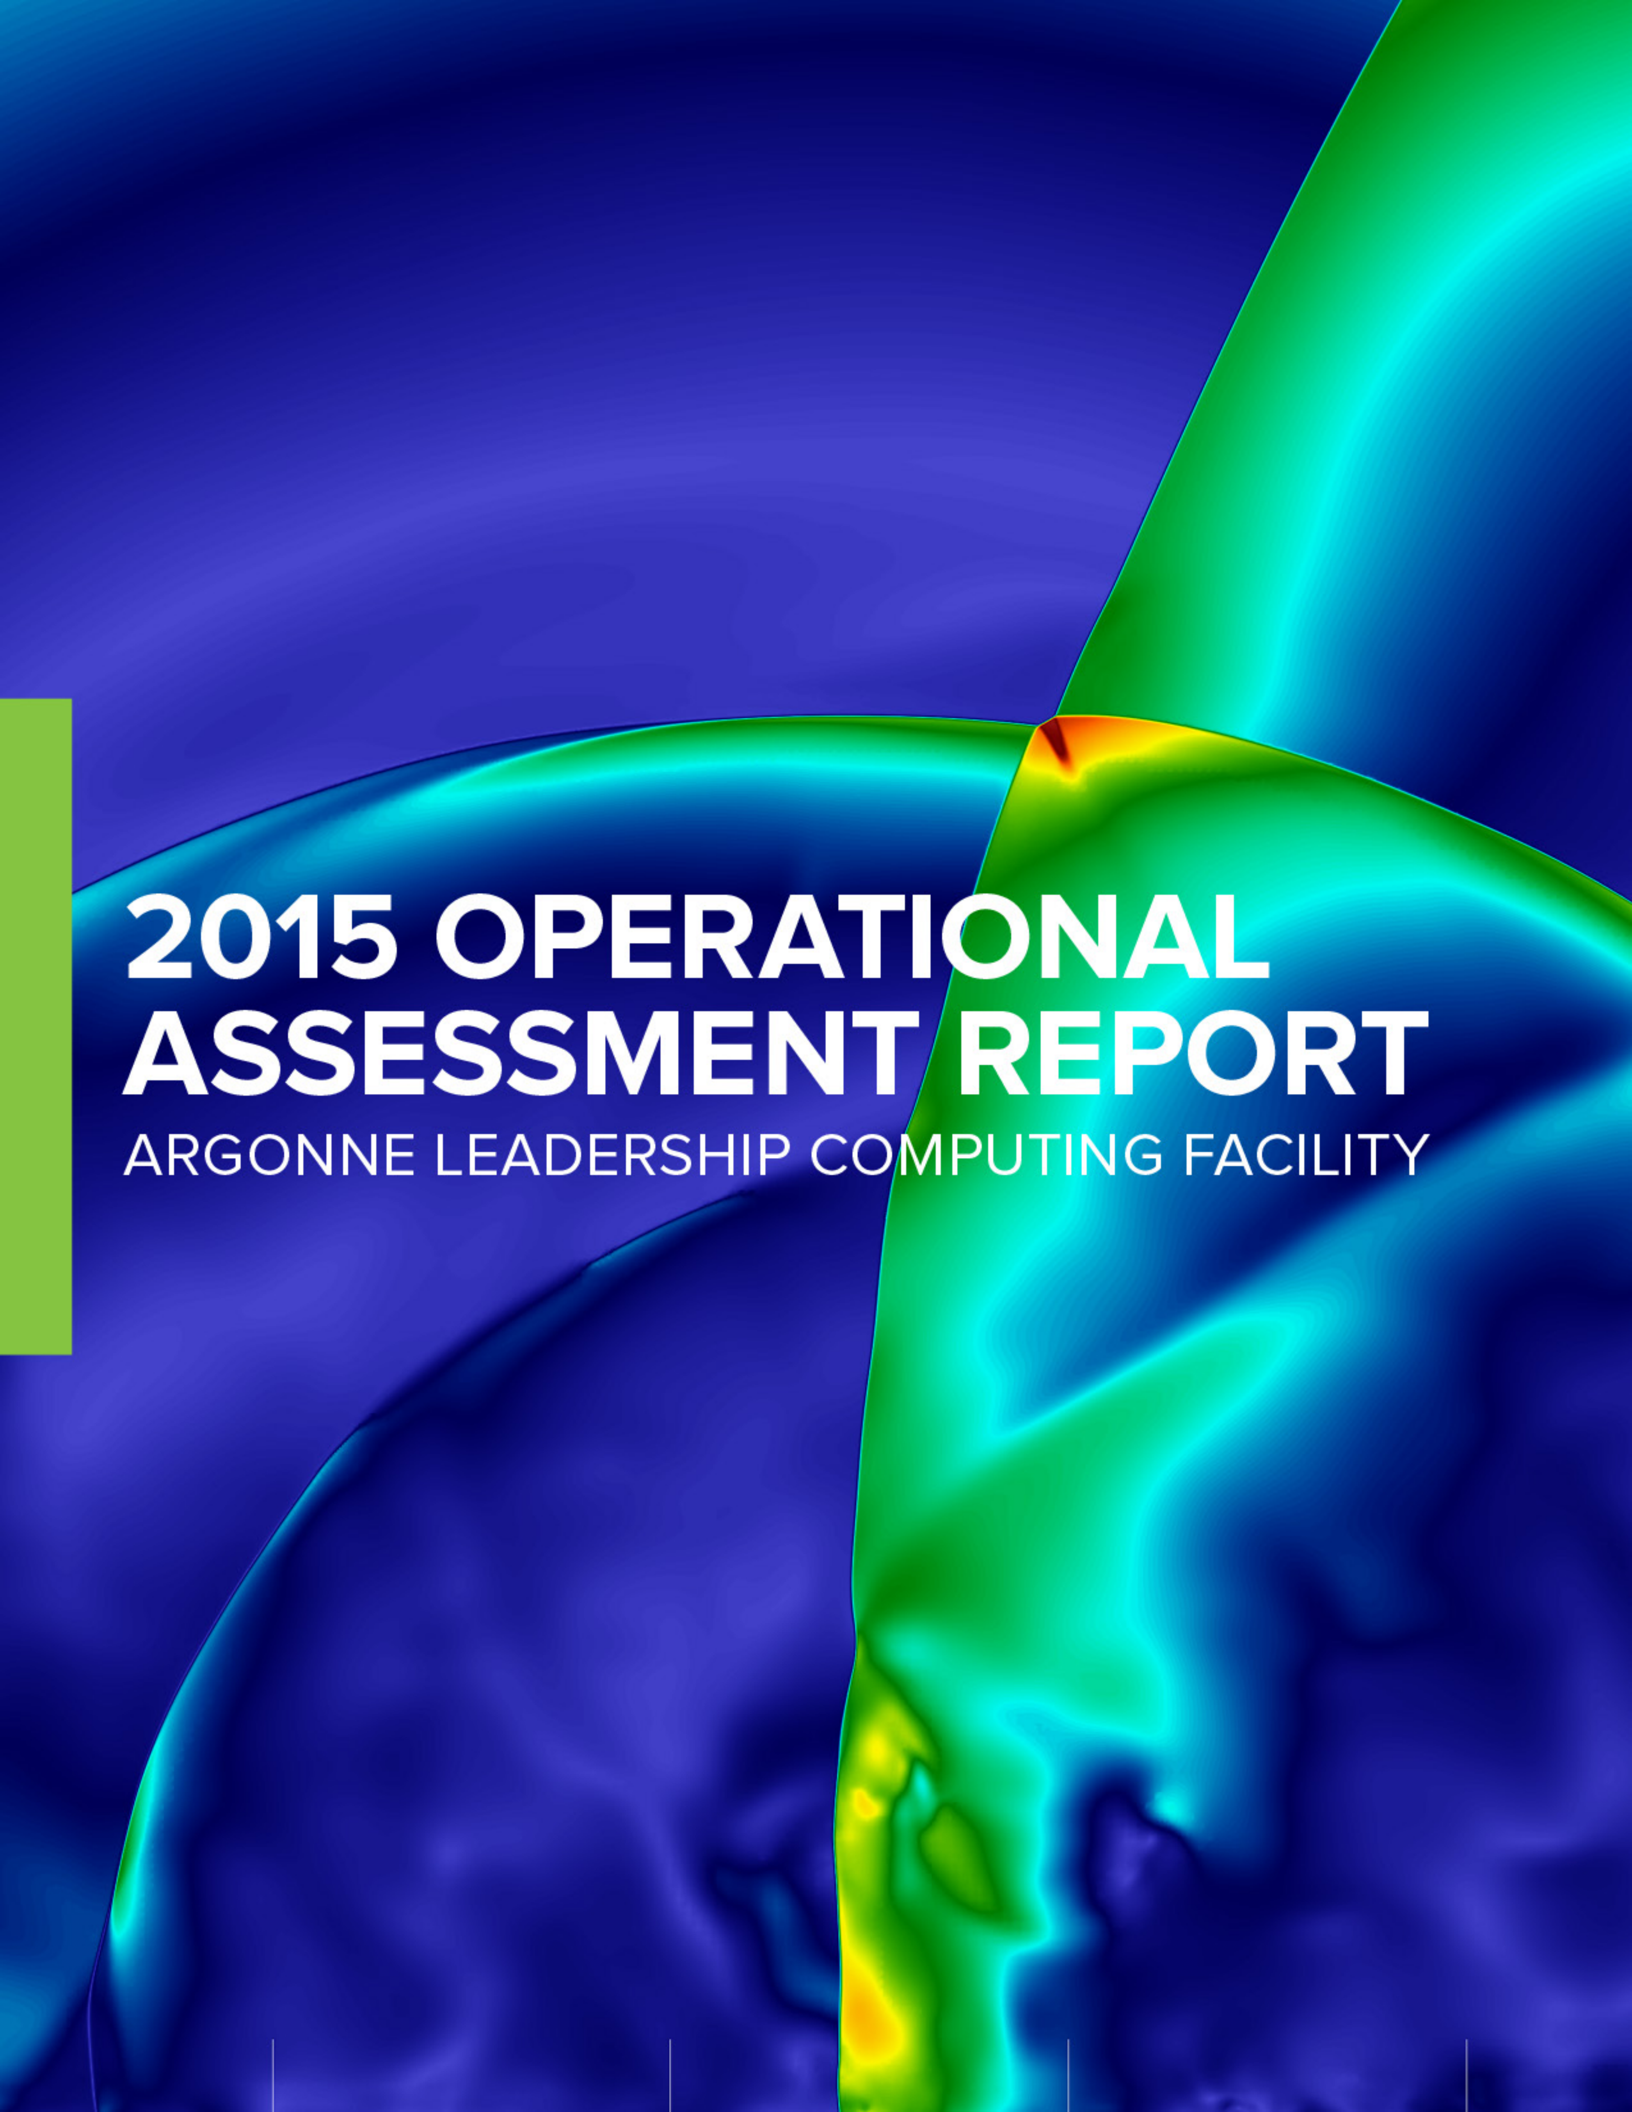

2015 OPERATIONAL ASSESSMENT REPORT

ARGONNE LEADERSHIP COMPUTING FACILITY

On the cover:

A snapshot from a visualization that depicts weak ignition behind a reflected shock in a $2H_2+O_2$ mixture at initially atmospheric pressure.

2014-2016 INCITE Project: First-Principles Simulations of High-Speed Combustion and Detonation
PI: Alexei Khokhlov, The University of Chicago
Image: Charles Bacon, Marta García, and Joseph A. Insley, Argonne National Laboratory; Alexei Khokhlov, The University of Chicago; Joanna Austin and Andrew Knisely, University of Illinois at Urbana-Champaign.

Contents

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Figures

Tables

Executive Summary

The U.S. Department of Energy's (DOE's) Advanced Scientific Computing Research (ASCR) program founded the Leadership Computing Facility over a decade ago to lead the world in open scientific computing. It was a huge investment in the nation's scientific and technological future, inspired by a growing demand for capability computing and its impact on science and engineering.

The Leadership Computing Facility operates two world-class centers, at Argonne National Laboratory and at Oak Ridge National Laboratory, and deploys two diverse petascale machines that are 10 to 100 times more powerful than systems typically available for open scientific research. Strategically, it ranks among the top U.S. major scientific facilities delivering impactful science, and the work being done at these centers helps inform policy decisions and advance innovation in far-reaching topics such as energy assurance, ecological sustainability, and global security.

These petascale machines run fully allocated, seven days a week, 52 weeks a year—much to the credit of their world-class technical staffs. From an operational standpoint, the consistently high level of service that these centers provide, and the great science that they produce, justify their continued existence to the Office of Science and to Congress.

This Operational Assessment Report describes how the Argonne Leadership Computing Facility (ALCF) met or exceeded every one of its goals as an advanced scientific computing center.
In 2015, ALCF's primary resource, Mira, delivered 4 billion core-hours to 37 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects and 1.6 billion core-hours to ASCR Leadership Computing Challenge (ALCC) projects (21 projects for the 2014–2015 ALCC year; 24 projects for the 2015–2016 ALCC year), and supported more than 190 Director's Discretionary projects that ranged from near-time computing needs to preparation work for other programs. Furthermore, Mira had its best year yet in terms of overall availability (96.3 percent), scheduled availability (99.2 percent), and utilization (89.2 percent).

Moreover, ALCF's user community published 164 papers in high-quality, peer-reviewed journals. The ALCF outreach team held several user workshops and tutorials, and ALCF supported the third highly successful Argonne Training Program on Extreme-Scale Computing to train the next generation of computational scientists.

As the Leadership Computing Facility prepares to roll out its massively scaled-up next-generation systems at Argonne and Oak Ridge (Argonne's future system, Aurora, is expected to deliver 180 petaflops), the teams at both centers are already busy working to improve critical software and to develop tools that will work at these unprecedented scales. When the pre-exascale systems arrive, the Leadership Computing Facility will once again stand ready to deliver science on Day One.

Table ES.1 Summary of the Target and Actual Data for the Previous Year (2014) Metrics

Area	Metric	2015 Targets	2015 Actuals
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.6/5.0
	User Survey – **Response Rate**	25%	45.9%
	% User Problems Addressed Within Three Working Days	80%	95.3%
Business Results	Mira Overall Availability	90%	96.3%
	Mira Scheduled Availability	90%	99.2%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072 – 786,432 cores)	40%	73.4%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144 – 786,432 cores)	10%	31.0%

Section 1. User Support Results

Are the processes for supporting the customers, resolving problems, and Outreach effective?

ALCF Response

ALCF has established processes in place for effectively supporting customers, resolving problems, and performing outreach. The 2015 survey measures satisfaction, user support, and problem resolution. It thereby provides input to ALCF about where improvements can be made (Table 1.1). The sections below document ALCF events and processes, the effectiveness of those processes, and what improvements to those processes were implemented during calendar year (CY) 2015.

Table 1.1 All 2015 User Support Metrics and Results[1]

		2014 Actual	2015 Target	2015 Actual
Number Surveyed		1,432	N/A	990
Number of Respondents (Response Rate)		430 (30.0%)	25.0%	454 (45.9%)
Overall Satisfaction	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.7
	Std Dev	0.7	N/A	0.8
Problem Resolution	Mean	4.5	3.5	4.6
	Variance	0.5	N/A	0.7
	Std Dev	0.7	N/A	0.8
User Support	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.7
	Std Dev	0.5	N/A	0.8
		2014 Actual	**2015 Target**	**2015 Actual**
% User Problems Addressed within Three Working Days[2]		96.0%	80.0%	95.3%

1 In September 2015, all ASCR facilities adopted a new definition of a facility user based on guidance from the Department of Energy (DOE). Under the new definition, a user must have logged into an ALCF resource during a given time period. The new criterion was used for the 2015 user survey and thereby decreases the number of "users" from the 2014 survey. The statistics shown in Table 1.1 for the two years are based on different populations, and should not be used for comparative purposes.

2 The population represented in this metric is tickets received from all users. This is a larger population than represented by the survey results.

Survey Approach

ALCF contracted with web survey hosting and consulting company Inquisium (formerly known as Cvent) to manage the 2015 user survey. The team incorporated lessons learned from previous surveys and internal feedback from various ALCF teams, ALCF leadership, the User Advisory Council, and DOE's Advanced Scientific Computing Research (ASCR) program. The result was a streamlined survey, improved questions, and a representative user response to the survey.

1.1 User Support Metrics

In September 2015, all ASCR facilities adopted a new user definition based on guidance provided by the DOE Office of Science. A total of 990 individual ALCF users met the new user definition and were invited to complete a user survey. Of those users, 454 responded, for a 45.9 percent response rate—excellent compared to generally accepted standards for survey response rates. ALCF surpassed all targets for the survey metrics.

Table 1.2 displays responses categorized by allocation program. While Director's Discretionary and INCITE users each reported higher average Overall Satisfaction than ALCC users, the results are not statistically significant. Other metrics are comparable, in that the variations are statistically insignificant.

Table 1.2 2015 User Survey Results by Allocation Program

2015 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		370	180	550	440	990
Number of Respondents		185	84	269	185	454
Response Rate		50.0%	46.7%	48.9%	42.0%	45.9%
Overall Satisfaction	Mean	4.5	4.3	4.4	4.5	4.5
	Variance	0.6	0.8	0.7	0.6	0.7
	Std Dev	0.8	0.9	0.9	0.8	0.8
User Support	Mean	4.6	4.4	4.5	4.5	4.5
	Variance	0.5	0.9	0.7	1	0.7
	Std Dev	0.7	0.9	0.8	1.0	0.8
Problem Resolution	Mean	4.7	4.5	4.6	4.5	4.6
	Variance	0.5	0.8	0.7	0.8	0.7
	Std Dev	0.7	0.9	0.8	0.9	0.8
All Questions	Mean	4.6	4.4	4.5	4.5	4.5
	Variance	0.5	0.9	0.7	0.9	0.7
	Std Dev	0.7	0.9	0.8	0.9	0.8

As Table 1.3 shows, in 2015 ALCF exceeded the Overall Satisfaction and User Support targets.

Table 1.3 2014 and 2015 User Support Metrics

Survey Area	2014 Target	2014 Actual	2015 Target	2015 Actual
Overall Satisfaction Rating	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0
Average of User Support Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0

1.2 Problem Resolution Metrics

As shown in Table 1.4, ALCF exceeded the target set for the percentage of problem tickets addressed in three days or less. ALCF defines a ticket as "addressed" once the following is true: a staff member has accepted the ticket; the problem has been identified; the user has received a notification; and the staff member is either working on or has found a solution.

Table 1.4 Tickets Addressed Metric

	2014 Target	2014 Actual	2015 Target	2015 Actual
% User Problems Addressed within Three Working Days[3]	80.0%	96.0%	80.0%	95.3%
Average of Problem Resolution Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.6/5.0

1.3 User Support and Outreach

1.3.1 Tier 1 Support

Phone and E-mail Support

ALCF answered 5,962 support tickets in 2015. The largest number of these tickets continued to fall under the Accounts category (see Table 1.5).

In 2015, the Automated E-mail Responses category had the largest reduction in tickets. This category encompasses both bounce messages from users due to retired email addresses on file at ALCF and out-of-office auto-responses that are triggered by reports sent from ALCF. The number of tickets in this category was high in 2014 due to an internal security audit and an internal mail routing issue. Both of these issues were addressed in late 2014.

ALCF observed an increase in tickets in the Accounts category. Looking deeper at the issue uncovered an increase in the number of account renewals. This is to be expected as the ALCF user base matures. Additionally, there was a rise in tickets categorized as Miscellaneous, which represents a large collection of tickets that do not easily fit into other categories. The increase does not correlate with any one given event or subject.

[3] The population represented in this metric is tickets received from all users. This is a larger population than represented by the survey results.

Table 1.5 Ticket Categorization for 2014 and 2015

Category	2014	2015
Access	1,085 (17%)	1,070 (18%)
Accounts	2,289 (36%)	2,532 (42%)
Allocations	690 (11%)	633 (11%)
Applications Software	323 (5%)	236 (4%)
Automated E-mail Responses	1,135 (18%)	553 (9%)
Data Transfer	55 (1%)	23 (0%)
I/O and Storage	153 (2%)	182 (3%)
Miscellaneous	155 (2%)	215 (4%)
Quota Management	44 (1%)	60 (1%)
System	487 (8%)	448 (8%)
Visualization	5 (0%)	10 (0%)
TOTAL TICKETS	6,421	5,962

Using Argonne's CIS Division for Authentication

ALCF worked with Argonne's Computing and Information Systems (CIS) Division and Acquia to integrate the laboratory's central authentication services with the ALCF website. This allows staff to use single sign-on (SSO) with their Argonne credentials to log into the website to edit and publish content. This adds extra security to the ALCF website by using federated logins.

ALCF upgraded and moved the CRYPTOCard authentication service, which is used as the mechanism to access ALCF compute and storage resources, to a new hosting environment hosted by CIS. Users were not impacted by the migration.

Enabling INCITE Computational Readiness Review with Confluence

ALCF staff worked closely with the INCITE Program Manager and her staff to support the INCITE computational readiness review process using Confluence, Argonne's newly adopted wiki technology. The previous wiki technology was being retired by Argonne's IT service. The team ensured that features critical to the review process were preserved when using the new wiki technology.

Website Feature Enhancements

ALCF staff made several feature enhancements on the ALCF website including a dynamically generated software and libraries page that provides users with accurate and near real-time data about new packages installed on the systems. Among the new features is a visualization of Mira usage on the home page that shows the jobs currently running on the machine and the percentage of Mira requested. This real-time pie chart includes information about the individual jobs on mouse-over. Another feature is a new machine status page with summary information about currently running jobs, jobs in the queue, and the number of upcoming machine reservations.

Revamping Trouble Ticket Categorization

ALCF staff has reviewed the existing categories used to sort trouble tickets and worked closely with individual groups within the division to propose a new list of categories that provides more specific options for classifying the tickets. For example, rather than using a broad category like Systems, the new list includes more detailed options such as File System and Network. The new categories are designed to improve the analysis that each group can perform on user trouble tickets and gain insights into user issues. This list will be used in beta mode with the Early Science Program support queue within the trouble ticket software to assess the merits of the new system before integrating it into ALCF's production environment.

1.3.2 Application Support

Facilitating Complex Workflows for JCESR's Electrolyte Genome Project

Led by Argonne, the Joint Center for Energy Storage Research (JCESR) is a DOE research partnership that integrates government, academic, and industrial researchers from multiple disciplines aimed at overcoming critical scientific and technical barriers to create breakthrough energy storage technology. One of the center's projects, the Electrolyte Genome, uses high-performance computers to sift through thousands of potential electrolyte materials to identify promising candidates for further research and development.

To perform the computations for this project, JCESR researchers are using an implementation of a Materials Project (https://www.materialsproject.org) virtual environment (MPenv) at ALCF. This MPenv has a front-end consisting of a python-based FireWorks workflow management system running Q-Chem's implementation of density functional theory (DFT), and a back-end consisting of a mongo database at NERSC which is populated with the computational results from ALCF via Fireworks. ALCF computational scientists Paul Coffman and William Scullin worked closely with the JCESR team to help them run the Q-Chem DFT implementation accurately and efficiently on ALCF systems. Work included identifying and fixing bugs in the application exposed by the Blue Gene/Q system and eliminating performance and scalability bottlenecks in I/O, the messaging layer, and on-node computation at the application and system levels. They achieved additional parallelism by employing MPI and OpenMP threading, enabling the application to efficiently scale to larger atomic systems.

Additionally, the ALCF team optimized performance by leveraging aspects of the Blue Gene/Q architecture, such as a distributed shared-counter implementation using the parallel active message interface (PAMI) layer, the use of transactional memory within OpenMP, and the use of the symmetric multi-processing (SMP) version of IBM's ESSL (Engineering and Scientific Subroutine Library). With these enhancements in place, the Q-Chem application performance was improved for Blue Gene/Q systems, going from 100x inefficiency (relative to results from NERSC's Edison supercomputer) to around 3x efficiency. Most of the improvements are being integrated into the Q-Chem code base with applicability to other current and future architectures.

The JCESR workflow also required high throughput of relatively small jobs without substantial modification to their FireWorks framework. The ALCF team helped ensure the efficient use of the computing systems by providing queue adaptors for the Cobalt scheduler and project-specific documentation. To further address their needs and better support users with similar high-throughput and ensemble workloads, ALCF staff developed a version of Cobalt that runs as a job within Cobalt, thereby simplifying sub-block partition management and improving the reliability and ease-of-use of the scheduler.

The ALCF contributions are helping the JCESR project evaluate vast libraries of electrolyte materials to inform and accelerate experimental efforts. JCESR scientists were able to successfully screen an initial pilot of 6,000 small molecules and now have the ability to virtually screen larger molecules thanks in part to the computing infrastructure developed by this collaboration.

Optimizing Quantum ESPRESSO

An accurate and detailed understanding of the microscopic structure of liquid water is important to a number of fields, ranging from biology and biochemistry to energy storage and electrochemistry. As part of the 2014–2015 ALCC project "Ion Solvation, Catalytic Interfaces, and Extreme Aqueous Environments: An *Ab Initio* Study of Liquid Water," principal investigator Robert A. DiStasio of Cornell University used Mira to carry out *ab initio* molecular dynamics (AIMD) simulations of liquid water and ionic solutions.

To improve the performance of the Quantum ESPRESSO code on Mira, ALCF computational scientist Álvaro Vázquez-Mayagoitia implemented an optimized multithreaded version of the code's numerical solution of a 3D Poisson equation, a demanding task in the evaluation of the exact exchange contribution in hybrid functionals related to quantum chemistry calculations. With this addition to the code and fine-tuning of parallel communications, the calculation time was reduced by up to 40 percent. This has enabled simulations that are providing detailed knowledge of the microscopic structure and equilibrium properties of liquid water and aqueous ionic solutions. The code now can be applied more efficiently to many types of periodic simulation cells, including other studies of complex condensed-phase systems.

Using Genetic Algorithms for Sustainable Energy Research

The quest for clean sustainable energy is driving the development of emerging technologies, such as organic and hybrid organic-inorganic solar cells. The functionality and efficiency of these devices often are determined by interactions at the interface between two materials. With the 2014–2015 ALCC project "Computational Design of Interfaces for Photovoltaics," principal investigator Noa Marom of Tulane University used Mira to conduct large-scale, massively parallel first-principles quantum mechanical and molecular dynamics simulations to probe the physical attributes of these critical interfaces.

ALCF computational scientist Álvaro Vázquez-Mayagoitia was instrumental in helping the team maximize their allocation time on Mira. He optimized FHI-aims, an all-electron, full-potential electronic structure code, for Mira using OpenMP directives and enabling safe multithreading of linear algebra libraries, which resulted in a 30 percent reduction in the run time. Vázquez-Mayagoitia also implemented a new parallelization scheme to allow linear scaling of GAtor, a versatile genetic algorithm package for structure prediction and design of molecular crystals.

These code improvements have helped the researchers to advance their first-principles and multiscale simulations, revealing details of these interfaces that are difficult to resolve experimentally.

Optimizing ACME and CESM Climate Modeling Codes

Multiple research teams use Mira for climate modeling projects aimed at better understanding changing climate conditions and their potential impacts. For the INCITE project "CESM Century-Scale Climate Experiments with a High-Resolution Atmosphere," principal investigator Warren Washington from the National Center for Atmospheric Research leads an effort to run sets of climate change simulations using the latest release of the Community Earth System Model (CESM). Mark Taylor of Sandia National Laboratories leads another INCITE project, "Accelerated Climate Modeling for Energy (ACME)," that employs the ACME climate modeling code, built using the same code base as CESM.

To advance the simulation work of these projects, ALCF computational scientist Paul Coffman helped overcome performance issues in the common parallel I/O (PIO) software component used by CESM and ACME, which comprised a significant portion of the simulation time on Mira. In collaboration with researchers from Argonne's Mathematics and Computer Science Division, Coffman helped identify and address the performance bottlenecks, driving solutions within the PIO code and down through the PnetCDF and MPICH MPI-I/O (ROMIO) libraries of the software stack.

This work included implementing a one-sided RDMA collective I/O aggregation algorithm within ROMIO, resolving memory utilization and performance issues that existed with the aggregation algorithm. This work has been made available to other applications using MPICH, as it was contributed and accepted into the MPICH open-source code base. Additionally, Coffman made enhancements to the MPICH PAMI device to provide further performance improvements on Mira. With these optimizations in place, INCITE researchers benefitted from I/O performance improvements of several orders of magnitude. Figure 1.1 depicts the increase in I/O bandwidth of the PIO library used by the ACME and CESM climate simulations when writing calculated data decompositions of various densities.

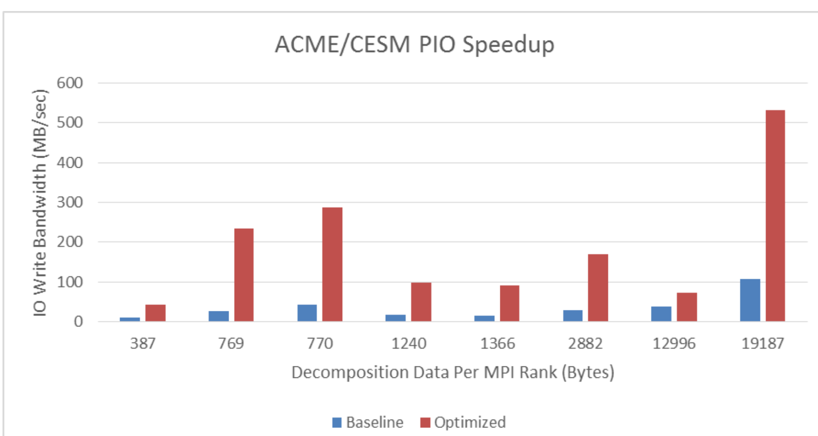

Figure 1.1 ACME/CESM PIO speedup for data decompositions of various densities.

Using Mira Boot Camp for Performance Improvements

The annual Mira Performance Boot Camp enables attendees to work directly with ALCF staff to tune scientific applications for Mira and demonstrate the use of leadership-class compute systems. At this year's event, ALCF catalysts and performance engineers applied their skills and expertise to facilitate significant improvements to the attendees' codes and make the enhancements available to the broader Mira user community where possible. Accomplishments from this year's event include:

- With the assistance from ALCF performance engineers, researchers from Missouri University of Science and Technology improved their code and I/O performance in preparation for an INCITE proposal submission. The group adopted ALCF-recommended compiler options to achieve the improvements. For I/O performance improvements, ALCF staff assisted by adjusting an environment variable being passed into the executable. By using HPC Toolkit to profile their code and identify performance bottlenecks, the group obtained a 9-times speedup over their baseline performance. This team's work is part of a larger effort aimed at improving supersonic aircraft design.
- ALCF staff and vendor experts collaborated to eliminate a bottleneck in an application used by a University of California–based INCITE team researching the evolution and present dynamical states of galaxies, stars, and other celestial bodies. While scaling beyond 131,072 cores on Mira, the team encountered a crash. Through the use of debugging tools ALCF staff identified an invalid value being passed into a decomposition parameter. By fixing code initialization problems and improving their use of MPI, the application was able to scale to 262,144 cores on Mira.
- Argonne's Virtual Engine Research Institute and Fuels Initiative (VERIFI) is a multidisciplinary team of scientists and engineers leveraging the breadth of the laboratory's state-of-the-art resources (including ALCF's leadership-class systems) to aid industry in next-generation engine design. At this year's Boot Camp, the VERIFI team compiled the latest version of the CFD software CONVERGE (2.3) and identified and addressed a bug related to the writing of the restart file. In addition, ALCF staff later resolved a hanging issue identified during the workshop. Altogether, the improvements allowed the code to use twice as many cores on Mira as before, jumping from 4,096-core runs to 8,192-core runs.
- Researchers at the University of Köln and TU Bergakademie Freiberg work on EXASTEEL, a project aimed at creating tools for simulating high-strength steels on exascale systems. The researchers had previously scaled their code to full-machine runs on the German JUQUEEN Blue Gene/Q machine, but the code was initially crashing during large run attempts at the Boot Camp. Using bgq_stack and DDT, ALCF staff helped the team identify the error in the source code preventing them from scaling beyond 262,144 cores. The project team was then able to apply a fix and scale the runs to the full Mira machine (786,432 cores).

