ANL/ALCF-20/06

Argonne NATIONAL LABORATORY

2019
Operational Assessment Report

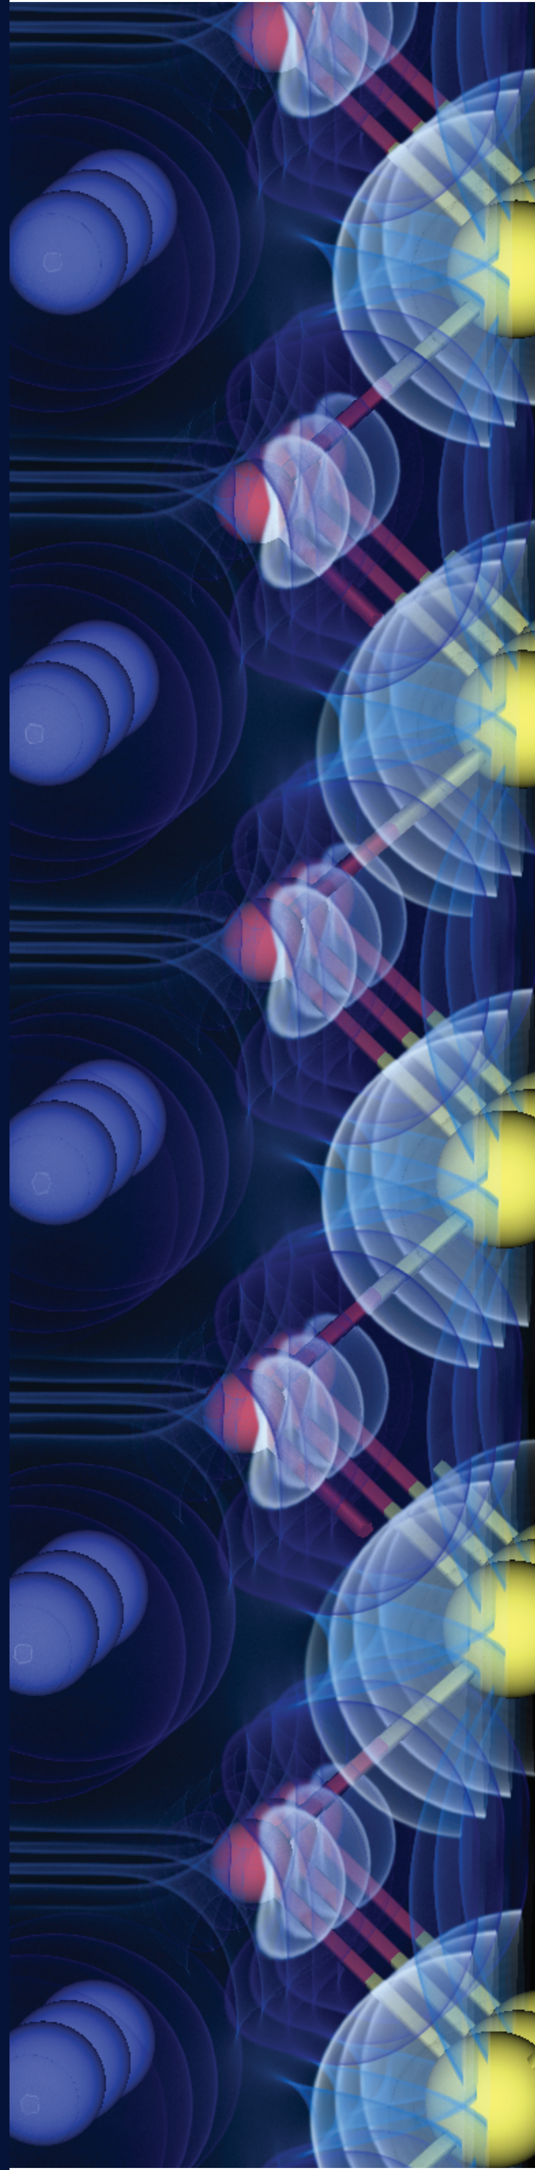

Argonne Leadership Computing Facility

On the cover: *A research team from the University of Illinois at Urbana-Champaign is modeling electronic stopping in three semiconductor materials with different band gaps, or degrees of electrical conductivity, and native defects. This project will establish a predictive computational framework for quantum-mechanical, first-principles modeling of dynamical response of electrons to charged-particle radiation in semiconductors and water/DNA.*

Image credit: *André Schleife, University of Illinois at Urbana-Champaign*

Computing time for this work was awarded through the U.S. Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.

Contents

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Figures

Figures (Cont.)

Tables

Tables (Cont.)

Executive Summary/Background

In 2004, the U.S. Department of Energy's (DOE's) Advanced Scientific Computing Research (ASCR) program founded the Leadership Computing Facility (LCF) with a mission to provide the world's most advanced computational resources to the open science community. The LCF is 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 LCF operates two world-class centers in support of open science at Argonne National Laboratory (Argonne) and at Oak Ridge National Laboratory (Oak Ridge) and deploys diverse petascale machines that are among the most powerful systems in the world today. Strategically, the LCF ranks among the top U.S. scientific facilities delivering impactful science. The work performed at these centers informs policy decisions and advances innovations in far-reaching topics such as energy assurance, ecological sustainability, and global security.

The leadership-class systems at Argonne and Oak Ridge run around the clock every day of the year. From an operational standpoint, the high level of service these centers provide and the exceptional science they produce justify their existence to the DOE Office of Science and the U.S. Congress.

This Operational Assessment Report describes how the Argonne Leadership Computing Facility (ALCF) met or exceeded every one of its goals for calendar year (CY) 2019 as an advanced scientific computing center.

In 2019, the ALCF operated two production resources: Theta, an Intel-based Cray system, and Mira, an IBM Blue Gene/Q, which ceased operation on December 31 after more than seven years and over 700 full-machine runs. At retirement, Mira had delivered nearly 40 billion core-hours to more than 800 projects and left a legacy in terms of system speed and energy efficiency.

Last year, Theta's workload continued to grow, as it delivered a total of 1.2 billion core-hours to 16 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects and 210.6 million core-hours to ASCR Leadership Computing Challenge (ALCC) projects (seven awards for the 2018–2019 ALCC year and 13 awards for the 2019–2020 ALCC year), as well as substantial support to Director's Discretionary (DD) projects (701.8 million core hours). As Table ES.1 shows, Theta performed exceptionally in terms of overall availability (95.0 percent), scheduled availability (99.8 percent), and utilization (97.0 percent; Table 2.1).

In Mira's final year in operation, it delivered 4.2 billion core-hours to 21 INCITE projects and 845.7 million core-hours to ALCC projects (16 awards for the 2018–2019 ALCC year), as well as significant support to a wide range of DD projects (1.1 billion core-hours). Yet again, as Table ES.1 shows, Mira had an excellent year in terms of overall availability (95.6 percent), scheduled availability (99.7 percent), and utilization (93.4 percent; Table 2.1).

Moreover, as of the submission of this document, the ALCF's user community has published 284 papers in high-quality, peer-reviewed journals and technical proceedings. The ALCF also offered a comprehensive program of high-performance computing (HPC) support services to help its community make productive use of the facility's diverse and growing collection of resources.

As the LCF prepares to enter the exascale era in the coming years, ALCF researchers are already leading and contributing to several strategic activities that aim to push the boundaries of what's possible in computational science and engineering. When the 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 (PY) (CY 2019) Metrics

Area	Metric	2019 Target	2019 Actual
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0
	User Survey – **User Support**	3.5/5.0	4.6/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.5/5.0
	User Survey – **Response Rate**	25.0%	41.9%
	% User Problems Addressed within Three Working Days	80.0%	93.3%
Business Results	Theta Overall Availability	80.0%	95.0%
	Theta Scheduled Availability	90.0%	99.8%
	Mira Overall Availability	90.0%	95.6%
	Mira Scheduled Availability	90.0%	99.7%
	% of INCITE core hours from jobs run on 20.0% or more of Theta (51,200–281,088 cores)	20.0%	70.6%
	% of INCITE core hours from jobs run on 60.0% or more of Theta (153,600–281,088 cores)	N/A	10.5%
	% of INCITE core hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	68.8%
	% of INCITE core hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	27.9%

Section 1. User Support Results

Are the processes for supporting the users, resolving users' problems, and conducting outreach to the user population effective?

ALCF Response

The Argonne Leadership Computing Facility (ALCF) has processes in place to effectively support its customers, to resolve problems, and to conduct outreach. The 2019 user survey measured overall satisfaction, user support, and problem resolution, and thereby served both to mark progress and to identify areas for improvement (Table 1.1). User satisfaction with ALCF services remains consistently high as evidenced by survey response data. The following sections describe ALCF events and processes; consider the effectiveness of those processes; and note the improvements that were made to those processes during calendar year (CY) 2019.

Table 1.1 All 2019 User Support Metrics and Results [a]

		2018 Actual	2019 Target	2019 Actual
Number Surveyed		954	N/A	1,246
Number of Respondents (Response Rate)		447 (46.9%)	25.0%	522 (41.9%)
Overall Satisfaction	Mean	4.5	3.5	4.5
	Variance	0.6	N/A	0.5
	Standard Deviation	0.8	N/A	0.7
Problem Resolution	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.5
	Standard Deviation	0.7	N/A	0.7
User Support	Mean	4.5	3.5	4.6
	Variance	0.5	N/A	0.4
	Standard Deviation	0.7	N/A	0.7
		2018 Actual	**2019 Target**	**2019 Actual**
% of Problems Addressed Within Three Working Days [b]		96.1%	80.0%	93.3%

[a] In September 2015, all Advanced Scientific Computing Research (ASCR) facilities adopted a new definition of a facility user based on guidance from the U.S. Department of Energy's (DOE's) Office of Science. Under this definition, a user must have logged in to an ALCF resource during a given time period. This definition of a user provides the basis for all survey results.

[b] The statistical population represented in this metric includes problem tickets submitted from all users. Note that this is a larger population than that of DOE users.

Survey Approach

In 2017, the ALCF worked with a consultancy to redesign its annual user survey. The new, shorter survey omitted workshop-related questions and retained those that comprise the DOE metrics for the Operational Assessment Report (OAR). (The facility now polls workshop attendees separately throughout the year.) The 2019 user survey was similar to the 2018 survey, with the addition of one new question to capture user preferences about features in a potential new ALCF dashboard; the removal of a question around multi-facility users that was no longer relevant; and minor modifications to the infrastructure, performance, and debugging tools/services that were listed in certain questions.

The 2019 survey was administered by Marketing Synergy, Inc., and contained 24 questions. The survey and companion e-mail campaign ran from November 22, 2019 to December 31, 2019 and again for 15 days in February 2020, reaching a smaller group of users who were not included in the first wave. Each reminder e-mail included a unique link to the online survey. Most respondents were able to complete the survey in 10 minutes or less.

Likert Scale and Numeric Mapping

Almost all Likert Scale questions in the ALCF user survey use a six-point scale. This is a standard way to rate user responses for surveys because (1) it provides a symmetric agree-disagree scale; (2) it can be mapped to a numeric scale; and (3) given a certain sample size, it can be used with a normal distribution to obtain useful statistical results. The method also allows for use of off-the-shelf statistics functions to determine variance and standard deviation.

The ALCF follows a standard practice and maps the Likert Scale in this way or similar:

Statement	Numeric
Strongly Agree	5
Agree	4
Neutral	3
Disagree	2
Strongly Disagree	1
N/A	(No Value)

Only the Overall Satisfaction question applied a different point scale, as follows:

Statement	Numeric
Excellent	5
Above Average	4
Average	3
Below Average	2
Poor	1

Beginning in 2017, some of the non-metric survey questions were revised to capture sentiments about various aspects of the ALCF's user services that used the options below:

Select all that apply.
Praise
Suggestions for Improvement
Average
Below Average
Poor

Comments

1.1 User Support Metrics

In CY 2019, 1,246 individuals met the definition of a facility user and were invited to complete the annual user survey. Of this group, 522 responded, for a 41.9 percent response rate. The ALCF surpassed all targets for the survey metrics.

Table 1.2 shows user survey results grouped by allocation program. While Innovative and Novel Computational Impact on Theory and Experiment (INCITE) users reported higher average Overall Satisfaction than ASCR Leadership Computing Challenge (ALCC) and Director's Discretionary (DD) users, the results are not statistically significant. Other metrics are comparable, in that the variations are statistically insignificant.

Table 1.2 2019 User Survey Results by Allocation Program

2019 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		361	192	553	693	1,246
Number of Respondents		172	95	267	255	522
Response Rate		47.6%	49.5%	48.3%	36.8%	41.9%
Overall Satisfaction	Mean	4.7	4.4	4.6	4.5	4.5
	Variance	0.3	0.7	0.5	0.5	0.5
	Standard Deviation	0.6	0.8	0.7	0.7	0.7
Problem Resolution	Mean	4.6	4.4	4.6	4.5	4.5
	Variance	0.4	0.6	0.5	0.5	0.5
	Standard Deviation	0.6	0.7	0.7	0.7	0.7
User Support	Mean	4.7	4.5	4.6	4.5	4.6
	Variance	0.3	0.5	0.4	0.5	0.4
	Standard Deviation	0.5	0.7	0.6	0.7	0.7
All Questions	Mean	4.7	4.5	4.6	4.5	4.5
	Variance	0.3	0.6	0.4	0.5	0.5
	Standard Deviation	0.6	0.7	0.7	0.7	0.7

As Table 1.3 shows, in 2019, the ALCF again exceeded the targets for overall satisfaction and user support.

Table 1.3 2018 and 2019 User Support Metrics

Survey Area	2018 Target	2018 Actual	2019 Target	2019 Actual
Overall Satisfaction Rating	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0
Avg. of User Support Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.6/5.0

1.2 Problem Resolution Metrics

Table 1.4 shows the target set for the percentage of problem tickets addressed in three days or less, which the ALCF exceeded. A ticket is defined as "addressed" once the following conditions have been met: (1) the ticket is accepted by a staff member; (2) the problem is identified; (3) the user is notified; and (4) the problem is solved, or it is in the process of being solved.

Table 1.4 Tickets Addressed Metric

	2018 Target	2018 Actual	2019 Target	2019 Actual
% of Problems Addressed Within Three Working Days [a]	80.0%	96.1%	80.0%	93.3%
Avg. of Problem Resolution Ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.5/5.0

[a] The statistical population represented in this metric includes problem tickets submitted from all users. Note that this is a larger population than that of DOE users.

1.3 User Support and Outreach

1.3.1 Tier 1 Support

1.3.1.1 Phone and E-mail Support

In 2019, the ALCF answered and categorized 5,582 support tickets. There was a significant decrease in the total number of support tickets in CY 2019 from the previous year. The biggest decreases were in Accounts, Access, and Bounces categories. There are four possible explanations for this change: First, with the introduction of the new account and project management system (UB3), some of the tickets that we would normally see were eliminated because of streamlined processes put in place. Second, in mid-2019, Mira was taken off of the ALCC program, shifting the projects awarded for the 2019–2020 ALCC allocation year to Theta alone. We normally expect to see an uptick in account project tickets associated with each new ALCC program allocation year, which begins July 1. With fewer new projects in the ALCC program, there were fewer onboarding tickets. The third reason is the rollout of mobile tokens to the ALCF users. Users who chose mobile tokens could enroll and activate the tokens on their own once the token was provisioned to them. The self-service nature of activating a mobile token eliminated the need for users to call the Support number to unlock their token and receive a temporary personal identification number (PIN). Nearly 76 percent (75.9 percent) of new users opted for mobile tokens. Last, UB3 required the use of the multifactor authentication token to log in to the account and project management portal instead of a password, which was the case for the older portal (UB1). As a result, tickets related to resetting portal passwords were completely eliminated.

Table 1.5 Ticket Categorization for 2018 and 2019

Category	2018	2019
Access	1,323 (18%)	1,010 (18%)
Accounts	2,948 (41%)	2,414 (43%)
Allocations	621 (9%)	629 (11%)
Applications Software	236 (3%)	197 (3%)
Automated E-mail Responses	640 (9%)	348 (6%)
Compilers	73 (1%)	23 (1%)
Data Transfer	37 (1%)	38 (1%)
Debugging and Debuggers	37 (1%)	11 (1%)
File System	221 (3%)	124 (2%)
HPSS [a] and Quota Management	211 (3%)	130 (2%)
Libraries	16 (0%)	36 (1%)
Miscellaneous	167 (2%)	87 (1%)
Network	21 (0%)	11 (0%)
Performance and Performance Tools	46 (1%)	42 (1%)
Reports	289 (4%)	278 (4%)
Scheduling	267 (4%)	196 (3%)
Visualization	19 (0%)	8 (0%)
TOTAL TICKETS	7,172 (100%)	5,582 (100%)

[a] HPSS = high-performance storage system.

1.3.1.2 Improved Account Renewal Communication

The automated account renewal e-mails now list all of the user's active projects and associated details for each, such as project name, PI name, PI e-mail, allocation type, and allocation end date. This information provides an at-a-glance summary of the user's projects and the associated account renewal request approver.

1.3.1.3 Allocation Expiration Notification E-mail

Project PIs receive automated e-mail notifications prior to their project allocation's expiration date. The automation scripts have been updated to eliminate such notifications in instances when the project has been renewed.

1.3.1.4 Automated Account Deactivation

Deactivation of user accounts, including expired Foreign National Access accounts without a renewal underway, is now performed automatically, rather than manually. Additional account verification checks have also been implemented to ensure compliance with DOE and Argonne National Laboratory (Argonne) policies for foreign national users.

1.3.1.5 New Account and Project Management System (Userbase3 or UB3)

In April 2019, ALCF replaced its legacy Userbase1 (UB1) account and project management system with Userbase3 (UB3). On the user-facing side, the new application has a modernized interface for better usability and transparency. Behind the scenes, UB3 contains more security features and presents a more complete view of user and project details than UB1, and has greatly improved the account setup process and reduced the number of repetitive workflows.

Overall, UB3 provides a very powerful single interface to accomplish most administrative tasks and eliminates the need to use various webpages and command line interface (CLI) commands. A dashboard view of all requests in the pipeline gives ALCF administrators a quick summary of pending tasks and an easy way to execute them.

Best practices and improvements that were implemented in UB3 include the following:

- Integration with Argonne's Foreign Assignment/Visit Online Report System (FAVOR), the Laboratory's internal database for approving and tracking foreign nationals. The previous application relied on manually entered (and often error-prone) data transfers between the ALCF and FAVOR and stored personally identifiable information (PII). The new application points users to forms within FAVOR, bypassing the entry of PII entirely, and permits ALCF staff to retrieve relevant data from the laboratory-wide system, via an application programming interface (API). Argonne's Business and Information Services (BIS) FAVOR team is recommending this approach to other divisions at Argonne. The ALCF will capture our process and lessons learned so we can help other divisions leverage this work as needed.
- Clear communication about subscription to facility updates.
- Software development best practices considered and implemented in UB3, which include the following:
    - All functionality is available via RESTful APIs allowing standardized access to the UB3 data from other applications and scripts.
    - Enhanced security by role-based security authentication and authorization. This protection extends to all RESTful API calls, allowing more granular security control.
    - Two-factor authentication, replacing passwords.
    - Responsive and reactive web design.

- Integration with FAVOR:
  - Ability to access and download form ANL-593 (Argonne Foreign National Registration) for users. See Section 7.3 for more detail on Argonne procedures for collecting information about foreign nationals who require site access or remote computer access.
  - Ability to activate/deactivate accounts in a timely manner without full manual intervention.

Below is a summary of the salient features of UB3's software:

- Multifactor authentication provides an additional security layer and eliminates the need for another separate password. The same MFA credentials can be used to access ALCF resources and the website.
- It is integrated with FAVOR. Relevant data that the ALCF needs are brought back to store and display within UB3 using APIs. The data are verified and updated multiple times a day by scheduled jobs.
- Users can now save and retrieve a partially completed account request prior to submission.
- Enhanced search function includes the type-ahead feature for easier project lookup, replacing the project database scrolling menu.
- Users can now opt for a physical token or a mobile token by changing their preference in UB3. This eliminates the need to send a support ticket.
- Project PIs can manage project membership much more effectively via the website in UB3, reducing the need to e-mail User Support.
- Addition of proxy designation. In UB1, proxies and PIs were stored and displayed in the same field. In UB3, the distinction is pronounced.
- System messaging conveys user actions, changes in job states, and errors to the user.
- Users are automatically added to the UNIX group and relevant systems upon project membership approval.
- Enhanced logging provides a log of not only user actions, but system actions too. Clear tagging in the logs makes it easier to search for error-related messages.
- UB3 was developed using RESTful APIs. All inputs to the server are type checked, value validated, and size checked. The signature is consistent for usage as CLI or Python API calls.
- Real-time notifications were built into UB3 that are used to trigger additional workflows. Previously, the ALCF team relied on scripts that were not part of UB3 but were set up to run daily to trigger the same workflows.
- UB3 has role-based access for administrators which allows for hierarchical access. UB1 had a single role for administrators which gave them read-write permissions to all of the interfaces.

1.3.1.6 DOE User Project Data Collection

The 2019 DOE User Project data collection survey was sent to all active users and PIs of active projects in the 2019 fiscal year (FY). The survey tool was developed by ALCF's Advanced Integration Group (AIG) leveraging the UB3 codebase and replaced the vendor-developed and supported survey model previously used. Bringing the collection method in-house allows the ALCF to be more agile and effective in soliciting responses from its users and greater control over when and how the survey is administered. In contrast to how survey data were collected, transferred, and stored in previous years, the 2019 survey responses were directly captured in UB3, which is ALCF's source of user and project information.

1.3.1.7 Getting Started Videos

The ALCF's Getting Started videos are available to all users, on-demand, via ALCF's website. The training videos cover system specifications, code building, storage, operating system (OS) and file systems, compilers, tools, queues, and other topics.

1.3.2 Application Support

1.3.2.1 Individual Projects

NAMD Code Enhancements for Biological Applications

This DD project is an ongoing ALCF-led NAMD code development effort funded by the National Institutes of Health (NIH). During 2019, a novel hybrid topology setup was implemented in NAMD for high-throughput free energy computations. In addition, a generic sampling enhancement code was implemented and then optimized on both central processing unit (CPU) (Theta) and graphical processing unit (GPU) architectures. This project has resulted in three publications: (1) Wei Jiang et al., "Reduced Free Energy Perturbation/Hamiltonian Replica Exchange Molecular Dynamics Method with Unbiased Alchemical Thermodynamic Axis," *J. Phys. Chem. B* 122, (2018): 9435; (2) Wei Jiang, "Accelerating Convergence of Free Energy Computations with Hamiltonian Simulated Annealing of Solvent (HSAS)," *J. Chem. Theory Comput.* 15, no. 4 (2019): 2179; and (3) Wei Jiang, Christophe Chipot, and Benoît Roux, "Computing Relative Binding Affinity of Ligands to Receptor: An Effective Hybrid Single-Dual-Topology Free-Energy Perturbation Approach in NAMD," *J. Chem. Inf. Model.* 59, no. 9 (2019): 3794.

Shedding Light on Singlet Fission for More Efficient Solar Cells

This DD project discovered the anisotropic singlet fission and its vibronic fingerprint in single-crystalline hexacene for more efficient solar energy conversion. These achievements further resulted in a new industry outreach project—a molecular single-crystal study—with Honda Research Institute. ALCF staff provided comprehensive support of this DD project, including CP2K code porting/usage and parallel performance benchmarking on Theta under various combinations of message passing interface (MPI) ranks and OpenMP threads. This project resulted in two publications: Gang-Hua Deng et al., "Vibronic Fingerprint of Singlet Fission in Hexacene," *J. Chem. Phys.* 151, 054703 (2019); and Dezheng Sun et al., "Anisotropic Singlet Fission in Single Crystalline Hexacene," iScience 19 (2019): 1079.

Tokamak Simulation Captures Important Phenomenon at Plasma Edge

The physics of the plasma in a tokamak configured with a divertor changes in important ways near the magnetic X-point. Recent plasma simulations using the 5D gyrokinetic particle code XGC1 revealed the development of an electrostatic structure around the X-point. This structure had been observed in experiments before and could have an important implication to fusion reactor performance, but its physical origin has not been understood. Realistic simulations at the fundamental physics level can provide the needed insight into these complicated issues. This 2017–2018 ALCC allocation supported the advancement of an ongoing study, using XGC1 on Theta and other leadership-class machines, of the fundamental physics underlying the heat deposition on the divertor target plates and the related pedestal structure issue in present tokamak devices and in ITER. Outcomes of this project will be higher confidence extrapolations of theory to practice, and increase the likelihood of success of large fusion experiments such as ITER. The team's findings resulted in one publication: Choong-Seock Chang, Seung-Ho Ku, and Robert M. Churchill, "X-point Ion Orbit Physics in Scrape-off Layer and Generation of a Localized Electrostatic Potential Perturbation around X-point," *Phys. Plasmas* 26, 014504 (2019). ALCF staff worked on code optimization for Theta and development for Aurora.

Convergence of Deep Learning (DL) and High-Performance Computing (HPC) for the Construction of Galaxy Catalogs in the Dark Energy Survey

This DD project developed a novel combination of DL methods to provide a highly accurate approach to classifying hundreds of millions of unlabeled galaxies. The team's findings were published in *Physics Letters B* for the construction of galaxy catalogs in the Dark Energy Survey (DES). The team's approach involved training deep neural networks using human-labeled, high-quality Sloan Digital Sky Survey (SDSS) galaxy images. This project leverages what SDSS has already done by using the Galaxy Zoo images as seed information to classify objects in DES. Furthermore, this project illustrates the use of deep transfer learning in cosmology by using a neural network model that is the state of the art for computer vision and pre-trained with millions of high-quality, real-object images divided into several thousand classes. ALCF staff provided comprehensive support including code porting/usage on Theta, parameter preparation for training, and visualization support. This study is the first of its kind in the literature in which deep transfer learning is combined with multi-node training to reduce the training stage from five hours to just eight minutes. This work resulted in one publication: Asad Khan et al., "Deep Learning at Scale for the Construction of Galaxy Catalogs in the Dark Energy Survey," *Physics Letters B* 795 (2019) 248.

Identification of Coupling Mechanisms Between Ultra-Intense Laser Light and Dense Plasmas

The success of petawatt (109-watt) laser facilities, which are presently under construction around the world and which aim at understanding and controlling these promising particle and light sources, will rely on the strong coupling between experiments and large-scale simulations with particle-in-cell (PIC) codes. In this 2018 INCITE work, the team solved one of the major issues of the last decades in this research field by shedding light on the coupling mechanisms at play in the interaction of ultra-intense laser pulses with dense plasmas. Understanding these coupling mechanisms is expected to give insight into improving the efficiency of elementary physical processes (e.g., ion acceleration) that are present in laser-solid interactions. ALCF staff worked with a team from Lawrence Berkeley National Laboratory (LBNL) and CEA (French Atomic Energy Commission) to resolve Python issues related to the installation and running of

WARP+PICSAR on Mira, assisted with bundling calculations into 8,192-node capability jobs to improve queue throughput, and helped to build a shared version of the HDF5 library. On Cooley, ALCF staff helped to compile and install a Visit plugin for visualization. This work resulted in two publications: Ludovic Chopineau et al., "Identification of Coupling Mechanisms between Ultraintense Laser Light and Dense Plasmas," *Phys. Rev. X* 9, 011050 (2019); and Henri Vincenti and Jean-Luc Vay, "Ultrahigh-order Maxwell Solver with Extreme Scalability for Electromagnetic PIC Simulations of Plasmas," *Comput. Phys. Comm.* 228, 22-29 (2018).

Large-Eddy Simulation of Commercial Transport Aircraft Model

This INCITE study examined the performance of wall-modeled large-eddy simulation (LES) in flows over realistic aircraft geometry using state-of-the-art methodologies for low-dissipation LES and an equilibrium wall model, particularly for the prediction of flow separation encountered at high angles of attack. In year one of the project, simulations of two relevant aircraft geometries at high angles of attack were considered: a high-lift Japan Aerospace Exploration Agency (JAXA) Standard Model (JSM) and the National Aeronautics and Space Administration (NASA) Common Research Model (CRM). Predictions of global forces and surface pressure distributions by two low-dissipation LES codes were shown to be in reasonable agreement with experimental measurements. The insight from this campaign will be applied to the fine grid calculations in the second year (i.e., 2019) of the INCITE allocation. The ALCF Catalyst provided the PI's team with inputs to organize their workflow in the queues, and in setting up reservations during the CTR-Stanford Summer program in July 2018 and before the APS-DFD Meeting held in November 2018, during which time, the comparisons between the Alya and CharLES flow solvers were performed for the JSM model.

Astrophysical Particle Accelerators: Magnetic Reconnection and Turbulence

Observations of highly energetic astrophysical systems show signatures of populations of extremely relativistic particles with vastly more kinetic energy than the ambient temperature of their surroundings would imply. Identifying the physical processes that transfer energy into those particles is an active area of research, but one suspected mechanism is the acceleration of charged particles by turbulence in plasma environments. Developing a good understanding of how turbulent plasmas can produce nonthermal particle acceleration drives the use of large-scale three-dimensional (3D) numerical simulations. However, even the largest supercomputers will be incapable of simulating real astrophysical system sizes in the foreseeable future; thus, identifying the scaling and convergence behavior is critical. This INCITE research team generated a series of simulations, using up to 32,768 nodes of Mira, to show that the energetic particle distributions do converge with increased system size, which justifies extrapolating the simulation results to astrophysical scales. They also found that the time required to converge depended on system scale in a manner consistent with Fermi acceleration, and concluded that turbulence could explain nonthermal spectra observed in systems like the Crab Nebula. ALCF staff have provided comprehensive support for this and the PI's previous INCITE work involving the same system and the same code to explore closely related science. This work resulted in one publication: Vladimir Zhdankin et al., "System-size Convergence of Nonthermal Particle Acceleration in Relativistic Plasma Turbulence," *ApJ* 867, no. 1, L18 (2018).

1.3.3 Resource Support

1.3.3.1 Cgroups on Theta MoM Nodes

ALCF staff has installed control groups (cgroups) on Theta to manage and restrict how many resources a user can consume on the service nodes used to launch jobs, or MoM nodes, where users often run them out of memory and kill the nodes where our Cobalt support processes are running. The result of this action is more reliable MoM nodes.

1.3.3.2 Containers

Singularity is fully supported and working across Theta and Iota. We are currently in the process of testing Singularity 3.0, the next major iteration of the container client. We also have an internal singularity registry, which allows users to store and retrieve their own container images.

1.3.3.3 Cobalt V2

After testing various resource managers, ALCF's Cobalt team approached Altair Engineering (the primary commercial entity behind PBS Pro) about the possibility of extracting and building on their resource management components (the MoM); however, after reviewing Altair's and ALCF's similar roadmaps, the Cobalt team opted to contribute to Cobalt V2 via the PBS Pro community's open source project.

For the ALCF, the biggest advantages of this strategy relate to risk mitigation: PBS Pro is an established, open-source project with a broad user base; Altair is already working on adding Shasta support—a requirement for Aurora—and PBS's Shasta support will be out and running on Shasta machines prior to Aurora; PBS Pro tests are extensive and run on a broad range of systems, so ALCF's contributions will be leveraged by a much wider audience; and PBS's documentation is excellent, whereas Cobalt's has traditionally been weak. While this approach will likely delay certain architectural changes that ALCF wants, it affords the ALCF more flexibility overall.

Our initial efforts will be to mitigate the impact of the changes to our code base and internal systems on our users, which include the following:

- Modifying the core PBS server to add event hooks (notifications of admin commands, node events, etc.), then writing hook event processors to push the data out in the machine-independent format we use to drive the rest of our reporting infrastructure.
- Integrating our scheduling algorithms as site customizations.
- Modifying the PBS Pro reservation system to more closely mimic Cobalt V1 behavior. In Cobalt V1, queues, jobs, and reservations are all separate, first-class entities, which are then bound together. Previously, a user could release a reservation without disrupting running jobs or queued jobs. In PBS Pro, the reservation queue, and all the jobs within it, are subordinate to the reservation; therefore, when a reservation ends, all running jobs are killed and the jobs left in the queue are lost.

- Developing migration tools. Cobalt and PBS Pro CLI and script directives are similar, but not identical, so existing scripts and CLI commands, particularly the more complicated ones, will require some modifications. We are working on scripts to translate CLI commands and scheduler scripts, as well as documenting the specific differences, to help in instances where the automated conversion fails.
- The first new “feature” (both to our users and to the PBS Pro community) will be the incorporation of data staging via Globus Online into job workflows. PBS Pro already provides data staging, but it is blocking and performed via SCP. We are working on how to use Globus Online and how to handle its asynchronous nature. We currently have a prototype and believe we can handle it via existing PBS Pro capabilities. Successful implementation will be a function of us providing non-compute node resources for running the monitoring tasks, and then documenting how to use that in concert with the PBS Pro dependency functionality.

1.3.3.4 Web Reports: INCITE Real-Time Data

ALCF’s Science team and scheduling committee makes job priority decisions given the most up-to-date data. Theta’s job data comes from three data sources: sbank, the operational data processing system (ODPS), and Cobalt. Sbank provides charged data, that is, up to 2 days out of date. ODPS data is up to 1 day out of date. Cobalt data is real time. Planning based on all three data sources is difficult.

ALCF’s Business Intelligence (BI) team captures and stores an extensive amount of job data from multiple sources including sbank, ODPS, and Cobalt. All data resides in the BI team’s data warehouse. The Theta 7-day web report combines all three data sources into one graph, which updates continuously. The following fields display data to aid decision makers:

- INCITE % Machine: The percent of the entire machine that INCITE jobs used over the past 7 days.
- INCITE % Needed: The percent of the entire machine that INCITE jobs *need* to use over the remainder of the calendar year to meet INCITE’s allocation.
- INCITE % of Usage: The percent of all the usage on the machine that INCITE jobs used over the past 7 days.
- INCITE % Queue: The percent of core-hours in the job queue that are INCITE core-hours.

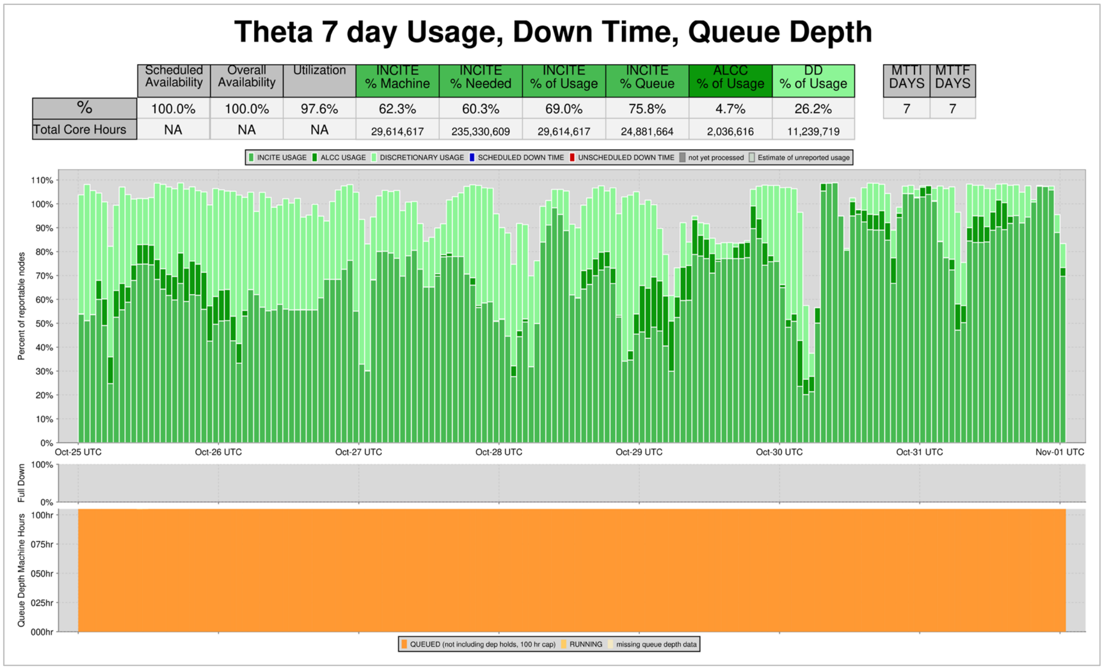

Figure 1.1 Screenshot of Theta usage, scheduled downtime, unscheduled downtime (7 days) on 2019-11-01.

In the above example from November 1, 2019, INCITE used 62.3% of the entire machine over the past 7 days; INCITE needs to use 60.3% of the machine for the rest of the year to meet its allocation goal; and 75.8% of the core hours in the job queue are INCITE core-hours.

