ARGONNE **LEADERSHIP COMPUTING** FACILITY

2014 **Operational Assessment Report**

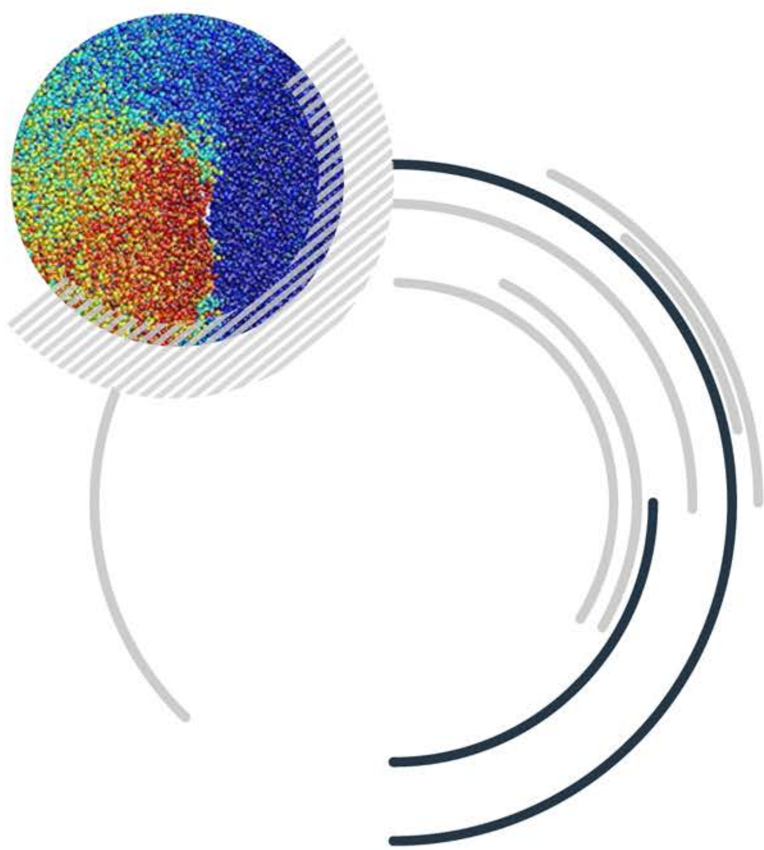

On the cover

Visualization taken from an 8-million-atom molecular dynamics simulation of a shockwave interacting with a 20-nm cylindrical void in a molecular crystal of high-explosive PETN, causing the void to collapse.

2014–2015 ALCC Project: Nanostructure-Enhanced Chemical Reactivity and Detonation in Energetic Materials
PI: Aidan Thompson, Sandia National Laboratories
Image credit: Tzu-Ray Shan and Aidan Thompson, Sandia National Laboratories

Contents

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Figures

Figures (Cont.)

Figures (Cont.)

Figures (Cont.)

Tables

Tables (Cont.)

Executive Summary

As one of two DOE Leadership Computing Facility (LCF) centers in the nation for open science, the Argonne Leadership Computing Facility (ALCF), supported by the DOE Advanced Scientific Computing Research (ASCR) Program, provides the computational science community with world-class computing capabilities, expertise, and assistance to ensure that every project achieves top performance on its resources.

In December, ALCF completed its first year of operations with Mira, a 48-rack IBM Blue Gene/Q system. ALCF also operated two smaller BG/Q computing resources: a 4-rack system (Cetus) that shares Mira's software stack and file systems and is used for tool and application porting, software testing and optimization, and systems software development; and a 2-rack system (Vesta) that is used for testing new versions of software prior to installation on Mira. Analysis and visualization are done on Tukey, a companion visualization cluster. Tukey shares the Mira network and parallel file system, enabling direct access to Mira-generated results.

During the past year, ALCF also decommissioned its ALCF-1 resource, Intrepid, and its two development systems. Several of the IBM BG/P racks found a new home at North Carolina State University, where they are being used in teaching and for research activities on campus. Intrepid's file system has been repurposed too, as both a testbed for future ALCF systems and for exploring data sharing with a wider audience than just ALCF users. The rest of the machine was processed and recycled.

ALCF has proudly met or exceeded all metrics set for the facility. ALCF delivered 5.8 billion core-hours of compute time in 2014 between January 1 and December 31, with 3.5 billion of those core-hours being used by capability jobs. In the same period of time, the science done on the machine produced more than 150 publications, in all major areas of interest to DOE.

LCF resources continue to address the computational and data science problems that the scientific community deems critical to the advancement of science and of the most benefit to the nation. Last year ALCF supported 1,432 DOE-defined users and engaged in more than 342 active projects from universities, national laboratories, and industry.

The annual Operations Assessment review of ALCF by ASCR provides the facility with an opportunity to receive external feedback on ways to improve the operation of the facility. The review takes into consideration agreed-upon metrics and reports describing the operation of the facility. The report is organized into eight sections. These sections address the 2014 OAR metrics and present User Support Results, Business Results, Strategic Results, Innovation, Risk Management, Safety, Cyber Security, and a Summary of the Proposed Metric Values for Future OARs.

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

Area	Metric	2014 Target	2014 Actual
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.5/5.0
	User Survey – **Response Rate**	25.0%	30.0%
	% User Problems Addressed within Three Working Days	80%	96.0%
Business Results	Mira Overall Availability	90.0%	95.7%
	Mira Scheduled Availability	95.0%	98.7%
	% of INCITE core hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	30%	64.5%
	% of INCITE core hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10%	33.1%

Responses to Recommendations from the Previous OA Review (2013 OAR conducted in Spring 2014)

1. Are the processes for supporting the customers, resolving problems, and outreach effective?

Comments:

- The efforts reported by ALCF to address users/customers' needs ensuring an effective utilization of the facilities is commendable. Improvements to their existing methods for communication, outreach, and engagement of existing and potential users are proactively pursued and manifest in a number of innovations that find direct application to several communication channels. AGREED
- Catalyst teams are an outstanding resource that provide the highest standard in terms of communication channels to support the users. In fact, their mission transcends those supporting activities and extends toward active collaboration. These activities not only drive the selection of INCITE proposals but also have a direct impact on the hiring process of new catalysts, based in part on the areas of knowledge in high demand according to users' requests. AGREED – See Appendix D.
- The implementation of checklist models and techniques looks interesting. Sharing this approach with other facilities could provide the immediate benefit. AGREED – ALCF will write a short summary to share.
- Synergies at the user and operation level with BG/Q systems in other National Labs (e.g. Sequoia, Vulcan at LLNL) could be leveraged (for example, to minimize impact of common problems with defective IBM driver rollouts). AGREED – We already have active collaborations with the BG/Q community, though informally, and via the Blue Gene Consortium.

Recommendations:

None

2. **Is the Facility maximizing the use of its resources consistent with its mission?**

   Comments:

   - ALCF met or exceeded all targets while both operating Intrepid and implementing Mira. AGREED
   - INCITE and ALCC are out of phase by 6 months. This is beneficial because it helps distribute onboarding activities throughout the year and causes end-of-year computational spikes to occur at different times. AGREED
   - Recruiting and retaining staff needs to remain a priority. AGREED

   Recommendations:

   - ALCF would benefit from more extensive and formal control over the TCS datacenter to ensure continued success of its mission, e.g.:
     - Formal SLA/MOU with the Trust, FMS, and other infrastructure support organization and co-tenants. AGREED – We have a formal SLA/MOU in place with the Trust, FMS, and co-tenants. ALCF participates in weekly/monthly meetings with all datacenter stakeholders (SLA/MOU attached).
     - Assurance of regularly scheduled and budgets preventive maintenance for electrical and mechanical systems. AGREED – We have standing maintenance contracts in place that are managed by the TCS Trust (e.g., Eaton performs preventive maintenance on ALCF circuit breakers).
     - Control of allocation of space, power/cooling, and changes. AGREED – As the major tenant of the TCS datacenter, ALCF's needs are prioritized and met first as generally accepted policy. Furthermore, the ALD office controls this space and ALCF's division director is the ALD office member in charge of space.

3. **Is the Facility enabling scientific achievements consistent with the Department of Energy strategic goals?** (Specifically applicable to Goal 2: "Maintain a vibrant U.S. effort in science and engineering as cornerstone of our economic prosperity with clear leadership in strategic areas." Goal 2 includes the targeted outcome: "Continue to develop and deploy high-performance computing hardware and software systems through exascale platforms." Sites may also include contributions to other goals and other targeted outcomes.)

   Comments:

   - The allocations structure is well aligned to the DOE strategic goals and has resulted in successful scientific results from each type of allocation. AGREED
   - The processes and priorities for the different types of allocations and support have been successfully implemented. AGREED
   - The user support model for each type of allocation is consistent with the intended goals and is being successfully applied in support of the projects. AGREED
   - The Catalyst program, in which each person works with about three INCITE projects, is very effective in advancing the research goals of the projects for each type of allocation. AGREED
   - Support efforts include algorithmic, porting and testing, optimization and tuning, and visualization. AGREED

- Long term (post award) follow up of results and effects is very important and should continue to be pursued and should be publicized. AGREED

Recommendations:

- The requirements to successfully support future ALCC allocations should be carefully evaluated, including technical readiness and staff support. AGREED, but the selection of ALCC projects is done by ASCR, so ALCF is very much in a reactive mode. ALCF will continue to support the ALCC users to the best of its ability.
- The Catalyst program should be fully staffed. AGREED – ALCF is actively working to fully staff the Catalyst program.
- Tracking, forecasting, and planning for potential new areas of science (e.g. geosciences) and technology (I/O and storage) should be enhanced to help shape planning for ALCF and for exascale. AGREED – This is a part of critical decision process. In addition, the ALCF is establishing a new annual publication to track user requirements and progress. The ALCF Application Book will track the generalization of domain/community codes, and user needs and achievements.

4. Have Innovations been implemented that have improved facility operations? This includes innovations adopted from, recommended to, or adopted by other facilities.

Comments:

- The installation of the secondary cooling loop supporting Mira is an excellent example of facilities/engineering best practice. This secondary loop significantly reduces dependencies on the plant-wide chilled water chemistry, temperature, and filtration characteristics, with a minimal investment and minimal impact to operating costs. AGREED
- Systems Management and Monitoring Tools (Application Performance Collection and Machine Time Overlay) are expected to yield valuable dividends in ensuring timely identification and remediation of both system- and user-related issues. AGREED
- The committee suggests that ALCF seek opportunities to leverage the tools/innovation/best practices work by other BG/Q sites, and to share the tools/innovations/best practices that ALCF has developed. AGREED – ALCF regularly talks with LLNL and other BG/Q sites.

Recommendations:

None

5. Is the Facility effectively managing risk?

Comments:

- The committee agrees that ALCF's restructuring of their risk management system is well considered and will lead to more effective management of key risks. AGREED

Recommendations:

- Staff overload, and thus recruitment and retention, is a major concern. ALCF management has done a good job reprioritizing and replanning work and retasking staff to meet key operational and project need. But in order to remedy the staffing situation, we believe that management needs to emphasize recruiting efforts over operational issues that can be deferred. ALCF should not compromise its hiring standards to meet any specific staff level goals. AGREED – ALCF hires the best people and never compromises its hiring standards to meet operational goals. While individual staff members may be overloaded at times, the division as a whole is committed to preserving these standards and to postpone hiring until the right person is found for every job.

6. **Has the Facility incorporated appropriate measures to protect the safety of workers and the public?**

   Comments:

   ALCF does not have ownership for safe operation of all elements that affect their operation or involve their people and should assert leadership over critical infrastructure. AGREED

   Recommendations:

   None

7. **Does the site have a valid cyber security plan and authority to operate?**

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

   Comments:

   - Capability metric is only against INCITE allocation, not all users. This is different from OLCF and NERSC. Consider working with ASCR and the other centers to develop a common definition. Neutral, we will ask ASCR for their opinion on this matter. Applying a capability metric to discretionary allocations may not make sense based on the role that these projects serve in the facility. The same can be said for ALCC allocations as "high-risk, high-reward" projects that undergo minimal readiness evaluation.
   - Raw v. allocated v. available v. used hours can be confusing. With ASCR, etc., report this information in a form that is more transparent. Neutral, we will ask ASCR for their opinion on this matter.

   Recommendations:

   None

9. **What is your overall assessment of the Facility Operational performance?**

   Additional Comments:

   The OA review committee commends the staff for their skill and commitment. Their efforts are enabling great science to be done. AGREED

This page intentionally left blank.

Section 1. User Support Results

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

ALCF Response

ALCF has established processes in place for effectively supporting customers, resolving problems, and performing outreach. The 2014 survey measures satisfaction, user support, and problem resolution. It thereby provides input to ALCF about where improvements can be made (Table 1.1). The sections below document ALCF events and processes, the effectiveness of those processes, and what improvements to those processes were implemented during calendar year (CY) 2014.

Table 1.1 All 2014 User Support Metrics and Results[1]

		2013 Actual	2014 Target	2014 Actual
Number Surveyed		1,150[2]	N/A	1,432[2]
Number of Respondents (Response Rate)		364 (31.7%)	25.0%	430 (30.0%)
Overall Satisfaction	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.5
	Std Dev	0.7	N/A	0.7
Problem Resolution	Mean	4.6	3.5	4.5
	Variance	0.5	N/A	0.5
	Std Dev	0.7	N/A	0.7
User Support	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.3
	Std Dev	0.7	N/A	0.5
% User Problems Addressed within Three Working Days		91.9%	80%	96.0%

[1] It is important to note that the definition of a user at a DOE facility has evolved over time and is not identical for each year. In past years, opt-outs and bounces were not included in the number surveyed. This year these two groups are part of the number surveyed based upon a final, standard definition of a DOE user. If the 2014 "number surveyed" is adjusted to account for this difference, the number surveyed changes to 1,304 and the response rate changes to 33.0 percent.

[2] The definition of an official user was made consistent across DOE user facilities by the program office. The timeframe for an official user is now reported as a fiscal year as opposed to a calendar year. In addition, staff members and vendors are considered users if they are members of a reviewed project.

Survey Approach

ALCF contracted with survey experts from Cvent, a web survey hosting and consulting company, to manage the 2014 survey. The team incorporated lessons learned from previous surveys and internal feedback from various ALCF teams, ALCF leadership, the User Advisory Council, and Advanced Scientific Computing Research (ASCR). The result was a streamlined survey, improved questions, and a representative user response to the survey.

Likert Scale and Numeric Mapping

Almost all Likert Scale questions in the ALCF user survey use a six-choice scale (for rating user responses). This is a standard for surveys because 1) it is easy for users to quickly place the response to a question within a range of options, 2) it can be mapped to a numeric scale and, 3) given a certain sample size, it can be used with a normal distribution. The method allows for use of off-the-shelf statistics functions to determine variance and standard deviation.

ALCF 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)

Some questions were not well suited to the six-point scale and a different scale was therefore used. The only question included as part of the OAR to which this applied was the overall satisfaction question. It used the following five-choice scale:

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

A non-metric question was revised on the 2014 User Survey that used the six-point scale below:

Statement	Numeric
Extremely Satisfied	5
Somewhat Satisfied	4
Neither	3
Somewhat Dissatisfied	2
Extremely Dissatisfied	1
N/A	(No Value)

1.1 User Support Metrics

In 2014, 1,432 individual ALCF users met the DOE user definition[3] and were invited to complete a user survey. Of those users, 430 responded for a 30 percent response rate, quite good compared to generally accepted standards for survey response rates. ALCF surpassed all targets for the survey metrics and there was no statistically significant change from 2013 results.

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

Table 1.2 2014 User Survey Results by Allocation Program

2014 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		591	148	739	693	1,432
Number of Respondents		195	46	241	189	430
Response Rate		33.0%	31.1%	32.6%	27.3%	30.0%
Overall Satisfaction	Mean	4.5	4.3	4.4	4.5	4.5
	Variance	0.5	0.8	0.5	0.5	0.5
	Std Dev	0.7	0.9	0.7	0.7	0.7
User Support	Mean	4.5	4.4	4.4	4.5	4.5
	Variance	0.3	0.2	0.3	0.3	0.3
	Std Dev	0.5	0.4	0.5	0.5	0.5

(continued on page 1-4)

[3] The definition of an official user was made consistent across DOE user facilities by the program office. The timeframe for an official user is now reported as a fiscal year as opposed to a calendar year. In addition, staff members and vendors are considered users if they are members of a reviewed project.

Table 1.2 2014 User Survey Results by Allocation Program (Cont.)

2014 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Problem Resolution	Mean	4.5	4.6	4.5	4.5	4.5
	Variance	0.4	0.3	0.4	0.5	0.5
	Std Dev	0.7	0.6	0.6	0.7	0.7
All Questions	Mean	4.5	4.4	4.5	4.5	4.5
	Variance	0.4	0.7	0.5	0.5	0.5
	Std Dev	0.6	0.8	0.7	0.7	0.7

In 2014, as Table 1.3 shows, ALCF exceeded the Overall Satisfaction and User Support targets.

Table 1.3 2013 and 2014 User Support Metrics

Survey Area	2013 Target	2013 Actual	2014 Target	2014 Actual
Overall Satisfaction Rating	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0
Average of User Support Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0

1.2 Problem Resolution Metrics

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

Table 1.4 Problem Resolution Metrics

	2013 Target	2013 Actual	2014 Target	2014 Actual
% User Problems Addressed within Three Working Days	80.0%	91.9%	80.0%	96.0%
Average of Problem Resolution Ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.5/5.0

1.3 User Support and Outreach

1.3.1 User Support

Phone and E-mail Support

ALCF answered 6,421 support tickets in 2014, a nearly 20 percent increase in ticket volume over 2013. The largest number of these tickets involved accounts (see Table 1.5).

Table 1.5 Ticket Categorization for 2013 and 2014

Category	2013	2014
Access	1,080 (20%)	1,085 (17%)
Accounts	1,934 (36%)	2,289 (36%)
Allocations	581 (11%)	690 (11%)
Applications Software	306 (6%)	323 (5%)
Automated E-mail Responses	397 (7%)	1,135 (18%)
Data Transfer	41 (1%)	55 (1%)
I/O and Storage	220 (4%)	153 (2%)
Miscellaneous	190 (4%)	155 (2%)
Quota Management	38 (1%)	44 (1%)
System	559 (10%)	487 (8%)
Visualization	11 (0%)	5 (0%)
TOTAL TICKETS	5,357	6,421

In 2014, the highest-percentage increase in number of tickets was in the "Automated E-mail Responses" category. Formerly known as "Bounce" (and renamed after reviewer feedback), this category reflects out-of-office responses to ALCF e-mail as well as undelivered mail messages from the mail server. The increase in automated e-mail responses was triggered by an internal security audit (400+ messages in April 2014) and an internal mail routing issue (identified in October 2014). The increases in account and allocation tickets are correlated with increases in users and requests for allocations. There has been a drop in the percentage of tickets categorized as "System"-related requests, as ALCF now supports Mira only.

Continuous Improvement Efforts in User Support

In order to ensure that ALCF met user expectations and target metrics for problem resolution, the support team met and reviewed open tickets on a weekly basis. Issues were then followed up with subject matter experts for specific tickets. Adopting this approach resulted in a steady increase in the number of tickets resolved month over month when compared to 2013 through the month of August. In the third quarter of 2014, ALCF added two staff members to the User Experience team in order to provide even better support to users.

Migrating Scripts out of Decommissioned Intrepid Hardware

ALCF migrated key internal reports and external communications from the decommissioned Intrepid support infrastructure to a centralized service using Jenkins software.

Jenkins provides continuous integration services for software development. It is a server-based system running on a web server. It supports source code management tools such as Subversion, Git, etc., and can execute projects as well as arbitrary shell scripts via scheduled jobs. Jobs can be triggered by cron-like mechanisms, updates to configuration files, and commits to code. The most valuable feature of Jenkins is how it aggregates information about a scheduled process, centralizes logging, and provides rich notifications about the jobs.

An advantage of using Jenkins over cron jobs is seen in its monitoring of externally run jobs, even those that are run on a remote machine. For example, with cron, all script owners receive regular e-mails that capture the output, and it is up to the script owner to look at them diligently and notice when something is broken. Jenkins keeps those outputs and makes it easy to notice when something has gone wrong.

The internal reports provide a system of warnings and informational messages to staff about the status of tickets in the ticketing system. They help provide checks and balances for membership between projects and Unix groups – the method used to provide access to the projects data, internal audit of user agreements for new additions to INCITE and ALCC projects, and a daily report of expiring user accounts for the accounts team to process. The external communication provides users and principal investigators (PIs) with a system of warning messages about upcoming user account expirations, including foreign national security paperwork expirations.

Standard Operating Procedure for Technical Support at Events

ALCF provides accounts and access to attendees of workshops and training events organized by the facility. This support effort includes setting up projects and allocations (request, approval, creation), tracking attendee accounts (creation, membership, login access), ensuring access to systems (reservation, special queue), providing assistance during the event (distributing tokens, login issues), establishing response rates with subject matter experts for potential issues, and performing final wrap up (collecting tokens, disabling accounts, revoking access, etc.).

While all of these tasks were performed in past years, in 2014 ALCF established a standard operating procedure for all of the above tasks, thereby supporting events based on prior years' experience. The procedure defined required timelines for each activity within the preceding list, along with the tasks associated with each activity in the form of checklists. The new procedure was used for the first on-site 2014 event, Mira Boot Camp, and was refined with subsequent events.

Processing "Pre-Approved" Accounts

As part of the communication to PIs leading up to the start of ALCC 2014 allocation awards at ALCF, ALCF set up a new step in the process to speed up account requests for an awarded ALCC project. In the message to PIs, ALCF asks for a pre-approved list of project members and their e-mails. Receiving this list speeds the account request process because it enables ALCF staff to begin the account process immediately. The process was refined further during the ramp-up of INCITE 2015 projects. This process contributed towards the success of each INCITE 2015 project having at least one member capable of running jobs on the machine on day one of the allocation.

1.3.2 Outreach Efforts

User Advisory Council

ALCF convenes the User Advisory Council (UAC) to comment on key technology upgrades, advise on messaging and communication, and provide feedback on user-centric metrics. ALCF is grateful for the time, effort, and contributions provided by this advisory body.

The seven-member UAC represents all three allocation programs (INCITE, ALCC, and Director's Discretionary). Meetings are held monthly with subsequent meetings scheduled as part of the current meeting.

In 2014, the UAC helped beta-test ALFC's annual survey and helped consider questions to be revised or replaced. They commented on the file system cache upgrade, the visualization cluster upgrade, and the upgrade of Cetus to four racks and the scheduling policy associated with it. They advised ALCF on how to communicate key Director Discretionary policies and provided feedback on how users approach applying for Director Discretionary allocations and changes in the scheduling policy on Tukey, ALCF's existing visualization cluster. All of the input helped ALCF to provide more useful and efficient service to users of the facility.

Industry Outreach

ALCF encourages industry to use ALCF resources, attend ALCF training, and collaborate with ALCF staff. The ALCF Industry Outreach Lead manages interactions with industry to help ensure that the appropriate people and resources are engaged. Integral to this work are the strong relationships with other Argonne divisions.

One example is ALCF's work with Caterpillar, where modeling work is being done in the Energy Systems division. Smaller-scale simulation runs are being done at the Laboratory Computing Resource Center, with software scaling and larger-scale runs accomplished at ALCF. The work led to the creation of the Virtual Engine Research Institute and Fuels Initiative (VERIFI) that to date has introduced more than 20 engine manufacturers to Argonne resources and expertise.

Standards Organizations

ALCF actively participates in and contributes to several standards organizations relevant to the scientific communities it supports. Organizations with ALCF staff as members are:

- OpenMP ARB, Language Comm
- SPEC HPG
- OpenSFS, OFA
- ScicomP

Other Support-focused Collaborations

ALCF expertise was tapped by the National Renewable Energy Laboratory (NREL) in developing their first survey for their high-performance computing (HPC) center. ALCF recommended several possible vendors and provided an annotated bibliography on the development of a survey.

2014 Workshops and Webinars

ALCF conducted workshops and webinars to support the efforts of users and their project teams (Table 1.6). The workshops are highly rated by those attending, as evidenced by feedback received in the annual user survey. ALCF also collaborates with peer DOE institutions to develop training opportunities, explore key technologies, and share best practices that improve the user experience. In addition, ALCF, Oak Ridge Leadership Computing Facility (OLCF), and the National Energy Research Scientific Computing Center (NERSC) are working together to develop strategies to ensure that key applications are ready for the scale and architecture of the next generation of DOE supercomputers. Using ALCF's Early Science Program for Mira as a model, this effort is intended to better understand the needs of users and increase user participation across all facility designs.

Table 1.6 2014 Workshops and Webinars

Event	Description	Dates (2014)
Getting Started Videoconference	Small-group telepresence event that ran several times over a two-week period for users new to Mira and ALCF.	January, July, August, and December
Mira Performance Boot Camp	Tutorials on scaling and performance tuning codes for projects applying for 2015 INCITE awards.	May
Joint Facilities Forum on Data-Intensive Computing	Held at NERSC, ALCF was among the facilities assisting with this effort by chairing panels and presenting on data-intensive topics.	June
QMC (quantum Monte Carlo) Training	Hosted and ran training program on QMC simulations of real materials funded by DOE Office of Science's Basic Energy Sciences (BES) program, and The National Science Foundation (NSF).	July
ATPESC	Training program targeted for students and postdocs.	August
ParaView Training	Developed cross-laboratory pilot training program to explore opportunities and discover issues with cross-lab training efforts.	November
VERIFI Workshop	Supported and ran a tutorial for a workshop funded by Argonne, Convergent Science, and Cray.	November

Getting Started Videoconference

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

Getting Started class sizes are kept small to emulate an intimate classroom experience that includes hands-on exercises. ALCF uses the Blue Jeans videoconferencing technology to create a virtual presence among participants. Session offerings are timed to coincide with INCITE and ALCC award announcements when many projects add new users. This just-in-time training maximizes research on ALCF resources by facilitating and encouraging early utilization of allocations. To further accommodate new user needs, a video presentation of this popular training course is also available for viewing on the ALCF website. A total of 57 participants attended this training.

Mira Performance Boot Camp

New and seasoned ALCF users convened at Argonne in May for the annual Mira Performance Boot Camp. For many, the primary goal of attending was to tap into the expertise of ALCF staff for assistance in improving their code's scalability to demonstrate computational readiness for a 2015 INCITE award. In support of this goal, ALCF's annual scaling workshop is especially geared for and timed to coincide with efforts of teams preparing INCITE proposals.

The bulk of the three-day event was devoted to hands-on tuning of applications. In addition, ALCF experts spoke on topics of interest, including Blue Gene/Q architecture, ensemble jobs, parallel I/O, and data analysis. Guest tool and debugger developers also provided information and individualized assistance to attendees.

With dedicated Boot Camp reservation queues, the 64 participants (comprising attendees, ALCF staff, and vendors) had quick, uninterrupted access to ALCF resources, allowing them to run nearly 400 jobs and use more than 1.7 million core-hours as they diagnosed code issues and tweaked performance.

Joint Facilities User Forum on Data-Intensive Computing

In June, leaders from several DOE facilities convened in Oakland, California, for the Joint Facilities User Forum on Data-Intensive Computing. Organized by Argonne, Oak Ridge, Lawrence Berkeley, Sandia, Lawrence Livermore, and Los Alamos national laboratories, the three-day workshop brought together 149 attendees (facility users and HPC center staff) to discuss the latest trends and techniques related to data management, analysis, and visualization. Building on the momentum of this successful conference, ALCF and its co-organizers are planning more user-focused events, including virtual workshops for DOE facility users and staff.