1.3.3 Resource Support

ALCF Network Upgrade

ALCF has undertaken a major networking upgrade, consisting of a Brocade BR-MLX-40GX4-M, 4-Port 40 Gigabit Ethernet (GbE) module installed in the ALCF core MLXe-32 router. This module is connected via 4x40 GbE to two Mellanox SX1710, 36-Port, 10/40/56 GbE switches with software defined networking (SDN) capability. These two switches will be configured to use multi-chassis link aggregation (MLAG) for redundancy and aggregation of the four 40 GbE uplinks. The SX1710 are, in turn, connected to the tiled display visualization wall switch, a Mellanox SX1024 using 2x56 GbE. The Data Transfer Node (DTN) Testbed and Petrel v2 upgrade are connected directly to the SX1710 switches.

Several DOE-funded projects are using this new 40/56 GbE-based network architecture, including:

- The DTN Testbed being used by the "Technologies and Tools for Synthesis of Source-to-Sink High-Performance Flows" project, which aims to develop a data transfer advisor that will provide recommendation on transport protocols, congestion control algorithms, number of parallel streams, and buffer size settings.
- The RAMSES file transfer predictive modeling project, which will develop models that can predict file transfer performance on 10 GbE, 40 GbE, and 100 GbE connections.
- The National Science Foundation-funded EPSON network reservation project, which aims to develop a network reservation module for GridFTP. This module will interface with ESnet OSCARS to reserve one or more paths for large-scale data transfers. They are also building a multi-path module for GridFTP that will utilize independent paths for multiple transport streams.
- The joint Globus and ALCF Petrel v2 Data Management and Sharing Pilot, which will utilize the additional network capacity to achieve near 100 Gbps file transfers.
- Finally, this upgrade will allow ALCF Operations staff to test and tune DTN configurations for the upcoming Theta cluster, which will utilize 40 GbE-connected servers.

Improved Visualization Capabilities

In June, ALCF launched Cooley, a new visualization and analysis cluster with nearly eight times the memory capacity of the facility's previous system, Tukey. This significant memory boost, along with Cooley's state-of-the-art hardware, is helping users to better analyze and explore the massive datasets that result from their simulations on Mira. Cooley will enhance many capabilities for users, including *in-situ* analysis, volume-rendered visualizations, meshing complex geometries, and uncertainty quantification analysis.

Cooley is with equipped 126 compute nodes, each with two Intel Xeon E5-2620 Haswell 2.4 GHz 6-core processors with 384 gigabytes of RAM, and an NVIDIA Tesla K80 graphics processing unit (GPU) with 24 gigabytes of memory. The system has a peak performance of 223 teraflops, an aggregate RAM of more than 48 terabytes, and an aggregate GPU memory of more than 3 terabytes. By contrast, Tukey had a peak performance of 99 teraflops and an aggregate RAM of 6 terabytes. Access to Cooley is provided by two login nodes, which provide compilation and job submission capabilities. ALCF's Cobalt scheduler provides job scheduling.

ALCF staff are demonstrating Cooley's capabilities to enable INCITE principal investigator Alexei Khokhlov to visualize data from his project. The data for each time step, in silo (visualization) format, is approximately 2.1 TB. There are about a dozen variables, loaded only one at a time. With very high-resolution grids, some of the resulting isosurfaces generated can be very large. ALCF has been using 60 Cooley nodes to work with this data. The analysis runs out of memory on 30 Cooley nodes, so it would have been beyond the capabilities of Tukey.

Figure 1.2 shows the visualization generated of weak ignition behind a reflected shock in a 2H2 + O2 mixture at initially atmospheric pressure. The figure shows a 3D pressure field in a tube with a square cross-section of 5 cm × 5 cm. The end wall of the tube is on the right. The reflected bifurcated shock is on the left and is moving to the left. The weak ignition took place approximately 60 microseconds after the shock reflection. The deflagration-to-detonation transition (DDT) happened several microseconds later. The tube cross-section is filled with acoustic (pressure) perturbations generated by the bifurcated reflected shock. Hot spots, which caused the ignition, were generated when the reflected shock strength was modulated by the sound waves. This mechanism represents a new discovery. Computations were performed using a very high-resolution simulation on Mira and visualization generated on Cooley.

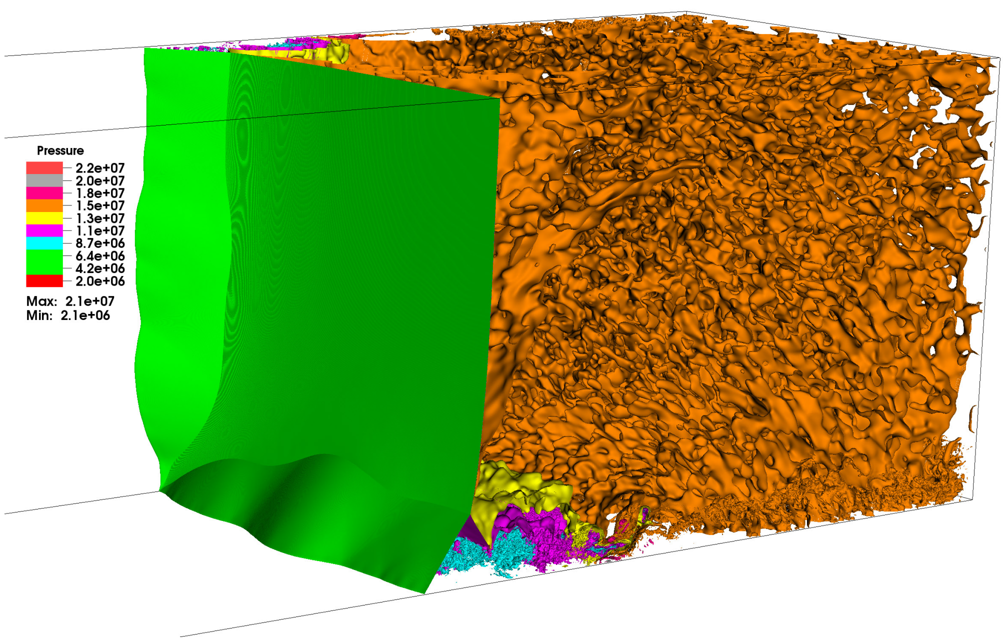

Figure 1.2 Visualization of first-principles simulations of high-speed combustion and detonation. Image: Charles Bacon, Marta García, and Joseph A. Insley, Argonne National Laboratory; Alexei Khokhlov, The University of Chicago.

Migration from LTO-4 Tapes to LTO-6 Tapes

As part of a storage infrastructure upgrade aimed at reducing the amount of time users have to spend on data management and I/O efforts, ALCF took advantage of new tape densities, replacing its LTO4 tape drives with LTO6 tape drives.

ALCF maintains a data archive and retrieval system based on HPSS, which is currently configured with disk and tape tiers. The disk tier has a capacity of 1.68 PB on a DataDirect Network SFA12K storage array. By default, all archived data is initially written to the disk tier. The upgraded tape tier consists of two SpectraLogic T950 robotic tape libraries equipped with 24 LTO6 tape drives and 10,000+ slots in each library. The total native storage capacity is about 50 PB, 125 PB compressed. Archived data is migrated to the tape tier at regular intervals and then deleted from the disk tier to create space for future archives.

The migration from LTO4 (capacity 800 GB uncompressed, 120 MB/s speed, or 1.6 TB, 240 MB/s compressed) to LTO6 (capacity 2.5 TB uncompressed, 160 MB/s speed, or 6.25 TB, 400 MB/s compressed) more than tripled ALCF's archival storage capacity. In addition, the effort has improved HPSS read/write performance when accessing files from tape.

1.3.4 Outreach Efforts

1.3.4.1 General Outreach

User Advisory Council

ALCF's User Advisory Council (UAC) provides valuable input on key technology upgrades, messaging and communication, and user-centric metrics. ALCF is grateful for the time, effort, and contributions provided by this advisory body. The seven-member UAC represents all three allocation programs (INCITE, ALCC, and Director's Discretionary). Meetings are held monthly.

ALCF staff presented software workflows to the UAC in March 2015. The agenda included a rough definition of software workflows, the common workflows used at ALCF, how these workflows manifest themselves in tools used by the UAC, and a discussion of how ALCF and the larger HPC community can help. ALCF also used the UAC as a sounding board for communication strategies related to DOE user data requests. In addition, the UAC helped ALCF by reviewing the changes to its user survey, with some UAC members providing a final end-to-end beta test.

All of the input helped ALCF to provide more useful and efficient service to its user community.

Industry Outreach

ALCF is actively engaged with industry to promote the use ALCF computing resources, ALCF training events, and opportunities to collaborate with ALCF staff. ALCF Industry Outreach Lead David Martin manages interactions with industry to help ensure that the appropriate people and resources are engaged. ALCF also acts as a gateway to other Argonne resources, linking domain expertise, computational science skills, and computing resources.

One example is the work with Brewer Science, a small business that develops and manufactures innovative materials, processes, and equipment for microelectronics and related industries. ALCF organized a day-long workshop for Brewer Science that included staff from ALCF, Argonne's Center for Nanoscale Materials, Argonne's Chemical Sciences and Engineering Division, Argonne's Laboratory Computing Resource Center (LCRC), and the University of Chicago's Institute for Molecular Engineering. This resulted in several current and planned projects with Brewer Science that will increase their competitiveness while contributing to open scientific inquiry.

Scientific Software Architecture for Portability and Performance

The DOE High-Performance Computing Operational Review (HPCOR) on Scientific Software Architecture for Portability and Performance brought together application developers, computing facilities, vendors, and library and tool developers to identify best practices for scientific software architecture, specifically to increase portability and performance. The event focused on best practices and their relevance in the coming decade. Participants were asked to review their existing and planned scientific software development efforts to identify successes, failures, and approaches that might be broadly applicable. A document with best practice recommendations for scientific software developers trying to use diverse architectures will be published as a result of this event.

DOE Exascale Requirements Reviews

ALCF is collaborating with DOE ASCR facilities, OLCF and NERSC, to hold a series of Exascale Requirements Reviews to determine the mission-critical computational science objectives for each of the six DOE Office of Science program offices through 2025. These workshops bring together key domain scientists and computational science experts to identify the requirements for developing an exascale ecosystem that includes computation, data analysis, software, workflows, HPC services, and other features. The reviews for High Energy Physics and Basic Energy Sciences were held in 2015, with the four remaining reviews planned for 2016. ALCF is coordinating and editing the reports for all six of the reviews.

1.3.4.2 Workshops and Webinars

ALCF conducted workshops and webinars to support the efforts of users and their project teams (Table 1.6). The workshops are highly rated by those attending, as evidenced by feedback received in the annual user survey. ALCF also collaborates with peer DOE institutions to develop training opportunities, explore key technologies, and share best practices that improve the user experience.

Table 1.6 2015 Workshops and Webinars

Event	Description	Dates (2015)
Getting Started Videoconferences	Small-group videoconferences providing new user training during the ramp-up periods for ALCC and INCITE allocation awards.	January/July/ August/December
INCITE Proposal Writing Webinars	Webinars designed to help attendees prepare effective proposals for INCITE.	April/May/June
Mira Performance Boot Camp	Tutorials on scaling and performance tuning codes for projects applying for 2015 INCITE awards.	May

Table 1.6 2015 Workshops and Webinars (Cont.)

Event	Description	Dates (2015)
Ensemble Jobs Videoconferences	Tutorial targeted at users whose workloads include multiple small jobs (<8K nodes) that are suitable to run concurrently.	July/ September
ATPESC	Training program on extreme-scale computing targeted at students and postdocs.	August

Getting Started Videoconferences

The ALCF Getting Started program continued to evolve in its second year using the interactive virtual videoconference format. With the convenient online format, new users from around the globe are able to log in remotely to learn about ALCF services and resources, obtain details about the IBM Blue Gene/Q architecture, and receive guided assistance in porting and tuning applications on Mira.

This year, ALCF used the online format to benefit some users registered for the Mira Performance Boot Camp. The intention was to identify registrants who had not used Mira and therefore may not be ready for the training event, which is designed for intermediate and advanced users. By engaging with six such registrants ahead of Mira Boot Camp, five of them decided to opt out of the event and instead attended a pre-Boot Camp Getting Started Videoconference tailored for beginners.

2016 INCITE Proposal Writing Webinar

In preparation for the open call for proposals, the INCITE program office conducts an INCITE Proposal Writing Webinar. These webinars took place on April 9, May 27, and June 1, hosting 29, 60, and 22 participants, respectively.

Free and open to the public, the webinars provided both prospective and returning users the opportunity to get specific answers to questions about writing an effective INCITE proposal. The INCITE Program Manager and representatives from each of the two Leadership Computing Facilities presented at the event.

Mira Performance Boot Camp

The annual Mira Performance Boot Camp once again drew new and experienced supercomputer users from around the globe. Now in its seventh year and a cornerstone of ALCF's user outreach program, the three-day on-site Boot Camp provides a timely opportunity for the community to tap into the expertise of assembled ALCF staff and invited guests for help ramping up their code's scalability as they prepare to submit a proposal for an INCITE award.

The bulk of this year's event was devoted to hands-on, one-on-one tuning of applications. In addition, ALCF experts spoke on topics of interest, including Blue Gene/Q architecture, ensemble jobs, parallel I/O, and data analysis. Guest speakers from tool and debugger vendors provided information and individualized assistance to attendees.

Reservation queues created specifically for the event gave the 39 participants quick, uninterrupted access to ALCF resources, allowing them to run 835 jobs and to use more than

18.8 million core-hours as they diagnosed code issues and tweaked performance. This year, with expert assistance and newly acquired knowledge, several groups were able to complete full-machine runs on Mira (786,432 cores) and generate plots to incorporate into their INCITE proposals.

Ensemble Jobs Videoconferences

In response to a number of small jobs being launched by some facility users, ALCF staff designed and executed an ensemble training video workshop. By sending a special invite to users with specific job profiles on the system, ALCF targeted projects and users who stood to benefit from this training. In a small-group videoconference setting limited to six to eight participants, ALCF staff worked directly with users and trained them in ensemble job techniques.

ATPESC 2015

For two weeks in August, a group of 65 students and early career researchers took up residence at Pheasant Run Resort's Gallery Hall in St. Charles, Ill., for an arduous training program designed to teach them the key skills and tools needed to efficiently use leading-edge supercomputers.

Packed with technical lectures, hands-on exercises, and dinner talks, the Argonne Training Program on Extreme-Scale Computing (ATPESC) addresses all aspects of high-performance computing with a curriculum that evolves each year to emphasize particular areas of interest. This year, the organizers incorporated more hands-on sessions and placed increased focus on the importance of performance portability across diverse computing architectures. The content was organized around seven core program tracks:

- Hardware Architectures
- Programming Models and Languages
- Numerical Algorithms and FASTMath
- Community Codes/Software Engineering
- Visualization and Data Analysis
- Toolkits and Frameworks
- Data Intensive Computing and I/O

In addition, participants gained access to hundreds of thousands of cores of computing power on some of today's most powerful supercomputing resources, including ALCF's Mira and Vesta systems, Oak Ridge Leadership Computing Facility's Titan system, and National Energy Research Scientific Computing Center's Edison system.

Selected from a highly competitive field of 170 applicants, program participants included doctoral students, postdocs, and computational scientists who have used at least one high-performance computing system for a reasonably complex application and are engaged in or planning to conduct research on large-scale computers. Their research interests span the disciplines that benefit from supercomputers, such as physics, chemistry, materials science, computational fluid dynamics, climate modeling, and biology.

As a result of this event, ALCF recorded 78 presentations and posted them to the Argonne training YouTube channel. ALCF staff also drafted and administered eight participation evaluation surveys and shared the results of these with track leads.

Train the Trainers

In an effort to continually improve user support and services, ALCF conducts multiple "Train the Trainers" workshops throughout the year on technologies that might benefit a larger audience. This year's workshops include:

- HPC Toolkit Staff Training – April 1–2
- Allinea DDT Staff Training – May 18
- TAU Staff Training – February 10

Theta Early Science Virtual Events

ALCF hosted several videoconferences to help prepare researchers for the Theta Early Science Program. These included a "kick-off" event on August 25 and a series of presentations about the program to the project teams on September 9.

1.3.4.3 Community Outreach

Introduce a Girl to Engineering Day

In February, several ALCF staff members participated in Argonne's Introduce a Girl to Engineering Day (IGED). The annual event gives eighth-grade girls a unique opportunity to discover STEM careers alongside Argonne scientists, engineers, and staff. The program consisted of a full day of interaction with mentors, an engineering expo, hands-on activities, and tours of ALCF and other Argonne facilities. ALCF participants included Liza Booker, Laural Briggs, Janet Jaseckas, Jini Ramprakash, Laura Ratcliff, and Emily Shemon.

Coding Camp

To contribute to regional computer science education, ALCF staff members Joseph Insley, Michael Papka, and Silvio Rizzi participated in a four-day coding camp. A total of 42 high-school students from the Chicago area attended the camp, which focused on problem-solving skills, developing code in Python, and a look at careers at a national lab or as a STEM professional.

INCREASE Workshop

In October, Argonne partnered with the Interdisciplinary Consortium for Research and Education and Access in Science and Engineering (INCREASE) for a two-day workshop aimed at increasing the participation in and diversity of the user base at the laboratory's scientific user facilities, with a focus on engaging minority-serving institutions. As part of the event, ALCF held a brainstorming session to identify and better understand the barriers to gaining access to ASCR user facilities and how to overcome them. The partnership between INCREASE and Argonne has established a foundation for growing the next generation of STEM professionals and HPC users at minority-serving institutions.

Hour of Code

As part of Code.org's annual Hour of Code event in December, eight ALCF staff members visited Chicago-area schools to spark interest in computer science and coding. Working with classrooms ranging from kindergarten to high school, the ALCF volunteers led a variety of activities designed to demystify code and show that anybody can learn the basics. The global outreach campaign aims to expand participation and increase diversity in computer science. ALCF participants included Liza Booker, Lisa Childers, Kevin Harms, Sunhwan Jo, David Martin, Michael Papka, Jini Ramprakash, and Khairi Reda.

1.3.5 Communications

Communications through Mailing Lists and Social Media

ALCF provided information to users through several electronic communication channels: weekly plain-text e-mails discussing user news, HTML-formatted monthly newsletters, HTML-formatted special announcements (ALCF and related organizations), intermittent social media postings, the ALCF website, and custom-tailored e-mail messages via scripts (Table 1.7). Users can opt out of the system notify and newsletter mailing lists. The target audiences for these channels are displayed in Table 1.8.

Table 1.7 2015 Primary Communication Channels

Channel Name	Description	When Used/Updated
Newsbytes	E-mail newsletter featuring science, facility news, events, and deadlines.	Monthly
Special Announcements	E-mail newsletter with information on conferences, training events, etc.—both ALCF and non-ALCF opportunities.	Ad hoc
Weekly Digest	Weekly e-mail to users to communicate enhancements to ALCF systems and software, key dates approaching, and training opportunities.	Weekly
Social Media	Social media used to repost ALCF news, and events.	Frequently
ALCF Website	An integrated information hub for user documentation, program and resources descriptions, user-centric events, feature stories about users, and related news.	Frequently
Custom E-mail Messages	Notification of machine status or facility availability typically formatted text-based per user and channel preference.	As needed

Table 1.8 2015 Target Audiences

Channel	Target Audience(s)
Newsbytes	Users, scientific communities, students, the general public
Special Announcements	Users, scientific communities, students, the general public
Weekly Digest	Current users on the systems with accounts
Social Media	Users, followers of ALCF, collaborators, students, scientific communities, the general public
ALCF Website	Users, collaborators, students, scientific communities, the general public
Custom E-mail Messages	Specific projects, user groups, principal investigators/proxies, individual users

Every month, ALCF's Newsbytes newsletter features science stories that highlight either the outcome of research carried out on ALCF resources or an advancement made by ALCF staff and researchers in the field. Critical deadlines, targeted announcements, key events, and relevant news stories are also included in the monthly publication. An example of a targeted special announcement was informing the recipients of the call for the Argonne Training Program on Extreme-Scale Computing.

Promotional Activities and Media Hits

The media team curates and publishes a list of ALCF media hits throughout the year and promotes the list on the ALCF website and in the monthly Newsbytes. In 2015, the facility posted 103 media hits to ALCF's website. In order for the team to post a media hit, it must be specifically about ALCF.

The media team uses Meltwater News public relations suite to help track media hits. This global online media monitoring company tracks articles from more than 200,000 news publications, as well as Twitter, YouTube, Facebook, and blogs. In 2015, Meltwater captured 332 mentions of the "Argonne Leadership Computing Facility" and "ALCF."

Other Publications

ALCF produced a variety of print publications used for promotion, education, and recruiting (Table 1.9). Argonne visitors receive an informational packet tailored to their particular area of interest.

Table 1.9 Publications Designed for Print in 2015

Publication	Frequency	When
INCITE Poster	Yearly	January
INCITE Brochure	Yearly	November
Fact Sheet	Yearly	November
Annual Report	Yearly	March
Science Highlights	Yearly	September
Press and Visitor Packets	As Needed	As Needed
Industry Brochure	Yearly	June

Conclusion

As a user facility, ALCF is focused on ensuring the success of all facility users and customers. During CY 2015, ALCF made website enhancements, improved the user survey, worked with users to improve application performance, upgraded computing resources, engaged in outreach activities, and enhanced communication efforts in various e-mail channels. As such, ALCF continues to help its users succeed by providing effective Tier 1 support, application support, resource support, outreach, and communications.

This page intentionally left blank.

Section 2. Business Results

Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?

ALCF Response

ALCF has exceeded the metrics target for system availability, INCITE hours delivered, and capability hours delivered. For the reportable areas, such as Mean Time to Interrupt (MTTI), Mean Time to Failure (MTTF), and system utilization, ALCF is on par with the other facilities and has demonstrated exceptional performance. To assist in meeting these objectives and to improve overall operations, ALCF tracks hardware and software failures and analyzes their impact on the user jobs and metrics as a significant part of its improvement efforts.

Table 2.1 summarizes all metrics reported in this section.

Table 2.1 Summary of All Metrics Reported in the Business Results Section

	Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	CY 2014		CY 2015	
	Target	Actual	Target	Actual
Scheduled Availability	95.0%	98.7%	90.0%	99.2%
Overall Availability	90.0%	95.7%	90.0%	96.3%
System MTTI	N/A	8.98 days	N/A	9.50 days
System MTTF	N/A	25.80 days	N/A	24.16 days
INCITE Usage	3.5B	3.9B[c]	3.5B	4.0B[e]
Total Usage	N/A	5.8B[d]	N/A	5.9B[f]
System Utilization	N/A	87.6%	N/A	89.2%
Mira INCITE Overall Capability[a]	30.0%	64.5%[c]	40.0%	73.4%[e]
Mira INCITE High Capability[b]	10.0%	33.1%[c]	10.0%	31.0%[e]

[a] Overall Capability = Jobs using ≥ 16.7 percent (8 racks, 131,072 cores) of Mira.

[b] High Capability = Jobs using ≥ 33.3 percent (16 racks, 262,144 cores) of Mira.

[c] Usage includes 15.0M core-hours from Cetus non-capability production jobs.

[d] Usage includes 15.0M core-hours from Cetus non-capability production jobs.

[e] Usage includes 9.9M core-hours from Cetus non-capability production jobs.

[f] Usage includes 30.2M core-hours from Cetus non-capability production jobs.

ALCF Resources

During CY 2015, ALCF operated one production resource, Mira. Mira is a 48K-node, 768K-core, 10 PF Blue Gene/Q with 768 TB of RAM. Mira mounts three General Parallel File System (GPFS) file systems with approximately 26.5 PB of usable space and has access to the facility-wide HPSS (high-performance storage system) tape archive. Mira has an associated visualization and analysis cluster called Cooley. ALCF operated two other Blue Gene/Q systems, Cetus and Vesta.

Cetus is a 4K-node, 64K-core Blue Gene/Q with 64 TB of RAM. Cetus shares file systems with Mira. Vesta is a 2K-node, 32K-core Blue Gene/Q with 32 TB of RAM. Vesta is an independent test and development resource and shares no resources with Mira or Cetus.

Starting in 2014, ALCF permitted select use of Cetus for INCITE projects with simulation runs that required non-traditional HPC workflows. This allowed Mira to continue to operate as designed and enabled a new class of leadership applications to be supported.

2.1 Resource Availability

Overall availability is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages. For HPC Facilities, scheduled availability is the percentage of time a designated level of resource is available to users, excluding scheduled downtime for maintenance and upgrades. To be considered a scheduled outage, the user community must be notified of the need for a maintenance event window no less than 24 hours in advance of the outage (emergency fixes). Users will be notified of regularly scheduled maintenance in advance, on a schedule that provides sufficient notification, and no less than 72 hours prior to the event, and preferably as much as seven calendar days prior. If that regularly scheduled maintenance is not needed, users will be informed of the cancellation of that maintenance event in a timely manner. Any interruption of service that does not meet the minimum notification window is categorized as an unscheduled outage. A significant event that delays a return to scheduled production will be counted as an adjacent unscheduled outage. Typically, this would be for a return to service four or more hours later than the scheduled end time. The centers have not yet agreed on a specific definition for this rare scenario.

This section reports on measures that are indicative of the stability of the system and the quality of the maintenance procedures.

2.1.1 Scheduled and 2.1.2 Overall Availability

Mira has been in full production since April 9, 2013. In consultation with the DOE Program Manager, ALCF has agreed to metrics of 90 percent overall availability and, for CY 2015, a new target of 90 percent scheduled availability (ASCR requested that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.2 summarizes the availability results.

Table 2.2 Availability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2014		CY 2015	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	95.0	98.7	90.0	99.2
Overall Availability	90.0	95.7	90.0	96.3

The remainder of this section covers significant availability losses, and responses to them, for both scheduled and overall availability data. Details on the calculations can be found in Appendix A.

Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability for the period January 1, 2015, through December 31, 2015, as annotated in Figure 2.1.

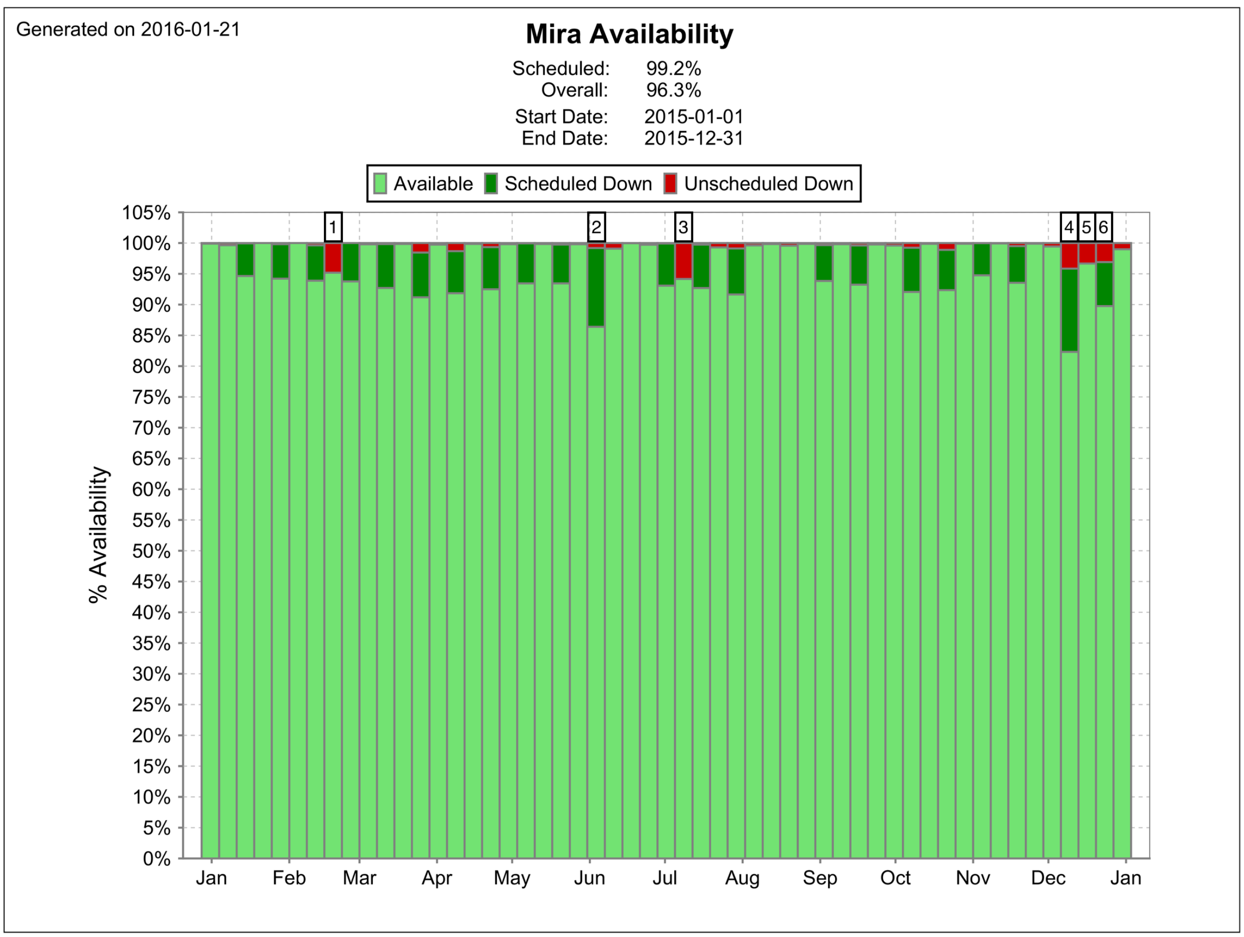

Figure 2.1 Mira Weekly Availability for CY 2015

Graph Description: Each bar in Figure 2.1 represents the average of seven days of core-hour usage. Each bar accounts for all of the time in one of three categories. The pale-green portion represents available core-hours; the darker green represents scheduled downtime for that week; and red represents unscheduled downtime. The numeric annotations are the significant losses. Each of these events is described in detail below.