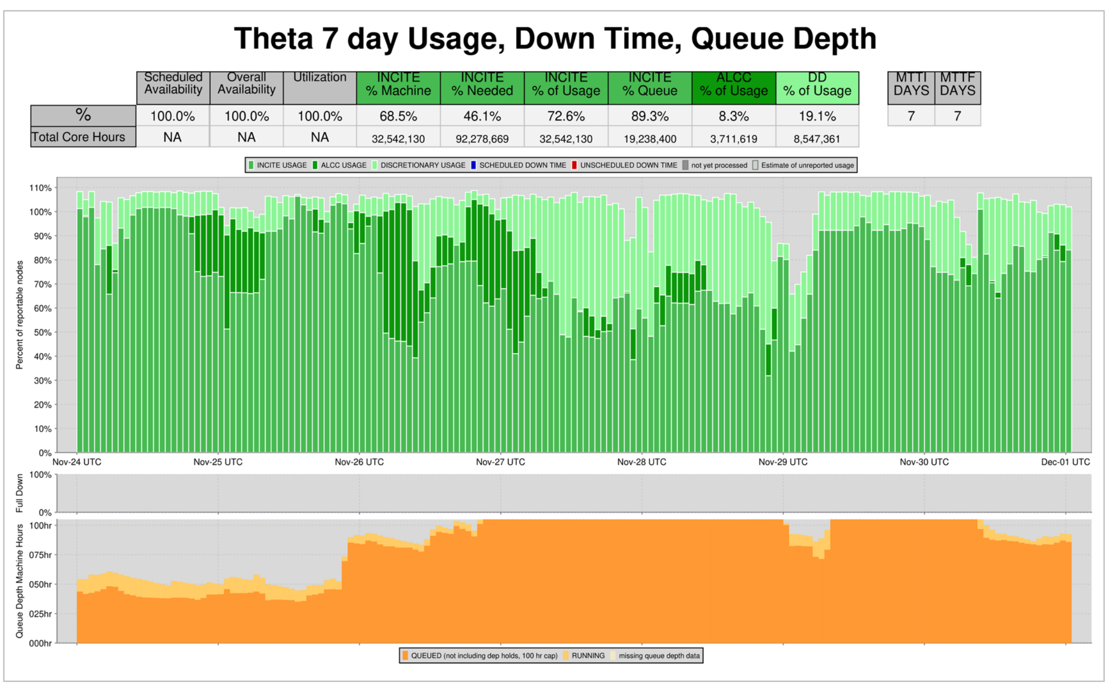

Figure 1.2 Screenshot of Theta usage, scheduled downtime, unscheduled downtime (7 days) on 2019-12-01.

The web report values helped the ALCF Science team and scheduling committee take actions over the month of November to increase INCITE usage. As a result, the INCITE % Needed of the machine decreased from 60.3% to 46.1% on December 1, 2019, as shown in Figure 1.2. The INCITE % Queue gave real-time updates about INCITE queue jobs.

1.3.3.5 RAM Area Network (RAN)

The ALCF has been conducting a small, multiyear research project with Chicago-based tech company Kove to look at disaggregated memory in the context of HPC. The ALCF has 6 TB of pooled random-access memory (RAM) across three appliances with a total of 14 FDR InfiniBand (IB) cards in them available to our Cooley visualization cluster. Disaggregated memory has a number of potentially significant impacts, which include the following:

- Reduction of aggregate system RAM and its associated purchase and operations costs.
- The ability to have virtually any amount of RAM on a node on demand.
- The decoupling (largely) of the purchase of CPUs, which we want to upgrade every few years to obtain improved processor performance, from RAM, whose performance has changed very little over the past decade.
- No need to have heterogeneous cluster and guess at the number of "big memory nodes" needed and have their RAM sit idle, consuming power, when not in use.

Work on RAN in 2019 had two major thrusts: performance improvements and evaluation of artificial intelligence (AI) workloads' performance when using RAN.

On the performance front, we had two significant outcomes. First, when evaluating a standard K-Means benchmark, it consistently showed better performance when using remote memory, which is highly counterintuitive. A doctoral research assistant discovered it was due to "false sharing" in the cache. Essentially, as written, the application was thrashing the cache, but the latency due to remote memory fetches significantly reduced this result and led to a net reduction in time to solution. The second outcome involved overall performance improvements in the Kove driver. In discussions with Kove about performance improvements, we suggested that due to the overhead associated with detailed tracking of least recently used (LRU) page state, a random algorithm might be more performant. In response, Kove implemented an algorithmic improvement that can double performance on certain workloads such as the STREAM benchmark. This improvement also has the added benefit of allowing the victim cache to be much smaller, meaning we realize better utilization of the local RAM. These improvements will be in the next production driver release.

AI workloads have consistently shown positive results when tested with RAN. A different Ph.D. student worked with ALCF staff on a more in-depth and rigorous investigation in this area. So far, she has evaluated the CANDLE benchmark suite and found results similar to previous work—that is, with 10 to 20 percent of RAM located locally and 80 to 90 percent of RAM located remotely, time to solution increases by only 10 to 15 percent. This approach has the potential to greatly improve the efficiency in allocating resources in clusters for AI workloads. A paper describing these results is in progress.

In addition, ALCF is working with Argonne's Mathematics and Computer Science (MCS) Division to use RAN as one of the underlying storage technologies they are evaluating with the Mochi project and with Julian Kunkel at the University of Reading on integrating RAN technology as a tier in a high-performance storage system. We have also updated our local ran_memctl utilities to support additional settings and per-connection settings.

Future work will be focused on broadening the AI workload analysis as well as trying to find an application group that is willing to collaborate with us on testing much larger, real-world workloads.

1.3.4 Outreach Efforts

1.3.4.1 General Outreach

User Advisory Council

The ALCF User Advisory Council was established in 2011 to help advise and assist ALCF leadership on matters related to resources, services, and policies that will strengthen opportunities for its users. The council members are all volunteers and hold expert knowledge of the tasks and requirements of specific applications or domain areas. In 2019, the council and ALCF's deputy division director met in February and in August, for discussions and updates on Theta's software policies and the 2018 annual survey results, among other topics.

Connection to Technology Commercialization and Partnerships Division

The ALCF interacts with Argonne's Technology Commercialization and Partnerships (TCP) Division and regularly supports TCP-organized meetings with potential industry partners. TCP in turn provides potential partners, even those without immediate HPC needs, with a complete picture of Argonne computing capabilities and partnership opportunities, including those at the ALCF.

1.3.4.2 Training Activities

The ALCF provides workshops and webinars to support the efforts of its users. The facility also collaborates with peer institutions to develop training opportunities, explore key technologies, and share best practices that improve the user experience. Below is a list of ALCF-sponsored and hosted training activities for 2019:

2019 ALCF Simulation, Data, and Learning Workshop

This annual workshop helps researchers improve the performance and productivity of simulation, data science, and machine learning (ML) applications on ALCF systems. Held at the ALCF from October 1–3, 2019, this workshop attracted 137 attendees. Participants worked directly with ALCF staff and industry experts during dedicated hands-on sessions to learn how to use available tools and frameworks to scale data-centric science on ALCF systems and improve productivity; how to set up workflows and use containers; and test and debug codes. Fifty percent of attendees (non-staff) have active allocations today in a mix of DD, ALCF Data Science Program (ADSP), Early Science Program (ESP), Exascale Computing Project (ECP), ALCC, and INCITE projects.

Aurora Programming Workshop

This invitation-only workshop focused on helping researchers, particularly ECP project teams, prepare applications and software technologies for ALCF's Aurora exascale system. The workshop was geared toward developers, and emphasized using the Intel software development kit to get applications running on testbed hardware. Teams were also given the opportunity to consult with ALCF staff and provide feedback. The workshop took place at Argonne from September 17–19, 2019, and attracted 163 participants.

ATPESC 2019

The Argonne Training Program on Extreme-Scale Computing (ATPESC) is an intensive, two-week summer program that teaches the skills, approaches, and tools needed to conduct computational science and engineering research on leading-edge supercomputers. The seventh annual ATPESC was held in St. Charles, near Argonne, from July 28–August 9, 2019, and attracted 73 graduate/postdoctoral students and early career researchers from 49 different institutions worldwide. The program is organized into eight core program tracks consisting of technical lectures, hands-on exercises, and dinner talks. To extend the reach of the program, 67 hours of video playlists of ATPESC presentations were uploaded to Argonne's YouTube training channel and promotion is ongoing. ATPESC organizers work diligently to continuously improve the program. In 2019, this effort included improvements to outreach efforts, logistics, and event management; new audio files of the lectures; and a new track for machine learning and deep learning topics. Participant and lecturer feedback on program content, structure, logistics, and pedagogy was very positive and emphasized the expertise of the lecturers, the quality of the program, and the benefits to the participants (and the institutions they represent) and the scientific community in general.

2019 ALCF Computational Performance Workshop

This annual workshop helps researchers achieve computational readiness on ALCF computing resources to prepare for a major allocation award (e.g., INCITE, ALCC, ADSP). The 2019 workshop was held at the ALCF from April 30–May 2, 2019, and attracted 114 attendees. Participants worked directly with ALCF staff and industry professionals during dedicated hands-on sessions to explore advanced techniques and tools to enhance code performance, expand their data science skills, and benchmark and debug codes. Roughly 56 percent of non-staff workshop attendees have active allocations today in a mix of DD, ADSP, ESP, ALCC, and INCITE projects.

Center of Excellence (COE) Aurora Early Science Program (ESP) Workshop

This workshop focused on applications and software development for Argonne's Aurora exascale system and was open to all ESP and select ECP project members. Held at Argonne from April 22–24, 2019, this workshop attracted 133 attendees, including 71 ESP project team members (8 of them project PIs) and 28 Intel staff members, as well as ALCF staff. All 20 Aurora ESP projects were represented. The workshop was the first introduction to the Intel CPU-GPU hardware and software for the ESP project teams. Talks centered around Intel GPU architecture, including present-day Gen9 integrated GPUs being used as a development platform for Aurora. Intel's software development tools and libraries for Aurora were introduced, and access to pre-beta software was provided on Argonne's Joint Laboratory for System Evaluation (JLSE) Gen9 systems (software later to be branded as Intel's oneAPI standard implementation).

A series of application experience talks illustrated DPC++, OpenMP, and OpenCL programming for Intel GPUs anticipating the use of Aurora's X^e hardware.

ALCF Webinars

The ALCF organizes and hosts monthly, live-presentation webinars for users on a wide range of topics. Each webinar consists of a 45-minute interactive talk followed by 15 minutes of Q&A. All talks are recorded and uploaded to ALCF's YouTube channel. The ALCF events website hosts all of the public training materials.

The ALCF hosted the following webinars in 2019:

- Boosting Power Efficiency of HPC Applications with GEOPM
- Run-to-Run Variability on Theta and Best Practices for Performance Benchmarking
- Preparing an Application for Hybrid Supercomputing Using Cray's Tool Suite
- Workflow Management with Balsam
- SENSEI Cross-Platform View of In Situ Analytics
- Singularity on Theta: How to Build and Scale Containers at the ALCF
- Quick Start: Using Apache Spark for Large-Scale Data Processing
- Cross-Platform Performance Engineering with Arm Allinea Studio
- Introduction to Globus: Research Data Management Software at the ALCF
- Building Deep Learning Applications for Big Data using BigDL and Analytics Zoo C++
- Best Practices 101: A miniQMC Case Study
- DeepHyper: Hyperparameter Search Package for Deep Neural Networks
- SYCL, Programming Language for Aurora ESP Webinar Series
- Effective Use of Python on Theta (Early Science Program)
- Performance Tuning Using Intel Advisor and VTune Amplifier (Early Science Program)
- Machine Learning with TensorFlow, Horovod, and PyTorch on HPC (Early Science Program)

1.3.4.3 Community Outreach

Facility Tours

The ALCF welcomes visitors to tour the facility year-round. A staff administrator fields all tour requests, including internal requests from various Argonne offices and outreach programs, and coordinates with ALCF staff accordingly. Tour stops include the machine room, data center, and visualization lab, and tour guides provide historical narrative along with information about current research activities. More than a thousand people took the tour in 2019, including middle school and high school students, government officials and Congressional staff members, government delegations, industry representatives, summer research program participants, visiting researchers, and journalists. Following is a selection of the groups and guests that visited in 2019.

Federal government officials and staff included U.S. Representative Sean Casten (IL 6th District); DOE ASCR program manager Bill Spotz; Assistant to the President Joe Grogan; Jeremy Ison, deputy chief of staff to Secretary Dabbar; and delegations from the office of U.S. Senator Tammy Duckworth and from the DOE Office of Enterprise Assessment. International delegations included government officials from Brazil, the Commonwealth Scientific & Industrial Research Organization (CSIRO) (Australia's national science agency), the Belgian Consulate, and Australia's Consul-General in Chicago.

Numerous school groups came through and included honors students from Purdue University's College of Engineering, business students from Northwestern University's Kellogg School and students from Northwestern's Medill School of Journalism, Davis-Bahcall Scholars from South Dakota universities, Notre Dame University undergraduates, attendees of Argonne's summer Creative Approaches to Problems in Science (CAPS) High School Computing Workshop, and After School Matters high school students.

Industry groups included representatives from General Motors, Exelon/ComEd, and the Korea Power Exchange.

Science, Technology, Engineering and Math (STEM) Activities

Introduce a Girl to Engineering Day (IGED)

IGED is an annual lab-sponsored event in February that pairs approximately 100 local eighth graders with Argonne engineers and scientists for a day of presentations and hands-on activities focused on STEM careers.

Science Careers in Search of Women (SCSW)

In April, Argonne hosted its 32nd Science Careers in Search of Women Conference, a one-day event that offers female high school students an opportunity to explore careers in STEM research. The students attended presentations and career panel discussions led by Argonne women scientists and engineers, facility tours, a poster session, and a career fair. ALCF staff members served in various roles at SCSW 2019.

CodeGirls

CodeGirls is a three-day STEM camp for 6th- and 7th-grade girls held in June that teaches the fundamentals of coding in the Python programming language. Campers also learn from women role models, including several ALCF staff members. The group also toured the ALCF machine room and visualization lab.

Creative Approaches to Problems in Science (CAPS) High School Computing Workshop

CAPS is a weeklong coding camp in July for local high school students taught by ALCF staff members. The camp curriculum promotes problem-solving and teamwork skills through hands-on coding activities, such as coding with Python and programming a robot. The camp is a joint initiative of the ALCF and Argonne's Educational Programs Office.

Big Data and Visualization Camp

This summer camp held in July aims to bridge the gap between high school and higher education and offers participants an intensive and immersive challenge in coding and data science.

Approximately a dozen current and recent high school students attended the camp, which is organized by Argonne's Educational Programs Office and taught by Argonne computer scientists. The students learn to program visualization tools in Python and work with data from the Array of Things (AoT) project, an Argonne–University of Chicago urban sensor project, to experiment firsthand with investigative research.

Hour of Code

As part of the national Computer Science Education Week (CSEdWeek), December 3–9, 2019, more than 50 Argonne computer scientists, including many ALCF staff members, visited Chicago and area suburban schools to assist teachers in celebrating the Worldwide Hour of Code initiative. Scientists gave short presentations on their career paths and work at Argonne, and interacted with students one-on-one while they participated in various Hour of Code activities. CSEdWeek was established by Congress in 2009 to raise awareness about the need to elevate computer science education at all levels.

1.3.5 Communications

New ALCF Website

A newly designed ALCF website was launched on December 9, 2019. The new web presence is Americans with Disabilities Act (ADA)-compliant, adapts to all device displays, and features a broad range of opportunities for showcasing multimedia content. ALCF's communications team worked with Sandbox Studios, an external design firm, to bring the site to fruition. The visual design, navigation, and content was informed by years of user feedback and reflects the priorities set by ALCF leadership.

Website Support Center Improvements

The ALCF formed an internal document committee to help design the Support Center portal on the new ALCF website. The Support Center contains information about current machine status, facility updates, training events, and various guides. Improvements to the portal include the ability to search and filter content in the new library of training videos and slides. These media are now classified by topic when entered into the database so they can be suggested to users. Another improvement allows code examples to be copied to clipboard and linked to GitHub source code.

Website Newsroom Improvements

The ALCF News Center on the new ALCF website now offers the ability to create media kits for marketing campaigns and features ALCF-related Twitter posts, or tweets, on the website landing page. "Posts" can now be added to the News Center to feature scientific visualizations on the homepage, as well as to highlight VIP tours, conference participation, community partnerships, outreach activities, and users' recent journal and conference publications.

Communications through Mailing Lists and Social Media

The facility maintains several communication channels, including direct e-mail campaigns, scriptable e-mail messages, social media postings, and website postings (Table 1.6; target audiences are identified in Table 1.7). Users can opt out of the system notify and newsletter mailing lists. In 2019, the ALCF communications team added two new social media channels, Twitter and LinkedIn, to help extend the reach of their campaigns.

Table 1.6 2019 Primary Communication Channels

Channel Name	Description	When Used/Updated
Newsbytes	HTML-formatted newsletter featuring science, facility news, recent news hits, and upcoming training events.	Monthly
Special Announcements	E-mail newsletter and text-format with information on conferences, training events, etc.—both ALCF and non-ALCF opportunities.	Ad hoc
Weekly Digest	Plain-text weekly rollup of events affecting ALCF systems and software, upcoming deadlines, and training opportunities.	Weekly
Social Media	Social media used to promote 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 in a text-based format per user and channel preference.	As needed

Table 1.7 2019 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 the 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, PIs/proxies, individual users

The e-mail newsletter Newsbytes highlights ALCF-supported research or advancements, advertises upcoming training events and allocation program announcements, and provides links to relevant news stories. Special e-mails advertising training and fellowship opportunities are sent throughout the year, as needed.

Promotional Activities and Media Hits

In 2019, the ALCF produced and placed 47 original science stories in various news outlets in coordination with Argonne's Communications, Education & Public Affairs (CEPA) Division and other ALCF direct relationships. The media team uses social monitoring software Meltwater to track media hits. In 2019, Meltwater reported 238 unique ALCF media hits, 127 of which were chronicled on the ALCF website, and an audience reach of 325.9 million.

Other Publications

The ALCF produces several high-quality publications throughout of the year that describe aspects of the facility's mission and summarize its research achievements (Table 1.8). Most of these documents are available for download on the ALCF website.

Table 1.8 Publications Designed for Print

Publication	Frequency	When
Press and Visitor Packets	As Needed	As Needed
Industry Brochure	As Needed	As Needed
Annual Report	Yearly	March
Science Report	Yearly	September
Fact Sheets	Yearly	November
INCITE Poster	Yearly	December

Conclusion

Our users are at the forefront of all of ALCF's interactions. As a user facility, the ALCF strives to improve our user experience processes and to help our customers make the most of their allocation time on our resources. In 2019, improvements included working with ALCC and INCITE project teams to transform and optimize their scientific codes and helping DD projects reach their scientific goals and obtain INCITE awards. The ALCF worked in partnership with other national laboratories to present sessions on our work at the annual meetings of major national scientific societies. The ALCF implemented new tools to help with data analysis and conducted a wide range of outreach events to teach best practices and help our users explore new technologies.

Section 2. Operational Performance

Did the facility's operational performance meet established targets?

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 Operational Performance Section

	Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				Theta (Cray XC40): 4,008-node, 251K-core,[a] 64 TB MCDRAM,[b] 770 TB DDR4			
	CY 2018		CY 2019		CY 2018		CY 2019	
	Target	Actual	Target	Actual	Target	Actual	Target	Actual
Scheduled Availability	90.0%	99.3%	90.0%	99.7%	90.0%	98.0%	90.0%	99.8%
Overall Availability	90.0%	96.5%	90.0%	95.6%	90.0%	94.3%	90.0%	95.0%
System MTTI	N/A	12.18 days	N/A	15.21 days	N/A	10.12 days	N/A	12.39 days
System MTTF	N/A	72.73 days	N/A	182.38 days	N/A	32.55 days	N/A	72.86 days
INCITE Usage	3.5B	4.1B [g]	3.5B	4.2B [i]	390.9M	529.4M	1.1B	1.2B
Total Usage	N/A	6.3B [h]	N/A	6.2B [j]	N/A	2.0B	N/A	2.1B
System Utilization	N/A	93.6%	N/A	93.4%	N/A	92.1%	N/A	97.0%
INCITE Overall Capability [c,e]	40.0%	70.8% [g]	40.0%	68.8% [i]	20%	64.4%	20.0%	70.6%
INCITE High Capability [d,f]	10.0%	22.0% [g]	10.0%	27.9% [i]	N/A	13.3%	N/A	10.5%

[a] For more information on performance metrics calculations, see Section A.7.

[b] MCDRAM = multi-channel DRAM; DDR4 = Double Data Rate 4 Synchronous Dynamic Random-Access Memory.

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

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

[e] Theta Overall Capability = Jobs using ≥ 20.0 percent (800 nodes, 51,200 cores) of Theta.

[f] Theta High Capability = Jobs using ≥ 60.0 percent (2,400 nodes, 153,600 cores) of Theta.

[g] Usage includes 0.3M core-hours from 2018 Cetus production jobs.

[h] Usage includes 11.1M core-hours from 2018 Cetus production jobs.

[i] Usage includes 1.3M core-hours from 2019 Cetus production jobs.

[j] Usage includes 29.1M core-hours from 2019 Cetus production jobs.

ALCF Resources

During CY 2019, the ALCF operated two production resources, Mira and Theta. Mira supported ALCC campaigns until June 30, and INCITE campaigns until December 31. Theta supported both ALCC and INCITE campaigns all year. Mira is a 48K-node, 768K-core, 10-petaflops (PF) Blue Gene/Q with 768 TB of RAM. Mira mounts three general parallel file systems (GPFSs) with approximately 26.5 petabytes (PB) of usable space. Theta is a 4,392-node, ~281K-core, 11.69-PF Cray XC40 with 891 TB of RAM. (For more information on performance metrics calculations, see Section A.7.) Theta mounts one GPFS file system and one Lustre file system with approximately 8.6 PB and 9.2 PB of usable space, respectively. The filesystems on both Mira and Theta are also mounted by the visualization and analysis cluster Cooley for analysis. Both Mira and Theta have access to the facility-wide HPSS (high-performance storage system) tape archive. ALCF operated two other Blue Gene/Q systems, Cetus and Vesta, and one other Cray XC40, Iota.

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. Iota is a 176-node, 11K-core Cray XC40 with 36.6 TB of RAM. Like Vesta, Iota is an independent test and development resource and shares no resources with Theta.

In 2014, ALCF began select use of Cetus for INCITE projects with simulation runs that required nontraditional HPC workflows and has continued implementing that usage approach ever since. This deployment of Cetus 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 systems and the quality of the maintenance procedures.

Theta

2.1.1 Scheduled and 2.1.2 Overall Availability

Theta entered full production on July 1, 2017. In consultation with the DOE Program Manager, ALCF has agreed to a target of 90 percent overall availability and a 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 for Theta

Theta (Cray XC40): 4,008-node, 251K-core, 64 TB MCDRAM, 770 TB DDR4				
	CY 2018		CY 2019	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	98.0	90.0	99.8
Overall Availability	90.0	94.3	90.0	95.0

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, 2019, through December 31, 2019, as annotated in Figure 2.1.

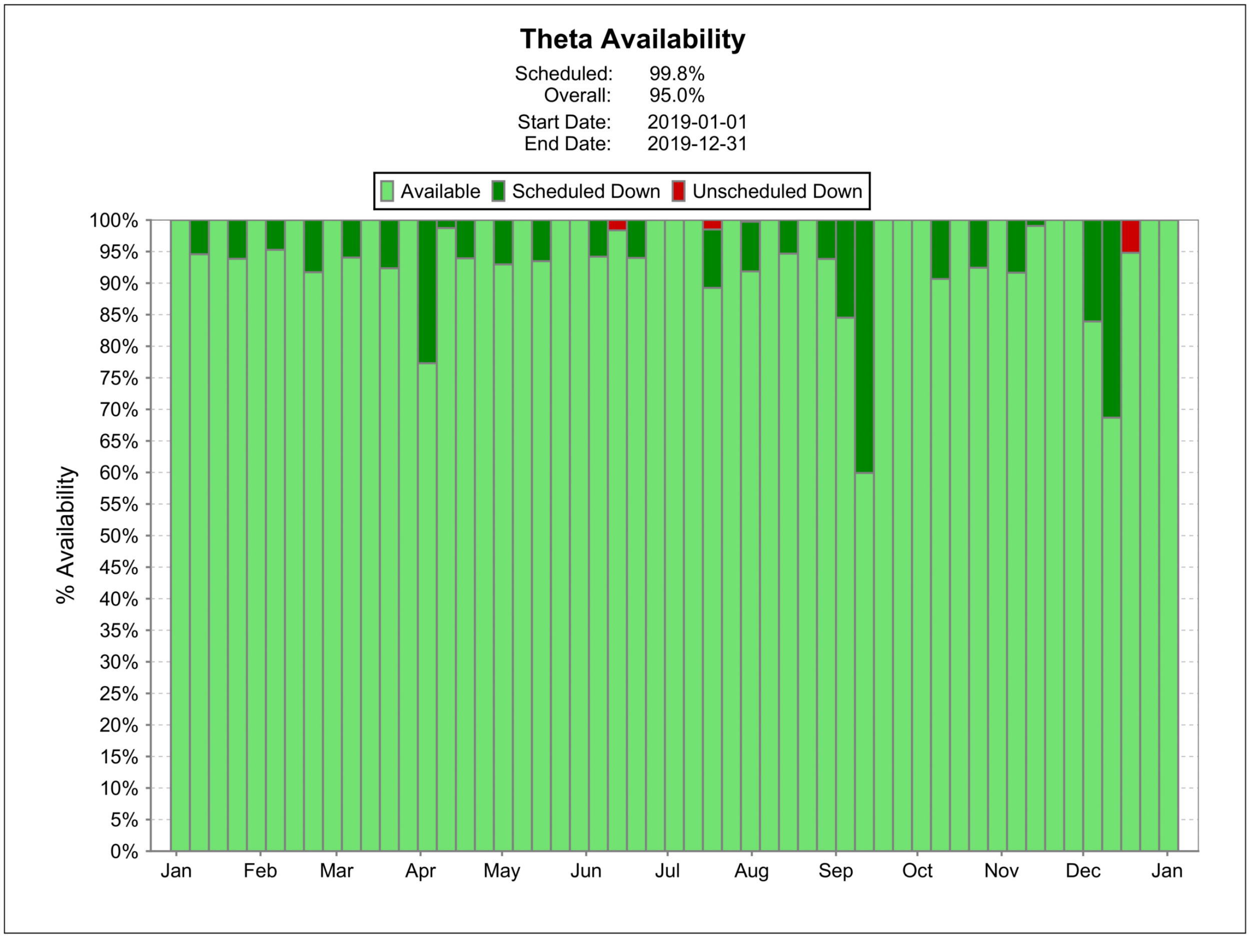

Figure 2.1 Theta Weekly Availability for CY 2019

Graph Description: Each bar in Figure 2.1 represents the percentage of the machine available for seven days. 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. Each of the significant loss events is described in detail below.

June 14, 2019: UPS unit failure

One of the building's two uninterruptable power supply (UPS) units failed briefly due to an over-temperature condition. The over-temperature condition occurred because of maintenance being performed on a rooftop cooling unit. Typically, when a UPS unit fails, it should fail to bypass so that power is not interrupted; however, that did not occur in this instance. This caused all load to shift briefly to the other UPS unit, which led to an overdraw on that unit. While Theta and its Lustre filesystem were not affected, upstream services—such as the Lightweight Directory Access Protocol (LDAP) and the GPFS filesystem that Theta relies on—were affected. This affected infrastructure briefly lost power since it is fed only from the UPS. The UPS was manually bypassed while it was investigated and brought back to service.

July 16, 2019: License server configuration extended planned maintenance
Following a scheduled preventative maintenance (PM), there was a routing issue between Theta and the license servers that prevented PM pre-release activities from succeeding. Once the routing issue was corrected, pre-release activities were able to complete and Theta was returned back to service.

September 6–10, 2019: Building pump, cooling tower, and UPS work
The TCS Building Owner (TCSB) installed a third water pump, recommissioned the air handler units' (AHUs) cooling tower, and installed two temporary 850-kVA UPS units. The third water pump ensures N+1 redundancy in all failure scenarios given the current and estimated future loads of the chilled water loop currently serving Mira and Theta. Bringing the cooling tower back online ensures that the AHUs have a separate water source from the chilled water plant. The temporary UPS units provided temporary N+1 UPS capacity until the three new units were installed in December (see below).

December 6–9: New UPS installation
The TCSB replaced the three aging 600 kVA UPS units with three 750 kVA units. This installation not only restored permanent N+1 redundancy, since one of the 600-kVA units had previously failed, but also provides 300 kVA of additional UPS capacity. Most of the UPS installation work was completed ahead of this outage window because of the previously installed temporary UPS units, with this outage used to make final, permanent connections to the new UPS units and remove the temporary UPS units and cabling.

December 15–16: Stale job
A job's Application Level Placement Scheduler (ALPS) reservation failed to release and stayed in a "pending cancel" state, which held the nodes. The next job in the queue was a large job that required some of the held nodes. Cobalt recognized the hung reservation and tried to release it unsuccessfully. Because Cobalt viewed those nodes as being in a "cleanup" state, it assumed a 5-minute backfill window, and this deadlock caused no new jobs to start. Because this event happened late on a Sunday evening, the situation was not discovered until later that night at which time the hung ALPS reservation was forcefully released and jobs began to run again.

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 targets. Table 2.3 summarizes the current MTTI and MTTF values.

Table 2.3 MTTI and MTTF Results

Theta (Cray XC40): 4,008-node, 251K-core, 64 TB MCDRAM, 770 TB DDR4				
	CY 2018		CY 2019	
	Target	Actual	Target	Actual
System MTTI	N/A	10.12 days	N/A	12.39 days
System MTTF	N/A	32.55 days	N/A	72.86 days

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

Theta (Cray XC40): 4,008-node, 251K-core, 64 TB MCDRAM, 770 TB DDR4				
	CY 2018		CY 2019	
	Target	Actual	Target	Actual
System Utilization	N/A	92.1%	N/A	97.0%

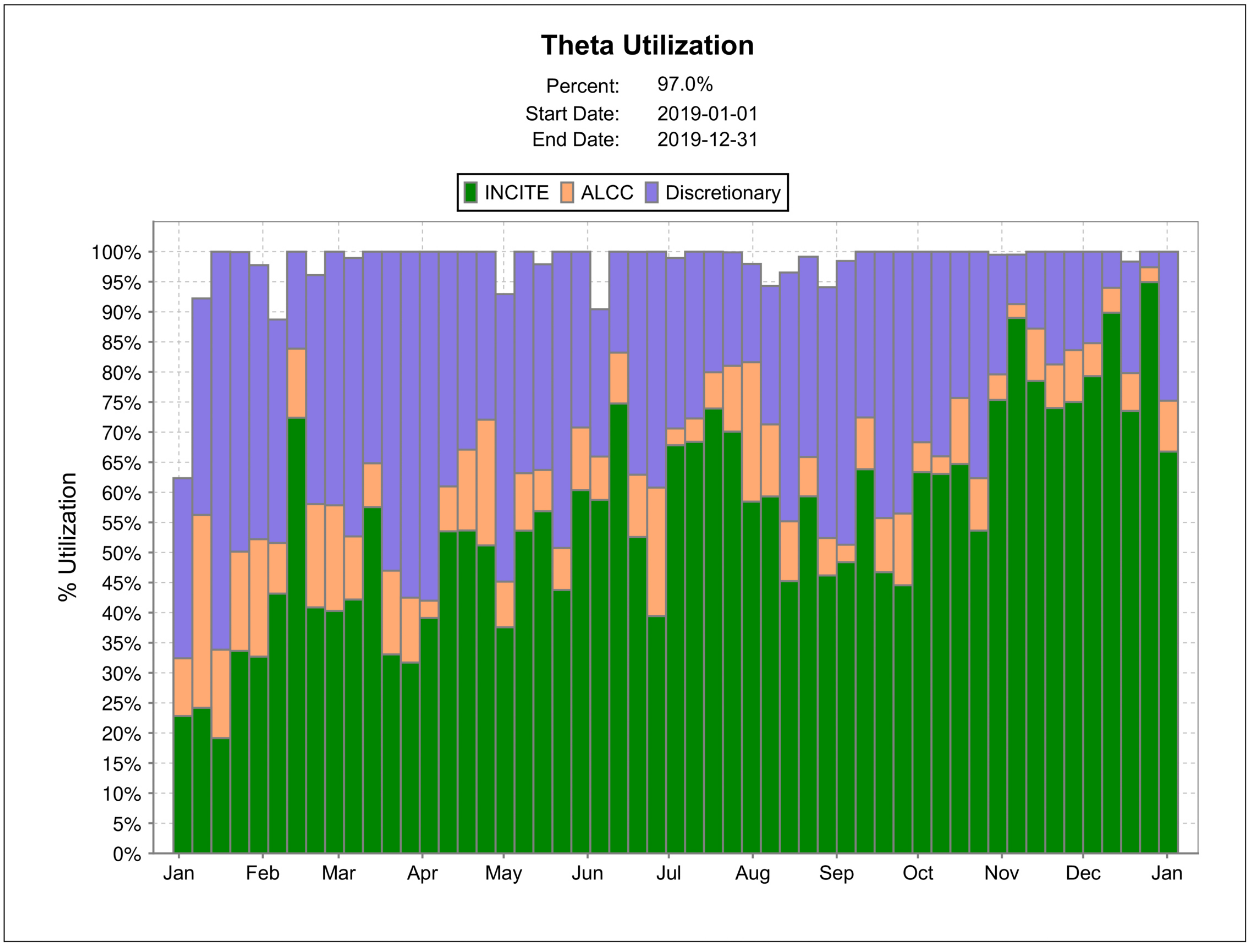

Figure 2.2 System Utilization over Time by Program

The system utilization for Theta was 97.0 percent for its 2019 production period of January 1, 2019, through December 31, 2019.

Table 2.5 shows how Theta'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 were allocated to the INCITE program, up to 20 percent to ALCC program allocations, and 20 percent to DD allocations. 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 3.1.2. In CY 2019, Theta delivered a total of 2.1 billion core-hours.

Table 2.5 Core-Hours Allocated and Used by Program

Theta (Cray XC40): 4,008-node, 251K-core, 64 TB MCDRAM, 770 TB DDR4						
	CY 2018			CY 2019		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	% of used core-hours	Core-hours	Core-hours	% of used core-hours
INCITE	390.9M	529.4M	27.1	1.1B	1.2B	56.5
ALCC	372.0M	508.2M	26.0	359.5M	210.6M	10.0
DD	1.5B	916.6M	46.9	1.1B	701.8M	33.5
Total	2.3B	2.0B	100.0	2.6B	2.1B	100.0

Summary: For CY 2019, 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 Theta, capability is defined as using greater than 20 percent of the machine, or 800 nodes, and high capability is defined as using greater than 60 percent of the machine, or 2,400 nodes. See Table A.2 in 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. Figure 2.4 shows the three programs' utilization of total core hours (from Table 2.6) over time, and Figure 2.5 shows the overall distribution of job sizes over time.