QMC Training

In July, ALCF hosted a training workshop on quantum Monte Carlo (QMC), an increasingly popular computational method for chemistry, materials science, and physics research.

Designed to grow the user base of QMC, the weeklong workshop, which was supported by the DOE Office of Science and the National Science Foundation, brought in 34 participants consisting of scientists, professors, and graduate students.

The training program introduced participants to the fundamentals of QMC theory, as well as the latest developments in QMC methods and applications. The attendees also received hands-on training using QMCPACK, an open source QMC code, on Mira.

ATPESC 2014

ALCF organized Argonne Training Program on Extreme-Scale Computing (ATPESC), an intense two-week program that provides in-depth training on the key skills, approaches, and tools needed to conduct research on the high-end computing systems in use today and emerging tomorrow.

This year's event marked the second installment of the ATPESC. With HPC experts serving as the lecturers, the 62 attendees learned the ins and outs of effectively designing, implementing, and executing large-scale computer science and engineering applications across a variety of supercomputing platforms—including methods expected to be applicable to future systems.

As part of the hands-on training, the participants were provided access to DOE supercomputing resources, including Argonne's IBM Blue Gene/Q systems, Vesta and Mira; the OLCF Titan system; and the Edison system at NERSC.

For those unable to participate, videos of the 2014 ATPESC lectures are publicly available at extremecomputingtraining.anl.gov/2014-videos.

Getting Started with ParaView for DOE Facility Users

As part of a pilot online series following the Joint Facilities User Forum in June 2014, ALCF volunteered to offer the first of the group's inter-facility training events for users of DOE high-performance computing centers. ALCF elected to conduct a tutorial on the fundamentals of ParaView—an open-source, multi-platform data analysis and visualization application available at many DOE high-performance computing centers.

Offered in the format of a highly interactive videoconference, the two, two-hour sessions were held November 5, and registration was limited to eight individuals per session to maximize interaction between instructors and participants. Users from ALCF, OLCF, LLNL, and NERSC attended. The technology that facilitated these virtual sessions was Blue Jeans, a cloud-based video collaboration service.

The tutorial led current DOE high-performance computing center users through hands-on exercises to explore ParaView's user interface, interact with data, and apply some of the application's main post-processing filters. A total of 12 participants launched ParaView in server mode on computing resources at participating DOE centers and connected to a ParaView client on their local workstations. Instructors also covered tips, tricks, and best practices with ParaView.

VERIFI Workshop

In November, ALCF staff teamed up with other scientists and researchers at Argonne's Virtual Engine Research Institute and Fuels Initiative (VERIFI) to organize a hands-on session at the Workshop for High-Performance Computing-Enabled Engine Simulations. The workshop, which was sponsored by Convergent Science, Inc., and Cray, was held at the Advanced Photon Source Conference Center, with the hands-on session conducted at the Theory and Computational Science Conference Center.

Bringing together an unequaled, multidisciplinary team of experts in high-performance computing, fuel chemistry, combustion science, and engine performance with some of the world's fastest supercomputers, most diverse engine labs, and world's brightest X-ray beams, VERIFI offers a "Dream Team" environment at Argonne National Laboratory in which to answer complex engine questions, verify the uncertainties associated with those answers, and shrink development timescales — at reduced cost.

As part of the hands-on session, 17 participants received access to Mira, along with step-by-step instructions on job submission, an overview of visualization techniques, and details on ALCF code improvements for scaling up CONVERGE™, the Computational Fluid Dynamics software from Convergent Science, Inc. The participants ran several piston engine simulations on 2,048 cores and realized the benefits of using HPC tools for engine simulations.

Outreach at Supercomputing 2014

ALCF staff engaged in outreach activities at Supercomputing 2014 (SC '14) to recruit new talent through the career fair, diversify the field of HPC by organizing and participating in Broader Engagement, and encourage student interest in supercomputing by sponsoring a student cluster challenge. ALCF also used SC '14 as a venue to promote the ATPESC program.

1.3.3 Communications

Communications through Mailing Lists and Social Media

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

Table 1.7 2014 Primary Communication Channels

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

Table 1.8 2014 Target Audiences

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

ALCF uses online software called MailChimp to manage and send HTML-formatted e-mail. Every month one or two "science stories" are featured. These stories are either the outcome of research carried out on ALCF resources or an advancement made by ALCF staff and researchers in the field. Critical deadlines, targeted announcements, key events, and significant news stories are also included in the monthly publication. An example of a targeted announcement was informing the audience of the call for the Argonne Training Program on Extreme-Scale Computing (ATPESC).

In 2013, ALCF changed the rate of release for newsletters from quarterly to monthly and moved the newsletter from a PDF-based format to HTML-formatted e-mail. This newsletter was distributed in two formats for more than nine months. In 2014, after several messages to users, ALCF removed the option for text-based e-mails to HTML-only formatted e-mails. This change resulted in only a few "unsubscribes" by the text-based newsletter users.

Promotional Activities and Media Hits

The media team curates and publishes a list of ALCF media hits throughout the year and promotes the list on the ALCF website and in the monthly Newsbytes. In 2014, the facility posted 59 media hits to ALCF's website. In order for the team to post a media hit, it must be specifically about ALCF. The media team also "weeds out" publications that simply reprint an already published story.

In addition to the media hits we track on our website, the media team uses Meltwater News public relations suite to track media hits. Established in 2001, Meltwater News suite is a global online media monitoring company that tracks articles from more than 200,000 news publications, as well as Twitter, YouTube, Facebook, and blogs. In 2014, Meltwater captured 142 mentions of the "Argonne Leadership Computing Facility" and "ALCF."

Custom Communication

Custom e-mail messages to users include allocation end reminders to INCITE and ALCC PIs to encourage them to use their remaining hours; a series of e-mails that initiate the start of INCITE and ALCC projects; ALCC and INCITE quarterly report reminders; special announcements to targeted groups; and decommissioning messages to Intrepid and Surveyor users. Users cannot opt out of these custom e-mails.

Starting with the 2014 ALCC project start-up e-mails to PIs, ALCF split the messages between the PIs' personal user account management and their project management tasks to increase clarity with separation of concerns and decreased complexity, and to increase the number of projects that were up and running on the first day of the allocation. ALCF adopted and further refined this process for 2015 INCITE project start-up e-mails (Table 1.9). For each engagement, ALCF staff focused on one or two changes that were meant to encourage a faster adoption.

Table 1.9 Key Changes and Outcomes Due to Efforts

Year and Program	Change	Outcome
2014 ALCC	Separated e-mail messages to PIs into account and project management	Reduced the number of steps for PIs to take per e-mail sent along with conceptually gathering those steps
	Tracked accounts that were "pre-approved" by PI in response to e-mail	Reduced the need to approve user accounts for the same member multiple times
2015 INCITE	Refined this separation further and incorporated scheduling welcome calls	Reduced the number of channels used to contact PIs for multiple purposes Long-standing PIs without user accounts in the past successfully set up their account
	Tracked accounts that were "pre-approved" by PI in response to e-mail	Every single project had at least one user who could run jobs on the first day of the allocation

ALCF Website Improvements

ALCF made multiple improvements to its website, including enhanced content publishing workflow, improved hosting environment, and module upkeep.

To improve managers' ability to track published content changes, ALCF staff created a new dashboard view that provides the option to sort content by date, see who made changes and when, know the content URL, export to .csv or .xls, and print a report.

Another feature enhancement is the option to "pass the pen." In this feature, an ALCF staff member could use the content management system to alert an ALCF colleague that it was his or her turn to edit the page or approve edits to that page. This alert would ensure that the right people would be part of an approval chain of edits.

ALCF switched hosting services from a web server maintained at Argonne to the Acquia Cloud hosting platform. This change provides access to Acquia's developer performance tools, 24×7 monitoring, and notifications of security updates. ALCF also began using Acquia's Remote Administrators to stay current with core modules updates and assist on critical security updates.

Other Publications

ALCF produced a variety of print publications used for promotion, education, and recruiting (Table 1.10). Visitors are provided with assembled packets tailored to their particular area of interest.

Table 1.10 Publications Designed for Print in 2014

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

Conclusion

As a user facility, ALCF is focused on ensuring the success of all facility users and customers. During CY 2014, ALCF made website improvements, revised the user survey, engaged in outreach activities, and enhanced communication efforts in various e-mail channels. As such, ALCF continues to support customers effectively, resolve problems, and perform outreach.

Section 2. Business Results

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

ALCF Response

ALCF has exceeded the metrics of 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 acceptable 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 Summary of All Metrics Reported in the Business Results Section

	Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	2013 (Apr 9–Dec 31)		CY 2014	
	Target	**Actual**	**Target**	**Actual**
Scheduled Availability	85.0%	95.5%	95.0%	98.7%
Overall Availability	80.0%	90.6%	90.0%	95.7%
System MTTI	N/A[a]	4.23 days	N/A	8.98 days
System MTTF	N/A	11.29 days	N/A	25.80 days
INCITE Usage	2.1B	2.4B	3.5B	3.9B
Total Usage	N/A	3.6B	N/A	5.8B
System Utilization	N/A	79.4%	N/A	87.6%
Mira INCITE Overall Capability[b]	20.0%	60.7%	30.0%	64.5%[d]
Mira INCITE High Capability[c]	5.0%	33.3%	10.0%	33.1%

[a] N/A = Not applicable.

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

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

[d] INCITE usage includes 15.0M core hours from Cetus production jobs.

ALCF Resources

During CY 2014, ALCF operated one production resource, Mira. Mira is a 48K-node, 768K-core, 10 PF Blue Gene/Q with 768 TB of RAM. Mira mounts two General Parallel File System (GPFS) file systems with approximately 26 PB of usable space and has access to the facility-wide HPSS (high-performance storage system) tape archive. Mira has an associated visualization and analysis cluster called Tukey. ALCF operated two other Blue Gene/Q systems, Cetus and Vesta.

Cetus is a 4k-node, 64K-core Blue Gene/Q with 64 TB of RAM. Cetus shares the file system 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.

In 2014, ALCF permitted select use of Cetus for INCITE projects with uncommon requirements that would have been more disruptive when attempting to fit onto Mira. Cetus was not a generally available resource for this purpose; ALCF assessed individual project requirements before allowing production runs by the projects.

2.1 Resource Availability

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

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

2.1.1 Scheduled and 2.1.2 Overall Availability

Mira has been in full production since April 9, 2013. ALCF has agreed, in consultation with the DOE Program Manager, to metrics of 90 percent overall availability and 95 percent scheduled availability. Table 2.2 summarizes the availability results.

Table 2.2 Availability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	2013 (Apr 9–Dec 31)		CY 2014	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	85.0	95.5	95.0	98.7
Overall Availability	80.0	90.6	90.0	95.7

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

Explanation of Significant Availability Losses

This section provides a brief description of the causes of major losses of availability for the period January 1, 2014, through December 31, 2014, as annotated in Figure 2.1.

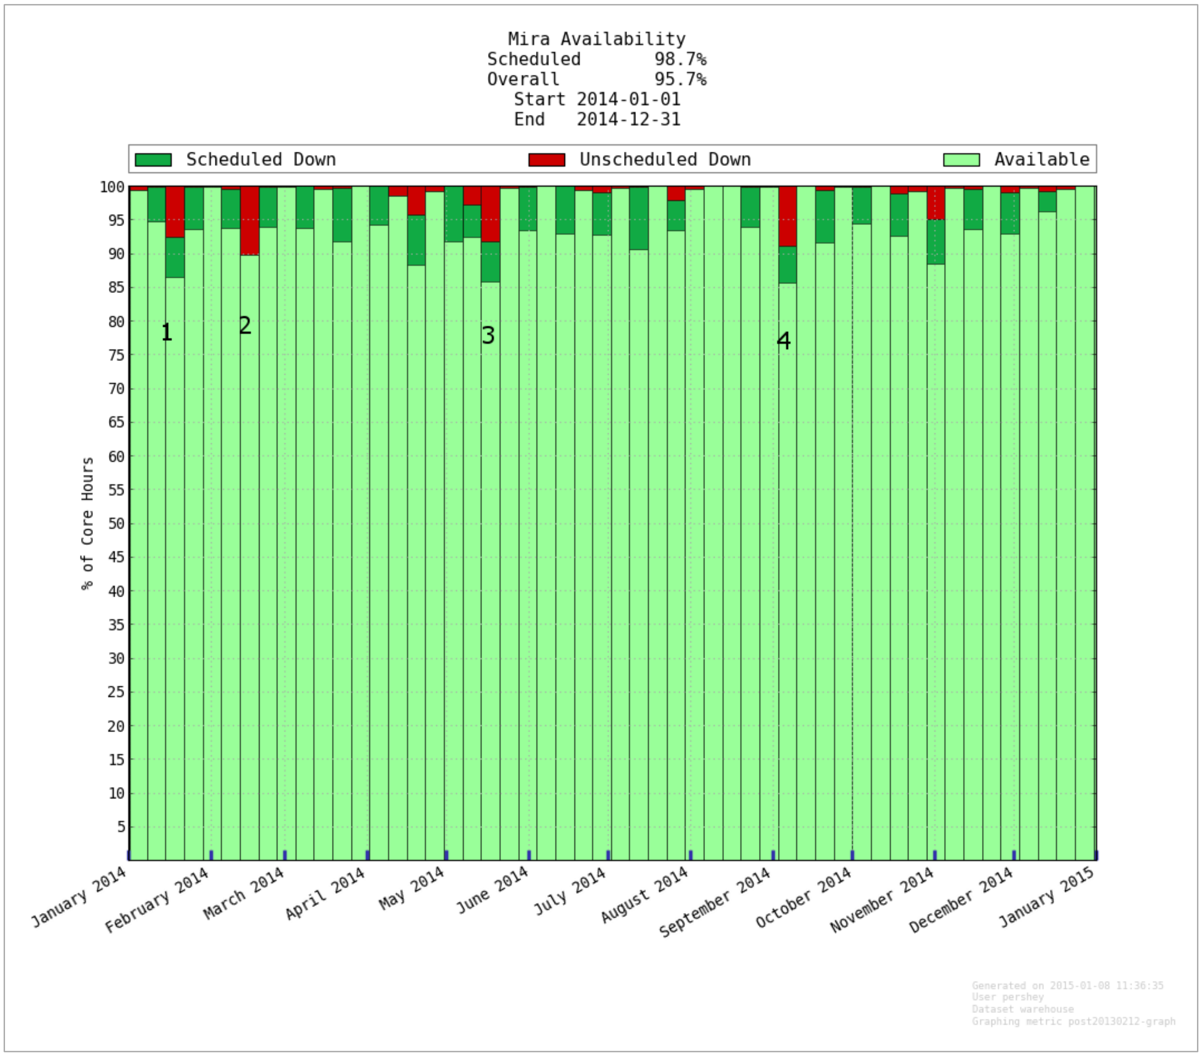

Figure 2.1 Mira Weekly Availability for CY 2014

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

Item 1:
A DB2 database outage contributed to the machine going down. The master of the High Availability Disaster Recovery (HADR) pair failed, and the backup did not properly fail-over using Tivoli Storage Manager (TSM). The scheduler, Cobalt, and supporting scripts depend on DB2 for its operation. The backup node could not confirm that the master was down, so it could not determine what to do. The backup node did not take over automatically because the TSM/DB2 configuration was not set up correctly for this scenario. After the fail-over was manually completed, some databases were still not working and had HADR-related problems, so they had to be started manually. This configuration error has been corrected.

Item 2:
An upgrade of the GPFS Storage Server (GSS) caused mount failures. GSS was mounted on Mira production nodes in order to perform scaling tests. A bug in the new GSS code being tested caused clients to crash during the remount of the file system after the upgrade. The remote mount was removed, and all future tests were limited to reservations with the GSS file systems mounted only on those clients being tested.

Item 3:
In mid-May, an electrical panel was found to be loose. Because this situation represented a potential fire hazard, ALCF undertook the repair on very short notice.

Item 4:
In early September, a power outage occurred on the data center A bus. As designed, it immediately failed over to the redundant B bus, but the system still went down. Two problems were discovered:

- On the cold/supply side of the heat exchanger, a pump drive was misconfigured, and it did not automatically restart after the power blip. This configuration has been corrected, and in future similar situations, this event would not cause an outage.

- On the hot/load side of the heat exchanger, the pump did restart; however, because pumps are mechanical devices, they are limited in how fast they can accelerate. The loss of flow to the Blue Gene was significant enough that the Blue Gene control system detected low flow and immediately initiated an emergency shutdown to protect the hardware. The only way to ameliorate this situation is to put a UPS (uninterruptible power supply) on those pumps/drives. ALCF is investigating this option, and once pricing is determined, the facility will decide whether doing so is a cost-effective option.

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

Through a series of discussions between ALCF, Oak Ridge Leadership Computing Facility (OLCF), and National Energy Research Scientific Computing Center (NERSC), all three sites agreed to a common approach for System MTTI and System MTTF. MTTI and MTTF are reportable values; however, no specific target has been set. Table 2.3 summarizes the current MTTI and MTTF values.

Table 2.3 MTTI and MTTF Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	2013 (Apr 9–Dec 31)		CY 2014	
	Target	Actual	Target	Actual
System MTTI	N/A*	4.23 days	N/A	8.98 days
System MTTF	N/A	11.29 days	N/A	25.80 days

* N/A = Not applicable.

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, and OS upgrades, etc.

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.

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

Table 2.4 System Utilization Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	2013 (Apr 9–Dec 31)		CY 2014	
	Target	Actual	Target	Actual
System Utilization	N/A*	79.4%	N/A	87.6%

* N/A = Not applicable.

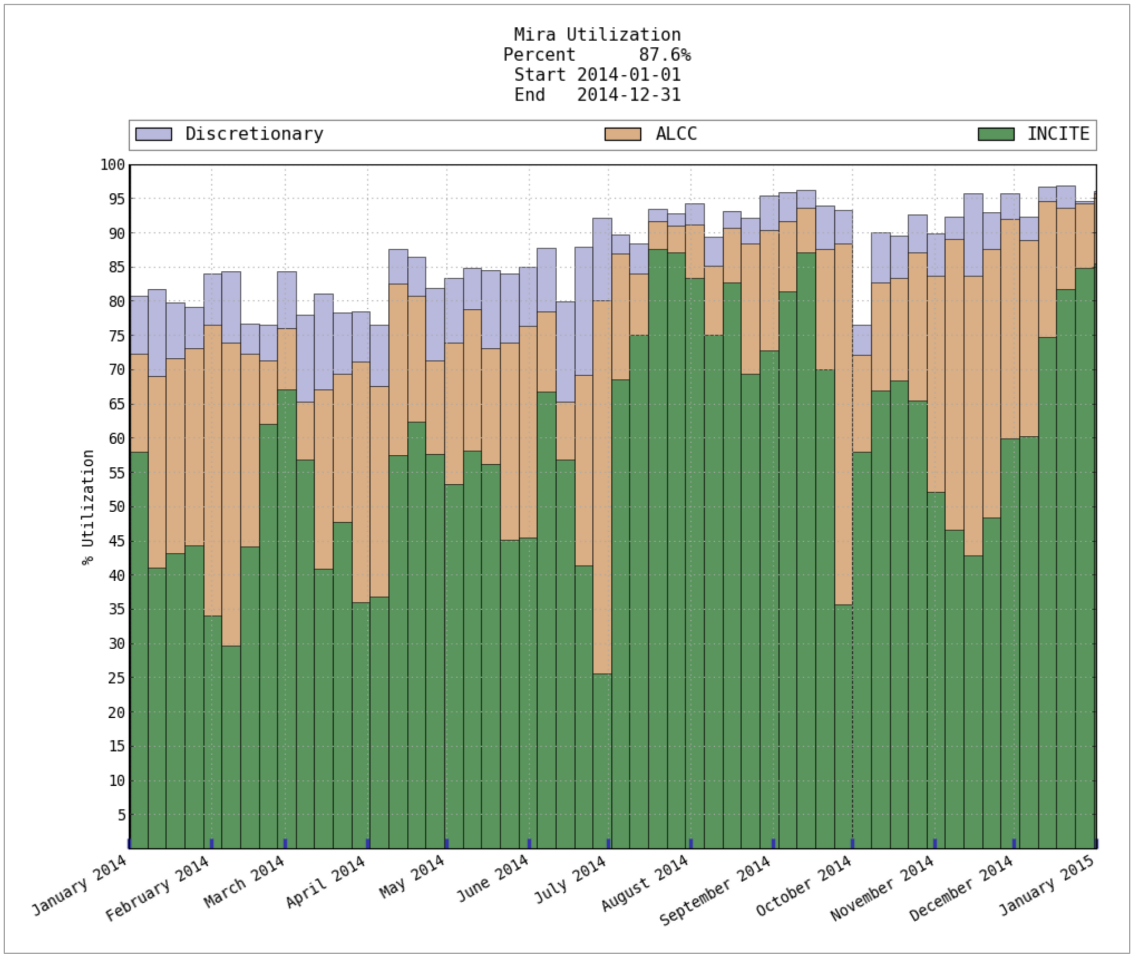

Figure 2.2 System Utilization over Time by Program

The system utilization for Mira was 87.6 percent for its 2014 production period of January 1, 2014, through December 31, 2014.

Table 2.5 shows how Mira's system hours were allocated and used by allocation source. Multiplying the theoretical hours by availability and utilization values to which we agreed upon with our DOE program manager determines the hours available. Of the hours available, 60 percent is allocated to the INCITE program, up to 30 percent is available for ALCC program allocations, and 10 percent is available for Director's Discretionary (DD) allocations. The ALCC program runs from July through June, so to arrive at allocated values for the calendar year, half of the hours are assigned to each year. The allocated values for the DD allocations appear higher than expected because they represent a rolling allocation. Because a majority of DD projects are exploratory investigations, the time allocations are often not used in full. DD allocations are discussed in detail in Section 3.3. In CY 2014, ALCF successfully delivered a total of 5.8 billion core-hours on Mira.

Table 2.5 Core-Hours Allocated and Used by Program

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	2013 (Apr 9–Dec 31)			CY 2014		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	%	Core-hours	Core-hours	%
INCITE	2.1B	2.4B	66.8%	3.5B	3.9B	67.0%
ALCC	555.5M	348.4M	9.6%	1.6B	1.4B	24.9%
DD	2.9B	851.3M	23.6%	1.1B	468.9M	8.1%

Summary: The DD use in 2013, as described in the 2013 OAR, includes some Mira Early Science work that was completing analysis. For CY 2014, the system usage and utilization values are in line with general expectations. Utilization gradually improved over the course of the year. 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. On Intrepid, capability was defined as using greater than 20 percent of the machine. However, 20 percent of Mira would be 9.6 racks, which is not a viable configuration. Hence, the Mira capability metric is split into two parts. Overall Capability requires that a minimum of 30 percent of the INCITE core-hours be run on eight racks or more (16.7 percent), and High Capability requires a minimum of 10 percent of the INCITE core-hours be run on 16 racks or more (33.3 percent). See Appendix A for more detail on the capability calculation. Table 2.6 and Figure 2.3 show that ALCF has substantially exceeded these metrics set for INCITE. Although no targets are set, data are also provided in the table for ALCC and DD projects as reference, and Figure 2.4 shows the overall distribution of job sizes over time.

Table 2.6 Capability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	2013 (Apr 9–Dec 31)			CY 2014		
Capability Usage	**Total Hours**	**Capability Hours**	**Percent Capability**	**Total Hours**	**Capability Hours**	**Percent Capability**
INCITE Overall	2.4B	1.5B	60.7%	3.9B	2.5B	64.5%
INCITE High	2.4B	803.1M	33.3%	3.9B	1.3B	33.1%
ALCC Overall	348.4M	134.0M	38.5%	1.4B	787.6M	54.8%
ALCC High	348.4M	54.9M	15.8%	1.4B	124.7M	8.7%
Director's Discretionary Overall	851.3M	291.3M	34.2%	468.9M	179.7M	38.3%
Director's Discretionary High	851.3M	190.8M	22.4%	468.9M	101.8M	21.7%
TOTAL Overall	3.6B	1.9B	52.3%	5.8B	3.5B	60.0%
TOTAL High	3.6B	1.0B	29.0%	5.8B	1.5B	26.1%

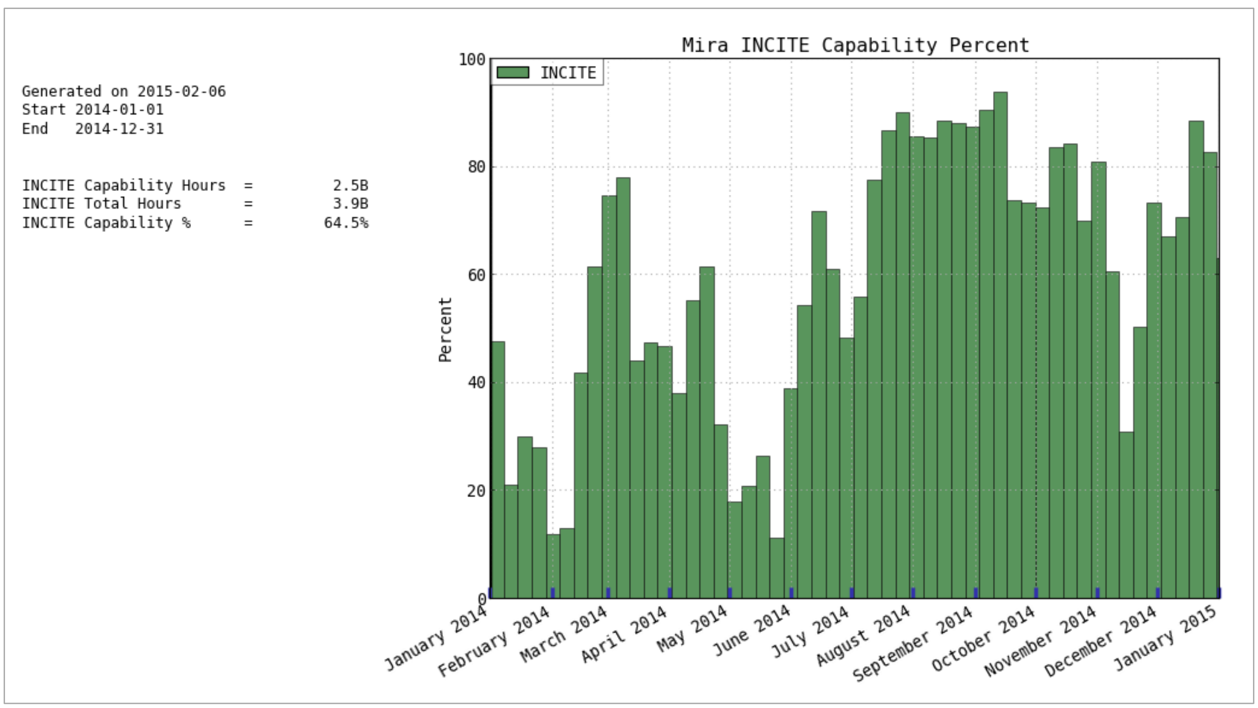

Figure 2.3 Mira INCITE Overall Capability

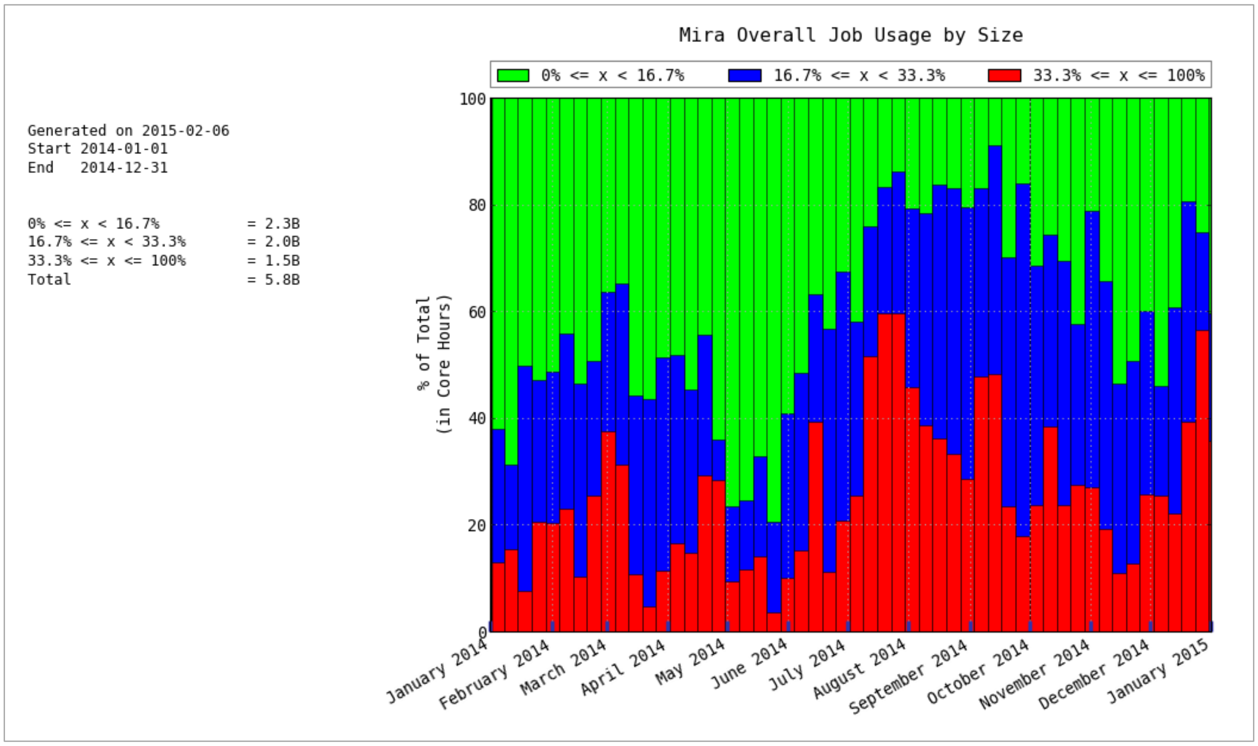

Figure 2.4 Mira Job Usage by Size

2.4 Intrepid Decommissioning

On December 31, 2013, after many years of service, the 40-rack Blue Gene/P supercomputer Intrepid was decommissioned. Because Intrepid was still a potentially valuable resource, potential users for parts of it were sought. ALCF was able to donate three racks to North Carolina State University, enabling their facility to run one rack with spare parts. ALCF was prepared to send two racks to Brookhaven National Laboratory for parts; however, they ultimately chose to decommission their system instead. Lawrence Livermore National Laboratory, the University of Notre Dame, Rice University, Texas A&M University, Montana Tech University, The University of District Columbia, and the Living Computer History Museum were also contacted. All of these institutions either declined the offer or had a resource constraint related either to staffing or the facility, which prevented each from accepting the offer. A recycling company was engaged to take the rest of the machine.