Item 1: Network spine failure – February 16
A spine module on one of the QDR IB switches failed. The problem initially manifested itself as severe "GPFS waiter" issues. ALCF Operations staff was able to spot-fix the file system for a couple of days, resulting in some slow I/O; meanwhile, a root-cause analysis was begun. On February 16, Operations staff took an unscheduled outage to engage in more explicit debugging. They experienced significant issues restarting GPFS. It was noted that a majority of the issues were happening on DDN hardware connected to a particular IB switch. Staff checked the spine modules and discovered a failed module. Mellanox QDR gear considered this particular failure mode "internal" and did not log the error messages where they had been monitoring. The faulty spine was replaced and the systems brought back into operation. The Operations team subsequently added more monitoring for this situation.

Item 2: Scheduled building power maintenance – June 1
On June 1, the Building 240 Data Center engaged in scheduled maintenance of power facilities. ALCF Operations staff performed a controlled shutdown of Mira and all supporting storage and infrastructure. This provided an opportunity to test manual and automated shutdown procedures. The subsequent start-up process allowed staff to identify weaknesses in and improve documentation of both its shutdown and start-up processes. While not an unscheduled outage, this maintenance is called out because it was scheduled for a longer-than-usual window.

Item 3: HVAC cooling failure – July 6
Mid-afternoon on July 6, the HVAC system providing cooling to the Building 240 Data Center (which includes Mira's supporting systems) experienced a temporary failure, resulting in an ambient air temperature increase upwards of +15°F. ALCF Operations staff preemptively shut down air-cooled systems in order to protect them from potential heat-related damage. They were able to implement changes to procedures for shutdown and start-up that were practiced on, and improved after, the June 1 scheduled maintenance.

Item 4: Filesystem upgrades: mira-fs1, mira-home – December 8
Following an upgrade to the firmware and the GridScalar software on mira-fs1's DDN (Data Direct Networks) couplets during a planned preventive maintenance outage, a large number of page allocation errors occurred, the mira-fs1 file system became sluggish, and some mira-fs1 nodes ran out of memory. ALCF Operations staff isolated mira-fs1 and scheduled a testing period with the DDN and IBM to perform root-cause analysis. A change in the vendor-specific version of OpenFabrics Enterprise Distribution (OFED) installed during the GridScalar upgrade caused significant increase memory usage when using OFED "connected mode" on mira-fs1 storage nodes, resulting in "out of memory" errors. The memory usage only became visible due to the size and configuration of mira-fs1. At the suggestion of the vendor, a switch to using

"datagram" mode was made. After this change was implemented, all tests ran correctly and mira-fs1 ran without further problems.

Item 5: Cooling loop – December 14
On the morning of December 14, 10 cabinets of Mira shut down due to low supply temperature fault detection. Around 8:00 AM, Building Maintenance staff from Argonne's Facilities Management and Services (FMS) Division was enabling the free-cooling system in Building 528. This is a standard maintenance procedure that has been performed without issue multiple times. However, this time, one valve that was scheduled to be switched to an open position became stuck in the closed position; this was not immediately realized by FMS staff. The resulting reduction in available water led to the aforementioned low supply temperature fault detection.

Operations was able to restore seven of the 10 cabinets to service that same day (December 14). The other three cabinets presented hardware issues that required additional work on the following day (December 15) to restore to service.

To prevent a future occurrence, FMS installed new programming in the Metasys system to send alerts if the 60°F process water flow drops below 750 gallons per minute (GPM) (typical operation is in the 1,700 GPM range). After an internal investigation into this incident, FMS believes that they have identified the root cause and the indications that resulted from it.

Item 6: Scheduler draining error – December 22
On December 22, a previously unidentified error in a scheduler policy resulted in 16 cabinets of Mira being unavailable for processing. Cobalt, the scheduling software, began draining hardware in order to run a 32K node job. However, due to a small hardware failure within that block, the 32K job would have been unable to run. Normally, Cobalt will route jobs around this failed hardware, allowing smaller jobs to backfill in around the failed hardware. Due to a bug in Cobalt, however, it failed to identify the potential alternate job routing in this particular edge case. The Cobalt development team tested and deployed a bug fix in early 2016.

2.1.3 System Mean Time to Interrupt (MTTI) and 2.1.4 System Mean Time to Failure (MTTF)

MTTI = Time, on average, to any outage on the system, whether unscheduled or scheduled. Also known as MTBI (Mean Time Between Interrupt).
MTTF = Time, on average, to an unscheduled outage on the system.

ALCF MTTI and MTTF Summary
MTTI and MTTF are reportable values with no specific target. Table 2.3 summarizes the current MTTI and MTTF values.

Table 2.3 MTTI and MTTF Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2014		CY 2015	
	Target	Actual	Target	Actual
System MTTI	N/A	8.98 days	N/A	9.50 days
System MTTF	N/A	25.80 days	N/A	24.16 days

Mira currently functions on a biweekly maintenance schedule. ALCF takes the machine out of service every other Monday to perform Blue Gene driver upgrades, hardware replacements, OS upgrades, etc. Further, while Mira is out of service, ALCF uses that opportunity to perform other potentially disruptive maintenance such as facilities power and cooling work, and storage systems upgrades and patching. ALCF's biweekly maintenance schedule caps MTTI at 14 days, but does not directly impact MTTF.

2.2 Resource Utilization

The following sections discuss system allocation and usage, total system utilization percentage, and capability usage. For clarity, *usage* is defined as resources consumed in units of core-hours. *Utilization* is the percentage of the available core-hours used (i.e., a measure of how busy the system was kept when it was available).

2.2.1 Total System Utilization

Total System Utilization is the percent of time that the system's computational nodes run user jobs. No adjustment is made to exclude any user group, including staff and vendors.

Utilization is a reportable value with no specific target. A rate of 80 percent or higher is generally considered acceptable for a leadership-class system. Table 2.4 summarizes ALCF utilization results, and Figure 2.2 shows system utilization over time by program.

Table 2.4 System Utilization Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2014		CY 2015	
	Target	Actual	Target	Actual
System Utilization	N/A	87.6%	N/A	89.2%

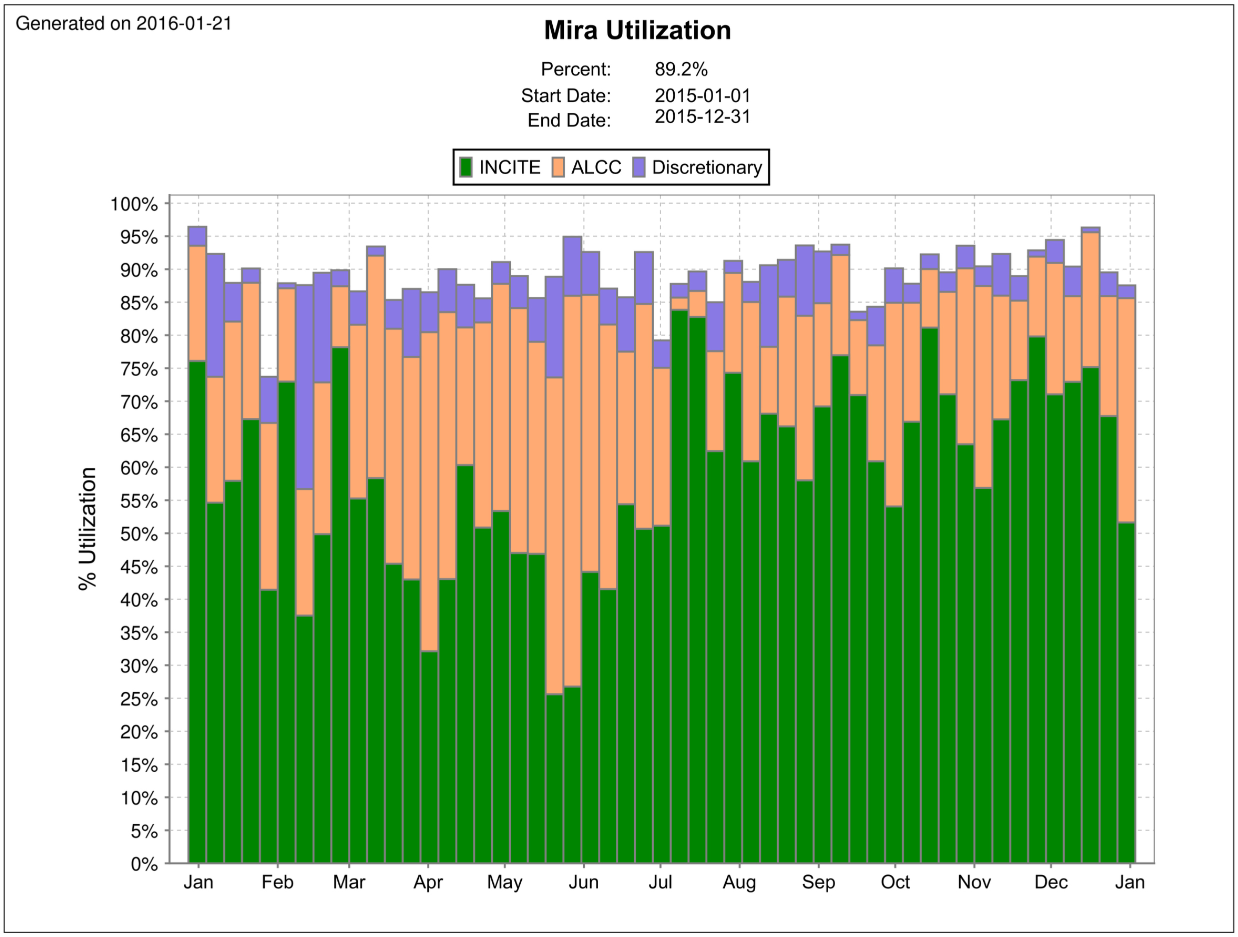

Figure 2.2 System Utilization over Time by Program

The system utilization for Mira was 89.2 percent for its 2015 production period of January 1, 2015, through December 31, 2015.

Table 2.5 shows how Mira's system hours were allocated and used by allocation source. Multiplying the theoretical hours by availability and utilization values that were agreed upon with ALCF's DOE Program Manager determines the hours available. Of the hours available, 60 percent is allocated to the INCITE program, up to 30 percent is available for the ALCC program allocations, and 10 percent is available for Director's Discretionary (DD) allocations. The ALCC program runs from July through June, so to arrive at allocated values for the calendar year, half of the hours are assigned to each year. The allocated values for the DD allocations appear higher than expected because they represent a rolling allocation. A majority of DD projects are exploratory investigations, so the time allocations are often not used in full. DD allocations are discussed in detail in Section 3.3. In CY 2015, ALCF delivered a total of 5.9 billion core-hours on Mira.

Table 2.5 Core-Hours Allocated and Used by Program

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2014			CY 2015		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	%	Core-hours	Core-hours	%
INCITE	3.5B	3.9B	67.0%	3.6B	4.0B[a]	66.7%
ALCC	1.6B	1.4B	24.9%	1.7B	1.6B[b]	27.0%
DD	1.1B	468.9M	8.1%	858.3M	373.2M[c]	6.3%

[a] Usage includes 9.9M core-hours from Cetus non-capability production jobs.

[b] Usage includes 10.4M core-hours from Cetus non-capability production jobs.

[c] Usage includes 9.9M core-hours from Cetus non-capability production jobs.

Summary: For CY 2015, the system usage and utilization values were in line with general expectations. The calculations for utilization are described in Appendix A.

2.3 Capability Utilization

The Facility shall describe the agreed upon definition of capability, the agreed metric, and the operational measures that are taken to support the metric.

On Mira, capability is defined as using greater than 16.7 percent of the machine. Historically, capability has been defined as using greater than 20 percent of the machine. However, 20 percent of Mira would be 9.6 racks, which is not a viable configuration. Hence, the Mira capability metric is split into two parts. Overall Capability requires that a minimum of 30 percent of the INCITE core-hours be run on eight racks or more (16.7 percent), and High Capability requires a minimum of 10 percent of the INCITE core-hours be run on 16 racks or more (33.3 percent). See Appendix A for more detail on the capability calculation. Table 2.6 and Figure 2.3 show that ALCF has substantially exceeded these metrics set for INCITE. Although no targets are set, data are also provided in the table for ALCC and DD projects as reference, and Figure 2.4 shows the overall distribution of job sizes over time.

Table 2.6 Capability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2014			CY 2015		
Capability Usage	Total Hours	Capability Hours	Percent Capability	Total Hours	Capability Hours	Percent Capability
INCITE Overall	3.9B	2.5B	64.5%	4.0B	2.9B	73.4%
INCITE High	3.9B	1.3B	33.1%	4.0B	1.2B	31.0%
ALCC Overall	1.4B	787.6M	54.8%	1.6B	809.9M	50.5%
ALCC High	1.4B	124.7M	8.7%	1.6B	298.5M	18.6%
Director's Discretionary Overall	468.9M	179.7M	38.3%	373.2M	202.4M	54.2%
Director's Discretionary High	468.9M	101.8M	21.7%	373.2M	127.2M	34.1%
TOTAL Overall	5.8B	3.5B	60.0%	5.9B	3.9B	66.0%
TOTAL High	5.8B	1.5B	26.1%	5.9B	1.6B	27.8%

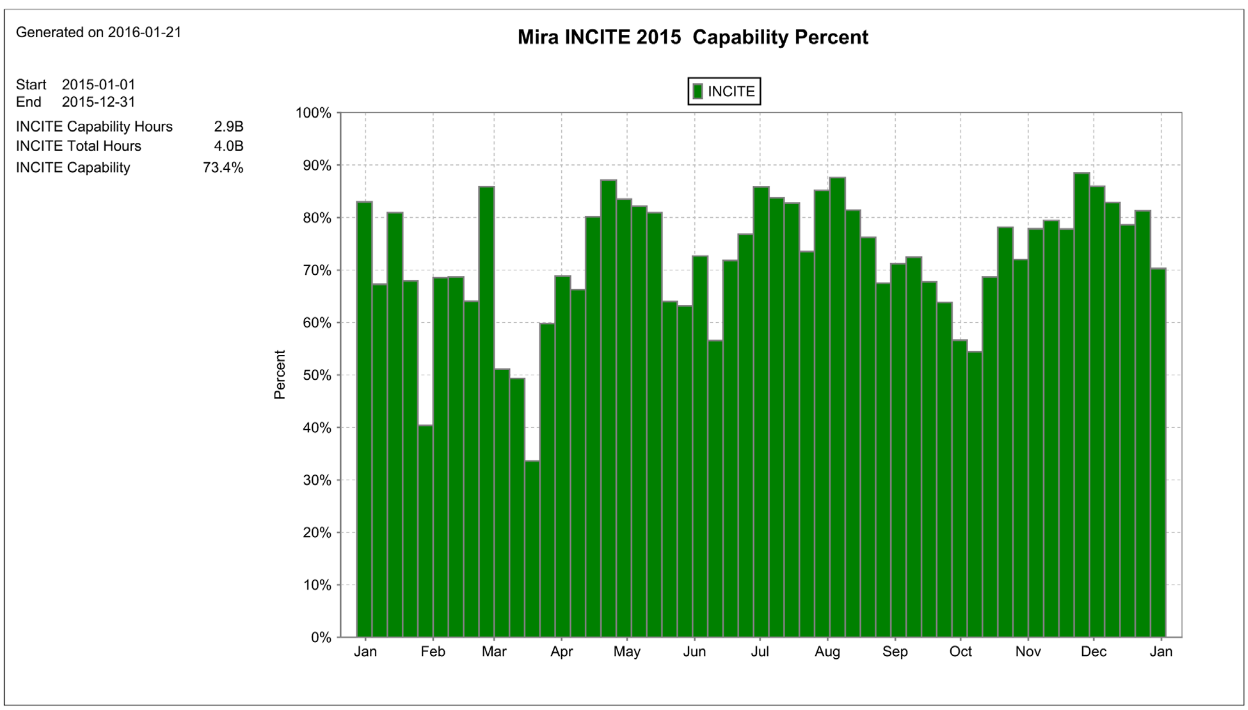

Figure 2.3 Mira INCITE Overall Capability

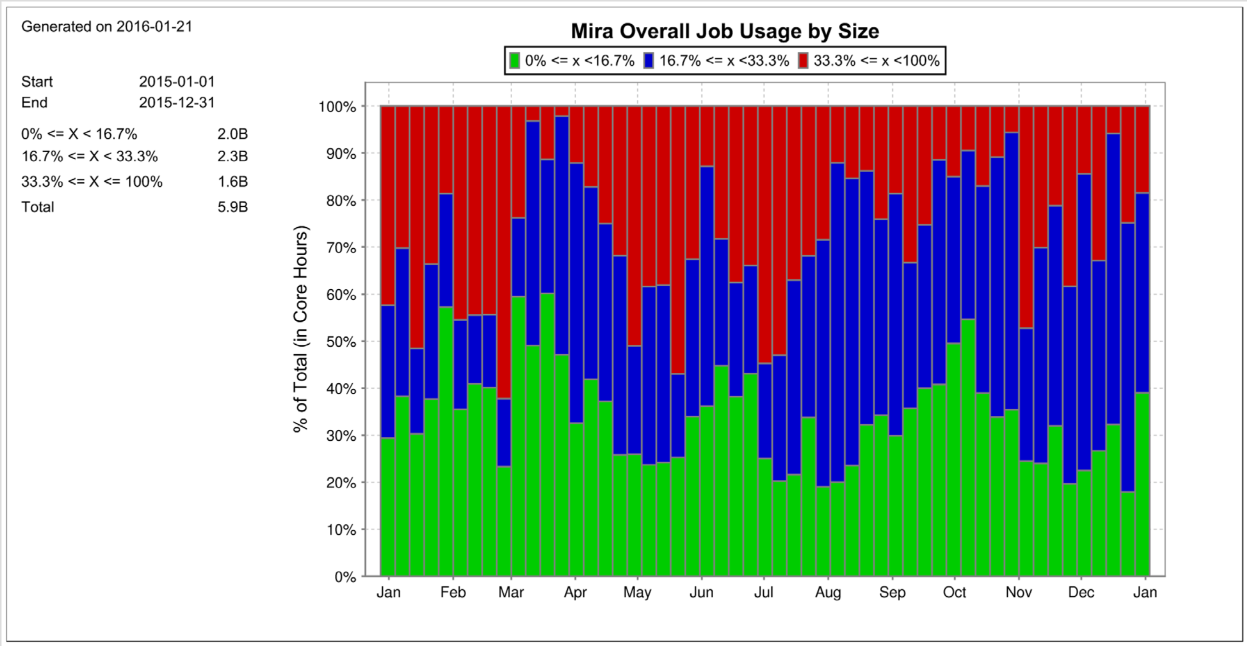

Figure 2.4 Mira Job Usage by Size

Conclusion

ALCF is maximizing the use of its HPC systems and other resources consistent with its mission. We have exceeded the metrics of system availability, INCITE hours delivered, and capability hours delivered. For the reportable areas—MTTI, MTTF, and utilization—ALCF is on par with OLCF and NERSC, and the values reported are reasonable. These measures are summarized in Table 2.7.

ALCF closely tracks hardware and software failures and their impact on user jobs and metrics. These data are used as a significant factor in the selection of troubleshooting efforts and improvement projects. In CY 2015, this regular failure analysis has driven code changes to Cobalt, ALCF's job scheduler, and has provided details to support debugging of storage system problems, as noted in Section 2.1.1 under “Explanation of Significant Availability Losses.”

Table 2.7 Summary of All Metrics Reported in the Business Results Section

	Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	CY 2014		CY 2015	
	Target	Actual	Target	Actual
Scheduled Availability	95.0%	98.7%	90.0%	99.2%
Overall Availability	90.0%	95.7%	90.0%	96.3%
System MTTI	N/A	8.98 days	N/A	9.50 days
System MTTF	N/A	25.80 days	N/A	24.16 days
INCITE Usage	3.5B	3.9B[c]	3.5B	4.0B[e]
Total Usage	N/A	5.8B[d]	N/A	5.9B[f]
System Utilization	N/A	87.6%	N/A	89.2%
Mira INCITE Overall Capability[a]	30.0%	64.5%[c]	40.0%	73.4%[e]
Mira INCITE High Capability[b]	10.0%	33.1%[c]	10.0%	31.0%[e]

[a] Overall Capability = Jobs using ≥ 16.7 percent (8 racks, 131,072 cores) of Mira.

[b] High Capability = Jobs using ≥ 33.3 percent (16 racks, 262,144 cores) of Mira.

[c] Usage includes 15.0M core-hours from Cetus non-capability production jobs.

[d] Usage includes 15.0M core-hours from Cetus non-capability production jobs.

[e] Usage includes 9.9M core-hours from Cetus non-capability production jobs.

[f] Usage includes 30.2M core-hours from Cetus non-capability production jobs.

This page intentionally left blank.

Section 3. Strategic Results

Is the Facility enabling scientific achievements consistent with the Department of Energy strategic goals?

ALCF Response

The science accomplishments of INCITE, ALCC, and Director's Discretionary (DD) projects clearly demonstrate ALCF's impact in supporting scientific breakthroughs. ALCF staff has worked effectively with individual project teams to adapt their simulation codes to run efficiently in a high-performance computing environment and has enabled scientific achievements that would not have been possible otherwise.

In this section, ALCF reports:

- Science Output
- Scientific Accomplishments
- Allocation of Facility Director's Reserve Computer Time

3.1 Science Output

The Facility tracks and reports the number of refereed publications written annually based on using (at least in part) the Facility's resources. For the LCFs, tracking is done for a period of five years following the project's use of the Facility. This number may include publications in press or accepted, but not submitted or in preparation. This is a reported number, not a metric. In addition, the Facility may report other publications where appropriate. ESnet will report an alternate measure, e.g., based on transport of experimental data.

Table 3.1 shows the breakdown by journal of refereed publications based (at least in part) on the use of ALCF resources. The Nature Journals entry refers to one publication each in *Nature Physics* and *Nature Communications.* In addition to these journals, ALCF users published in journals such as *Science, PNAS, Physical Review Letters,* and *SC '15 Proceedings*.

Table 3.1 Summary of Refereed Publications

Science	Nature Journals	PNAS	Physical Review Letters	SC '15	Total 2015 Publications
1	2	2	9	4	164

3.2 Scientific Accomplishments

The Facility highlights a modest number (top five) of significant scientific accomplishments of its users, including descriptions for each project's objective, the implications of the results achieved, the accomplishment itself, and the facility's actions or contributions that led to the accomplishment. The accomplishment slides should include the allocation, amount used, and a small bar graph indicating size of jobs.

LCFs should include tables/charts comparing time allocated to time used by projects. NERSC should include a chart summarized by SC program.

Each science highlight includes a box with a bar graph. The top line indicates the machine used, program (INCITE, ALCC, DD) and the year of the allocation. The second line lists the total core-hours allocated to the project and, in parentheses, the core-hours used. The graph shows the core-hour breakdown for each project, by the percentage of the machine used. The breakdown is based on the ALCF capability metric detailed in Section 2. The bar on the left represents runs below capability, the middle bar (where present) represents runs at the first capability threshold up to but not including the second threshold, and the bar on the right (where present) represents runs at the highest capability.

Computations for the Development of the Nanoporous Materials Genome

J. Ilja Siepmann, University of Minnesota

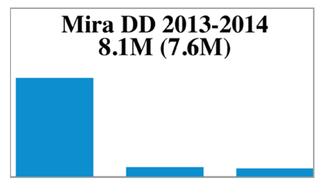

With their ability to selectively screen and promote chemical transformations, zeolites play numerous important roles in energy-relevant processes (e.g., production of fuels and chemical feedstocks from renewable resources). To date, 213 framework types have been synthesized and a sizeable collection of thermodynamically accessible zeolite structures predicted (>330,000). Searching such a large pool of candidates for optimal performance of a specific task is intractable using traditional trial-and-error methods in the lab. Supported by the DOE-funded Nanoporous Materials Genome Center, researchers at the University of Minnesota and Rice University developed computational tools to accelerate discovery and investigation of nanoporous materials for energy-relevant processes (Figure 3.1).

To overcome difficulties associated with solution non-idealities (due to hydrogen-bonding and aggregate formation) and conformational sampling (due to confined channels within zeolite), scientists employed a novel combination of grid tabulated interaction potentials and advanced Monte Carlo sampling algorithms using the MCCCS-MN software package developed by the Siepmann group. In collaboration with ALCF staff, the MCCCS-MN code was ported to Mira along with an ensemble driver to enable high-throughput calculations capable of utilizing 100 percent of the system.

Two main achievements of this DD project were the identification of separate zeolite candidates with improved selectivity of 1) ethanol over water and 2) linear vs. slightly branched alkanes. The zeolite discovered with the ability to exceed the ethanol/water azeotropic concentration in a single separation step from fermentation broths used in biofuel production has the potential to replace an energy-intensive distillation process. In the same study, scientists experimentally validated this candidate zeolite's ability to purify ethanol over a broad range of compositions. The second zeolite identified has up to two orders of magnitude better adsorption capability than current technology for linear alkanes with 18-30 carbon atoms, demonstrating the potential to advance the hydroisomerization processes used to dewax lubricants. The results of this DD project were published in *Nature Communications* in January 2015. The project was also featured as one of the University Research Highlights on DOE Office of Science website.

IMPACT: This project is providing knowledge and computational tools for high-throughput screening of nanoporous materials of significant interest for chemical, biorenewable, and petrochemical industries with the potential for significant societal benefits. Using predictive modeling for high-throughput screening of materials, the discovery and design of materials specifically tailored for energy-related applications has the potential to significantly advance technologies used today and those imagined for tomorrow. The vast amount of data generated will be made available to the public via the *Nanoporous Materials Explorer*, an app that is part of the DOE-funded *Materials Project* (www.materialsproject.org).

ALCF Contribution: ALCF computational scientist Chris Knight collaborated with the team, helping to port the MCCCS-MN code to Mira, adding OpenMP support for hybrid MPI/OpenMP parallelism, and designing and implementing an MPI-based driver for ensemble calculations.

Publication: Bai, P., M. Y. Jeon, L. Ren, C. Knight, M. W. Deem, M. Tsapatsis, and J. I. Siepmann. “Discovery of Optimal Zeolites for Challenging Separatons and Chemical Transformations Using Predictive Materials Modeling.” *Nature Communications*, January 2015, no. 5912.

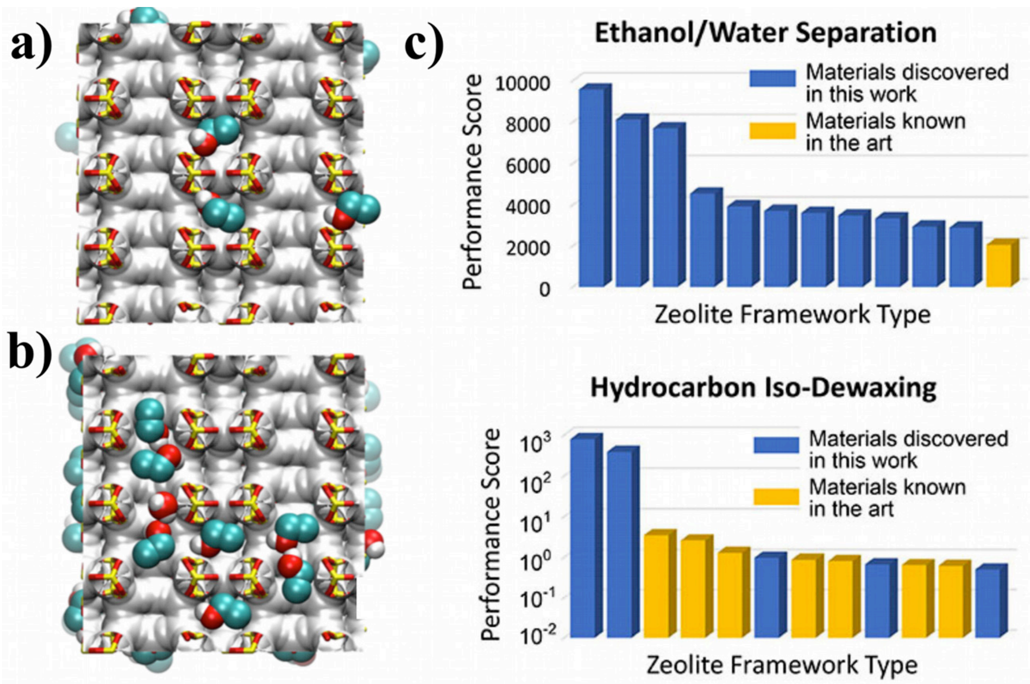

Figure 3.1 Snapshots of representative ethanol/water configurations inside the FER zeolite at low (a) and high (b) loading. *[Nat. Commun.,* 6, 5912 (2015)] Silicon, oxygen, carbon, and hydrogen atoms are shown as yellow, red, cyan, and white spheres, respectively. Isosurfaces equidistant to zeolite atoms are shown as grey. c) Performance scores from high-throughput screening for the two applications targeted in this study. Image: College of Science and Engineering at the University of Minnesota.

Reactive MD Simulations of Electrochemical Oxide Interfaces at Mesoscale

Subramanian K. R. S. Sankaranarayanan, Argonne National Laboratory

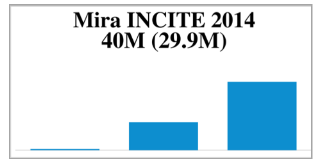

Lubricants play a critical role in dissipating heat and reducing the wear and tear of mechanical assemblies found in such everyday objects as airplanes and automobiles. Friction is the result of the kinetic energy used to slide surfaces past one another, transforming into heat as the surfaces resist that motion due to imperfections. The discovery of a lubricant eliminating essentially all friction would have profound economic and environmental implications. In a collaborative experimental/modeling investigation by a multidisciplinary team at Argonne National Laboratory, scientists identified such an ideal lubricant in the combination of three chemical forms of carbon: graphene sheets, diamond nanoparticles, and a protective coating of diamond-like carbon (DLC).

Using an INCITE award at ALCF, researchers were able to corroborate the results of an extensive experimental characterization at the Center for Nanoscale Materials (CNM) with large-scale reactive molecular simulations. Multi-million atom molecular simulations using eight racks of Mira provided the microscopic insight to understand the controlling factors for this superlubricity phenomenon, in which friction is essentially absent. Key to the success of these large simulations was a collaboration of staff at ALCF, Lawrence Berkeley National Laboratory (LBNL), and IBM to optimize the performance of the reactive simulation method achieving a two-fold speedup in the LAMMPS code.

In very much the same way that ball bearings maintain separation between surfaces and reduce friction, the diamond nanoparticles—when wrapped in graphene nanosheets—significantly reduced the friction when a large steel ball coated in DLC was slid across a silicon dioxide surface (Figure 3.2). This phenomenon is the result of graphene-wrapped nanoparticles providing transient surfaces for the DLC ball to glide across. As hypothesized experimentally, the friction coefficients computed from the simulations dropped to near-zero levels only upon formation of the graphene nanoscrolls. The molecular simulations also provided enlightenment as to the microscopic mechanism for suppression of superlubricity in an ambient humid environment being due to water inhibiting formation of the protective nanoscrolls.

An important aspect of this discovery is that superlubricity can be maintained over a range of conditions and the materials involved might very well be compatible with sustained performance in industrial applications at engineering length scales. This work is a great example of how the state-of-the-art resources available at CNM and ALCF can be utilized in tandem experimental/computational studies to discover and characterize materials that are highly beneficial to the public.

IMPACT: Friction and mechanical wear are primary modes of energy loss in moving mechanical assemblies. For example, one-third of the fuel used in automobiles is forfeited to friction, thus significant economic gains are possible with even modest improvements in the performance of lubricants. The observation of superlubricity in this work is a demonstration that this unusual, but highly desirable property can be realized over a range of conditions and at engineering length scales.