Table 2.6 Capability Results

Theta (Cray XC40): 4,008-node, 251K-core, 64 TB MCDRAM, 770 TB DDR4						
	CY 2018			CY 2019		
Capability Usage	Total Hours	Capability Hours	Percent Capability	Total Hours	Capability Hours	Percent Capability
INCITE Overall	529.4M	340.8M	64.4	1.2B	835.6M	70.6
INCITE High	529.4M	70.2M	13.3	1.2B	124.7M	10.5
ALCC Overall	508.2M	139.8M	27.5	210.6M	55.5M	26.3
ALCC High	508.2M	18.3M	3.6	210.6M	0.0M	0.0
Director's Discretionary Overall	916.6M	337.4M	36.8	701.8M	279.2M	39.8
Director's Discretionary High	916.6M	98.9M	10.8	701.8M	48.6M	6.9
TOTAL Overall	2.0B	818.0M	41.9	2.1B	1.2B	55.8
TOTAL High	2.0B	187.4M	9.6	2.1B	173.3M	8.3

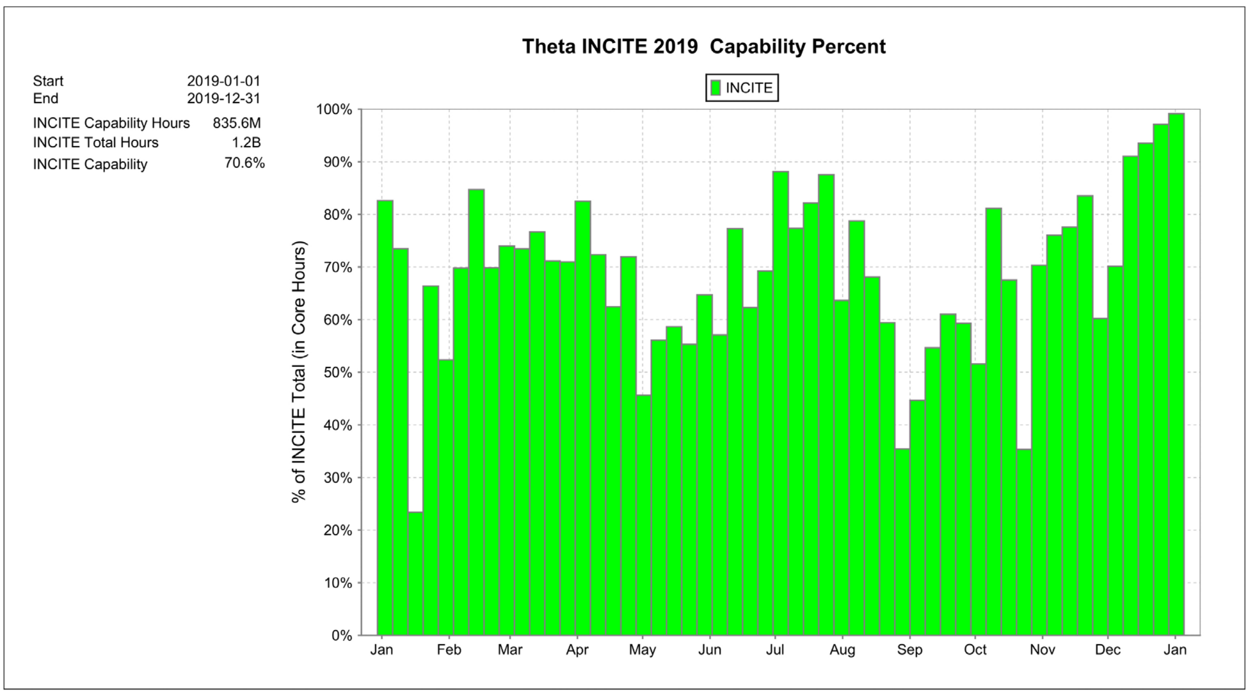

Figure 2.3 Theta INCITE Overall Capability

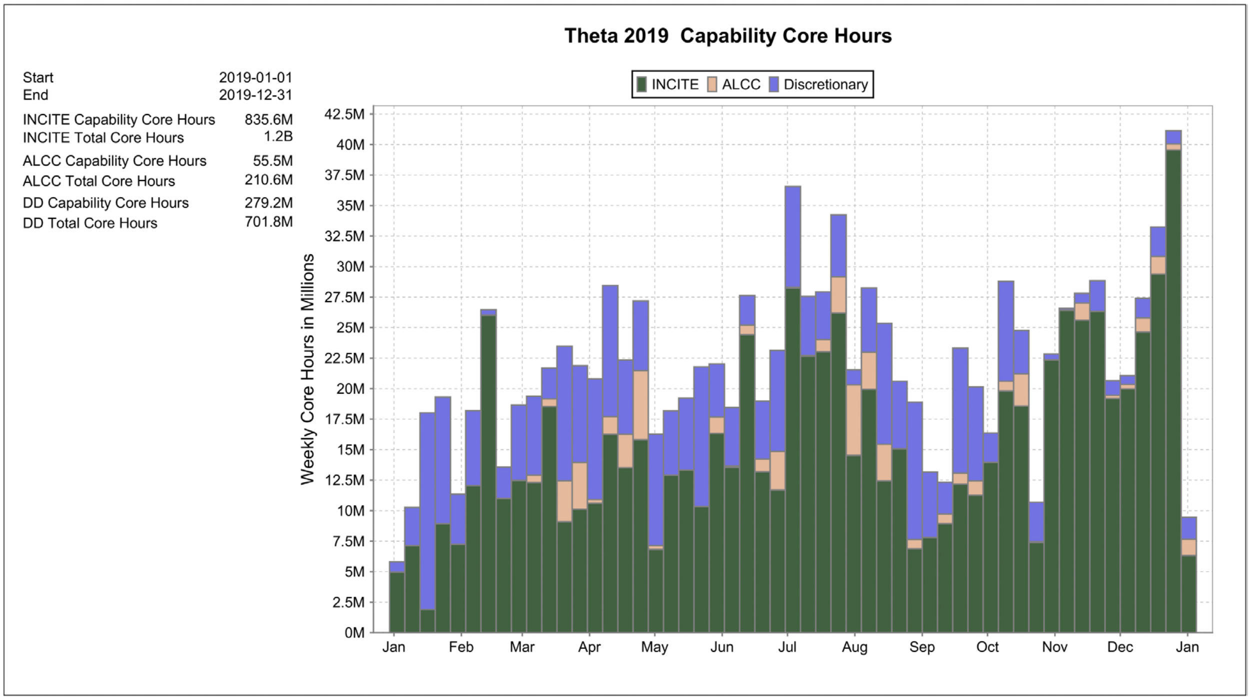

Figure 2.4 Theta Capability Core-Hours by Program

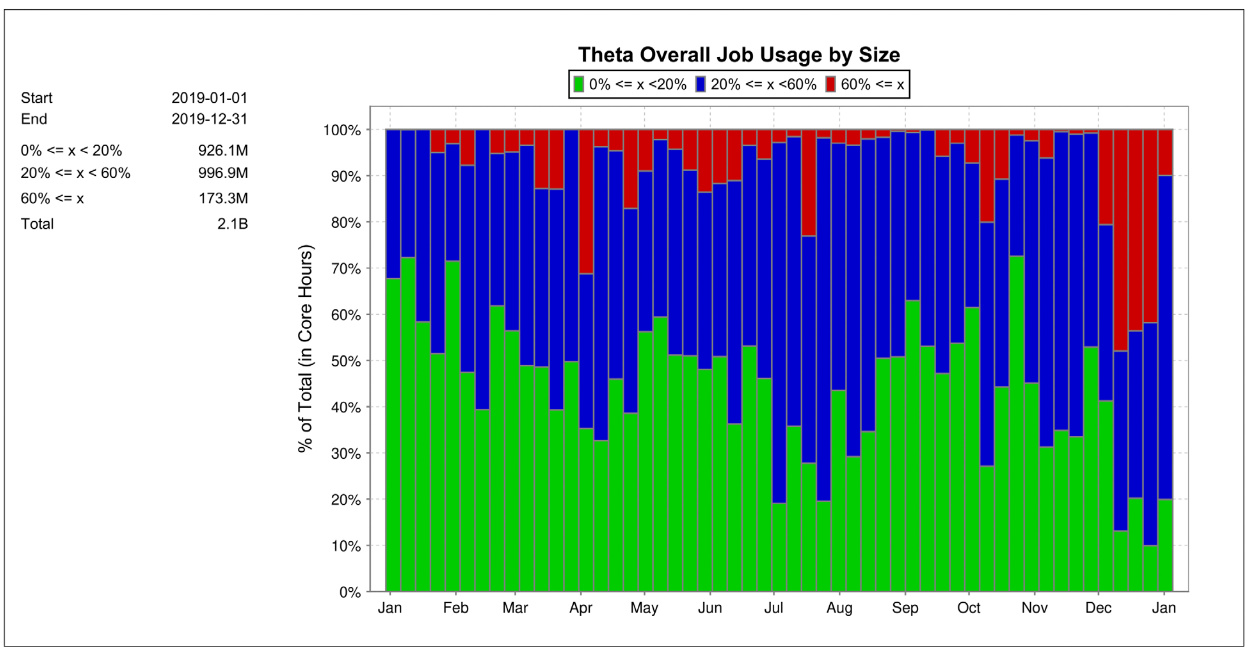

Figure 2.5 Theta Job Usage by Size

Mira

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, the ALCF has agreed to metrics of 90 percent overall availability and a target of 90 percent scheduled availability (in response to ASCR's request that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.7 summarizes the availability results.

Table 2.7 Availability Results

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2018		CY 2019	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	99.3	90.0	99.7
Overall Availability	90.0	96.5	90.0	95.6

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 of Mira for the period January 1, 2019, through December 31, 2019, as annotated in Figure 2.6.

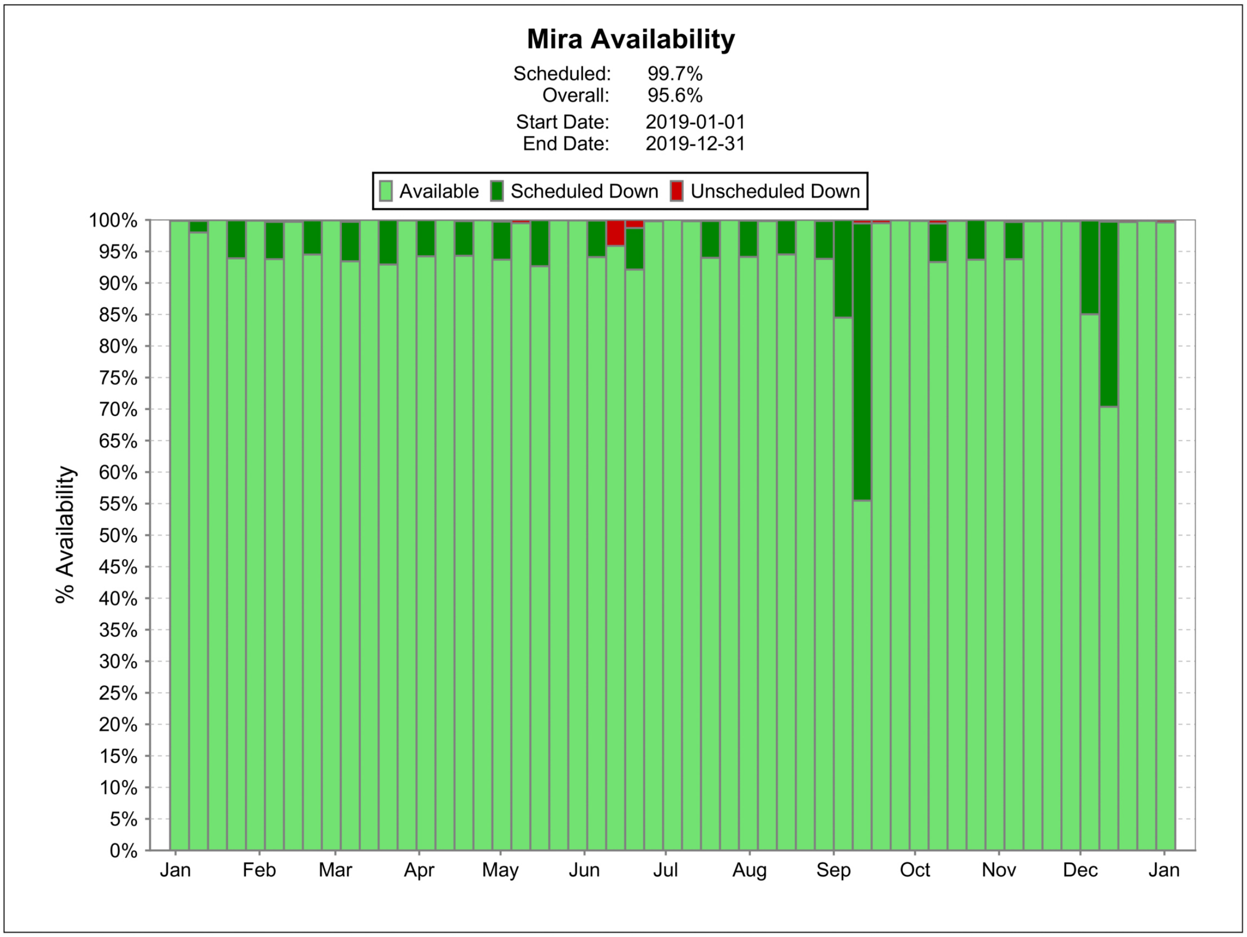

Figure 2.6 Mira Weekly Availability for CY 2019

Graph Description: Each bar in Figure 2.6 represents the percentage of the machine available for seven days. 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. Each of the significant loss events is described in detail below.

June 14, 2019: UPS unit failure

One of the two building UPS units failed briefly due to an over-temperature condition. The over-temperature condition occurred because of maintenance being performed on a rooftop cooling unit. Typically, when a UPS unit fails it should fail to bypass so that power is not interrupted; however, that did not occur in this instance. This caused all load to shift briefly to the other UPS unit, which led to an overdraw on that unit. All of the mira-fs and mira-home filesystems, as well as the Mira management infrastructure, briefly lost power since they are fed only from the UPS. The UPS was manually bypassed while it was investigated and brought back to service.

September 6–10, 2019: Building pump, cooling tower, and UPS work
The TCSB installed a third water pump, recommissioned the AHUs' cooling tower, and installed two temporary 850-kVA UPS units. The third water pump ensures N+1 redundancy in all failure scenarios given the current and estimated future loads of the chilled water loop currently serving Mira and Theta. Bringing the cooling tower back online ensures that the AHUs have a separate water source from the chilled water plant. The temporary UPS units provided temporary N+1 UPS capacity until the three new units were installed in December (see below).

December 6–9: New UPS installation
The TCSB replaced the three aging 600-kVA UPS units with three 750-kVA units. This not only restored permanent N+1 redundancy—since one of the 600-kVA units had previously failed—but also provides 300 kVA of additional UPS capacity. Most of the UPS installation work was completed ahead of this outage window because of the previously installed temporary UPS units, with this outage used to make final, permanent connections to the new UPS units and remove the temporary UPS units and cabling.

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 targets. Table 2.8 summarizes the current MTTI and MTTF values.

Table 2.8 MTTI and MTTF Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2018		CY 2019	
	Target	Actual	Target	Actual
System MTTI	N/A	12.18 days	N/A	15.21 days
System MTTF	N/A	72.73 days	N/A	182.38 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 concurrently with Theta, 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 tends to cap MTTI at 14 days, but does not directly affect MTTF. However, due to a few skipped scheduled maintenance days in July, September, and November of 2019 for various reasons, the system MTTI was, on average, better than the 14-day cap.

2.2 Resource Utilization

The following sections discuss system allocation and usage, total system utilization percentage, and capability usage. *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 2.9 summarizes ALCF utilization results, and Figure 2.7 shows system utilization over time by program.

Table 2.9 System Utilization Results

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2018		CY 2019	
	Target	Actual	Target	Actual
System Utilization	N/A	93.6%	N/A	93.4%

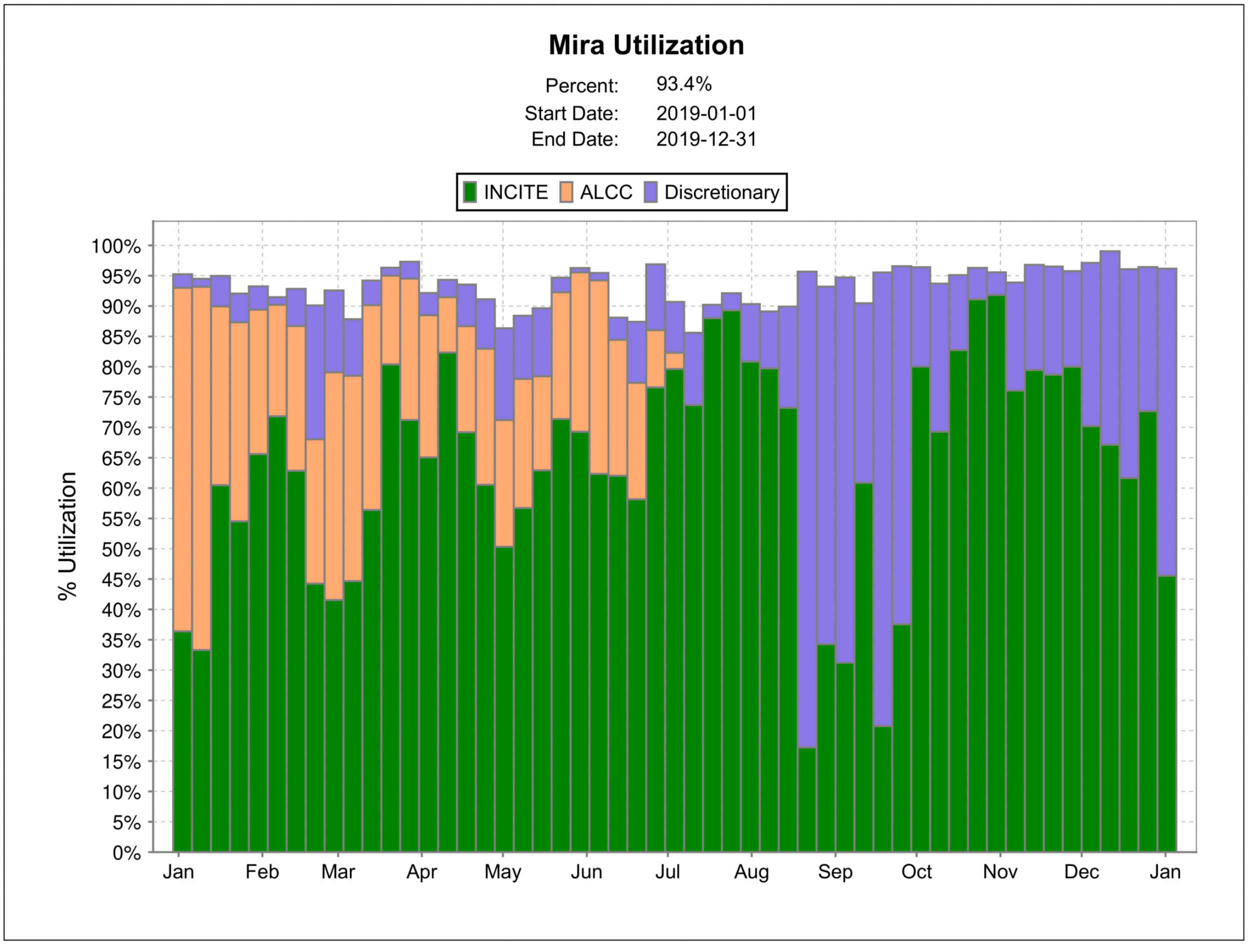

Figure 2.7 System Utilization over Time by Program

The system utilization for Mira was 93.4% for its 2019 production period of January 1, 2019, through December 31, 2019.

Table 2.10 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. In the first half of CY 2019, 60 percent of available hours went to the INCITE program, 20 percent went to the ALCC program, and 20 percent to the DD allocations. In the second half of the year, 60 percent went to INCITE, 20 percent to DD, and the previous ALCC 20 percent went to two projects chosen in collaboration with the Office of Science. The 2019 ALCC program was Mira's final ALCC allocation year due to the machine's retirement in December 2019. 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 allocations are often not used in full. The values for core-hours used for the DD allocations are higher because several projects were given allocations to perform final runs on Mira. DD allocations are discussed in detail in Section 3.1.2. In CY 2019, Mira delivered a total of 6.2 billion core-hours.

Table 2.10 Core-Hours Allocated and Used by Program

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2018			CY 2019		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	% of used core-hours	Core-hours	Core-hours	% of used core-hours
INCITE	3.4B	4.1B [a]	65.8	3.5B	4.2B [e]	68.4
ALCC	1.5B	1.7B [b]	26.6	599.0M	845.7M [f]	13.7
DD	1.4B	472.2M [c]	7.6	1.8B	1.1B [g]	17.9
Total	6.2B	6.3B [d]	100.0	6.3B	6.2B [h]	100.0

[a] Usage includes 0.3M core hours from 2018 Cetus production jobs.

[b] Usage includes 0.5M core hours from 2018 Cetus production jobs.

[c] Usage includes 10.4M core hours from 2018 Cetus production jobs.

[d] Usage includes 11.1M core hours from 2018 Cetus production jobs.

[e] Usage includes 1.3M core hours from 2019 Cetus production jobs.

[f] Usage includes 0.2M core hours from 2019 Cetus production jobs.

[g] Usage includes 27.6M core hours from 2019 Cetus production jobs.

[h] Usage includes 29.1M core hours from 2019 Cetus production jobs.

Summary: For CY 2019, 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 2.11 and Figure 2.8 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. Figure 2.9 shows the three programs' utilization of total core hours (from Table 2.11) over time, and Figure 2.10 shows the overall distribution of job sizes over time.

Table 2.11 Capability Results

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2018			CY 2019		
Capability Usage	Total Hours	Capability Hours	Percent Capability	Total Hours	Capability Hours	Percent Capability
INCITE Overall	4.1B [a]	2.9B	70.8	4.2B [e]	2.9B	68.8
INCITE High	4.1B [a]	904.4M	22.0	4.2B [e]	1.2B	27.9
ALCC Overall	1.7B [b]	888.6M	53.3	845.7M [f]	460.9M	54.5
ALCC High	1.7B [b]	261.4M	15.7	845.7M [f]	65.2M	7.7
Director's Discretionary Overall	472.2M [c]	254.7M	53.9	1.1B [g]	732.0M	66.3
Director's Discretionary High	472.2M [c]	66.6M	14.1	1.1B [g]	442.9M	40.1
TOTAL Overall	6.3B [d]	4.1B	64.9	6.2B [h]	4.1B	66.4
TOTAL High	6.3B [d]	1.2B	19.7	6.2B [h]	1.7B	27.3

[a] Usage includes 0.3M core hours from 2018 Cetus production jobs.

[b] Usage includes 0.5M core hours from 2018 Cetus production jobs.

[c] Usage includes 10.4M core hours from 2018 Cetus production jobs.

[d] Usage includes 11.1M core hours from 2018 Cetus production jobs.

[e] Usage includes 1.3M core hours from 2019 Cetus production jobs.

[f] Usage includes 0.2M core hours from 2019 Cetus production jobs.

[g] Usage includes 27.6M core hours from 2019 Cetus production jobs.

[h] Usage includes 29.1M core hours from 2019 Cetus production jobs.

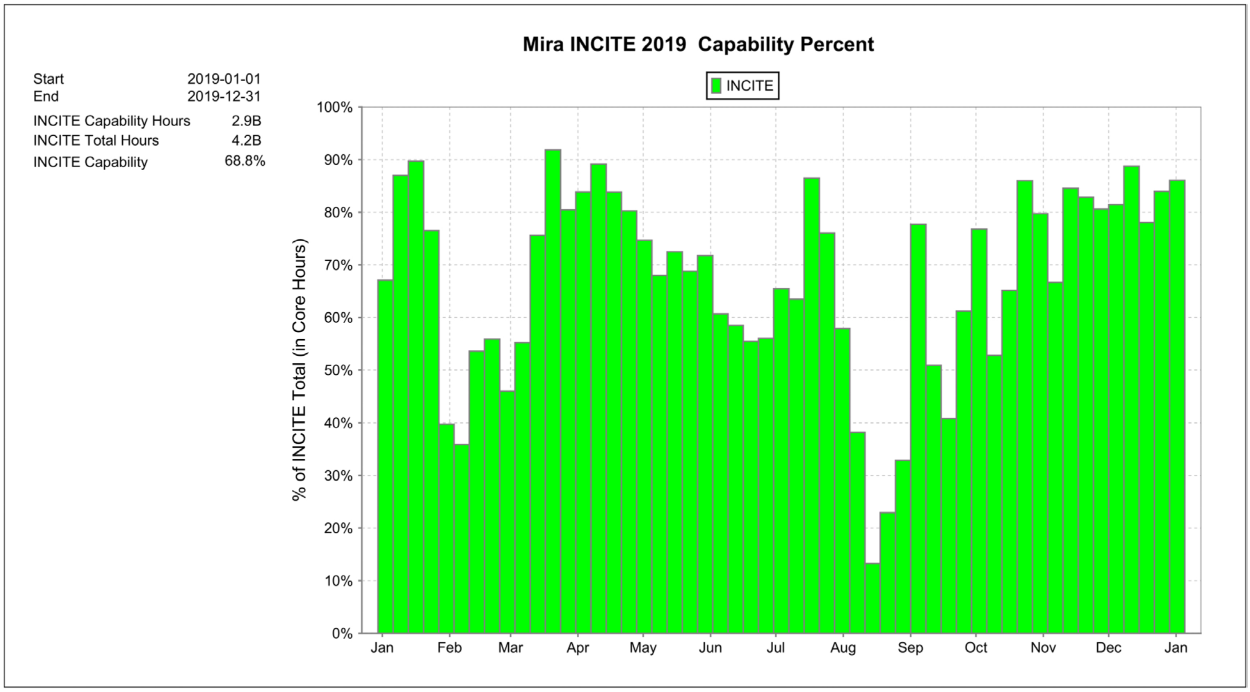

Figure 2.8 Mira INCITE Overall Capability

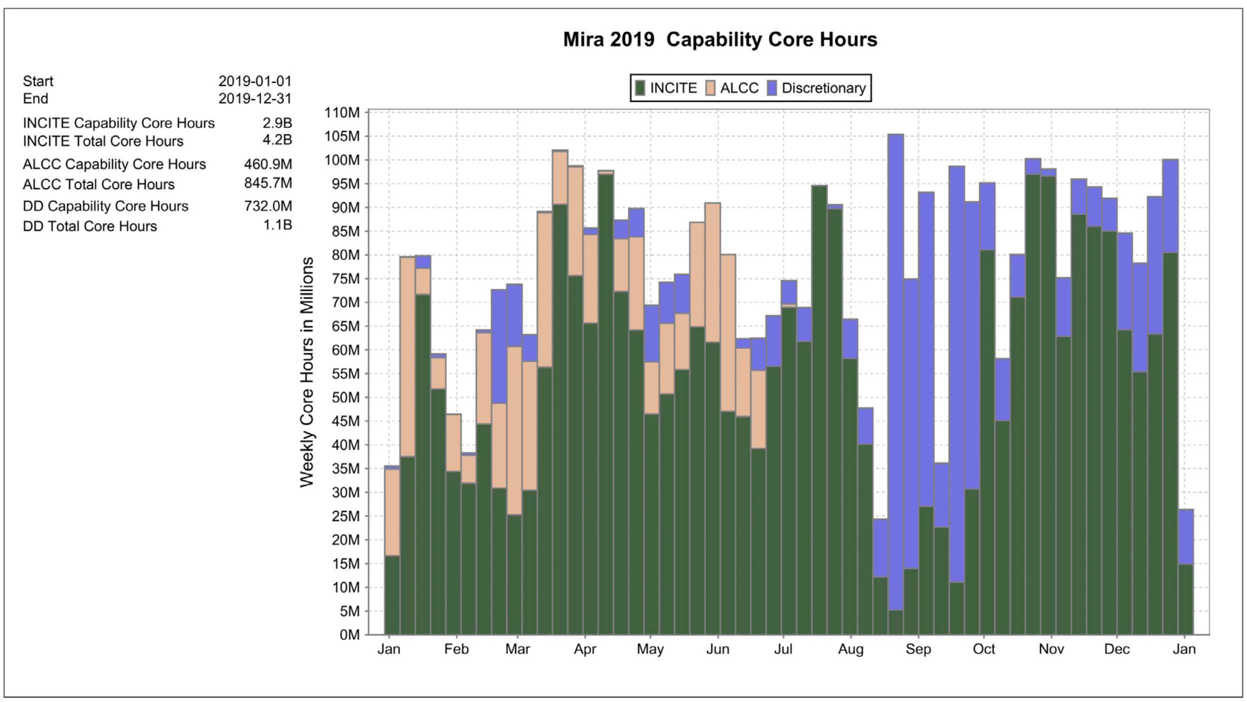

Figure 2.9 Mira Capability Core-Hours by Program

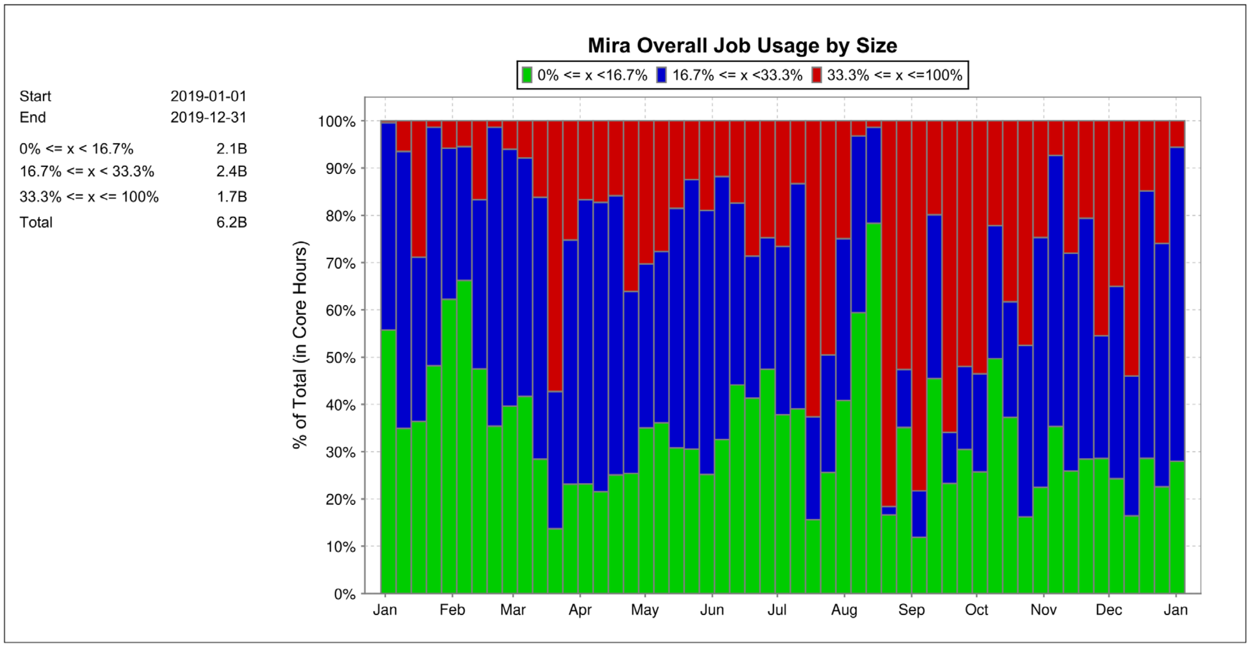

Figure 2.10 Mira Job Usage by Size

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 2.1.

ALCF closely tracks hardware and software failures and their impact on user jobs and metrics. These data are used as a factor in the selection of troubleshooting efforts and improvement projects. In CY 2019, this regular failure analysis has continued to drive code changes to Cobalt, ALCF's job scheduler, and has provided details to support debugging of storage system problems for Theta and Mira.

This page intentionally left blank.

Section 3. Allocation of Resources

Is the allocation of resources reasonable and effective?

ALCF Response

As our results show in Section 8, these are a reasonable allocation of resources. Here are a few data points the ALCF looks at when analyzing usage statistics for the various allocation programs.

3.1 Allocation of Resources

3.1.1 Usage of the INCITE and ALCC Hours

The INCITE 2019 program allocated 3.5 billion core-hours on Mira and 1.1 billion core-hours on Theta. The allocation usage for Mira is shown in Figure 3.1. The allocation usage for Theta is shown in Figure 3.2. A total of 4.2 billion core-hours was delivered to INCITE on Mira (Table 3.1). The contribution from Cetus to the total INCITE hours delivered was 0.03 percent. Of the 21 INCITE projects on Mira, 2 projects used less than 50 percent of their allocation, 4 projects used more than 90 percent of their allocation, and 15 projects used their entire allocation (or more), including 3 projects using over 150 percent. These projects used the extra core-hours to achieve additional milestones. A total of 1.2 billion core-hours was delivered to INCITE on Theta (Table 3.2). Of the 16 INCITE projects on Theta, 2 projects used less than 50 percent of their allocation, 3 projects used more than 70 percent of their allocation, and 11 projects used more than 90 percent of their allocation, including 8 projects that used their entire allocation (or more). These projects used the extra core-hours to achieve additional milestones. The overuse of Mira and Theta 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.

For the 2018–2019 ALCC year, 16 projects had allocations on Mira for a total of 845.7M core-hours (Table 3.3). The allocation usage is shown in Figure 3.3. Four of these projects used less than 50 percent of their allocation, 3 of them used more than 50 percent of their allocations, and 9 of them used 90 percent or more of their allocations. Cetus usage accounted for less than 0.02 percent of the total ALCC 2018–2019 usage. For the 2018–2019 ALCC year, 7 projects had allocations on Theta for a total of 167.5M core-hours. The allocation usage is shown in Figure 3.4. Two of these projects used less than 50 percent of their allocation, 1 of them used more than 80 percent of their allocation, and 4 of them used their entire allocations or more.

The 2019–2020 ALCC year is approximately halfway through its allocation cycle. So far, the 13 projects have received allocations of 192 million core-hours on Theta. The projects have used a total of 77.1 million core-hours from July 1, 2019, through December 31, 2019. The allocation usage is shown in Figure 3.5.

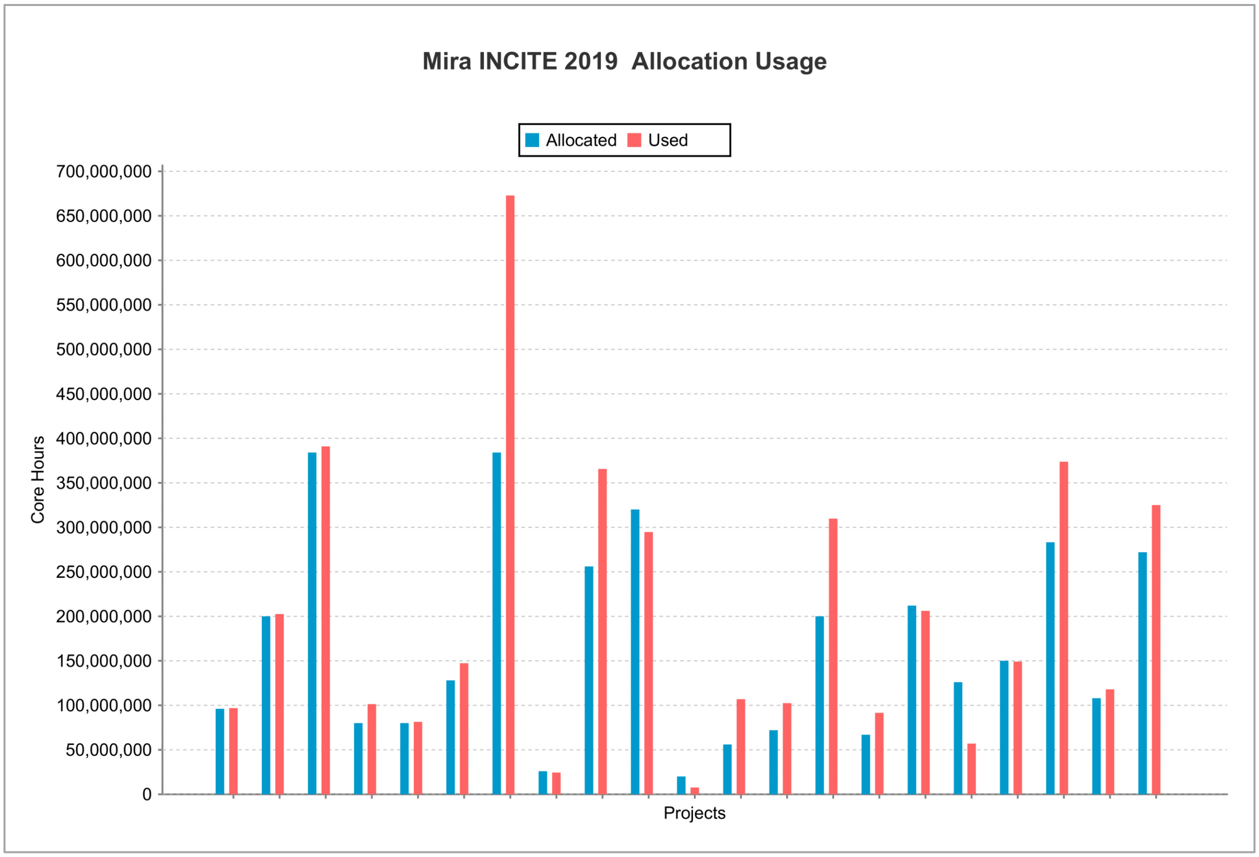

Figure 3.1 Mira INCITE 2019 Allocation Usage (Note: Projects are randomly ordered.)

Table 3.1 INCITE 2019 Time Allocated and Used on Mira

Projects	Mira
Allocated Core-Hours	3.5B
Used Core-Hours	4.2B [a]

[a] Usage includes 1.3M core-hours from Cetus production jobs.

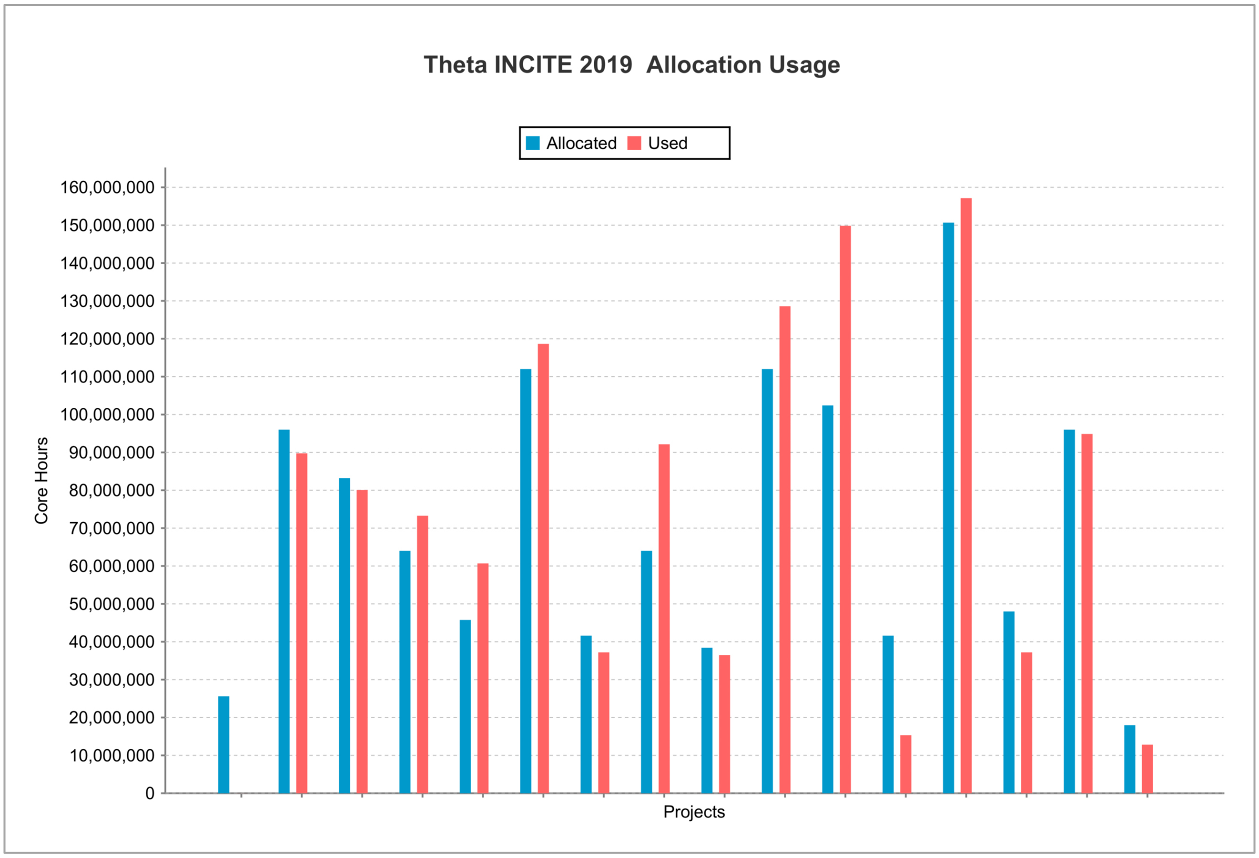

Figure 3.2 Theta INCITE 2019 Allocation Usage (Note: Projects are randomly ordered.)