2.4.1 Intrepid Decommission — Process Failure

During the shutdown of the Blue Gene/P systems in Building 369 (ISSF), ALCF experienced a failure to follow documented processes. This failure did not result in damage, injury, or any other notable effect beyond the generation of a report and process improvements.

Two gaps in the process were noted:

1. The Laboratory work planning control (WPC) software failed to include key information in the worker package. The worker package did not list the authorized workers, and the individuals performing the work—while appropriately trained to do the work—had not been authorized to do it.
2. Workers were not wearing protective gloves during an administratively added second Zero Energy Verification (ZEV). The contributing factors were determined to be shift handoff (this being the second shift working on the task) and improper mindset. Because the racks had been powered off for several months and workers had performed the "normal" ZEV check, they did not regard the equipment as electrically energized. No injury or damage occurred. The defense-in-depth method of multiple checks on the environment prevented workers from being exposed to dangerous energy.

As a result of this incident, the following steps were taken:

- A stop work was called immediately upon discovery.
- An investigation was conducted and report written.
- The Work Planning and Control (WPC) system was updated so that the Worker Package lists the authorized workers.
- Updates to the Lockout/Tagout (LO/TO) training course were implemented, and staff members were retrained.
- Training in the use of the WP&C system was provided to the Operations Team.

The ALCF Operations team enjoys an excellent safety culture. In the history of the facility, only two safety incidents had previously occurred: in each case, a cut requiring a Band-Aid. This excellent record notwithstanding, there is always room for improvement. By avoiding complacency and taking even such minor incidents seriously, the likelihood of incurring a major incident is reduced.

2.5 Cetus Upgrades

Cetus is intended primarily as a debug machine. Its scheduling policy favors short, small jobs, and provides fast turnaround times. IBM announced that it was discontinuing the Blue Gene line and started offering racks at greatly reduced prices to sell off excess inventory. ALCF purchased four additional BG/Q racks for Cetus. Cetus was upgraded to have four racks of Blue Gene/Q online, leaving one rack offline to be used for spare parts. This provides a significant cost savings because the cost of the rack was much lower than the cumulative cost of the parts, if bought individually.

In addition, this move ensures an adequate parts supply, which was an issue on Blue Gene/P and is expected to be an issue on Blue Gene/Q, particularly if Mira's lifetime is extended to 2019—a strategy currently under consideration. Increasing Cetus to four racks also provides a better resource for running very long, small jobs. Running such jobs on Cetus avoids blocking large jobs on Mira for long periods of time. If needed, later in the life of Mira, ALCF will reduce the size of Cetus a mid-plane at a time and use them for spare parts. Implementing this longer-term approach was part of the agreement with ALCF's DOE program manager when ALCF purchased the additional racks.

Conclusion

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

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

	Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	2013 (Apr 9–Dec 31)		CY 2014	
	Target	Actual	Target	Actual
Scheduled Availability	85.0%	95.5%	95.0%	98.7%
Overall Availability	80.0%	90.6%	90.0%	95.7%
System MTTI	N/A[a]	4.23 days	N/A	8.98 days
System MTTF	N/A	11.29 days	N/A	25.80 days
INCITE Usage	2.1B	2.4B	3.5B	3.9B
Total Usage	N/A	3.6B	N/A	5.8B
System Utilization	N/A	79.4%	N/A	87.6%
Mira INCITE Overall Capability[b]	20.0%	60.7%	30.0%	64.5%[d]
Mira INCITE High Capability[c]	5.0%	33.3%	10.0%	33.1%

[a] N/A = Not applicable.

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

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

[d] INCITE usage includes 15.0M core hours from Cetus production jobs.

ALCF closely tracks hardware and software failures and their impact on user jobs and metrics. These data are used as a significant factor in the selection of troubleshooting efforts and improvement projects.

This page intentionally left blank.

Section 3. Strategic Results

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

ALCF Response

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

In this section, the Facility reports:

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

3.1 Science Output

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

Table 3.1 shows the breakdown by journal of refereed publications based (at least in part) on the use of ALCF. In addition to broad-audience journals such as *Nature, Physical Review Letters,* and *SC '14,* ALCF users published in a two-part special issue of *Computing in Science & Engineering* (CiSE) (Advances in Leadership Computing) that focuses on science enabled by leadership computing. Drs. James Hack and Michael Papka, directors of OLCF and ALCF, respectively, served as guest editors for this special edition of CiSE.

Table 3.1 Summary of Refereed Publications

Nature	CiSE	Physical Review Letters	SC '14	Total 2014 Publications
1	5	7	6	157

3.2 Scientific Accomplishments

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

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

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

Intensity-Dependent Dynamics in Fermilab and CERN Accelerators

James Amundson, Fermilab

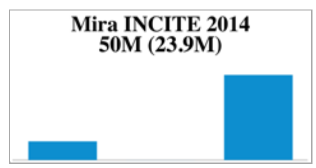

Particle accelerators are an enabling technology for both basic research (e.g., studying the fundamental constituents of matter and the structure of nuclei) and the applied sciences (e.g., probing the structure of materials). Fermilab researchers are using ALCF resources to perform complex accelerator simulations aimed at reducing the risks and costs involved in developing the world's highest-intensity particle beams.

The future of high-energy physics requires running today's accelerators at intensities that are higher than ever. Both Fermilab and CERN strive to accurately understand intensity-dependent effects in their accelerator complexes. Such understanding requires detailed numerical modeling that goes beyond the capabilities of desktop machines and simple clusters. This project is working to simulate these accelerators with unprecedented fidelity, which will enable new discoveries in particle physics.

The Fermilab Recycler and Main Injector form the final high-energy stage of the Fermilab accelerator complex. During each acceleration cycle, the Recycler receives protons in six batches from the Booster. Through a radio frequency manipulation procedure known as slip-stacking, batches are combined to increase the instantaneous intensity (Figure 3.1). The protons are then transferred to the Main Injector where they are accelerated from 8 GeV to 120 GeV. In the future, the Recycler will be used to slip-stack at much higher intensities than previously reached. It is unknown whether the energy difference combined with the space charge effects of higher-intensity beams will result in increased particle losses. Slip-stacking simulations of the Fermilab Recycler aim to understand these intensity-dependent effects and provide the Fermilab neutrino program with high-intensity beams. The current simulations have been able to capture theoretically predicted instabilities in these beams, including demonstrating the existence of a multiple-bunch instability in the Fermilab Booster.

IMPACT: These simulations will help scientists understand the dynamics of intense particle beams, which will be used to create neutrino sources for the Long Baseline Neutrino Experiment (LBNE) at Fermilab and the upgrade to the Large Hadron Collider (LHC) at CERN. Beyond the realm of particle physics, the accelerator technology being developed at Fermilab has potential impact in basic energy applications, such as the transmutation of nuclear waste.

ALCF Contributions: Staff ported the team's hybrid C++-Python code to the BG/Q platform (William Scullin) and identified optimizations specific to the BG/Q architecture (James Osborn).

Presentations:

- "Accelerator Beam Dynamics on Multicore and GPU and MIC Systems," SIAM Conference on Parallel Processing for Scientific Computing, February 18–21, 2014, Portland, Oregon.
- James Amundson, "Transverse Instabilities in the Fermilab Booster," 16th Advanced Accelerator Concepts Workshop (AAC 2014), July 2014, San Jose, California.
- Eric Stern, "Developments in Synergia Space-Charge-Induced Resonance Trapping," 16th Advanced Accelerator Concepts Workshop (AAC 2014), July 2014, San Jose, California.

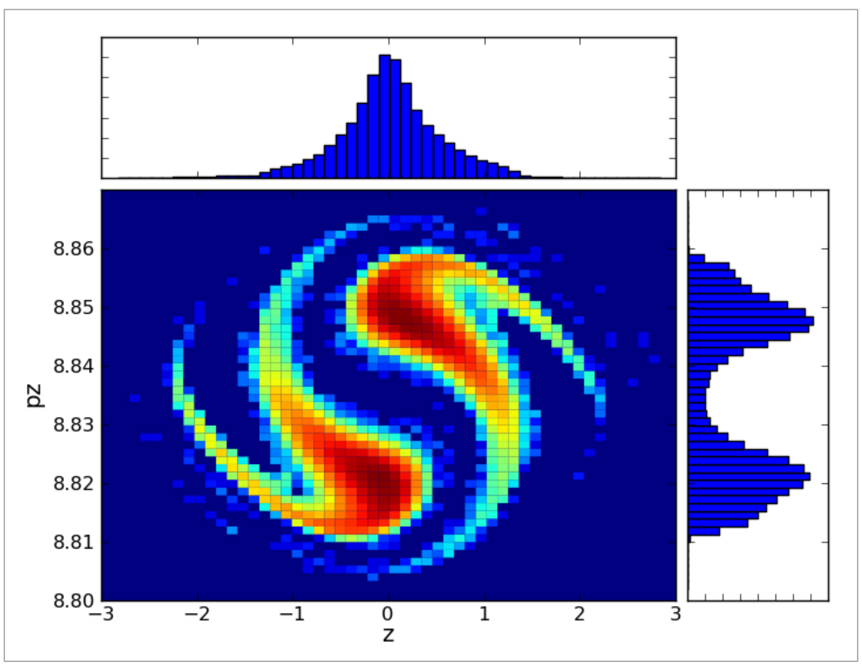

Figure 3.1 Slip-stacking simulation in the Fermilab Recycler. Two injected batches of bunches are combined to form a batch of high-intensity super bunches.

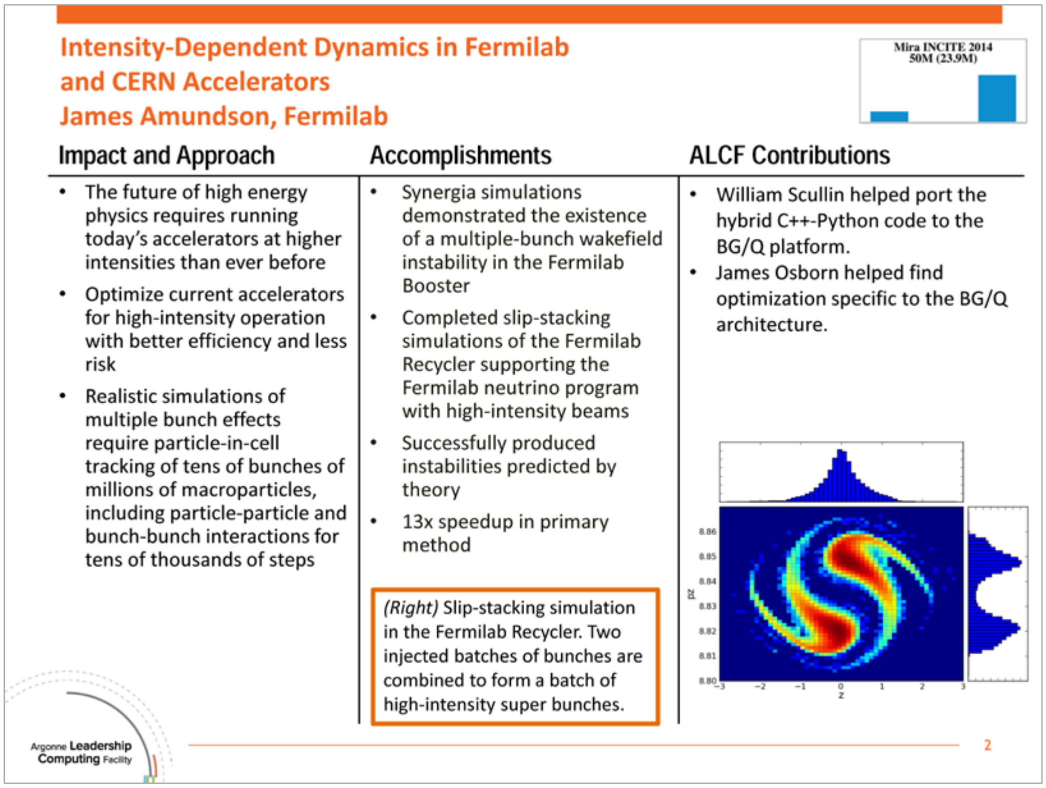

Ab Initio Quantum Liquid Water and Aqueous Ionic Solutions

Robert A. DiStasio, Princeton University

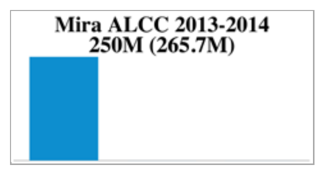

Obtaining a highly accurate and detailed understanding of the microscopic structure of liquid water would be critical to a number of fields, including energy storage, biochemistry, and the environmental sciences. For this project, researchers from Princeton University are using an innovative computational approach on Mira to study liquid water and aqueous ionic solutions at the atomic level with unprecedented accuracy.

At present, no experimental methodology exists that has directly obtained the microscopic structure of liquid water; however, computer-based simulations can furnish such structural information. Unfortunately, although ab initio molecular dynamics (AIMD) based on density functional theory is the most accurate and widely used computational methodology for modeling condensed phase systems, this approach has severe limitations when applied to liquid water.

To overcome these issues, the Princeton research team is combining novel algorithmic and theoretical developments to perform highly accurate atomistic simulations of liquid water on Mira. The team is also applying this approach to aqueous ionic solutions that are relevant to the design of clean energy materials, such as electrolytes for aqueous-ion batteries.

The researchers have completed several production-level path integral (PI) AIMD simulations of liquid water at the Perdew–Burke–Ernzerhof (PBE), PBE+vdW (van der Waals), and PBE0+vdW levels of theory. These simulations allowed them to accurately capture the main experimental isotope effects in the oxygen-oxygen, oxygen-hydrogen, and hydrogen-hydrogen radial distribution functions of liquid water. The team also carried out large-scale PI-AIMD simulations to investigate the proton transfer mechanism in the fundamental hydronium (H_3O+) and hydroxide (OH-) ionic solutions (the models of acidic and basic conditions, respectively). The proton transfer rates from these simulations were found to be in excellent agreement with experimental data (Figure 3.2).

In addition, the team's work has determined the importance of exact exchange and non-local vdW-related interactions in predicting the density ordering between crystalline ice and liquid water—an anomalous property of water that has presented a substantial challenge for theory to date. Researchers have also investigated the importance of vdW/dispersion interactions in predicting the structure of weakly bound molecular crystals at finite temperatures and pressures.

IMPACT: These simulations are providing detailed knowledge of the microscopic structure of liquid water and aqueous ionic solutions with unprecedented accuracy. Stored in a publicly available structural database, the results will serve as a valuable resource to inform further theoretical investigations and to advance clean energy applications, such as the rational design of aqueous-ion batteries.

ALCF Contributions: Staff collaborated with DiStasio's team on developing a new multi-threaded version of a 3D Poisson solver (Alvaro Vazquez-Mayagoitia). The solver is computationally expensive because of the exact exchange contribution of hybrid functionals like PBE0.

Publication: R. A. DiStasio, Jr., B. Santra, Z. Li, X. Wu, and R. Car, J. Chem. Phys. **141**, 085402 (2014).

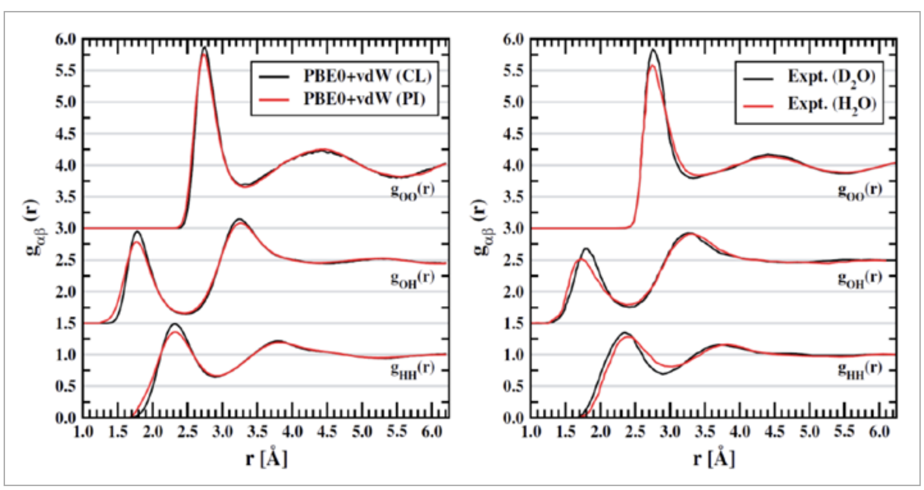

Figure 3.2 Radial distribution functions of liquid water from theory (classical [CL] and path integral [PI] AIMD simulations) and experiment.

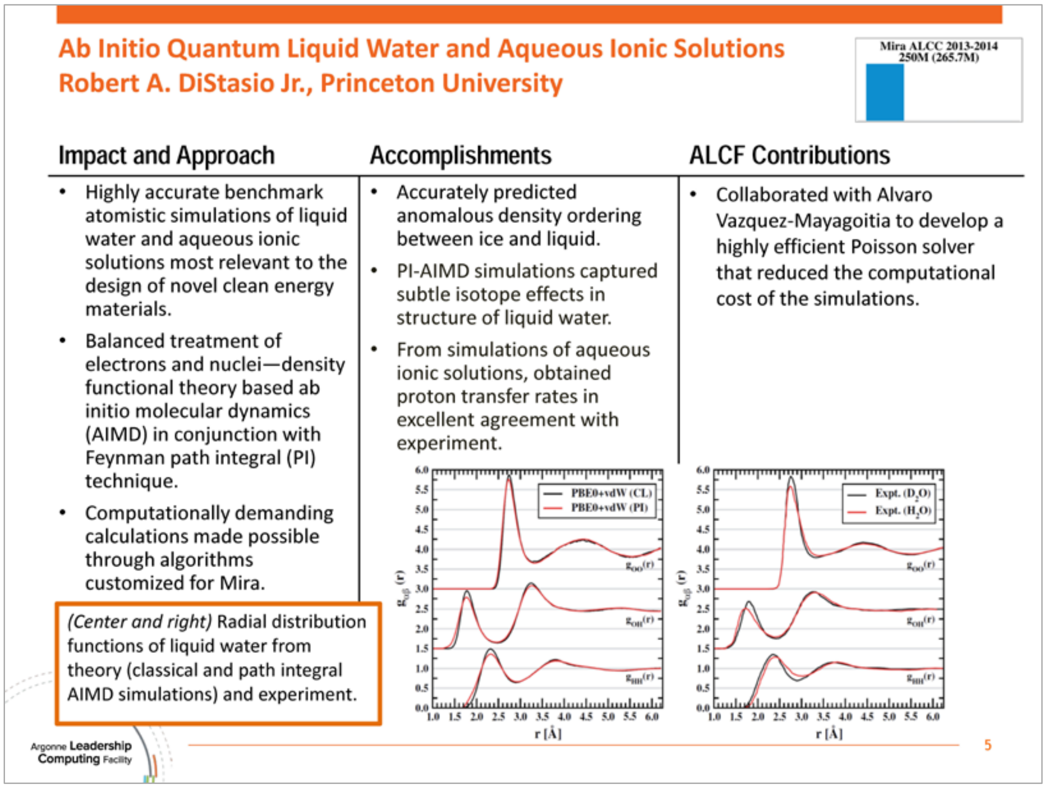

Cosmological Simulation for Large-Sky Surveys

Salman Habib, Argonne National Laboratory

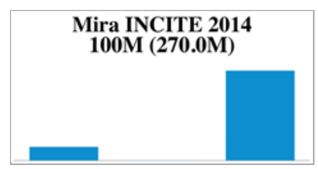

Modern precision cosmology substantially contributes to our investigation of dark matter and dark energy, the as-yet-unidentified fundamental constituents of our universe. This science depends on large-scale simulations to interpret data from large-scale cosmological surveys. These surveys measure the positions of faraway galaxy clusters over large portions of the observable universe; however, extracting fundamental physics knowledge from the statistics of the spatial distribution of these objects requires comparing the statistics to data from simulations. Furthermore, the development and testing of the software necessary to perform this data analysis, and the planning of the survey itself, require mock (that is, synthetic) sky catalogs generated by large simulations.

The completion of the OuterRim simulation—to date, the largest high-resolution simulation ever carried out in the world—meets the requirements of several upcoming surveys. The simulation was performed using the Hardware/Hybrid-Accelerated Cosmology Code (HACC), an N-body simulation code specifically optimized for large-scale simulations on today's largest supercomputers. Over one trillion particles were evolved by the simulation, using 32 racks of Mira for hundreds of millions of core hours over several years. Only a leadership-class machine has the resources to deliver a simulation at this scale. And even with current hardware, this simulation pushes the limit of what is possible with modern technology.

IMPACT: The OuterRim simulation is currently being analyzed to generate mock object catalogs for several different upcoming cosmological surveys. The simulation has already been used to create a synthetic catalog for the Dark Energy Spectroscopic Instrument (DESI) collaboration for survey optimization. In addition, it has been used to create strong lensing maps (Figure 3.3) to study galaxy clusters as observed by the Hubble Space Telescope.

ALCF Contributions: This simulation would not have been possible without the assistance of ALCF staff. The staff improved short-range-force evaluation algorithms and adaptive time stepping, greatly optimized the primary short-range-force kernel, and substantially improved HACC's I/O infrastructure (Hal Finkel, Adrian Pope, Venkatram Vishwanath, Vitali Morozov). Staff also helped produce visualizations from the simulation outputs; these visualizations greatly enhanced researchers' understanding of simulation results and improved many of the automated analysis routines (Joe Insley, Silvio Rizzi).

Publication: K. Heitmann, S. Habib, H. Finkel, N. Frontiere, A. Pope, V. Morozov, S. Rangel, E. Kovacs, J. Kwan, N. Li, S. Rizzi, J. Insley, V. Vishwanath, T. Peterka, D. Daniel, P. Fasel, and G. Zagaris, Comput. Sci. Eng. **16(5)**, 14 (2014).

Presentations:

- Salman Habib, "Extreme Scaling and Performance with HACC across Diverse Architectures: Future Challenges," University of California, Berkeley, May 7, 2014.
- Salman Habib, "Computing the Universe: How to Stuff a Supercomputer into a Laptop," Carnegie Mellon University Astrophysics Seminar, April 9, 2014, Pittsburgh, Pennsylvania.
- Katrin Heitmann, "Simulating the Universe with HACC," Higgs Symposium 2015, January 7, 2015, Edinburgh, Scotland.
- Katrin Heitmann, "HACCing the Universe: Or How I Learned to Stop Worrying and Love Supercomputers," Carnegie Mellon University Astrophysics Seminar, April 8, 2014, Pittsburgh, Pennsylvania.
- Katrin Heitmann, "Cosmological Simulations," Joint LSST-DES Workshop, Fermilab, March 26, 2014, Batavia, Illinois.
- Katrin Heitmann, "Large-Scale Structure Formation with Massive Neutrinos," Astrophysics Seminar, Fermilab, February 24, 2014, Batavia, Illinois.
- Katrin Heitmann, "Particles to Cosmos: Simulating the Universe with HACC," Computation Institute, Inside the Discovery Cloud series, March 19, 2014, Chicago, Illinois.
- Adrian Pope, "HACCing the Universe," Meeting of the American Physical Society, April 8, 2014, Savannah, Georgia.