ALCF Contribution: ALCF computational scientists Wei Jiang and Nichols Romero collaborated with researchers at IBM and Lawrence Berkeley National Laboratory to optimize the Reax/C implementation in the LAMMPS molecular simulation code. A two-fold speedup over baseline performance was achieved by adding OpenMP support and MPI collectives. The optimized Reax/C code has been delivered to the LAMMPS developers and will soon be accessible to the LAMMPS user community.

Publication: Berman, D., S. A. Deshmukh, S. K. R .S. Sankaranarayanan, A. Erdemir, A. V. Sumant. "Macroscale Superlubricity Enabled by Graphene Nanoscroll Formation." *Science*, May 2015, vol. 348, 1118.

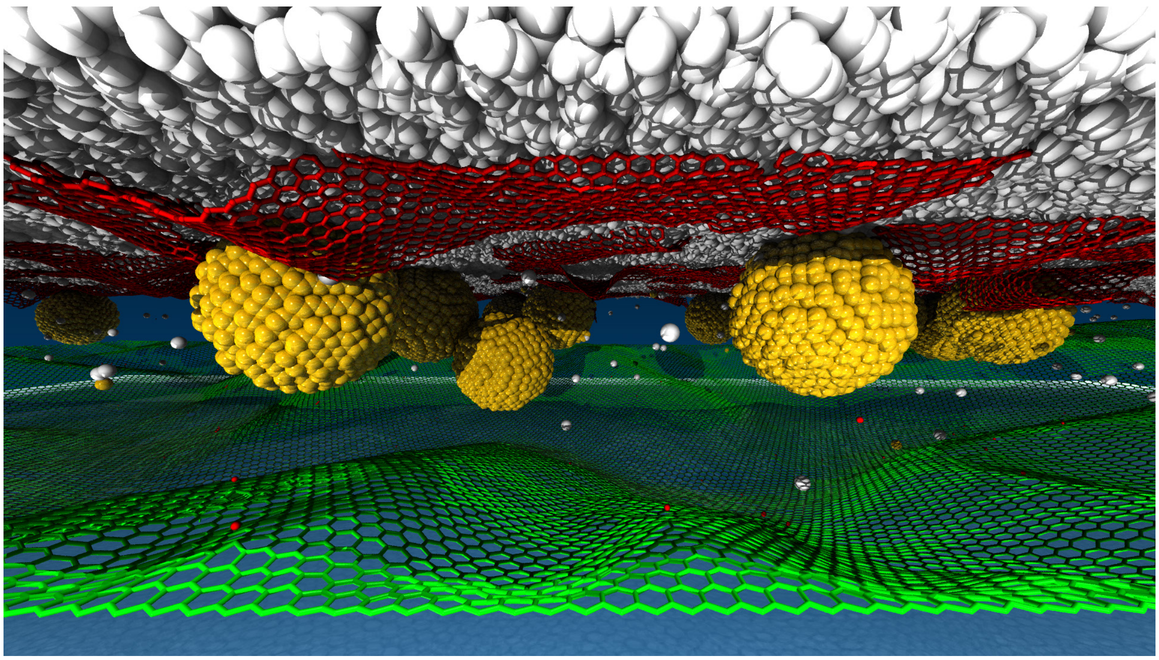

Figure 3.2 Snapshot from a large-scale molecular dynamics simulation of the initial stage of superlubricity where diamond nanoparticles (large yellow spheres) start being wrapped by patches of graphene (red sheets) to form nanoscrolls that have reduced contact area with the DLC surface (shown in white) and underlying graphene (shown in green). Image: Sanket Deshmukh, Joseph A. Insley, and Subramanian Sankaranarayanan, Argonne National Laboratory.

Simulation of Large Hadron Collider Events Using Leadership Computing

Tom LeCompte, Argonne National Laboratory

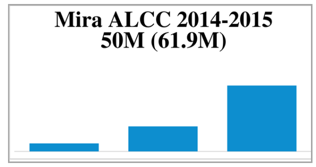

Scientists from across the world use the Large Hadron Collider (LHC) to explore the behavior of matter, energy, space, and time at the smallest scales ever probed (1/40,000 of the size of a proton). Most notably, experiments at two LHC particle detectors—ATLAS and CMS—confirmed the existence of the Higgs Boson in 2012. Simulations are key to enabling scientists to understand the response of the LHC detectors to particle collisions at the facility. However, some LHC events are so complex that it would take weeks to complete the calculations on LHC's computing grid, a system of 100,000 PC-like computers distributed all over the world. In addition, the LHC's computing needs are expected to grow by at least a factor of 10 in the next several years.

An expensive but ubiquitous background event class, W/Z+Jets, are computed using the ALPGEN program, and are particularly difficult to simulate using the current grid infrastructure (Figure 3.3). Because these events produce a neutrino for Z events or a charged lepton for W events (particles that generally escape the detector), understanding the statistics of these events is a critical dependency for all searches that involve new undetected particles (supersymmetry, dark matter, etc.). A few of these events occur during every second of detector operation, making them much more common than the hypothesized events from more-novel "new physics" interactions.

Currently, about 7 percent of the ATLAS grid activity has moved to Mira. Using HPC resources like Mira is new for this community. As HPC resources get better integrated into the LHC's workflow, a much larger fraction of simulations could eventually be shifted to supercomputers.

IMPACT: All the W/Z+Jets ALPGEN events required by ATLAS for the next two years were generated in a matter of weeks, freeing significant grid resources for other work. By improving the code's I/O performance and reducing its memory usage, the research team was able to scale ALPGEN to run on the full Mira system and help the code perform 23 times faster than it initially did. If this project were a country, it would have been the seventh-largest contributor of cycles to ATLAS's computing grid.

ALCF Contribution: Mostly through the work of ALCF researcher Tom Uram, the facility assisted with adapting the ALPGEN grid workflow to Mira, improving ALPGEN's I/O to allow it to scale to all of Mira, and reducing ALPGEN's memory footprint to allow for a configuration using 64 processes per node.

Publications: The papers that will be published based, in part, on the simulated events are awaiting the real data for comparison. This work was the subject of two CHEP presentations:

- Childers, J. T., T. D. Uram, T. J. LeCompte, M. E. Papka, and D. P. Benjamin. "Simulation of LHC Events on a Million Threads." *21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015), Journal of Physics: Conference Series* 664, 9 (2015), 092006.

- Uram, T. D., J. T. Childers, T. J. LeCompte, M. E. Papka, and D. P. Benjamin. "Achieving Production-Level Use of HEP Software at the Argonne Leadership Computing Facility." *21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015), Journal of Physics: Conference Series* 664, 6 (2015), 062063.

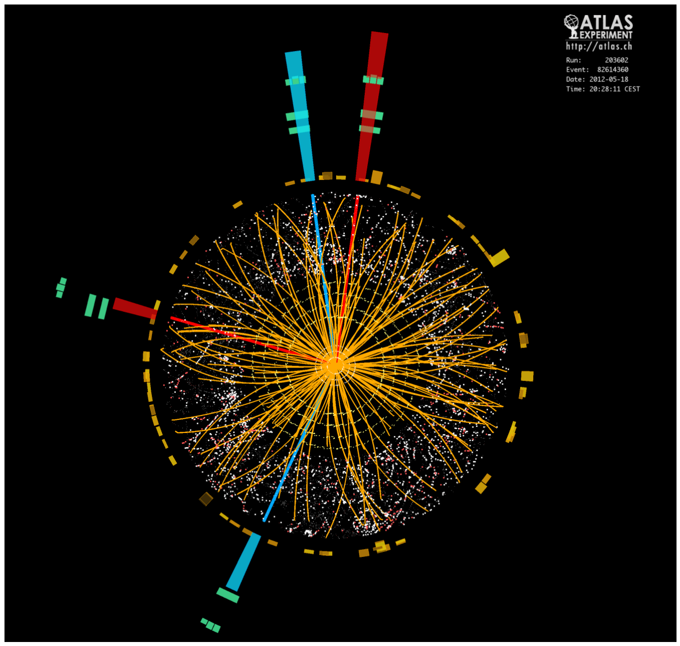

Figure 3.3 A Higgs -> ZZ* candidate event from the ATLAS detector, showing the particle tracks and jets that require simulations to interpret. Each Z goes to muons—the blue ones at 12 and 7 o'clock form on Z, and the red ones at 10 and 1 o'clock form the other. Image: ATLAS Collaboration.

SiO_2 Fracture: Chemomechanics with a Machine Learning Hybrid QM/MM Scheme

James Kermode, University of Warwick

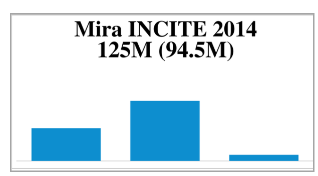

Fracturing minerals requires a tremendous amount of energy. In mining, the energy spent fracturing represents almost 85 percent of such activity. The World Business Council estimates that mining is responsible for nearly 5 percent of total human energy consumption, and annually releases millions of tons of carbon dioxide.
Understanding the physical process of grinding and crushing minerals under different environments has a technological and scientific relevance to reduce the ecological footprint of mining. The goal of this project is to study the physics behind the fractures in rocks with computational methods.

Essentially, fracturing rocks implies breaking and making new bonds. To simulate these complex microscopic processes, advanced techniques are required. First, theoretically to simulate fractures, it is necessary to use molecular dynamics (MD), but this ubiquitous simulation tool is limited to tens of picoseconds and 10^4 atoms in supercomputers. And second, accurately estimating the energy changes in breaking bonds is a many-body problem that requires quantum mechanics (QM) to be solved, which could be difficult to obtain for large systems.

A collaborative team of King's College London, the University of Warwick, and Argonne National Laboratory endeavors to advance the techniques required to understand the intricate mechanisms of fracturing materials, and overall to overcome the challenges of simulating large materials and complex physical processes at atomic resolution (Figure 3.4).

The team is integrating the methodology into a "Machine Learning on-the-Fly" (MLOTF) method that uses a multi-scale approach to predict structural changes and thermodynamic properties of amorphous materials. This methodology involves using a database of reference configurations every time an energy and force evaluation are required. The database is expanded on-the-fly with QM calculations; as the simulation progresses, fewer QM calculations are required. Bayesian statistics are applied to help guarantee optimal transferability of parameters obtained from the database.

IMPACT: The MLOTF scheme implemented through this initiative efficiently combines classical and quantum mechanics to study bond breaking of amorphous systems. It has been demonstrated that MLOTF works on the fracture of real (3D) materials such as silicon, silicon carbide, and silica, common elements in rocks. In practical tests, MLOTF has shown a speedup of a factor of around 10^3 to 10^4.

ALCF Contribution: ALCF computational scientist Álvaro Vázquez-Mayagoitia helped advance this project by debugging and porting the CP2K code and enabling an optimized version of the open-source linear algebra library ELPA for large QM simulations. In addition, ALCF staff resolved interoperability issues between their code and the sockets-mode version of CP2K. With these optimizations in place on Mira, the research team benefitted from a 10-fold increase in sampling time in their simulations.

Publication: Li, Z., J. R. Kermode, and A. De Vita. "Molecular Dynamics with On-the-Fly Machine Learning of Quantum-Mechanical Forces." *Physical Review Letters*, December 2015, vol. 114, no. 9.

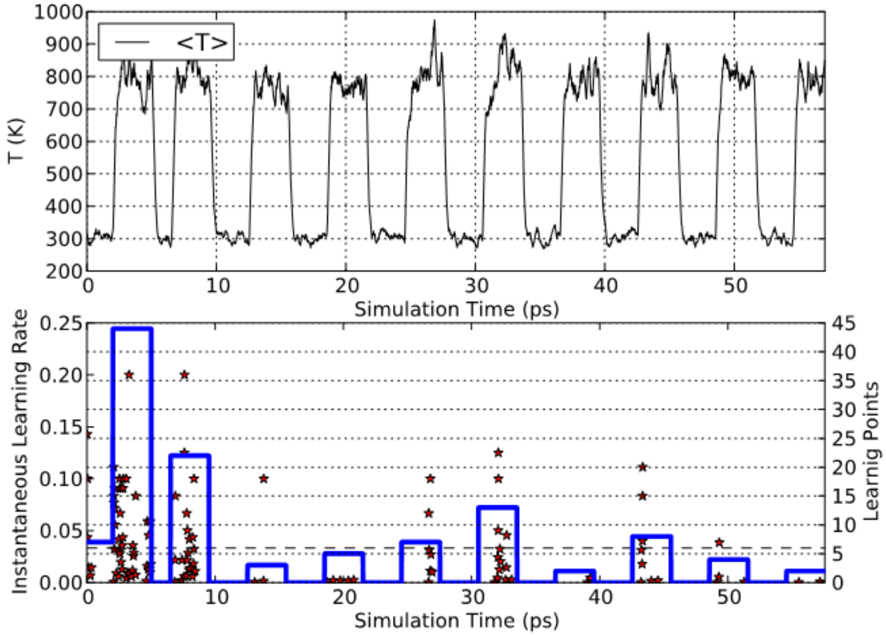

Figure 3.4 Simulation temperatures oscillate from 300K to 800K. Numbering of learning points (red stars) decreases as the simulation time progresses.

High-Fidelity Simulation of Tokamak Edge Plasma Transport

C. S. Chang, Princeton Plasma Physics Laboratory

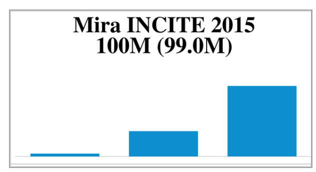

While magnetic fusion promises the possibility of clean energy for millions of years, many difficult scientific problems first must be understood and overcome before commercially profitable fusion reactors can be realized. One fundamental and difficult physics problem is understanding and predicting the self-generated electrical current density in steep edge pedestal plasmas. The self-generating current, known as "bootstrap" current, yields a more efficient reactor because it helps create a barrier called the "steep edge plasma pedestal" between the highly energized core plasma and the edge plasma. Maintaining a stable barrier at the edge pedestal is a critical challenge in experimental fusion.

High-current density is critical for the formation of the edge pedestal plasma, which leads to higher fusion efficiency. However, an overly strong edge current density can make the edge plasma catastrophically unstable, sending damaging plasma energy to the material wall. The edge pedestal plasma exists in a nonlinearly self-organized state among scale-inseparable multiscale kinetic physics in non-thermal equilibrium. The only way to understand and predict its behavior appears to be through extreme-scale simulations of the kinetic physics.

The XGC gyrokinetic code used has a unique capability in studying the nonlinear kinetic physics in realistic edge geometry. In this study, researchers used one-third of Mira to perform simulations that shed light on the fundamental physics of the self-generated electrical current in steep edge pedestal plasmas. The team discovered that the physics of the self-generated current is different in the edge plasma from the 40-year-old description in textbooks.

In tokamak geometry, the magnetic field is stronger at the inner part of the torus than at the outer part. Particles traveling at a slower speed along the magnetic field cannot penetrate the higher magnetic field region and are trapped at the outboard part of the torus. These are called "trapped" particles, while other particles passing through the high magnetic field region are called "passing" particles.

While textbooks suggest that the bootstrap current is carried mostly by the passing electrons, the INCITE team's discovery shows that trapped particles carry most of the bootstrap current at the edge (Figure 3.5, left). Many simulations for different geometric and plasma conditions led to creation of a new unified analytical formula that describes the edge current accurately, as well as the core current (Figure 3.5, right). It is expected that this new formula will be widely used by the international fusion community.

IMPACT: The present study offers resolution of a long-standing question on the predictability of the self-generated electrical current in a steep edge pedestal of a fusion reactor. The edge electrical current is critically important in stabilizing catastrophic edge instabilities and in setting the edge pedestal height that enables efficient fusion production in the burning core.

ALCF Contribution: Earlier this year, the research team experienced an issue with MPI collective operations on integers with large message sizes on Mira. Large XGC runs were freezing irreproducibly on Mira, and the simulations could not continue. The XGC team worked with ALCF computational scientist Timothy Williams, who devised a workaround method and resolved the issue.

Publication: Hager, R. and C. S. Chang. "Bootstrap Current in the Edge Pedestal of Tokamak Plasmas." To be submitted to *Physics of Plasmas* (2015).

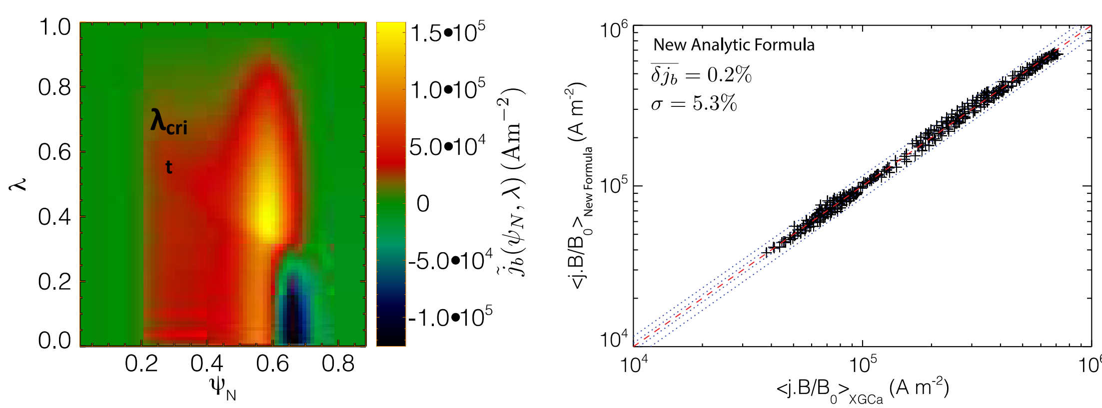

Figure 3.5 ***Left:*** **Edge bootstrap current distribution in the velocity (λ) and tokamak minor radius (Ψ_N) space. Strong current (yellow) in the trapped phase space above the λ_{crit}–curve can be seen in the edge region 0.5<Ψ_N<0.7. In the passing space below the λ_{crit}–curve, the current can even be negative (blue).** ***Right:*** **The new analytical current formula (vertical axis), created from numerous large-scale XGCa simulations using one-third of Mira, agrees excellently with the simulation results (horizontal axis).**
Image: Robert Hager, Princeton Plasma Physics Laboratory.

Usage of the INCITE and ALCC Hours

The INCITE 2015 program allocated 3.6 billion core-hours on Mira. The allocation usage is shown in Figure 3.6. Of the 37 INCITE projects, 25 projects used more than 90 percent of their allocation. Fifteen of these used their entire allocation (or more), including two projects using over 150 percent. These projects used the extra core-hours to achieve additional milestones. The overuse of Mira was made possible through the use of the backfill queue (low priority) and an “overburn” policy that permitted projects to continue running capability-sized jobs after their allocation was completely exhausted. The duration of the “overburn” policy was chosen to be July 1 through November 30, 2015.

Of the remaining 22 projects with outstanding allocations, 19 projects used more than 50 percent of their time and of those 19, 12 projects used more than 75 percent of their time. Three projects used less than 50 percent of their allocation. One of these projects was delayed due to code not being ready in time. The other two projects waited until too late in the year to complete their campaigns. A total of 4.0 billion core-hours were delivered to INCITE. The total number of INCITE hours delivered include 9.9 million core-hours carried out on Cetus because the scientific campaign of one of the INCITE projects could not be easily accommodated on Mira. The contribution from Cetus to the total INCITE hours delivered was about one quarter of 1 percent.

For the 2014–2015 ALCC year, 21 projects had allocations on Mira for a total of 1.8 billion core-hours. The allocation usage is shown in Figure 3.7. Fourteen of these projects used 90 percent or more of their allocation. One project also used Cetus for production runs that were not easily accommodated on Mira. Cetus usage accounted for about 10 percent of their usage and less than 1 percent of the total ALCC 2014-2015 usage. Three projects used less than 50 percent of their allocation. Of those three, one project had to modify code due to out of memory errors which then caused performance issues. They changed plans and tried to incorporate a newer version of the code in their production runs, but due to unforeseen manpower issues with the code development team, it was not ready in time. Another project had workflow issues and engaged the facility too late in the allocation year for staff to have an impact. The last project had to change milestones due to the failure of a third party to deliver promised code enhancements.

The 2015–2016 ALCC year is approximately halfway through its allocation cycle. So far, 24 projects have received allocations of 1.7 billion core-hours. They have used a total of 571 million core-hours from July 1 through December 31, 2015. The allocation usage is shown in Figure 3.8. One of the projects has already used up its allocation and is running in the backfill queue, while another two projects have used up 50 percent or more of their allocation. Again, one of the projects is using Cetus for runs that are not easily accommodated on Mira.

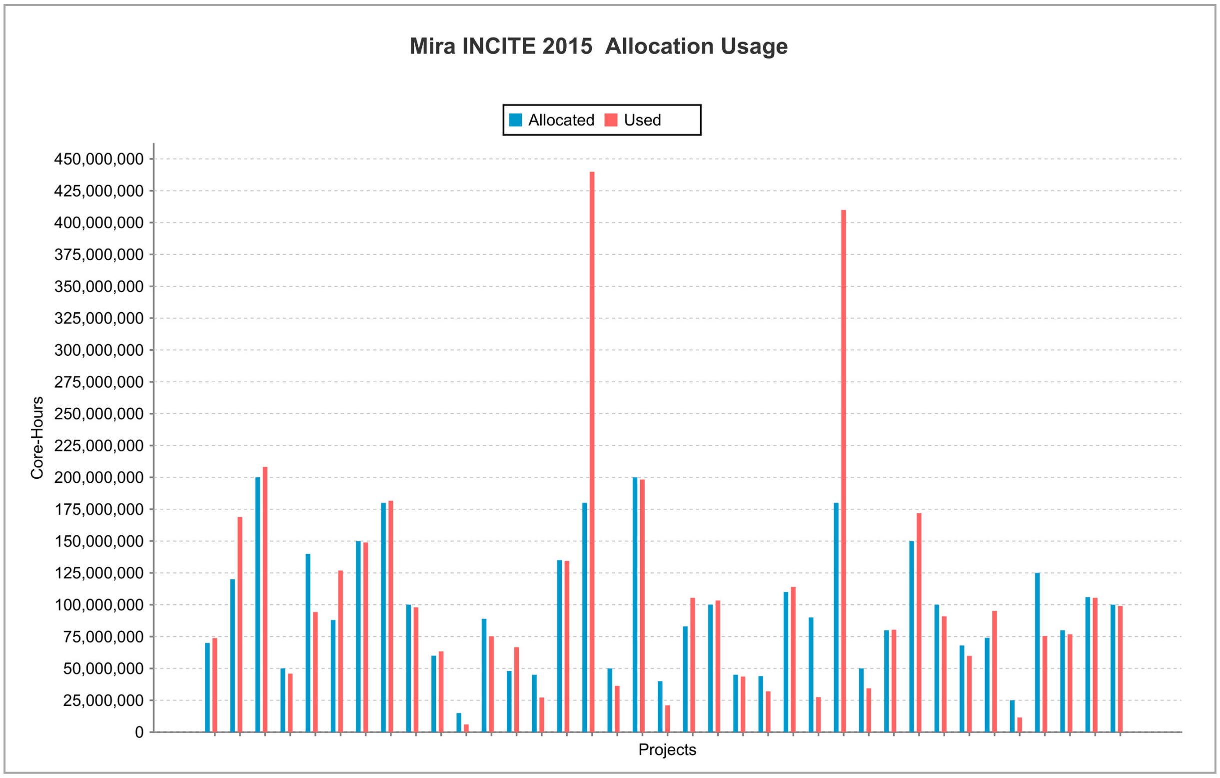

Figure 3.6 Mira INCITE 2015 Allocation Usage

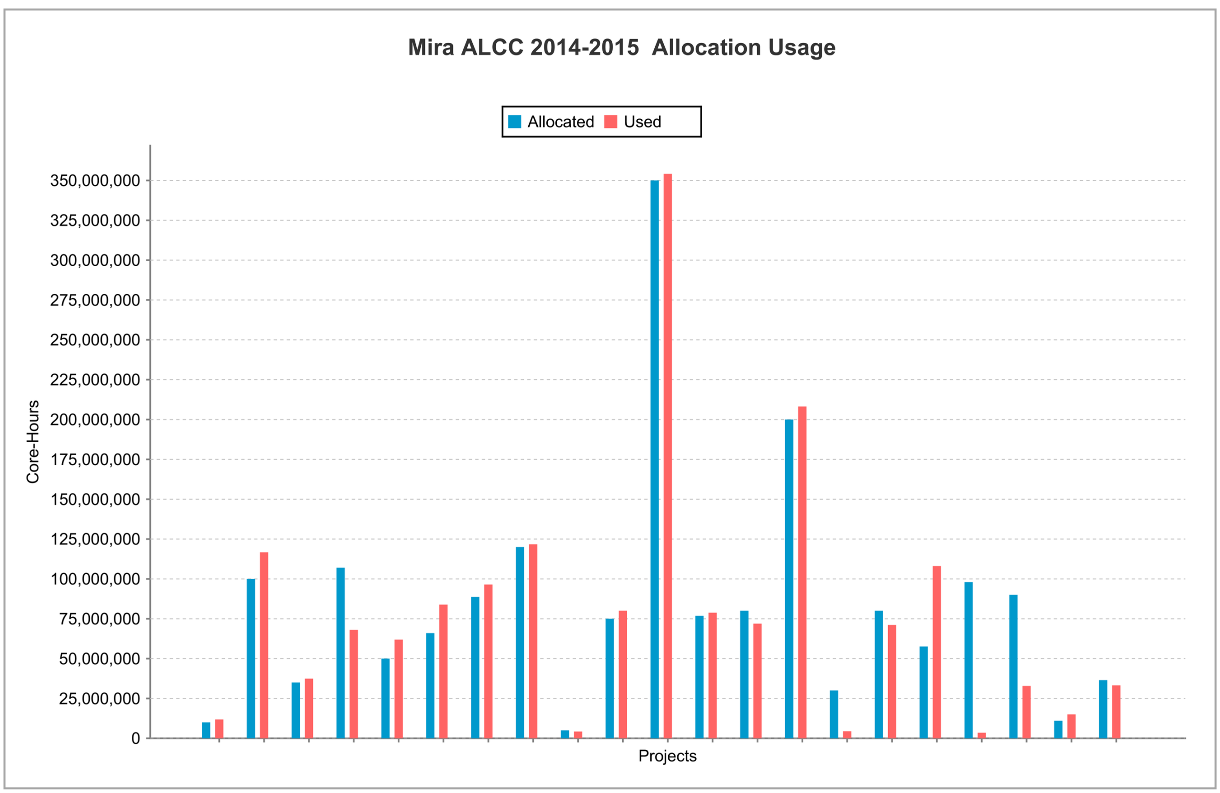

Figure 3.7 Mira ALCC 2014–2015 Allocation Usage

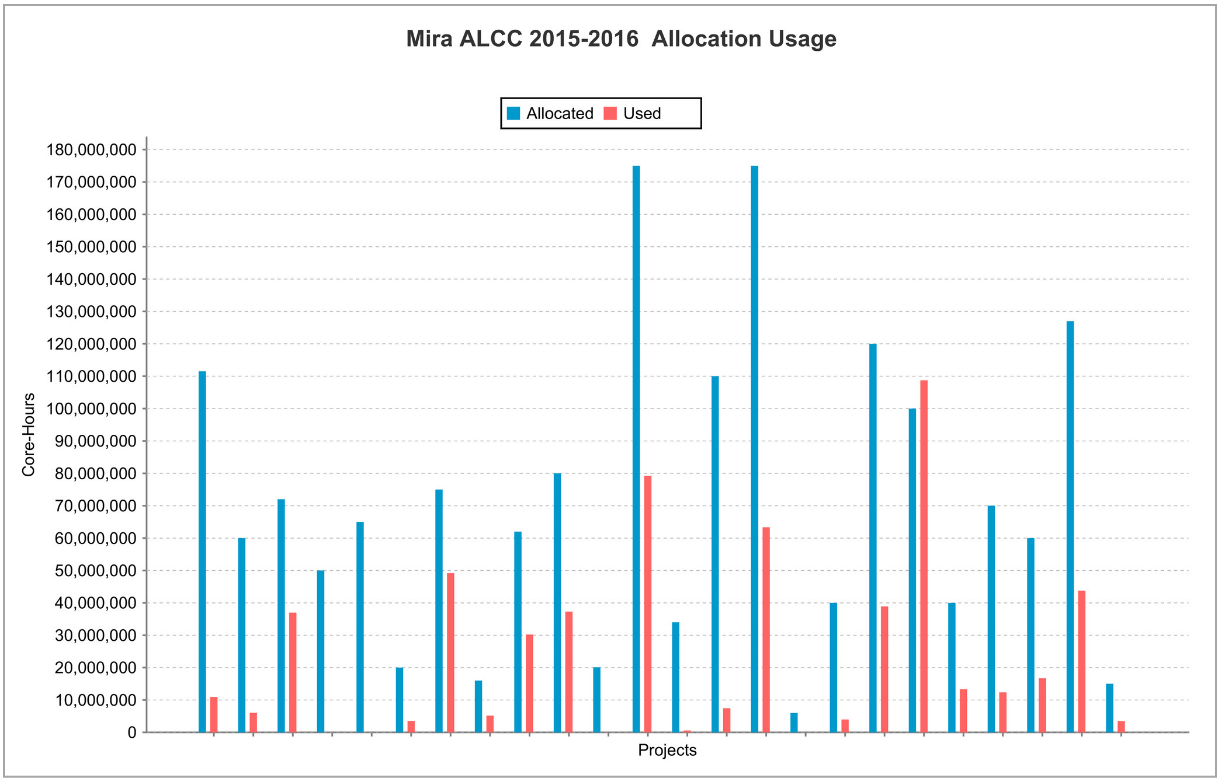

Figure 3.8 Mira ALCC 2015–2016 Allocation Usage

3.3 Allocation of Facility Director's Reserve Computer Time

In this section we are interested in the strategic rationale behind use of Director's Discretionary time. The Facility should describe how the Director's Discretionary reserve is allocated and list the awarded projects, showing the PI name, organization, hours awarded, and project title.

The Director's Reserve, or Director's Discretionary (DD) program, serves the HPC community interested in testing science and applications on leadership-class resources. Projects are allocated in four categories:

1) INCITE or ALCC proposal preparation
2) Code support and/or development
3) Strategic science
4) Internal/support

INCITE and ALCC proposal preparation allocations are offered for projects that are targeting submission of an ALCC or INCITE proposal. These projects can involve short-term preparation (e.g., a run of scaling tests for their computational readiness) or longer-term development and testing.

Code support and/or development allocations are used by teams porting and optimizing codes or projects developing new capabilities. This category includes the development, testing, and runs required for competitions such as the Gordon Bell Prize. Projects in this category have been responsible for bringing new capabilities to ALCF. For example, the PARTS project has supported multiple libraries and software packages. This effort has fueled many successful INCITE proposals and papers.

ALCF also allocates time to projects that might still be some time away from an INCITE award, or that offer a "strategic science" problem worth pursuing. Examples include supporting projects from DOE's Scientific Discovery through Advanced Computing (SciDAC) program, industry research efforts, and emerging use cases, such as coupling experimental and computing facilities.