Table 3.2 INCITE 2019 Time Allocated and Used on Theta

Projects	Theta
Allocated Core-Hours	1.1B
Used Core-Hours	1.2B

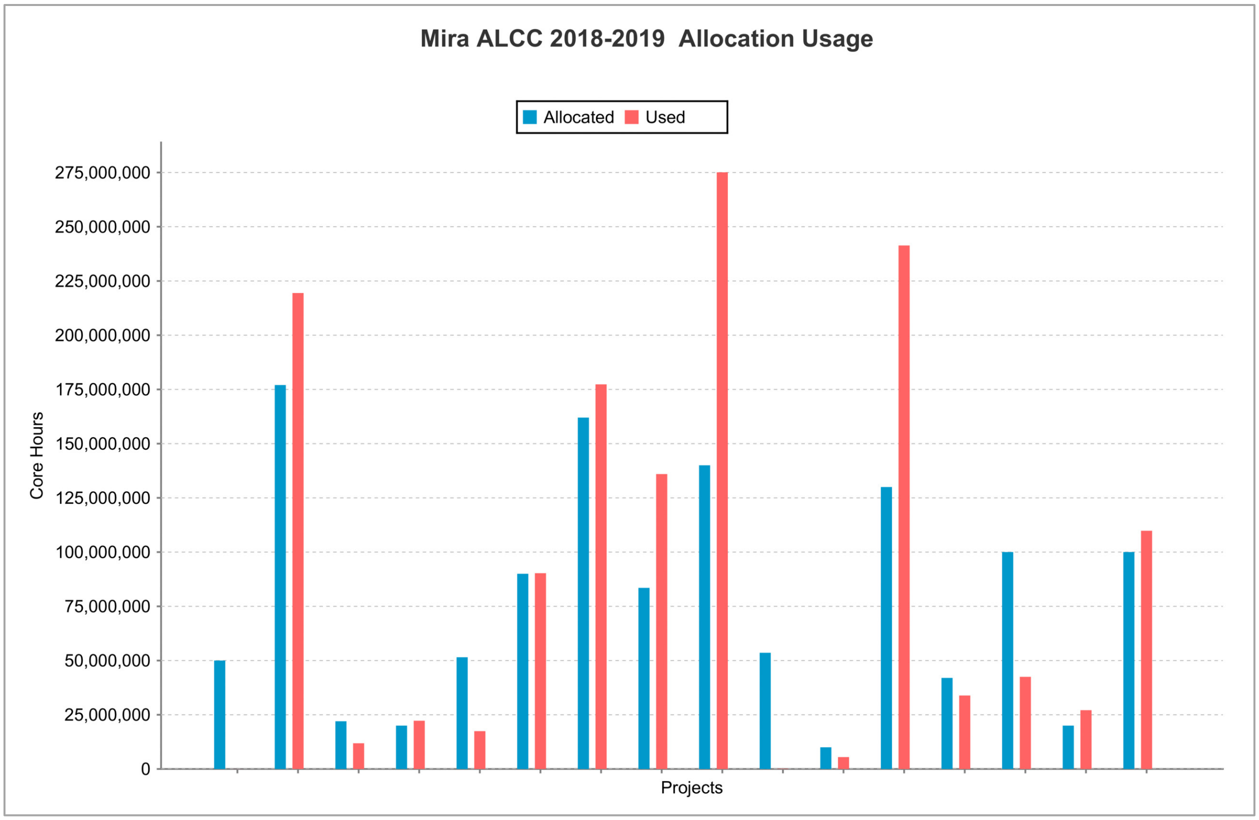

Figure 3.3 Mira ALCC 2018–2019 Allocation Usage (Note: Projects are randomly ordered.)

Table 3.3 ALCC Time Allocated and Used on Mira in CY 2019

Projects	Mira
Allocated Core-Hours	599.0M
Used Core-Hours	845.7M [a]

[a] Usage includes 0.2M core-hours from Cetus production jobs.

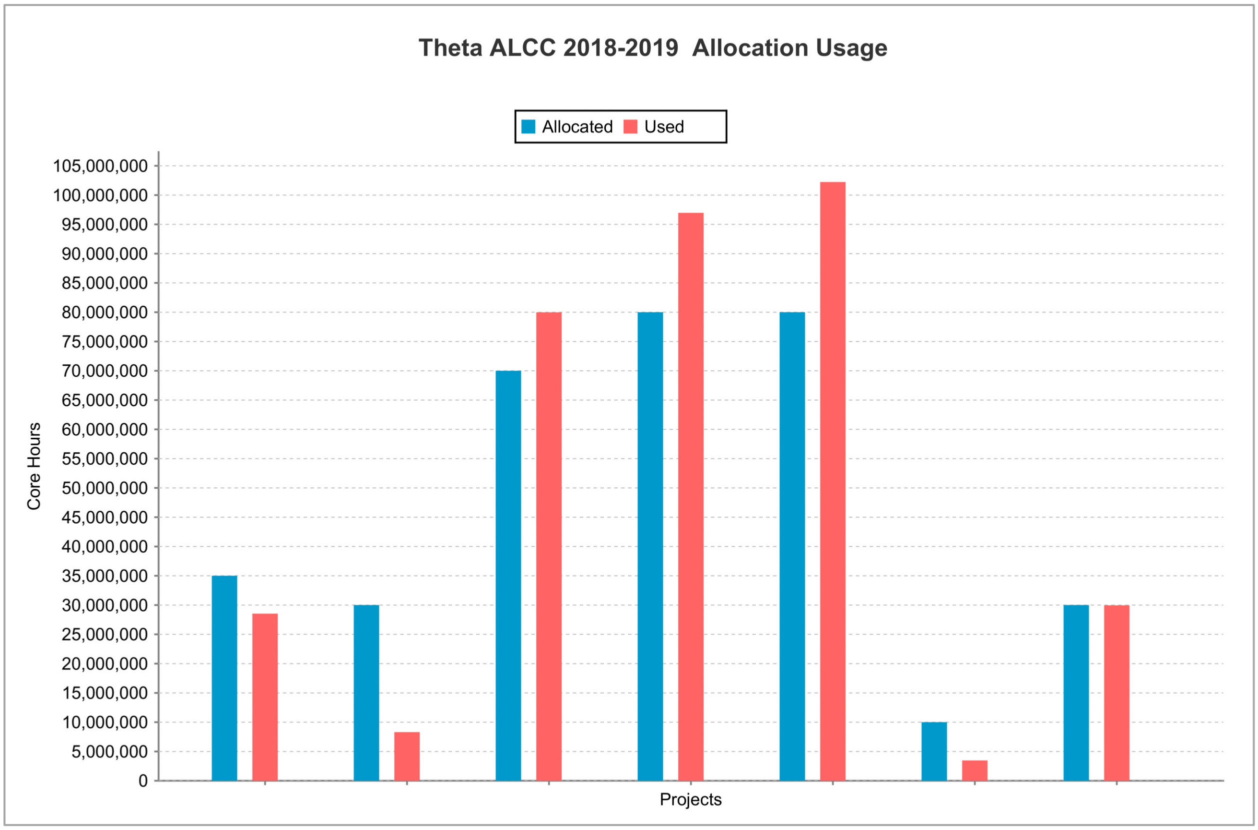

Figure 3.4 Theta ALCC 2018–2019 Allocation Usage (Note: Projects are randomly ordered.)

Table 3.4 ALCC Time Allocated and Used on Theta in CY 2019

Projects	Theta
Allocated Core-Hours	359.5M [a]
Used Core-Hours	210.6M [b]

[a] Allocation total is the sum of two different ALCC allocation years split in half.

[b] Usage total is the number of core-hours charged for jobs run against any ALCC allocation in CY 19.

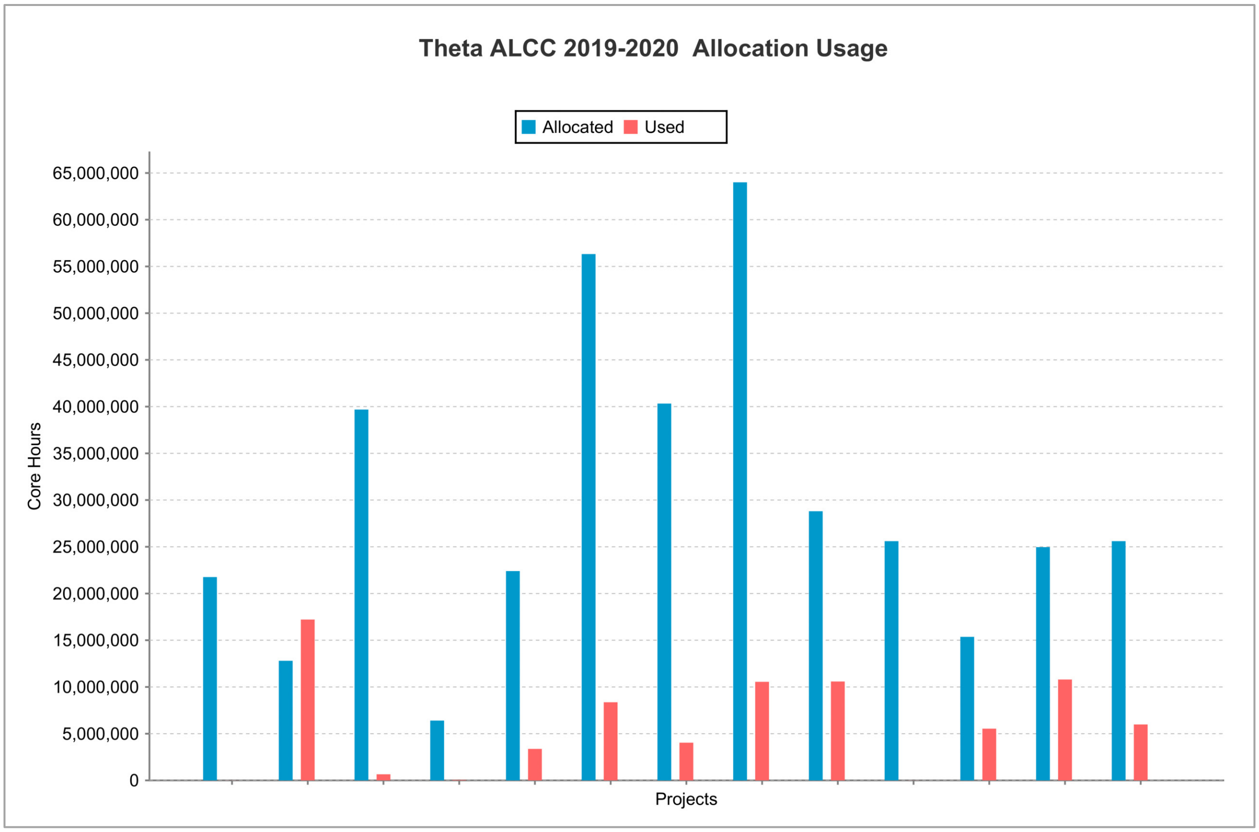

Figure 3.5 Theta ALCC 2019–2020 Allocation Usage (Note: Projects are randomly ordered.)

3.1.2 Facility Director's Discretionary Reserve Time

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

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

ALCF also allocates time to projects that might still be some time away from proposing for and receiving 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. The ALCF Data Science Program (ADSP) is allocated through the DD pool and is a program focused on developing the technical capabilities of data-driven projects that need leadership-class resources.

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 the staff support the users and maintain the system, such as diagnostics and testing of tools and applications.

As of the mid-point of 2018, the discretionary pool grew to 20 percent of the system to support the needs of the ECP program. ECP and the computing facilities run a Resource Allocations Council (RAC) that meets on a monthly basis to discuss the computing needs for ECP. Allocation needs are discussed and decided in the council, allocating up to 10 percent of the system.

Allocations are requested through the ALCF website and are reviewed by the Allocations Committee (which includes representatives from Operations, User Experience, 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 3.5 and Table 3.6 show the number of projects and total time allocated in the DD program on Mira and Theta, respectively, during 2019. By its very nature, the DD program is amenable to over-allocation because 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 3.5 DD Time Allocated and Used on Mira, 2019

Projects	Mira
Allocated Core-Hours	1.8B
Used Core-Hours	1.1B [a]

[a] Usage includes 27.6M core-hours from Cetus production jobs.

Table 3.6 DD Time Allocated and Used on Theta, 2019

Projects	Theta
Allocated Core-Hours	1.1B
Used Core-Hours	701.8M

A list of the CY 2019 DD projects on Mira, including title, PI, institution, and hours allocated, is provided in Appendix B.

Figures 3.6 and 3.7 provide a breakdown of the CY 2019 DD allocations by domain for Mira and Theta, respectively.

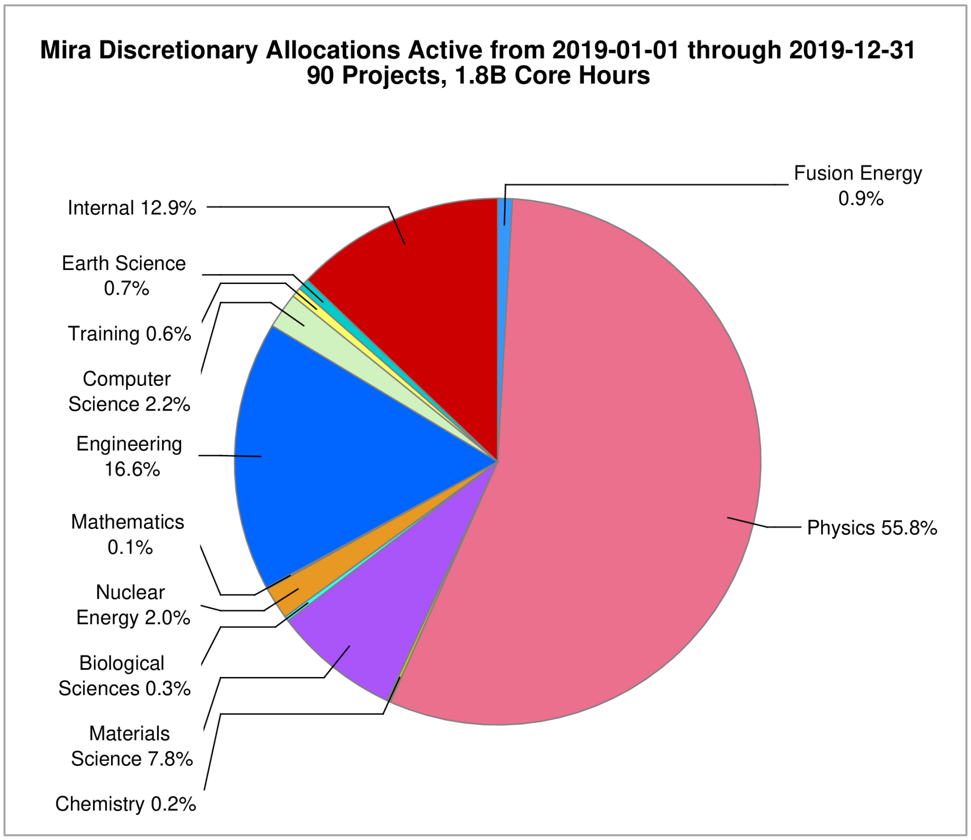

Figure 3.6 Mira CY 2019 DD Allocations by Domain

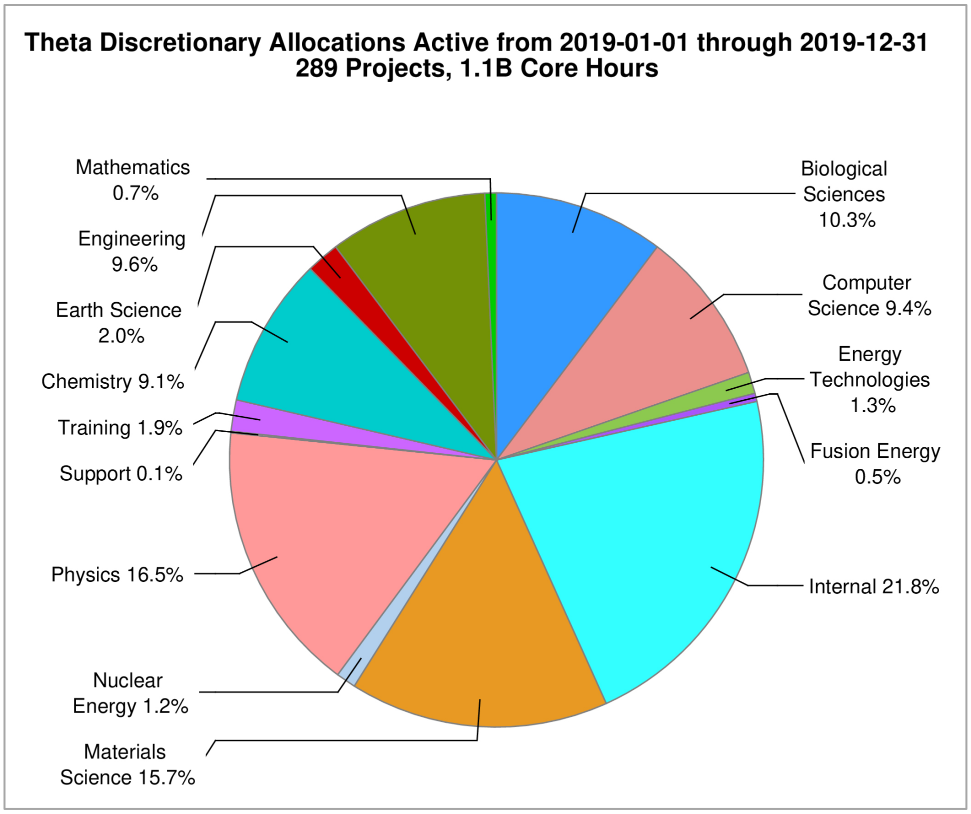

Figure 3.7 Theta CY 2019 DD Allocations by Domain

Conclusion

The ALCF delivered the following core-hours to the allocation programs in CY 2019: 5.4 billion to INCITE, 1 billion to ALCC, and 1.8 billion 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 to achieve science goals that could not otherwise have been reached.

This page intentionally left blank.

Section 4. Innovation

(a) Have innovations been implemented that have improved the facility's operations?

(b) Is the facility advancing research, either intramurally or through external collaborations, that will impact next generation high performance computing platforms?

(c) Is the facility effectively utilizing their postdoctoral fellows?

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2019. ALCF innovations and best practices have helped to prepare for future systems, have enabled more efficient operations, and have strengthened collaboration and engagement, both across ASCR facilities and beyond.

4.1 Operational Innovation

ALCF has undertaken several projects to improve the operations of ALCF and to better respond to user needs.

4.1.1 Balsam Backfill Queue

Challenge: The minimum job sizes mandated by ALCF's Theta queueing policy create an opportunity for unused nodes when jobs that are less than the minimum job size are launched. Meanwhile, DOE experimental science facilities desire integration with ALCF HPC resources to offload compute-intensive analyses to our systems.

Approach: By running the Balsam edge service on Theta, users can remotely submit analysis tasks over a secure connection without writing any job scripts or having any prior knowledge of Theta. In response to a backlog of tasks, the Balsam edge service then acquires nodes on a dedicated backfill queue, thereby "filling the gaps" in underutilized nodes without interfering with higher-priority, leadership-scale computing jobs. Remotely submitted analysis tasks are streamed onto these compute nodes, enabling near-on-demand analysis for experimental science facilities. To close the loop, Balsam integrates tightly with Globus Online and other protocols in order to transfer results back to the experimental facilities and stage-in new data for analysis in a continuous fashion.

Impact/Status: Balsam has been used to process X-ray photon correlation spectroscopy data from three DOE synchrotron facilities in a simulated 48-hour continuous experiment (Figures 4.1 through 4.3). This work was presented at the XLOOP Workshop at the 2019 International Conference for High Performance Computing, Networking, Storage and Analysis (SC19).

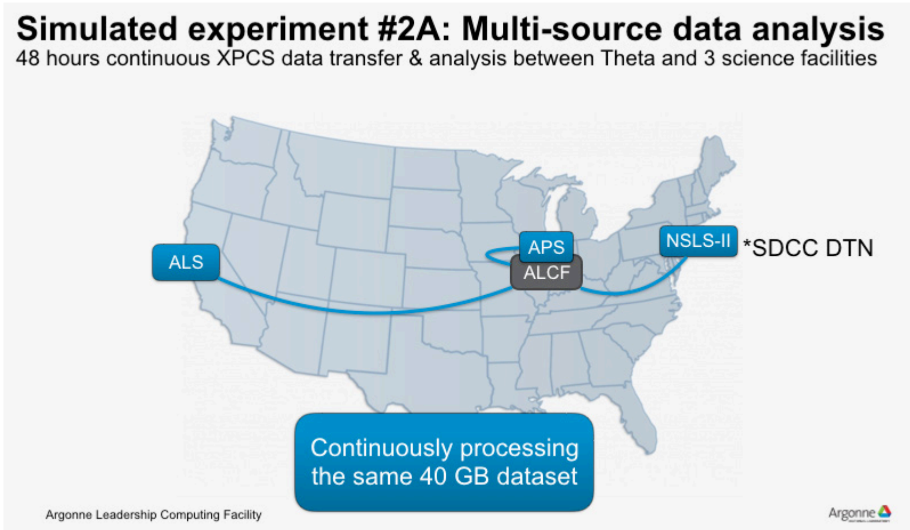

Figure 4.1 Three DOE synchrotron facilities process X-ray photon correlation spectroscopy data.

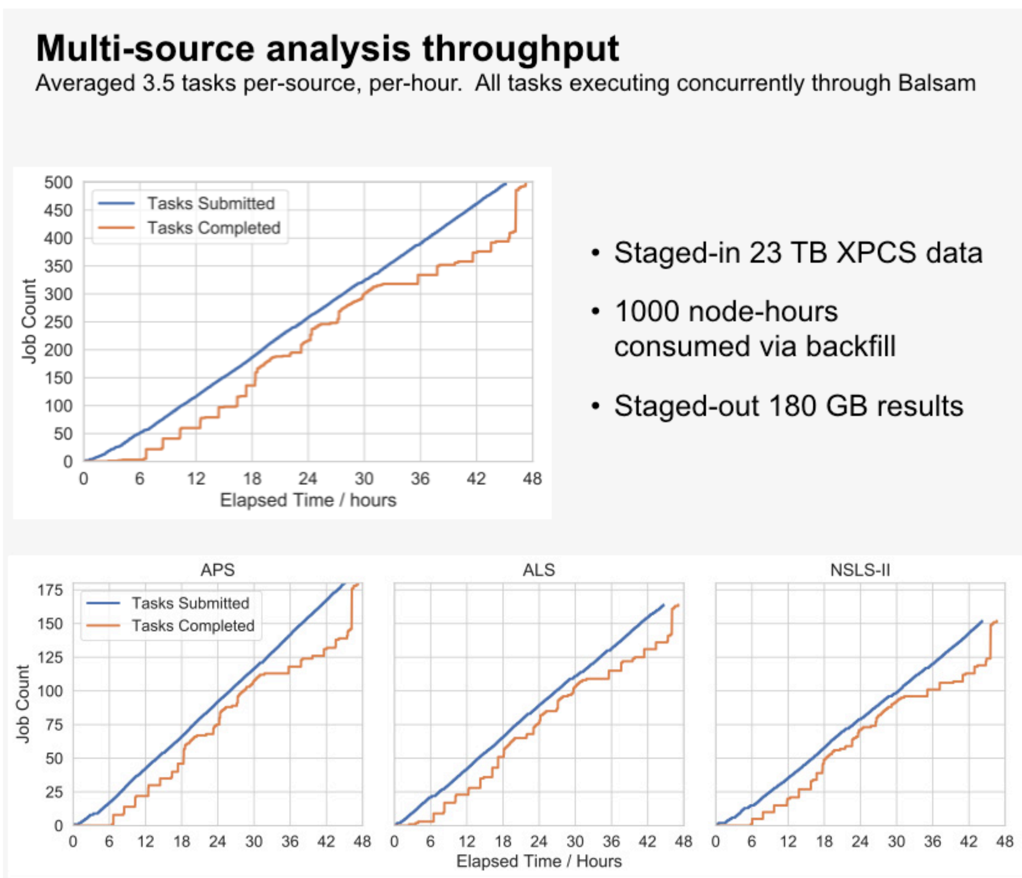

Figure 4.2 By running the Balsam edge service on Theta, users can remotely submit analysis tasks.

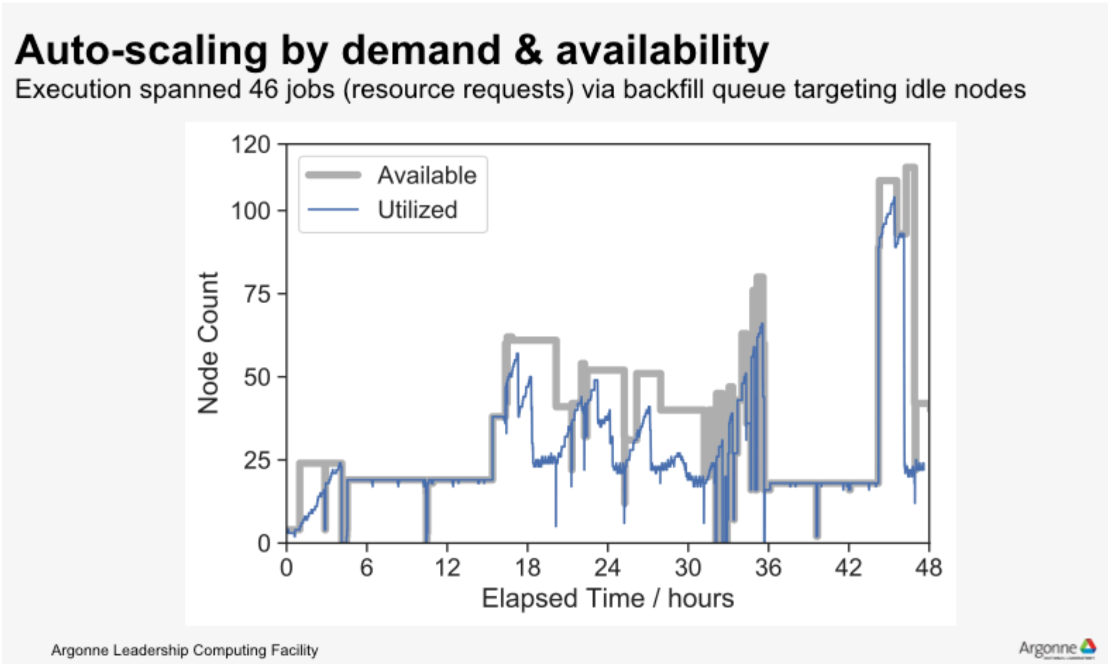

Figure 4.3 Remotely submitted analysis tasks are streamed onto underutilized compute nodes, enabling near-on-demand analysis for experimental science facilities.

4.1.2 Roll out of Mobile Tokens

Challenge: The ALCF has always issued new and replacement physical tokens to our users, which affects how quickly they can be assigned, transported, and enabled, especially for offsite or remote users.

Approach: After piloting mobile tokens internally in 2018, the ALCF began a phased rollout of "soft" tokens to all users in April 2019.

Results/Impact: The response has been overwhelmingly positive, with more than 80 percent of new users opting for the mobile version over the physical one. The chief advantage of mobile tokens is how quickly they can be assigned and enabled, which is especially useful when multiple user accounts for workshops and training events are needed. Whereas mobile tokens can be used as soon as an account is created, physical tokens must be mailed out by ALCF staff and unlocked via a Help Desk phone call, and requires the Help Desk to set a temporary PIN and the user to set a new PIN after initial use.

4.1.3 User Base Development

Data Accuracy

Challenge: ALCF's legacy account and project management system, UB1, relied on input in the form of free-form text; therefore, typos and inconsistent identifying terminology (ANL, Argonne, etc.) affected overall data quality.

Approach: Several approaches to improve data quality were applied to the new software, including:

- A collaborative effort involving members of the Business Intelligence (BI) team, the Advanced Integration Group (AIG), and User Experience team ran multiple iterations through the data-cleaning "pipeline" to make sure that the data were accurate and consistent while migrating from the legacy system.
- Increasing the amount of data validation at the time of input and providing drop-down menus where possible to avoid errors previously introduced by free-form text fields.
- Storing multiple institutions as individual records and using relations rather than a concatenated string.
- Identifying individual PIs, co-PIs, and proxies for a project within the system. UB1 could only designate an ALCF account holder as a "PI" in the system.

Results/Impact: Quality of the data imported into the new software was improved. It is now much easier to keep the institutions accurate, and while this does not eliminate periodic review of the data entered in the system, it reduces the inconsistencies introduced. In the legacy system, a PI who did not have an ALCF account had to designate a proxy and the proxy's name appeared as the PI for the project. In addition, UB1 did not have the ability to store co-PIs. UB3 eliminates these problems.

Security Improvements

Challenge: UB1 had only password-level security, forcing users and staff to maintain a separate password that was easily forgotten with little use. UB1 also stored Personally Identifiable Information (PII) and had no defined administrator roles to protect privileged data.

Approach: UB3's authentication mechanism is configurable. It currently supports local two-factor authentication, as it is used elsewhere in ALCF. UB3 has defined roles, such as limited user vs. authenticated user or admin (read-only) vs. admin (read/write). PII is never stored in UB3. UB3 also has extensive logging, and "triggers" can be defined that send an e-mail for activity deemed suspicious.

Results/Impact: UB3 uses the same two-factor authentication as required to access ALCF HPC resources, which eliminates the need for users to maintain a separate password for Userbase. UB3 has a more robust, role-based access control and integration with other Argonne systems, which means PII does not need to be stored in Userbase.

Improved User Interface

Challenge: UB1 was not designed with mobile devices in mind, and the user interface design was inconsistent with the rest of the ALCF web presence. Error messages and help prompts were confusing or absent.

Approach: Screen arrangement, fonts, sizes, and color were designed to indicate hierarchy and improve clarity. The entire interface was designed to be used on desktop and mobile devices with equal ease. The Help function, error messages, and confirmation prompts were improved to help users navigate and understand what input is needed.

Results/Impact: UB3 can be used equally well on desktop and mobile devices. Prompts and help screens are present to guide users and speed entry of needed data.

Automated Testing

Challenge: While the UB3 backend had some automated tests; the front end was largely tested by hand, following scripts.

Approach: The ALCF hired a dedicated test engineering contractor to help improve testing, both in terms of the quantity and quality of the tests, but also of best practices. The ALCF is now using Cypress to perform automated front-end testing, and tests are now required before a feature or bug fix can be merged into the develop branch. The ALCF also added back-end tests via Cypress by directly calling the various REST endpoints. Cypress is a next-generation, front-end testing tool built for the modern web, with an all-in-one testing framework and assertion library, with mocking and stubbing. It addresses the key pain points developers and quality assurance (QA) engineers face when testing modern applications.

Results/Status: Implementing automated testing has resulted in slower turn-around on new features and bug fixes; however, higher-quality code with fewer bugs will save ALCF time later and provide a better user experience. Writing end-to-end tests requires that a lot of different tools work together. With Cypress, ALCF gets multiple tools in one installation.

4.1.4 Archiving Software on Theta

Challenge: The need existed to shorten the list of available software presented to users on Theta to improve user experience while also preserving the ability to have all previous software versions available, and to speed provisioning at boot time.

Approach: Library tracking data from XALT was used to determine the usage of libraries and compilers. From that, the versions used most often must be identified to be kept as is, while unused versions must be moved to an archive that is not the default view but is still very easy to access and use.

Results/Status: The archiving process initiated in late summer archived 11 Intel compilers, 9 GNU Compiler Collection (GCC) compilers, and 2,613 Cray-built libraries and third-party packages, thereby reducing the default list of software by roughly 50 percent (Figure 4.4). This process arguably improves the experience of users when they are looking for software, as there are fewer screenfuls of information to digest. Yet, all of the older versions that were archived were put in an "archive" that is easy to use if a user still requires an older version. A user simply does "module load archive" and all versions are available if needed. This approach also allows more focused user support as a core set of compilers and libraries are now used.

Much of the deployed software environment was moved out of the boot image and into the archive, which is available on an NSF-like, shared GPFS filesystem. Operationally, this move benefits the ALCF since the software image that has to be pushed out to each of the nodes is much smaller now, which helps reduce machine reboot times. Also, with this new archive area, software can actually be installed to the shared area first rather than installing it in Cray's system image and then moving it later, which has the benefit of requiring less staff time to manage the installs as well as reducing chances for mistakes; the ALCF controls what is seen by default and what is in the archive with the use of modules, which are easy to move when archived.

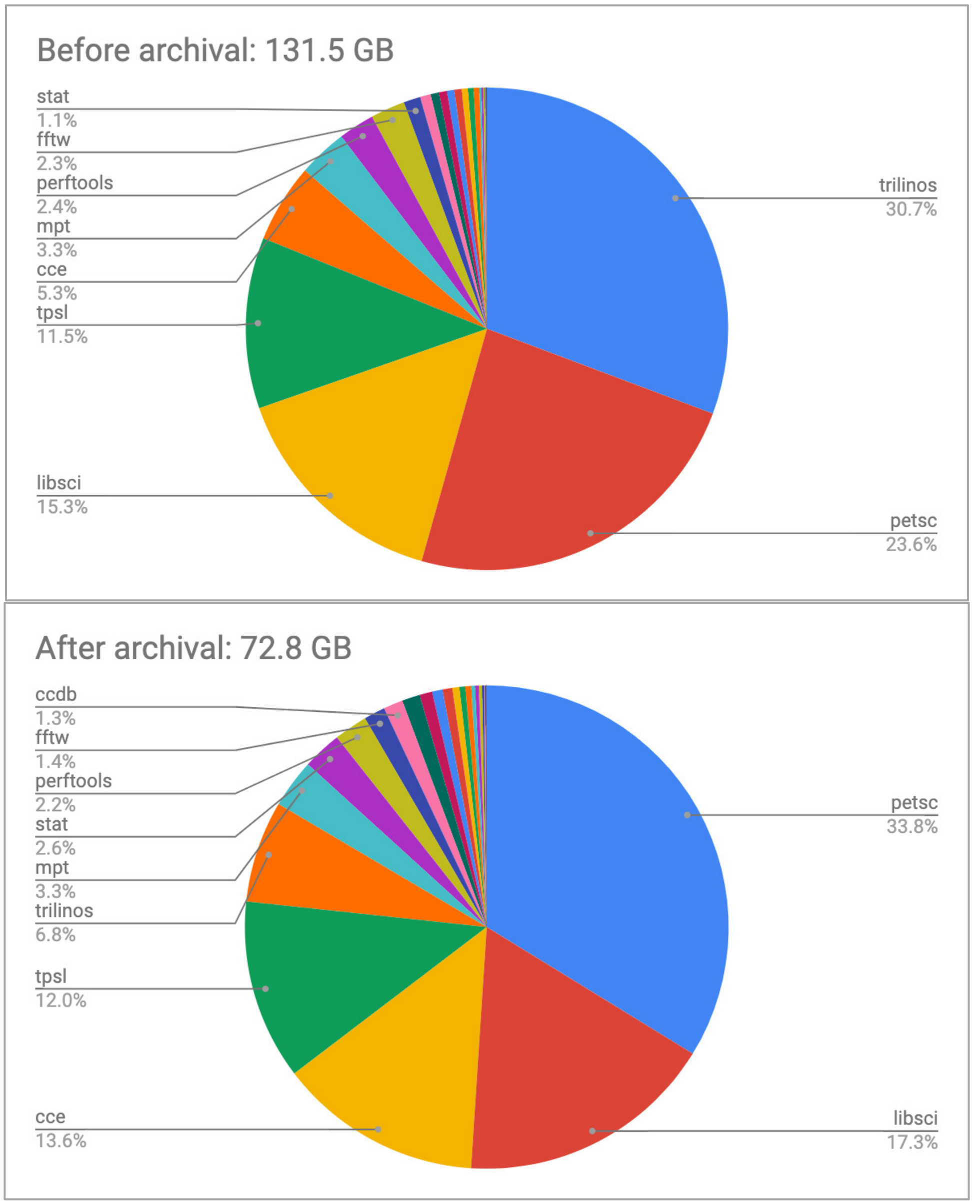

Figure 4.4 Graphical breakdown of available software on Theta before and after the archiving process initiated last fall. The effort reduced the default list of software by roughly 50 percent.

4.1.5 UPS Replacement Project

Challenge: All three of ALCF's existing Chloride uninterruptible power supply (UPS) units were aging, inefficient, and used lead acid batteries. In addition, one system was not reliable when operating on battery power.

Approach: The ALCF replaced the three 750kVA Chloride UPS units with three 750kVA Galaxy VX units. Although the new Galaxy units are not a unity power factor type, they are more efficient and stable than the old units. In addition, the new units have scalable, three-phase power protection with flexible operating modes and a new form of eco-mode, called *ECOnversion*. The units also have lithium-ion batteries replacing the former lead acid units.

Results/Status: The new Galaxy VX units appear to be 95 percent efficient at this load and possibly more efficient at higher loads. As Figure 4.5 shows, the new Galaxy UPS is consuming 274kW on the input and delivering 261kW on the output, showing 95 percent efficiency at this load. ECOnversion mode ensures ultra-high efficiency while charging the batteries, conditioning the load power factor, and ensuring a Class 1 output voltage regulation.