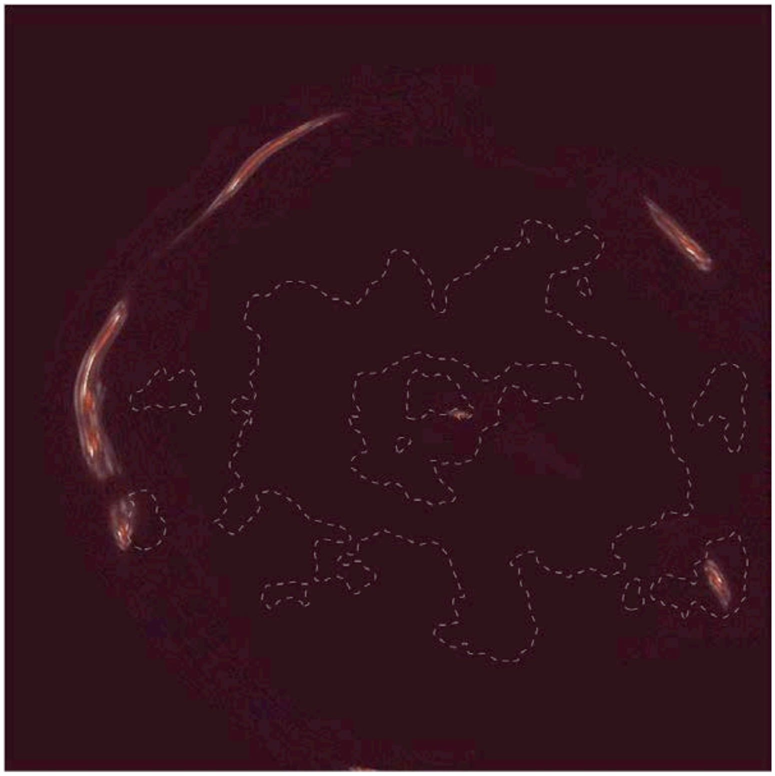

Figure 3.3 Simulated strong lensing image. White contours show the mass distribution, whereas the arcs are lensed images of background galaxies. The lens is a very large halo from the simulations. Image by Nan Li.

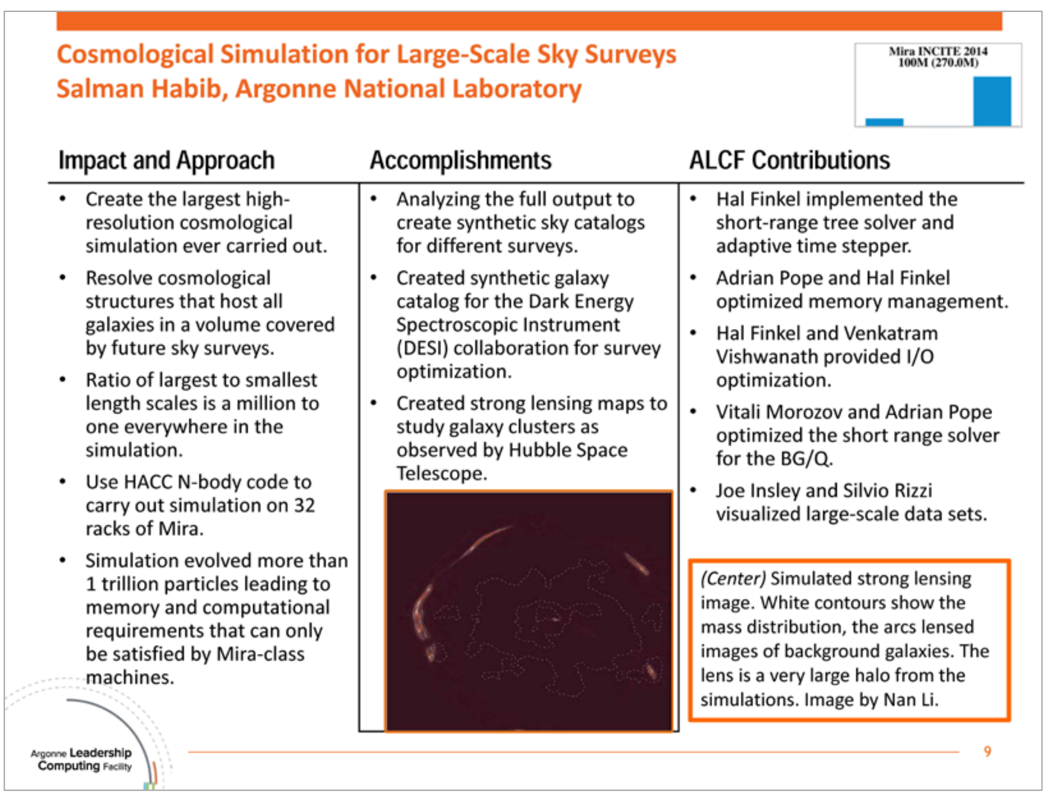

SiO_2 Fracture: Chemomechanics with a Machine Learning Hybrid QM/MM Scheme

James Kermode, King's College

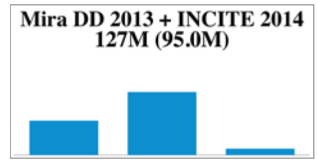

Fracture and cracking of brittle materials belongs to chemomechanical phenomena that scientists do not yet fully understand (despite their obvious relevance for realizing energy savings in such widespread industrial applications as mining). To develop predictive modeling and simulation tools for these processes, atomistic simulation offers an attractive avenue for rigorously incorporating the underlying physics. Unfortunately, these phenomena occur at length scales that are too significant such that it is impossible to treat all atoms on the same footing using quantum mechanics. A promising way to avoid this issue involves using hybrid quantum mechanics/molecular mechanics (QM/MM) methods that combine less expensive and less accurate force-field potentials for most atoms, with more expensive and more accurate quantum mechanics-derived potentials for a few reactive atoms. While multiple spatial scales can thereby be treated in a consistent fashion, multiple timescales can be accounted for through interpolation of atomic forces using statistical methods, such as machine learning. The latter approach, commonly known as the "Learn on the Fly" (LOTF) scheme, is of particular interest because it has never been used for the systematic study of larger systems.

Within an INCITE allocation at ALCF, King's College London researchers have been using the BG/Q supercomputer to model several hypothesized SiO_2 crack propagation processes. Access to Mira enabled a ten-fold increase in sampling time, significantly strengthening the numerical evidence for the findings. The project results represent an important advance toward a deepened understanding of stress corrosion cracking. In particular, scientists learned that silicon can crack when applying less energy than expected in the presence of oxygen (Figure 3.4). Experiments have confirmed the simulation results, and are of such scientific relevance that this work has already been published in *Physical Review Letters*. Upcoming results are currently being prepared for submission.

IMPACT: This study has deepened our understanding of crack propagation and clarified the role of oxygen in Si cracking. Simulation results have also been confirmed by experiment. New algorithms have been implemented and deployed successfully in the context of high-performance computing.

ALCF Contributions: Outreach activities helped the team's principal investigators understand the usefulness of the INCITE program (O. Anatole von Lilienfeld). Staff members helped port code and libraries as well as improve the scalability of the LOTF method (Alvaro Vazquez-Mayagoitia, Kevin Harms). Members of the London research team attended ALCF's Mira Performance Boot Camp in 2013.

Publication: A. Gleizer, G. Peralta, J. R. Kermode, A. De Vita, and D. Sherman, "Dissociative Chemisorption of O_2 Inducing Stress Corrosion Cracking in Silicon Crystals," Phys. Rev. Lett. **112**, 115501 (2014).

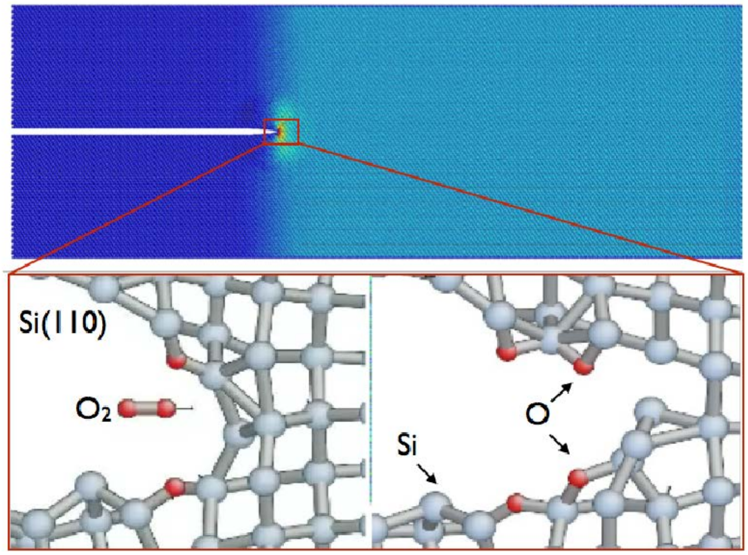

Figure 3.4 *(Top)* **Stress corrosion cracking in the Si (110) cleavage system. Model system, colored by principal stress, with red regions corresponding to high tensile stress and dark blue regions to zero; the red rectangle represents the region shown in the bottom panel.** *(Bottom)* **Crack tip structures obtained by QM-MM geometry optimization.**

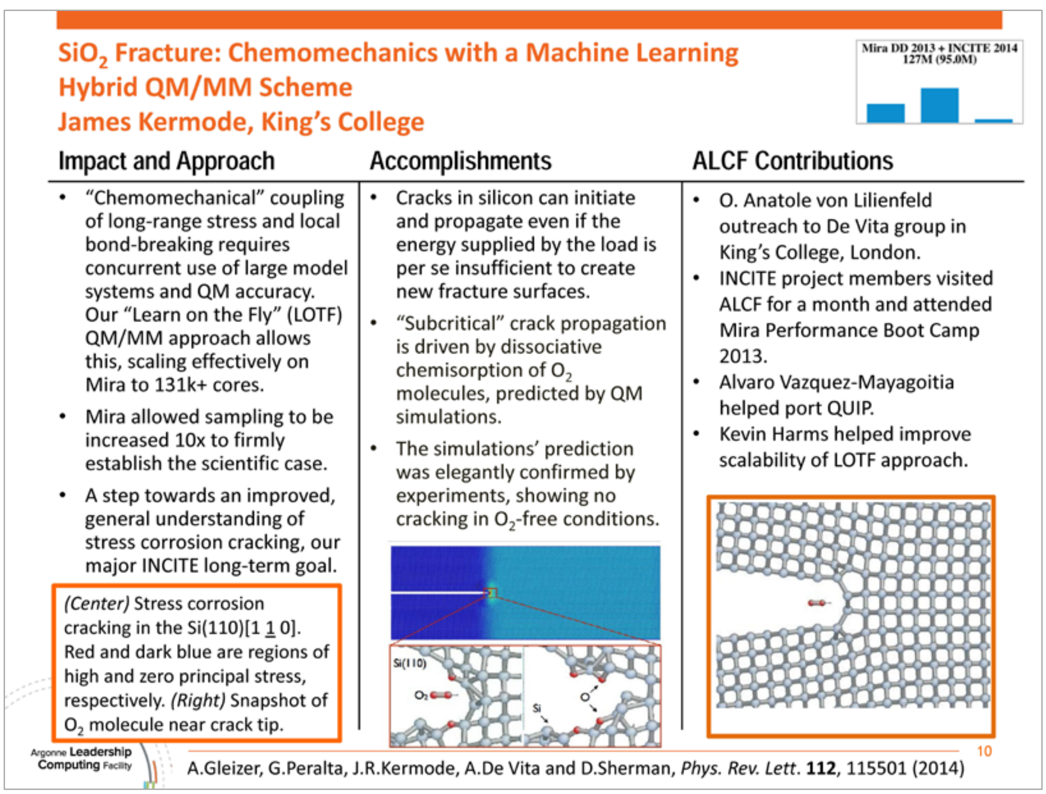

Direct Numerical Simulations of High Reynolds Number Turbulent Channel Flow

Robert Moser, University of Texas, Austin

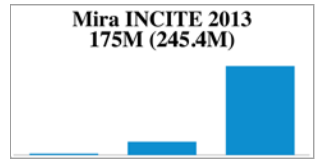

A substantial amount of the energy consumed by a moving vehicle is accounted for by the drag and dissipation of energy caused by turbulence as it moves through air or water. The same forces are at work as air or liquid moves through ducts or pipes. An estimated 10-20 percent of the world's energy consumption can be traced to this turbulence and the energy it dissipates. Leveraging the computational power of the Mira supercomputer at Argonne, Robert Moser and his University of Texas team have conducted the largest-ever direct numerical simulations of this fluid dynamics problem—simulations at a Reynolds number (the dimensionless ratio of inertial forces to viscous forces) of 5200 on a 15360 × 1536 × 11520 mesh. These runs generate a new reference data set that will remain useful for turbulence research for many years.

Of particular interest to Moser's team is the overlap region, where the viscous near-wall turbulence interacts with the outer-layer turbulences (Figure 3.5). The region is currently not well understood because simulations to date have not allowed for a high-enough Reynolds number to obtain the scale separation needed to shed light on the complexity of this multiscale turbulent structure. Researchers believe the simulations currently being performed have a high-enough Reynolds number to generate sufficient scale separation.

During the Mira Early Science Program period, Moser's team was able to get their code ported to Blue Gene/Q architecture with little effort. However, the code had a fixed global communication pattern and its data structures were not well suited for multi-threading. To capitalize on the Blue Gene/Q's capabilities, researchers modified their code.

IMPACT: As a result of this study, the most highly resolved (to date) channel flow data are now available to the public for calibrating turbulence models, developing new models, and validating existing ones. The data will allow faster, more effective studies of turbulence in a wide range of fields, including air flow over vehicles, improved vehicle surfaces, and reduced-drag piping and ducts. Data are available at http://turbulence.ices.utexas.edu/.

ALCF Contributions: Staff members facilitated the improved management of cache and execution threads, resulting in performance two times greater than that of the old code (Ramesh Balakrishnan, Jeff Hammond). By minimizing the inter-memory access between OpenMP threads, they achieved near-perfect (99 percent) OpenMP scalability.

Publications:

- M. Lee, N. Malaya, and R. D. Moser, "Petascale Direct Numerical Simulation of Turbulent Channel Flow on up to 786K Cores," in *Proceedings of SC13: International Conference for High Performance Computing, Networking, Storage and Analysis,* November 17–21, 2013, Denver, Colorado.
- M. Lee, R. Ulerich, N. Malaya, and R. D. Moser, Comput. Sci. Eng. **16(5)**, 24 (2014).

Presentation: Myoungkyu Lee, Nicholas Malaya, and Robert D. Moser, "Direct Numerical Simulation for Incompressible Channel Flow at Re_{tau} = 5200," American Physical Society Division of Fluid Dynamics, November 24–26, 2013, Pittsburgh, Pennsylvania.

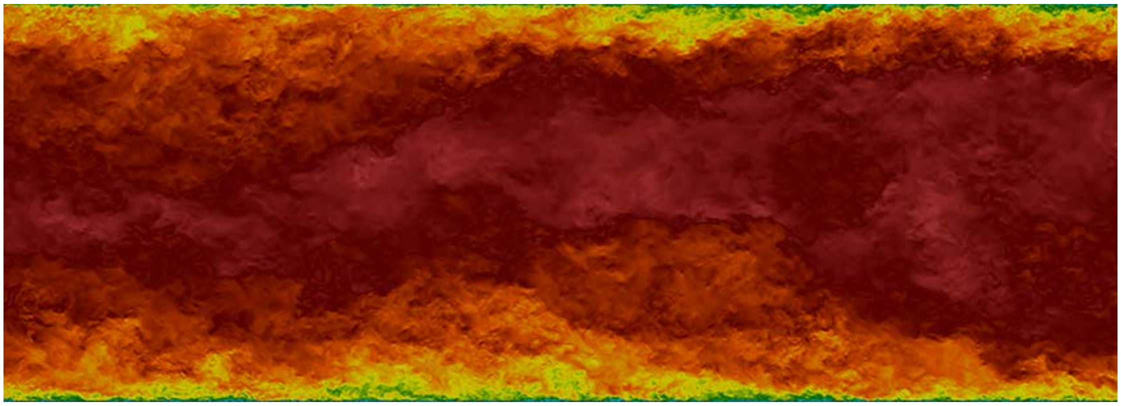

Figure 3.5 A visualization of the instantaneous streamwise velocity component over a section of the simulated channel.

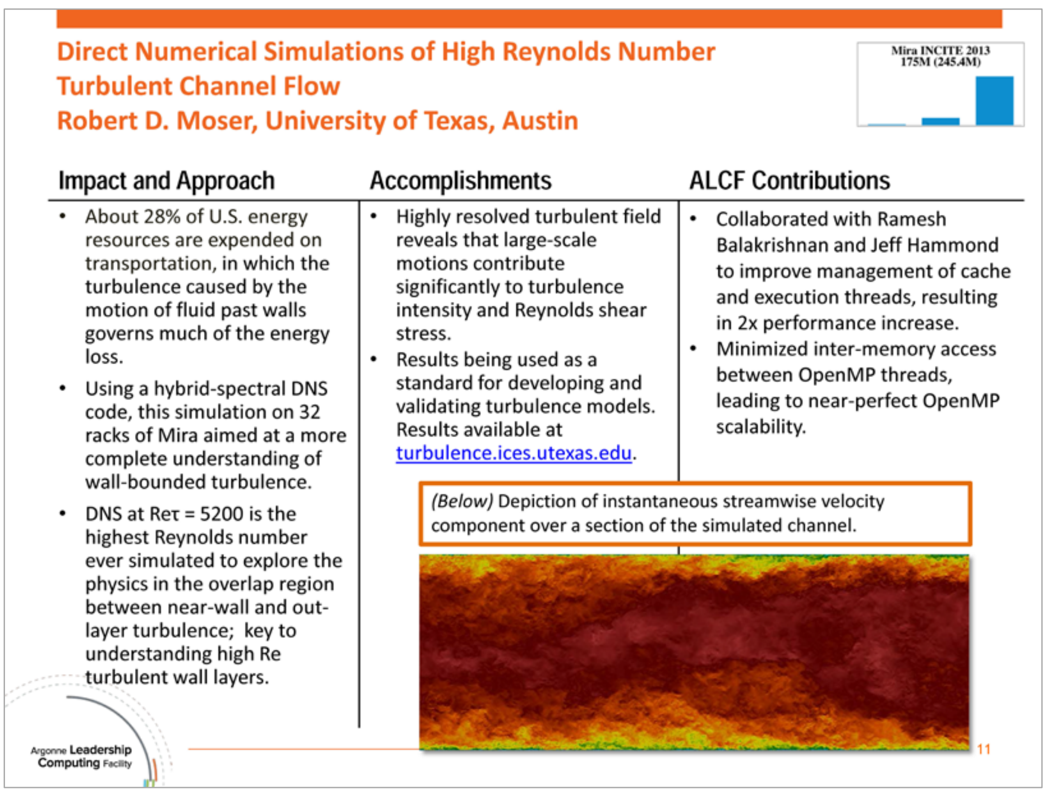

Usage of the INCITE and ALCC Hours

The INCITE 2014 program allocated 3.5 billion core-hours on Mira. The allocation usage is shown in Figure 3.6. Of the 40 INCITE projects on Mira, 18 used their entire allocation (or more). Four projects used more than 150 percent of their allocation. These projects used the extra core-hours to achieve additional milestones. The overuse of Mira was made possible through the use of the backfill queue (low priority) and an "overburn" policy that permitted projects to continue running capability-sized jobs after their allocation was completely exhausted. The duration of the "overburn" policy was chosen to be June 9, 2014 – December 1, 2014. Of the remaining 22 with outstanding allocations, 19 projects used more than 50 percent of their time and of those 19, 12 projects used more than 75 percent of their time. Three projects used less than 50 percent of their allocation. One of these projects accomplished most of its milestones with substantially less time. Two projects had unexpected computational difficulties. ALCF worked very closely with the teams in question to resolve their issues; however, as a result, these projects were substantially delayed in using their allotted time. A total of 3.9 billion core-hours were delivered to INCITE. The total number of INCITE hours delivered include 15.0 million core-hours that were delivered on Cetus because the scientific campaign of one of the INCITE projects could not be easily accommodated on Mira. The contribution from Cetus to ALCF total INCITE hours delivered was less than half a percent.

The 2013–2014 ALCC year was different from previous years. ALCC allocations were given on Intrepid from July 1, 2013, through December 31, 2013 (the end date for production runs on Intrepid). ALCF and ASCR decided to migrate Intrepid projects with unused allocations to Mira to provide them with an opportunity to continue their work and become familiar with the new systems. Currently, 18 projects have allocations on Mira with a total of 1.2 billion core-hours, including those that were migrated with an end date of June 30, 2014. The allocation usage is shown in Figure 3.7. Of these 18 projects, 11 used 90 percent or more of their allocation. Four projects used less than 50 percent of their allocation. Of those four, two experienced staffing issues and one had technical issues with PETSc that was passed onto the PETSc team. The remaining underutilized project was a joint project between Argonne National Laboratory and industry that was significantly delayed because of a large change in scope.

The 2014–2015 ALCC year is approximately halfway through its allocation cycle. So far, 21 projects have received allocations of 1.8 billion core-hours. Two of these projects were late ALCC submissions and received their awards in late November 2014. The allocation usage is shown in Figure 3.8. One of the projects has already used up its allocation and is running in the backfill queue, while another four projects have used up 50 percent or more of their allocation.