Internal/support projects are devoted to supporting the ALCF mission. ALCF does not reserve core-hours for division activities. All activities come out of the DD allocation pool. This category regularly includes projects that help staff support the users and maintain the system, such as diagnostics and testing of tools and applications.

Allocations are requested through the ALCF website and are reviewed by the Allocations Committee (which includes representatives from Operations, User Services, and the Catalyst teams). The committee collects additional input from ALCF staff, where appropriate. Allocations are reviewed on their readiness to use the resources and their goals for the allocations and are awarded time on a quarterly basis. The DD allocation pool is under great demand, and often the requested amount cannot be accommodated.

Table 3.2 shows the number of projects and total time allocated in the DD program during 2015. By its very nature, the DD program is amenable to over-allocation since often time is left unused; however, it should be noted that these totals do not represent open allocations for the entire calendar year. A project might have a 1-million core-hour allocation that only persists for three months, but that 1-million core-hour allocation is counted entirely in the annual total core-hour number. Projects are not guaranteed the allocated time; rather, the time is provided on a first-come, first-served basis. DD projects run at a lower priority than INCITE or ALCC projects, which reduces the amount of time available for their use. Exceptions are made for some internal projects that support acceptance of new hardware or support of users, which is in line with the ALCF core mission.

Table 3.2 DD Time Allocated and Used on Mira, 2015

Projects	Mira
Allocated Core-Hours	858.3M
Used Core-Hours	373.2M[a]

[a] Usage includes 9.9M core-hours from Cetus non-capability production jobs.

A list of the CY 2015 DD projects, including title, PI, institution, and hours allocated, is provided in Appendix B.

Figure 3.9 provides a breakdown of the CY 2015 DD allocations by domain.

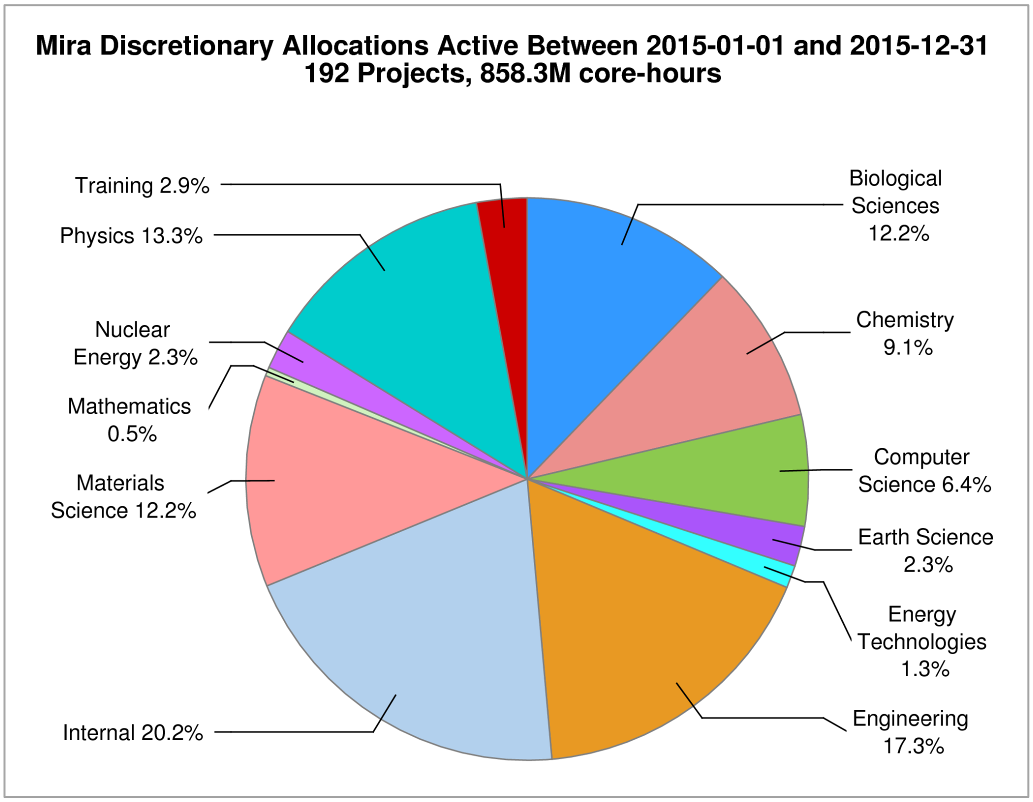

Figure 3.9 CY 2015 DD Allocations by Domain

Conclusion

ALCF continues to enable scientific achievements, consistent with DOE's strategic goals for scientific breakthroughs and foundations of science, through projects carried out on facility machines. Researchers participating in projects using ALCF resources published 164 papers in CY 2015. ALCF projects have had success in a variety of fields, using many different computational approaches. They have been able to reach their scientific goals and successfully use their allocations. A number of the projects and PIs have subsequently received awards or have been recognized as achieving significant accomplishments in their fields.

ALCF delivered the following core-hours to the allocation programs in CY 2015: 4.0 billion to INCITE, 1.6 billion to ALCC, and 373 million to DD. The DD Reserve has been used not only to develop INCITE and ALCC proposals but also to conduct real science of strategic importance and to drive development and scaling of key INCITE and ALCC science applications. Excellent ALCF support and solid, high-performing ALCF resources have enabled INCITE and ALCC projects to run simulations efficiently on HPC machines and achieve science goals that could not otherwise have been reached.

This page intentionally left blank.

Section 4. Innovation

Have innovations been implemented that have improved Facility operations? This includes innovations adopted from, recommended to, or adopted by other Facilities.

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2015. ALCF innovations and best practices have helped to prepare for future systems, enabled more efficient operations, and strengthened collaborations across Argonne.

4.1 Preparing for Future Systems

Future leadership-class systems will continue to increase node concurrency, requiring code that scales to millions of ranks. ALCF has worked to improve key software and develop tools that work at scale. In addition, ALCF and MCS have jointly developed a testbed for evaluating new HPC technologies.

4.1.1 OpenFOAM Improvements

Challenge: OpenFOAM is a free, open-source computational fluid dynamics (CFD) software package. Though widely used in science and engineering research, OpenFOAM is difficult to scale and tune for emerging high-performance computing architectures. While much of OpenFOAM's current user base is in Europe, the U.S. user base is expected to increase rapidly as ALCF transitions from the PowerPC to Intel Xeon-Phi architecture.

Approach: Based on user input, ALCF researchers have improved OpenFOAM performance and scaling in several domains. Computational scientist Ramesh Balakrishnan has led the facility's efforts, enabling improvements in memory utilization per core, integration of the National Renewable Energy Laboratory's Simulator for Wind Farm Applications (SOWFA) model into OpenFOAM-2.4.x, and the in-house development of improved subgrid models (namely, the WALES model). ALCF computational scientist Hal Finkel has also contributed to OpenFOAM improvements by developing a memory-logging tool to determine which objects occupy much of the memory, especially for simulations on large core (processor) counts.

ALCF has validated OpenFOAM for a wide range of flow scenarios via large eddy simulation (LES) and detached eddy simulation (DES) implementations in the software. In addition to simulating flows over complex terrain, the improved version of OpenFOAM has been applied to LES of flows over airfoils that are representative of wind turbine sections.

Figure 4.1 captures separated turbulent flow over a thick airfoil at a fairly high angle of attack, under inflow conditions that are representative of flows over wind turbine blades. The image shows the detail of wall shear stress on the surface of the airfoil with a background velocity plot. The DES, with the Spalart-Allmaras RANS wall model, is for a Reynolds number of 1.5 million (based on the airfoil chord length). Balakrishnan carried out the simulations as part of a Director's Discretionary project aimed at assessing and validating the computational performance and modeling capability of OpenFOAM for wind energy applications.

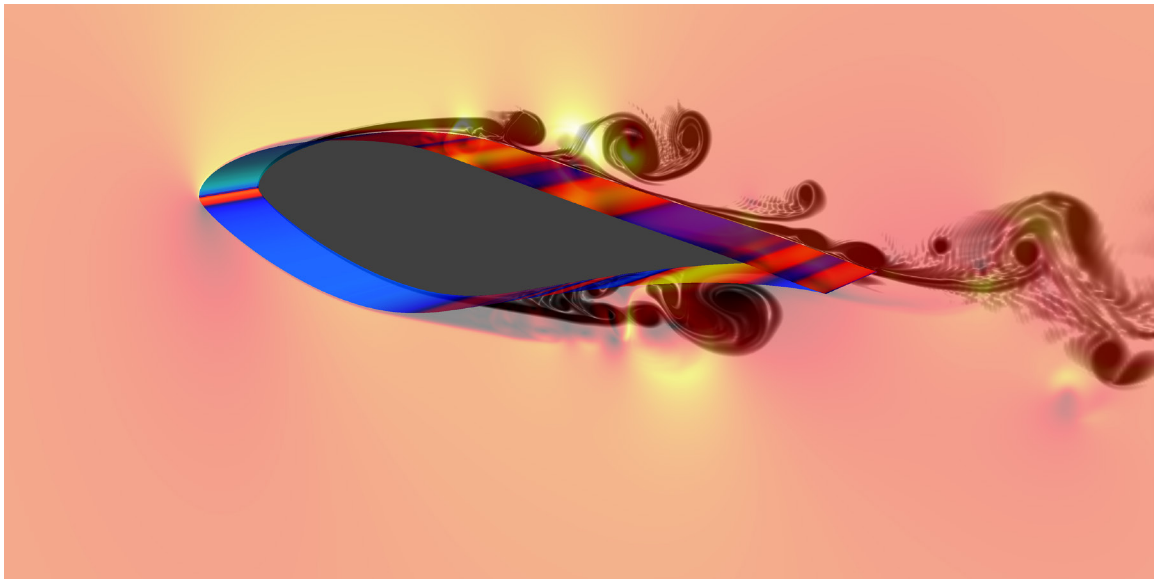

Figure 4.1 Separated turbulent flow over a S809 airfoil at an angle of attack of 10.2 degrees under inflow conditions that are representative of flows over wind turbine blades. Image: Ramesh Balakrishnan, Argonne National Laboratory.

Status/Impact: ALCF has continued working with OpenFOAM, validating and improving its capabilities, and rewriting some of the classes in OpenFOAM to further improve performance. As a result of the efforts:

- ALCF has improved OpenFOAM's performance and suitability for LES and DES of atmospheric boundary layers, as well as improving its overall performance on the Blue Gene/Q platform.
- ALCF has shared the improvements with the OpenFOAM community through presentations and contributions to the OpenFOAM code base.
- Several industry users (such as Vestas Wind Systems) have begun Director's Discretionary projects using OpenFOAM.
- Researchers at the National Renewable Energy Laboratory are pushing the limits of OpenFOAM performance on their own resources and Mira. They are working with ALCF on improving and validating their SOWFA module for wind energy applications.
- Researchers at Stanford University are working with ALCF to compare the performance of OpenFOAM with the open-source CFD package SU2, and improve SU2 by building

some of OpenFOAM's capabilities into it. As part of their efforts, the Stanford team was chosen to participate in ALCF's Early Science Program for Theta.

- The Wind and Water Power Technologies Office within DOE's Office of Energy Efficiency and Renewable Energy funded an Argonne project (with Balakrishnan as co-principal investigator) to simulate the atmospheric boundary layer for some canonical cases. Some of this funding was used to improve OpenFOAM performance. These efforts include improvements to the Clang/LLVM Blue Gene/Q compiler to better vectorize C++ codes and development of a memory logging software package (libmemlog) that monitors OpenFOAM's memory use patterns.

4.1.2 LAMMPS Reax/C Development

Challenge: The LAMMPS code did not perform well for researchers carrying out reactive molecular simulations on Mira.

Approach: In a collaborative effort including researchers from Michigan State University (MSU), IBM, and Sandia National Laboratories, ALCF staff improved performance of the Reax/C module in the LAMMPS code for reactive molecular simulations. This work largely focused on efficient hybrid MPI/OpenMP implementation of the original MPI-only Reax/C code. The team also improved the performance of LAMMPS by replacing MPI point-to-point communication with MPI collectives and leveraging MPI I/O to take advantage of network topology optimizations. Atoms and their inter-particle interactions were scheduled across threads in a balanced way employing thread-privatization of force arrays to prevent race conditions and subsequent tallying across threads. Additionally, a dual conjugate-gradient solver for charge equilibration was implemented to simultaneously solve two sparse systems of equations for determination of atomic charges, whereby multiple OpenMP loops were fused to assign larger per-thread workloads and reduce the number of OpenMP parallel regions. The separate MPI messages for each system were bundled together into larger messages improving communication throughput and reducing synchronization overheads during the iterative solve.

As shown in Figure 4.2, with all performance improvements combined, speedups up to 4.5x were observed for a modest sized system of 32,000 particles on 1,024 BG/Q nodes and 2–3x speedups on larger systems containing 16.6 million particles on 16,384 BG/Q nodes with an overall weak scaling parallel efficiency of 91 percent on Mira.

This software development capability was initiated by ALCF staff working with the INCITE project led by Subramanian Sankaranarayanan (Argonne) to investigate electrochemical oxide interfaces at the mesoscale. The INCITE team used the hybrid Reax/C code to help determine the microscopic mechanism for superlubricity whereby the friction coefficient measured from two sliding surfaces is significantly reduced to negligible levels. This important discovery was reported in the journal *Science* and has the potential to greatly benefit industry and consumers alike that lose energy and money due to overcoming friction loss.

ALCF computational scientists Chris Knight, Wei Jiang, and Nichols Romero along with H. Metin Aktulga (Lawrence Berkeley National Laboratory/MSU), Tzu-Ray Shan (Sandia), and Paul Coffman (IBM/ALCF) collaborated on these performance optimizations.

Status/Impact: The code is currently supporting science runs for ALCF projects. The ALCF team is preparing a manuscript describing the hybrid Reax/C package for publication and the code will soon be distributed with the public version of LAMMPS.

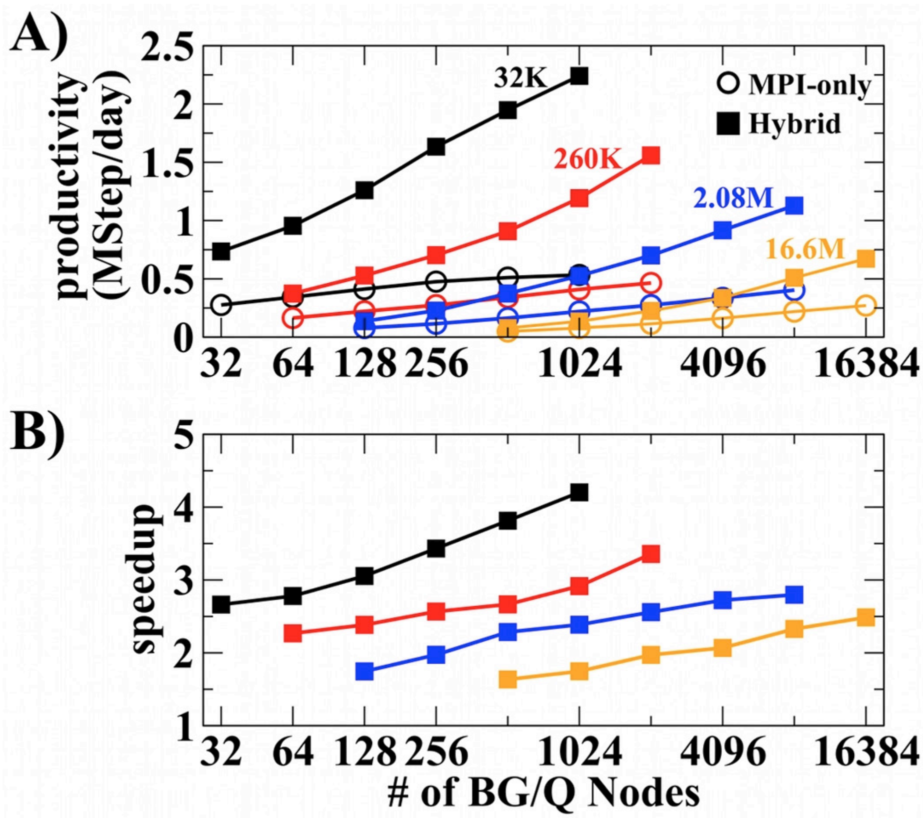

Figure 4.2 A) Measured performance of original MPI-only (open circles) and newest hybrid (squares) LAMMPS Reax/C implementations reported as millions of timesteps (MSteps) per wallclock day for four system sizes of a PETN crystal benchmark: 32,480 (black), 259,840 (red), 2,078,720 (blue), and 16,629,760 (orange) atoms. B) Relative speedups of hybrid vs. MPI-only implementations.

4.1.3 VSVB

Challenge: Traditional methods of quantum chemistry struggle to combine computational scalability with chemical generality. That is, the ability to treat any class of chemical problem together with the ability to expedite large problems in a reasonable amount of time, making efficient use of large processor counts.

Approach: ALCF computational scientist Graham Fletcher has developed the Variational Subspace Valence Bond (VSVB) method based on ideas from his doctoral thesis (1994) and papers. VSVB is a highly scalable and rigorous *ab initio* electronic structure model built from naturally localized and chemically intuitive objects such as bonds and lone pairs (Figure 4.3). VSVB can model any kind of chemical problem, and is strongly compute-bound and trivially parallel, with linear memory requirements and sub-cubic complexity. In addition, VSVB is easily threaded and vectorized, making it ideally suited to harnessing large-scale accelerator-based architectures for the study of complex chemical problems.

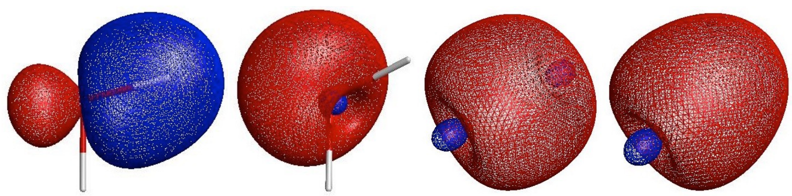

Figure 4.3 Overlapping orbitals (left to right): OH-bond and oxygen σ-lone pair in water; carbon-carbon σ-bond; CH bond.

The key to achieving this lies in VSVB's use of overlapping linear combinations of atomic orbitals (OLCAO) rather than the orthogonal "molecular" orbitals used in mainstream quantum chemistry methods. To represent the shell structure of chemistry arising from the "exclusion principle" of quantum mechanics, the OLCAO must be provided with variational support against collapse during optimization of the wave function with respect to minimizing the total energy. To do this, VSVB uses special subsets of the basis set—the 'variational subspaces'—in which to expand each OLCAO in the problem.

VSVB's use of variational subspaces greatly ameliorates the exponential scaling of valence bond methods, while the flexible OLCAO forms (e.g., bonds and lone pairs) greatly simplify complex chemical problems. In this way, VSVB solves the long-standing problem in quantum chemistry of how to combine computational scalability with chemical generality. In addition, VSVB makes it possible to combine OLCAO saved from previous calculations to generate a qualitatively correct and accurate wave function for a new problem at negligible cost.

In contrast to conventional quantum chemistry methods, the cost of VSVB is dominated by computing determinants (for the density) rather than multi-center integrals. Determinants have the advantage of being a relatively straightforward application of linear algebra, unlike the complex integrals, and are therefore much easier to multi-thread and vectorize. As an example of VSVB's high scalability, (Figure 4.4, left) shows a calculation exceeding 2 million ranks on Mira with 85 percent efficiency. On the right, the figure shows the calculated problem with an unprecedented 728 atoms and 1,638 electrons compared to traditional valence bond methods which can handle around 10 electrons. An example of VSVB's chemical generality is its ability to predict the correct dipole orientation in carbon monoxide using a simple form of wave function (known as a "single reference" wave function), compared to conventional methods that need forms with much higher complexity.

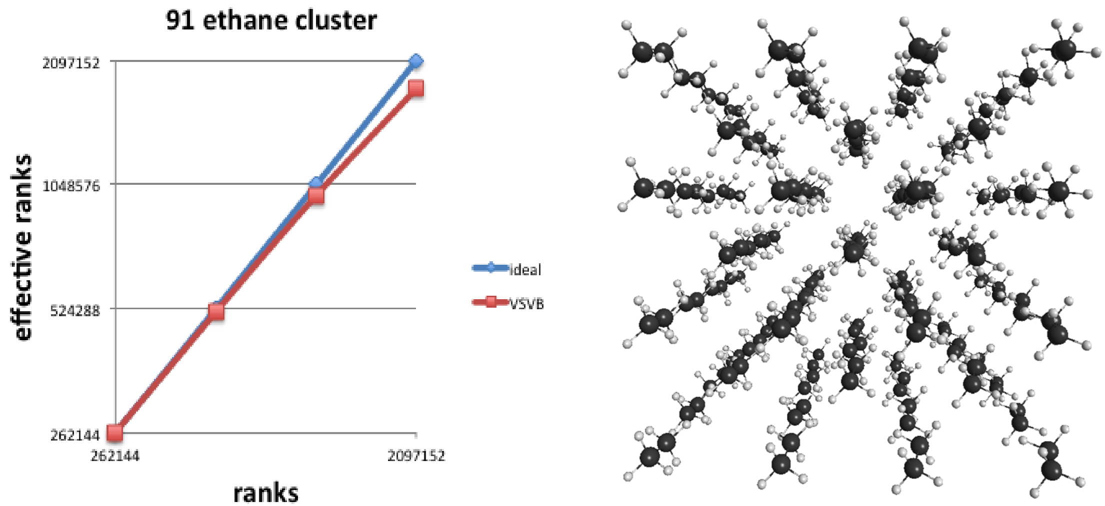

Figure 4.4 Over 2 million ranks with 85% efficiency simulating 728 atoms and 1,638 electrons.

VSVB+QMC
Fletcher and fellow ALCF computational scientist Anouar Benali are currently developing and improving workflow tools to allow QMCPACK to be used in conjunction with VSVB. By using a VSVB wave function as input, a subsequent QMC calculation needs both less memory (due to fewer determinants) and less arithmetic (due to localized orbitals), while the high scalability of VSVB also allows larger problems to be treated with the combined method. VSVB+QMC is currently the most accurate and scalable way to model complex chemical problems.

Impact/Status: VSVB is being used for QMC work in several projects. Both an INCITE proposal on transition metal complexes for battery technology and an ALCC project on exciton models for photovoltaic materials use VSVB. ALCF is currently advertising a postdoctoral student position to further VSVB+QMC work.

4.1.4 Evolving Leadership Computing with Workflows

Challenge: The use-model for leadership computing is evolving from the traditional large campaign. Moving forward, science use-models are expected to make increasingly complex demands. Coupling experimental facilities with real-time data analysis, deep learning methods, complex workflows, and other factors are expected to be regular use patterns. Understanding how the facility needs to adapt to serve these emerging use patterns is critical.

Approach: For several years, ALCF staff has collaborated with projects to adapt their workflows for ALCF computing resources to expand scientific capabilities. ALCF continued this work in 2015 and many of the lessons learned are helping prepare the facility for next-generation systems.

Impact/Status: ALCF projects used a wide variety of workflows in 2015, ranging from ensembles to complex work patterns for managing data as well as simulations. Examples include:

- Tom LeCompte's ALCC project performed workflow-driven event generation with ALPGEN. ALCF software developer Tom Uram played a key role in moving this work forward, helping the research team enable automation of serial phase-space integration using ALPGEN on an HEP cluster, GridFTP-based movement of this data to Mira, submission of event-generation jobs on Mira, and GridFTP-based movement of generated events back to the HEP cluster for transmission to the ATLAS experiment.
- With assistance from ALCF researchers Ketan Maheshwari, Marta García, and Kevin Harms, Sibendu Som's ALCC project used Swift to run the computational fluid dynamics (CFD) software Converge to conduct a global sensitivity analysis on the key CFD model inputs for gasoline compression ignition.
- Mainak Mookherjee's DD project collaborated with ALCF staff member Ketan Maheshwari to conduct about 900 independent VASP runs on Mira using the Swift sub-block technique.
- Ilja Siepmann's ALCC project used an ensemble driver to mine material databases and run 100,000+ molecular simulations to find materials with optimal performance for targeted applications. ALCF computational scientist Chris Knight designed an MPI-based framework to the manage ensemble calculations.
- ALCF computational scientists Paul Coffman and William Scullin worked closely with the Joint Center for Energy Storage Research's (JCESR's) DD project to enable the Fireworks workflow software and improve application performance. Their work allowed the JCESR team to run high-throughput virtual screening to identify properties for candidate molecules for electrolytes using Q-Chem and for cathodes using VASP.
- Aytekin Gel's ALCC project performed uncertainty quantification running thousands of streaming experimental data to ALCF from the Advanced Photon Source for real-time analysis. The project team collaborated with ALCF computational scientist William Scullin to find and implement the right workflow package.
- Noa Marom's ALCC project used a genetic algorithm to search for the optimized crystal structure of a given molecule with a custom Python framework. ALCF computational scientist Alvaro Vazquez-Mayagoitia collaborated with project researchers to develop the GAtor code for use on Mira.
- James Kermode's INCITE project used a server-client model to create and improve systematically large training datasets for a "Machine Learning on-the-Fly" (MLOTF) method for quantum mechanics/molecular mechanics simulations. This model enabled their project to compute thousands of small VASP calculations on >16K Mira nodes. ALCF computational scientist Alvaro Vazquez-Mayagoitia helped port the application code QUIP and ALCF software developer Kevin Harms improved scalability of the team's MLOTF approach.
- Warren Washington's INCITE project used a long series of ensemble jobs to perform multiple simulations concurrently. Multiple ALCF staff members, including Yuri Alekseev and Ray Loy, collaborated with the team to help build the workflow for Mira.

- ALCF staff member William Scullin used Lawrence Livermore National Laboratory's Cram tool (https://github.com/LLNL/cram) to run sample UQ workloads with scaling only seemingly limited by I/O considerations.

4.1.5 Joint Laboratory for Systems Evaluation

Challenge: ALCF needed a testbed to evaluate new systems and their impact on ALCF software.

Approach: To systematically test new hardware, ALCF and MCS established the Joint Laboratory for System Evaluation (JLSE), a collaborative effort aimed at evaluating future high-performance computing platforms. JLSE was set up to manage the ALCF and MCS computing and computational science activities at Argonne specifically aimed at researching next-generation hardware and software platforms. ALCF staff and other JLSE members have access to various processor architectures including ARM64 (X-Gene), x86-64 (Haswell), IBM Power, Xeon Phi (Knights Corner), and Nvidia GPUs to perform their development work.

To prepare for the scale and architecture of future systems, ALCF staff has begun analyzing important application benchmarks and mini-apps using JLSE computing resources to understand the performance and characteristics of next-generation platforms. By considering architecture characteristics, such as instruction issue rates, memory and cache hierarchy, and vector instruction operations, the team is working to identify the capabilities and limitations of future architectures, best practices for improving the performance of applications, and, ultimately, a roadmap for users to effectively adapt and tune their codes for new leadership-class systems. Examples of ongoing JLSE projects include work involving the spectral element Poisson solver used by the Nek5000 application and open-source libraries commonly used in quantum chemistry and computational physics, such as libxsmm, libMADMTXM (in MADNESS), and ELSI project (ELPA, libOMM and PEXSI).

Impact/Status: JLSE is an active testbed that has provided ALCF, MCS, and other Argonne researchers with access to new and experimental hardware. Several JLSE activities are underway with the goal of addressing ALCF and MCS needs in a variety of areas, including:

- Improving science productivity on current and future leadership computing platforms.
- Investigating alternative approaches to current and future deployments (both hardware and software) within ALCF.
- Maintaining a range of hardware and software environments for testing MCS research ideas.
- Helping to drive standards in standard forums on benchmarks, programming models, programming languages, memory technology, etc.

4.2 Enabling More Efficient Operations

As part of ALCF's effort to continually analyze and improve on existing operations, the facility implemented new approaches that are helping to better understand system usage, enhance communications on key system details, and streamline how ALCF-developed codes are shared with the community.

4.2.1 Compiler and Library Tracking

Challenge: Data on compiler and library usage is needed to better understand user behavior and improve user support.

Approach: In the summer of 2015, ALCF deployed two experimental software components designed to track compiler and library usage on our systems: Trackdeps and Tracklib. This work is in an early, investigative stage.

The overall approach is to enable tracking when executables are built; the input files used to build the executable are recorded. Later at runtime, executables are scanned and runtime information is logged. Over time, historical data will be available for analysis and can be used to inform decisions on user support and research priorities.

The first stage in the library-tracking process involves the Trackdeps component. This component is loaded into the user's environment and injects itself into invocations of compilers, linkers, and other programs involved in the build process. Trackdeps records paths to all of the inputs to the build process that contribute to the final output, including compiler identity, header files, Fortran module files, and libraries. This information is stored in a special log directory along with a uniquely identifying hash of the resulting executable. This hash is later used to correlate the input dependencies recorded by Trackdeps with user/project information when the executable is eventually run.

Tracklib, the second stage in the library-tracking process, is a set of tools used to maintain a history of the libraries, object files, and functions incorporated into or used by applications. At present, Tracklib records sufficient information to track each execution of a unique application build. This information may also be used to extract details about the build from Trackdeps.

A future version of Tracklib will include tools for annotating existing and newly created libraries and object files. In part, these annotations would establish the usage of libraries and object files by an application when other mechanisms might not be successful. Such annotations could also be used to list which applications used a library, determine which routines are most frequently used, or find all applications utilizing a routine in particular library build that has a known problem.

Impact/Status: Trackdeps and Tracklib are currently deployed and undergoing tests on ALCF production resources. The data warehouse is receiving the experimental usage data, and while preliminary, early results are promising. Figure 4.5 depicts a sample overview of library usage enabled by this suite of tools.

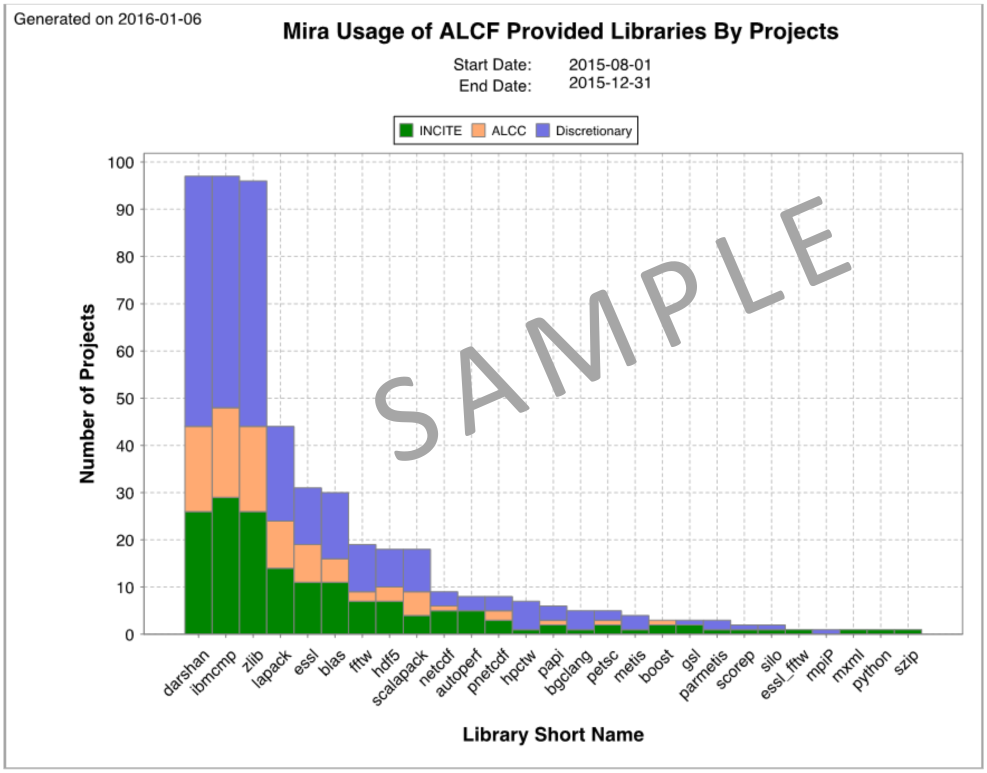

Figure 4.5 Library Usage by Project

4.2.2 AutoPerf

Challenge: Collecting hardware performance counter and MPI information on the Blue Gene platform is difficult without specialized knowledge and libraries.