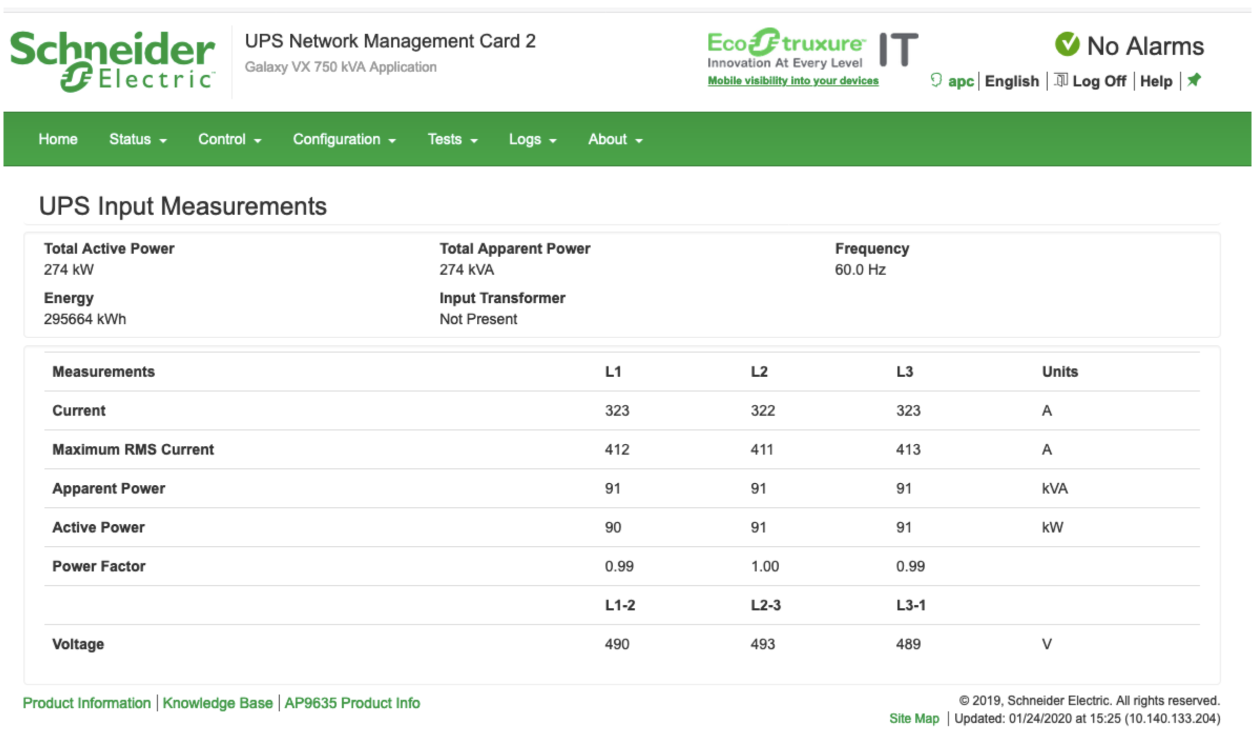

UPS Input Measurements

Total Active Power	Total Apparent Power	Frequency
274 kW	274 kVA	60.0 Hz
Energy	Input Transformer	
295664 kWh	Not Present	

Measurements	L1	L2	L3	Units
Current	323	322	323	A
Maximum RMS Current	412	411	413	A
Apparent Power	91	91	91	kVA
Active Power	90	91	91	kW
Power Factor	0.99	1.00	0.99	
	L1-2	L2-3	L3-1	
Voltage	490	493	489	V

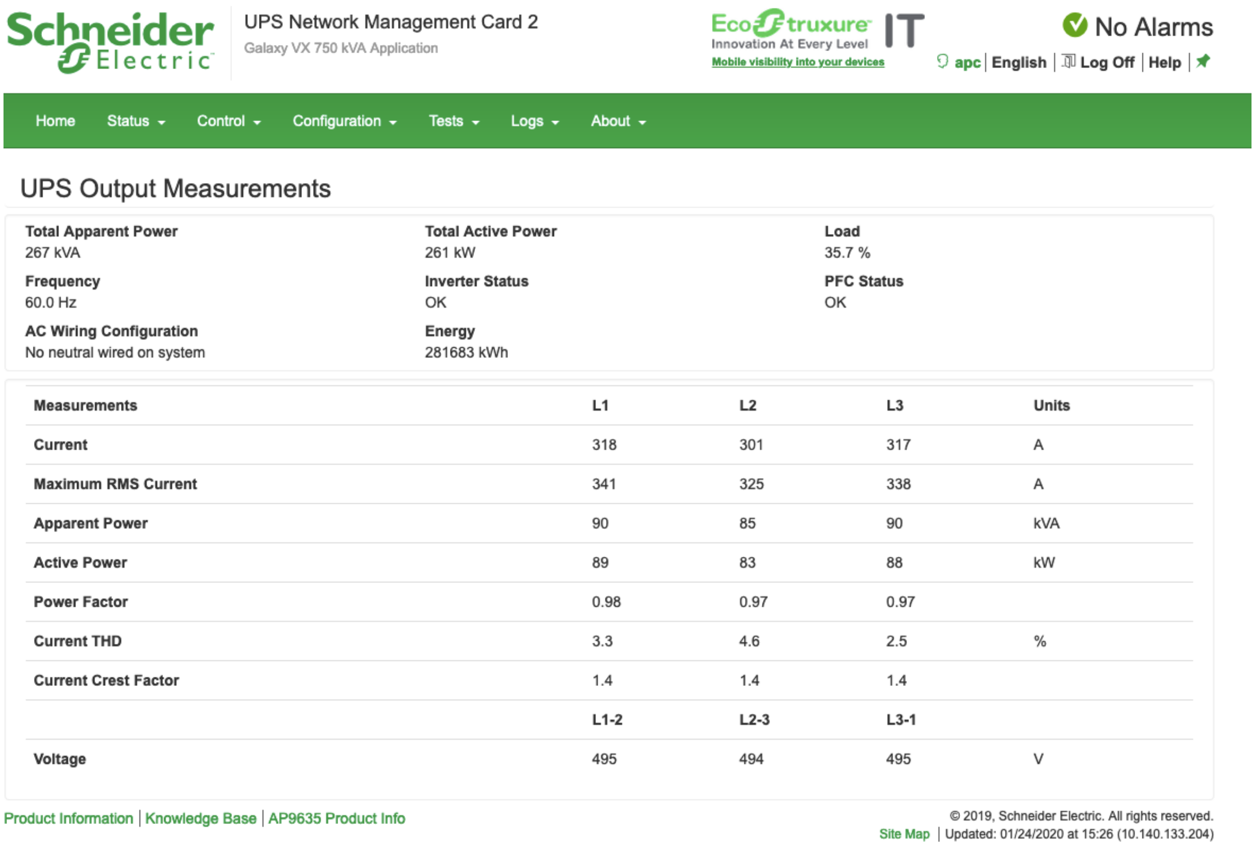

Measurements	L1	L2	L3	Units
Current	318	301	317	A
Maximum RMS Current	341	325	338	A
Apparent Power	90	85	90	kVA
Active Power	89	83	88	kW
Power Factor	0.98	0.97	0.97	
Current THD	3.3	4.6	2.5	%
Current Crest Factor	1.4	1.4	1.4	
	L1-2	L2-3	L3-1	
Voltage	495	494	495	V

Figure 4.5 Screen captures showing the new Galaxy VX UPS's input and output measurements. The new UPS is consuming 274kW on the input and delivering 261kW on the output, showing 95 percent efficiency at this load.

4.2 Research Activities for Future Operations

The ALCF works to prepare for next-generation systems through collaboration with vendors and other DOE facilities. The Facility reports on participation in research projects in Section 1.3.2.1 and professional community activities in Section 8.3.1.

4.2.1 Cray Benchmark for Congestion

Challenge: Applications running on Theta, a Cray XC-40, can experience varying performance due to perceived network performance. This variation is caused by network traffic on the system from other applications running during the same time period. Although the property of variable performance related to shared resources is widely known, there has not been a good systematic benchmark to evaluate the effects of the variance and to characterize between different network designs.

Approach: ALCF, in collaboration with Cray and NERSC, developed and executed a new benchmark specifically to evaluate network variability induced by competing workloads. The benchmark consists of a systematic set of measurement "probes" and "congestor" workloads. The congestor workloads emulate common communication patterns that can cause disruption to other workloads on the system. The measurement probes evaluate the effect of these congestors to the performance perceived by the probe.

Results: The benchmark and results of the study on some common systems were published at SC19. With this publication, the benchmark is available for future interconnect providers to evaluate the performance of their networks for these properties. Future procurements may also use this work to guide machine design and evaluation.

Cray, in conjunction with NERSC and ALCF, developed a suite of benchmarks representing the real workload characteristics to study the effects of congestion and its management on various networks. Experiments using these benchmarks under controlled conditions on various networks demonstrated the limitations and promise of current- and future-generation networks, respectively.

4.2.2 Automatic Testing of Compilers

Challenge: For applications to run successfully on Aurora, an optimized, validated compiler is essential. Tracking the specific compiler functionality needed for applications as the compiler is developed helps ensure that the necessary applications can run on Aurora as soon as possible.

Approach: The ALCF developed a test set and automated testing framework to help improve and ensure the quality of the new Intel compilers for Aurora. For example, for each compiler bug reported to Intel, we added a test to check when the bug is resolved and to ensure the same bug does not occur in future versions of the compiler. We chose to use a lightweight and simple testing framework, bats-core, to lower the barrier of entry and make it easy for application developers to add tests.

The test set is in a local git repository in the Joint Laboratory for System Evaluation (discussed in Section 4.2.3). ALCF staff can add tests either by adding an entry to a table on the JLSE confluence page or by cloning the test set, adding a test to it, and pushing it back to the repository. The bugs are maintained in a table on the JLSE confluence page. As part of the collaboration with Intel for Aurora, colleagues at Intel add Intel internal bug report numbers for each bug we report.

The test set is organized by programming model, tool, or library (e.g., OpenMP, DCP++, OpenCL, MLK) to allow the bugs to be sorted and organized more easily. To add a new test, someone simply has to add a new subdirectory under the programming model, tool, or library with their code, and add a short .bats file. The .bats file is composed of bash functions, which can contain any content that the application developer would like to test, such as building and running the code. If the exit code of a test is 0, then it passes and otherwise it fails. This approach allows the person reporting the bug flexibility in how to express the bug (compilation error, runtime failure, etc.). A driver script runs all the tests and saves the output by date.

We additionally incorporated the OpenMP Validation and Verification Suite from the Scaling OpenMP Via low-level virtual machine (LLVM) for Exascale Performance and Portability (SOLLVE) project.

Impact/Status: The test set is available on JLSE so that ALCF staff can add tests to it. It continually evolves with ALCF staff adding tests based on bugs they find in the compiler when porting their code to Aurora. The testing framework is run with each new Aurora software development kit (SDK) drop, and it allows us to monitor and provide feedback to Intel as quickly as possible. At this time, more than 40 tests/bugs reports were added by more than a dozen people.

4.2.3 Joint Laboratory for System Evaluation

Argonne's JLSE enables researchers to assess and improve next-generation computing platforms of interest to the DOE. Established by the CELS computing divisions and run by the ALCF, the JLSE centralizes Argonne's research activities aimed at evaluating future extreme-scale computing systems, technologies, and capabilities. JLSE users leverage existing infrastructure and next-generation hardware and software to explore low-level experimental computer and computational science, including operating systems, messaging, compilers, benchmarking, power measurements, input/output (I/O), and new file systems. By providing access to leading-edge computing resources and fostering collaborative research, the JLSE enables researchers to address Argonne and DOE needs in a variety of areas, including by:

- Improving science productivity on future hardware and software platforms.
- Providing an avenue for Argonne researchers to work collaboratively with HPC vendors on prototype technologies for petascale and beyond.
- Investigating alternative approaches to current and future system deployments.
- Maintaining a range of hardware and software environments for testing research ideas.
- Helping to drive standards on benchmarks, programming models, programming languages, and memory technologies, etc.

The joint laboratory provides users with access to several diverse testbeds, including:

- An Intel Xeon Phi (Knights Landing) and Xeon (Skylake) cluster
- An NVIDIA DGX-1 System (Volta graphical processing units [GPUs])
- A Skylake Xeon cluster with 4 V100 SMX2 GPUs in each node
- An HPE Apollo Cluster with 72 ARM ThunderX2 processors
- An Atos Quantum Learning Machine
- A Kubernetes testbed
- A Lustre testbed
- An IBM AC922 air-cooled Summit-like node with 2x 20core P9 processors and 4x Nvidia V100 GPUs with 16GB HBM
- An Intel Skylake E3-1585v5 Xeon cluster with Gen9 integrated GPUs for preparing for Aurora (accessible by ESP and ECP projects)
- A Ceph 3.2PB filesystem named Petrel v3
- Intel Distributed Asynchronous Object Storage (DAOS) testbed for preparing for Aurora

In 2019, the JLSE supported over 500 users spanning more than 70 projects. These projects ranged from application portability to software development to tools and compiler development for an ALCF Early Science Project. Teams from within the ECP's Application Development and Software Technology groups, developing applications and software for Aurora, have had access to the JLSE Gen9 nodes and the Aurora SDK for their development work. The following summaries represent a sampling of current JLSE projects:

ALCF Data Science Program: Application teams from the ALCF Data Science Program use JLSE resources to explore and improve data science techniques, such as data mining, graph analytics, machine learning, and complex and interactive workflows.

ALCF Early Science: Application teams from the ALCF Early Science Program use JLSE resources to prepare and optimize applications for the next-generation supercomputers in advance of the systems becoming available. For example, researchers from the Aurora ESP projects access the Xeon Skylake Iris nodes with Intel's integrated GPUs and the early versions of OneAPI software to develop and test their applications for Aurora.

PetrelKube: This small Kubernetes testbed is managed with Rancher and hosts back-end services supporting web portals running outside of the JLSE and HPC workflows. The testbed utilizes Petrel for its primary data storage, initially staging data over the network and eventually having direct access to Petrel via a POSIX interface.

Neuromorphic: Researchers explore software approaches to power efficient implementation of machine learning algorithms on specialized hardware like Neuromorphic chips.

Argo: Argo is a new exascale operating system and runtime system designed to support extreme-scale scientific computation. Researchers from the Argo project used JLSE resources to prototype the GlobalOS distributed resource management and to evaluate the performance of NodeOS. They also used the laboratory's testbeds to develop and optimize a lightweight, low-level threading and task framework for OpenMP and other programming models (Cilk, Quark, Charm++).

Big Data: Researchers are using JLSE testbeds to study the layering of HPC programming models beneath big data programming models. Specifically, they are researching the development of a software environment with a Spark user interface (Java and Scala) that can run on a supercomputer, cluster, or cloud with a back end for executing data-intensive communication patterns.

CANDLE: Using the NVIDIA DGX-1 system and other JLSE computing resources, researchers are developing the CANcer Distributed Learning Environment (CANDLE), a computational framework designed to facilitate breakthroughs in the fight against cancer.

Deep Learning: Multiple projects are using JLSE systems to investigate the potential of deep learning. One research team is focused on understanding how deep learning can be used to improve lossy compression of scientific data from simulations and instruments. Another effort is exploring the performance of different machine learning frameworks that have implemented deep learning and neural networks on KNL systems.

LLVM: Researchers used the JLSE's IBM power systems to advance LLVM compiler development. The instruction set architecture for these systems is the same as for the IBM Blue Gene/Q system, with the only difference being in vectorization. LLVM and CLANG builds were carried out on the Intel Xeon Phi systems for QA purposes. Researchers can complete these builds in 10 minutes using JLSE resources (compared to hours on a laptop).

MPI: A number of MPI Chameleon (MPICH) improvements were tested on JLSE systems, including the memory scalability of MPI communicators by exploiting regular patterns in rack-address mapping, enhanced threading support through locking optimizations, and communication-aware thread scheduling.

Quantum Computing: A research team is using the JLSE's Atos Quantum Learning Machine and other resources to develop an artificial neural network for spectral analysis called Spectranne. This tool will automate the analysis of vast amounts of data being produced by state-of-the art, chirped-pulse spectroscopy experiments.

4.3 Postdoctoral Fellows

The ALCF supports a steady-state postdoctoral fellowship program. Within the purview of this program, the ALCF supports one named postdoctoral fellow, the ALCF's Margaret Butler for Computational Science Fellow. Postdocs are awarded 1-year appointments that have an option to be renewed for an additional year (this is typically the case), with a similar option to renew for a third year. The major goal of the program is to either convert the postdoc to a regular staff appointment, place them at another DOE laboratory, or support their efforts to find an academic or industry appointment. The objective, in all cases, is that these postdocs continue to be lifelong users of DOE compute resources.

The general process is that applications for postdoctoral positions are handled by Argonne's Postdoctoral Program Office. In CY 2019, the ALCF hired one new postdoctoral researcher (fellow) from a pool of 17 applicants, and supported a total of eight; representing a range of scientific domain backgrounds. The ALCF steady-state postdocs work on various research topics, including physics-informed machine learning; heterogeneous computing and program modeling; data-intensive workflows and hyperparameter optimization; uncertainty quantification; development of machine-learned surrogate models; integration of data science into simulation environments; analysis of distributed training of Bayesian neural networks on HPC systems; and developing multidimensional free energy computational methods for membrane protein systems.

Once hired, each postdoc is assigned both a direct research supervisor and an Argonne staff mentor. The mentor, initially selected by the division or the supervisor, can be changed by the postdoc. The supervisor meets with the postdoc weekly and engages in postdoc's research effort. The supervisor then evaluates the progress and completes a standardized review process that is submitted to the ALCF management for review, including the division director, who reviews and authorizes all appointment renewals. The mentor is responsible for meeting with the postdoc to discuss career development milestones and personal goals; this happens as needed, but no less than once a quarter. The guidance for these discussions includes key skills the postdoc could focus on over the next year; opportunities for development; and, if entering the third year, what will be most beneficial to enabling a smooth career transition. The division director also meets monthly with the postdocs as a group to hear progress updates, address any issues specific to the postdoc community, and to solicit feedback in general.

The ALCF supported the following postdoctoral researchers in CY 2019:

Romit Maulik (Ph.D., mechanical & aerospace engineering, Oklahoma State University). Margaret Butler Fellow. **Hired:** January 2019. **Research area:** Physics-informed machine learning. **Current projects:** 1.) development of machine-learned surrogate models for computational physics and computational chemistry; 2.) development of turbulence closure models for large eddy and Reynolds-averaged Navier Stokes simulations; and 3.) integration of data science into simulation environments for in-situ inference and training. **Accomplishments include**: five publications (two under review, one accepted in-press, two published); four invited talks; one invited poster, and three workshops as an invited participant.

Zheming Jin (Ph.D., computer science & engineering, University of South Carolina). **Hired:** January 2017. **Research area:** heterogeneous computing and program modeling. **Projects:** Aurora ECP. **Accomplishments include**: 14 publications, including numerous case studies using Intel FPGA, CPU, and GPU platforms. **Next employer:** TBD in June 2020.
Pankaj Rajak (Ph.D., mechanical & aerospace engineering, University of Southern California). **Hired:** October 2018. **Research area:** physics-informed machine learning. **Current projects:** 1.) Aurora Early Science Program projects; NAQMD-RMD. **Accomplishments include**: four publications, including one in *MRS Advances*, one in *Phys. Rev. B*, one in *J. Phys. Chem. Lett.*, and one in *SoftwareX*; and three invited talks.

Misha Salim (Ph.D., chemistry, University of Illinois at Urbana-Champaign). **Hired:** October 2017. **Research area:** Data-intensive workflows; hyperparameter optimization; machine learning. **Projects:** 1.) Balsam (ALCF's workflow management system and edge service); 2.) Balsam+DOE light source integration (infrastructure and testing capability to process XPCS (or similar) data from beamlines at NSLS-II, APS, ALS); 3.) DeepHyper: Framework for neural architecture search and hyperparameter optimization of deep learning models at scale; 4.) 2D magnetic materials ADSP project (designing workflow to run campaign of density functional theory (DFT) optimization and magnetic moment calculations); 5.) AI for Materials ADSP project (machine-learning the screened Coulomb interaction to accelerate many-body perturbation theory calculations for materials). **Accomplishments include**: Paper titled "Scalable-reinforcement learning based Neural Architecture Search for Cancer Deep Learning Research" accepted at SC19; awarded Best Presentation at the XLOOP Workshop at SC19 for a talk on Balsam. **Status update:** Reclassified in October 2019 to Assistant Computational Scientist.

Himanshu Sharma (Ph.D., mechanical engineering, Iowa State University). **Hired:** March 2019. **Research area:** HPC, machine learning, uncertainty quantification, fluid mechanics, urban modeling and building systems. **Current projects:** 1.) analyzing distributed training of Bayesian neural networks on HPC, 2.) Framework development for embedding and interfacing machine learning models with scientific numerical solvers; and, 3.) attending vendor interaction meetings and participating in activities of developing and deploying software for Aurora. **Accomplishments include**: journal publication (in preparation); two posters (one given, one accepted); one invited talk; and implemented a neural network model for gravitational wave parameter estimation for an ADSP project.

Ganesh Sivaraman (Ph.D., engineering physics, University of Stuttgart). **Hired:** July 2017. **Research area:** computational science; machine learning. **Current Projects:** code developments for various projects, including: 1.) developed an Apache-Spark-based distributed machine learning workflow for chemical property prediction. This workflow is being used by the ALCF data science team for Spark A21 benchmark tests; 2.) implemented an active learning workflow for automating the mapping of AIMD trajectory to ML force-field generation for LAMMPS; 3.) implemented VASP support for BALSAM workflow; and 4.) provided MongoDB support and documentation for ALCF. **Accomplishments:** Three publications, one invited talk, and one contributed talk.

Jonathan Thirman (Ph.D., chemistry, University of California-Berkeley). **Hired:** August 2017 (terminated in August 2019). **Research area:** chemistry. **Projects:** NAMD. **Accomplishments include:** 1.) significantly contributed to the development of the Replica Exchange Solute Tempering (REST2) module in NAMD; 2.) developed multidimensional free energy computation methods for membrane protein systems, which are the focus of biophysical research and pharmaceutical industry. **Next employer:** TBD.

Antonio Villarreal (Ph.D., physics, University of Pittsburgh). **Hired:** August 2018. **Research area:** physics; cosmology. **Current projects:** 1.) ALCC image simulation campaign for LLST Dark Energy Science Collaboration (DESC); 2.) study of halo assembly bias in Mira-Titan universe simulations and impact on cosmological inference; 3.) development of halo modeling tools for HPC, including development of algorithms running on GPU platforms. **Accomplishments include**: two publications in progress; development of workflow pipeline for DESC image simulations and development of GPU algorithms for halo model evaluations.

Section 5. Risk Management

Is the Facility effectively managing operational risks?

ALCF Response

The ALCF has clearly demonstrated successful management of both its project risks and operational risks (Section 5.1) in CY 2019. As part of the ALCF's Risk Management Plan (RMP), all risks (proposed, open, and retired) are tracked, along with their triggers and mitigations (proposed, in progress, and completed), in a risk register managed by risk managers. All risk ratings in this report are post-mitigation ratings. The ALCF currently has **38** open risks, with **two** high operational risks: (1) Funding/Budget Shortfalls, which is managed by careful planning with the DOE program office and the implementation of austerity measures as necessary; and (2) Staff Recruitment Challenges, which is managed by ongoing recruiting and re-tasking of current staff as needed. The major risks tracked for the past year are listed in Section 5.2, along with the details of these risks in Table 5.1. The risks that occurred and the mitigations for those risks are described in greater detail in Section 5.3. Section 5.6 and Table 5.2 provide details on the major risks that will be tracked in CY 2020.

Of primary interest here is a description of the most significant operational risks and the Risk Management Plan's effect on the Facility's day-to-day operations.

The Facility should provide:

- *A brief overview of the risk management process employed by the Facility, including the cycle for identifying, mitigating, and retiring risks;*
- *A brief summary of the key risks and their mitigations, including:*
    - *The 3–5 most important operational risks for the review year;*
    - *Any significant new operational risks since the previous review;*
    - *The operational risks that were retired during the review year;*
    - *The major risks that will be tracked in the next year; and*
    - *For the risk events that occurred, how the Risk Management Plan was implemented and an assessment of the mitigations' effectiveness.*

5.1 Risk Management Process Overview

The ALCF uses the documented risk management processes, first implemented in June 2006 and outlined in its RMP, for both operational and project risk management. This RMP is a strategic plan that is annually reviewed, and updated as needed throughout the year, to reflect changes, to incorporate new risk management techniques as they are adopted, and to incorporate best practices from other facilities. Risk management is part of ALCF's culture, and RMP processes are part of normal operations and all projects, such as the ALCF-3 project launched in CY 2013.

Risk management is an iterative process that includes identifying and analyzing risks, performing response planning, and monitoring and controlling risks as shown in Figure 5.1.

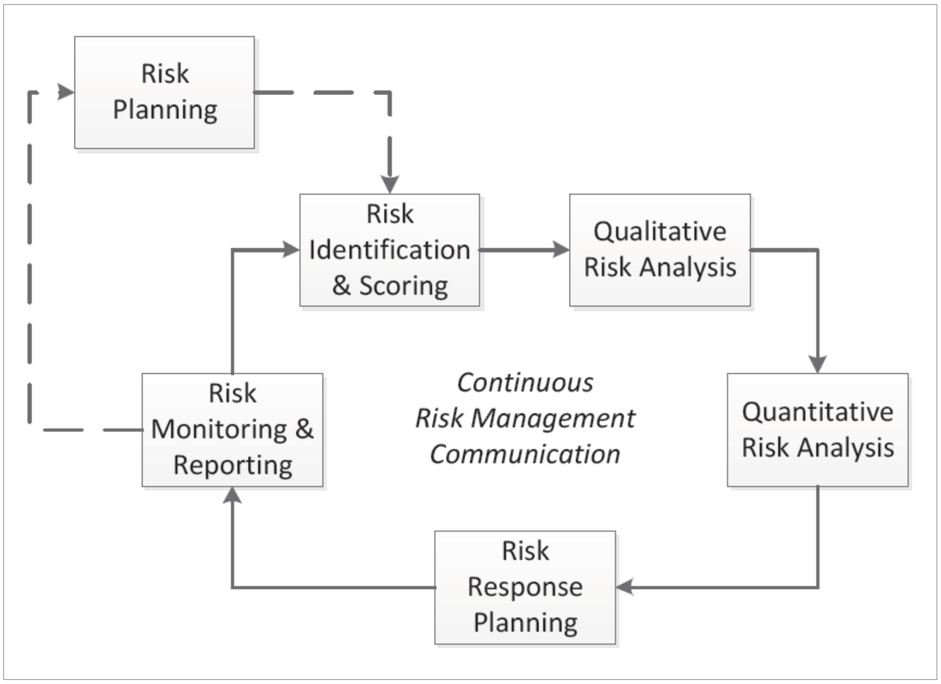

Figure 5.1 Illustration of the Iterative Risk Management Process from the ALCF's Risk Management Plan

The ALCF risk management process consists of the following steps, which are performed on a continuous basis in all normal operations and in all ALCF projects:

1. Plan, implement, and revise the RMP.
2. Identify threats and opportunities to the cost, schedule, and technical objectives.
3. Analyze the impact of identified threats and opportunities to the cost, schedule, and technical baselines; and develop risk management strategies to manage the risks.
4. Monitor risks, mitigation plans, and management reserve and contingency until the risks are retired or a project is closed.

The objective of this process is to identify potential threats and opportunities as early as possible so that the most critical risks can be assessed, the triggers effectively monitored, and the amount of management reserve/contingency needed to moderate the risks determined.

Risks are tracked using a secure and shared cloud-based storage system, and risk forms and the risk register are formatted using Excel. Risk owners continuously monitor the risks they own and submit monthly reports on all risks through the ALCF online risk reporting form.

5.1.1 Risk Review Board

The ALCF employs a five-person Risk Review Board with representatives from senior management, the operations team, the science team, industry outreach, and the financial services team to serve in an advisory capacity to ALCF management. The board meets as needed and offers recommendations regarding steady-state risk management issues. The RMP is consulted at all risk meetings. At each meeting, the board:

- Reviews proposed new risks and makes recommendations on adding a proposed risk to the steady-state 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 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 strategies for the risk and considers whether any of the strategies need updating for any reason, including as a result of changes in the technology landscape.
  - Works with the risk owner to modify the risk statement should any risk information indicate a need for changes to risk mitigation strategies, risk triggers, or risk scope.
- Reviews and identifies any risks to retire.
- Reviews the risks encountered in the past 18 months to discuss potential actions.
- Discusses risks encountered at other facilities and identifies any that apply to the ALCF.

5.1.2 Risk Management in Day-to-Day Operations

ALCF currently has **38** 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 and individual risk meetings and the Risk Review Board meetings, the ALCF holds many informal risk discussions. Risks are identified and evaluated, and mitigation actions are developed, for all changes that occur 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 developed in consultation with safety and subject matter experts. JHQs are used for all staff and all contractors and cover all work. During planning meetings for any activities, staff members 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 and then documented in the WPC. The WPC is then used during the activity to direct the work.

Beyond the operations of the machines, risk management is used in such diverse ways as in 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.

5.1.3 Continuation of the ALCF-3 Project

The project to procure and deploy the ALCF's next supercomputer, known as ALCF-3, continued in CY 2019. Risk Register managers continue to maintain a project risk register and track a set of detailed risks. 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.

5.2 Major Risks Tracked for the Review Year

The ALCF has been operating both Mira and Theta during CY 2019, and planned the growth of both the staff and the budget in order to bring the facility to full strength. As such, the ALCF continues to monitor a large number of major risks for the facility. No major risks were retired during CY 2019.

Four major operations risks were tracked for CY 2019, **two** with a risk rating of High, **one** with a risk rating of Moderate, and **one** with a risk rating of Low. None of these were encountered and all of them were managed. The **four** major operational risks are described in Table 5.1. All risk ratings shown are post-mitigation ratings. The risks are color-coded as follows:

- Red risks are Moderate or High risks.
- Orange risks are Low risks.

Table 5.1 Major Risks Tracked for CY 2019

ID	Title	Encountered	Rating	Notes
1059	Funding/Budget Shortfalls	No	High	The ALCF regularly worked with the program office to plan a budget for handling the impact of a Continuing Resolution in FY 2019, new hires, and changes in the laboratory indirect expense rate. This risk remains a major concern as the facility moves forward with Theta and ALCF-3 in CY 2020.
25	Staff Recruitment Challenges	No	High	The ALCF added 6 new hires overall this year, plus 2 internal reclassifications. Four of the 6 new hires are funded by ECP. The ALCF continues to have staff available who can be re-tasked as needed. With difficulty competing with industry for new hires, staff hiring remains a concern.
1049	Staff Retention	No	Moderate	The ALCF lost **4** staff members during CY 2019, 2 of whom were funded by ECP. Budget concerns at Argonne and the growth in high-paying industry jobs for system administrators and programmers with HPC expertise make staff retention in out years a continuing concern.
1091	Injury to Workers/Overall Safety of the Division	No	Low	

5.3 Risks Encountered in the Review Year and Their Mitigations

The ALCF encountered **one** risk in CY 2019. The risk owner is identified below, along with an assessment of the risk's probability and impacts, a brief description of what transpired, and how the risk was ultimately managed. This risk has a rating of Low.

5.3.1 Facility Power Interruptions

0031: Facility Power Interruptions	
Risk Owner	Mark Fahey
Probability	Low
Impact	Cost: Low; Technical Scope: Low
Risk Rating	Low
Primary Management Strategies	The ALCF's director of operations participates in the Data Center management group. The ALCF pays part of the cost of an Argonne Data Center liaison. Improve power system bus transfer by investigating the practicality of modifications to the power system to weather bus transfers without interruption.
Triggers	Electrical failure; multiple events related to power quality; scheduled power outages.

Description

One of the building's two UPS units briefly failed due to an over-temperature condition caused during maintenance of a rooftop cooling unit. All mira-fs and mira-home filesystems, as well as Mira's management infrastructure, briefly lost power as they are fed only from the UPS units.

Evaluation

Typically, when a UPS unit fails, it should fail to bypass so that power isn't interrupted. However, that did not occur, and the load shifted to the other UPS unit, which caused an overdraw on that unit.

Management

The UPS was manually bypassed while it was investigated and brought back to service.

5.4 Retired Risks

There were no risks retired in CY 2019.

5.5 New and Recharacterized Risks since the Last Review

There are no new risks and one recharacterized risk to report since the last review. The details of the recharacterized risk appear below, including the risk owner, an assessment of the risk's probability and impacts, a description of the changes, and what management strategies have been developed to address the risk. This risk has a rating of Low.

5.5.1 The Facility Where ALCF Stores Critical System Backup Resources Becomes Unavailable

1076: The Facility Where ALCF Stores Critical System Backup Resources Becomes Unavailable	
Risk Owner	Mark Fahey
Probability	Very Low
Impact	Cost: Low; Technical Scope: Very Low
Risk Rating	Low
Primary Management Strategies	Explore alternative locations onsite. Explore alternative locations offsite (e.g., another lab to host ALCF disaster recovery resources). Accept the risk of a major loss of service or data. Raise concerns with Argonne upper management. Develop plan to move equipment.
Triggers	Information from Argonne management that current location must be vacated.

The risk was originally titled: If the ISSF [Interim Supercomputer Support Facility] is decommissioned, ALCF will not have an appropriate facility to host disaster recovery resources.

The risk was encountered when it was triggered, that is, Argonne management set the date for decommissioning the ISSF. With careful planning, the ALCF vacated the facility. However, the risk was re-characterized to cover any facility where ALCF stores critical system backup resources becoming unavailable.

5.6 Projected Major Operating Risks for the Next Year

Table 5.2 lists the current top operating risks projected for CY 2020, along with the current risk rating and management strategies for each risk. These are the risks that experience has shown are most likely to be encountered in any fiscal year.

Table 5.2 Projected Operating Risks for CY 2020

ID	Title	Rating	Management Strategies
1059	Funding/Budget Shortfalls	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	Staff Recruitment Challenges	High	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	Moderate	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	Low	Promote safety culture at all levels of the division. Follow Argonne Integrated Safety Management Plan. Monitor work areas for potential safety concerns. Enforce use of personal protective equipment.

Conclusion

The 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 the ALCF's culture and applies equally to all staff, from senior management to summer students. A formal risk assessment is performed for every major activity within the ALCF, with informal assessments used for smaller activities. Risks are monitored and tracked using a secure and shared cloud-based storage system, along with risk forms and a risk register that are both formatted using Excel. Beyond these activities, many tools are used to manage risks at the ALCF, particularly in the area of safety. The ALCF's effective risk management plan has contributed to the successful management of all significant risks encountered in the past year.

Section 6. Environment, Safety, and Health

(a) Does the Facility exhibit a culture of continual improvement in Environment, Safety, and Health (ES&H) practices to benefit staff, users, the public, and the environment?

(b) Has the Facility implemented appropriate Environment, Safety, and Health measures?

ALCF Response

The ALCF has an exemplary safety record and maintains a safe and healthy workplace. To date, the facility has not experienced a single lost time incident. In 2019, ALCF experienced zero recordable injuries, zero near misses, one first aid incident, and zero pause/stop work events. The division's commitment to safety also extends to office work, with 13 ergonomic evaluations performed on ALCF employees in 2019.

The ALCF produced a local work planning and control manual for executing work according to Integrated Safety Management guiding principles and various approaches from DOE-HDBK-1211-2014: *Activity-Level Work Planning and Control Implementation*. The division uses Aware, the laboratory's work planning and control online application. The ALCF analyzes hazard and control measures to ensure that work is performed while mitigating risks to the safety, environment, mission, security, and health of the users, public, and workers.

A formal skill-of-the-worker document is in place and has been implemented to cover routine low-rigor system administrator tasks. (Skill-of-the-Worker refers to laboratory-wide Argonne Procedure LMS-PROC-65, Rev. 7 that establishes the process for (1) identifying work that can be performed by a worker with an appropriate level of proficiency to independently identify hazards, implement controls, and safely complete work tasks; and (2) identifying such workers. Formal task-based work control documents 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. Management actively monitors work scope and updates hazard analyses and revises work planning and control documents for emergency work or whenever there is an unexpected change to previously planned work. Worker health and safety documentation that cannot be housed within Aware is retained either within the Laboratory's Training Management System or the Laboratory's official cloud resource, Box.

The ALCF practices electrical safety in accordance with the Laboratory's electrical safety program, which encompasses electrical hazards and electrical safe work practices. All electrical work within ALCF is planned with an Argonne electrical safety Subject Matter Expert (SME) and performed by a Laboratory Qualified Electrical Worker (QEW). All work control documents currently in use for electrical work have been reviewed and walked-through by a Laboratory electrical SME. Laboratory electrical SMEs are involved early in the work-planning process and consulted for review before large work projects begin, such as the three-day electrical shutdown in September 2019. The ALCF conducts pre and post work planning meetings with all involved workers in accordance with their work planning manual to continually improve work by learning from what went well or what needed improvement.

ALCF uses formal management assessments, such as the 2019 Personal Protective Equipment (PPE) Use in the Data Center Assessment and the 2019 Lock-out/Tag-out (LOTO) Assessment, to seek continuous improvement within the division. The 2019 PPE assessment led to a hands-on training session with the Laboratory Noise SME on the proper use of hearing protection for the operations team. The 2019 LOTO assessment led to an additional voluntary training session for two ALCF QEWs with the Laboratory electrical SME to review switching techniques. In addition to formal management assessments, ALCF's division director and ESH coordinator conduct biannual health and safety inspections. These inspections serve as an opportunity to observe work, seek worker feedback, and identify areas for improvement. In 2019, these inspections resulted in changes to the data center signage pertaining to noise.