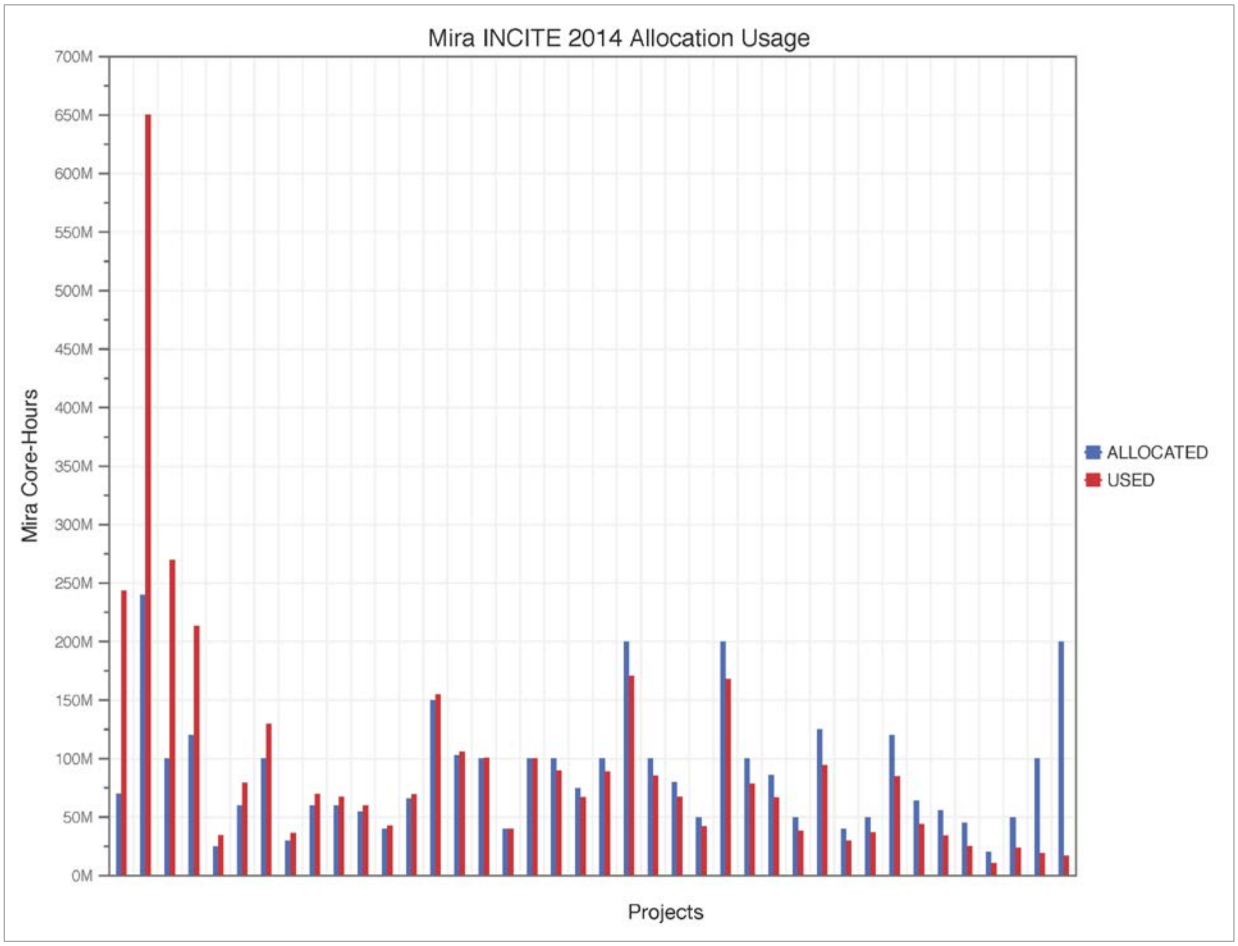

Figure 3.6 Mira INCITE 2014 Allocation Usage

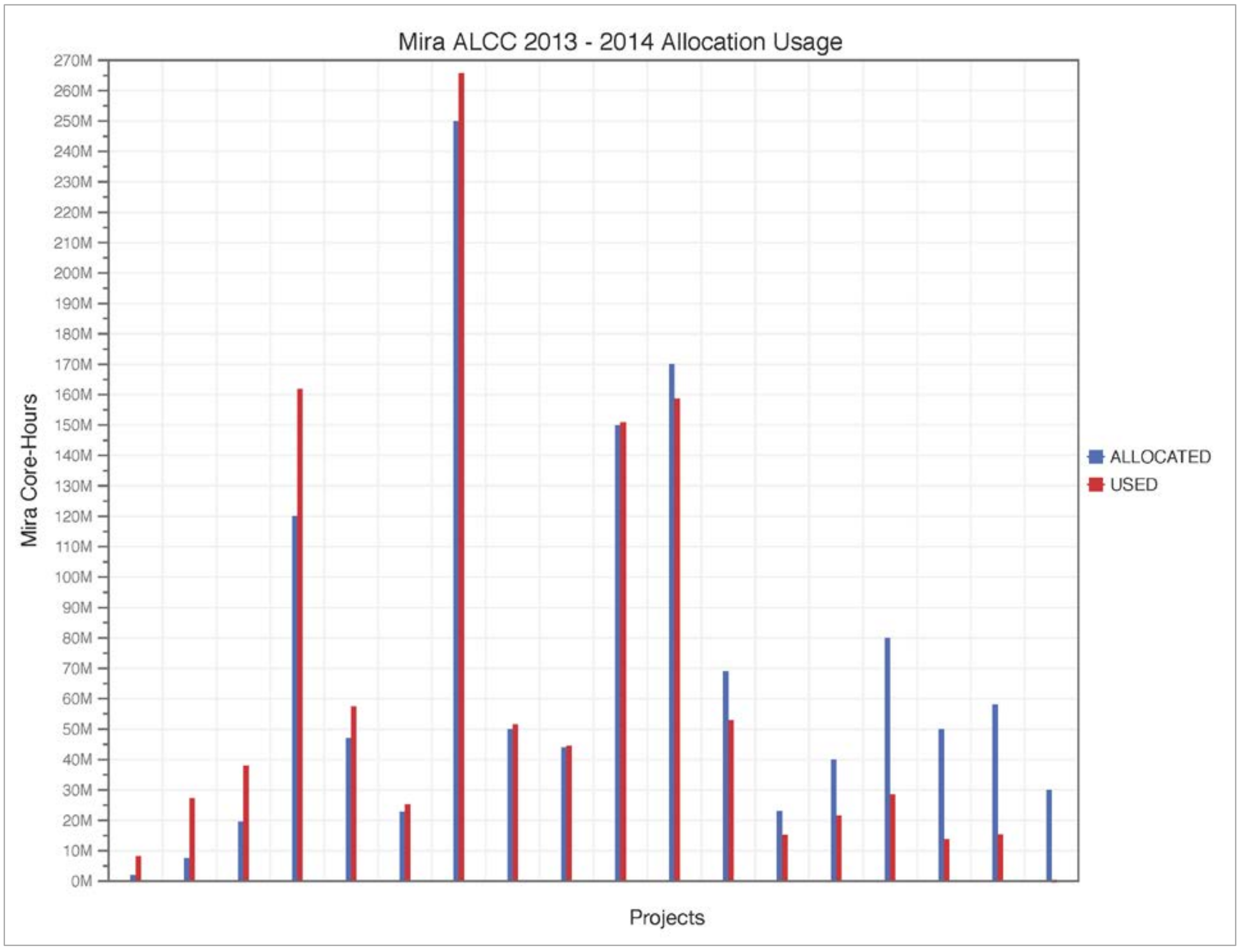

Figure 3.7 Mira ALCC 2013–2014 Allocation Usage

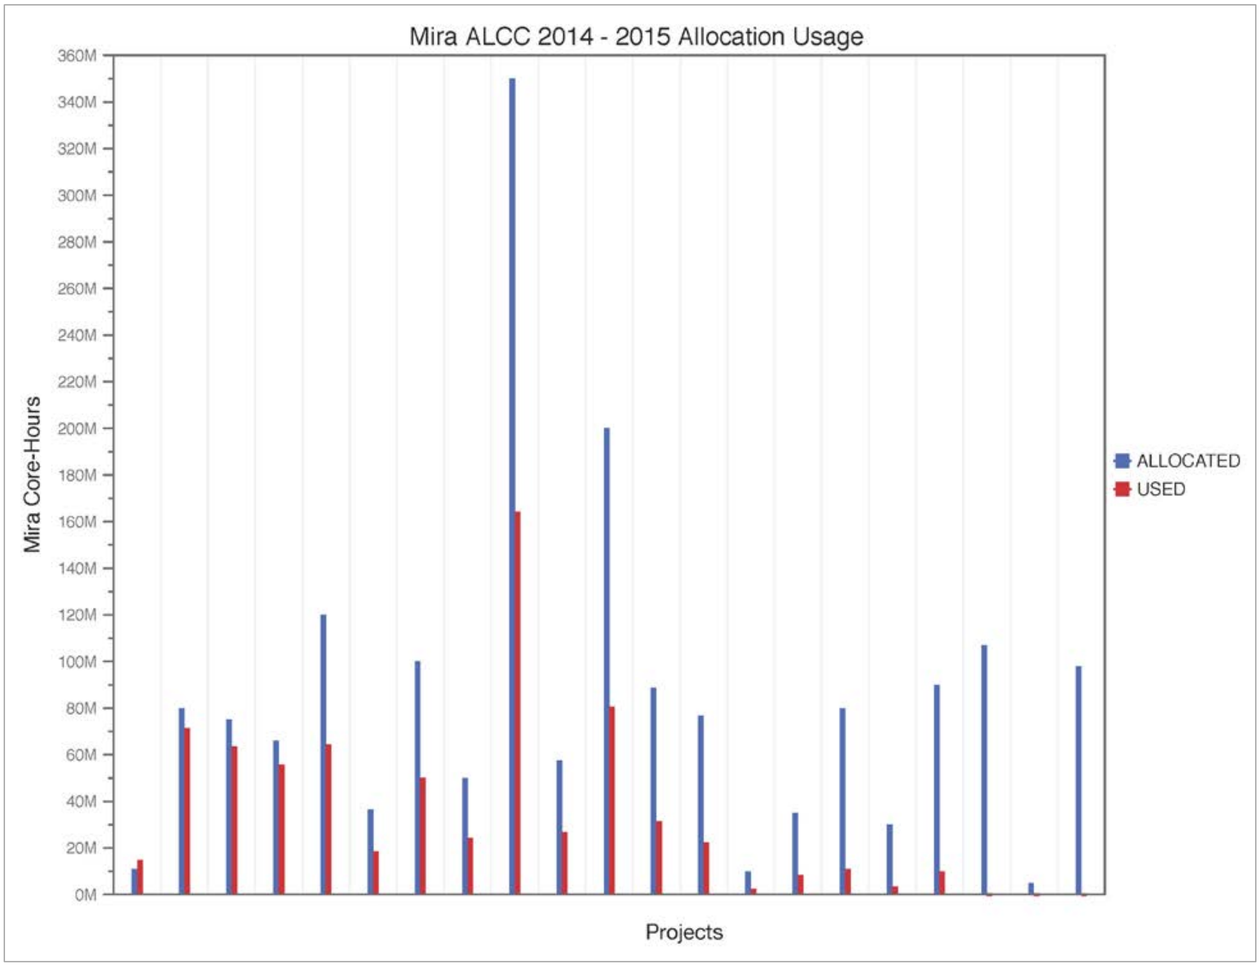

Figure 3.8 Mira ALCC 2014–2015 Allocation Usage

3.3 Allocation of Facility Director's Reserve Computer Time

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

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

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

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

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

ALCF also allocates time to projects that might still be some time away from an INCITE award, or that offer a “strategic science” problem worth pursuing.

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

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

Table 3.2 shows the number of projects and total time allocated in the DD program during 2014. By its very nature, the DD program is amenable to over-allocation since often time is left unused; however, it should be noted that these totals do not represent open allocations for the entire calendar year. A project might have a 1-million core-hour allocation that only persists for three months, but that 1-million core-hour allocation is counted entirely in the annual total core-hour number. Projects are not guaranteed the allocated time; rather, the time is provided on a first-come, first-served basis. DD projects run at a lower priority than INCITE or ALCC projects, which reduces the amount of time available for their use. Exceptions are made for some internal projects that support acceptance of new hardware or support of the users, which is in line with ALCF core mission.

Table 3.2 DD Time Allocated and Used on Mira, 2014

Projects	Mira
Allocated Core-Hours	1.1 B
Used Core-Hours	468.9 M

A list of the CY 2014 DD projects, including title, PI, institution, and hours allocated, is provided in Appendix B.

Figure 3.9 provides a breakdown of the CY 2014 allocations by the standard INCITE science domains.

Mira Discretionary Allocation Active in 2014
197 Projects, 1.1 B Core-Hours

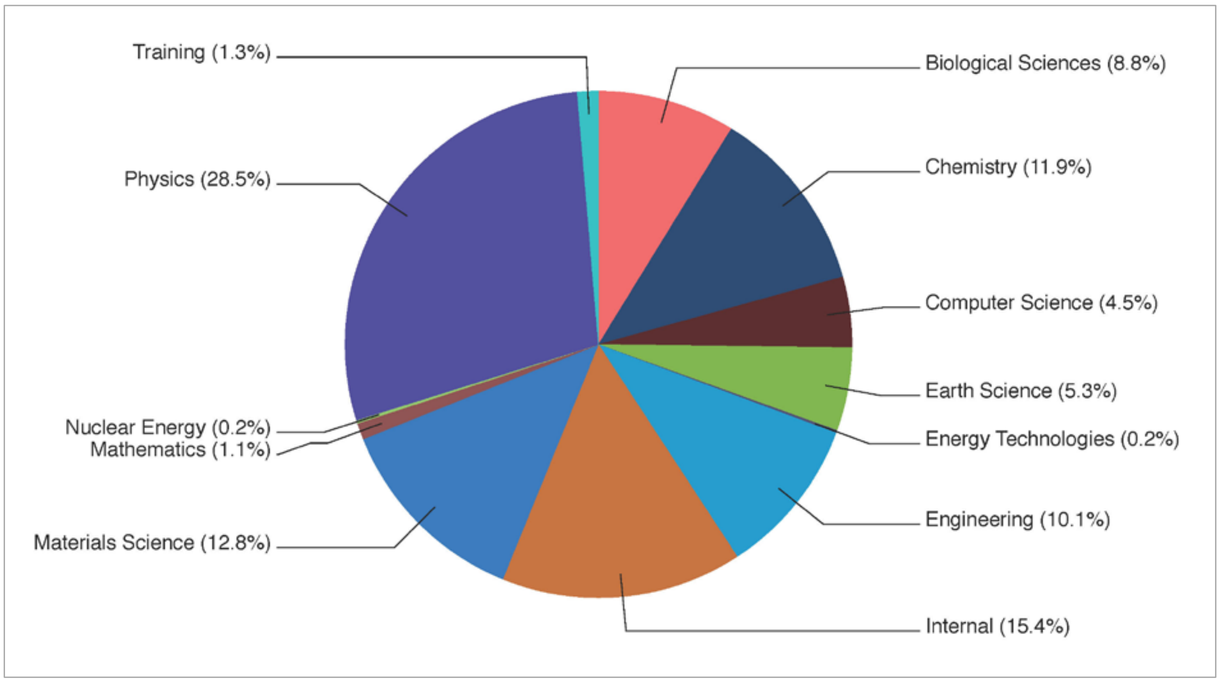

Figure 3.9 CY 2014 DD Allocation by Standard INCITE Science Domains

Conclusion

ALCF continues to enable scientific achievements, consistent with DOE's strategic goals for scientific breakthroughs and foundations of science, through projects carried out on facility machines. Researchers participating in projects using ALCF resources published 157 papers in CY 2014. ALCF projects have had success in a variety of fields, using many different computational approaches. They have been able to reach their scientific goals and successfully use their allocations. A number of the projects and PIs have subsequently received awards or have been recognized as achieving significant accomplishments in their fields.

ALCF delivered the following core-hours to the allocation programs in CY 2014: 3.9 billion to INCITE, 1.4 billion to ALCC, and 468.9 million to DD. The DD Reserve has been used not only to develop INCITE and ALCC proposals but also to conduct real science of strategic importance and to drive development and scaling of key INCITE and ALCC science applications. Excellent ALCF support and solid, high-performing ALCF resources have enabled INCITE and ALCC projects to run simulations efficiently on HPC machines and achieve science goals that could not otherwise have been reached.

Section 4. Innovation

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

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2014.

4.1 Improvements in Software Infrastructure

ALCF has worked to improve software that is of broad interest to the ALCF user community through innovative applications of code optimization, parallelization, and software engineering.

4.1.1 LAMMPS Modernization

Challenge: LAMMPS (Large-Scale Atomic/Molecular Massively Parallel Simulator) software was inefficient at leadership scale.

Approach: From approximately February 2013 through August 2014, ALCF, in partnership with IBM, has engaged in several efforts to improve the performance of the LAMMPS software at leadership scale in the areas of computation, communication, and input/output (I/O). The collaboration involved several scientists at Argonne, as well as researchers from other institutions, working under ALCF direction and leadership. Of all of the efforts attempted, implementation of the following four improvements resulted in significant performance improvements. IBM worked closely with ALCF on the computational and communication optimizations and with Argonne's Mathematics and Computer Science division on the I/O optimizations, leveraging work performed under both the Scientific Discovery through Advanced Computing (SCIDAC) program's Scalable Data Management, Analysis, and Visualization (SDAV) and the DOE Robust Analytical Models for Science at Extreme Scale (RAMSES) projects.

Improvement #1. The first successful target of optimization was intra-node computational performance, in which a technique used for OpenMP (Open Multi-Processing) thread data reduction was enhanced to utilize the processor cache more effectively, on average cutting the data reduction time in half. Because most pair-wise force fields and the particle-particle-particle mesh (PPPM) long-range electrostatic solver use these thread data reductions extensively, it was a common bottleneck.

Improvement #2. The second successful target of optimization was the communication bottleneck for the PPPM, where the legacy point-to-point communication was replaced with a message passing interface (MPI) collective. The baseline performance degraded nonlinearly at scale, so for simulations that needed to do PPPM computations at significant scale, most of the wall time was in the PPPM. Significant code changes were needed to use the collective effectively; however, once successfully implemented, the wall time speedup was nearly 3x for some large simulations (e.g., 16 racks of Mira).

Improvement #3. The third successful target of optimization was the introduction of MPI-I/O. The LAMMPS legacy code made use of both a serial and a simple multi-file parallel native I/O implementation for restart and dump files, which could be a significant portion of the wall time based on the size of the simulation and the resolution a user wanted for analyzing the evolution of the system over time. The benefit of adding MPI-I/O is the improvement it delivers in I/O performance as well as the convenience of a single output file, avoiding the post-processing time needed when the native multi-file parallel native I/O is used. Several restart and dump style functions required code changes; these were made more complex by the need to pre-allocate the file for optimum general parallel file system (GPFS) performance. Furthermore, for text-based dump styles, a general change to all text formatting to parallelize the computation resulted in potentially huge speedups for the I/O (up to 400×) relative to the size of the simulation and number of processors. Specifically, for Mira and its use of the GPFS file system, the ROMIO implementation of MPI-I/O in MPI Chameleon (MPICH) had significant performance problems on Mira with its use of GPFS, so several code optimizations were made to ROMIO to speed up I/O significantly. In summary, in terms of the overall performance improvement to I/O in LAMMPS, there is very significant improvement over the serial code; however, at this point, there is only modest improvement over the native parallel code because of an ongoing issue with parallel file opens on Mira's GPFS cluster.

Improvement #4. The fourth successful target of optimization was the implementation of OpenMP threads for the REAX/C force field. This force field has the ability to model bond formation and destruction—but only at a very high degree of computational complexity and cost. Therefore, there has been extensive demand for improvement to the node computation time; nevertheless, because of the computational complexity and the number of disparate loops needed to be threaded because of different atomic interaction functions, this threading had not yet been completed. By avoiding shared variable conflicts among threads, performing thread data reductions, reworking algorithms, and utilizing the existing USER-OMP package within LAMMPS, dramatic speedups were achieved. This optimization involved extensive collaboration at the code-level with Sandia National Laboratories and Lawrence Berkeley National Laboratory. For some simulations, a threading efficiency of up to 80 percent and a time-to-solution speedup of nearly 3× were measured on Mira.

Impact/Status: Improvements to the LAMMPS source code changes were accepted into the open-source repository at Sandia National Laboratories. Changes to speed up the ROMIO implementation of MPI-I/O in MPICH were accepted as part of MPICH 3.1.2. Both are available for others to utilize, including IBM's product development.

4.1.2 Scaling and Validation of OpenFOAM

Challenge: OpenFOAM (Open Source Field Operation and Manipulation) did not scale well when computing incompressible Navier-Stokes equations.

Approach: The OpenFOAM Computational Fluid Dynamics (CFD) Toolbox is a free, open-source CFD software package that has a large user base across many areas of engineering and science, in both commercial and academic organizations. Based on user input, ALCF improved the scaling and validating of OpenFOAM for large eddy simulations running on Mira. OpenFOAM now scales up to 8192 cores when running large eddy simulations of the incompressible Navier-Stokes equations. Figure 4.1 is a visualization of the resulting simulation of a periodic hill, where the computational domain consists of 512 × 256 × 256 finite volumes. The flow is from the left with a mean inflow velocity of 5.6 m/s, and it accelerates along the length of the domain. The color plot in the volume shows the mean stream-wise velocity with the two recirculation zones. The central plane (normal to the z-axis) shows the viscosity due to the sub-grid scale model; the plane at the top of the domain (normal to the y axis) shows the turbulent kinetic energy; and the plane that is close to the outflow (normal to the x-axis) shows the fluctuating velocity component.

Several enabling technologies and contributions have made this work possible:

- ALCF use of IBM's libhpm and libmpihpm libraries for dynamically linked executables, allowing for profiling with Blue Gene hardware counters;
- The Clang/LLVM (low-level virtual machine) compilers, built by ALCF staff with improved vectorization capabilities;
- A test node (with 1 TB of RAM) made available by ALCF operations for developing large-scale runs with 64-bit indexing for partitioning the meshes and for mapping fields from coarse meshes to fine meshes;
- The ALCF staff's investigation and resolution system-level issues;
- The SANDIA/Trilinos team's work with ALCF on implementing the Trilinos-based ML solvers in OpenFOAM;
- The DOE Atmosphere to Electrons (A2e) High Fidelity Modeling Group, which encouraged and informed this work by making demonstration of the scaling and validation of OpenFOAM a high priority; and
- The team at Vestas, which provided motivation and input base on their interest in using OpenFOAM in their INCITE project.

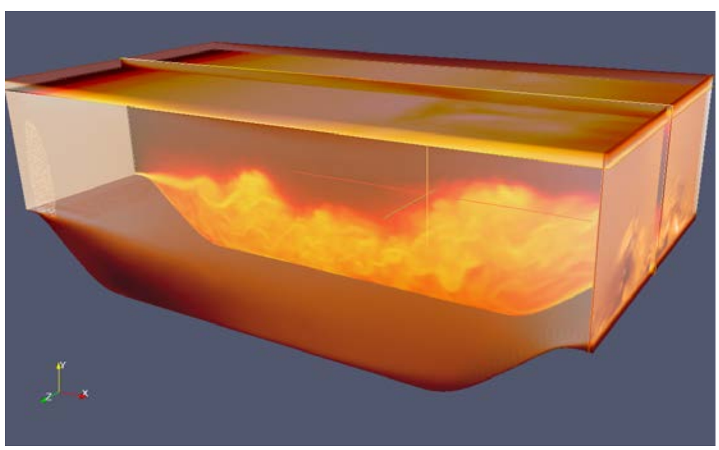

Figure 4.1 Periodic Hill Generated by OpenFOAM

Impact/Status: All of the OpenFOAM improvements are available to ALCF users and will be fed into the open-source effort. ALCF is evaluating the OpenFOAM GAMG Geometric Agglomerated Algebraic Multigrid Solver) solvers for possible inclusion. ALCF is also working to scale the same simulation to 16K and 32K cores of Mira.

4.1.3 CONVERGE™ Improvements

Challenge: CONVERGE™, a widely used commercial computational fluid dynamics (CFD) code for engine simulation, did not run on high-performance computing (HPC) architectures, including Blue Gene.

Approach: The Argonne Energy Systems division has been working with Convergent Science, Inc. (CSI), developers of CONVERGE, for several years to improve engine and fuels simulations. About three years ago, ALCF began working with Argonne scientists and CSI industrial partners to improve the CONVERGE code. The aim was to improve CONVERGE so that it could run efficiently on high-performance computing architectures like IBM's Blue Gene.

The size of normal production runs of internal combustion engine simulations using CONVERGE on cluster computers is of the order of 64 to 256 cores, using grid sizes between two and five million cells. However, as simulation requirements grow in size and complexity, the use of leadership computing resources can play a vital role in the engine design process by reducing the timescales to obtain and analyze simulation results.

In working with CSI, ALCF was able to profile and detect the main bottlenecks preventing scaling, and ALCF staff initiated a plan to solve the more restrictive ones (memory bugs at scale, long I/O times, etc.) progressively. Initial tests were performed on Surveyor, Intrepid and Vesta and continue this year on Mira under several Director's Discretionary awards. CSI granted access to a number of source codes, which enabled a much easier and faster analysis than possible in the early days of CONVERGE development when code access was restricted. An LDRD award involving the Energy Systems division and the CELS directorate funded some of this work.

Initially, ALCF staff focused on improving the I/O performance of the CONVERGE solver, a major bottleneck when trying to run CONVERGE at scale on massively parallel machines. First tests covered runs on a small case with a 20,000-cell grid and on a larger case with a 35-million-cell grid. In both cases, ALCF observed that I/O performance was mainly driven by the time required to write the output for post-processing visualization and the restart files. Before ALCF modifications, this I/O could take a half-hour or more (for a normal restart file). After ALCF improvements, it takes two to four seconds. First, ALCF optimized Blue Gene I/O performance by having only a single MPI rank read the restart data and then broadcasting it to all other ranks. This resulted in over 20x speedup when restarting a typical case on 2,048 cores. Second, ALCF implemented MPI I/O in order to speed up writing of restart and post processing files, resulting in an over 600x speedup in the write operations on 2048 cores and bringing write times down from minutes to seconds.

Another major achievement was ALCF's implementation of an improved algorithm for load-balancing chemical kinetics calculations. In the original version of CONVERGE, the chemistry work was distributed based on equalizing the number of CFD cells being solved across MPI ranks. This arrangement resulted in significant load imbalances since the cells have vastly differing computational times associated with chemical kinetics calculations, particularly during ignition. ALCF developed a novel scheme wherein CFD cells are weighted by the computational effort associated with the chemical kinetics. The software then performs the load balancing in such a way that the overall computational effort associated with chemistry for each MPI rank is nearly equal. This scheme was implemented by CSI and resulted in an over 8x improvement in load balancing and about 3.4 times faster time to solution near ignition for a test case run on 2,048 cores. With this improvement to load balancing, the CONVERGE code now scales with over 90 percent efficiency on 2,048 processors, and with 80 percent efficiency on 4,096 processors, with a mesh size of around nine million cells.

Additional ALCF modifications focused on the introduction of OpenMP directives to take advantage of the hardware support for multiple threads per core on Mira nodes. Although preliminary work has led to a modest five percent speedup, work is ongoing to move to a hybrid version of the code that will take advantage of high core/node architectures and multi-threading through more extensive implementation of OpenMP. ALCF staff also wrote a tool that reads code from little-endian machines. The tool has been very helpful in preparing for a public workshop and will be useful in general to enhance code portability between Linux clusters (which use little-endian format) and Blue Gene systems (which use big endian format) through standardizing the format of restart files. Finally, ALCF identified load-balancing issues and worked with CSI staff to introduce additional input variables to improve load balancing at scale.

In the area of visualization, ALCF staff identified several areas for improvement in the post-convert tool used to convert data from its native format to the VTK (visualization toolkit) format. In one case, for parcel data, ALCF discovered that the visualization tool ParaView could not access the variables associated with each parcel in the VTK file. A simple change in the converted file to designate the data arrays as point data rather than cell data corrected this problem. ALCF also worked with the developers to include additional information, such as crank angle and simulation time, into the VTK files in order to better visualize that information in ParaView. In addition, while investigating VTK data produced with CONVERGE's post-convert utility, ALCF discovered a bug in the unstructured grid reader in the latest version of ParaView (4.2.0) when run in parallel. ALCF submitted a bug report to the ParaView developers, who are working on a fix. Further improvements to the post-convert tool to enable faster I/O have been discussed with the CSI team, including who will implement the suggested changes.

Impact/Status: ALCF has improved the CONVERGE software and visualization of CONVERGE simulations in ParaView. Both sets of improvements are available to the community at large through CSI. Future ALCF efforts are focused on the application of the improved code to run 35-million-cell and 200-million-cell grids.

4.1.4 Optimized GEMM Library

Challenge: IBM's ESSL GEMM library is optimized for large square matrices; lower performance is obtained for long-skinny rectangular matrices. However, GEMM operations on long-skinny rectangular matrices are common in many scientific applications such as Density Functional Theory.

Approach: Under the guidance of ALCF staff, the team at IBM's Thomas J. Watson Research Center developed an optimized GEMM library for Blue Gene/Q that is called JAG_GEMM. (JAG is the main developer's initials; GEMM stands for general matrix multiply.) This work was funded by ALCF. On average, JAG_GEMM outperforms the IBM ESSL library by a factor of two for medium to large matrices. The library supports both real and complex double-precision data types. This large performance boost is available for most of the MPI modes on Blue Gene/Q. The performance is obtained through the use of assembly code and by leveraging the L2 prefetcher API.

Algorithms that make heavy use of matrix multiplies on medium-to-large matrices are good candidates for using JAG_GEMM. These candidates include codes that make use of parallel dense linear algebra such as planewave (PW) density functional theory (DFT) codes. One success story is the PW DFT code of a new 2014 INCITE project, Petascale Simulations of Self-Healing Nanomaterials. Through the use of JAG_GEMM and a few other libraries, the project obtained 50 percent of peak performance on Mira for its linear scaling PW DFT code. Another code that has leveraged JAG_GEMM is GPAW (general projector-augmented wave), which is being used in an INCITE project. GPAW is a real-space DFT code that also spends a significant amount of execution time in parallel dense linear algebra.

Impact/Status: Although the library is not officially supported by ALCF, it is available on a per-request basis. Several additional users have expressed interest in using JAG_GEMM, including the lead developer of Qbox and a developer of CPMD. Qbox and CPMD are both PW DFT codes used by an INCITE 2014 project. For the INCITE 2015 allocation year, there is one project that will likely leverage GPAW enhanced by JAG_GEMM.

4.1.5 Enhancing Performance Analysis

Challenge: The HPC Rice Toolkit did not support in-depth analysis of MPI + OpenMP applications and did not scale well on Mira.

Approach: ALCF is working with Rice University to enhance the capabilities of the Rice HPCToolkit for leadership computing applications. HPCToolkit is an integrated suite of multi-platform tools for measuring and analyzing program performance on scalable parallel systems. HPCToolkit has been used on ALCF Blue Gene systems over the last several years to analyze and improve application performance.

Based on this experience, several areas for enhancement have been identified and are being implemented in the upcoming versions of HPCToolkit. This work includes utilizing the recently proposed OpenMP Tools interface to support in-depth analysis of MPI + OpenMP applications, advanced support for analysis of codes with in-lined and templated functions, implementation of a novel blame-shifting analysis to attribute lock contention and load balance problems to the responsible threads, and enhanced scalability of HPCToolkit data collection and trace viewer.

Status/Impact: All ALCF improvements are fed back into the HPCToolkit code base and released to all HPCToolkit users, thus benefiting many outside of the ALCF user community.

4.2 Improvements in ALCF Operations

ALCF has implemented several improvements in systems that monitor and manage computing, storage, and networking resources.

4.2.1 Analyzing Machine State Using the Machine Time Overlay Graphic Representation

Challenge: Analyzing the state of Mira over time for characteristics like voltage, reservations, and failures was difficult without having a visual representation.

Approach: ALCF has developed a graphic representation of machine state called a machine time overlay (MTO). The MTO tool has been used for more than a year to represent the time and exit state of jobs running on Mira. This advance was reported in the 2013 OAR. The use of MTO has been expanded to show:

- Current
- Voltage
- Power (calculated from current and voltage)
- Task usage
- Reservations
- Various failures: reliability, availability and serviceability (RAS) errors, boot issues, and collisions
- Draining

Figures 4.2–4.4 show examples. These graphics are automatically generated and available daily to ALCF operations staff and are reviewed in a weekly operations meeting. The visual representation helps ALCF staff find unusual events quickly and gives the Operations staff a timeframe for investigating further.

ALCF Operations staff reviews the job usage version of the MTO graph on a weekly basis to find anomalies in scheduling and to correlate any outages with jobs that may be affected. ALCF used this graph when adding the four new racks of Cetus to ensure that overall coolant flow was not drastically reduced and that excessive hot spots did not develop. The graph can easily show jobs that are generating vast amounts of heat compared to others (which will aid in debugging problems if scheduling is changed) to taking power usage into account. Another use for this graph is to show not only where a job executed but also where and when the machine drained for that job (Figure 4.4). This graphic revealed a scheduler anomaly where it toggled inefficiently between scheduling jobs, causing larger drain windows. In investigating this issue, ALCF was able to tune the scheduler to reduce the toggling and improve utilization.

Impact/Status: ALCF is just starting to scratch the surface in terms of the data that can be interpreted by these new graphs. Already, the graphs have provided insight into how Mira is working. For example, Figure 4.3 shows the effect of a coolant failure and the recovery. It would be difficult to understand how individual nodes recover at different rates without this representation.