Approach: ALCF performance engineer Scott Parker developed AutoPerf, a new tool for the automatic collection of application performance data on Blue Gene/Q. AutoPerf records performance information for all applications that are recompiled or relinked with the tool and run on a supported system. The library transparently collects performance data from running jobs and saves it into files at jobs completion. AutoPerf output is in plain text and includes MPI usage and performance information indicating which MPI routines were called, how many times each routine was called, the time spent in each routine, and the number of bytes sent or received if applicable. Data from the hardware performance counters is also collected and written (Figure 4.6).

The collection of performance data and the generation of performance data files requires the program use MPI, call MPI_Init() and MPI_Finalize(), terminate without error, and not use other performance tools or libraries that use the PMPI_ interface and the BGPM performance counter API. The measured impact of data collection on program runtime has been observed to be less than 1 percent.

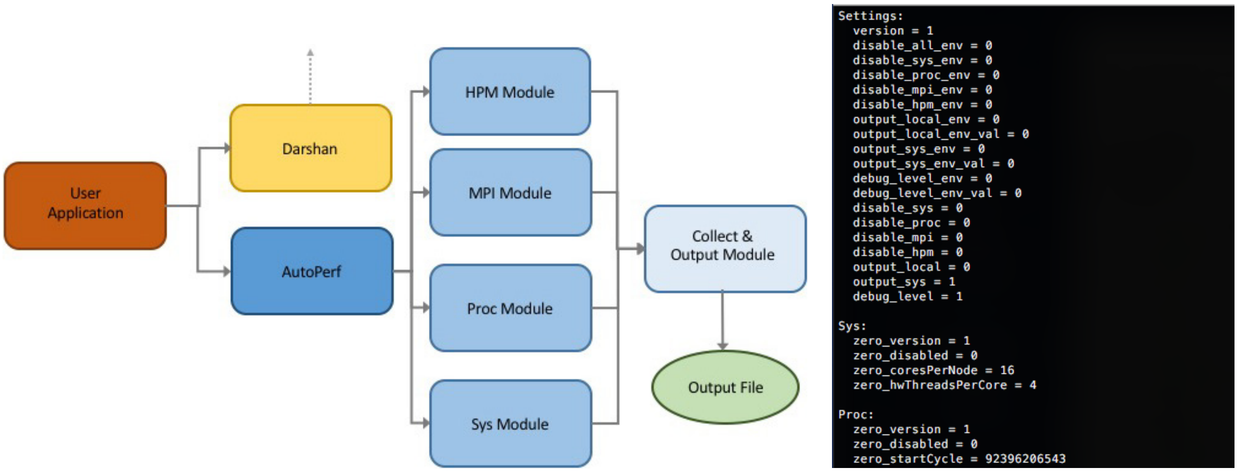

Figure 4.6 The structure of the AutoPerf software modules (left); Typical output from an AutoPerf run (right).

Status/Impact: AutoPerf has been enabled on ALCF systems by default since September 2015. Through the end of 2015, AutoPerf data has been collected for over 85,000 jobs from over 175 users and 500 distinct binary names. No issues or problems have been reported and in-depth analysis of the collected data is beginning.

4.2.3 ALCF on the Move

Challenge: Status information on ALCF resources and events was only available on the ALCF website through a traditional web browser on a desktop or laptop computer. To quickly respond to changes in resources and events, ALCF staff and users need real-time information that is easily accessible and quickly customized.

Approach: ALCF software developer Janet Knowles created ALCF On The Move (OTM), a mobile productivity application that provides staff with real-time information on ALCF resources and events. Staff can easily access resource information anywhere and anytime, freeing them from the traditional desktop or laptop experience that is dependent on a Wi-Fi or Internet connection.

The four main sections of the app are:

Machines Status Table

The opening screen of the app is the Machines Status Table, which displays an overview of each ALCF resource in different table cells. Each machine has a doughnut graphic indicating the percentage of the machine currently in use and statistics on running and scheduled jobs.

Map

The Map displays the physical layout of a machine and indicates running jobs. Each group of blocks of the same color represents the blocks in use by a particular job. Jobs can be selected to show all nodes that are processing that job and further detail is available in a pop-up window.

Machine Jobs Table
The Machine Jobs Table is a tabbed pane that displays all running and queued jobs on a machine, along with any reservations. Users can utilize this screen to request notifications about particular jobs. For each watched job, the user will receive notifications when a job changes status (e.g., when a job's status changes from queued to running).

News, Events, and Announcements
The News, Events, and Announcements screen shows a brief summary of major news, events, and announcements to keep users informed of activity within ALCF.

In addition to the main sections, the Settings screen allows users to customize OTM. If activated, notifications can be set for each machine to let the user know when the job queue drops below a specified number of jobs. The refresh interval to upload new data can also be set in this screen.

As a brief example that illustrates the utility of OTM to ALCF staff, an ALCF computational scientist used OTM to request a notification when the queue was low on the visualization cluster. She was out of the office when she received a push notification on her phone that the queue was low. She checked the app and saw there were no reservations scheduled. She then quickly contacted a user who was doing visualization work, informed him that the cluster had low activity, and recommended that he schedule jobs. The staff member was able to aid a user by acting quickly on real-time information, rather than waiting until she had a moment to access a laptop or time to scan the web for queue information.

OTM currently supports Apple's iOS platform. It was released in May 2015 through Apple's Enterprise Distribution system to ALCF staff. An April 2016 release is planned for general availability through the Apple App Store.

Impact/Status: By providing an easy way to quickly check or be notified of the status of resources and jobs, OTM enables staff (and, eventually, users) to rapidly react to situations that will affect their work. Rather than having to manually extract specific information from a website, OTM users can leverage push notifications to inform them when an important event occurs.

OTM also provides ALCF with increased exposure in two dimensions (Figure 4.7). First, OTM highlights ALCF news, events, and announcements. Second, it adds a valuable new dimension to the ALCF community by offering users another platform to stay connected with the facility. OTM is currently available to ALCF staff. A half dozen ALCF staff use it daily and ALCF plans to publicize it more widely to the user community in the coming year.

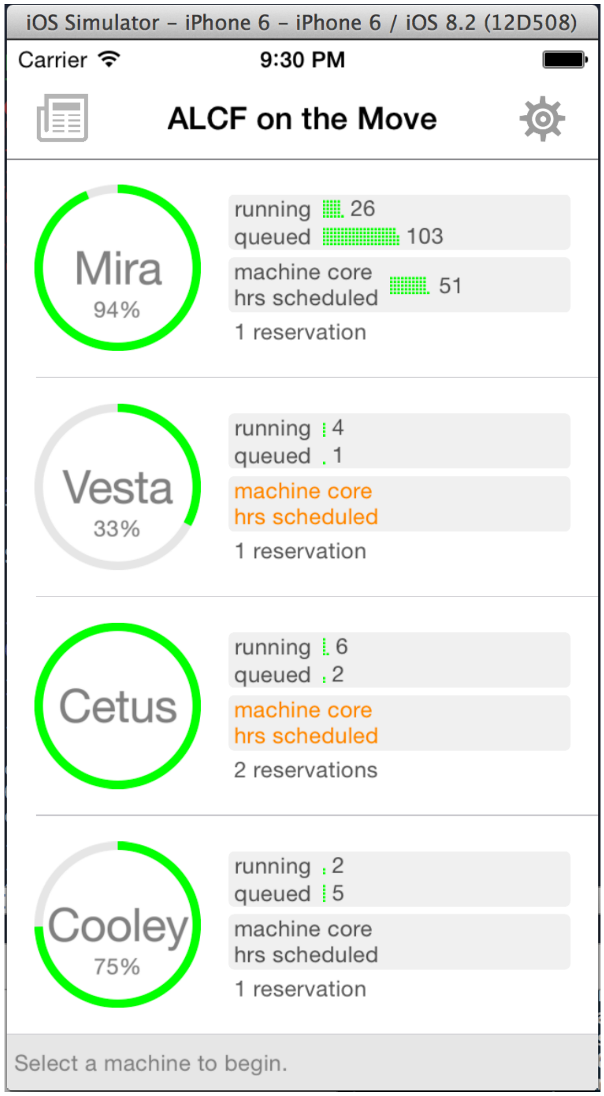

Figure 4.7 ALCF on the Move screen capture.

4.2.4 Improved Open Source Procedure

Challenge: Argonne's procedure for declaring code to be open source led to confusion and delays of many months.

Approach: ALCF-developed code is often of use outside Argonne, so ALCF makes codes available through open-source distribution. In the past, the procedure for doing this was time consuming and poorly documented. Argonne's Technology Development and Commercialization (TDC) Division managed many of the steps but lacked the domain knowledge to collect the needed information and the right set of inputs. Working with TDC and Argonne's legal department, ALCF developed a well-documented set of steps to streamline the process and ensure most of the analysis work is done within ALCF, resulting in more informed decisions and speeding the overall process. For example, ALCF does the intellectual property review internally, calling on outside experts as needed. While this process formerly required about one month for TDC to coordinate, ALCF can complete it in a matter of days. TDC now showcases ALCF's process as a model for other Argonne divisions, such as Mathematics and Computer Science.

Also in coordination with TDC and Argonne's legal department, ALCF established a standard BSD-style software license that strikes the right balance between preserving recognition for Argonne and allowing unlimited modification and use. This license is now used across all ALCF open-source projects.

Impact/Status: ALCF can now process staff requests for open source in a few weeks with much less effort. ALCF's procedure is being used as a model for other Argonne divisions.

4.3 Strengthening Collaborations Across Argonne

ALCF has strengthened collaborations with colleagues from across Argonne to help broaden the impact of ALCF computing resources and expertise. This includes working with other divisions to improve data management, repurposing Intrepid hardware to groups that needed additional storage, and partnering with other divisions to improve industry outreach efforts.

4.3.1 Petrel Project

Challenge: Researchers at ALCF and Argonne need to share large datasets with internal and external collaborators.

Approach: The Petrel data service pilot project provides a mechanism for Argonne researchers and ALCF users to store their data and share with collaborators without the burden of local account management. Researchers from ALCF and Globus are developing the system collaboratively.

Petrel leverages ALCF's storage and infrastructure and Globus's transfer and sharing services to provide a mechanism for researchers to transfer data into the system, manage data on the file system, and share and transfer data to other locations. Authentication and identity to access the system is provided through Globus and users can access Petrel using their campus or institution federated login credentials.

The pilot system consists of 32 file servers and 1.7 PB of usable GPFS storage. The backing storage consists of four DDN S2A9900 storage systems. The file servers also serve as GridFTP data transfer nodes (DTNs), with a 1 GbE wide-area network (WAN) connection per file server. GPFS traffic utilizes a dedicated 10 Gbps Clos network. GPFS was benchmarked with IOR at 8301.27/6059.07MB/s R/W using all 32 file servers as clients. Single client performance was benchmarked with IOR at 1215.48/736.42 MB/sec R/W. Networking performance has been measured with nuttcp between Mira and Petrel at 26 Gbps using all 32 file servers to a test Mira DTN with a 4x10 GbE aggregate.

In one exemplary use of the Petrel project, ALCF, MCS, and APS partnered together to improve the data management of APS experimental data. Key criteria were as follows:

- Not part of the mission critical infrastructure of the APS (the beamlines could continue to operate without it)
- Provide sufficient retention time to complete analysis and move data to home institutions
- Provide robust, "fire and forget" data movement mechanisms to minimize scientist effort (Globus Online)
- Allow guest scientists to enable access to the data by other scientists who do not have Argonne accounts (Globus sharing)
- Ideally, though not yet addressed: provide transparent access to this data from other Argonne compute facilities.

Impact/Status: Petrel is enabling researchers to share large datasets with high performance, security, and reliability. Petrel is deployed for researchers at APS and will be expanded to other users in the near future.

4.3.2 Repurposing of Intrepid Storage Infrastructure

Challenge: After decommissioning Intrepid, ALCF had DDN9900 storage arrays and file servers that were still functional, but no longer part of the production infrastructure. ALCF also had 272 sockets of GPFS server licenses that IBM agreed to continue licensing at the original (below current market) rates. In CY 2015, ALCF sought to extend the benefits of this infrastructure beyond its intended use where possible.

Approach: ALCF has worked with other Argonne divisions to identify where these resources might be repurposed at a benefit to the laboratory.

Impact/Status: ALCF has deployed surplus storage infrastructure to a number of divisions at Argonne, increasing collaboration, reducing costs, and extending the value of the existing investment. The following uses have been implemented or are under discussion:

- Prototyping Support:
  - ALCF has repurposed four DDNs and their file servers for the Petrel project.
- Additional Storage:
  - DOE's Midwest Integrated Center for Computational Materials (MICCoM) is using two DDNs, the associated servers, and GPFS licenses.
  - Argonne's Advanced Photon Source will use four DDNs, the associated servers, and GPFS licenses.

- Argonne's High Energy Physics Division will use one DDN, the associated servers, and GPFS licenses on their compute cluster.
- Argonne's Center for Nanoscale Materials will use one DDN for additional storage, if space and power allow.

4.3.3 Collaborative Industry Outreach

Challenge: ALCF interacts with many commercial entities interested in using ALCF resources, but often did not introduce them to the other collaboration opportunities at Argonne.

Approach: ALCF has many interactions with industry representatives interested in working with ALCF and its high-performance computing resources. When this interaction reaches a level in which concrete collaborations are discussed, ALCF has begun a new approach that involves introducing them to other parts of Argonne that may be relevant to their research. In particular, ALCF presents an integrated offering of computing resources that includes the Laboratory Computing Resource Center (LCRC). ALCF also involves the Technology Development and Commercialization (TDC) Division early in the process to help structure meetings and present Argonne's vast capabilities. In addition, ALCF engages Argonne domain scientists to present a more complete picture of the laboratory's resources and encourage additional collaborations. These efforts have resulted in broader engagements with a number of companies. In the case of Brewer Science, for example, the company initially approached ALCF to discuss opportunities for collaboration. Through this new approach, the company was also introduced to Argonne's Materials Science Division (MSD) and the LCRC. Brewer Science ultimately established overlapping projects with ALCF, MSD, and LCRC and was very appreciative of the coordinated response to its needs.

Impact/Status: ALCF's efforts to showcase the full breadth of Argonne's capabilities have resulted in deeper collaboration with industry partners and better coordination between laboratory resource providers. All future industry engagements will follow this model.

Conclusion

ALCF has identified innovations and best practices that have helped to prepare for future systems, enabled more efficient operations, and strengthened Argonne collaborations in CY 2015. This includes preparing for ALCF's next-generation supercomputers by developing critical software and computing tools that work at scale; deploying tools to gather detailed information on system usage; and forging partnerships across the laboratory to broaden the impact of ALCF work.

Section 5. Risk Management

Is the Facility effectively managing risk?

ALCF Response

ALCF has clearly demonstrated successful risk management in the past year for both project and operation risks. The risk management strategy is documented in the ALCF Risk Management Plan (RMP), which is reviewed and updated regularly to incorporate new ideas and best practices from other facilities. Risk management is a part of ALCF culture, and the RMP processes have been incorporated into both normal operations and all projects, such as the ALCF-3 project launched in CY 2013. Risks (proposed, open, and retired) are tracked, along with their triggers and mitigations (proposed, in progress, and completed), in a risk register managed by two Risk Co-Managers. All risk ratings in this report are post-mitigation ratings. ALCF currently has 40 open risks, with one High operational risk: funding uncertainty, which is managed by careful planning with the DOE program office and the continuation of austerity measures as necessary. The major risks tracked for the past year are listed below, with the risks that occurred and the mitigations for those risks described in more detail, along with new and retired risks, as well as the major risks that will be tracked in CY 2016.

Discuss how the Facility uses its RMP in day-to-day operations, how often the RMP is reviewed or consulted, and what happens when a risk occurs. For this review the focus is on Operational risks, not Project risks.

The Facility should highlight various risks to include:

- *Major risks that were tracked for the review year;*
- *Any risks that occurred and the effectiveness of their mitigations;*
- *A discussion of risks that were retired during the current year;*
- *The mechanism used to track risks and trigger warnings;*
- *Any new or recharacterized risks since the last review; and*
- *The major risks that will be tracked in the next year, with mitigations as appropriate.*

Note: *This is a high level look at the risks, not a deep dive into the risk registry.*

5.1 ALCF Risk Management

ALCF uses the documented risk management processes, first implemented in June 2006 and outlined in its RMP, for both operations and project risk management. ALCF reviews and updates the plan annually. The plan is also updated as needed during the year to reflect changes and to incorporate new risk management techniques as they are adopted by the facility. The RMP is consulted at all monthly and individual risk meetings. Details of the RMP, including the attributes of each risk managed by ALCF, have been described in past reports and will not be discussed further here. Risks are tracked in a risk register; as described below, in 2015, ALCF changed to a new risk register form.

Continuation of the ALCF-3 Project

The ALCF-3 project—procuring and deploying the next ALCF supercomputer—continued in 2015. A project risk register is maintained and a set of detailed risks is tracked. Risk mitigation costs on the project side are developed using a bottom-up cost analysis, then are input to the commercial project risk analysis tool Oracle Primavera Risk Analysis (OPRA) to set the contingency pool utilizing the OPRA integration with the Primavera project management tool. These risks are not included in the risk numbers covered in this document and are not discussed further.

New Steady-State Risk Register Format and Risk Form

During CY 2015, the ALCF transitioned the steady-state risk register from using OPRA to using an Excel spreadsheet format, with the official copies of the risk register entries stored in a Box folder. Box is a secure, cloud-based storage system used throughout Argonne. It records a version history for all files, so the history of all risks is automatically recorded.

In conjunction with the move to Box, a new steady-state risk Excel spreadsheet form was created that risk owners can use to easily create a proposed new risk or update an existing risk. Formatting the risk form as an Excel spreadsheet has the added advantage that a complete set of instructions for populating the form could be incorporated directly into the form on a separate page of the spreadsheet, so the instructions are always available to risk owners.

Finally, because Excel is simply a platform and not a project management tool, it is completely flexible, and the risk register can be perfectly designed to fit the requirements of operating a research user facility, which was not possible with OPRA due to the limitations of a project management tool.

The new risk register format and risk form were adopted from the risk register and risk form that were first developed by the ALCF-3 project. The project risk register and risk form were adapted to versions suitable for steady-state risk management, and the use of Box as a storage system was carried over from the project to steady-state operations.

The new blank risk form and instructions are shown in Figure 5.1.

Risk Form

(1) Risk Title: ______ (2) ID: ______

(3) Owner: ______ (4) Risk Status: ______

(5) Description of possible event (Risk): (6) Risk Type: ______

(7) Date Risk Proposed: ______ (9) Date Risk Recharacterized: ______

(8) Date Risk Approved: ______ (10) Anticipated Risk Closure Date: ______

(11) Cause: (What is the gap that allows the Risk event to be possible?)

(12) Effect: (What will happen if the Risk event occurs?)

(13) Triggers: (What event will happen that will mean the potential for the Risk event is imminent?)

Pre Mitigated:

(14) **Probability:** (18) Explanation:

(15) Probability Thresholds

VL	L	M	H	VH
<10% Risk Event Almost Never Occurs	10%-25% Risk Event Rarely Occurs	26%-74% Risk Event Occurs on Occasion	75%-90% Risk Event Occurs Often	>90% Risk Event Almost Always Occurs

Impact:

(16) Technical Scope (18) Explanation:

(17) Cost (18) Explanation:

(19) Pre Mitigated Score: ______

(20) Impact Thresholds

	VL	L	M	H	VH
Technical Scope Includes degradation to 1) computation, 2) security, 3) storage, and/or 4) data, as well as machine downtime. When applicable, it includes impacts on user experience, safety, and cyber security	No measurable impact on OAR metric, not required to inform DOE	OAR metric above target but drops by > 0.5% of target, DOE informed, but well below reportable	OAR metric above target but drops by > 2% of target, DOE informed, but somewhat below reportable	OAR metric above target but drops by > 4% of target, DOE informed, incident close to reportable	Drop below OAR metric target, or a reportable incident
Cost ($M)	< 0.1	0.1 - 0.5	0.5 - 1.0	1.0 - 2.0	> 2.0

(21) Response Type: ______

Figure 5.1 New ALCF Steady-State Risk Form and Instructions

(22) Management Strategy: (Planned strategies to lessen probability or impact of possible risk)

(23) Actions: (Planned actions if the risk occurs)

(24) Post Mitigated:

Probability: Explanation:

Impact:

Technical Scope Explanation:

Cost Explanation:

(25) Post Mitigated Score:

Risk Encountered: (27) Date of Encounter:

(26) Description of Encounter:

(28) Actions taken to manage the risk during encounter: (including start and end dates)

(29) Revised Mitigations:

Risk Closure: (31) Date of Closure:

(30) Description of Closure:

(32) Notes:

Figure 5.1 New ALCF Steady-State Risk Form and Instructions (Cont.)

(33) Revisions:

Figure 5.1 New ALCF Steady-State Risk Form and Instructions (Cont.)

Risk Form Instructions

(numbers in red correspond to the red item numbers in the risk form above)

Pre-Mitigated – Risk is identified, no actions have been taken.
Post-Mitigated – If actions (mitigations) are taken, what is the residual/remaining probability and impacts?

1 **Risk Title:** Short description of risk.
2 **ID:** Will be issued by the Risk Manager (sequential as new risks are opened).
3 **Owner:** The person who is responsible for monitoring the risk and implementing the mitigation.
4 **Risk Status:** Choose Proposed, Open, Impacted (Closed), Managed (Closed), or Rejected (Closed).

"Proposed" is used for new risks until approved by the Risk Review Board and ALCF management. Once approved, the status will be converted to Open. If Closed option is selected, you must fill out Risk Closure section.

- Impacted (closed) – The risk occurred, impacted operations, and is no longer an active risk.
- Managed (closed) – The risk was successfully managed as to not affect operations, whether it occurred or did not occur. The risk is no longer active.
- Rejected (closed) – The risk was deemed by the Risk Review Board and ALCF management as not needing to be tracked for steady-state operations.

5 **Description:** Detailed explanation of the risk event.
6 **Risk Type:** Select "Threat" for a risk that would have a negative impact or "Opportunity" for a risk that would have a positive impact, should the risk occur.
7 **Date Risk Proposed:** The date on which the risk owner submits the risk to the Risk Review Board for evaluation.
8 **Date Risk Approved:** The date on which the risk is approved by the Risk Review Board.
9 **Date Risk Recharacterized:** The date on which a significant change was made to the risk, such as changing the scope of the risk, changing the risk score, or changing the risk response type.
10 **Anticipated Risk Closure Date:** The date on which the risk is anticipated to end. "Ongoing" is acceptable for risks with no known end date.
11 **Cause:** The cause (preferably the root cause) of the risk.
12 **Effect:** The impact or effect of the risk event, should it occur, on the facility.
13 **Triggers:** One or more events that indicate that the risk event is imminent; early warnings.

Pre-Mitigated Section

This section is for probability and impacts based upon no mitigations; Risk is identified, what are the probability and impacts?

Figure 5.1 New ALCF Steady-State Risk Form and Instructions (Cont.)

14 **Probability:** The probability that the risk event will occur.

15 **Probability Thresholds:** Very High, High, Moderate, Low, or Very Low as defined by the table.

16 **Technical Scope:** Is there any degradation to computation, security, storage, and/or data as well as machine downtime? When applicable, it includes impacts on user experience, safety, and cyber security.

17 **Cost:** Is there any added cost including equipment, material, staff? All costs need to be expressed in dollars, not just equipment descriptions, amount of staff time, etc.

18 **Explanations:** Explanations of probability and impacts are there to help you account for the rationale behind the rankings provided and later decide if a revision to the assessment is needed. Each probability and impact must have an explanation.

19 **Pre Mitigated Score:** The score/ranking of the risk before any Management Strategies are enacted.

20 **Impacts Thresholds:** Very High, High, Moderate, Low, or Very Low as defined by the table.

21 **Response Type:** Accept, Reduce, Transfer, Avoid, Enhance, Exploit, Share.

Choose the response type that corresponds to the response/mitigation that has the greatest effect on the probability or impacts.

22 **Management Strategy:** Planned strategies to lessen the probability of impact of a risk.

23 **Actions:** What management strategies are planned to put into action if the risk happens?

Post-Mitigated Section

After responses from pre-mitigated section are applied, use the same thresholds provided in the Pre-Mitigated section to provide the updated probability and impacts. If Accept response – probability and impacts would be the same as pre-mitigated

24 **Probability and Impacts:** This is the residual probability and impacts; after mitigations but before risk event. Assumes all mitigations are implemented.

The same thresholds from Pre-Mitigation are used and an explanation for each must be completed.

25 **Post Mitigated Score:** The score/ranking of the risk after any Management Strategies are enacted.

Risk Encountered

26 **Description of Encounter:** Events that led to the risk occurring.

27 **Date of Encounter:** Date in which the risk occurs.

28 **Actions taken to manage risk during encounter:** What management strategies were planned if the risk happened and what was implemented to resolve the situation? Include start and end dates as appropriate.

29 **Revised Mitigations:** Any further responses that may need to be implemented even after the risk has occurred.

Figure 5.1 New ALCF Steady-State Risk Form and Instructions (Cont.)

Risk Closure

30 **Description of Closure:** Explanation of why the risk can or should be closed.
31 **Date of Closure:** Date in which the risk is no longer possible or risk event has completed.
32 **Notes:** Any added information that you feel should be documented regarding any of the fields above.
33 **Revisions:** Dates and descriptions of all the changes to the risk.

Figure 5.1 New ALCF Steady-State Risk Form and Instructions (Cont.)

New 5x5 Steady-State Risk Matrix

ALCF has revised the 5x5 steady-state risk matrix impacts to better reflect the impacts of risks that can be encountered in steady-state operations. The 5x5 matrix adopted in CY 2013 for steady-state risks was patterned after the ALCF-3 project 5x5 risk matrix. The ALCF Risk Review Board revisited the 5x5 matrix in CY 2015 and concluded that 1) the Technical Scope Impact needed to be broadened and more directly linked to OAR metrics in order to be applicable to the full range of steady-state risks and 2) the Schedule Impact, which was being interpreted for the steady state as machine downtime, should be considered as one part of Technical Scope and not a separate impact. Keeping the Schedule Impact had been a holdover from the project format, where it does need to be considered as a separate impact. Technical Scope is now based on impact on OAR metrics and whether encountering the risk constitutes 1) a reportable event, 2) an event that DOE simply needs to be informed about, or 3) an event with very low consequences that need not be reported at all.

The board then redesigned the 5x5 steady-state risk matrix and discussed it with ALCF management, who approved the changes. The new matrix, along with explanatory notes, is shown in Figure 5.2.

IMPACTS			INCREASING PROBABILITY →				
Technical Scope*	Cost ($M)		Risk Event Almost Never Occurs < 10%	Risk Event Rarely Occurs 10% - 25%	Risk Event Occurs On Occasion 26% - 74%	Risk Event Occurs Often 75% - 90%	Risk Event Almost Always Occurs > 90%
			VL	L	M	H	VH
Drop below OAR metric target, or a reportable incident	> 2.0	VH	5	10	15	20	25
OAR metric above target but drops by > 4% of target, DOE informed, incident close to reportable	1.0 – 2.0	H	4	8	12	16	20
OAR metric above target but drops by > 2% of target, DOE informed, but somewhat below reportable	0.5 – 1.0	M	3	6	9	12	15
OAR metric above target but drops by > 0.5% of target, DOE informed, but well below reportable	0.1 – 0.5	L	2	4	6	8	10
No measurable impact on OAR metric, not required to inform DOE	< 0.1	VL	1	2	3	4	5

INCREASING IMPACT ↑

*"Technical Scope" includes degradation to 1) computation, 2) security, 3) storage, and/or 4) data, as well as machine downtime. When applicable, it includes impacts on user experience, safety, and cyber security.

Scoring: VL=1, L=2, M=3, H=4, VH=5. Overall risk score is the product of the probability score and the greater of the two impact scores.

Cell colors corresponding to the risk score (Very Low/Low=green, Moderate=yellow, Severe=orange, Critical=red) are used in the steady-state Risk Register to better distinguish the steady-state Risk Matrix risk levels.

Figure 5.2 ALCF Steady-State 5x5 Risk Matrix, revised October 2015

ALCF Risk Review Board

ALCF employs a five-person Risk Review Board to serve in an advisory capacity to ALCF management. The board meets quarterly and makes recommendations to ALCF management regarding steady-state risk management issues. At each meeting, the board:

- Reviews proposed new risks and makes recommendations on adding a proposed risk to the steady-state ALCF risk register.
- Monitors open risks and, for each open risk, reviews any new information on the risk provided by the risk owner and/or the ALCF steady-state risk managers and:
  - Determines whether the risk needs to be recharacterized.
  - Considers whether the risk has been managed and should be closed.
  - Reviews the mitigation actions for the risk and considers if any of the actions need updating.

Risk Management in Day-to-Day Operations

ALCF currently has 40 open risks in the facility operations risk register and uses the post-mitigated risk scoring to rank the risks. These risks include general facility risks (such as funding uncertainties, staffing issues, and safety concerns) and specific risks (such as system component failures, availability of resources, and cost of electricity). On the operations side, subject matter experts estimate risk mitigation costs and use them to inform management reserves.

In addition to formal monthly and individual risk meetings and the Risk Review Board quarterly meetings, ALCF has many informal risk discussions. Risks are identified and evaluated, and mitigation actions developed, for all changes at the facility, from installing a new piece of hardware, to changing the scheduling policy, to upgrading software. If the risks identified are short-term or minor, they are not added to the registry. New significant risks identified during the activity planning are added to the registry and reviewed at the next risk meeting.

Other tools beyond the risk register are used for managing risks in day-to-day operations. An example is the use of Work Planning and Controls (WPCs) and Job Hazard Questionnaires (JHQs) to manage risks for activities where safety is a potential concern. WPCs are primarily used for any non-routine work and are developed in consultation with safety and subject matter experts. JHQs are used for all staff and all contractors and cover all work, both routine and non-routine. During planning meetings for non-routine activities, staff review the planned actions and evaluate possible safety concerns. If a potential risk is identified, detailed discussions with the safety experts are scheduled, and procedures for mitigating the risks are developed, then documented in the WPC. The WPC is then used during the activity to direct the work.