Section 7. Security

(a) Does the Facility exhibit a culture of continual improvement in cyber security practices?

(b) Does the Facility have a valid cyber security plan and Authority to Operate?

(c) Does the Facility have effective processes for compliance with applicable national security policies related to Export Controls and foreign visitor access?

ALCF Response

The ALCF works to continually improve its cyber security practices by developing and maintaining relationships between facility personnel and Argonne's internal security personnel, and between the ALCF and the Argonne Cybersecurity Program Office (CSPO).

7.1 Continual Improvement in Cyber Security Practices

The ALCF also works to maintain a strong cyber security posture. In 2019, the ALCF deployed a new user management application, called UserBase3 (UB3). The application, which was developed in-house, relies on a modern web-application stack that supports rapid development and enhanced usability and security. This application leverages object-relational mapping, a well-trusted authentication framework, to prevent SQL injection attacks, and standard HTTP security headers to help protect users' information. The ALCF application development team worked in coordination with cyber security personnel throughout the development cycle and routinely included ALCF's Cyber Security Program Representative on pull requests to provide additional input on the code's security features. In addition, several people from the CSPO agreed to assess the security of the application.

In CY 2019, there were zero cyber security incidents on ALCF-managed systems. ALCF's cyber security personnel take a proactive approach to problem management. Examples of proactive measures taken by our security personnel include the following:

1. Performed privileged access reviews across the environment to make sure no one has access that they should not have.
2. Performed reviews of how and where data is stored to make sure it is only accessible by those with the proper authorization.
3. Educated users and staff on how to store and use passwords in applications to prevent accidental exposure.
4. Educated developers on secure coding best practices via internal discussions/reviews and external training courses.
5. Updated the National Institute of Standards and Technology (NIST) Certification package including NIST 800-53, 800-34, 800-30, and 800-18 compliance documents.
6. Deployed Duo as the new multifactor authentication method in the JLSE.
7. Archived and deleted data that is no longer needed.

8. Verified that passwords are being rotated on a regular basis, and codified policies for when rotating passwords is required on ALCF systems.
9. Monitored new vulnerabilities and evaluated whether they affect ALCF systems.
10. Performed penetration testing of internal- and external-facing web applications to evaluate what improvements can be made and how we can better protect our users and data.
11. Developed tools to verify permissions on home directories to follow ALCF best practices.
12. Deployed Elastic Beats agent to collect and store machine logs in Elasticsearch for short-term analysis.
13. Deployed new security scanners and evaluated the integration of those scanners into the CSPO's new security scanning platform.
14. Worked with Argonne's Export Control Office to update documentation and practices to allow specific types of low-sensitivity Export Control data onto ALCF systems.

Some of these activities revealed security vulnerabilities that were promptly addressed and fixed, usually within days of initial discovery. Immediately upon detection, ALCF staff would also investigate all relevant logs to determine whether the security vulnerability had been exploited. In CY 2019, none of the issues investigated were found to have been exploited. Examples of the security issues that were detected and their ensuing mitigations are as follows:

1. **Issue:** A large user base was given access to a documentation space that contained sensitive information. **Mitigation:** Authorization was corrected and the access and authentication logs were analyzed. It was verified that only authorized individuals had accessed the data.
2. **Issue:** SVN repositories hosted by another Argonne division were exposing ALCF SVN repositories. **Mitigation:** The service was immediately taken offline until appropriate permissions settings were restored and all repositories were properly secured. Old and unneeded repositories were also identified and removed.
3. **Issue:** Vendor support staff were found to be performing administrative actions on ALCF resources without using Homeland Security Presidential Directive 12 (HSPD-12) badges to authenticate. **Mitigation:** All vendor support staff that have regular and continued access are required to use, and are in the process of obtaining, HSPD-12 badges to be able to authenticate with an Identity Assurance Level (IAL) 3 Authenticator Assurance Level (AAL) 3 Multifactor Authentication (MFA) (NIST 800-63-3). Some vendor support staff require access on an ad-hoc, temporary basis (e.g., to cover a regular vendor staff member who is on extended leave or to help with a specific problem). For these cases, exceptions are made and documented and alternative strong MFA is still required.
4. **Issue:** Passwords in some applications were found to have been stored insecurely. **Mitigation:** Subsequent evaluation of the application data integrity determined that no unauthorized access to the associated databases had occurred. Passwords were also rotated as a precaution against potential exploit.

5. **Issue:** Security risks were identified in pre-production applications under development. **Mitigation:** The development instances were locked down while the developers addressed the issues. The fixes were verified by cyber security personnel prior to application deployment.

The ALCF will continue to proactively investigate security issues and to monitor and respond to all vulnerabilities. Plans for improving the security of ALCF resources include the following:

1. Retiring services and data that are no longer needed.
2. Verifying that strong encryption is used everywhere in the environment and that plain text protocols are not used for production needs.
3. Improving real-time log analysis techniques.

The CSPO performs an annual internal assessment with all Divisions called the Division Site Assist Visit (DSAV). The DSAV is an audit of the Divisions' compliance with NIST-800-53 controls. Each year the DSAV assessment covers roughly one third of the controls. In 2019, the DSAV targeted 66 controls from 13 of the 18 NIST control families for assessment—16 of these target controls were not applicable to ALCF. A maturity score from 0-5 is assigned to each control assessed. Table 7.1 lists ALCF results from the 2019 DSAV. The ALCF continues to have a very mature and predictable environment with over 75 percent of the targeted controls assessed fully implemented in a way that is uniform. In 2019 the CSPO identified 3 opportunities for improvement that the ALCF will address during CY2020:

1. Update our disk shredding exemption.
2. Automatically deactivate ALCF staff accounts when their relationship with the laboratory and the facility is terminated.
3. Verify with CSPO that they continue to receive all logs that they require.

The CSPO also complimented the ALCF on the following areas for their strength and maturity:

1. The ALCF is one of the leaders at the lab in contingency planning.
2. The ALCF has excellent account management processes. The processes are well understood and documented and tightly enforces ALCF policies through technical controls.
3. ALCF does well at distinguishing guests/regular users/admins and maintains the separation of user privileges in their environment.
4. The ALCF is one of the leading divisions at the lab in leveraging HSPD-12 badges as authentication for system administrators and is forward thinking in the use of multi-factor technologies to protect systems and data.
5. Software development processes use Continuous Integration/Continuous Deployment (CI/CD) practices. Using automated CI/CD practices to help minimize deployment mistakes and do unit testing for Quality Assurance (QA) is of particular note.

Table 7.1 2019 DSAV Maturity Scores

Maturity Score	Number of Controls	Description
0	0	Unable to Assess
1	0	Planned Implementation for Control
2	5	Ad-Hoc Implementation of Control
3	7	Managed Implementation of Control
4	4	Established Implementation of Control
5	34	Predictable Implementation of Control

7.2 Cyber Security Plan

The Argonne Authority to Operate (ATO) includes the ALCF as a major application and was granted on January 22, 2018. 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 is included at the end of this section.

7.3 Foreign Visitor Access and Export Controls

The ALCF follows all national security policies and guidelines related to Export Controls and foreign visitor access.

Argonne is a controlled access facility, and anyone entering the site or accessing the resources remotely must be authorized. The ALCF follows Argonne procedures for collecting information about foreign nationals who require site access or remote (only) computer access. All foreign nationals are required to have an active and approved ANL-593 in order to have an active ALCF account. Users can access ALCF resources only with an active ALCF account.

To apply for an ALCF account, the user fills out a secure webform in the ALCF Account and Project Management system (UB3) where they provide details such as legal name, valid e-mail address, work address, phone number and country of citizenship. They also identify the ALCF project they are associated with. In addition, all foreign nationals (non-U.S. citizens) are required to fill out their personal, employer, demographic, and immigration/USCIS information in Argonne's Visitor Registration system, which is integrated with UB3. After the user submits their account application request, an e-mail is sent to the user's project PI for their approval. Once the ALCF Accounts team receives the approval from the project PI, if the user is a foreign national, the user's details are electronically attached to an ANL-593 form and submitted to Foreign Visits and Assignments (FV&A) office for review. *The FV&A Office is responsible for overseeing compliance within the laboratory and ensuring compliance with the DOE.* The ANL-593 form records the type of work the user will be performing, including the sensitivity of the data used and generated. The ANL-593 has to be approved by Argonne Cyber Security, FV&A, the Argonne Office of Counterintelligence, and the Argonne Export Control Office. Argonne's foreign visitor and assignments process integrates with the DOE Foreign Access Central Tracking System (FACTS), which documents and tracks access control records of international visits, assignments, and employment at DOE facilities and contractor sites. Once the ANL-593 form for the user is approved, UB3 is automatically updated with the user's ANL-593 start and end dates. The ALCF Accounts team then creates the user account and notifies the user. Any changes to the ANL-593 dates are automatically updated in UB3. Accounts are suspended if the user does not have an active 593.

The ALCF allows only a limited subset of Export Control data on our systems. The ALCF works closely with Argonne's Export Control Office to complete a detailed security plan for what Export Control classifications are allowed and what security measurements are required for each instance of Export-Controlled data. If, at any time, the ALCF would like to allow new classifications of Export Control data on its systems, a new security plan must be created and approved by Argonne's Export Control Office and Argonne Cyber Security.

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

JAN 2 2 2018

Dr. Paul K. Kearns
Director, Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Kearns:

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

Reference: Letter, J. Livengood to P. Littlewood, dated November 21, 2016, Subject: Authority to Operate for the Argonne National Laboratory Information Technology Infrastructure

Over the past year, 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. The Laboratory has revised system security documentation to incorporate NIST SP800-53 Revision 4 security controls and has been testing at least 60 security controls annually on a rotating basis as part of the self-assessment program. 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 the Authority to Operate (ATO) for the General Computing – Low enclave and the General Computing – Moderate enclave. The IT Infrastructure continues to contain the following sub-component major applications, which have components in both enclaves:

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

This ATO 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.

A component of the Office of Science

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

Joanna M. Livengood
Manager

cc: S. Hannay, ANL-BIS
M. Skwarek, ANL-BIS
M. Kwiatkowski, ANL-BIS
B. Helland, SC-21
N. Masincupp, SC-OR
F. Healy, SC-CH

This page intentionally left blank.

Section 8. Strategic Results

(a) Are the methods and processes for monitoring scientific accomplishments effective?

(b) Has the Facility demonstrated effective engagements with strategic stakeholders (i.e., beyond the user population)?

(c) Is the Facility operating in a manner that enables delivery of facility mission and Department of Energy mission including maintaining a vibrant US effort in science and engineering?

ALCF Response

The science accomplishments of INCITE, ALCC, and DD projects clearly demonstrate ALCF's impact in supporting scientific breakthroughs. ALCF staff members have worked effectively with individual project teams to adapt their simulation codes to run efficiently in an HPC environment and have enabled scientific achievements that would not have been possible otherwise.

In this section, ALCF reports:

- Scientific Accomplishments;
- Scientific Highlights; and
- Stakeholder Engagement.

8.1 Scientific Accomplishments

The ALCF employs various methods and processes for monitoring its science accomplishments. Monthly scientific highlights mostly originate from the catalyst team and are based on an outcome documented in a quarterly report. The determination and coordination of scientific highlights is performed by ALCF's applications team, made up of members of both the catalyst team and the performance engineering team, and in consultation with ALCF's director of science. Other sources of scientific highlights include technical communications between ALCF staff members and a project PI or co-PI; significant findings reported in a high-impact publication or conference presentation; and a catalyst's own involvement in a publication.

The Facility also annually tracks and reports the number of peer-reviewed publications resulting (in whole or in part) from use of the Facility's resources. For the LCFs, tracking takes place for a period of five years following the project's use of the Facility. This number may include publications in press or accepted but does not include papers submitted or in preparation. This is a reported number, not a metric. In addition, the Facility may report other publications where appropriate. Methods used for gathering publication data include asking users to verify or update ALCF's online publications database, and conducting Google Scholar and Crossref searches. The ALCF also collects approximately one-third of its users' ORCID iDs in any given year and has been investigating ways to use this method to collect more user publication data.

Table 8.1 shows the breakdown of refereed publications based, in whole or in part, on the use of ALCF resources, and highlights those appearing in major journals and proceedings. These include six publications in *Nature/Springer Nature*, one in *Nature Chemistry*, one in *Nature Reviews Physics*, one in *Nature Physics*, and three in *Nature Communications* (combined in the *Nature* journals category in the table below); the journal *Science*, *Proceedings of the National Academy of Sciences (PNAS)*, *Physical Review Letters*, and the proceedings of the *2019 International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. Table 8.2 shows updated publication counts from prior years and are based on new information after the prior year's OAR deadline.

Table 8.1 Summary of Users' Peer-Reviewed Publications in CY 2019

Nature Journals	Science	PNAS	Physical Review Letters	SC	Total 2019 Publications
12	1	2	8	4	284

Table 8.2 Summary of Users' Peer-Reviewed Publications for 5-year Moving Window

OAR Year	CY 2015	CY 2016	CY 2017	CY 2018	CY 2019
Total Publications	181	199	225	276	284

8.2 Scientific Highlights

Scientific highlights are short narratives that illustrate the user facility's contribution to advancing DOE strategic goals. Highlights may describe a research accomplishment or significant finding from either a current project or from a project originating in a previous year, as data analysis may occur several months after the computational campaign has been completed.

Each project highlight includes a figure and a bar graph showing time allocated and time used. The first number in the graph title is the allocation total and the second is what the project used. The individual bars represent the percentage of time used on the fraction of the machine shown below the bar, which are "no capability," "low capability," and "high capability" from left to right.

8.2.1 Creating Proteins That Mimic DNA

The Science

Protein design is a massive combinatorial problem that requires consideration of how amino acids interact with each other resulting in 3D protein structure. The even more complicated task is to use protein design to find proteins that perfectly complement each other. Using advanced computational methods to find working designs, researchers successfully created proteins that mimic deoxyribonucleic acid (DNA).

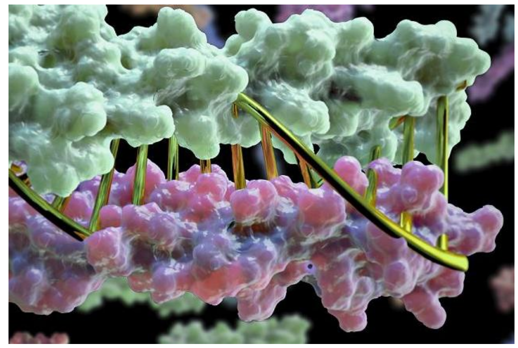

Figure 8.1 **A designed coiled-coil heterodimer, with halves colored green and purple. A DNA double-helix, for which particular sequences have the same property of forming specific, orthogonal heterodimers, is superimposed to show scale.**

Image courtesy of Institute for Protein Design.

The Impact

Engineered pairs of proteins that bind exclusively to one another would enable sophisticated cellular control logic in living systems for bioengineering applications—with potentially large impacts for medicine and biomaterial production. This work provides a path forward for researchers to computationally design specific, programmable binding, previously a property found only in the DNA and ribonucleic acid (RNA) world, into proteins.

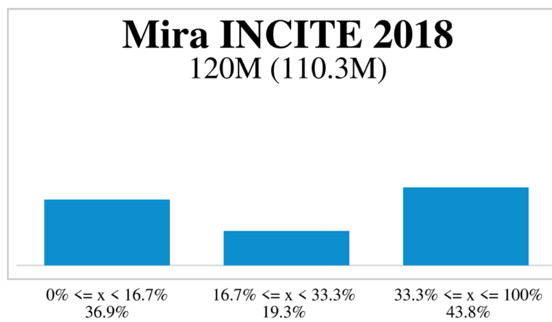

Summary

DNA is a widely used building material at the nanoscale because it is simple and predictable: A pairs with T and C pairs with G. Because of this, DNA strands can be programmed to click together into precise and increasingly complex structures. But DNA has drawbacks. It is not as bioactive as RNA, and not nearly as active as proteins. Bioactive *protein* assemblies run cells (kinetochores, polymerases, proteasomes, etc.). What if designing them was as easy as clicking together DNA?

Using computational design, a multi-institutional team of researchers led by scientists at the University of Washington created heterodimeric proteins that form double helices with hydrogen-bond-mediated specificity (Figure 8.1). When a pool of these new protein zippers gets melted and then allowed to refold, only the proper pairings form. They are all-against-all orthogonal. With these new tools in hand, the team can now begin constructing large protein-based machines that self-assemble in predictable ways.

ALCF Contribution: ALCF staff provided guidance for improving Rosetta with efficient MPI parallelization and multi-threaded execution.

Contact

David Baker
University of Washington
dabaker@uw.edu

Publications

Zibo Chen et al. "Programmable Design of Orthogonal Protein Heterodimers," *Nature* **565**, 106–111 (2019). https://doi.org/10.1038/s41586-018-0802-y.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

August 2019

8.2.2 Shedding Light on Molecules' Structure and Motion in Cells

The Science

To understand why very-large molecules behave the way they do inside cells, scientists must first understand the relationship between these molecules' structure and motion. Engineers created algorithms that provide the physics backbone for a new "microscope in a computer." This tool produces high-resolution microscopic images (Figure 8.2). It also made possible the development of a new technique, partial wave spectroscopic (PWS) microscopy. Scientists can use this technique to detect structures and the movement of molecules in living cells at scales smaller than what should be physically possible with a specific microscope.

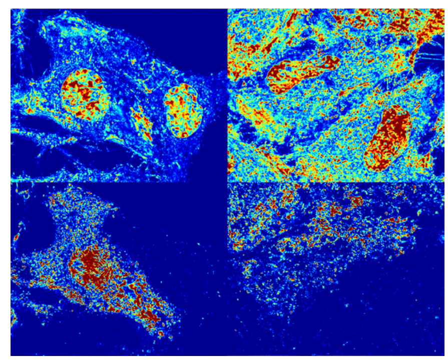

Figure 8.2 Representative mass maps are shown both before (top) and after (bottom) cellular paroxysm for cytoskeletal (left) and membrane (right) disruption.

Image reproduced from Figure 7 of *Nat. Commun.* **10**, 1652 (2019).

The Impact

Detecting cancer at an early stage is one of the most critical factors to successfully treating it and preventing death. Developing a low-cost, high-throughput microscopy technique that can sense changes in large molecules in living cells is a key step toward this goal. Using PWS's high resolution with the new algorithm opens the door for more accurate and more efficient early-stage cancer screenings. Now it may be possible to detect seven different human cancers with minimal false positives: lung, colon, ovarian, esophageal, pancreatic, thyroid, and prostate.

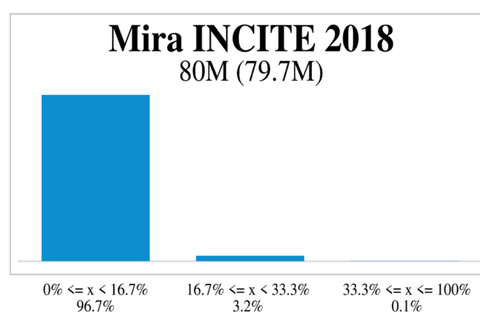

Summary

The finite time-difference time-domain code, Angora, is a computational microscopy code that models the four components of an optical imaging system (illumination, scattering, collection, and refocusing). It then synthesizes microscope images to validate theoretical results and predict scattering behaviors. Simulations with Angora were critical in this project to validate the analytic theory developed to extract physical properties from the measured spectral interference. The team then applied PWS to explore the nanoscale structural and dynamic changes of stem cells. Experiments probed the spatiotemporal response of macromolecular assemblies to ultraviolet (UV) light. Scientists discovered a near instantaneous burst of motion that occurs early in the process of UV-induced cell death. This process called cellular paroxysm was found to occur synchronously within a cell, as far as 30 micrometers apart in under 35 milliseconds, but asynchronously between cells. The team notes that this effect may play a role in the earliest stages of a unique form of cell death. However, the exact mechanism of cell death is unclear. The team also developed a machine learning model to classify cells based solely on physical properties of the nucleus with an accuracy of 80% and indicated that potential improvements could come by way of larger datasets.

ALCF Contribution: A member of ALCF's Catalyst team helped the project team to fix parallel bugs related to Mira's architecture and to improve code performance with efficient parallel I/O.

Contact

Vadim Backman
Northwestern University
v-backman@northwestern.edu

ASCR Allocation PI: Allen Taflove, Northwestern University

Publications

Scott Gladstein et al. "Multimodal Interference-based Imaging of Nanoscale Structure and Macromolecular Motion Uncovers UV-induced Cellular Paroxysm," *Nat. Commun.* **10**, 1652 (2019). https://doi.org/10.1038/s41467-019-09717-6.

James A. Winkelmann et al. "Spectral Contrast Optical Coherence Tomography Angiography Enables Single-Scan Vessel Imaging," *Light Sci. Appl.* **8**, 7 (2019). https://doi.org/10.1038/s41377-018-0117-7.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

June 2019

8.2.3 Impact of Grain Boundary Defects on Hybrid Perovskite Solar Absorbers

The Science

Hybrid organic-inorganic perovskites (HOIPs) are an emergent technology that rely on low-cost materials and simple manufacturing processes. Nevertheless, under regular operational conditions, these cells could be exposed to a wide range of temperatures, causing repeated contractions and expansions of the HOIP structure and promoting the irreversible degradation of the solar cells. Recently, experimental studies also confirmed that the energy conversion efficiency decreases linearly with the temperature increase. Some low-dimension models fail to explain temperature effects in solar cells in the low- or high-temperature extremes being off by almost 3x.

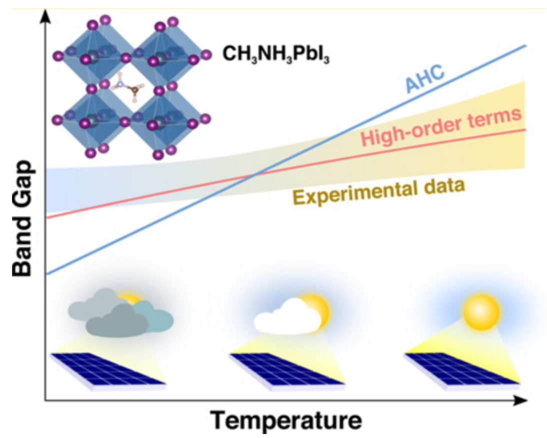

Figure 8.3 Schematic plot of how the ability of a hybrid perovskite material, $CH_3NH_3PbI_3$, to absorb energy (band gap) changes with temperature experimentally and calculated with two levels of approximation.

Image courtesy of Wissam Saidi, University of Pittsburgh.

The Impact

The design of new and reliable solar panels may help many to transition from fossil fuel to solar energy resources. Increasing the durability of solar cell materials is critical as their temperature dependency has profound implications in the lifespan of solar panels. This project contributes to advancing the understanding of solar cell efficiency degradation and highlights the importance of electron-phonon coupling on the properties of the hybrid perovskite (Figure 8.3).

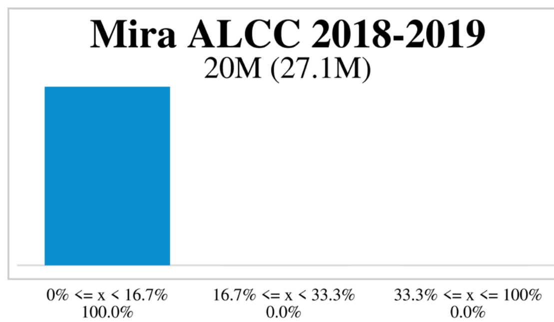

Summary

In this project, researchers studied changes in electronic properties of HOIP materials with respect to temperature using electronic structure calculations that included high-order electron-phonon terms, capturing correctly the physics over a wide range of temperatures. The results of this project are in excellent agreement with experimental evidence, a remarkable contrast with previous estimations that have been off by almost 3x.

ALCF Contribution: Members of ALCF's Catalyst team helped the project team to prepare ensemble jobs and facilitated access to the latest binaries of the FHI-aims and CP2K codes on Mira.

Contact

Wissam Saidi
University of Pittsburgh
alsaidi@pitt.edu

Publication

Wissam A. Saidi and Ali Kachmar, "Effects of Electron-Phonon Coupling on Electronic Properties of Methylammonium Lead Iodide Perovskites," *J. Phys. Chem. Lett*. **9**, 24, 7090–7097 (2018). https://doi.org/10.1021/acs.jpclett.8b03164.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

September 2019

8.2.4 Accelerating the Development of Commercially Viable Nuclear Fusion

The Science

Fusion energy offers the prospect of a safe, carbon-neutral, and sustainable energy source. Researchers from TAE Technologies used the ALCF to support their fusion research. The company is working to develop the world's first fusion device that can generate electricity and is commercially viable. TAE is using a unique machine to research fusion called a field-reversed configuration (FRC) device. This study used computer models to understand how plasma loses heat in this type of device.

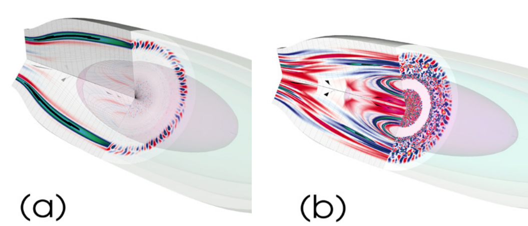

Figure 8.4 Snapshots in time of the simulation of turbulence in fusion plasma confined in an advanced field-reversed configuration (FRC).

Image courtesy of Calvin Lau, TAE Technologies, Inc.

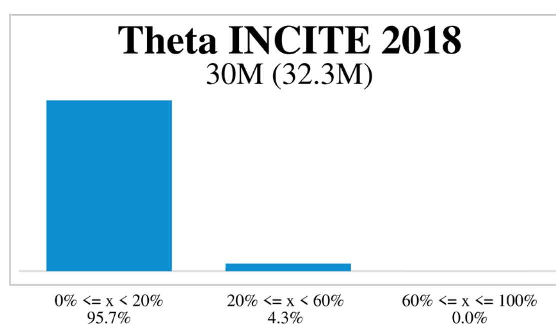

The Impact

Starting with the foundational ideas is helping this team to answer fundamental questions of both science and technology. Why does the FRC device's ability to contain the plasma's energy improve as plasma gets hotter (Figure 8.4)? What is the best configuration to produce a very high-pressure fusion plasma? Answering these questions will accelerate progress toward TAE Technologies' goal of designing a commercially viable fusion-based power plant driven by an aneutronic hydrogen-boron (p-B11) fuel cycle. Developing such a machine could help transform the international energy market.

Summary

To understand energy confinement in FRCs, the research team is performing first principles simulations of turbulent heat losses in FRCs using the ANC kinetic micro-turbulence code. In 2018, the team used Theta to perform and publish a systematic study of global transport due to electrostatic driftwave turbulence in the exterior and interior of the FRC. It was found that short wavelength modes are linearly unstable outside of the FRC. In the nonlinear case, an archetypal inverse turbulent cascade occurs, shifting energy from short to long wavelengths. No source of instability was found in the interior of the FRC itself. These simulations were validated against previous experimental DBS (Doppler backscattering) observations and were found to agree well. In 2019, these studies were extended to compare against coming experimental DBS measurements on the Norman experiment. Theta is also being used to study global configuration stability of the FRC using the FPIC code developed at TAE Technologies. Here, the effects of external actuators, including neutral beam injection, electrode biasing, and magnetic field shaping, are included to simulate the effects on plasma self-organization. These results will be used to inform the operating states and feedback and control strategies that are being developed for Norman and for future FRC devices.

ALCF Contribution: Upgrades of both the ANC and FPIC codes from the MPI paradigm to the hybrid OpenMP/MPI paradigm were informed by discussions with ALCF scientists during the 2017 ALCF Performance Workshop, and miscellaneous workflow and optimization improvements were informed by discussions with ALCF scientists during the 2018 ALCF Simulation, Data, and Learning Workshop.

Contact

Sean Dettrick
TAE Technologies, Inc.
sean@tae.com

Publication

Calvin K. Lau et al. "Cross-separatrix Simulations of Turbulent Transport in the Field-reversed Configuration," *Nucl. Fusion* **59**, 6 (May 2019). https://doi.org/10.1088/1741-4326/ab1578.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

April 2019

8.2.5 A New Twist on Controlling Magnetic Properties

The Science

Computer storage devices often use magnetic materials printed on very thin films. In this study, researchers rotated cobalt-iron alloy thin films relative to an applied magnetic field (Figure 8.5). Unexpectedly, depending on the rotation angle, a sizeable change—up to 400 percent—was seen in how well the material holds on to energy (referred to as “damping anisotropy”). The key difference appeared to be that these materials are usually heated to high temperatures in making computer disk drives, resulting in a well-ordered atomic structure. In this research, unbaked alloys with more random arrangements of the atoms were used. To explain the physics of the observation, the electronic and magnetic structures were calculated. The calculations showed that the variable separation of the disordered atoms altered the interactions among the spins and thus changed their magnetic effects.

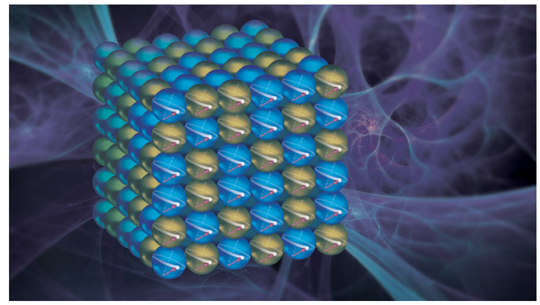

Figure 8.5 Interactions between electron spin and an external magnetic field was found to give rise to a sizeable magnetic damping anisotropy as observed in thin film samples of a cobalt-iron alloy (blue and yellow spheres, respectively), a material commonly found in computer disk drives.

Image courtesy of Argonne National Laboratory.

The Impact

Control of intrinsic magnetic properties in a single device is a long-standing challenge for scientists and engineers. This study highlights a key mechanism that could be used to control a major magnetic property in these thin-film magnetic metals. A device using this technique could be four times more energy efficient than current devices. Such capability is important for exploration of new technologies and improving future devices, such as those used in computer memory, electric motors, generators, and magnetic bearings.

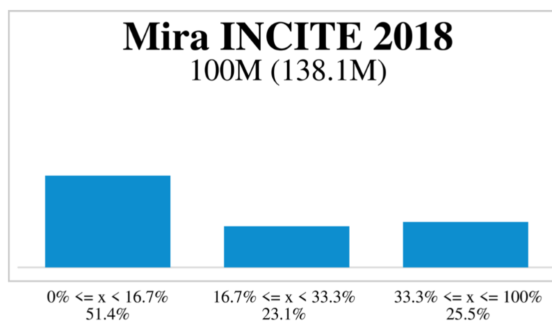

Summary

A phenomenon called “Gilbert damping” explains the relationship between the energy required and speed at which a magnetic device can operate. Higher damping of the energy leads to faster relaxation. The ability to tune and explore magnetic damping in ferromagnetic materials used in magnetic memory and spintronic devices is critical to designing improved devices. Exploring the changes in the intrinsic damping induced by rotating the magnetization orientation along crystalline axes is thus far one of the most promising approaches to study how to design and control the damping for an individual magnetic device. In this research, an international team led by Argonne National Laboratory combined experiment and theory to make, characterize, measure, and calculate physical properties of epitaxially thin Co50Fe50 metal films. Rotating the metallic films in magnetic field caused up to 400 percent variation in the magnetic damping. In order to investigate the origin of this unexpected anisotropy, a series of magnetic, non-collinear DFT calculations of the electronic and magnetic structures were performed by using the Quantum Espresso software available for HPC at the ALCF. Calculations incorporated electron spin-orbit coupling effects for three materials with differently oriented cobalt-iron structures. The researchers found that spin-orbit, coupling-induced changes were dominant and affected the anisotropy of magnetic damping more significantly in disordered alloys. This finding suggests that the origin of the experimental observation of a giant Gilbert damping anisotropy in a metallic ferromagnet can be explained by the variation of the spin-orbit coupling for different magnetization orientations in the cubic lattice. This integrated experimental and theoretical study provides a promising strategy for tailoring the damping and dynamic properties within a single magnetic device.

ALCF Contribution: An ALCF staff member ran Quantum Espresso calculations on Cooley using the non-collinear DFT method with spin-coupling effects included; while relatively small calculations, these require longer runtimes than normal due to slow convergence. Another ALCF staff member contributed OpenMP performance improvements that are now available in public distribution of Quantum Espresso.

Contact

Olle G. Heinonen
Argonne National Laboratory
heinonen@anl.gov

ASCR Allocation PI: Paul Kent, Oak Ridge National Laboratory

Publication

Yi Li et al. "Giant Anisotropy of Gilbert Damping in Epitaxial CoFe Films," *Phys. Rev. Lett.* **122**, 117203 (2019). https://doi.org/10.1103/PhysRevLett.122.117203.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

June 2019

8.2.6 Hadronic Light-by-Light Scattering and Vacuum Polarization Contributions to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions

The Science

The Standard Model has been the definitive description of particle physics for the past 50 years, but it is being challenged by a discrepancy between the measured and predicted values of the magnetic moment of the muon, a heavier relative of the electron. The largest uncertainty in the theoretical calculation comes from particles that interact through the strong force, quantum chromodynamics (QCD), known as hadronic contributions. This project aims to quantify and reduce the largest uncertainties associated with this value in order to obtain the most precise calculation of the anomaly. These results, combined with data soon to come from Fermilab's Muon g-2 experiment, may lead to the discovery of new physics beyond the Standard Model.

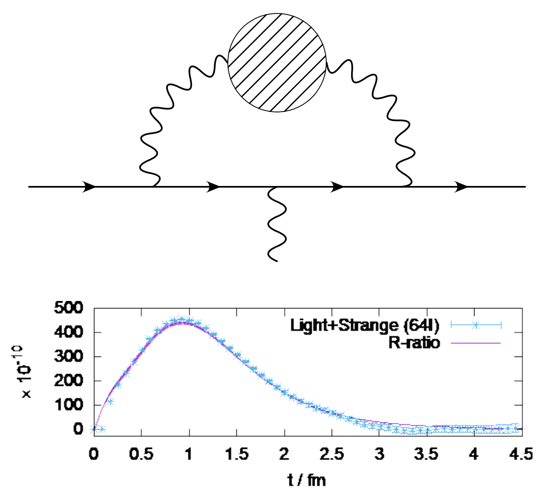

Figure 8.6 Hadronic vacuum polarization contribution to the muon anomaly. Upper: muon (horizontal line) interacting with hadrons (shaded circle) through emission of virtual photons (upper wavy lines). Lower: contribution to the anomaly (crosses) with distance between the virtual photons in femtometers.

Image courtesy of Tom Blum, University of Connecticut.

The Impact

Using Mira, the team performed large-scale lattice QCD simulations to compute the hadronic contributions (Figure 8.6). Their new result for the contribution due to the hadronic vacuum polarization has been combined with a previous result from experiment, to produce the most accurate prediction for this part of the anomaly to date. Once experimental measurements from Fermilab's Muon g-2 experiment come out later this year, this new theoretical result can be used in the comparison, potentially leading to the discovery of new physics.

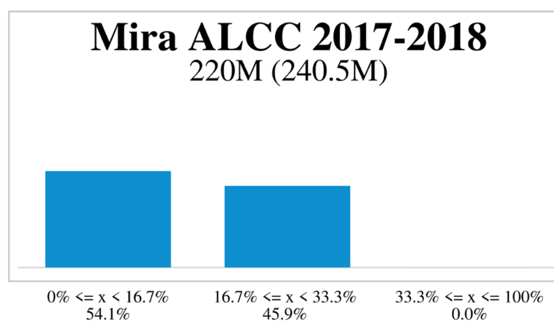

Summary

The muon's magnetic moment describes how this fundamental particle interacts with a magnetic field; it depends on all particles that can couple to the muon—including as-yet-undiscovered particles. The muon moment has been both measured in experiments and calculated theoretically, but those two values do not quite match up—hinting at the existence of new physics. A higher-precision measurement could help resolve the discrepancy, and is the goal of a new experiment in progress at Fermilab. However, the success of the experiment also relies on reducing the errors in the theoretical calculation to match the improved experiment.

The largest uncertainty in the calculation comes from particles that interact through the strong force, QCD, known as hadronic contributions. Because these contributions cannot be solved with conventional perturbative methods at low-energy scales, researchers previously resorted to experimental data or approximations with large uncertainties.

In a paper recently published in *Physical Review Letters*, the team detailed their successful determination of the most precise prediction of the leading-order hadronic vacuum polarization contribution to the muon's magnetic moment. An earlier result published by the team in *Physical Review Letters* reported the first-ever calculation of the sub-leading hadronic light-by-light contribution with physical masses, also computed on Mira. Ongoing calculations on Mira will remove and reduce systematic and statistical errors of these results.

ALCF Contribution: ALCF staff have assisted on many occasions with optimizing and debugging code on Mira.

Contact

Thomas Blum
University of Connecticut
Thomas.blum@uconn.edu

Publication

Thomas Blum et al. "Calculation of the Hadronic Vacuum Polarization Contribution to the Muon Anomalous Magnetic Moment," *Phys. Rev. Lett.* **121**, 022003 (July 2018). https://doi.org/10.1103/PhysRevLett.121.022003.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

March 2019

8.2.7 Machine Learning Helps Create Detailed, Efficient Models of Water

The Science

How water acts affects everything from storm clouds to ice sheets. Computer scientists want to model water's various properties. Accurate and computationally efficient molecular-level descriptions of large samples of ice-water systems are difficult to build. Why? The numerous molecules and various timescales remain a challenge despite advances in computing hardware. Now, a team has developed machine-learning–based water models that correctly predict water's key features, such as the melting point of ice (Figure 8.7). The team's models use a fraction of the computational cost of the best atomistic water models available today.