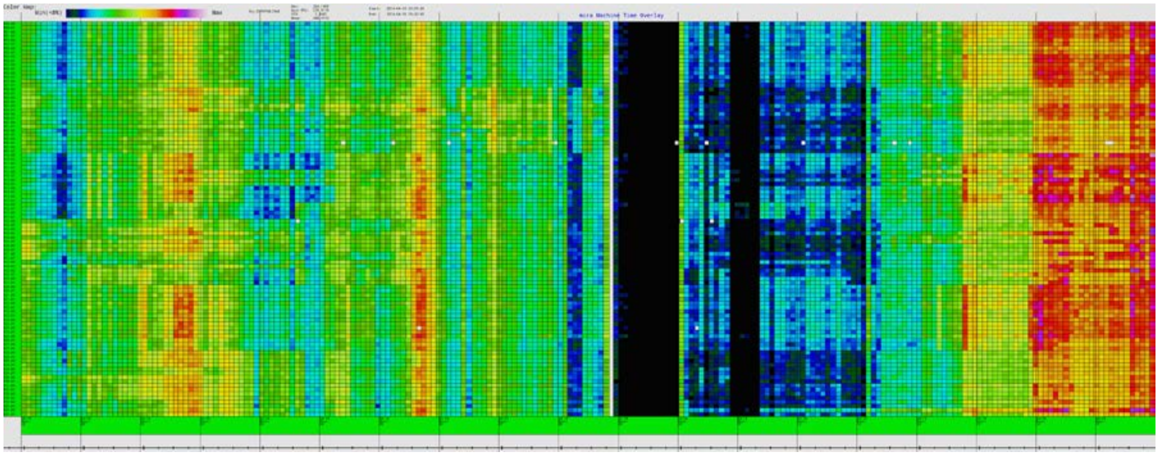

Figure 4.2 Input Voltage

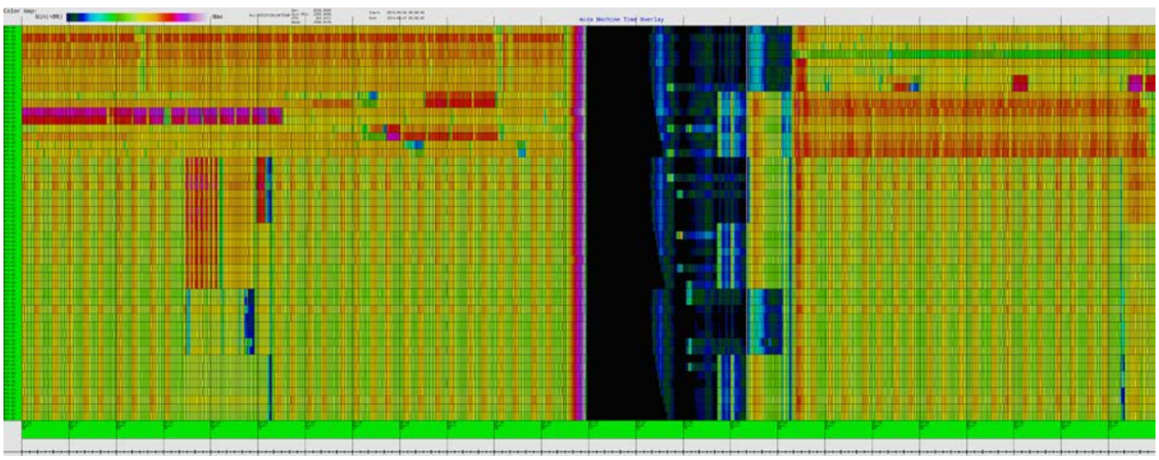

Figure 4.3 Output Coolant Temperature (Coolant Failure of 2014-09-05)

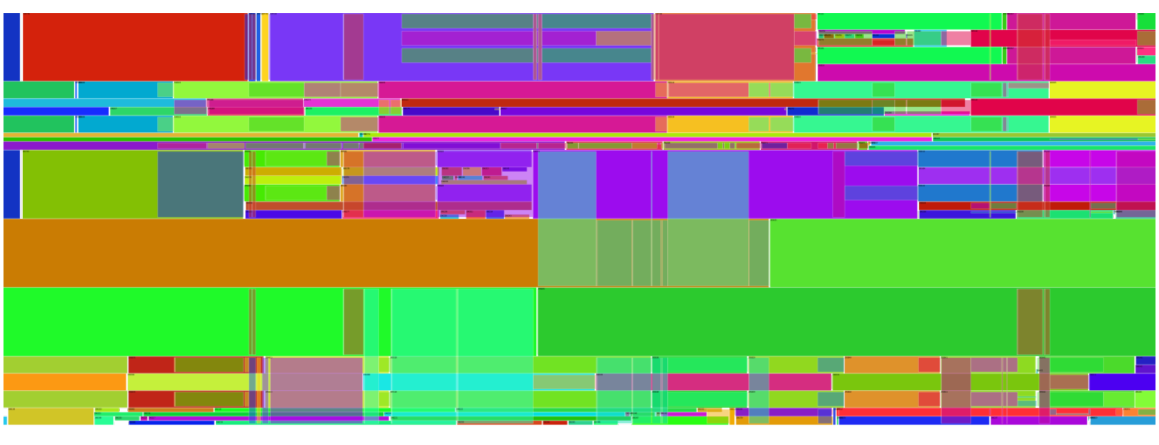

Figure 4.4 Draining Graph

4.2.2 Burn Rate Fan-Out Graph

Challenge: Analysis and projection of an INCITE project's usage was manual and time-consuming.

Approach: ALCF developed a system to generate graphs that show an INCITE project's projected usage based on a number of scenarios. This graph is used in ALCF's biweekly scheduling meeting to get a better idea of how close each project's usage and the overall INCITE usage are tracking to the INCITE target line and also to predict how projects and overall usage will track in the future. The various lines show projections of the slope over various periods to help understand the project usage characteristics. ALCF can then tune the scheduling algorithms to encourage usage that will direct the burn line closer to the goal. If the slope of the burn rate goes too high, it indicates that non-INCITE jobs are probably suffering, so staff can adjust the INCITE queue priorities. If burn rate is too low, ALCF can boost INCITE priority to encourage usage.

Figure 4.5 provides an example. The fan-out lines show different scenarios where:

- Red is the target line.
- Blue is the current burn rate.
- @60 and @80 are projections if the machine were loaded at, respectively, 60 percent or 80 percent INCITE jobs.
- @curr is the slope of the overall period.
- @week is the slope over the last week.
- @month is the slope over the last month.

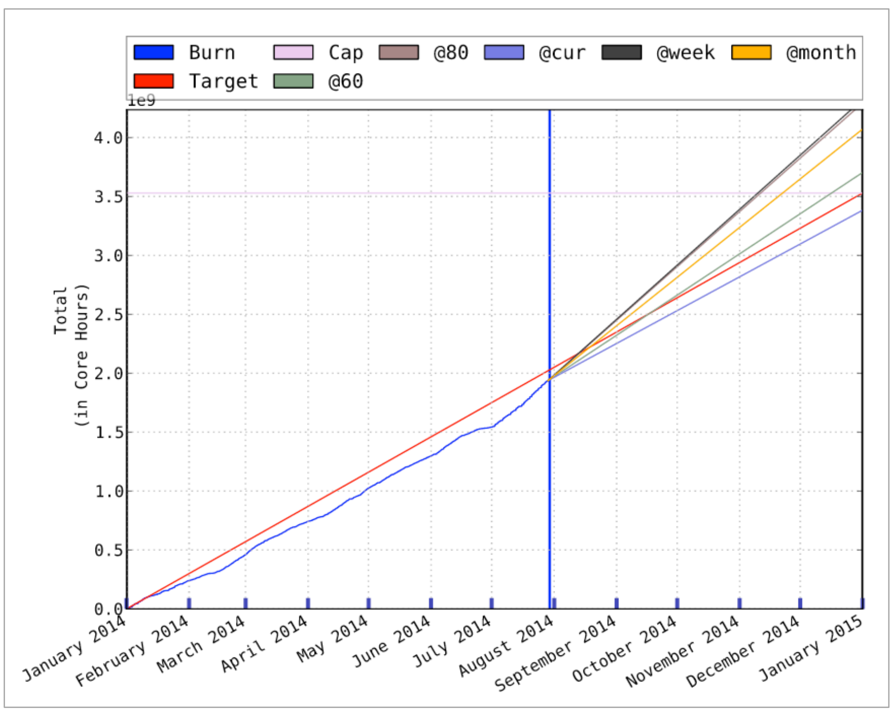

Figure 4.5 Example of a Burn Rate Fan-Out Graph

Impact/Status: The graph has enabled earlier detection of potential usage issues with INCITE projects. Because (as noted above) the graph for each INCITE project is examined in a biweekly scheduling meeting, its use has resulted in improvements to the scheduling policy.

4.2.3 Job Failure Analysis Database

Challenge: Analysis of job failure was time-consuming given that many sources of data had to be examined by hand.

Approach: Each week, ALCF analyzes all jobs that may have exited with an error. To help in the analysis, ALCF has developed a system to allow interactive querying of job data and annotation of job status. It re-uses much code from Intrepid that drove the original Job Failure Analysis (JFA). This system helps ALCF staff determine the fate of every job with a non-zero exit code. Instead of many people talking to the one person who has specific job information, the system allows staff to save the data in a database in real time. The system also removed much of the manual log parsing that individuals were doing to find out what really happened to a particular job.

This screen shot (Figure 4.6) of JFA's graphical user interface demonstrates some of the system features:

- A real-time interface for collaboration;
- Live editing of annotations;
- Automatic and assisted lookups for issues like RAS, interrupts, GPFS problems, etc.;
- Quick view of some job information like user, project, and event times; and
- Retrieval of output and errors from the various log files.

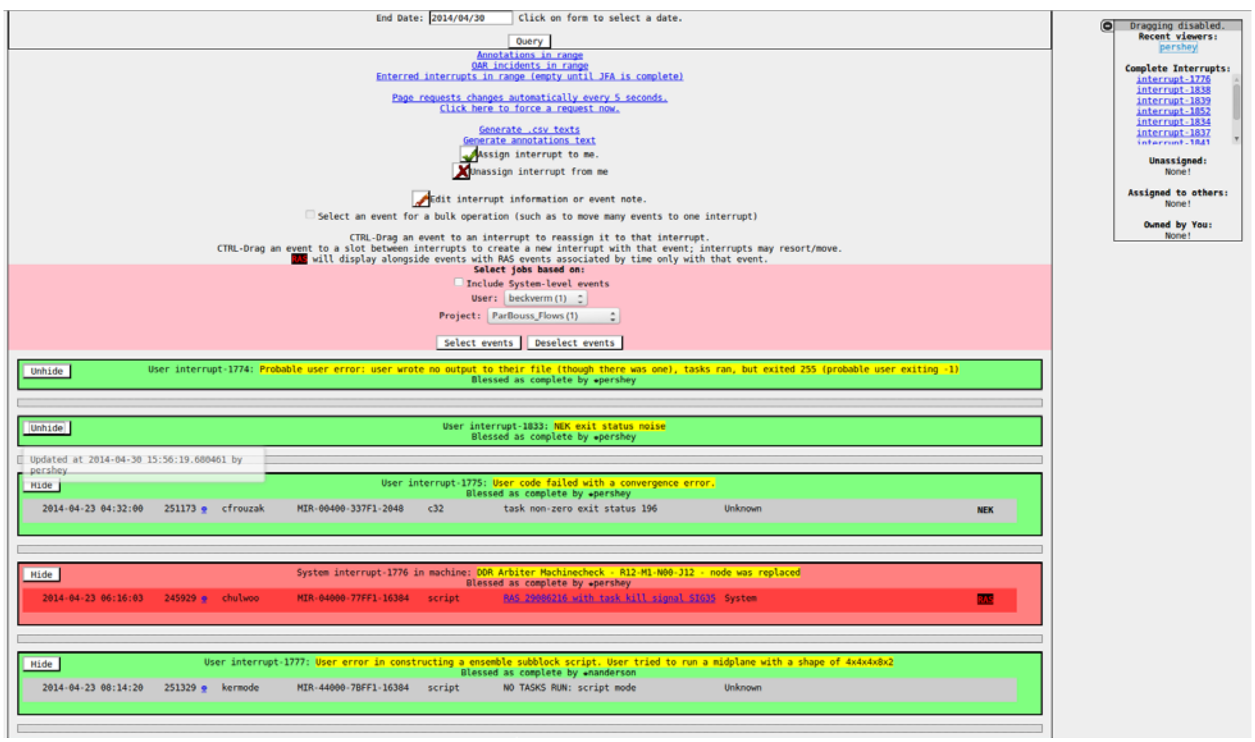

Figure 4.6 The JFA Interface

Impact/Status: The JFA database is now an integral part of the weekly operations of Mira and is under active development to improve its functionality. It saves 10 people each around two hours per week.

4.2.4 Automating and Monitoring Background Business Intelligence Job Execution

Challenge: Human monitoring of thousands of background jobs is time-consuming and error-prone.

Approach: The business intelligence team executes more than 6,300 background jobs per month, with the execution time of each job varying from a few seconds to several hours. Some jobs execute a single time during the night, whereas others execute several times per day. While most jobs perform processing to update and maintain the data in the data warehouse, some jobs perform data verification, system testing, and validation.

Human monitoring and verification of all of these jobs required several hours each day. ALCF implemented an automated system using Jenkins, an open-source continuous integration server. This automation replaces human effort by monitoring each job and reporting failures or timeouts immediately through e-mail or Jabber and visually through the Jenkins Web interface.

The business intelligence team can then quickly analyze errors by reading the logs stored in Jenkins. The Jenkins Web interface also allows a particular job to be analyzed, changed, and restarted. Jenkins processing allows job chaining; that is, one job can start other jobs upon completion, allowing restart of a complex set of tasks.

Impact/Status: The key improvement is that information can be vetted quickly. Bugs and errors that would not have been detected through manual monitoring are quickly revealed with the Jenkins automation.

4.3 Staff and Financial Management Improvements

ALCF has implemented management, organizational, and financial improvements that have resulted in more efficient operations.

4.3.1 Implementation of Three-Level Management Hierarchy

Challenge: The growth of ALCF staffing resulted in some managers having many more than the recommended five people per supervisor.

Approach: In 2015, ALCF fully implemented a plan to move from a two-level to a three-level management hierarchy. Previously, most managers reported directly to the ALCF Director, and most staff reported to a manager. Now, however, most staff report to team leads, and team leads report to director-level management, who in turn report to the ALCF Director.

Impact/Status: The three-level hierarchy has opened up new management opportunities for ALCF staff. It has also provided for clear backup of managers and has reduced the people-management workload, allowing managers to focus more on technical issues and long-term strategy.

4.3.2 Rotation of OAR Author Responsibilities

Challenge: Each year, the same ALCF staff members were charged with assembling and writing the same OAR sections.

Approach: ALCF has implemented a policy of rotating responsibility for the sections of the OAR. Each section has a primary author, a backup author, and a resource person. In the subsequent year, the backup author becomes the primary author, and a new backup author is identified. The primary author becomes the resource person. The team that assembles the OAR operates the same way.

Impact/Status: The rotation has several positive benefits, including the following:

- The primary author has motivated and engaged helpers, which improves the gathering of data and presentation of results.
- The backup author gains valuable experience before being tasked to provide primary authorship.
- Understanding of OAR reporting and the clear documentation that is required is spread more widely through ALCF staff, which improves data collection and writing.
- Staff members do not become fatigued with the same task year after year.

4.3.3 Improving HPC Financing

Challenge: Several national laboratories obtain high-performance computing equipment under lease arrangements and there is a need to better estimate future interest rates and understand the impact of upcoming changes in lease accounting rules.

Approach: In conjunction with participants from DOE's Advanced Scientific Computing Research, National Energy Research Scientific Computing Center, and Oak Ridge Leadership Computing Facility, ALCF remains an ongoing coordinator and participant in the Financing Improvement Team (FIT) best practices group. All of the organizations share information in order to improve leasing and negotiating practices, with the end goal of obtaining the best rates possible for all facilities. General areas of work discussed by the group include lease options, requests for proposals, evaluation templates, processes, and overall best practices. The impact has been significant, with more than $2M savings in interest and fees for the ALCF-1 lease and $3M savings in interest and fees for the ALCF-2 lease.

In 2014, Argonne broke new ground with the following two white papers:

- "Estimating Future Interest Rates," and
- "FASB–IASB Lease Rule Discussion Summary."

The first white paper formally develops a method for estimating future interest rates using current market data. Summarizing the method:

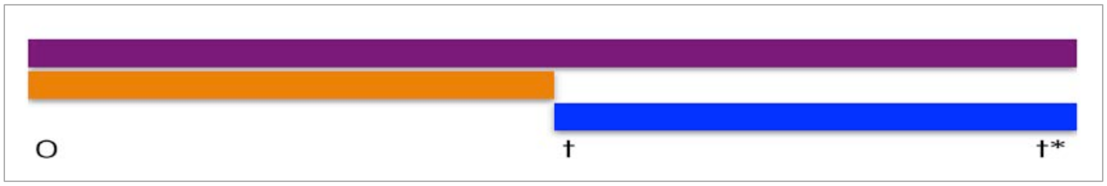

Figure 4.7 Interest-Rate Time Frame Analysis

As shown in Figure 4.7, the interest rate from time 0 through time t* comprises two pieces: 0 through t, and t through t*.

Because the bond market provides data for the annual and total returns represented by the purple (0 through t*) and the orange (0 through t) time periods, we use that information to solve for the blue (t through t*). The second white paper provides a preliminary assessment of potential changes in FASB and GASB lease rules.

Impact/Status: The white papers were shared among the FIT participants, and the analysis will impact new leases at several sites. The issues raised will be monitored and discussed more fully by the FIT best practices group in 2015 as more information from laboratories and financial markets becomes available.

Conclusion

ALCF has identified innovations and best practices that have improved the software infrastructure, operations, and management of ALCF in CY 2014.

- ALCF made improvements and enhancements to software that is critical to several user communities and has shared code back to the communities.
- ALCF improved staff understanding of machine behavior through automated tools and unique graphical representations.
- ALCF strengthened staff and financial management through an improved organizational structure, a more collaborative OAR process, and focused financial analysis.

This page intentionally left blank.

Section 5. Risk Management

Is the Facility effectively managing risk?

ALCF Response

ALCF has clearly demonstrated successful risk management in the past year for both project and operation risks. The risk management strategy is documented in the ALCF Risk Management Plan (RMP), which is reviewed and updated regularly to incorporate new ideas and best practices from other facilities. Risk management is a part of ALCF culture, and the RMP processes have been incorporated into both normal operations and all projects, such as the ALCF-3 project launched in CY 2013. Risks (proposed, open, and retired) are tracked, along with their triggers and mitigations (proposed, in progress, and completed), in a risk register managed by the Risk Manager. All risk ratings in this report are post-mitigation ratings. ALCF currently has 41 open risks, with one high operational risk – funding uncertainty, which is managed by careful planning with the DOE program office and the continuation of austerity measures as necessary. The major risks tracked for the past year are listed below, with the risks that occurred and the mitigations for those risks described in more detail, along with new and retired risks, as well as the major risks that will be tracked in CY 2015.

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

The Facility should highlight various risks to include:

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

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

5.1 ALCF Risk Management

ALCF uses the documented risk management processes, first implemented in June 2006 and outlined in its RMP, for both operations and project risk management. ALCF reviews and updates the plan annually. The plan is also updated as needed during the year to reflect changes and to incorporate new risk management techniques as they are adopted by the facility. The RMP is consulted at all monthly and individual risk meetings. Details of the RMP, including the attributes of each risk managed by ALCF, have been described in past reports and will not be discussed further here. Risks are tracked in a risk register using the commercial management tool Oracle Primavera Risk Analysis (OPRA, formerly known as PertMaster), which integrates with the Primavera project management tool used to manage all large ALCF projects.

Continuation of the ALCF-3 Project

The ALCF-3 project – procuring and deploying the next ALCF supercomputer – continued in 2014. A project risk register is maintained and a set of detailed risks is tracked. Risk mitigation costs on the project side are developed using a bottom-up cost analysis, then are input to OPRA to set the contingency pool utilizing the OPRA integration with Primavera. These risks are not included in the risk numbers covered in this document and are not discussed further.

Risk Review Board

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

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

Risk Management in Day-to-Day Operations

ALCF currently has 41 open risks in its operations risk register. These risks include general facility risks (such as funding uncertainties, staffing issues, and safety concerns) and specific risks (such as system component failures, availability of resources, and cost of electricity). On the operations side, subject matter experts estimate risk mitigation costs and use them to inform management reserves.

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

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

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

5.2 Major Risks Tracked for the Review Year

Since Q4 of FY 2010, ALCF has experienced several eventful years as a result of Mira's transition to operations in FY 2013 and the planned growth of both ALCF staff and budget in order to bring the facility to full strength. As such, ALCF was monitoring, and continues to monitor, a large number of major risks for the facility. These risks are described in Table 5.1. All risk ratings shown are post-mitigation ratings. There were twelve major operations risks tracked for CY 2014, one with a risk rating of *High* and eleven with a risk rating of *Moderate.* Of these, four were encountered and managed. No major risks were retired during CY 2014. The risks are color-coded as follows to assist with reading the table:

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

Table 5.1 Major Risks Tracked for CY 2014

ID	Title	Encountered	Rating	Notes
1059	Funding/budget uncertainties	Yes	High	Uncertainty plus the need for facility growth in order to operate Mira combined to make this a significant challenge in the past year. This risk remains a major concern as the facility moves forward with ALCF-3.
25	Staffing challenges	Yes	Mod	With the ongoing budget uncertainties and difficulties with retaining staff in the face of a recovering economy, this will continue to be a concern.
31	Facility power interruptions	Yes	Mod	One power outage occurred during CY 2014 that caused a Mira outage; this event is discussed as item 4 in Sec. 2, "Business Results" subsection 2.1.1.
1049	Staff retention	Yes	Mod	Between budget concerns at Argonne and the growth in high paying industry jobs for system administrators and programmers with HPC expertise, ALCF has lost several staff members during CY 2014. This remains a concern.

(continued on page 5-4)

Table 5.1 Major Risks Tracked for CY 2014 (Cont.)

ID	Title	Encountered	Rating	Notes
1018	INCITE and ALCC users are not provided adequate support by ALCF	No	Mod	ALCF staff frequently solicit feedback from project members on service. ALCF staff manage the support expectations of the project members.
1050	ALCF has insufficient disk space to support science needs	No	Mod	Storage upgrade. Established and enforced quotas. Require users to more aggressively move data to tape. Pre-negotiate space. Install additional disk. Review allocation proposals for storage needs and evaluate existing storage for space.
1054	Catastrophic failure of home file system	No	Mod	Tar backup files prior to archiving them. Mirror the data in the home file system. Replicate metadata. Improve performance of the tape system with purchase of additional hardware. Scripts in place to tar the backup files prior to archiving them.
1056	System stability issues due to upgrades	No	Mod	Perform upgrades on non-critical systems first when feasible. Have a rollback plan in place. Monitor performance closely following upgrade. Work with the vendor to understand the upgrade(s) and the quality control processes.
1076	If the ISSF is decommissioned, we will not have an appropriate facility to host our disaster recovery resources	No	Mod	Explore alternative locations on site. 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 upper lab management; contact other laboratories. Develop plan to move equipment.
1085	Diagnostic suite and utilities fail to detect hardware problems	No	Mod	Track and monitor job and hardware failures; correlate. Work with vendor to resolve issues and improve diagnostic suite.
1091	Injury to workers/overall safety of the division	No	Mod	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.
1099	INCITE and ALCC do not use all allocated core-hours	No	Mod	Generally occurs with only a few projects. Addressed through catalyst monitoring, adjusting of scheduler priorities, and catalyst/project communications.

5.3 Risks Encountered in the Review Year and Their Mitigations

The top risks encountered in the last 12 months are discussed below, along with the risk owner, its probability and impacts, a description of the actual problem that occurred, and the management of the risk.

5.3.1 Funding/Budget Uncertainties

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

Description

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

Evaluation

During the past year, the Funding/Budget Uncertainties risk was ALCF's highest risk, and it was also one of the risks encountered. Although ALCF was supposed to be in a growth phase, the facility was required to operate with moderate austerity measures during the early part of the year, which carried a significant impact.

Management

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

ALCF continues to closely monitor budget information for FY 2015 and beyond in case of a reduction in funds from the plan of record. Moderate austerity measures remain in place, with spending being prioritized, and these measures may be augmented depending on the budget through FY 2015.

5.3.2 Staffing Recruitment and Retention Challenges

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

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

Description

This is a period of necessary growth for ALCF as it continues to staff up to operate Mira and prepare for ALCF-3. An aggressive staff ramp up, originally planned for FY 2010 through FY 2012, was extended because of budget reductions. An ALCF risk evaluation identified two key risks associated with this ramp up, and both occurred in CY 2014 as a result of industry competition for retention of existing employees and potential new hires. The risks have been combined for this discussion, as they are related:

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

Evaluation

As the economy continues to recover, more industry jobs open up for ALCF staff. As a result, in the past year five ALCF staff left, three of them for higher-paying jobs in industry. Seven new full-time and one part-time staff were added during CY 2014, for a net gain of +2.5 ALCF staff for the year. Thus, ALCF has made progress on adding priority new hires.

Management

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

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

By carefully and judiciously managing both risks, ALCF has successfully operated the facility and moved ahead with the ALCF-3 project. The facility staffing level is still not in alignment with the staffing plan that was developed several years ago, and there continues to be high demand for the skills of ALCF staff. Thus both staff recruitment and staff retention will remain a concern for ALCF.

5.4 Retired Risks

Three risks were retired during the past year, as indicated in Table 5.2.

Table 5.2 Risks Retired in CY 2014

ID	Title	Rating	Management Strategies	Notes
26	Unable to recruit qualified staff within required timeframes	Low	Prioritized staffing needs. Replanned work. Retasked staff.	This risk was determined to be covered by risk 25 and so was retired.
1077	Cost of decommissioning and removing Intrepid unplanned	Very low	Approach other sites to see if they would like to take part of Intrepid. Explore nonstandard options. Use management reserves.	With the decommissioning of Intrepid completed, this risk has been retired.
1101	Facility that provides cooling is insufficient	Low	Ongoing communications with Facilities Management and Services Division (FMS, provide electrical power to the building) and Theory and Computer Science (TCS) building management. Regular meetings with FMS and TCS building management.	This risk was determined to be covered by risk 30 and so was retired.

5.5 New and Recharacterized Risks since the Last Review

Staff operating within the ALCF risk culture regularly identify new risks and recharacterize existing risks. Because only one major production machine (Mira) operated in CY 2014 and no new machine was deployed during this year, no new risks were added in CY 2014. Table 5.3 lists the five risks recharacterized during the year.