Beyond the operations of the machine, risk management is used in such diverse ways as evaluating and managing INCITE and ALCC proposal risks (the risk of too few proposals, the risk of a lack of diversity across science domains, the risk of too few capability proposals, etc.), safety risks in staff offices, leasing risks, support risks (including the opportunity risk that electricity costs could be lower than budgeted), etc.

5.2 Major Risks Tracked for the Review Year

Since Q4 of FY 2010, ALCF has experienced several eventful years as a result of Mira's transition to operations in FY 2013 and the planned growth of both ALCF staff and budget in order to bring the facility to full strength. As such, ALCF monitored—and continues to monitor—a large number of major risks for the facility. These risks are described in Table 5.1. All risk ratings shown are post-mitigation ratings. Twelve major operations risks were tracked for CY 2015, one with a risk rating of High and eleven with a risk rating of Moderate. Of these, four were encountered and managed. No major risks were retired during CY 2015. The risks are color-coded as follows to assist with reading the table:

- Red risks were encountered and remain Moderate or High risks.
- Orange risks were not encountered but remain Moderate or High risks.

Table 5.1 Major Risks Tracked for CY 2015

ID	Title	Encountered	Rating	Notes
1059	Funding/budget uncertainties	Yes	High	ALCF worked with the program office to plan a budget for handling the impact of a Continuing Resolution, new hires, and changes in laboratory indirect expense rate. This risk remains a major concern as the facility moves forward with ALCF-3.
25	Staffing challenges	Yes	Mod	ALCF added two priority hires and five new hires overall this year, plus one internal transfer. ALCF continues to have staff available who can be re-tasked as needed. With ongoing budget uncertainties and difficulty competing with industry for new hires, staff hiring remains a concern.
1049	Staff retention	Yes	Mod	Between budget concerns at Argonne and the growth in high-paying industry jobs for system administrators and programmers with HPC expertise, ALCF lost two staff members during CY 2015. This remains a concern.
1056	System stability issues due to upgrades	Yes	Mod	An upgrade to FS1 file system resulted in sluggish performance of FS1. This event is discussed in detail in Sec. 5.3.5 below.
1018	INCITE and ALCC users are not provided adequate support by ALCF	No	Mod	ALCF staff is proactive about limiting the chance of encountering this risk by 1) frequently soliciting feedback from project members about the service ALCF provides and 2) actively managing the support expectations of the project members.
1050	ALCF has insufficient disk space to support science needs	No	Mod	Recent storage upgrades have increased space and improved enforcement of quotas.

(continued on page 5-12)

Table 5.1 Major Risks Tracked for CY 2015 (Cont.)

ID	Title	Encountered	Rating	Notes
1054	Catastrophic failure of home file system	No	Mod	This risk has a low probability; in addition, recent file system upgrades mirror the home file system and would allow system restoration in a day or two, which reduces the potential impact.
1076	If the ISSF is decommissioned, we will not have an appropriate facility to host our disaster recovery resources	No	Mod	At present, there is no planned shutdown of the ISSF. Alternative storage locations are being explored.
1085	Diagnostic suite and utilities fail to detect hardware problems	No	Mod	ALCF continues to track and monitor job and hardware failures and correlate the information.
1091	Injury to workers/overall safety of the division	No	Mod	ALCF continues to promote a safety culture at all levels of the division and to follow the Argonne Integrated Safety Management plan. The facility monitors work areas for potential safety concerns and enforces the use of personal protective equipment.
1099	INCITE and ALCC do not use all allocated core-hours	No	Mod	As in the past, only a few projects did not use their full allocation, usually because of staffing issues at the user's institution. Through proactive techniques such as monitoring usage throughout the allocation period, adjusting of scheduler priorities, and frequent communications with the users, the catalysts ensure most projects use their full allocation.

5.3 Risks Encountered in the Review Year and Their Mitigations

The six risks encountered during CY 2015 are discussed below, along with the risk owner, its probability and impacts, a description of the actual problem that occurred, and the management of the risk. The ratings of the risks encountered were as follows: one High, three Moderate, one Low, and one Very Low.

5.3.1 Funding/Budget Uncertainties

1059: Funding/Budget Uncertainties	
Risk Owner	Michael Papka
Probability	High
Impact	Cost: Very Low; Technical Scope: High
Risk Rating	High
Primary Management Strategies	Develop austerity measures. Work closely with DOE sponsors to manage expectations and scope. Plan carefully, in conjunction with program office, for handling Continuing Resolution, leasing costs, and hires. Forward-pay lease to reduce overall leasing costs.
Triggers	ASCR provides funding scenario for budget exercise that is less than planned. Information from DOE indicating a likely extended Continuing Resolution. Argonne laboratory management calls for austerity measures.

Description

The Office of Science might not fund the ALCF budget as planned, or could reduce the ALCF budget below previous funding levels. An extended or full-year Continuing Resolution (CR) could prevent ALCF from receiving planned funding. These scenarios could result in the inability to pay leases, contracts, staff, and to deploy future machines.

Evaluation

During the past year, the Funding/Budget Uncertainties risk was ALCF's highest risk, and it was also one of the risks encountered. The facility was required to operate with moderate austerity measures during the early part of the year. ALCF plans for carry-forward funds each year, with the intention of starting each fiscal year using carry-forward funding from the previous fiscal year while waiting for the first allotment of current-year funding to arrive. This was done in FY 2016, but the funding uncertainty was large enough to also require some budget reprioritizing of purchases and new hires. In addition, the laboratory indirect expense rate model was changed at the start of FY 2016, resulting in an unexpected increase to the burdened cost of materials and subcontracts (M&S) effective October 1, 2015 (see discussion below of Risk 1090).

Management

In conjunction with the DOE-ASCR Budget Deep Dive, ALCF prepared for a full-year CR and reduced budget scenarios. To assure that adequate funds were available to operate Mira and prepare for ALCF-3, ALCF continued moderate austerity measures to provide maximum flexibility for the coming fiscal year.

ALCF continues to closely monitor budget information for FY 2016 and beyond in case of a reduction in funds from the plan of record. Moderate austerity measures remain in place, with spending being prioritized, and these measures may be augmented, depending on the budget.

5.3.2 Staffing Recruitment and Retention Challenges

25: Staffing Recruitment Challenges	
Risk Owner	Michael Papka
Probability	Moderate
Impact	Cost: Very Low; Technical Scope: Moderate
Risk Rating	Moderate
Primary Management Strategies	Evaluate possible additional recruiting avenues. Prioritize staffing needs. Adjust work planning. Retrain staff to meet ALCF needs. Retask staff as needed.
Triggers	Lack of response to job postings. Rejection of job offers. Staff turnover.

1049: Staff Retention	
Risk Owner	Michael Papka
Probability	Moderate
Impact	Cost: Very Low; Technical Scope: Moderate
Risk Rating	Moderate
Primary Management Strategies	Make salaries as competitive as feasible. Identify promotion opportunities. Develop flexible work schedules. Implement flexibility in work assignments.
Triggers	Staff resignations. Staff reports of receiving outside offers.

Description

This is a period of necessary growth for ALCF as it continues to staff up to operate Mira and prepare for ALCF-3. An aggressive staff ramp-up, originally planned for FY 2010 through FY 2012, was extended because of budget reductions. An ALCF risk evaluation identified two key risks associated with this ramp-up, and both occurred in CY 2015 as a result of industry competition for retention of existing employees and potential new hires. The risks have been combined for this discussion, as they are related:

- 25: Challenges encountered in hiring new qualified HPC staff
- 1049: Unable to retain staff due to increased demand for staff with compute expertise and staff worries about DOE funding

Evaluation

As the economy continues to recover, more industry jobs open up for ALCF staff. As a result, in the past year, two ALCF staff left for higher-paying jobs in industry. Five new full-time staff and one staff transfer from another Argonne division were added during CY 2015, for a net gain of +four ALCF staff for the year. Key ALCF leadership positions, the Director of Science and the Director of Operations, were filled during CY 2015. Thus, ALCF has made good progress on adding priority new hires.

Management

Because of industry competition for potential new hires, a limited pool of experienced and available HPC staff, and the fact that candidates do not come out of universities trained for HPC work, it can be very challenging to hire experienced HPC staff. For these reasons, several years ago the ALCF risk management team began preparing to execute mitigations in advance of the occurrence of these risks. When the risks occurred, ALCF was able to continue supporting existing projects successfully even while understaffed.

ALCF has continued to use mitigations to manage both risks over the past year. Facility management continues to replan work as needed, sometimes delaying both planned improvements and lower-priority work. Other mitigation strategies that have been used to address staffing issues include retasking staff, dropping lower-priority tasks, and, when possible, matrixing in staff expertise from other divisions.

By carefully and judiciously managing both risks, ALCF has successfully operated the facility and moved ahead with the ALCF-3 project. However, open positions are often difficult to fill, despite aggressive efforts to find and attract qualified candidates, and there continues to be high demand for the skills of ALCF staff. Thus both staff recruitment and staff retention will remain a focus for ALCF.

5.3.3 Interruptions to Facility That Provides Cooling

30: Interruptions to Facility That Provides Cooling	
Risk Owner	Mark Fahey
Probability	Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Increased redundancy. Site cooling piping is now interconnected to enable chiller plants to provide backup to each other.
Triggers	Temperatures of equipment in the Data Center start rising; planned maintenance; monitoring notification of outage or rising temperatures on equipment.

Description

Data Center cooling provided and maintained by Argonne's FMS Division and TCS Building Management was lost on July 6, 2015 for 2 hours and 16 minutes. One hour after cooling was lost, the ambient air temperature had risen by 15°F, and ALCF Operations staff shut down machine operations to protect the equipment. In total, the ALCF machines were unavailable 9 hours and 19 minutes. For further details, see Item 3 in Section 2.1.1.

Evaluation

Building engineering continues to work with the vendor of the chiller to try to determine what caused the failure. No subsequent chiller failures have occurred. Because the laboratory recently upgraded the site-wide chiller system to interconnect chiller plants around the site, the chiller plants can back each other up, and facilities can recover quickly from the loss of one chiller. In addition, experience indicates that this type of outage occurs only rarely, much less than 10 percent of the time. Thus the overall rating for this risk remains Very Low.

Management

During this event, ALCF Operations staff was able to implement changes to procedures for shutdown and start-up that were practiced and improved on after the June 1 scheduled maintenance. Once ALCF had been informed that there was a malfunction of the chiller plant, Operations preemptively shut down air-cooled systems in order to protect them from potential heat-related damage, meaning that jobs had to be interrupted. E-mail was sent out to all users explaining the event, and ASCR was notified**.** After being notified that the chiller plant was back functioning normally, ALCF staff waited about 30 minutes to ensure that cooling was stable and then began the process of bringing equipment back online. Once all equipment was back in operation, another e-mail was sent out to users, informing them that service to Mira had been restored. ASCR was also notified that ALCF operations had returned to normal.

5.3.4 Changes in Laboratory Indirect Expense (IE) Cost Model or Increase in IE Rates Result in Increased Cost to ALCF

1090: Changes in Laboratory Indirect Expense (IE) Cost Model or Increase in IE Rates Result in Increased Cost to ALCF	
Risk Owners	Michael Papka, Darin Wills
Probability	Very Low
Impact	Cost: Moderate; Technical Scope: Very Low
Risk Rating	Low
Primary Management Strategies	Engage with laboratory management. Track carefully with DOE and laboratory management. Ensure adequate management reserves. Develop increased austerity plan. Work with laboratory management and DOE to request exemptions and assessment modifications. Monitor indirect expense rate recovery.
Triggers	Monthly indirect expense rate recovery report indicates shortfall. Advance notice from CELS office of possible indirect rate increase. Official notification by laboratory management of an increase in the indirect rates. Laboratory funding is decreasing without a corresponding decrease in laboratory indirect costs.

Description

DOE Order 413.2C, Laboratory Directed Research and Development (LDRD), was revised in October 2015. The major impact to ALCF is that M&S costs are now assessed the LDRD indirect expense burden effective October 1, 2015. Prior to FY 2016, M&S costs were not assessed the LDRD IE burden. The result is an increase in the IE costs for ALCF over what had been planned for FY 2016 and future fiscal years.

Evaluation

The impact of the revision to DOE Order 413.2C will result in additional cost to ALCF, estimated to be in excess of $1M for FY 2016. It is expected that current management reserves and project contingency funds can cover the FY 2016 increased cost. Long-term operations funds will be impacted in the out years, resulting in prioritization of ALCF budget line items. For example, decreased available funding may cause the deferral or cancellation of planned equipment purchases, and the scope of some work may be reduced.

Historically, this risk occurs less than 10 percent of the time and is considered very low probability. While the cost of encountering this risk can exceed $1M, mitigation strategies such as applying management reserves and re-budgeting for the out years reduce this amount to between $0.5M and $1M. The risk would, at most, possibly postpone planned purchases or hires, so its impact on facility operation is low. The overall risk rating is therefore Low.

Management
ALCF has held discussions with laboratory management and the DOE to evaluate the budget impact. Management reserves and project contingency funds will be used in FY 2016 to cover the additional cost, if necessary. On February 16, 2016, ALCF presented its annual Budget Deep Dive to the ASCR program office. An out-year budget (thru FY 2021) was prepared, taking into account the new LDRD assessment for FY 2016 and beyond. Per ASCR's request, a number of different funding profile scenarios were presented during the Deep Dive. As in past years, ALCF will continue to work closely with ASCR to monitor spending against available funding levels.

5.3.5 System Stability Issues Due to Upgrades

1056: System Stability Issues Due to Upgrades	
Risk Owner	John Reddy
Probability	Low
Impact	Cost: Very Low; Technical Scope: Moderate
Risk Rating	Moderate
Primary Management Strategies	Perform upgrades on non-critical systems first when feasible. Have a rollback plan in place. Monitor performance closely following upgrade. Work with the vendor to understand the upgrade(s) and the quality control processes. Deep test on Test and Development Systems.
Triggers	Planned system upgrades. System instability observed following system upgrades.

Description
Following upgrade of the firmware and GridScalar on FS1 couplets during a planned preventive maintenance outage, a large number of page allocation errors occurred, and the FS1 file system was sluggish, with some nodes running out of memory. FS1 was therefore isolated and a testing period was scheduled with the vendor and IBM to determine and correct the problem. The problem was identified and corrected, and FS1 was then released back to the users. For further details, see Item 4 in Section 2.1.1.

Evaluation
The basic problem was that GPFS servers were running out of memory. This was caused by using the "connected" mode setting that had been used prior to the upgrade. The OFED package in the new version of GridScalar uses large amounts of memory in this mode, resulting in the servers running out of memory. At the suggestion of the vendor, a switch was made to

using "datagram" mode. After this change, all tests ran correctly, and FS1 has run without problem going forward.

Management

FS1 was kept isolated until the problem was resolved. There were 57 jobs submitted from FS1 that were kept in a held status until the filesystem was put back in production. The affected users were notified by e-mail, and ASCR was notified as well. ALCF worked with the vendors (IBM and DDN) to determine the source of the problem and implement a fix.

Before upgrading the full file system, the upgrade was tested on a test and development system and encountered no problems. However, this smaller system did not use large enough amounts of memory to trigger the memory problem with the OFED package.

5.4 Retired Risks

No risks were retired during the past year.

5.5 New and Recharacterized Risks since the Last Review

Staff operating within the ALCF risk culture regularly identify new risks and recharacterize existing risks. No new risks were added and one risk, Risk 1090, was recharacterized in CY 2015. Table 5.2 lists this risk.

Table 5.2 New and Recharacterized Risks from CY 2015

ID	Title	Rating	Management Strategies	Notes
1090	Laboratory indirect rates could increase	Low	Engage with laboratory management. Track carefully with DOE and laboratory management. Ensure adequate management reserves. Develop increased austerity plan. Work with laboratory management and DOE to request exemptions and special rates. Monitor indirect expense rate recovery.	Risk was reevaluated and recharacterized after being encountered.

This risk was recharacterized because the scoring of both impacts was changed. The scores for probability, technical scope, and cost had previously all been Very Low, resulting in an overall risk scoring of Very Low. The risk has been encountered less frequently than once every ten years, so the probability continued to be scored as Very Low. However, with the estimated cost impact of this encounter exceeding $1M prior to any mitigation, the cost score was changed to High. Encountering this risk could postpone planned purchases or hires, so the technical scope impact on facility operation was increased to Low. These scores combined to an overall pre-mitigated score of Low, with a numeric score of 4. After applying management strategies, the technical scope and cost scores were both reduced one level, to Very Low and Moderate, respectively. This produced a post-mitigated overall score of Low for the risk, with a lowered numeric score of 3.

5.6 Projected Major Operating Risks for the Next Year

Table 5.3 lists the current top operating risks projected for CY 2016 along with the current risk rating and management strategies for the risk. These are the risks that experience has shown are most likely to be encountered in any fiscal year.

Table 5.3 Projected Operating Risks for CY 2016

ID	Title	Rating	Management Strategies
1059	Funding/Budget Uncertainties	High	Develop austerity measures. Work closely with DOE sponsors to manage expectations and scope. Plan carefully, in conjunction with program office, for handling Continuing Resolution, leasing costs, and hires. Forward-pay lease to reduce overall leasing costs.
25	Staffing Recruitment Challenges	Mod	Evaluate possible additional recruiting avenues. Prioritize staffing needs. Adjust work planning. Retrain staff to meet ALCF needs. Re-task staff as needed.
1049	Staff Retention	Mod	Make salaries as competitive as feasible. Identify promotion opportunities. Develop flexible work schedules. Implement flexibility in work assignments.
1091	Injury to Workers/Overall Safety of the Division	Mod	Promote safety culture at all levels of the division. Follow Argonne ISM. Monitor work areas for potential safety concerns. Enforce use of personal protective equipment.

Conclusion

ALCF uses a proven risk management strategy that is documented in its RMP. This document is regularly reviewed and updated to reflect the dynamic nature of risk management as well as new lessons learned and best practices captured from other facilities. Risk management is a part of ALCF culture and applies equally to all staff, from senior management to summer students. A formal risk assessment is performed for every major activity within ALCF, with informal assessments used for smaller activities. Risks are monitored and tracked using the Box cloud-based storage system and risk forms formatted using Excel. Over the past year, no risks were retired, no new risks were added, and one risk was recharacterized. Beyond this, many tools are used to manage risks at ALCF, particularly in the area of safety. ALCF's effective risk management plan has contributed to the successful management of all significant risks encountered in the past year.

This page intentionally left blank.

Section 6. Safety

Has the site implemented measures for safety of staff and the public that are appropriate for HPC/networking facilities?

ALCF Response

ALCF has an exemplary safety record. Since the division's inception in 2006, ALCF has never experienced a lost time incident.

ALCF employs appropriate work planning and control principles. A formal "skill of the worker" document is used for routine tasks. Formal specific procedures are in place for more complex tasks, such as changing out the Blue Gene/Q power supplies (thermal hazard) and node boards (very mild chemical hazard due to water treatment chemicals, weight, and potential damage to hardware), as well as medium-voltage electrical maintenance. The facility performs hazard analysis and creates work planning and control documents for emergency work or when there is an unexpected change to previously planned work.

This page intentionally left blank.

Section 7. Cyber Security

Has the site been certified to operate (cyber-security)?

ALCF Response

Yes. The Argonne Authority to Operate (ATO) includes ALCF as a major application, and it was granted on November 2, 2015. It is valid as long as Argonne National Laboratory maintains robust, continuous monitoring of the Cyber Security Program as detailed in the letter. A copy of the ATO letter follows.

DOE Mail 0815-021 Response
Received: November 2, 2015
SDH (cc: PKK, JPQ, MAS)

Department of Energy
Argonne Site Office
9800 South Cass Avenue
Argonne, Illinois 60439

NOV 0 2 2015

Dr. Peter B. Littlewood
Director, Argonne National Laboratory
President, UChicago Argonne, LLC
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Littlewood:

SUBJECT: RENEWAL OF AUTHORITY TO OPERATE FOR THE ARGONNE NATIONAL LABORATORY INFORMATION TECHNOLOGY INFRASTRUCTURE

References:
1. Letter, J. Livengood to P Littlewood dated August 27, 2015. Subject: Renewal of Authority to Operate for the Argonne National Laboratory Information Technology Infrastructure
2. Letter, J. Livengood to E. Isaacs dated February 8, 2012. Subject: Approval of Authority to Operate for the Argonne National Laboratory Information Technology Infrastructure

Argonne National Laboratory (Argonne) has satisfactorily operated its Information Technology (IT) Infrastructure within Argonne's risk management framework and its FIPS-199 categorization of Moderate. The Laboratory has conducted regular continuous monitoring briefings and has kept me informed of changes in cyber security risk in accordance with the Risk Management Framework. Updates to the Laboratory's Risk Assessment have incorporated changes in threat and risk, as well as the interaction of Argonne's systems with cloud based services. This has demonstrated that the Laboratory's IT Infrastructure is operating at an acceptable level of risk and I am therefore, as the Authorizing Official, renewing Argonne's Authority to Operate (ATO). The ANL IT Infrastructure consists of Argonne's General Computing Enclave, the following five sub-component systems, and all IT investments.

- Accelerator Control Systems (APS and ATLAS)
- Argonne Leadership Computing Facility
- Business Systems
- Sensitive Information
- Cyber Federated Model (CFM)

A component of the Office of Science

This renewal maintains the enclave's FIPS-199 level of Moderate. This ATO clarifies the scope of the ATO issued on August 27, 2015 and will remain in effect as long as the Laboratory carries out continuous monitoring under the Risk Management Framework and there are no significant changes to Argonne's IT Infrastructure. The Laboratory should retain a copy of this letter with the security authorization package. If I can be of any assistance, please contact me or have your staff contact Francis Healy at (630) 252-2827 or e-mail frank.healy@science.doe.gov.

Sincerely,

Joanna M. Livengood
Manager

cc: S. Hannay, ANL-CIS
M. Skwarek, ANL-CIS
M. Kwiatkowski, ANL-CIS
F. Healy, SC-CH
W. Dykas, SC-3

This page intentionally left blank.

Section 8: Summary of the Proposed Metric Values for Future OARs

Are the performance metrics used for the review year and proposed for future years sufficient and reasonable for assessing Operational performance?

ALCF Response

The ALCF and the DOE have agreed to the 2016 metrics and targets as proposed in the March 2015 OAR report with one exception. ASCR has requested that all user facilities use a target of 90 percent for Scheduled Availability for the lifetime of the production resources. ALCF has modified its 2015 and 2016 metric targets to meet this request (see tables 8.1 and 8.2). The proposed metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2017, the proposed metrics and targets for the current production resources remain the same as for 2016. A new system, Theta, is expected to enter production during 2017, and metrics and targets have been added to the 2017 proposed metrics and targets to cover this new system.

The facility should provide a summary table of the metrics and targets agreed upon for the review of Calendar Year 2016 and include the target and actual values of similar metrics used for 2015 for comparison. The facility should also provide metrics and targets under consideration for CY 2017. Those will be finalized later in the year.

The facility should discuss the rationale and use of proposed metrics and targets. This is also a place where a facility can suggest any long term changes in the metrics and targets used for Operational Assessments.

8.1 Overview

The ALCF metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. ASCR has requested that all user facilities use a target of 90 percent for Scheduled Availability for the lifetime of the production resources. ALCF has modified its targets to meet this request. For 2017, the proposed metrics and targets for the current production resource, Mira, will remain the same as for 2016. Appropriate metrics and targets have been proposed for Theta, a resource expected to enter production in 2017. The 2016 metrics are covered in Section 8.2 and the 2017 metrics are covered in Section 8.3.

8.2 ALCF 2015 OA Performance Metrics

The OA performance metrics, 2015 targets and actuals, and agreed upon 2016 targets are presented in Table 8.1. The target for Scheduled Availability has been revised to 90 percent.

Table 8.1 Performance Metrics: 2015 Targets, 2015 Actuals, and Agreed-Upon 2016 Targets

Area	Metric	2015 Targets	2015 Actuals	2016 Targets
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.6/5.0	3.5/5.0
	User Survey – **Response Rate**	25%	45.9%	25%
	% User Problems Addressed Within Three Working Days	80%	95.3%	80%
Business Results	Mira Overall Availability	90%	96.3%	90%
	Mira Scheduled Availability	90%	99.2%	90%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072 – 786,432 cores)	40%	73.4%	40%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144 – 786,432 cores)	10%	31.0%	10%

8.3 ALCF Proposed 2017 OA Performance Metrics

The OA performance metrics, agreed-upon 2016 targets, and 2017 proposed targets are shown in Table 8.2.

Table 8.2 Performance Metrics: Agreed-Upon 2016 Targets and Proposed 2017 Targets

Area	Metric	2016 Targets	Proposed 2017 Targets
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	3.5/5.0
	User Survey – **Response Rate**	25%	25%
	% User Problems Addressed within Three Working Days	80%	80%
Business Results	Mira Overall Availability	90%	90%
	Mira Scheduled Availability	90%	90%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072 – 786,432 cores)	40%	40%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144 – 786,432 cores)	10%	10%
	Theta Overall Availability	N/A	80%
	Theta Scheduled Availability	N/A	90%
	% of INCITE core-hours from jobs run on 20% or more of Theta	N/A	15%

8.4 ALCF Reportable Only Metrics (No Targets)

ALCF has a set of metrics that have no targets and are only reported. These are shown in Table 8.3.

Table 8.3 ALCF Reportable Only Metrics

Area	Metric (No Targets)
User Support Results	Summarize training events and provide examples of in-depth collaborations between facility staff and the user community
Business Results	Report MTTI, MTTF, Utilization, and Usage for the past CY
INCITE Management	Report reviewer survey responses and the proposal allocation results (# of proposals, # of awards, % awarded, # hours requested/awarded, oversubscription) to DOE
Science Results	Track and report the number of publications written annually (projects are tracked for five years after award). Report on at least five significant scientific accomplishments, and the DD awards.
Innovation	Report on innovations that have improved operations

Conclusion

The agreed-upon 2016 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2017, the proposed metrics and targets will remain the same as for 2016 for the current production resources. For the future system Theta, expected to enter production in 2017, a set of metrics and targets in line with past new systems have been proposed for 2017. Achieving the agreed-upon 2016 and the proposed 2017 targets will indicate that the facility is performing up to stakeholder expectations. ALCF anticipates being able to meet all metric targets.

This page intentionally left blank.

Appendix A – Calculations

A.1 Scheduled Availability Calculation Details

Scheduled availability is the percentage of time a designated level of resource is available to users, excluding scheduled downtime for maintenance and upgrades. To be considered a scheduled outage, the user community must be notified of the need for a maintenance event window no less than 24 hours in advance of the outage (emergency fixes). Users will be notified of regularly scheduled maintenance in advance, on a schedule that provides sufficient notification, no less than 72 hours prior to the event and preferably as much as seven calendar days prior. If the regularly scheduled maintenance is not needed, users will be informed of the cancellation of the maintenance event in a timely manner. Any interruption of service that does not meet the minimum notification window is categorized as an unscheduled outage.

A significant event that delays a return to scheduled production will be counted as an adjacent unscheduled outage. Typically, this designation would be assigned for a return to service four or more hours later than the scheduled end time. The centers have not yet agreed on a specific definition for this rare scenario.

Formula:

$$SA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period - time\ unavailable\ due\ to\ scheduled\ outages\ in\ period}\right) \times 100$$

Where:

time in period = *start time* – *end time*
start time = end of last outage prior to reporting period
end time = start of first outage after reporting period (if available) or start of the last outage in the reporting period

A.2 Overall Availability Calculation Details

Overall availability is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages.

Formula:

$$OA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period}\right) \times 100$$

A.3 ALCF Availability Calculations

Here is a simple example.

If a machine had 14 hours of scheduled maintenance and two hours of downtime due to unexpected failures, it had 8 hours of availability (24−14−2 = 8), resulting in 33.3 percent overall availability (8/24). Even though ten hours were scheduled to be available (24−14 = 10), only 8 hours were actually available, and the machine's scheduled availability was 80 percent (8/10).

In its calculations, ALCF tracks availability at the core-second level. The Blue Gene architecture allows an individual node card, containing 32 nodes, to be taken off line to replace one node while the rest of the machine continues to run. However, in calculating availability, ALCF takes into account the ALCF scheduling policy for its large production systems, which does not allow jobs smaller than those using 512 nodes (8,192 cores) to run, which means that 512 nodes is the smallest number of nodes that will be allocated.

Therefore, if a single node were to fail for exactly one hour, it would be recorded as

8,192 cores × 3,600 seconds = 7,372,800 core-seconds of down time.

ALCF has multiple production "scratch" file systems for Mira. Therefore, if all of them are down, the entire machine is considered to be down. If any one of them is available, there are users who can run and therefore, the machine is considered to be available.

The following exception exists. Sometimes, jobs can run successfully even when hardware is considered "down." Examples are test jobs run during a maintenance outage, or a job that was running during a file system outage that didn't attempt any input/output (I/O) while the file system was down, and therefore was able to complete successfully. When this happens, ALCF credits back the core-seconds for those jobs that occurred during the downtime. This credit is made to prevent reporting greater than 100 percent utilization.

To produce the actual numbers, ALCF calculates the scheduled and overall availability on a daily basis. The grand averages for a period are a straight average of the daily results. To produce the bar graph, daily values for the overall availability and the scheduled availability are arithmetically averaged over seven-day intervals, and each bar in the graph represents one of those averages. So, for instance, the first bar in the chart is the average of days January 1–January 7, the second data point is the average of January 8–January 14, etc. If the number of days is not an even multiple of 7, the last data point is handled as follows: If there are more than half (four or more) of the data points, a final data point is calculated from those values and plotted. If not (three or fewer), those values are included in the previous data point, which becomes an average of between 8 and 11 data points. This treatment is performed to avoid significant deviations of the last point because of a small average.

A.4 MTTI Calculation Details

MTTI (Mean Time to Interrupt) is defined as time, on average, to any outage on the system, whether unscheduled or scheduled. It is also known as MTBI (Mean Time Between Interrupt).

Formula:

$$MTTI = \frac{time\ in\ period - (duration\ of\ scheduled\ outages + duration\ of\ unscheduled\ outages)}{number\ of\ scheduled\ outages + number\ of\ unscheduled\ outages + 1}$$

A.5 MTTF Calculation Details

MTTF (Mean Time to Failure) is defined as the time, on average, to an unscheduled outage on the system.

Formula:

$$MTTF = \frac{time\ in\ period - (duration\ of\ unscheduled\ outages)}{number\ of\ unscheduled\ outages + 1}$$

A.6 ALCF Utilization and Capability Calculations

Calculating these values is fairly straightforward. ALCF finds any availability loss as described in the availability section that is for the whole machine; determines how long the loss lasted by wall-time, and whether it was scheduled or not; and then plugs all such losses into the guidance formulas.