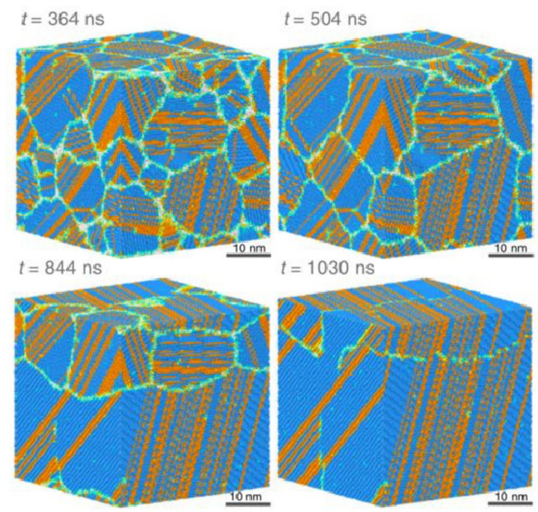

Figure 8.7 Time series of ML-BOP$_{dih}$ simulation snapshots spanning ~1 microsecond showing evolution of grain boundaries (green) between regions of hexagonal (blue) and cubic (orange) ice.

Image courtesy of Subramanian Sankaranarayanan, Argonne National Laboratory.

The Impact

The team devised a way to better model water's properties. They developed a machine-learning workflow. It offers accurate and computationally efficient models. Researchers can apply the A.I.-based approach to other material models to improve predictions.

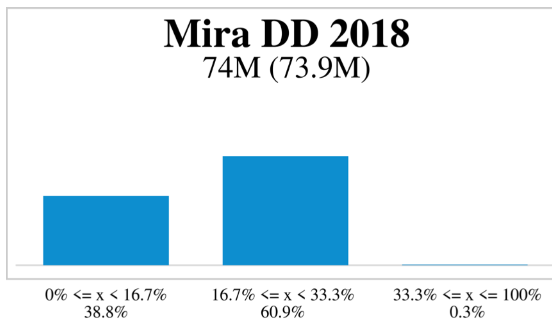

Summary

The team developed three machine-learning–based, coarse-grained water models that accurately describe the structure and thermodynamic anomalies of both water and ice at larger, mesoscopic scales based on molecular-level interactions. All three models are two orders of magnitude cheaper—in terms of computational cost—than today's atomistic models. In these models, individual water molecules are modeled as single particles that interact. In a major departure from traditional approaches, the team trained the models against first-principles calculations, experimental results, and temperature-dependent properties from molecular simulations. To overcome optimization challenges, the team employed a multi-level hierarchical workflow using global and local optimization algorithms coupled with on-the-fly molecular simulations. They used the global algorithm (Genetic Algorithm) to generate a broad survey of the parameter landscape, which is followed by multiple refinements using a local optimization algorithm with tens of milliseconds of simulation data generated on-the-fly to compute temperature-dependent properties for candidate models.

Two models successfully predicted a range of thermodynamic properties, including the density maximum of liquid water, the melting transition temperature, and the self-diffusion coefficient (a dynamic property). The team used the models on Mira to investigate homogeneous nucleation of supercooled water leading up to the formation and growth of ice grains and also made discoveries about the energy involved in forming ice. In addition, the team examined how their ML workflow could be used to improve existing water models.

ALCF Contribution: Earlier ML-BOP simulations with 8 million water molecules made use of a 2-week reservation on Mira as part of a 2017 INCITE project. ALCF staff also assisted by identifying and fixing a long-standing functionality bug in the LAMMPS BOP model that prevented the team from running multiple concurrent simulations via Python drivers.

Contact

Subramanian Sankaranarayanan
Argonne National Laboratory
ssankaranarayanan@anl.gov

Publications

Henry Chan et al. "Machine Learning Coarse Grained Models for Water," *Nat. Commun.* **10**, 379 (2019). https://doi.org/10.1038/s41467-018-08222-6.

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

March 2019

8.3 Stakeholder Engagement

To help ensure that the ALCF delivers on its mission of delivering breakthrough science, staff outreach needs to closely engage with domain science and keep a close eye on directions for supercomputing technologies. The ALCF provides a crucial balance of understanding how production science applications can move into new and exciting machine architectures in the near term and in the future.

ALCF staff support a wide range of computer science and domain science projects, and work in close collaboration with the project teams to advance their use of both production resources and future resources alike. Additionally, staff members participate in community and domain activities, especially at the annual SC conference, where ALCF staff members lead numerous sessions including technical talks, Birds of a Feather, workshops, and tutorials. Staff members are part of the organizing committee for the annual Performance Portability Workshops. Staff also visit universities and participate in town hall conversations to discuss areas of need with researchers. These activities focus not just on sharing ALCF work but collecting requirements and building collaborations. Staff members support DOE mission needs by serving on review committees, advisory boards, and participating and organizing DOE and broader community workshops.

ALCF staff are regular participants in DOE and National Science Foundation (NSF) workshops and reviews. Staff are engaged in standards committees and boards for both future and current software and hardware technologies.

8.3.1 Outreach

Engagement in Standards and Community Groups

The ALCF participates in a number of HPC standards and community groups in order to promote ALCF interests, educate the community about ALCF resources, and increase collaboration with ALCF staff. These include the following: HPC User Forum; Cray User Group; Intel Xeon Phi User Group (IXPUG); C++ Standards Committee; and MPI Forum.

Industry Lunch at SC Conference

The ALCF organizes an annual lunch for industry users of HPC at the annual SC conference. The 2019 lunch, organized by the ALCF Industry Partnerships and Outreach Manager, brought together senior ALCF staff and industry leaders to discuss how to work together more effectively. The nine companies present were a mixture of former users, current users, and potential users. The companies ranged in size from Fortune 50 to startups. This was the first year that an independent software vendor joined, which gave a valuable new perspective. The meeting included discussions of deploying new heterogeneous machines, workforce development, and some novel ideas about how industry can work more collaboratively with ALCF and other divisions at Argonne.

Quantum Computing Training

The spring 2019 Quantum Computing Tutorial was held at Argonne on May 14, 2019. Organized by ALCF staff, the tutorial provided hands-on experience with programming quantum simulators and computers. Participants learned the basics of programming IBM quantum computers with the Qiskit framework and how to use the Atos QLM35 quantum simulator in Argonne's JLSE. The tutorial brought together about 50 students, postdoctoral fellows, and staff from Argonne and the University of Chicago.

Summer Student Research Programs

Every summer, ALCF staff members mentor undergraduate and graduate students on real-world research projects through DOE's Science Undergraduate Laboratory Internship (SULI) program and Argonne's Research Aide Appointments. In 2019, 29 students worked on projects ranging from system administration and data analytics to computational science and performance engineering.

8.3.2 Summary of Engagements with the Exascale Computing Project

Argonne is a core laboratory of the ECP, and several members of ALCF's leadership team are engaged in the ECP project. Katherine Riley is a level-3 lead for the Application Integration area. Susan Coghlan and David Martin is a part of the ECP leadership team: Coghlan is deputy director of Hardware and Integration (HI) and Martin is co-executive director of the Industry Council. Other leadership team members participate in the various working groups and projects, including Mark Fahey (Facilities), Jini Ramprakash (Facilities), and Scott Parker (Software Technology and Hardware Integration). ALCF Division Director Michael E. Papka regularly participates in teleconferences with the ECP project director and other facility directors. In addition, numerous other ALCF staff members have roles in the projects and working groups listed above, as well as PathForward and CI efforts.

Last year, 26 ALCF staff attended the ECP Annual Meeting held January 14–18, 2019, in Houston, Texas, to participate in technical conversations, project discussions, and facility-specific breakouts. In addition, the ALCF participated in several planning meetings with ECP and the other computing facilities (NERSC, OLCF) to augment and execute the ECP/Facilities engagement plan, and worked with the ECP training lead to promote ECP training activities to ALCF users.

The ALCF ECP/HI Applications Integration effort made great strides in 2019 by hiring multiple performance engineers and training staff to support application readiness. There are 17 staff members funded at various levels to work with ECP Application Development projects, 7 of those were new hires for the ALCF.

The ALCF hired a staff member to engage with ECP on their CI needs, and, in 2019, moved that system closer to deployment. An additional staff member was hired in 2019 to work on ECP's planned packaging tool, Spack.

GitLab Continuous Integration

Since 2018, Onyx Point, Inc. has been working in collaboration with ECP partners, including the ALCF, to deliver a CI solution for DOE HPC facilities. The IT company is building consolidated testing systems that will provide users with easy access to automated builds at HPC facilities, and will allow scientists to test open source projects and internal projects automatically, and securely, on DOE HPC machines. This work is intended to dramatically increase the reliability, quality, and performance of the HPC software ecosystem; which is one of the key goals of the ECP.

In 2019, the ECP team completed initial development efforts with Onyx Point to deliver the core features required for CI to enable HPC workflows and federation for multi-site run features. A preview environment was deployed at ALCF and connected to Theta with local site capabilities. Science teams were given access to the preview environment to construct CI pipelines for their specific project uses and have provided feedback to the ECP CI team. A proof-of-concept effort was also completed with the HPC Software Integration team to demonstrate deploying HPC software packages using their planned packaging tool, Spack, and leveraging automated CI workflows in the ECP preview environment at ACLF.

Spack Software Package Management

Working with ECP, ALCF staff installed, and improved, package management tool Spack to manage ECP and all ALCF software going forward. Spack supports multiple versions and configurations of software on a wide variety of platforms and environments, including supercomputers. ALCF staff architected a multi-user deployment of Spack, intended to make a single instance of Spack available to all users, as well as packagers and publishers of software via Linux groups. The installation was intended to put Spack-built software in users' hands as quickly as possible, allow CI deployments, leverage components built by others, and maintain system settings in an organized manner—all while providing easy traceability of package ownership simply by looking at the user-owner of installed packages.

As the principal site attempting to use Spack in this fashion enabled the ALCF to report bugs relating to file and group ownership of Spack-internal cache, databases, etc. One important bug that ALCF staff discovered, and which was subsequently fixed, related to where builds were being cached (previously, it required a portion of the Spack installation to be world-writable). Other successes include being able to build a significant portion of the extreme-scale scientific software development kit xSDK on Theta to a separate Spack installation, and reproducing an early containerized prototype of the Kitware Spack build pipeline, which may serve as the basis for future Spack-driven deployments at the ALCF.

A focus of effort for 2020 is to define a 'facility Spack stack' to be used as input for this pipeline and to architect a production version of the pipeline (whether containerized or bare metal) that will address cross-compilation, scheduler integration, and integration with existing vendor (Cray PE, Intel, etc.) and ALCF software deployments. Additional goals are to clean up and curate the current software environment, develop a coherent deployment for Python-based software, and improve the viability of container-based workflows.

Communication between the ALCF and the ECP Resource Allocation Council

In 2018, the ECP ALCC allocation completed and the compute facilities and ECP switched to the RAC (discussed in Section 3.1.2) to support ECP computing needs. This group, composed of representatives of the facilities and the ECP, meets monthly to review project progress and to assess new project needs.

To help automate how the RAC consumes this data, the ALCF also sends a .CSV file (comma-separated values file) report to the ECP mailing list each week. While the web version is visually appealing, the .CSV format can be merged into any other workflow to process.

Conclusion

The 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. In CY 2019, researchers participating in projects using ALCF resources published 284 papers in high-quality conferences and journals. 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.

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 outage** time 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 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 the 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:

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

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*:**

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

A.3 System Mean Time to Interrupt (MTTI) Calculation Details

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

Formula:

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

A.4 System Mean Time to Failure (MTTF) Calculation Details

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

Formula:

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

A.5 Utilization Calculation Details

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. Jobs that ran during an outage are excluded.

Formula:

$$\mathbf{Utilization} = \left(\frac{\boldsymbol{corehours\ used\ in\ period}}{\boldsymbol{corehours\ available\ in\ period}}\right) * \mathbf{100}$$

A.6 Capability

Capability is an attribute assigned to user jobs that meet the capability definition for a machine. **High Capability** is an attribute assigned to user jobs that meet the high capability definition for a machine.

Tables A.1 and A.2 show the capability definitions for each reportable machine.

Table A.1 Capability Definitions for Mira

Mira				
Capability	High Capability	Range	Minimum Nodes/Cores	Maximum Nodes/Cores
No	No	0% <= x < 16.7%	1 / 16	8,191 / 131,056
Yes	No	16.7% <= x < 33.3%	8,192 / 131,072	16,383 / 262,128
Yes	Yes	33.3% <= x <= 100%	16,384 / 262,144	49,152 / 786,432

Table A.2 Capability Definitions for Theta

Theta				
Capability	**High Capability**	**Range**	**Minimum Nodes/Cores**	**Maximum Nodes/Cores**
No	No	0% <= x < 20.0%	1 / 64	799 / 51,136
Yes	No	20.0% <= x < 60.0%	800 / 51,200	2,399 / 153,536
Yes	Yes	60.0% <= x	2,400 / 153,600	See: A.7 Theta Nodes

Capability also refers to a calculation. The capability calculation is the percentage of core-hours of jobs with the capability attribute versus the total core-hours of all jobs. The calculation can be applied to a class of jobs. For example: Innovative and Novel Computational Impact on Theory and Experiment (INCITE) capability is the percentage of core-hours of INCITE jobs with the capability attribute versus the total core-hours of all INCITE jobs for a time period.

***Formula*:**

$$\textbf{OVERALL CAPABILITY} = \left(\frac{\textbf{\textit{Capability Core Hours Consumed}}}{\textbf{Total Core Hours Consumed}}\right) * \mathbf{100}$$

$$\textbf{HIGH CAPABILITY} = \left(\frac{\textbf{\textit{High Capability Core Hours Consumed}}}{\textbf{Total Core Hours Consumed}}\right) * \mathbf{100}$$

A.7 Theta Nodes

The number of reportable nodes on Theta is fewer than the total number of nodes. The total node count for Theta changed during 2017, as shown in Table A.3.

Table A.3 Total and Reportable Nodes for Theta

Theta		
Data Range	**Total Nodes**	**Reportable Nodes**
07/01/2017 – 12/12/2017	3,624	3,240
12/13/2017 – 12/31/2017	4,392	3,240
01/01/2018	4,392	4,008

The reportable node count is used in the following calculations:

- Scheduled Availability: Affects the scheduled outage and unscheduled outage calculations when the node count in the outage was fewer than the total number of nodes.
- Overall Availability: Affects the scheduled outage and unscheduled outage calculations when the node count in the outage was fewer than the total number of nodes.

- Utilization: The calculation capped the daily utilization at 100 percent of reportable nodes. The number of core-hours for each day was calculated as the minimum of the core-hours used and the core-hours possible.
- Overall Capability: 20 percent of the reportable nodes.
- High Capability: 60 percent of the reportable nodes.

Appendix B – Argonne Leadership Computing Facility's Director's Discretionary Projects: Mira

January 1, 2019 – December 31, 2019
Director's Discretionary (DD) Projects on Mira

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Allinea	Raymond M. Loy, Kalyan Kumaran	Argonne National Laboratory	Improved debugging memory usage for BG/Q	Internal	2,000,000
AMASE	Rajkumar Kettimuthu	Argonne National Laboratory	Architecture and Management for Autonomous Science Ecosystems	Computer Science	977,795
ART_Physics_DD	Hussein Aluie	University of Rochester	Multiscale physics of the Ablative Rayleigh-Taylor Instability	Fusion Energy	15,000,000
ATPESC19_Instructors	Marta Garcia Martinez, Raymond M. Loy	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing for ALL Instructors	Computer Science	100,000
ATPESC2019	Marta Garcia Martinez, Raymond M. Loy	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing 2019	Training	10,000,000
AutomatedBench	Kevin Harms	Argonne National Laboratory	Automated Benchmarking	Computer Science	3,000,000
bloodflow_dd	Jifu Tan	Northern Illinois University (NIU)	Microfluidic design and optimization for cell separation	Engineering	2,000,000
Bubblecollapse	Eric Johnsen	University of Michigan	Dynamics and heat transfer in bubble collapse near solid surfaces	Engineering	1,500,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Carbon_composites	Hendrik Heinz	University of Colorado-Boulder	Rational design of ultra-strong composites	Materials Science	3,300,000
Catalyst	Katherine M. Riley, Christopher Knight, James Osborn, Timothy Williams	Argonne National Laboratory	Catalyst	Internal	10,000,000
CEED_ECPAD	Gustaaf Bernardus Jacobs	Argonne National Laboratory	CEED	Mathematics	1,030,000
CFD-TPM	Surya Vegendla	Argonne National Laboratory	Validation of two-phase flow models and critical heat flux prediction for the highly scalable CFD code NEK-2P	Nuclear Energy	7,000,000
CharmRTS	Laxmikant Kale, Abhinav Bhatele, Juan Jose Galvez-Garcia	University of Illinois at Urbana-Champaign	Charm++ and its applications	Computer Science	3,000,000
climate_severe	Vittorio Gensini	Northern Illinois University (NIU)	Anticipating Severe Weather Events via Dynamical Downscaling	Earth Science	1,000,000
CNTmetallization	Iman Salehinia, Michael E. Papka	Northern Illinois University (NIU)	Metallization of CNT for thermal and structural applications	Materials Science	1,000,000
CobaltDevel	Paul Michael Rich, William Edward Allcock	Argonne National Laboratory	Cobalt Development	Internal	10,000,000
CombustOpt_DD	Sibendu Som, Chao Xu, Prithwish Kundu	Argonne National Laboratory	Heavy-duty Diesel Engine Combustion Optimization for Reduced Emissions, Reduced Heat Transfer, and Improved Fuel Economy	Engineering	10,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CORALDev	Scott Parker	Argonne National Laboratory	CORAL Development and Testing	Internal	1,000,000
critical_perf	Scott Parker, Raymond M. Loy	Argonne National Laboratory	Critical Debugging project	Internal	10,000,000
CrystalsADSP	Alexandre Tkatchenko, Alvaro Vazquez Mayagoitia	University of Luxembourg	Constructing and Navigating Polymorphic Landscapes of Molecular Crystals	Materials Science	51,000,000
CSC249ADCD04	Tzanio Kolev, Misun Min, Paul Fischer	Lawrence Livermore National Laboratory (LLNL)	1.2.5.3.04 Center for Efficient Exascale Discretizations (CEED)	Computer Science	1,784,000
CSC249ADCD05	Mahantesh Halappanavar	Pacific Northwest National Laboratory (PNNL)	1.2.5.3.05 ADCD05-ExaGraph	Computer Science	720,000
CSC249ADCD502	Kenneth John Roche	Pacific Northwest National Laboratory (PNNL)	1.2.5.02 ADCD502 Application Assessment	Computer Science	880,000
CSC249ADSE03	Andreas Kronfeld, Norman Howard Christ, Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	1.2.1.03 ADSE03-LatticeQCD	Physics	52,480,000
CSC249ADSE08	Steven Hamilton, Paul Kollath Romano	Oak Ridge National Laboratory (ORNL)	1.2.1.08 ADSE08 ExaSMR	Nuclear Energy	2,560,000
CSC249ADSE09	Paul Kent, Anouar Benali	Oak Ridge National Laboratory (ORNL)	1.2.1.09 QMCPACK	Materials Science	1,920,000
CSC249ADSE16	Mark S Gordon	Ames Laboratory	2.2.1.03 ADSE16-GAMESS	Chemistry	2,696,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADSE17	Charlie Catlett, Melissa Allen, Rajeev Jain, Scott A Ehling	Argonne National Laboratory	2.2.4.01 Multiscale Coupled Urban Systems	Computer Science	2,080,000
CSC250STDA05	Kenneth Moreland	Sandia National Laboratories, California	2.3.4.13 STDA05-ECP/VTK-m	Computer Science	3,200
CSC250STDM08	Franck Cappello, Scott A Ehling	Argonne National Laboratory	2.3.4.06 STDM08-Ez: Fast, Effective, Parallel Error-bounded Exascale Lossy Compression for Scientific Data	Computer Science	3,360,000
CSC250STDM10	Surendra Byna, Venkatram Vishwanath	Lawrence Berkeley National Laboratory (LBNL)	2.3.4.15 ExaIO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Unify	Computer Science	32,000
CSC250STMS07	Barry Smith, Hong Zhang, Satish Balay, Todd S. Munson	Argonne National Laboratory	1.3.3.07 STMS07-PETSc/TAO	Mathematics	8,000
CSC250STPM11	George Bosilca, Earl Carr, Jack Dongarra, Thomas Herault	The University of Tennessee at Knoxville	1.3.1.11 STPM11-Distributed Tasking (ParSEC)	Computer Science	800,000
CSC250STPM17	Paul Hamilton Hargrove	Lawrence Berkeley National Laboratory (LBNL)	1.3.1.17 STPM17-PGAS	Computer Science	16,000
CSC250STTO09	Jack Dongarra, Anthony Danalis, Earl Carr, Heike Jagode	The University of Tennessee at Knoxville	1.3.2.09 STTO09-ExaPAPI	Computer Science	1,296,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STTO11	John Michael Mellor-Crummey	Rice University	2.3.2.08 STTO11 - Extending HPCToolkit for Exascale	Computer Science	240,000
CSC251HIHE05	Scott Pakin	Los Alamos National Laboratory (LANL)	2.4.2.05 HIHE05- Analytical Modeling	Computer Science	320,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	2,000,000
DNS3D	Ramesh Balakrishnan	Argonne National Laboratory	Direct Numerical Simulation of Three-Dimensional Turbulence	Engineering	89,000,000
DNS_gravity	Som Dutta	Utah State University	Investigating the Coherent Structures and Transport Processes within Gravity Currents using High-Fidelity CFD	Engineering	10,000,000
DSEM_SUPERSONIC	Farzad Mashayek	University of Illinois at Chicago	Simulation of Supersonic Combustion	Engineering	858,229
Dynstall_DD	Anupam Sharma	Iowa State University (ISU)	Analysis of Dynamic Stall and its Mitigation in Energy Machines using Bioinspired Blade Designs	Engineering	20,000,000
ExM	Justin Michael Wozniak, Michael Joseph Wilde	Argonne National Laboratory	Extreme many-task computing with Swift	Computer Science	5,000,000
EZ	Franck Cappello	Argonne National Laboratory	Ez: Fast, Effective, Parallel Error-bounded Exascale Lossy Compression for Scientific Data	Computer Science	1,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
GFMC_responses	Alessandro Lovato	Argonne National Laboratory	Green<ESCDQ>s function Monte Carlo calculations of neutrino-nucleus scattering	Physics	24,000,000
GrainBoundaries	Wissam Saidi	University of Pittsburgh	Structure and Properties of Grain Boundaries in Materials for Energy Applications	Materials Science	12,000,000
HadronicLight_DD	Thomas Blum	University of Connecticut	Hadronic light-by-light scattering contribution to the muon anomalous magnetic moment from lattice QCD with chiral fermions	Physics	20,000,000
HetMat_DD		The University of Chicago (UChicago)	Large scale simulations of heterogeneous materials for energy conversion applications	Materials Science	20,000,000
HHPMT_PostProcess	Karl Daniel Hammond	University of Missouri-Columbia	Post-Processing of HHPMT_3 files	Fusion Energy	250,000
HighReyTurb_PostProc	Robert D. Moser, Myoungkyu Lee	The University of Texas at Austin	Data analysis of turbulent Channel Flow at High Reynolds number	Engineering	15,000,000
HTI_MIRA	Charles Moulinec	Science and Technology Facilities Council (STFC)	Heat Transfer Impairment due to Carbon Deposition on Riblets of Fuel Assemblies	Nuclear Energy	16,400,000
ICE_BOP	Subramanian Sankaranarayanan	Argonne National Laboratory	Phase transitions in water-ice-vapor system	Materials Science	10,000,000
JCESR	Larry Curtiss	Argonne National Laboratory	Development of High Throughput Methods	Materials Science	10,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
LastJourney	Katrin Heitmann	Argonne National Laboratory	Extreme-scale Cosmology - The Last Journey	Physics	800,000,000
LatticeQCD-CSD	Norman Howard Christ	Columbia University	Fourier acceleration of the Hybrid Monte Carlo algorithm for Lattice QCD	Physics	20,000,000
LQCDdev	James Osborn	Argonne National Laboratory	Lattice QCD development	Physics	1,000,000
LTC_Aramco	Sibendu Som, Muhsin Mohammed Ameen, Pinaki Pal, Yuanjiang Pei	Argonne National Laboratory	Investigation of a Low Octane Gasoline Fuel for a Heavy-Duty Diesel Engine in a Low-Temperature Combustion Regime	Engineering	30,000,000
MagnetismHPC_PostPro	Jonathan M Aurnou	University of California-Los Angeles	Frontiers of Planetary and Stellar Magnetism	Earth Science	1,000,000
magnetotail	Samuel Richard Totorica	Princeton University	Kinetic simulations of the dynamic magnetotail	Physics	10,000,000
Maintenance	Mark R Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	150,000,000
MDHTProp	Mark Christian Messner	Argonne National Laboratory	Assessing the scalability of direct MD calculations of structural material properties	Engineering	1,000,000
MPICH_MCS	Kenneth James Raffenetti, Pavan Balaji	Argonne National Laboratory	MPICH - A high performance and widely portable MPI implementation	Computer Science	12,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
NCAlloy_TM_Stability	Garritt Tucker	Colorado School of Mines	Ascertaining the Thermo-Mechanical Mechanisms of Solute-Stabilized Nanocrystalline Alloys	Materials Science	1,000,000
Nek_VHTR	Masahiro Kawaji, Aleksandr V. Obabko	City College of New York	Scalability and Validation of Nek5000 for VHTR challenge problem of pipe flow relaminarization	Engineering	4,000,000
Omega-NIF_Exp	Petros Tzeferacos	The University of Chicago (UChicago)	Mira Simulations of High-Intensity Laser Experiments to Study Turbulent Amplification of Magnetic Fields	Physics	20,000,000
Operations	Mark R Fahey	Argonne National Laboratory	Systems administration tasks	Internal	32,000,000
P3DFFT	Dmitry Pekurovsky	University of California-San Diego	Performance studies of three-dimensional Fast Fourier Transforms using overlap of communication with computation	Computer Science	1,000,000
Performance	Scott Parker, Raymond M. Loy	Argonne National Laboratory	Performance	Internal	10,000,000
PHASTA_dd	Kenneth Edward Jansen	University of Colorado-Boulder	PHASTA_dd	Engineering	20,000,000
PHASTA_NCSU	Igor A Bolotnov	North Carolina State University (NCSU)	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	10,000,000
PlasmaReconnection	Andrey Beresnyak	Naval Research Laboratory	Studying plasma reconnection and diffusion	Physics	3,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
PlasticNeocortex_DD	Giuseppe Chindemi	École Polytechnique Fédérale de Lausanne (EPFL)	Biophysical principles of learning and memory in the neocortex	Biological Sciences	5,000,000
radix-io	Philip Carns	Argonne National Laboratory	System software to enable data-intensive science	Computer Science	1,235,865
RCM_4km	Jiali Wang	Argonne National Laboratory	Generation of a next level dataset for regional scale climate modeling: convective resolving spatial scales	Earth Science	10,000,000
REI_Flares	Marc Cremer	Reaction Engineering International	Leveraging the UCF for Simulation of Industrial Flares	Chemistry	750,000
SDXServiceDTN	Jim Chen	Northwestern University	Software Defined Network Exchange for Data Transfer Node	Computer Science	1,000
SherpaProd	John Taylor Childers	Argonne National Laboratory	Sherpa Parton Generator Production	Physics	250,000
SolarWindowsADSP	Jacqueline Cole	University of Cambridge	Data-driven molecular engineering of solar powered windows	Materials Science	29,500,000
SoPE_4_DD	Thomas Jackson	University of Florida	Demonstration of Scalability of Programming Environments by Simulating	Engineering	5,000,000
SPMV_Scalablity		University of Notre Dame	Strong Scalability of Hybrid Sparse Problems	Computer Science	10,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SprayWall_UMassD	Roberto Torelli	Argonne National Laboratory	Direct Numerical Simulations of Droplet-Wall Impingement under High-Pressure Spray Conditions	Engineering	7,117,265
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	2,000,000
TotalView	Peter Thompson, Raymond M. Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene P	Internal	100,000
TurbulentLiquidDrop	Arne Pearlstein	University of Illinois at Urbana-Champaign	Computation of Transitional and Turbulent Drop Flows for Liquid Carbon Dioxide Drops Rising in Seawater	Engineering	500,000
UltraFlux	Frederico Fiuza	SLAC National Accelerator Laboratory	Unveiling the 3D Physics Behind Compact Ultrahigh Flux Neutron Sources	Physics	50,000,000
User_Services	Haritha Siddabathuni Som, Sreeranjani Ramprakash	Argonne National Laboratory	User Services	Internal	20
visualization	Joseph Insley, Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	2,000,000
VTR	Dillon Shaver	Argonne National Laboratory	Thermal Hydraulic Simulations for the Versatile Test Reactor	Nuclear Energy	9,255,061
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	70,500,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
XGC_mira	Timothy Williams	Argonne National Laboratory	Testing Modern XGC Code on Mira	Fusion Energy	500,000
				Total Mira DD	**1,789,330,435**

This page intentionally left blank.

Appendix C – Argonne Leadership Computing Facility Director's Discretionary Projects: Theta

January 1, 2019 – December 31, 2019
Director's Discretionary (DD) Projects on Theta