Table 5.3 Recharacterized Risks from CY 2014

ID	Title	Rating	Management Strategies	Notes
25	Staff recruitment challenges	Mod	Evaluate possible additional recruiting avenues. Prioritize staffing needs. Adjust work planning. Retrain staff to meet ALCF needs. Retask staff as needed.	Risk recharacterized, combined with risk 26.
1018	INCITE and ALCC users are not provided adequate support by ALCF	Mod	Frequently solicit feedback from project members on service. Hire additional staff as required. ALCF staff manage the support expectations of the project members.	Risk recharacterized to cover both INCITE and ALCC users.
1038	Incorrect answers come out of the computers	Low	Run validation on the production codes. Use the CPS test. Build a regression test with benchmarks and applications. Engage pertinent vendor(s).	Risk recharacterized to generalize to all allocations and any pertinent machine vendors.
1044	Unable to meet high end-of-year demand for INCITE and ALCC cycles	Low	Continuously monitor usage and inform the PIs of their status and impact of project status via bimonthly reports and Catalyst weekly reports. Encourage projects to use time earlier. Enable scheduler changes to improve throughput.	Risk recharacterized to cover both INCITE and ALCC users.
1099	INCITE and ALCC do not use all allocated core-hours	Mod	Clear communications between catalysts and projects on importance of time. Ongoing monitoring by catalysts. Adjusting of scheduler priorities to encourage project workflows. Catalysts work with project PI's on any problems with running codes, monitor project usage of allocations.	Risk recharacterized to cover both INCITE and ALCC users.

5.6 Projected Major Operating Risks for the Next Year

Table 5.4 lists the current top operating risks projected for CY 2015 along with the current risk rating and management strategies for the risk.

Table 5.4 Projected Operating Risks for CY 2015

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

Conclusion

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

This page intentionally left blank.

Section 6. Safety

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

ALCF Response

ALCF has an exemplary safety record. Since the division's inception in 2006, ALCF has never experienced a lost time incident. The facility had one minor first-aid case in CY 2013 and a near miss in CY 2014, which is described in the section on the decommissioning of Intrepid. As noted there, root cause analysis was conducted and improvements, corrections, and additional training were put in place as a result. ALCF is very proud of its deeply ingrained safety culture and continually seeks additional ways to improve it. ALCF employs appropriate work planning and control principles. A formal "skill of the worker" document is used for routine tasks. Formal and specific procedures are in place for more complex tasks, such as changing out the Blue Gene/Q power supplies (which can pose a thermal hazard) and node boards (which pose a very mild chemical hazard owing to use of water treatment chemicals, weight, and potential damage to hardware), as well as medium-voltage electrical maintenance. The facility performs hazard analysis and creates work planning and control documents for emergency work or in case of an unexpected change to previously planned work.

This page intentionally left blank.

Section 7. Cyber Security

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

ALCF Response

Yes. The Argonne Authority to Operate (ATO) includes ALCF as a major application, and it was granted on February 8, 2012. It is valid as long as Argonne National Laboratory maintains robust, continuous monitoring of the Cyber Security Program as detailed in the letter. A copy of the ATO letter follows.

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

OFFICE OF THE DIRECTOR
ARGONNE NATIONAL LABORATORY
2012 FEB 21 PM12:52

To: MAS
Copy: PKK
Action: Yes No
Due Date: —
Response Copy to: doemail@anl.gov
OTD File #: —

February 8, 2012

Dr. Eric D. Isaacs
Director, Argonne National Laboratory
President, UChicago Argonne, LLC
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Isaacs:

0212-022

SUBJECT: APPROVAL OF AUTHORITY TO OPERATE FOR THE ARGONNE NATIONAL LABORAORY (ANL) INFORMATION TECHNOLOGY INFRASTRUCTURE

References:
1. Letter, Skwarek to Livengood dated January 23, 2012, Subject: Authority to Operate for the Argonne National Laboratory Information Technology Infrastructure
2. Letter, Livengood to Isaacs dated May 28, 2010. Subject: Extension of Argonne National Laboratory (ANL) Unclassified Cyber Authority to Operate (ATO)

The submitted security authorization package has been reviewed for the General Computing Enclave and its major applications:

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

This enclave and the listed major applications include all IT investments at ANL. Based on the submitted documentation and the results of the FY 2011 Office of the Inspector General (OIG) evaluation of the Unclassified Cyber Security Program and the November 2011 Review conducted by the Office of Science, Chicago Office, I am granting these systems an Authority to Operate (ATO) at the FIPS-199 level of Moderate. This ATO will be maintained via robust continuous monitoring for the Cyber Security Program under the auspices of the Contractor Assurance Systems (CAS). The initial expectations for the continuous monitoring program will be provided under separate cover.

A component of the Office of Science

This ATO supersedes the Extension of Authority to Operate issued on May 28, 2010. The Laboratory should retain a copy of this letter with the security authorization package.

If I can be of any assistance, please contact me or have your staff contact Francis Healy at 630- 252-2827 or e-mail Frank.Healy@ch.doe.gov.

Sincerely,

Dr. Joanna M. Livengood
Manager

cc: M. Skwarek, ANL
M. Skwarek, ANL
M. Kwiatkowski, ANL
W. Dykas, SC-31.3 GTN

This page intentionally left blank.

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

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

ALCF Response

ALCF and DOE have agreed to the 2015 metrics and targets as proposed in the March 2014 OAR report. These metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2016, the proposed metrics and targets for the current production resources remain the same as for 2015.

The facility should provide a summary table of the metrics and targets agreed upon for the review of Calendar Year 2015 and include the target and actual values of similar metrics used for 2014 for comparison. The facility should also provide metrics and targets under consideration for CY 2016. Those will be finalized later in the year.

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

8.1 Overview

The 2015 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2016, the proposed metrics and targets for the current production resource, Mira, will remain the same as for 2015. The 2015 metrics are covered in Section 8.2, and the 2016 metrics are covered in Section 8.3.

8.2 ALCF 2015 OA Performance Metrics

The OA performance metrics, 2014 targets and actuals, and agreed-upon 2015 targets are presented in Table 8.1.

Table 8.1 Performance Metrics – 2014 Targets, 2014 Actuals, and Agreed-Upon 2015 Targets

Area	Metric	2014 Target	2014 Actual	2015 Target
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **Response Rate**	25%	30.0%	25%
	% User Problems Addressed within Three Working Days	80%	96.0%	80%
Business Results	Mira Overall Availability	90%	95.7%	90%
	Mira Scheduled Availability	95%	98.7%	95%
	% of INCITE core hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	30%	64.5%	40%
	% of INCITE core hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10%	33.1%	10%

8.3 ALCF Proposed 2016 OA Performance Metrics

The OA performance metrics, agreed-upon 2015 targets, and 2016 proposed targets are shown in Table 8.2.

Table 8.2 Performance Metrics – Agreed-Upon 2015 Targets and Proposed 2016 Targets

Area	Metric	2015 Target	Proposed 2016 Target
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	3.5/5.0
	User Survey – **Response Rate**	25%	25%
	% User Problems Addressed within Three Working Days	80%	80%
Business Results	Mira Overall Availability	90%	90%
	Mira Scheduled Availability	95%	95%
	% of INCITE core hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40%	40%
	% of INCITE core hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10%	10%

8.4 ALCF Reportable Only Metrics (No Targets)

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

Table 8.3 ALCF Reportable Only Metrics

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

Conclusion

The agreed-upon 2015 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2016, because Mira will have been in a stable state for multiple years, the proposed metrics and targets for the current production resources remain the same as for 2015. Achieving the agreed-upon 2015 targets and the proposed 2016 targets will indicate that the facility is performing up to stakeholder expectations. ALCF anticipates being able to meet all metric targets.

This page intentionally left blank.

Appendix A – Calculations

A.1 Scheduled Availability Calculation Details

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

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

Formula:

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

where:

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

A.2 Overall Availability Calculation Details

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

Formula:

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

A.3 ALCF Calculations

Here is a simple example.

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

In its calculations, ALCF tracks availability at the core-second level. The Blue Gene architecture allows an individual node card, containing 32 nodes, to be taken off line to replace one node while the rest of the machine continues to run. However, in calculating availability, ALCF takes into account the ALCF scheduling policy for its large production systems, which does not allow jobs smaller than those using 512 nodes (8,192 cores) to run, which means that 512 nodes is the smallest number of nodes that will be allocated. Therefore, if a single node were to fail for exactly one hour, it would be recorded as 8,192 cores × 3600 seconds = 7,372,800 core-seconds of down time. ALCF has multiple production "scratch" file systems for Mira. Therefore, if all of them are down, the entire machine is considered to be down. If any one of them is available, there are users who can run and therefore, the machine is considered to be available.

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

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

A.4 MTTI Calculation Details

MTTI (mean time to interrupt) is defined as time, on average, to any outage on the system, whether unscheduled or scheduled. It is also known as MTBI (mean time between interrupt).

Formula:

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

A.5 MTTF Calculation Details

MTTF (mean time to failure) is defined as the time, on average, to an unscheduled outage on the system.

Formula:

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

A.6 ALCF MTTI/MTTF Calculations

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

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

$$Utilization = \frac{Core\ hours\ consumed}{Total\ core\ hours\ that\ were\ available} * 100$$

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

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

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

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

Overall Capability represents the total of all the jobs in the first two categories; High Capability represents all the jobs in the third category. On Intrepid, capability was defined as using greater than 20 percent of the machine. However, 20 percent of Mira would be 9.6 racks, which is not a viable configuration. Hence, the Mira metric was defined in two parts. Overall Capability requires that a minimum of 30 percent of INCITE core-hours be run on eight racks or more (16.7 percent), and High Capability requires that a minimum of 10 percent of INCITE core-hours be run on 16 racks or more (33.3 percent). These targets increase to 40 percent and 10 percent, respectively, for CY 2015.

Appendix B – ALCF Director's Discretionary Projects

Mira DD Allocations by Project Name, January 1, 2014–December 31, 2014

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
0vbbqrpa	University of Tsukuba, Japan	Jun Terasaki	Calculation of Nuclear Matrix Element of Neutrinoless Double-Beta Decay	Physics	8,400,000
3DTransition	University of Michigan (Ann Arbor) and Johns Hopkins University	Karthik Duraisamy	Direct Numerical Simulations and Machine Learning for Swept Wing Transition	Engineering	1,000,000
Acoustic_Modelling	Queen Mary University of London	Sergey Karabasov	Acoustic Modelling of Jet-Wing Interaction	Engineering	1,572,864
ACRT-EJECT	University of Tunis El Manar, Department of Physics	Ghassen Rezgui	Thermal Conduction Effect on the Accretion Ejection Disk around Young Stars	Physics	2,000,000
ALCFCERFACS_CFD	Argonne National Laboratory/Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS)	Gabriel Staffelbach	Scaling and Baseline Study of Coupled Simulations Using A Very Big Program (AVBP) on O(B) Finite Volumes on the Blue Gene/Q	Chemistry	5,000,000
ALCF_Getting_Started	Argonne National Laboratory	Chel Heinzel	ALCF Getting Started	Training	10,100,000
ALCF_SANDIA_RAPTOR	Sandia National Laboratories, New Mexico	Joseph C. Oefelein	Scaling the RAPTOR Combustion LES Code on Mira	Engineering	13,000,000
ALK	Western University of Health Sciences	Yun Luo	Rational Design of Selective ALK2 Inhibitors	Biological Sciences	8,000,000
Allinea	Allinea Software	Ray Loy	Improved Debugging Memory Usage for BG/P	Internal	2,000,000
alpha-nek	The University of Chicago (UChicago)/Fermi National Accelerator Laboratory (Fermilab)	Maxwell Hutchinson	DNS of Multi-Mode Rayleigh-Taylor Instability	Engineering	4,000,000
ARL-KSDFT	U.S. Army Research Laboratory	Alexander Breuer	Scalable All-Electron Structure Calculations	Chemistry	100,000
ATLASQ	Argonne National Laboratory	Thomas J. LeCompte	Grid-Enabling High Performance Computing for ATLAS	Physics	3,000,000
ATPESC2014	Argonne National Laboratory	Paul Messina	Argonne Training Program on Extreme Scale Computing	Training	7,000,000
Autoignition_esp	Eidgenössische Technische Hochschule Zürich (ETH Zurich)	Christos Frouzakis	Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow	Chemistry	10,000,000
BGQ_Energy_Profiling	Argonne National Laboratory	Venkatram Vishwanath	Understanding and Characterizing the Power Consumption and Energy Efficiency of Applications on BG/Q	Computer Science	1,000,000
BIG_MAC	Northwestern University	Monica Olvera de la Cruz	Effective Interactions in Coulombic Systems with Highly Disparate Particle Sizes	Physics	3,145,728
bio_stress	Purdue University	Lyudmila Slipchenko	Mechanistic Understanding of Oxidative Stress in Biological Membranes from Quantum Chemical Calculations	Chemistry	9,000,000
Boeing_NTS	Boeing Commercial Airplanes	Philippe Spalart	Simulation of a Shock-Boundary-Layer Interaction	Engineering	2,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
Boot_Camp_2014	Argonne National Laboratory	Richard Coffey	Boot Camp	Training	50,000
bulk_properties_dd14	Argonne National Laboratory	Graham Fletcher	High Accuracy Predictions of the Bulk Properties of Water	Chemistry	3,000,000
CABARET_Noise	Queen Mary University of London	Sergey Karabasov	High-Resolution Acoustics-Coupled Computational Fluid Dynamics for Community Noise Modelling	Engineering	2,752,512
CALFIRE	Desert Research Institute	Dr. Timothy Brown	The Development of Fire Map 1: Fire Threat to Utilities	Earth Science	6,000,000
Camellia	Argonne National Laboratory	Nathan Roberts	Camellia for Discontinuous Petrov-Galerkin Simulations of Incompressible Flow	Physics	5,000,000
CARDIO-PAR	University of Milan, Italy	Luca F. Pavarino	Scalable Domain Decomposition Methods for Computational Cardiology and Isogeometric Analysis	Mathematics	800,000
Catalyst	Argonne National Laboratory	Katherine Riley	Catalyst	Internal	20,000,000
CESM_Security	Fermi National Accelerator Laboratory (Fermilab)	Alfred Tang	Simulating Crop Yield and Water supplies Using CESM	Earth Science	20,000
charmm_zmod	Harvard University/Harvard Medical School	Robert J. Petrella	Highly Parallel Macromolecular Conformational Searches and Energy Evaluations with the CHARMM Program	Biological Sciences	17,000,000
ChemiCAD	The University of Chicago (UChicago)	Rick Stevens	Parallel Multi-Objective Optimization Algorithm for de Novo Ligand Design	Biological Sciences	1,000,000
cig_dynamo	University of California-Davis	Nicholas Featherstone	Development and Testing of Exascale Dynamo Simulation Code	Earth Science	1,500,000
CIT	Cornell University	Olivier Desjardins	Computational Study of Cluster-Induced Turbulence	Engineering	5,000,000
ClimateUncertainty1	Argonne National Laboratory	Ian Foster	Investigation of Initial Condition Uncertainty in Climate Models	Earth Science	10,000,000
cluster_mergers	National Aeronautics and Space Administration (NASA)/Goddard Space Flight Center	John ZuHone	MHD Simulations of Galaxy Cluster Mergers	Physics	1,000,000
CMT	Argonne National Laboratory	Scott Parker	Compressible Multiphase Turbulence	Engineering	10,000,000
CobaltDevel	Argonne National Laboratory	Narayan Desai	Cobalt Development	Internal	10,000,000
CompBIO	Argonne National Laboratory/The University of Chicago (UChicago)	Rick Stevens	Multiscale Simulations in Biology: Evolution and Ecology of Microbes	Biological Sciences	3,000,000
ComplexCrystalSearch	Argonne National Laboratory	Carolyn Phillips	Design Space Searching for Complex Crystals	Materials Science	28,151
ComplexTurbulence	University of Minnesota	Krishnan Mahesh	High-Fidelity Simulations of Complex Turbulent Flows	Engineering	2,000,000
Concerted_Flows	Argonne National Laboratory	Venkatram Vishwanath	Topology-Aware Data Movement and Flow Coordination for Supercomputing	Computer Science	1,000,000
CONVERGE-BGQ-LDRD	Argonne National Laboratory	Marta Garcia Martinez	Performance Improvement of CFD Code CONVERGE™ on BG/Q Systems	Engineering	5,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
Coral_rfp_prep	Argonne National Laboratory	Kalyan Kumaran, Vitali Morozov	CORAL RFP Preparatory Work	Internal	5,000,000
Coronary_Bifurcation	University of Canterbury	Tim David	Massively Parallel Models of Cellular Activity in Coronary Arteries	Biological Sciences	2,000,000
critical_perf	Argonne National Laboratory	Kalyan Kumaran	Critical Debugging Project	Internal	50,000,000
CROC	Fermi National Accelerator Laboratory (Fermilab)/The University of Chicago (UChicago)	Nickolay Y. Gnedin	Cosmic Reionization on Computers	Physics	3,500,000
Cryst_Eng_2013	The University of Texas at Austin/Fritz-Haber-Institut der Max-Planck-Gesellschaft (FHI)/Argonne National Laboratory	Noa Marom, Volker Blum, Jeff Hammond	Toward Crystal Engineering from First Principles	Materials Science	2,000,000
CSM_PE_LBM	Colorado School of Mines	Feng Xiao	Lattice Boltzmann Simulations in Complex Porous Medium Geometries	Earth Science	5,000,000
ctrsp2014	Stanford University	Curtis Hamman	Studying Turbulence Using Numerical Simulation	Engineering	5,000,000
DetailedChem	Purdue University	Prof. William Anderson	LES Simulation of Combustion Instability with Detailed Chemical Kinetics	Engineering	984,000
DiffLight_APS	Argonne National Laboratory	Michael Borland	Diffraction-Limited Light Source for APS Upgrade	Physics	10,000,000
Direct_Noise	General Electric Company (GE)/Global Research	Umesh Paliath	LES Simulation of Fan and Exhaust Aerodynamics & Acoustics for Propulsion Systems	Engineering	5,500,000
DiscoveryEngines	Argonne National Laboratory	Justin M. Wozniak	Integrating Simulation and Observation: Discovery Engines for Big Data	Materials Science	5,000,000
DNS_Clustering_HTSF	Cornell University	Lance R. Collins	DNS of Particle Clustering in Homogeneous Turbulent Shear Flow at Higher Reynolds Numbers	Engineering	2,557,600
duanl	Missouri University of Science and Technology	Lian Duan	Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers	Engineering	2,000,000
duanlianx	Missouri University of Science and Technology	Lian Duan	Direct Numerical Simulations of Laminar-Turbulent Transition in Swept-Wing Boundary Layers	Engineering	4,000,000
E2_MD_trajectories	Structural Biophysics Laboratory, National Cancer Institute	R. Andrew Byrd	Correlating Experimentally Measured Molecular Dynamics with Computational Trajectories: Understanding Dynamic Allostery in Ubiqu	Biological Sciences	5,000,000
Ebola	Argonne National Laboratory	Charles Macal	Agent-based Behavioral Modeling of Ebola Spread in Chicago and Other Large Urban Areas	Computer Science	500,000
ECN-DNS	Arizona State University	Marcus Herrmann	Primary Atomization DNS of ECN's Spray A	Engineering	1,000,000
elastic	Buffalo University (SUNY)	Jaroslaw Zola	Similarity Graphs from Large-Scale Biological Sequence Collections	Biological Sciences	2,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
EnergySystems	Argonne National Laboratory	Cosmin Petra	Optimization of Stochastic Power Grid Systems	Mathematics	2,500,000
esumcfd	Rensselaer Polytechnic Institute (RPI)	Cameron Smith	Extreme Scale Unstructured Mesh CFD Workflow	Engineering	5,000,000
ExaHDF5	Argonne National Laboratory	Venkatram Vishwanath	ExaHDF5: Advancing HDF5 HPC I/O to Enable Scientific Discovery	Computer Science	3,000,000
ExM	Argonne National Laboratory	Justin M. Wozniak	Extreme Many-Task Computing with Swift	Computer Science	7,500,000
FE2SIM	University of Cologne AND TU Bergakademie Freiberg	Axel Klawonn, Oliver Rheinbach	Parallel Multiscale Simulations of Advanced Steel Materials	Mathematics	1,000,000
Femtomagnetism	Indiana State University	Guoping Zhang	First-Principles Calculation of Laser-Induced Ultrafast Magnetism	Physics	3,000,000
FFTsim	King Abdullah University of Science and Technology (KAUST)	Benson Muite	Large and Accurate Numerical Solutions of Partial Differential Equations	Mathematics	3,000,000
FLASH_combustion	University of North Carolina at Charlotte	Praveen Ramaprabhu	Numerical Simulations of Turbulent Combustion Using the FLASH Code	Chemistry	1,200,000
FPMC	Mississippi State University	Neeraj Rai	First-Principle Monte Carlo Algorithm Development and Implementation in CP2K	Chemistry	5,000,000
fpnmd	University of Illinois at Urbana-Champaign	Andre Schleife	Electronic Response to Particle Radiation in Semiconductor Systems	Materials Science	10,000,000
gamra	Earth Observatory of Singapore	Sylvain Barbot	Géodynamique Avec Maille Rafinée Adaptivement	Earth Science	2,000,000
GFDL_dd	Geophysical Fluid Dynamics Laboratory (GFDL)/National Oceanic Atmospheric Administration (Princeton, NJ)	Christopher L. Kerr	Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model	Earth Science	15,000,000
graph500	Indiana University (IU)	Andrew Lumsdaine	Graph500 Benchmark Run on Intrepid	Computer Science	6,300,000
GRChombo	Argonne National Laboratory	Hal Finkel	Early-Universe Phase Transitions with Strong Gravity and Instabilities in Higher-Dimensional Black Holes	Physics	1,500,000
GTRI_IBM2M_Init	Argonne National Laboratory	Micheal A. Smith	GTRI and NEAMS-Related Production Tests and Runs	Nuclear Energy	10,000,000
haiboyu	University of Wollongong	Haibo Yu	Computing the Binding Affinities between PTP1B and Allosteric Inhibitors	Biological Sciences	4,000,000
HBNDMC	University of Illinois at Urbana-Champaign	N. R. Aluru	h-BN-Water Interaction Using Quantum Monte Carlo Calculation	Materials Science	1,000,000
HEP-ANL	Argonne National Laboratory	Sergei Chekanov	Simulation of High-Energy Particle Collisions Using Monte Carlo Models	Physics	8,000,000
Heterocrystals	University of California-Los Angeles	Christian Ratsch	High-Throughput Computational Design of Heterocrystals	Materials Science	1,700,000
HighReyTurb_DD2014	The University of Texas at Austin	Robert Moser	Direct Numerical Simulations of High Reynolds Number Turbulent Channel Flow	Engineering	1,000,000
HighReyTurb_PostProc	The University of Texas at Austin	Robert D. Moser	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	10,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
HPCTuning	Lawrence Berkeley National Laboratory	Khaled Ibrahim	HPC Challenge Tuning	Computer Science	10,000,000
HPqC	California Institute of Technology (Caltech)	Jeff Amelang	High Performance Quasicontinuum Proof of Concept	Engineering	1,000,000
HP_Li	Argonne National Laboratory	Ying Li	High-Performance Li-Air Battery	Materials Science	5,000,000
Hydra_Test	Rolls Royce plc	Luigi Capone	Hydra Test for INCITE Application	Engineering	1,000,000
Hydro_model	Western Michigan University	Mohamed Sultan	Use of GRACE, Remote Sensing, and Traditional Data Sets for Modeling Time-Dependent Water Partitioning on Continental Scales	Earth Science	500,000
HyPar-Scalability	Argonne National Laboratory	Debojyoti Ghosh	Scalability of Weighted, Non-Linear Compact Schemes	Engineering	5,500,000
H_MHD	Los Alamos National Laboratory (LANL)	Andrey Beresnyak	Hall-MHD Turbulence and Reconnection	Physics	500,000
IBM-performance	Argonne National Laboratory	Kalyan Kumaran	BG/P Performance Runs Carried out by IBM Employees	Internal	2,100,000
IBM-Software-Testing	Argonne National Laboratory	Susan Coghlan	IBM Testing of Early Release Software for the BG/Q	Internal	50,000,000
IBMGSSTest	Argonne National Laboratory	Kevin Harms	IBM GSS Scalability Testing	Computer Science	6,553,600
IBM_CORAL	IBM Corporation/Watson Research Center	Bob Walkup	IBM Coral Benchmarks	Internal	6,553,600
ice-nucleation-ffs	Princeton University	Pablo G. Debenedetti	Using Advanced Sampling Techniques in Studies of Ice Nucleation Under Atmospherically Relevant Conditions	Chemistry	1,000,000
icee	Lawrence Berkeley National Laboratory	John Wu	ICEE	Computer Science	1,500,000
Illustris-TNG	Massachusetts Institute of Technology (MIT)	Mark Vogelsberger	Large-Scale Cosmological Hydrodynamic Simulations of Structure Formation in the Universe	Physics	500,000
IME_BlockCoPolymers	Argonne National Laboratory	Venkatram Vishwanath	Scalable Data Analysis of Soft X-Ray Scattering for APS Beamline Experiments	Materials Science	100,000
Intel-Performance	Argonne National Laboratory	Kalyan Kumaran	Intel Employees Performance Testing on Mira	Computer Science	50,000
InterX	InterX Inc.	Boris Fain	A Comprehensive Survey of Molecular Interactions at Chemical Accuracy	Biological Sciences	3,000,000
Inverse_Design	TeraDiscoveries	Edwin R. Addison	Inverse Design of Molecules and Drugs	Biological Sciences	2,000,000
LAMMPSopt	IBM Corporation	Paul Coffman	LAMMPS Performance Optimization	Materials Science	4,000,000
LaSco	CERFACS	Gabriel Staffelbach	Large-Scale Combustion Preparatory Access	Chemistry	10,000,000
LASPT	George Mason University	Adam Cadien	Liquid and Amorphous Structural Phase Transitions	Materials Science	552,960
LLVM	Argonne National Laboratory	Hal Finkel	LLVM Compiler Tools for the BG/Q	Internal	10,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
LPS-AMPs	The University of Kansas	Wonpil Im	Interactions of Lipopolysaccharide Bilayers and Antimicrobial Proteins	Biological Sciences	3,000,000
LSU3shell	Louisiana State University (LSU)	Tomas Dytrych	Ab Initio Modeling of Medium-Mass Atomic Nuclei	Physics	600,000
Maintenance	Argonne National Laboratory	Cheetah Goletz	LCF Operations System Maintenance	Internal	20,000,000
MCMD_Nano	Argonne National Laboratory	Phay J. Ho	Radiation Damage in Nanoparticles by Intense X-Ray Interactions	Physics	4,000,000
MembraneTransporters	University of Illinois at Urbana-Champaign	Emad Tajkhorshid	Characterizing Large-Scale Structural Transitions in Membrane Transporters	Biological Sciences	1,000,000
Meso_CCS_DD13	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS)	Roberto Paoli	Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus	Earth Science	9,000,000
ML4TMO	Argonne National Laboratory/Columbia University	Alejandro Lopez-Bezanilla	Machine Learning for Transition Metal Oxides	Physics	5,000,000
MM-MEDE	California Institute of Technology (Caltech)	Mauricio Ponga	Multiscale Modeling of Materials under Extreme Dynamic Environments through Large-Scale Computer Simulations	Materials Science	5,000,000
MMF-IPHOC	Science Systems and Applications, Inc. (SSAI)/National Aeronautics and Space Administration (NASA)	Anning Cheng	Climate Sensitivity Experiments Using a Multi-Scale Modeling Framework with a Higher-Order Turbulence Closure in Its CRM	Earth Science	595,200
ModSS	RWTH Aachen University	Ahmed E. Ismail	Molecular Dynamics Simulations of Superspreaders	Materials Science	5,000,000
moose	Idaho National Laboratory (INL)	Derek Gaston	Full Core Reactor Modeling with MOOSE	Nuclear Energy	1,000,000
MRCCSN	The University of Chicago (UChicago)/Fermi National Accelerator Laboratory (Fermilab)	Sean M. Couch	3D Simulations of Magnetorotational Core-Collapse Supernovae	Physics	30,000,000
MSMClimate_DD13	The National Center for Atmospheric Research (NCAR)	Greg Holland	Multi-Scale Modeling of Climate and High-Impact Weather	Earth Science	15,000,000
NanoInterfaces	The University of Chicago (UChicago)	Giulia Galli	Large-Scale Calculations on Nanostructured Heterogeneous Interfaces	Materials Science	25,000,000
NanoKappa	Argonne National Laboratory	Maria Chan	Prediction of Thermal Conductivity of Nanostructured Materials from Molecular Dynamics Simulations	Materials Science	1,000,000
NekCEM	Argonne National Laboratory	Misun Min	Electromagnetics	Physics	1,000,000
NekLBM	Argonne National Laboratory	Misun Min	Lattice Boltzamann Simulations for Fluids	Engineering	1,000,000
NEUP-FSI	The George Washington University	Elias Balaras	Large Eddy Simulations of Slender, Flexible Cylinders in Axial Flow	Engineering	1,500,000
NIMROD_xMHD	Tech-X Corporation	Jacob King	Extended Magnetohydrodynamic Simulations for Burning Plasma Experiments	Physics	2,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
NMGC-Mira-2013	University of Minnesota	J. Ilja Siepmann	Computations for the Development of the Nanoporous Materials Genome	Materials Science	8,100,000
noskov	University of Calgary	Sergei Noskov	ATP Transport in VDAC	Biological Sciences	2,000,000
NRCM_DD	Argonne National Laboratory	V. Rao Kotamarthi	Dynamic Downscaling of Climate Models	Earth Science	7,500,000
NUMA	Naval Postgraduate School	Andreas Mueller	Scalability Study for NUMA (Non-hydrostatic Unified Model of the Atmosphere)	Earth Science	3,000,000
OpenFOAM-ALCF	Argonne National Laboratory	Ramesh Balakrishnan	OpenFOAM-Based Computational Fluid Dynamics Simulations at the Argonne Leadership Computing Facility	Engineering	3,000,000
Operations	Argonne National Laboratory	Bill Allcock	Systems Administration Tasks	Internal	20,000,000
Operations_Test	Argonne National Laboratory	Cheetah Goletz	Operations Infrastructure Testing	Internal	10,000,000
OPV	The University of Chicago (UChicago)	Kenley Pelzer	Mesoscale Modeling of Charge Transport in Organic Photovoltaics	Materials Science	2,000,000
Outerrim	Argonne National Laboratory	Katrin Heitmann	Outer Rim Analysis	Physics	25,000,000
P3DFFT	University of California-San Diego/San Diego Supercomputer Center	Dmitry Pekurovsky	Performance Studies of Three-dimensional Fast Fourier Transforms Using Overlap of Communication with Computation	Computer Science	1,000,000
parallelQD	Texas Tech University	Bill Poirier	Massively Parallel Quantum Dynamics	Chemistry	14,000,000
ParamStudyBouss	Los Alamos National Laboratory (LANL)	Susan Kurien	Parameter Studies of Rotating Stratified Boussinesq Flows	Engineering	2,250,000
ParaOpt	Massachusetts Institute of Technology (MIT)	Qiqi Wang	Parallel Optimization of Turbulent Flow Simulations	Engineering	2,000,000
PARTS	Argonne National Laboratory	Kalyan Kumaran	Parallel Run-Time Systems	Computer Science	1,000,000
Performance	Argonne National Laboratory	Kalyan Kumaran	Performance	Internal	20,000,000
PESurfChem	University of Minnesota	Donald Truhlar, Osanna Tishchenko	Potential Energy Surfaces for Simulating Complex Chemical Processes	Chemistry	4,915,200
Peta_Spec	The University of Chicago (UChicago)	Marco Govoni	Petascale Spectroscopy	Physics	1,500,000
Petrel	The University of Chicago (UChicago)	Mike Papka	Petrel	Computer Science	100,000
PETSc	Argonne National Laboratory	Barry Smith	PETSc	Mathematics	1,000,000
PHASTA_NCSU	North Carolina State University (NCSU)	Igor A. Bolotnov	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	2,500,000
PIP_Regulation_PTEN	Johns Hopkins University	Thomas B. Woolf	How Does a PIP Control a State?: Simulations of the PI3K/PTEN Micro-Network	Biological Sciences	5,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
PLATINGSTRIPPING	Massachusetts Institute of Technology (MIT)	Gerbrand Ceder	Understanding the Mechanism of Ion Stripping and Depositions at the Anode Materials in Novel Multivalent Ion Batteries	Materials Science	1,000,000
Poly_Self_Assembly	Argonne National Laboratory	Dr. Derrick Mancini	Meso-scale Modelling of Self-Assembly of Polymer Grafted Nanoparticles	Materials Science	10,000,000
PorousHEDM_Shock	Purdue University	Alejandro Strachan	Shock Initiation of Energetic Materials via Reactive Molecular Dynamics	Materials Science	1,000,000
PPI_Entropy	The University of Chicago (UChicago)	Benoit Roux	Quantifying Protein-Protein Binding with Greatly Scalable Multiple-Copy Algorithms of NAMD	Biological Sciences	14,500,000
ProteinStruct	University of Washington	David Baker	Towards Breakthroughs in Protein Structure Calculation and Design (INCITE); Computational Protein Structure and Protein Design	Chemistry	20,000,000
PSU-CWF	The Pennsylvania State University	Balaji Jayaraman	Towards High-Fidelity 'Cyber' Computational Experiments of Atmospheric Turbulence-Driven Wind Turbine Aerodynamics and Wakes	Engineering	2,500,000
QBOX	University of California-Davis	Francois Gygi	QBOX	Materials Science	15,000,000
QDiags	Argonne National Laboratory	William Scullin	Blue Gene/Q Diagnostic Runs	Internal	83,361,792
QLG-Turbulence	The College of William & Mary	George Vahala	Quantum Lattice Algorithm for Quantum Turbulence	Engineering	950,000
qmcsc	University of Illinois at Urbana-Champaign	Lucas K. Wagner	First-Principles Quantum Monte Carlo for Superconducting Materials	Physics	500,000
QMC_Hyperoxides	Argonne National Laboratory	John J. Low	Quantum Monte Carlo Applied to Lithium Hyperoxides in Li-Air Batteries	Chemistry	10,000,000
RayBenard_DD	Occidental College	Janet Scheel	Turbulent Rayleigh-Benard Convection at High Rayleigh and Low Prandtl Numbers	Engineering	1,000,000
Repast	Argonne National Laboratory	Michael North	Exascale Agent-Based Modeling System	Computer Science	3,100,000
RESTMD_DD14	Boston University/Louisiana State University	Tom Keyes	Replica Exchange Statistical Temperature Molecular Dynamics in LAMMPS	Chemistry	2,000,000
RIKEN-Mira	RIKEN Advanced Institute for Computational Science (AICS)	Keiji Yamamoto	Collaborations between ALCF and AICS	Computer Science	1,000,000
rlins	Universidade Federal de Pernambuco	Roberto D. Lins	A Microscopic Perspective on Outer Membrane Remodeling and Antimicrobial Peptide Resistance	Biological Sciences	3,000,000
RNA-stcalculation	National Cancer Institute, National Institutes of Health	Yun-Xing Wang	Mira Request – Computing Three-dimensional Structures of Large RNA from Small-Angle X-ray Scattering Data and Secondary Structure	Biological Sciences	4,000,000
rtflames	Northwestern University	Elizabeth P. Hicks	DNS Simulations of Turbulent Rayleigh-Taylor Unstable Flames Using Nek5000	Physics	6,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
Scalable-IMEX	Argonne National Laboratory	Emil Constantinescu	Scalable Implicit-Explicit (IMEX) Algorithms and Software for Time Dependent Multimodel PDEs	Mathematics	800,000
Scalable_CRWENO	Argonne National Laboratory	Debojyoti Ghosh	Scalable Implementation of Weighted, Non-Linear Compact Scheme	Engineering	2,000,000
SDAV	Argonne National Laboratory	Mike Papka	SciDAC Scalable Data Management Analysis and Visualization	Computer Science	500,000
SEGMEnT_HPC	The University of Chicago (UChicago)	R. Chase Cockrell	Anatomic Scale Modeling with a Spatially Explicit General-Purpose Model of Enteric Tissue (SEGMEnT_HPC)	Biological Sciences	2,000,000
SEMO	Argonne National Laboratory	Murat Keceli	Fast Electronic Structure Methods for Rapid Reaction Screening for Inorganic Materials Synthesis and Particle Formation	Chemistry	500,000
shan-2013-shock	Sandia National Laboratories (New Mexico)	Tzu-Ray Shan	Nanostructure-Enhanced Detonation in Energetic Materials	Materials Science	7,500,000
shearbands	Columbia University	Haim Waisman	Robust Preconditioners for Shear Bands	Engineering	500,000
Shell_Cat	Shell International E&P, Inc.	Leonardo Spanu	Investigation of Catalytic Properties of Nanoclusters	Chemistry	2,000,000
SimAnneal	Rutgers University	Eric Altschuler	Simulated Annealing Methods on a Class of NP-Complete Problems	Mathematics	2,000,000
Singlet_Fission	The George Washington University	Hanning Chen	Probing the Free Energy Surface of Spin Separation in Singlet Fission	Chemistry	4,000,000
SMARTDESIGN	University of Cambridge (UK)	J.M. Cole	Smart Material Design for Optoelectronic Applications	Materials Science	5,000,000
Solitons	Argonne National Laboratory/University of British Columbia	Benjamin Gutierrez	Solitons in Field Theory and Applications	Physics	1,000,000
SOWFA	National Renewable Energy Laboratory (NREL)	Matt Churchfield	Large Eddy Simulations of a 36-Turbine Offshore Wind Plant	Energy Technologies	5,000,000
spectra_biomolecules	RIKEN Advanced Institute for Computational Science (AICS)	Shinichiro Nakamura, Ph. D.	Ab Initio-Level Calculations of Infrared and Raman Spectra of Biomolecules	Chemistry	5,000,000
StallNoise	Stanford University	Sanjiva K. Lele	Large Eddy Simulation of Stall Noise	Engineering	4,500,000
SteamGenerator	The University of Utah	Jeremy Thornock, Jennifer Spinti	Assessing Steam Generator Readiness for Compliance with Future NO_x and CO_2 Regulation	Engineering	3,884,449
SWIPE2014	Canadian Space Agency; Natural Sciences and Engineering Research Council of Canada (NSERC); University of Waterloo	Dr. Lucian Ivan, Prof. Hans De Sterck	Large-Scale Simulation of Planetary Environments Using Cubed-Sphere Grids: Software Package Preparation	Physics	500,000
TACOMA	General Electric Company (GE)	Brian E. Mitchell	TACOMA Porting and Scaling Study	Engineering	2,250,000
TBL_Poggie	Air Force Research Laboratory	Jonathan Poggie	DNS of Compressible Turbulent Boundary Layers	Engineering	5,000,000
TECA_Scaling	Argonne National Laboratory	Venkatram Vishwanath	Scaling the Toolkit for Extreme Climate Analysis on Blue Gene/Q Systems	Earth Science	4,000,000