***ALCF Utilization Calculation Detail*:** The Cobalt job scheduler writes out job records to the Cobalt database. Each night these data are loaded and processed into a warehouse database that is used to generate usage reports. This warehouse database records the time, date, duration, user, project, and various other system parameters for every job run in the facility. Attributes (INCITE, Discretionary, type of science, etc.) are associated with each project. To calculate the utilization, queries are run against the warehouse database to determine the daily total hours delivered to the various attribute classes and the total hours delivered. Hours for jobs that cross day boundaries are appropriately apportioned to the days. Combining this data with the availability data described in the availability section, the following value is computed on a daily basis:

$$SU = \frac{Core\ hours\ used\ in\ period}{Core\ hours\ available\ in\ period} \times 100$$

Capability Calculation Detail: Except for the data displayed in Figure 2.4 (Mira Job Usage by Size), all data are the sum of the core-hours for qualifying jobs, with the plots showing daily values. Each bar in Figure 2.4 depicts one week's worth of data. Data are summed by type and then divided by the total for the week to determine the percentage. Figure 2.4 has three categories of core-hour usage:

- 0% <= × < 16.7% = Jobs run using up to 16.7 percent of Mira (0 to 131,072 cores);
- 16.7% <= × < 33.3% = Jobs run using from 16.7 percent to 33.3 percent of Mira (131,072 to 262,144 cores); and
- 33.3% <= × <= 100% = Jobs run using 33.3 percent or more of Mira (262,144 to 786,432 cores).

The metrics that exist for Capability are for INCITE only and are explained below:

- Overall Capability = % of INCITE core-hours from jobs run using 16.7 percent or more of Mira (131,072 to 786,432 cores); and
- High Capability = % of INCITE core-hours from jobs run using 33.3 percent or more of Mira (262,144 to 786,432 cores).

Historically, capability has been defined as using greater than 20 percent of the machine. However, 20 percent of Mira would be 9.6 racks, which is not a viable configuration. Hence, the Mira metric was defined in two parts. Overall Capability represents the total of all the jobs in the two categories of Capability jobs. High Capability represents the subset of Overall Capability that is over 33.3 percent of Mira.

Appendix B – ALCF Director's Discretionary Projects

Mira DD Allocations by Project Name, January 1, 2015–December 31, 2015

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
0vbbqrpa	Jun Terasaki	University of Tsukuba	Calculation of Nuclear Matrix Element of Neutrinoless Double-Beta Decay	Physics	3,800,000
3DTransition	Karthik Duraisamy	University of Michigan/John Hopkins University	Direct Numerical Simulations and Machine Learning for Swept Wing Transition	Engineering	1,000,000
ACMEcalibration	Steven Ghan	Pacific Northwest National Laboratory	Multi-Fidelity Calibration of a Climate Model	Earth Science	40,000
Acoustic_Modelling	Sergey Karabasov	Queen Mary University of London	Acoustic Modelling of Jet-Wing Interaction	Engineering	1,572,864
ae-Imperial	Spencer Sherwin	Imperial College London	Pioneering Scale-Resolving Simulations of Flow over Complex Automotive Geometries	Engineering	800,000
ALCF_Getting_Started	Chel Lancaster	Argonne National Laboratory	ALCF Getting Started	Training	150,000
Allinea	Ray Loy	Argonne National Laboratory	Improved Debugging Memory Usage for BG/Q	Internal	2,000,000
alpha-nek	Maxwell Hutchinson	The University of Chicago	DNS of Multi-Mode Rayleigh-Taylor Instability	Engineering	5,048,576
Angora_scaling_study	Allen Taflove	Northwestern University	Angora Scaling Study	Biological Sciences	1,000,000
APS_UBeam_ Dynamics	Michael Borland	Argonne National Laboratory	Beam Dynamics Simulations for the Advanced Photon Source Upgrade	Physics	22,000,000
ARL-KSDFT	Alexander Breuer	U.S. Army Research Laboratory	Scalable All-Electron Structure Calculations	Chemistry	100,000
ATLASQ	Thomas J. LeCompte	Argonne National Laboratory	Grid-Enabling High Performance Computing for ATLAS	Physics	3,000,000
ATPESC15_Instructors	Jini Ramprakash	Argonne National Laboratory	Argonne Training Program on Extreme Scale Computing for All Instructors	Training	2,000,000
ATPESC2015	Paul Messina	Argonne National Laboratory	Argonne Training Program on Extreme Scale Computing	Training	20,000,000
aurora_app	William Scullin, Kevin Harms	Argonne National Laboratory	Aurora Application Enablement	Computer Science	200,000
AytekinALCCPrep	Aytekin Gel	National Energy Technology Laboratory	AytekinALCCPrep	Energy Technologies	100,000
Bachalo_Johson_DNS	Philippe R. Spalart	Boeing	Direct Numerical Simulation of Bachalo-Johnson Transonic Separated Flow	Engineering	6,000,000
backscatter_purdue	Carlo Scalo	Purdue University	Inter-Scale Energy Transfer in High Reynolds Number Turbulent Premixed Flames	Engineering	5,000,000
BGQ_Energy_Profiling	Venkatram Vishwanath	Argonne National Laboratory	Understanding and Characterizing the Power Consumption and Energy Efficiency of Applications on BG/Q	Computer Science	1,000,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
BHAccretion	James Stone	Princeton University	Magnetohydrodynamic Studies of Black Hole Accretion	Physics	500,000
BigDFT4Q	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Simulation of Large Molecular Systems with BigDFT	Chemistry	3,000,000
BIG_MAC	Monica Olvera de la Cruz	Northwestern University	Effective Interactions in Coulombic Systems with Highly Disparate Particle Sizes	Physics	3,145,728
BlumALCCPrep	Thomas Blum	University of Connecticut	Hadronic Light-by-Light Scattering Contribution to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions	Physics	500,000
Boundary_layer_DNS	Gary Coleman	National Aeronautics and Space Administration, Langley	DNS of Turbulent Boundary Layers	Engineering	500,000
CAMD_Toolkit	Dee Dickerson	The Dow Chemical Company	Computer Aided Molecular Design Toolkit	Chemistry	50,000
Camellia	Nathan Roberts	Argonne National Laboratory	Camellia for Discontinuous Petrov-Galerkin Simulations of Incompressible Flow	Physics	5,000,000
CARDIO-PAR	Luca F. Pavarino	The University of Milan	Scalable Domain Decomposition Methods for Computational Cardiology and Isogeometric Analysis	Mathematics	1,300,000
Catalyst	Katherine Riley	Argonne National Laboratory	Catalyst	Internal	20,000,000
cesar_transport	Ronald Rahaman	Argonne National Laboratory	Neutron Transport Performance Studies for CESAR	Nuclear Energy	10,000,000
CESM_Security	Alfred Tang	Fermilab	Simulating Crop Yield and Water Supplies Using CESM	Earth Science	20,000
charmm_zmod	Robert J. Petrella	Harvard University/Harvard Medical School	Highly Parallel Macromolecular Conformational Searches and Energy Evaluations with the CHARMM Program	Biological Sciences	10,000,000
CharmRTS	Laxmikant V. Kale	University of Illinois at Urbana-Champaign	Charm++ and Its Applications	Computer Science	1,000,000
CIBA	Ying Li	Argonne National Laboratory	Collective I/O and Bond Analysis Code Development on SiC Nanoparticle Oxidation	Materials Science	5,000,000
ClimateUncertainty1	Ian Foster	Argonne National Laboratory	Investigation of Initial Condition Uncertainty in Climate Models	Earth Science	5,000,000
cmsframemini	Elizabeth Sexton-Kennedy	Fermi National Accelerator Laboratory (Fermilab)	CMS Framework MiniApp	Physics	10,000
CMT	Scott Parker	Argonne National Laboratory	Compressible Multiphase Turbulence	Engineering	6,000,000
CobaltDevel	Narayan Desai	Argonne National Laboratory	Cobalt Development	Internal	10,000,000
CombDynGTE	Frank Ham	Cascade Technologies, Inc.	Combustion Dynamics in Gas Turbine Engines	Engineering	30,000,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CompBIO	Rick Stevens	Argonne National Laboratory/The University of Chicago	Multiscale Simulations in Biology: Evolution and Ecology of Microbes	Biological Sciences	2,000,000
Compressible_RT	Hussein Aluie	University of Rochester	Multi-Scale Coupling in Rayleigh-Taylor Flows	Engineering	500,000
conquest	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory/London Centre for Nanotechnology/National Institute for Materials Science (Japan)	Linear Scale DFT with Conquest	Materials Science	1,000,000
CONVERGE-BGQ-LDRD	Marta García Martínez	Argonne National Laboratory	Performance Improvement of CFD Code CONVERGE on BG/Q Systems	Engineering	5,000,000
CORALDev	Scott Parker	Argonne National Laboratory	CORAL Development and Testing	Computer Science	2,000,000
CORALtestApps	James Osborn	Argonne National Laboratory	Preparing Test Applications for CORAL Machines	Computer Science	100,000
coreneuron	Fabien Delalondre	École Polytechnique Fédérale de Lausanne (EPFL), Blue Brain Project	Towards Coreneuron Scaling on Full Mira System	Biological Sciences	5,000,000
critical_perf	Kalyan Kumaran, Ray Loy	Argonne National Laboratory	Critical Debugging Project	Internal	50,000,000
DEGAS	Katherine Yelick	Lawrence Berkeley National Laboratory	X-Stack: DEGAS - Dynamic Exascale Global Address Space	Computer Science	10,000,000
DetailedKinetics	William Anderson	Purdue University	Combustion Instability Simulations with Detailed Kinetics for Direct Comparison with Experiments	Chemistry	2,000,000
DiscoveryEngines	Justin M. Wozniak, Rajkumar Kettimuthu	Argonne National Laboratory	Integrating Simulation and Observation: Discovery Engines for Big Data	Materials Science	6,866,782
DNAorigami	Aleksei Aksimentiev	University of Illinois at Urbana-Champaign	Characterization of Self-Assembled DNA Systems	Biological Sciences	5,025,000
DNSGeo	Christos Frouzakis	University of Western Macedonia, Greece	DNS of Forced- and Auto-Ignition in Spherical and Engine-Like Geometries	Chemistry	450,000
DNS_Poggie	Jonathan Poggie	Purdue University	Direct Numerical Simulation of Compressible, Turbulent Flow	Engineering	250,000
DRE_Transition_DNS	Ali Uzun	National Institute of Aerospace	Aircraft Fuel Burn Reduction Using Minute Roughness Elements	Engineering	1,000,000
drugER	Sichun Yang	Case Western Reserve University	Binding Affinity Calculations of Estrogen Receptor against FDA-Approved Drugs with Scalable FEP/λ-REMD Simulations Assisted by a Novel Sampling-Boost Algorithm	Biological Sciences	8,000,000
duanl	Lian Duan	Missouri University of Science and Technology	Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers	Engineering	2,000,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
E2DynamicsNMR	R. Andrew Byrd	National Cancer Institute/National Institutes of Health	Correlating Experimentally Measured Molecular Dynamics with Computational Trajectories: Understanding Dynamic Allostery in Ubiqu	Biological Sciences	5,000,000
Ebola	Charles Macal	Argonne National Laboratory	Agent-Based Behavioral Modeling of Ebola Spread in Chicago and Other Large Urban Areas	Computer Science	500,000
ECN-DNS	Marcus Herrmann	Arizona State University	Primary Atomization DNS of ECN's Spray A	Engineering	1,500,000
elastic	Jaroslaw Zola	University at Buffalo	Similarity Graphs from Large-Scale Biological Sequence Collections	Biological Sciences	2,000,000
ElasticEarthCrust	Mainak Mookherjee	Cornell University	High Pressure Elasticity of Crustal Mineral Albite	Earth Science	2,000,000
ertekin-qmc	Elif Ertekin	University of Illinois at Urbana-Champaign	Many-Body Stochastic Analysis of Semiconductor Bulk and Defect Properties	Materials Science	2,000,000
esumcfd	Cameron Smith	Rensselaer Polytechnic Institute	Extreme Scale Unstructured Mesh CFD Workflow	Engineering	5,000,000
es_tddft	Yosuke Kanai	University of North Carolina at Chapel Hill	First-Principles Simulation of Electronic Excitation Dynamics in Liquid Water and DNA under Proton Irradiation	Chemistry	3,000,000
ExaHDF5	Venkatram Vishwanath	Argonne National Laboratory	ExaHDF5: Advancing HDF5 HPC I/O to Enable Scientific Discovery	Computer Science	5,000,000
EXCEL	Francisco Doblas-Reyes	Institut Català de Ciències del Clima	EXtreme Climate Event Attribution Using Dynamical Seasonal Predictions	Earth Science	300,000
ExM	Justin M. Wozniak	Argonne National Laboratory	Extreme Many-Task Computing with Swift	Computer Science	3,000,000
ExtremeComputing	Sadasivan (Sadas) Shankar	Harvard University	Class on Extreme Scale Computing	Materials Science	5,000,000
Extreme_Scale_TS	William M. Tang	Princeton Plasma Physics Laboratory	Extreme Scale Electromagnetic Kinetic Simulation of Burning Plasmas	Physics	35,000,000
FE2SIM	Axel Klawonn, Oliver Rheinbach	University of Cologne/TU Bergakademie Freiberg	Parallel Multiscale Simulations of Advanced Steel Materials	Mathematics	3,000,000
FLASH_combustion	Praveen Ramaprabhu	University of North Carolina at Charlotte	Numerical Simulations of Turbulent Combustion Using the FLASH Code	Chemistry	2,000,000
FokkerPlanck	Debojyoti Ghosh	Argonne National Laboratory	Fokker Plank Model of a 3-Bus Power System	Energy Technologies	2,621,440
FPMC	Neeraj Rai	Mississippi State University	First Principle Monte Carlo Algorithm Development and Implementation in CP2K	Chemistry	9,000,000
fpnmd	Andre Schleife	University of Illinois at Urbana-Champaign	Electronic Response to Particle Radiation in Semiconductor Systems	Materials Science	15,999,196
Framework_dynamics	Karl Andrew Wilkinson	University of Cape Town	Dynamic Properties of Porous Frameworks upon the Absorption of Gas Molecules	Chemistry	2,000,000
gamra	Sylvain Barbot	Earth Observatory of Singapore	Géodynamique Avec Maille Rafinée Adaptivement	Earth Science	2,000,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
GAtor	Noa Marom	Tulane University	GAtor: A Cascade Genetic Algorithm for Crystal Structure Prediction	Materials Science	10,000,000
GenomeOrganization	Alexey V. Onufriev	Virginia Polytechnic Institute and State University	Investigation of Genome Compaction and Organization with All-Atom and Coarse-Grained MD Simulation	Biological Sciences	1,800,000
GEPW_DYN	Yonduck Sung	General Electric Company	Combustion Instability Study in GE Industrial Gas Turbine	Engineering	1,000,000
GlobalViewResilience	Andrew A. Chien	The University of Chicago/Argonne National Laboratory	Exploiting Global View for Resilience	Computer Science	300,000
graph500	Andrew Lumsdaine	Indiana University	Graph500 Benchmark Run on Intrepid	Computer Science	6,000,000
Gravito-turbulence	Andrew MacFadyen	New York University	Moving Mesh Simulations of Gravito-turbulence in Global Proto-Planetary Disks	Physics	250,000
GRChombo	Hal Finkel	Argonne National Laboratory	Early-Universe Phase Transitions with Strong Gravity and Instabilities in Higher-Dimensional Black Holes	Physics	3,500,000
GTRI_IBM2M_Init	Micheal A. Smith	Argonne National Laboratory	GTRI and NEAMS Related Production Tests and Runs	Nuclear Energy	10,000,000
H3DandVPIC	Ari Le	Los Alamos National Laboratory	Preliminary Tests for Global Kinetic Simulations of Space and Laboratory Plasma Systems	Physics	2,000,000
haiboyu	Haibo Yu	University of Wollongong	Computing the Binding Affinities between PTP1B and Allosteric Inhibitors	Biological Sciences	2,000,000
HEP-ANL	Sergei Chekanov	Argonne National Laboratory	Simulation of High-Energy Particle Collisions Using Monte Carlo Models	Physics	3,000,000
Heterocrystals	Christian Ratsch	University of California-Los Angeles	High-Throughput Computational Design of Heterocrystals	Materials Science	2,000,000
HighAspectRTI	Maxwell Hutchinson	The University of Chicago	Direct Numerical Simulation of the High-Aspect Rayleigh-Taylor Instability	Engineering	2,600,000
HighReyTurb_PostProc	Robert D. Moser	The University of Texas at Austin	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	10,000,000
HPC-Exp-Cancer	Eric Stahlberg	Frederick National Laboratory for Cancer Research	HPC Explorations Supporting Cancer Research	Biological Sciences	5,000,000
HPCTuning	Khaled Ibrahim	Lawrence Berkeley National Laboratory	HPC Applications Tuning	Computer Science	5,000,000
HP_Li	Ying Li	Argonne National Laboratory	High-Performance Li-Air Battery	Materials Science	5,000,000
Hydra_Test	Luigi Capone	Rolls Royce plc	Hydra Test for INCITE Application	Engineering	1,000,000
Hydro_model	Mohamed Sultan	Western Michigan University	Use of GRACE, Remote Sensing and Traditional Data Sets for Modeling Time-Dependent Water Partitioning on Continental Scales	Earth Science	500,000
HyPar-Scalability	Debojyoti Ghosh	Argonne National Laboratory	Scalability of Weighted, Non-Linear Compact Schemes	Engineering	5,500,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
IBM-performance	Kalyan Kumaran	Argonne National Laboratory	Performance	Internal	13,000,000
IBMGSSTest	Kevin Harms	Argonne National Laboratory	IBM GSS Scalability Testing	Computer Science	6,553,600
IME_BlockCoPolymers	Venkatram Vishwanath	The University of Chicago/Argonne National Laboratory	Scalable Data Analysis of Soft X-Ray Scattering for APS Beamline Experiments	Materials Science	100,000
Inverse_Design	Edwin R. Addison	Cloud Pharmaceuticals, Inc.	Inverse Design of Molecules and Drugs	Biological Sciences	1,000,000
KokkosReaxFF2015	Tzu-Ray Shan	Sandia National Laboratories	Porting the Reactive Force Field (ReaxFF) in LAMMPS to Kokkos	Computer Science	1,000,000
LAMMPSopt	Paul Coffman	Argonne National Laboratory	LAMMPS Performance Optimization	Materials Science	3,000,000
LaSco	Gabriel Staffelbach	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique	Large Scale Combustion Preparatory Access	Chemistry	12,000,000
LASPT	Adam Cadien	George Mason University	Liquid & Amorphous Structural Phase Transitions	Materials Science	552,960
LESOIF_2015	Francesco Grasso	DynFluid Laboratory - Arts et Métiers Paris Tech	LES of Shock-Wave Boundary Layer Interaction in Internal Flow with Corner Effects	Engineering	5,800,000
Li-rich	Jeffrey C. Grossman	Massachusetts Institute of Technology, Department of Material Science and Engineering	Simulating Li-Rich Layered Oxide Materials via Quantum Monte Carlo Method	Materials Science	1,000,000
Maintenance	Mark Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	20,000,000
Meso_CCS_DD13	Roberto Paoli	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique	Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus	Earth Science	1,500,000
Metal-sulfur_protein	Sergey Varganov	University of Nevada, Reno	Modeling Nonadiabatic Spin-Forbidden Reaction Mechanisms in Metal-Sulfur Proteins	Chemistry	4,000,000
MHD-turb	Pui-kuen Yeung	Georgia Institute of Technology	Strained and Magnetohydrodynamic Turbulence	Engineering	1,000,000
MiraBootCamp2015	Chel Lancaster	Argonne National Laboratory	MiraBootCamp2015	Training	3,000,000
MM-MEDE	Mauricio Ponga	California Institute of Technology	Multiscale Modeling of Materials under Extreme Dynamic Environments through Large-Scale Computer Simulations	Materials Science	1,000,000
MPACT	Brendan Kochunas	University of Michigan	Michigan Parallel Characteristics-Based Transport	Nuclear Energy	100,000
NanoInterfaces	Giulia Galli	The University of Chicago	Large Scale Calculations on Nanostructured Heterogeneous Interfaces	Materials Science	25,000,000
NanoscaleCombustion	Li Qiao	Purdue University	MD Simulations of Combustion Dynamics in Nanoscale Environments	Chemistry	2,457,600

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
NIMROD_xMHD	Jacob King	Tech-X Corporation	Extended Magnetohydrodynamic Simulations for Burning Plasma Experiments	Physics	500,000
NRCM_DD	V. Rao Kotamarthi	Argonne National Laboratory	Dynamic Downscaling of Climate Models	Earth Science	2,500,000
NUMA	Andreas Mueller	Naval Postgraduate School	Scalability Study for NUMA (Non-hydrostatic Unified Model of the Atmosphere)	Earth Science	5,000,000
Omega-NIF_Exp	Don Q. Lamb	The University of Chicago/Argonne National Laboratory	Mira Simulations of High-Intensity Laser Experiments to Study Turbulent Amplification of Magnetic Fields	Physics	10,000,000
OpenAtom	Laxmikant Kale	University of Illinois at Urbana-Champaign	Scaling and Science Studies of CPAIMD and BOMD	Physics	2,000,000
OpenFOAM-ALCF	Ramesh Balakrishnan	Argonne National Laboratory	OpenFOAM Based Computational Fluid Dynamics Simulations At the Argonne Leadership Computing Facility	Engineering	5,000,000
Operations	Mark Fahey	Argonne National Laboratory	Systems Administration Tasks	Internal	20,000,000
OPV	Kenley Pelzer	Argonne National Laboratory	Mesoscale Modeling of Charge Transport in Organic Photovoltaics	Materials Science	1,000,000
OSCon	Andreas Glatz	Argonne National Laboratory	Optimizing Superconductor Performance through Large-Scale Simulation	Materials Science	150,000
P3DFFT	Dmitry Pekurovsky	University of California-San Diego	Performance Studies of Three-Dimensional Fast Fourier Transforms Using Overlap of Communication with Computation	Computer Science	4,000,000
parallelQD	Bill Poirier	Texas Tech University	Massively Parallel Quantum Dynamics	Chemistry	10,000,000
ParaOpt	Qiqi Wang	Massachusetts Institute of Technology	Parallel Optimization of Turbulent Flow Simulations	Engineering	2,000,000
ParBous_ProcTrans	Susan Kurien	Los Alamos National Laboratory	New Post-Processing Diagnostics of Boussinesq Flow Data	Engineering	500,000
Particle_Flow	Brian Helenbrook	Clarkson University	Examination of Particle-Wall Collision Models in Turbulent Particle Laden Flows	Engineering	1,000,000
Performance	Kalyan Kumaran, Ray Loy	Argonne National Laboratory	Performance	Internal	20,008,000
Petrel	Mike Papka	Argonne National Laboratory	Petrel	Computer Science	100,000
PHASTA_NCSU	Igor A. Bolotnov	North Carolina State University	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	7,000,000
pl_binding_with_fmo	Casper Steinmann Svendsen	University of Bristol	Scaling of the FMO Method for Heterogeneous Systems	Chemistry	1,000,000
PPI_Entropy	Benoit Roux	The University of Chicago	Quantifying Protein-Protein Binding with Greatly Scalable Multiple Copy Algorithms of NAMD	Biological Sciences	8,000,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
PTF_high_ReKa	Guillaume Blanquart	California Institute of Technology	Direct Numerical Simulations of High Karlovitz and High Reynolds Number Turbulent Flames	Engineering	2,000,000
PTPases	Andrei Karginov	University of Illinois at Chicago	Study of Protein Tyrosine Phosphatases	Biological Sciences	5,000,000
Py8HPC	Stephen Mrenna	Fermilab	Py8HPC	Physics	100,000
QBOX	Francois Gygi	University of California-Davis	QBox	Materials Science	5,000,000
QCDHPC	Radja Boughezal	Argonne National Laboratory	Predicting the Terascale On-Demand with High Performance Computing	Physics	4,000,000
QMC-Prep	Anouar Benali	Argonne National Laboratory	QMC Umbrella	Chemistry	10,000,000
qmcspinvdwsolid	Lubos Mitas	North Carolina State University	Quantum Monte Carlo for Spin-Orbit Interactions, Spintronic and Van Der Waals Systems	Physics	2,000,000
QMC_Hyperoxides	John J. Low	Argonne National Laboratory	Quantum Monte Carlo Applied to Lithium Hyperoxides in Li-Air Batteries	Chemistry	10,000,000
QuanPol	Hui Li	University of Nebraska-Lincoln	QuanPol QMMM Style MP2 Simulation Methods	Chemistry	1,000,000
RayBenard_DD	Janet Scheel	Occidental College	Turbulent Rayleigh-Benard Convection at High Rayleigh and Low Prandtl Numbers	Engineering	1,000,000
ReactingRT	Praveen Ramaprabhu	University of North Carolina at Charlotte	New pathways to stability and instability in Rayleigh-Taylor non-premixed flames	Engineering	2,000,000
relRecon	Dmitri Uzdensky, Greg Werner	University of Colorado-Boulder	Kinetic Simulations of Relativistic Radiative Magnetic Reconnection	Physics	1,000,000
RESTMD_DD14	Tom Keyes	Boston University/Louisiana State University	Replica Exchange Statistical Temperature Molecular Dynamics in LAMMPS	Chemistry	2,000,000
rlins	Roberto D. Lins	Universidade Federal de Pernambuco	A Microscopic Perspective on Outer Membrane Remodeling and Antimicrobial Peptide Resistance	Biological Sciences	1,000,000
RNA-stcalculation	Yun-Xing Wang	National Cancer Institute/National Institutes of Health	Computing Three-Dimensional Structures of Large RNA from Small Angle X-ray Scattering Data and Secondary Structures	Biological Sciences	20,000,000
rtflames	Elizabeth P. Hicks	Epsilon Delta Labs	DNS Simulations of Turbulent Rayleigh-Taylor Unstable Flames Using Nek5000	Physics	2,725,000
ScalableMachineLearn	Abhinav Vishnu	Pacific Northwest National Laboratory	Scalable Machine Learning and Data Mining Using MaTEx	Computer Science	2,000,000
SDAV	Mike Papka	Argonne National Laboratory	SciDAC Scalable Data Management Analysis and Visualization	Computer Science	3,250,000
SEGMEnT_HPC	R. Chase Cockrell	The University of Chicago	Anatomic Scale Modeling with a Spatially Explicit General-Purpose Model of Enteric Tissue	Biological Sciences	2,000,000
SENSEI	Venkat Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	3,000,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
shearbands	Haim Waisman	Columbia University	Robust Preconditioners for Shear Bands	Engineering	500,000
Shell_Cat	Leonardo Spanu	Shell International E&P, Inc.	Investigation of Catalytic Properties of Nanoclusters	Chemistry	2,000,000
SiepmannALCCPrep	J. Ilja Siepmann	University of Minnesota	Computations for the Development of the Nanoporous Materials Genome	Materials Science	100,000
SIPs	Murat Keceli	Argonne National Laboratory	Shift-and-Invert Parallel Spectral Transformation Eigensolver	Chemistry	500,000
SMARTDESIGN	J M Cole	University of Cambridge	Smart Material Design for Optoelectronic Applications	Materials Science	5,000,000
sncore	Adam Burrows	Princeton University	3D Core-Collapse Supernova Simulations	Physics	2,000,000
Solar-Eruptions	Ward Manchester	University of Michigan	Simulating Magnetic Solar Eruptions	Physics	500,000
SOWFA	Matt Churchfield	National Renewable Energy Laboratory	Large-Eddy Simulations of a 36-Turbine Offshore Wind Plant	Energy Technologies	5,000,000
SSSPScalability	David Padua	University of Illinois at Urbana-Champaign	Iterative Stepping: A Faster Algorithm and Implementation for Single-Source Shortest Path and Betweenness Centrality	Computer Science	1,000,000
StructuralBiology	Ruth Nussinov	National Cancer Institute, Frederick	Long Time MD Simulation of Protein Structural Function	Biological Sciences	2,000,000
SWIPE2014	Lucian Ivan, Hans De Sterck	University of Waterloo	Large-Scale Simulation of Planetary Environments Using Cubed-Sphere Grids: Software Package Preparation	Physics	500,000
swmf_epic	Daniel Welling	University of Michigan	SWMF MHD-EPIC Simulations of the Terrestrial Magnetosphere	Physics	1,500,000
TACOMA	Brian E. Mitchell	General Electric Company	TACOMA Readiness	Engineering	250,000
TBL_Poggie	Jonathan Poggie	Air Force Research Laboratory	DNS of Compressible Turbulent Boundary Layers	Engineering	5,000,000
Ti-ox_transport	Olle Heinonen	Argonne National Laboratory	Titanium Oxides Transport	Materials Science	10,000,000
TMDC_BILAYERS	Nicholas Hine	University of Cambridge	Transition Metal Dichalgonide Bilayer Heterostructures: Interaction Energies and Interlayer Couplings	Physics	2,500,000
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	15,000,000
TopologyMapping	Zhiling Lan	Illinois Institute of Technology	Topology Mapping of Irregular Applications	Physics	819,200
TotalView	Peter Thompson, Ray Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene/ P	Internal	1,000,000
TRG	James Osborn	Argonne National Laboratory	Tensor Renormalization Group	Physics	100,000
TTC	Tomasz Plewa	Florida State University	Thermonuclear Turbulent Combustion	Physics	500,000
two-phase-flow	Madhusudan Pai	General Electric Company	Towards Petascale First-Principles Simulations of Complex Two-Phase Flow Systems	Engineering	5,900,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
UdSaeroacoustics	Marlene Sanjose	Université de Sherbrooke	LES of Turbulent Jet Noise	Engineering	1,612,800
Uintah_Safety	Martin Berzins	The University of Utah	Solving Petascale Public Health & Safety Problems Using Uintah	Chemistry	1,000,000
umn_crackle	Joseph Nichols	University of Minnesota	Large Eddy Simulation of Crackling Supersonic Jets	Engineering	10,000,000
US-REST2	Sunhwan Jo	Argonne National Laboratory/University of Kansas	Studying Protein-Protein Interaction Using Umbrella Sampling and Solute Tempering	Biological Sciences	4,073,374
VarRhoFlow	Paul E. Dimotakis	California Institute of Technology	Variable-Density Fluid Dynamics	Engineering	3,000,000
VDVAT	Daniel Livescu	Los Alamos National Laboratory	Variable-Density Turbulence under Variable Acceleration	Engineering	1,500,000
Vestas_Park_LES	Gregory Oxley	Vestas Wind Systems A/S	LES Investigation of Stability Enhanced Wake Losses on Large Wind Parks	Energy Technologies	2,750,000
VirtualEye	Marco L. Bittencourt	University of Campinas (UNICAMP)	Computational Modeling of the Human Eye	Engineering	1,000,000
Viruscalculations	Yuri Alexeev	Argonne National Laboratory	Virus Calculations with FMO	Biological Sciences	15,000,000
visualization	Mike Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	2,000,000
Viz_Support	Bill Allcock	Argonne National Laboratory	Visualization Support	Computer Science	8,000
VSL3D	Trung Bao Le	University of Minnesota	Large Eddy Simulations of River Flows	Earth Science	500,000
WATER_SPLITTING	Hanning Chen	The George Washington University	Photocatalytic Water Splitting by TiO_2 Semiconductors	Chemistry	374,684
wdmerger	Maximilian Katz	Stony Brook University	White Dwarf Mergers on Adaptive Meshes	Physics	2,000,000
WindFarmLES	Fotis Sotiropoulos	University of Minnesota	High Fidelity Modeling of Wind Farms in Complex Terrain	Energy Technologies	500,000
windPowerUQ	Larry Berg	Pacific Northwest National Laboratory	UQ Studies in Support of Wind Energy Applications	Energy Technologies	250,000
xFDBenchmarking	Tushar Shethaji	Caterpillar Inc.	Engine Combustion CFD Tool Performance Benchmarking for HPC Readiness	Engineering	800,000
				Total Mira DD	858,334,804