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
3DChromatin	Jie Liang	University of Illinois at Chicago	Large Ensemble Model of Single-Cell 3D Genome Structures	Biological Sciences	2,000,000
3M_Next	Ian Foster	Argonne National Laboratory	Next Generation Nonwovens Manufacturing based on Model-driven Simulation Machine Learning Approach	Engineering	3,000,000
Acoustic_bubbles	Parviz Moin	Stanford University	Acoustic characterization of oceanic bubbles in turbulent environments using direct computation and modeling methods	Engineering	2,000,000
AD_Brain_Imaging	Jiook Cha	Columbia University	Computational Analysis of Brain Connectomes for Alzheimer's Disease	Biological Sciences	3,000,000
AIASMAAR		Argonne National Laboratory	Artificial Intelligence Assisted Safety Modeling and Analysis of Advanced Reactors	Nuclear Energy	2,000,000
aimaterialsADSP	Marco Govoni	The University of Chicago (UChicago)	Advanced Materials Characterization with AI-informed Computation	Materials Science	4,000,000
AlgLA	Oded Schwartz	The Hebrew University	Algorithmic Linear Algebra	Computer Science	500,000
Allinea	Raymond M. Loy, Kalyan Kumaran	Argonne National Laboratory	Improved debugging memory usage for BG/Q	Internal	100,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
AMASE	Rajkumar Kettimuthu	Argonne National Laboratory	Architecture and Management for Autonomous Science Ecosystems	Computer Science	100,000
AMEWS_DD	Giulia Galli	The University of Chicago (UChicago)	Advanced Materials for Energy-Water Systems Center	Materials Science	15,000,000
AM_modelling	Lianghua Xiong	Argonne National Laboratory	Advanced modelling of phase transformation in metal additive manufacturing	Materials Science	500,000
APS-Linac	Yine Sun	Argonne National Laboratory	APS Linac Operation Optimization with AI	Physics	200,000
APSDataAnalysis	Nicholas Schwarz	Argonne National Laboratory	APS Beamline Data Processing and Analysis	Computer Science	200,000
ARPA-E-NPM-2019		University of Colorado-Boulder	Thermal Conductivity of Doped Nanophononic Metamaterials Using Massively Parallel Molecular Dynamics Simulations	Materials Science	2,520,000
AstroAccelTheta	Dmitri Uzdensky	University of Colorado-Boulder	Extreme Particle Acceleration in Astrophysical Environments	Physics	500,000
AstroGK_turb	Yohei Kawazura	University of Oxford	Energy flow and heating in astrophysical plasma turbulence	Physics	1,000,000
athena_performance	Brian William O'Shea	Michigan State University	Scaling and performance enhancement of an astrophysical plasma code	Physics	2,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
atlasMLbjets	Rui Wang	Argonne National Laboratory	Using ML in b-jet identification at ATLAS	Physics	4,000,000
atlas_aesp	Walter Hopkins	Argonne National Laboratory	Simulating and Learning in the ATLAS detector at the Exascale	Physics	8,000,000
ATPESC19_Instructors	Marta Garcia Martinez, Raymond M. Loy	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing for ALL Instructors	Computer Science	100,000
ATPESC2019	Marta Garcia Martinez, Raymond M. Loy	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing 2019	Training	3,000,000
AuPM6	Emilie Brigitte Guidez	University of Colorado-Denver	Reparametrization of the PM6 method for gold nanoclusters	Chemistry	4,474,368
AutoBEM	Joshua New	Oak Ridge National Laboratory (ORNL)	Automatic Building Energy Modeling and analysis	Energy Technologies	1,000,000
AutomatedBench	Kevin Harms	Argonne National Laboratory	Automated Benchmarking	Computer Science	900,000
BIP167	Philip Kurian	Howard University	Computing van der Waals many-body dispersion effects for MD simulations and THz spectroscopy of biomacromolecules in water	Physics	1,000,000
bloodflow_dd	Jifu Tan	Northern Illinois University (NIU)	Microfluidic design and optimization for cell separation	Engineering	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
BrainImaging	Walter Scheirer	University of Notre Dame	BrainImaging ADSP Bridge to INCITE Proposal	Biological Sciences	5,000,000
BrainImagingML	Thomas David Uram	Argonne National Laboratory	Large-scale Brain Imaging and Reconstruction	Biological Sciences	400,000
Bubblecollapse	Eric Johnsen	University of Michigan	Dynamics and heat transfer in bubble collapse near solid surfaces	Engineering	1,500,000
Bubble_Collapse	Charlotte Barbier	Oak Ridge National Laboratory (ORNL)	High Definition Simulation of cavitating bubble near a wall with a shear flow	Engineering	1,000,000
BuoyantTurbulence	Aleksandr V. Obabko	Argonne National Laboratory	Testing a new theory of density effects in turbulent transport and mixing	Earth Science	400,000
candle_aesp	Rick Lyndon Stevens, Thomas Brettin, Venkatram Vishwanath	Argonne National Laboratory	Virtual Drug Response Prediction	Biological Sciences	20,000,000
Carbon_composites	Hendrik Heinz	University of Colorado-Boulder	Rational design of ultrastrong composites	Materials Science	2,000,000
catalysis_aesp	David Bross	Argonne National Laboratory	Exascale Computational Catalysis	Chemistry	16,600,000
Catalyst	Katherine M Riley, Christopher Knight, James Osborn, Timothy Williams	Argonne National Laboratory	Catalyst	Internal	10,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CESM_Highres_Testing	Gerald Meehl, Susan C Bates	The National Center for Atmospheric Research (NCAR)	Porting and Benchmarking CESM 1.3 on Theta	Earth Science	2,000,000
cfdml_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Data Analytics and Machine Learning for Exascale CFD	Engineering	13,000,000
CharDynStall	Anupam Sharma	Iowa State University (ISU)	Characterization of Dynamic Stall Onset	Engineering	2,000,000
CharmRTS	Laxmikant Kale, Abhinav Bhatele, Juan Jose Galvez-Garcia	University of Illinois at Urbana-Champaign	Charm++ and its applications	Computer Science	3,000,000
climate_severe	Vittorio Gensini	Northern Illinois University (NIU)	Anticipating Severe Weather Events via Dynamical Downscaling	Earth Science	1,998,117
CobaltDevel	Paul Michael Rich, William Edward Allcock	Argonne National Laboratory	Cobalt Development	Internal	2,000,000
Coin	Chukwunwike Iloeje	Argonne National Laboratory	DMDII-Argonne CNC optimization through machine learning	Engineering	100,000
CompCatalysis	Rajeev Surendran Assary	Argonne National Laboratory	Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning	Chemistry	8,000,000
Comp_Perf_Workshop	Raymond M. Loy	Argonne National Laboratory	ALCF Computational Performance Workshop	Training	7,500,000
connectomics_aesp	Nicola Ferrier, Thomas David Uram	Argonne National Laboratory	Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience	Biological Sciences	30,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CONUS-Carbon	Jinxun Liu	U.S. Geological Survey (USGS)	Terrestrial ecosystem carbon cycle of the conterminous U.S.	Earth Science	4,000,000
Cray	Torrance Ivan Leggett, Mark R Fahey, Susan Marie Coghlan, Timothy Williams	Cray Inc.	Cray Installation	Internal	10,000,000
CrystalsADSP	Alexandre Tkatchenko, Alvaro Vazquez Mayagoitia	University of Luxembourg	Constructing and Navigating Polymorphic Landscapes of Molecular Crystals	Materials Science	8,000,000
CSC249ADCD01	Ian Foster	Argonne National Laboratory	1.2.5.3.01 ADCD01-CODAR	Computer Science	1,536,000
CSC249ADCD02	Susan Mniszewski, Timothy C. Germann	Los Alamos National Laboratory (LANL)	2.2.6.04 ADCD02-COPA: Co-Design Center for Particle Applications	Physics	128,000
CSC249ADCD04	Tzanio Kolev, Misun Min, Paul Fischer	Lawrence Livermore National Laboratory (LLNL)	1.2.5.3.04 Center for Efficient Exascale Discretizations (CEED)	Computer Science	2,265,088
CSC249ADCD05	Mahantesh Halappanavar	Pacific Northwest National Laboratory (PNNL)	1.2.5.3.05 ADCD05-ExaGraph	Computer Science	1,504,000
CSC249ADCD08	Francis Alexander	Brookhaven National Laboratory (BNL)	2.2.6.08 ADCD08-ExaLearn	Physics	768,000
CSC249ADCD502	Kenneth John Roche	Pacific Northwest National Laboratory (PNNL)	1.2.5.02 ADCD502 Application Assessment	Computer Science	1,344,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADCD504	David Richards, Shirley Victoria Moore	Lawrence Livermore National Laboratory (LLNL)	2.2.6.01 ADCD504-Proxy Applications	Computer Science	198,400
CSC249ADOA01	Rick Lyndon Stevens, Thomas Brettin	Argonne National Laboratory	1.2.3.01 ADOA01 CANDLE	Biological Sciences	3,584,000
CSC249ADSE01	Salman Habib, Katrin Heitmann	Argonne National Laboratory	1.2.1.01 ADSE01-ExaSky	Physics	640,000
CSC249ADSE03	Andreas Kronfeld, Norman Howard Christ, Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	1.2.1.03 ADSE03-LatticeQCD	Physics	19,840,000
CSC249ADSE04	Danny Perez	Los Alamos National Laboratory (LANL)	1.2.1.04 ADSE04-EXAALT - Molecular dynamics at the exascale	Nuclear Energy	2,560,000
CSC249ADSE05	David Trebotich	Lawrence Berkeley National Laboratory (LBNL)	1.2.1.05 ADSE05-Subsurface	Earth Science	128,000
CSC249ADSE08	Steven Hamilton, Paul Kollath Romano	Oak Ridge National Laboratory (ORNL)	1.2.1.08 ADSE08 ExaSMR	Nuclear Energy	1,984,000
CSC249ADSE09	Paul Kent, Anouar Benali	Oak Ridge National Laboratory (ORNL)	1.2.1.09 QMCPACK	Materials Science	1,344,000
CSC249ADSE11		University of Washington	2.2.1.02 ADSE11-NWChemEx: Tackling Chemical, Materials, & Biomolecular Challenges in Exascale	Chemistry	704,000
CSC249ADSE12	Amitava Bhattacharjee	Princeton Plasma Physics Laboratory (PPPL)	1.2.1.12 ADSE12 WDMAPP	Computer Science	2,368,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADSE14	Jacqueline Chen	Sandia National Laboratories, California	2.2.2.02 ADSE14-Combustion-Pele: Transforming Combustion Science & Technology with Exascale Simulations	Engineering	8,640,000
CSC249ADSE16	Mark S Gordon	Ames Laboratory	2.2.1.03 ADSE16-GAMESS	Chemistry	13,171,200
CSC249ADSE17	Charlie Catlett, Melissa Allen, Rajeev Jain, Scott A. Ehling	Argonne National Laboratory	2.2.4.01 Multiscale Coupled Urban Systems	Computer Science	1,984,000
CSC249ADSE18	Daniel Kasen	Lawrence Berkeley National Laboratory (LBNL)	1.2.1.18/ADSE18 Astro (Exastar)	Physics	704,000
CSC249ADSE22	Andres Marquez, Zhenyu Huang	Pacific Northwest National Laboratory (PNNL)	1.2.1.22 ADSE22-ExaSGD	Energy Technologies	4,544,000
CSC249ADTR02		Oak Ridge National Laboratory (ORNL)	2.4.6.02 ADTR02-Productivity	Computer Science	192,000
CSC250STDA04	James Ahrens, Terece Louise Turton	Lawrence Livermore National Laboratory (LLNL)	2.3.4.12 STDA04-ALPINE: Algorithms and Infrastructure for In Situ Visualization and Analysis	Computer Science	353,280
CSC250STDA05	Kenneth Moreland	Sandia National Laboratories, California	2.3.4.13 STDA05-ECP/VTK-m	Computer Science	76,800
CSC250STDM07	Franck Cappello	Argonne National Laboratory	2.3.4.05 STDM07-VeloC: Very Low Overhead Transparent Multilevel Checkpoint/Restart	Computer Science	1,984,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STDM08	Franck Cappello, Scott A Ehling	Argonne National Laboratory	2.3.4.06 STDM08-Ez: Fast, Effective, Parallel Error-bounded Exascale Lossy Compression for Scientific Data	Computer Science	1,344,000
CSC250STDM09	Kathryn Mohror	Lawrence Livermore National Laboratory (LLNL)	2.3.4.07 STDM09-UCNIFYCR: A Checkpoint/Restart File System for Distributed Burst Buffers	Computer Science	128,000
CSC250STDM10	Surendra Byna, Venkatram Vishwanath	Lawrence Berkeley National Laboratory (LBNL)	2.3.4.15 ExaIO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Unify	Computer Science	800,000
CSC250STDM11	Scott Klasky, Norbert Podhorszki	Oak Ridge National Laboratory (ORNL)	2.3.4.09 STDM11-ADIOS Framework for Scientific Data on Exascale Systems	Computer Science	1,600,000
CSC250STDM12	Robert B. Ross, Robert J Latham	Argonne National Laboratory	2.3.4.10 STDM12-DataLib: Data Libraries and Services Enabling Exascale Science	Computer Science	1,280,000
CSC250STDM14		Argonne National Laboratory	2.3.4.14 STDM14 - VeloC: Very Low Overhead Transparent Multilevel Checkpoint/Restart/SZ: Fast, Effective, Parallel Error-bounded Exascale Loss....	Computer Science	576,000
CSC250STDM15		Lawrence Livermore National Laboratory (LLNL)	2.3.4.15 STDM15 - HDF5/UnifyCR	Computer Science	64,000
CSC250STDM16		Lawrence Livermore National Laboratory (LLNL)	2.3.4.16 STDM16-ALPINE/ZFP	Computer Science	64,000
CSC250STDT10	Jeffrey S Vetter	Oak Ridge National Laboratory (ORNL)	2.3.2.10 STDT10-PROTEAS/Y-Tune	Computer Science	64,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STML13	Jack Dongarra	The University of Tennessee	2.3.3.13 STML13 - CLOVER	Computer Science	64,000
CSC250STMS05	Ulrike Meier Yang, Satish Balay	Argonne National Laboratory	2.3.3.01 STMS05- Extreme-scale Scientific xSDK for ECP	Mathematics	2,240,000
CSC250STMS07	Barry Smith, Hong Zhang, Satish Balay, Todd S. Munson	Argonne National Laboratory	1.3.3.07 STMS07- PETSc/TAO	Mathematics	1,440,000
CSC250STMS08	Xiaoye Sherry Li	Lawrence Berkeley National Laboratory (LBNL)	2.3.3.07 STMS08- STRUMPACK/SuperLU: Factorization Based Sparse Solvers and Preconditioners for Exascale	Mathematics	64,000
CSC250STNS01	David Rogers, Terece Louise Turton	Los Alamos National Laboratory (LANL)	2.3.6.01 – STNS01 -- LANL ATDM ST Projects	Computer Science	64,000
CSC250STPM07	Bruce Joseph Palmer	Pacific Northwest National Laboratory (PNNL)	1.3.1.07 STPM07 xGA	Computer Science	128,000
CSC250STPM11	George Bosilca, Earl Carr, Jack Dongarra, Thomas Herault	The University of Tennessee at Knoxville	1.3.1.11 STPM11- Distributed Tasking (ParSEC)	Computer Science	1,984,000
CSC250STPM17	Paul Hamilton Hargrove	Lawrence Berkeley National Laboratory (LBNL)	1.3.1.17 STPM17- PGAS	Computer Science	230,400
CSC250STPR19	Peter Hugh Beckman	Argonne National Laboratory	2.3.1.19 STPR19- Argo/Power Steering	Computer Science	64,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STPR27	David Bernholdt	Oak Ridge National Laboratory (ORNL)	2.3.1.17 STPR27-OMPI-X: Open MPI for Exascale	Materials Science	96,000
CSC250STSS08	Peter Hugh Beckman	Argonne National Laboratory	STSS08-Argo: Operating System and Resource Management for Exascale	Computer Science	64,000
CSC250STTO09	Jack Dongarra, Anthony Danalis, Earl Carr, Heike Jagode	The University of Tennessee at Knoxville	1.3.2.09 STTO09-ExaPAPI	Computer Science	832,000
CSC250STTO10	Mary Hall	University of Utah	2.3.2.07 STTO10-Autotuning Compiler Technology for Cross-Architecture Transformation and Code Generation	Computer Science	448,000
CSC250STTO11	John Michael Mellor-Crummey	Rice University	2.3.2.08 STTO11 - Extending HPCToolkit for Exascale	Computer Science	832,000
CSC251HIHE05	Scott Pakin	Los Alamos National Laboratory (LANL)	2.4.2.05 HIHE05-Analytical Modeling	Computer Science	364,800
CSC251HISD01		Los Alamos National Laboratory (LANL)	2.4.4.01 HISD01-Software Integration	Computer Science	70,400
Cu_CO2_ethanol		Brookhaven National Laboratory (BNL)	Theoretical Study of Cu-Based Catalysts Ethanol Synthesis	Chemistry	400,000
darkskyml_aesp	Salman Habib	Argonne National Laboratory	Dark Sky Mining	Physics	6,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Databricks	Xiaoyong Jin	Databricks Inc.	Evaluations of Databricks analytics platform	Computer Science	1,000,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	110,000,000
DCOMP-USF	Ivan Oleynik	University of South Florida (USF)	Predictive simulations of dynamic compression of materials	Materials Science	900,000
DD-ALSDFT	Kaushik Bhattacharya	California Institute of Technology (Caltech)	Accelerated Linear Scaling Density Functional Theory: simulating defects in crystalline magnesium	Materials Science	5,000,000
DDICF-Dev	Duc Cao	University of Rochester	Direct-Drive Inertial Confinement Fusion Code Porting and Proposal Preparation	Fusion Energy	2,500,000
debug_dd	Colleen Bertoni	Argonne National Laboratory	Scalability Issue on Theta for ESP	Chemistry	160,000
DetonationEngines	Venkatramanan Raman	University of Michigan	High-fidelity Simulations of Detonation Engines	Engineering	1,000,000
DFT-FE	Phani Sudheer Motamarri	University of Michigan	Large-scale real-space electronic structure calculations for understanding energetics of complex defects in materials	Materials Science	8,000,000
DistrElStructCalcs	Marco Govoni	Argonne National Laboratory	For Everyone A21: Distributed Electronic Structure Calculations Using a Globus-Enabled Programmable Infrastructure	Materials Science	16,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
dist_relational_alg	Sidharth Kumar	The University of Alabama	Distributed relational algebra at scale	Computer Science	1,000,000
DLHMC	James Osborn	Argonne National Laboratory	Deep Learning HMC	Physics	1,100,000
DLILT	Krishnan Raghavan	Argonne National Laboratory	Deep learning for inverse Laplace transform	Computer Science	1,000,000
DMH	Xian-He Sun	Illinois Institute of Technology (IIT)	Utilizing Memory Parallelism for High Performance Data Processing	Computer Science	1,010,000
DMplex_Petsc	Oana Marin	Argonne National Laboratory	Unstructured domain support at scale in PETSc	Engineering	800,000
DNS_spray	Yue Ling	Baylor University	Direct numerical simulation of interfacial wave development and breakup in a gas-liquid mixing layer	Engineering	1,000,000
DNS_wavyWall	Sparsh Ganju	University of Kentucky	Direct Numerical Simulations of Turbulent Channel Flows with Sinusoidally Distributed Roughness	Engineering	1,000,000
DowAcclimation	William Joseph Edsall	The Dow Chemical Company	Dow HPC Team Acclimating to Theta	Chemistry	50,000
drx		Lawrence Berkeley National Laboratory (LBNL)	Local structure in Disordered Rocksalt Cathode Materials - DRX	Energy Technologies	1,000,000
DUNE-LBNF	Nikolai V. Mokhov	Fermi National Accelerator Laboratory (Fermilab)	MARS Energy Deposition and Neutrino Flux Simulations	Physics	5,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
DynamicCS	Huihuo Zheng	Argonne National Laboratory	Dynamic Compressed Sensing for Real-time Tomographic Reconstruction	Materials Science	600,000
Dynstall_DD_1		Iowa State University (ISU)	Characterization of Dynamic Stall Onset	Engineering	50,000
ecp-ci	Kyle Shaver	Argonne National Laboratory	ECP Operations Testing	Computer Science	50,000
ECP_SDK	Sameer Suresh Shende	University of Oregon	Deploying the ECP SDK software stack at ALCF	Computer Science	100,000
EE-ECP	Valerie Taylor, Xingfu Wu	Argonne National Laboratory	Energy efficient tradeoff among execution time, power, and resilience of two ECP applications	Computer Science	3,000,000
EHTModeling	Charles Gammie	University of Illinois at Urbana-Champaign	High Resolution Models of Event Horizon Telescope Sources	Physics	640,000
EngineDNS	Christos Emmanouil Frouzakis	Eidgenössische Technische Hochschule Zürich (ETH Zurich)	Towards reactive DNS in complex internal combustion engine geometries	Engineering	1,500,000
EvoSim	Ying Li	Argonne National Laboratory	Optimizing Deep Learning Models of the ATLAS LHC Detector Response using Evolutionary Algorithms	Physics	2,000,000
EWN-DP-Tool	Getnet Betrie	Argonne National Laboratory	A Tool for Energy-Water-Nexus Analysis: A Deep Learning Framework	Earth Science	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ExaChemPrep	Colleen Bertoni	Argonne National Laboratory	Benchmarking and Characterization of Electronic Structure Codes in Preparation for Exascale	Chemistry	700,000
EXASTEEL-THETA	Axel Klawonn, Oliver Rheinbach	TU Bergakademie Freiberg	Simulation of Dual-Phase Steel on Theta	Engineering	1,000,000
EXATHERM	David Bross	Argonne National Laboratory	Automated Generation of Customized Quantum Chemical Composite Methods Validated via the ATcT Database	Chemistry	1,000,000
ExM	Justin Michael Wozniak, Michael Joseph Wilde	Argonne National Laboratory	Extreme Many-task Computing with Swift	Computer Science	25,000
FFTBench	William Scullin	Argonne National Laboratory	FFT Benchmarking for Imaging Applications	Computer Science	2,250,000
fhi_etc1p	Thomas Alexander Reichmanis Purcell	Fritz-Haber-Institut der Max-Planck-Gesellscha ft (FHI)	Benchmarking a High-Throughput Framework for the Thermal Conductivity of Perovskites	Physics	3,500,000
FLoRIN	Walter Scheirer, Jeffery Dwayne Kinnison	University of Notre Dame	Large-Scale Segmentation and Reconstruction of Neural Microscopy	Biological Sciences	3,000,000
FreeEnMemPro_DD	Benoit Roux	Argonne National Laboratory	Free Energy Landscapes of Membrane Proteins	Biological Sciences	5,000,000
fusiondl_aesp	William Tang	Princeton University	Accelerated Deep Learning Discovery in Fusion Energy Science	Fusion Energy	2,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
G3DNet	John Taylor Childers	Argonne National Laboratory	3D GraphCNN for Particle Detectors	Computer Science	2,000,000
GA-bfg3d	Mark Kostuk	General Atomics	General Atomics Next Generation Skew-Symmetric Fluid Solver	Physics	1,250,000
GFDL_Ensemble	Thomas Edward Robinson	Geophysical Fluid Dynamics Laboratory (GFDL)	Ensemble-based Regression Tests for Climate Model Reproducibility	Earth Science	1,802,732
HACC_aesp	Katrin Heitmann	Argonne National Laboratory	Extreme-Scale Cosmological Hydrodynamics	Physics	1,000,000
HEPCloud-FNAL	Burt Holzman	Fermi National Accelerator Laboratory (Fermilab)	High Energy Physics Computing for Fermilab experiments via HEPCloud	Physics	500,000
HEP_on_HPC	Jim B Kowalkowski, Marc Francis Paterno	Fermi National Accelerator Laboratory (Fermilab)	HEP analysis workflows on HPC	Physics	1,000,000
hifiturbfsi	Ivan Bermejo Moreno	University of Southern California (USC)	High-fidelity simulation of supersonic turbulent flow-structure interaction and mixing	Engineering	4,000,000
high-beta_kinetics	Matthew Kunz	Princeton University	Multi-scale Dynamics of Kinetic Turbulence in Weakly Collisional, High-Beta Plasmas	Physics	256,000
HIV-PR	Ao Ma	University of Illinois at Chicago	Understanding the mechanism of ligand-induced conformational dynamics of HIV-1 protease and the effects of mutations	Biological Sciences	2,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
HNPballistics	Sinan Keten	Northwestern University	Engineering Nanocellulose based Hairy Nanoparticle Assemblies for High Ballistic Impact Performance	Engineering	2,000,000
HPC4SBND	Andrzej Michał Szelc	The University of Manchester	High Perfomance Computing for the Short Baseline Near Detector at Fermilab	Physics	500,000
hpcbdsm	Tanwi Mallick	Argonne National Laboratory	High-Performance Computing and Big Data Solutions for Mobility Design and Planning	Computer Science	2,000,000
HTQMC	Can Ataca	University of Maryland, Baltimore County	Accurate Simulations of Low Dimensional Materials for Energy and Device Applications	Materials Science	1,000,000
hydro	Eugene Yan	Argonne National Laboratory	Deep learning solution for flood forecasting	Earth Science	1,000,000
IBR2M	Alberto Talamo	Argonne National Laboratory	Transient Simulation of IBR2M	Nuclear Energy	3,000,000
ICE_Nek5000	Muhsin Mohammed Ameen	Argonne National Laboratory	Scalable Internal Combustion Engine Simulations using Nek5000	Energy Technologies	5,000,000
IL_elec	Zhengcheng Zhang	Argonne National Laboratory	Ionic liquid as a potential electrolyte of high performance lithium ion battery	Chemistry	8,000,000
ImageNanoX_DD	Phay Ho	Argonne National Laboratory	X-ray Imaging of Transient Structure and Dynamics of Nanoparticles	Physics	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
INSPIRE	Peter Coveney	University College London (UCL)	INtegrated and Scalable PredictIon of REsistanace	Chemistry	19,200,000
Intel	Kalyan Kumaran, Scott Parker, Timothy Williams, Venkatram Vishwanath	Argonne National Laboratory	Intel employees in support of Theta	Internal	10,000,000
interconnect_bench	Devi Sudheer Kumar Chunduri	Argonne National Laboratory	Collaboration with Cray on Interconnect related studies	Computer Science	400,000
iRF-LOOP	Daniel Jacobson	Oak Ridge National Laboratory (ORNL)	Explainable-AI Applications for Systems Biology	Biological Sciences	500,000
Job_Interference	Zhiling Lan	Illinois Institute of Technology (IIT)	Workload Interference Analysis on Theta	Computer Science	2,500,000
K-Ras-Model	Sunhwan Jo	Qulab, Inc.	Dynamic Model of K-Ras WT and Mutant	Biological Sciences	1,000,000
KineticTurbulence	Jason TenBarge	Princeton University	A Discontinuous Galerkin Fully Kinetic Vlasov Maxwell Study of Plasma Turbulence	Physics	2,000,000
Landscapes	Michael Prentiss	University of Cambridge	Exploring Protein Folding Energy Landscapes	Biological Sciences	1,000,000
Large-ScalePF	Jiamian Hu	University of Wisconsin-Madison	Phase-field Simulations and Atomistic Modeling of Large-Scale Spin-Lattice Dynamics	Materials Science	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
large3dxrayADSP	Chris Johnson Jacobsen	Northwestern University	X-ray microscopy of extended 3D objects: scaling towards the future	Biological Sciences	28,000,000
LatticeQCD_aesp	Paul Mackenzie, Norman Howard Christ	Fermi National Accelerator Laboratory (Fermilab)	Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics	Physics	1,000,000
LHCATLAS-scaling	Douglas Paul Benjamin	Argonne National Laboratory	Scaling of LHC ATLAS fine grain simulation of p-p collisions	Physics	1,000,000
LIGHTCONTROL	Sandra Biedron	University of New Mexico	Light sources and their control using AI techniques	Physics	100,000
LQCDdev	James Osborn	Argonne National Laboratory	Lattice QCD development	Physics	2,199,935
lqcdml_aesp	William Detmold	Massachusetts Institute of Technology (MIT)	Machine Learning for Lattice Quantum Chromodynamics	Physics	14,000,000
LQCD_VeloC	Chulwoo Jung	Brookhaven National Laboratory (BNL)	Lattice QCD with VeloC	Physics	1,583,000
LSSTADSP_DESC	Katrin Heitmann	Argonne National Laboratory	Realistic Simulations of the LSST Survey at Scale - DESC	Physics	14,000,000
LSSTADSP_HACC	Katrin Heitmann	Argonne National Laboratory	Realistic Simulations of the LSST Survey at Scale - HACC	Physics	2,000,000
magjet	Edison Liang	Rice University	Magnetized jet creation using a hollow ring of laser beams	Physics	2,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Magnetoconvection	Oleg Zikanov	University of Michigan	DNS of Rayleigh-Benard convection in a very strong magnetic field	Engineering	1,000,000
magstructsADSP	Trevor David Rhone, Prof. Kaxiras	Rensselaer Polytechnic Institute (RPI)	Machine learning magnetic properties of van der Waals heterostructures	Materials Science	36,000,000
Maintenance	Mark R. Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	10,000,000
matml_aesp	Noa Marom	Carnegie Mellon University	Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials	Materials Science	24,000,000
MCPSI_DD	Ahren Jasper	Argonne National Laboratory	High temperature kinetics and dynamics of fluxional molecules	Chemistry	1,500,000
mep	Whitney Armstrong	Argonne National Laboratory	Medium Energy Physics Data Processing	Physics	1,000,000
MERF_Data_Layer	Santanu Chaudhuri	Argonne National Laboratory	MERF Data Layer	Materials Science	2,000,000
MIT-MSEI-Partnership	Santanu Chaudhuri	Argonne National Laboratory	MIT-Argonne Manufacturing Collaboration	Materials Science	1,000,000
ML-ACTAR	Sandeep Madireddy	Argonne National Laboratory	Machine learning-based Adaptive-Congestion & Topology Aware Routing	Computer Science	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ML-GA	Pinaki Pal	Argonne National Laboratory	Machine Learning Workflow Tools for Rapid Optimization of Product Designs and Manufacturing Processes	Energy Technologies	2,000,000
ML4MPF	Gina Maureen Sforzo	Argonne National Laboratory	Enabling predictive simulations of reacting multiphase flows via data-driven emulation	Engineering	2,000,000
MLPerfHPC	Murali Krishna Emani	Argonne National Laboratory	Evaluation of Scientific ML Benchmarks on HPC Systems	Computer Science	1,000,000
mmaADSP	Eliu Huerta Escudero, Elise Jennings	University of Illinois at Urbana-Champaign	Deep Learning at Scale for Multimessenger Astrophysics through the NCSA-Argonne Collaboration	Physics	20,000,000
MOAB_App	Vijay Subramaniam Mahadevan	Argonne National Laboratory	MOAB Algorithmic Performance Portability	Mathematics	1,799,984
ModelingSoftMatter	Zhangli Peng	University of Illinois at Chicago	Multiscale Modeling of Soft Materials by Unifying Shell Elements and Anisotropic Coarse-Grained Force Fields in LAMMPS	Biological Sciences	1,000,000
Monochalcogenides	Salvador Barraza-Lopez	University of Arkansas	Structural phase transitions and catalitic properties of group-IV monochalcogenide monolayers	Materials Science	2,000,000
MOVES_Matrix	Randall Guensler	Georgia Institute of Technology (Georgia Tech)	MOVES-Matrix Interagency Collaboration	Earth Science	1,345,500
MPICH_MCS	Kenneth James Raffenetti, Pavan Balaji	Argonne National Laboratory	MPICH - A high performance and widely portable MPI implementation	Computer Science	20,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
MPI_Aurora_Intel	Devi Sudheer Kumar Chunduri	Argonne National Laboratory	MPI development for Aurora	Computer Science	100,000
MSAS_DD	Frank Xavier Vazquez	St. John's University	Multiscale Investigation of a-synuclein structural dynamics	Chemistry	2,000,000
mu2e_CRY	Lisa Goodenough	Fermi National Accelerator Laboratory (Fermilab)	CRY Cosmic Ray Simulation for Mu2e	Physics	1,500,000
multiphysics_aesp	Amanda Randles	Duke University	Extreme-scale In Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations	Engineering	10,000,000
MultiscaleADSP	Rathakrishnan Bhaskaran, Marc Edgar	General Electric Company (GE)	Enabling Multi-Scale Physics for Industrial Design using Deep Learning Networks	Engineering	6,000,000
n2ase	Graham Donald Fletcher	Argonne National Laboratory	Valence bond study of nitrogenase	Chemistry	2,000,000
NAMD_aesp	Benoit Roux, James C. Phillips	The University of Chicago (UChicago)	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	1,000,000
NAQMC_RMD_aesp	Aiichiro Nakano	University of Southern California (USC)	Metascalable Layered Materials Genome	Materials Science	1,000,000
Nek_VHTR	Masahiro Kawaji, Aleksandr V. Obabko	City College of New York	Scalability and Validation of Nek5000 for VHTR challenge problem of pipe flow relaminarization	Engineering	500,000
NERSC_DD_Project	Sreeranjani Ramprakash, Haritha Siddabathuni Som	Argonne National Laboratory	NERSC DD project at ALCF	Support	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
networkbench	Devi Sudheer Kumar Chunduri, Elise Jennings, Kevin Harms	Argonne National Laboratory	network benchmarking and modeling	Computer Science	200,000
NetworkScience	Mehmet H Gunes	University of Nevada, Reno	Analysis of complex networks on HPC: A case study of the Internet	Computer Science	2,800,000
neutrinoADSP	Andrzej Michał Szelc, Corey Adams, Diego Garcia-Gamez	The University of Manchester	Developing High Performance Computing Applications for Liquid Argon Neutrino Detectors	Physics	3,000,000
Next	Paola Ferrario	Donostia International Physics Center	NEXT	Physics	1,000,000
NWChemEx_aesp	Alvaro Vazquez Mayagoitia	Pacific Northwest National Laboratory (PNNL)	NWChemEx: Tackling Chemical, Materials & Biochemical Challenges in the Exascale Era	Chemistry	1,000,000
OctopusNeuro	Thomas David Uram	Argonne National Laboratory	Octopus Neuroscience	Biological Sciences	1,000,000
OF_ICING	Roberto Paoli	University of Illinois at Chicago	Aircraft icing simulations with OpenFoam	Engineering	1,000,000
OF_SCALING	Roberto Paoli	University of Illinois at Chicago	Scaling analysis of a new time-accurate OpenFOAM solver	Engineering	1,500,000
Operations	Mark R Fahey	Argonne National Laboratory	Systems administration tasks	Internal	10,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
pearlnecklace	Michael David Schneider	Lawrence Livermore National Laboratory (LLNL)	Pearl Necklace	Physics	20,000,000
PENA	Gabriel Staffelbach	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS)	Performance evaluation of new algorithms in AVBP	Chemistry	2,000,000
Performance	Scott Parker, Raymond M. Loy	Argonne National Laboratory	Performance	Internal	70,000,000
perf_research	Devi Sudheer Kumar Chunduri	Argonne National Laboratory	Performance group external facing research	Computer Science	200,000
PETScKrylov	Hannah Morgan, Todd S. Munson	Argonne National Laboratory	The influence of noise on Krylov methods	Computer Science	250,000
PhaseIdentification	Jianguo Wen	Argonne National Laboratory	Deep Learning for TEM Phase Identification	Materials Science	1,000,000
PHASTA_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	1,000,000
PHASTA_NCSU	Igor A Bolotnov	North Carolina State University (NCSU)	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	4,000,000
pieiclat	Ian Cloet	Argonne National Laboratory	Pion and Kaon Quark-Gluon Structure from Lattice QCD	Physics	4,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
plume_dns	Andrew Poje	The City University of New York (CUNY)	Direct Numerical Simulation of Near and Medium Field Evolution of Oil Plumes in the Oceanic Environment	Engineering	2,000,000
Polymersome19		The City University of New York (CUNY)	Multiscale Modeling of Polymersomes	Chemistry	888,960
PrincetonConnectome	Sebastian Seung	Princeton University	Petascale neural circuit reconstruction	Biological Sciences	26,000
PROTEUS_KNL	Yeon Sang Jung	Argonne National Laboratory	Porting high fidelity neutron transport code PROTEUS-MOC to KNL Platform	Nuclear Energy	1,000,000
psr001	Ronald Grover	General Motors Company	Electric Motor Thermal Management Analysis	Engineering	5,000,000
PtychoHPC	Junjing Deng	Argonne National Laboratory	High performance computation for ptychography reconstruction	Materials Science	1,000,000
qcarchive	Daniel G. A. Smith	The Molecular Sciences Software Institute	The MolSSI Quantum Chemistry Archive Project	Chemistry	1,000,000
QCSim	Hal Finkel	Argonne National Laboratory	Simulating Realistic Quantum Computers	Computer Science	16,000,000
QMC-EFMO	Federico Zahariev	Iowa State University (ISU)	Linearly scaling QMC through EFMO fragmentation	Chemistry	2,000,000
QMC-model	Lucas Kyle Wagner	University of Illinois at Urbana-Champaign	Quantum Monte Carlo derivation of effective Hamiltonians	Physics	100,002

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
QMCPACK_aesp	Anouar Benali	Argonne National Laboratory	Extending Moore's Law computing with Quantum Monte Carlo	Materials Science	1,000,000
QMC_ankit	Ankit Mahajan	University of Colorado-Boulder	Study of strongly correlated electrons	Chemistry	999,100
Quantum_ML	Yuri Alexeev	Argonne National Laboratory	Quantum Machine Learning	Mathematics	2,000,000
radix-io	Philip Carns	Argonne National Laboratory	System software to enable data-intensive science	Computer Science	1,000,000
RaptorX	Jinbo Xu	Toyota Technological Institute at Chicago (TTIC)	Protein Folding through Deep Learning and Energy Minimization	Biological Sciences	2,000,000
rare-events-dhuz	Amir Haji-Akbari	Yale University	Computational Investigation of Collective Phenomena in Soft Matter and Biological Systems	Materials Science	1,500,000
REI_Flares	Marc Cremer	Reaction Engineering International	Leveraging the UCF for Simulation of Industrial Flares	Chemistry	1,000,000
RL-fold	Arvind Ramanathan	Oak Ridge National Laboratory (ORNL)	Targeting intrinsically disordered proteins using artificial intelligence driven molecular simulations	Biological Sciences	1,000,000
SBC-UBEM	Ralph T Muehleisen	Argonne National Laboratory	Scalable Bayesian Calibration of Urban-Scale Building Energy Models	Energy Technologies	200,000
scalablemoose	Fande Kong	Idaho National Laboratory (INL)	MOOSE	Nuclear Energy	3,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SCO_thermal_conduct	Bilge Yildiz	Massachusetts Institute of Technology (MIT)	Thermal conductivity modulated by oxygen and hydrogen ions in SrCoOx	Materials Science	100,000
SCRAMJET_VULCAN	Roberto Paoli	Argonne National Laboratory	Scramjet simulations on NASA VULCAN	Engineering	1,000,000
SDL	Min Si	Argonne National Laboratory	Scalable deep learning	Computer Science	4,000,000
SDL_Workshop	Raymond M. Loy	Argonne National Laboratory	ALCF Simulation, Data, and Learning Workshop	Training	10,000,000
SeismicHazard_2	Christine Goulet	University of Southern California (USC)	Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling	Earth Science	7,680,000
SENSEI	Silvio Rizzi, Joseph Insley, Nicola Ferrier, Venkatram Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	5,820,067
shockturb	Johan Larsson	Stanford University	High-fidelity calculations of shock/turbulence interaction	Engineering	3,000,000
shockturb2	Johan Larsson	University of Maryland	High-speed compressible turbulent flows with shock waves	Engineering	3,000,000
SimMCSProce	Andreas Franz Prein	The National Center for Atmospheric Research (NCAR)	Using ARM Observations to Evaluate Process-Interactions in MCS Simulations Across Scales	Earth Science	400,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Singlet_Fission	Hanning Chen	The George Washington University	Probing the Free Energy Surface of Spin Separation in Singlet Fission	Chemistry	6,000,000
SI_IBM	Saurabh Chawdhary	Argonne National Laboratory	Scalable, efficient sharp-interface immersed boundary method for fluid-structure interaction problems	Engineering	1,000,000
SolarWindowsADSP	Jacqueline Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-Powered Windows	Materials Science	32,000,000
SolarWindTurb	Jean Carlos Perez	Florida Institute of Technology	Understanding Solar Wind Turbulence in the inner Heliosphere	Physics	1,000,000
Solv_extraction	Lynda Soderholm	Argonne National Laboratory	Liquid-liquid extraction from solvent extraction	Chemistry	4,000,000
SPMV_Scalablity		University of Notre Dame	Strong Scalability of Hyrbrid Sparse Problems	Computer Science	20,000
SPS	Vadim Roytershteyn	Space Science Institute	Spectral Kinetic Simulations of Plasmas	Physics	100,000
SSB_cathode	Gerbrand Ceder	University of California-Berkeley	Optimization of Composite Cathode in All-Solid-State Batteries	Materials Science	200,000
Stable-MABs	Deniz Cakir	University of North Dakota	The Stable MAB and MBenes for Spintronics and Battery Applications	Physics	750,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SU2_PadeOps_aesp	Sanjiva K Lele	Stanford University	Benchmark Simulations of Shock-Variable Density Turbulence and Shock-Boundary Layer Interactions with Applications to Engineering Modeling	Engineering	1,000,000
SweptWing_DD	Lian Duan	Missouri University of Science and Technology	Laminar-Turbulent Transition in Swept-Wing Boundary Layers	Engineering	2,000,000
TNContract	James Osborn	Argonne National Laboratory	Tensor Network Contractions for QIS	Physics	1,000,000
TomoDev	William Scullin	Argonne National Laboratory	Tomographic Software Development	Computer Science	750,000
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	2,000,000
TotalView	Peter Thompson, Raymond M. Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene P	Internal	100,000
TraceX	Tekin Bicer	Argonne National Laboratory	Time Sensitive Analysis of Synchrotron Light Source Data Streams	Computer Science	600,000
Transonic_Buffet	Kai Mbali Kruger Bastos	Duke University	The Buffet Phenomenon in Subsonic, Transonic and Hypersonic Flow Regimes	Engineering	2,401,022
TurbLowBeta	Luca Franci	Queen Mary University of London	Hybrid simulations of kinetic plasma turbulence at low electron beta	Physics	8,256,478

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
TurbNet	Romit Maulik	Argonne National Laboratory	TurbNet: Scaleable physics-informed deep learning for turbulence model development	Engineering	3,000,000
UA_TRANSITION	Hermann Fasel	The University of Arizona	Direct Numerical Simulations of Boundary-Layer Transition	Engineering	2,500,000
uboone		Fermi National Accelerator Laboratory (Fermilab)	MicroBooNE physics data processing	Physics	1,000,000
UINTAH_aesp	Martin Berzins, John Andrew Schmidt	University of Utah	Design and evaluation of high-efficiency boilers for energy production using a hierarchical V/UQ approach	Chemistry	1,000,000
ultrafast	Jin Wang	Argonne National Laboratory	Developing High-Fidelity Simulation and Ultrafast X-Ray Imaging Tools for the APS-Upgrade	Engineering	2,000,000
UltrafastINCITE	Aiichiro Nakano	University of Southern California (USC)	INCITE proposal: Ultrafast Control of Functional Materials	Materials Science	2,000,000
ultrafastxrayADSP	Jin Wang	Argonne National Laboratory	Developing High-Fidelity Dynamic and Ultrafast X-Ray Imaging Tools for APS-Upgrade	Physics	8,000,000
UMass_MultiPhaseFlow	David Schmidt	University of Massachusetts-Amherst	Fuel Spray Modeling and Simulation of Wave Loads on Offshore Wind Turbines	Engineering	1,226,707
US3DonTheta	Graham Vardy Candler	University of Minnesota	Porting the US3D Hypersonics Code to Theta	Engineering	1,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
User_Services	Haritha Siddabathuni Som, Sreeranjani Ramprakash	Argonne National Laboratory	User Services	Internal	20
VariantToMinibrain	Win Drongelen	The University of Chicago (UChicago)	Computational selection of genetic variants underlying epilepsy	Biological Sciences	200,000
VeloC	Bogdan Nicolae	Argonne National Laboratory	VeloC: Very Low Overhead Checkpointing System	Computer Science	110,000
visualization	Joseph Insley, Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	2,000,000
Viz_Support	Joseph Insley, William Edward Allcock	Argonne National Laboratory	Visualization Support	Computer Science	2,000,000
VST-sode	Olaseni Sode	California State University (CSU), Northridge	Development of vibrational structure techniques for van der Waals and hydrogen-bonded complexes	Chemistry	500,000
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	2,000,000
WaterHammer	Hong Zhang	Argonne National Laboratory	Water Hammer Simulation	Mathematics	20,000
WGSanalysis	Elizabeth McNally	Northwestern University	Large scale alignment and analysis of whole human and mouse genomes, with focus on realigning to HG 38 and harmonized workflows	Biological Sciences	500,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
wisc_atlas	Sau Lan Wu	University of Wisconsin-Madison	Apply quantum machine learning to High Energy Physics	Physics	800,000
XFELO	Henry Freund	University of New Mexico	Simulation of x-ray FEL Oscillators	Physics	5,000,000
XGC_aesp	Choongseok Chang, Timothy Williams	Princeton Plasma Physics Laboratory (PPPL)	High fidelity simulation of fusion reactor boundary plasmas	Fusion Energy	1,000,000
yakushin	Igor Yakushin	The University of Chicago (UChicago)	Using Deep Learning for Gravitational Lens classification and parameter regression	Physics	100,000
				Total Theta DD	**1,084,713,360**

Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Lemont, IL 60439-4832

www.anl.gov

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