Project Name	PI Institution	PI Name	Project Title	Science Field (Short)	Allocation Amount
Ti-ox_transport	Argonne National Laboratory	Olle Heinonen	Titanium Oxides Transport	Materials Science	10,000,000
TMDC_BILAYERS	University of Cambridge	Nicholas Hine	Transition Metal Dichalgonide Bilayer Heterostructures: Interaction Energies and Interlayer Couplings	Physics	2,500,000
Tools	Argonne National Laboratory	Scott Parker	ALCF Performance Tools	Internal	11,000,000
TotalView	Rogue Wave Software, Inc.	Peter Thompson, Ray Loy	TotalView Debugger on Blue Gene/P	Internal	2,000,000
TRG	Argonne National Laboratory	James Osborn	Tensor Renormalization Group	Physics	100,000
two-phase-flow	General Electric Company (GE)/Global Research	Madhusudan Pai	Towards Petascale First-Principles Simulations of Complex Two-Phase Flow Systems	Engineering	3,000,000
UdSaeroacoustics	Universite de Sherbrooke	Marlene Sanjose	LES of Turbulent Jet Noise	Engineering	1,612,800
Uintah_Safety	The University of Utah	Martin Berzins	Solving Petascale Public Health and Safety Problems Using Uintah	Chemistry	1,000,000
umn_crackle	University of Minnesota	Joseph Nichols	Large Eddy Simulation of Crackling Supersonic Jets	Engineering	10,000,000
urban_street_canyon	Illinois Institute of Technology (IIT)	Dietmar Rempfer	Characterization and Low-Dimensional Modeling of Urban Fluid Flow	Engineering	2,000,000
US-REST2	Argonne National Laboratory, University of Kansas	Sunhwan Jo	Developing Novel Umbrella Sampling/Solute Tempering Algorithms in NAMD	Biological Sciences	2,000,000
VERIFI_Workshop	Argonne National Laboratory	Sibendu Som	Visualizing the New Frontier: A VERIFI Workshop for High-Performance Computing-Enabled Engine Simulations	Engineering	2,000,000
Vestas_Park_LES	Vestas Wind Systems A/S	Dr. Gregory Oxley	LES Investigation of Stability-Enhanced Wake Losses on Large Wind Parks	Energy Technologies	12,500,000
visualization	Argonne National Laboratory	Michael E. Papka	Visualization and Analysis Research and Development for ALCF	Internal	183,333
Voth_DD14	Argonne National Laboratory/The University of Chicago (UChicago)	Gregory A. Voth	Multistate Reactive Molecular Dynamics: Development of Electrochemical and Fragment Molecular Orbital Methods	Chemistry	2,000,000
WATER_SPLITTING	The George Washington University	Hanning Chen	Photocatalytic Water Splitting by TiO2 Semiconductors	Chemistry	2,000,000
WLM	Syracuse University	Subas Dhakal	Computational Studies of the Topological Properties of Micellar Solutions	Physics	3,000,000
WLM-NP	Syracuse University	Abhinanden Sambasivam	Structure, Dynamics, and Rheology of Plasmonic Nanofluids Using Molecular Dynamics Simulations	Physics	4,000,000
xFDBenchmarking	Caterpillar, Inc.	Tushar Shethaji	Engine Combustion CFD Tool Performance Benchmarking for HPC readiness	Engineering	800,000
			Total Mira DD		**1,093,733,789**

Appendix C – Strategic Results Slides

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

Solving Petascale Public Health and Safety Problems Using Uintah

Martin Berzins, University of Utah

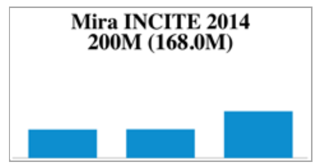

In 2005, a truck carrying 36,000 pounds of high explosives in Utah rolled over, caught fire and, within three minutes, underwent a transition from an ordinary thermal explosion to a fully developed detonation, destroying a substantial section of highway and adjacent railway. To help improve the safe transport of explosives, this project examines how different packing arrangements and storage techniques can be used to stop such transitions.

The research team uses simulations to better understand the mechanisms of large-scale deflagration-to-detonation transition (DDT) of packed explosives. Simulations of sufficiently high spatial and temporal resolution, to examine the impact of various explosive packing densities and configurations on DDT, required 32,768 nodes of Mira, and would not have been possible otherwise (Figure C.1).

For one study, the team investigated the reflection and coalescences of pressure waves in partially burned explosives. They determined that the pressure waves did not amplify as a function of mass burned, as was originally hypothesized. The researchers also determined how the pressure waves can transition into DDT and what other mechanisms could be responsible, depending on the geometry of the explosives and the packaging configuration.

In addition, the team has performed a large parametric study with simulations using 5,120 explosive cylinders to look at how packing density and spacing impacts the time to detonation. Results indicate a critical packing density of 40 percent for explosive cylinders with a radius of 0.027 m (the same radius as the cylinders involved in the 2005 truck accident). This means that if explosive cylinders are evenly spaced and occupy less than 40 percent of a container, they will not transition into detonation.

This work coincides with an ongoing National Science Foundation PetaApps project aimed at using simulation science to explore ways to use the Uintah computational framework to solve problems in public health and public safety. Uintah is a set of software components and libraries that facilitate the solution of partial differential equations on structured adaptive mesh refinement grids using hundreds to thousands of processors.

IMPACT: With a better understanding of the DDT process and potential mitigation strategies, this project stands to improve the safe storage and transport of explosive devices, which will ultimately improve the safety of our roads and railways. Additionally, the project is showing that it is possible to create the software needed to tackle such complex multiscale, multiphysics simulations with leadership-class supercomputers like Mira.

ALCF Contributions: Staff helped the team's job throughput on Mira by increasing the priority of the project's jobs in the queue, and with queue reservations, so the team could meet production deadlines (ALCF Scheduling Committee). Staff also helped the team improve the I/O performance in their application by using MPI-IO (Venkatram Vishwanath).

Presentations:

- J. Beckvermit, T. Harman, Andrew Bezdjian, Qingyu Meng, Alan Humprey, M. Berzins, and C. Wight, "Modeling Accidental Explosions and Detonations," SIAM Annual Meeting, July 2014, Chicago.
- J. Beckvermit, T. Harman, Andrew Bezdjian, Qingyu Meng, M. Berzins, and C. Wight, "Parallel Multiscale Modeling of Transportation Accidents Involving Explosives," XSEDE'14 Conference on High-Performance Computing, July 2014, Atlanta, Georgia.
- J. Beckvermit, "Macro-Scale Modeling of Deflagration to Detonation Transitions in Large Arrays of Explosives," Visualization Showcase, XSEDE'14 Conference on High-Performance Computing, July 2014, Atlanta, Georgia.
- Andrew Bezdjian, J. Beckvermit, T. Harman, M. Berzins, and C. Wight, "Critical Density Needed for a Deflagration to Detonation Transition," XSEDE'14 Conference on High-Performance Computing, July 2014, Atlanta, Georgia.

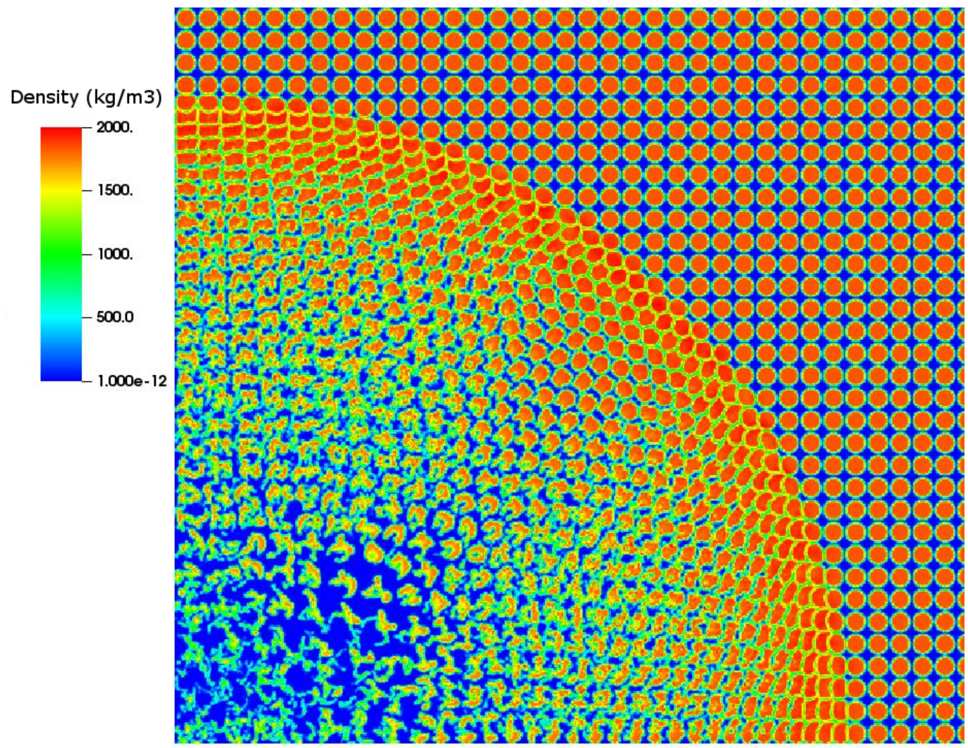

Figure C.1 Proposed DDT mechanism of inertial confinement. This image shows the density of the explosive cylinders late into the simulation.

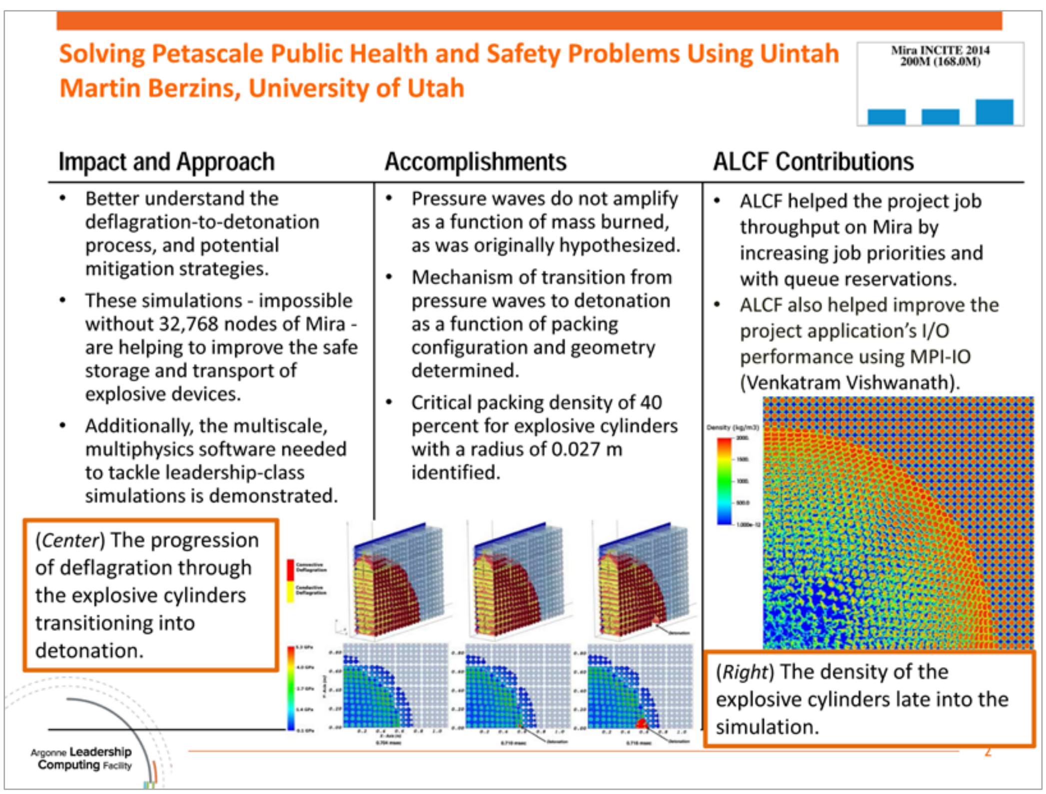

Molecular Modeling of Singlet Fission for Solar Energy Conversion

Hanning Chen, George Washington University

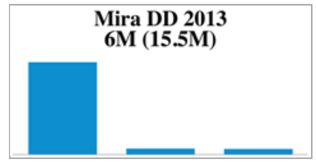

Singlet fission is a promising solution that may enable researchers to overcome the long-standing bottleneck in solar energy conversion by turning two singlet molecules into two triplet ones upon the absorption of a single incident photon. This project investigates the molecular mechanism of this photo-induced spin separation process in a rich library of crystalline materials by a novel functional mode electron transfer theory.

Research efforts have focused on large-scale molecular dynamics (MD) calculations to ascertain a full spectrum of vibrational normal modes of the 1-3-diphenylisobenzofuran (DPIBF) thin film before the spin separation process is examined by hybrid quantum mechanics/molecular mechanics (QM/MM) simulations. ALCF was integral in helping researchers skillfully run CP2K, the open-source molecular simulation package, at petascale on Mira, rendering a detailed picture of thermally driven spin transfer in molecular crystals.

Petascale simulations provide a free energy profile along the reaction coordinate of singlet fission in addition to the projection of the associated reorganization energy onto a large number of vibrational normal modes. It was found that a specific vibrational mode of DPIBF is strongly correlated with the photo-induced spin separation, also suggesting a very strong vibronic coupling during singlet fission in DPIBF due to the excessively large reorganization energy.

More importantly, the strong vibronic coupling in the DPIBF thin film indicates the feasibility of improving singlet fission efficiency by modifying the vibrational density of states by temperature (Figure C.2). For example, the pronounced anharmonicity of softened optical phonons at higher temperatures may result in a slower spin relaxation, which is responsible for the lower triplet quantum yield as observed by experiments.

IMPACT: These results deepen our understanding of singlet fission at the electronic structure level of theory—that is, the interplay between atomic displacements and electronic wave functions. Such detailed comprehension will speed our journey toward optimal solar energy conversion using environmentally friendly organic materials.

ALCF Contributions: Staff ported CP2K to BG/Q and helped implement the full vibrational spectrum QM/MM method by using a scalable ensemble algorithm (Wei Jiang). In addition, researchers were assisted with the efficient use of CP2K on BG/Q; these enhancements included improving on the default OpenMP runtime environment and improving the intranode parallelism of key routines using compiler directives (Wei Jiang).

Publications:

- H. Chen, J. Phys. Chem. B. **118**, 7586 (2014).
- J. E. Elenewski and H. Chen, Phys. Rev. B **90**, 085104 (2014).

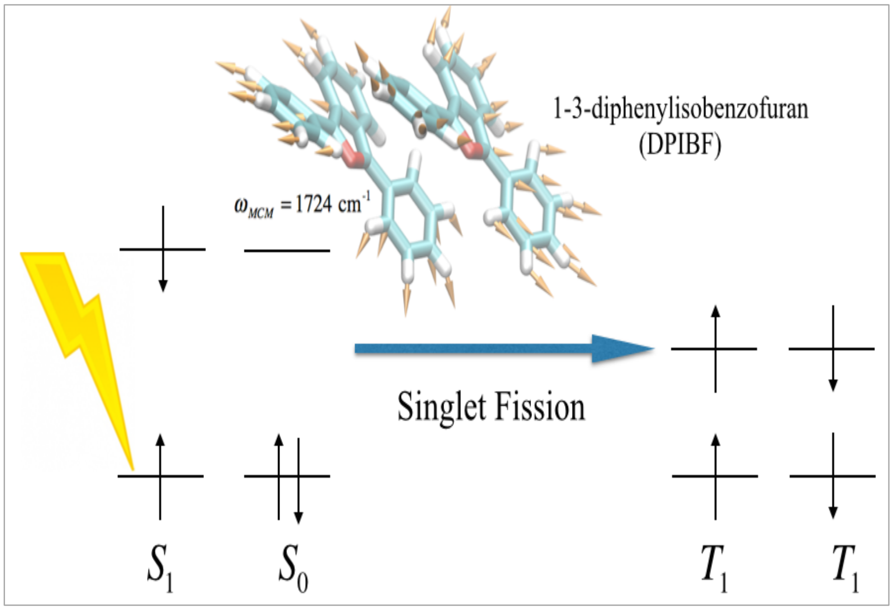

Figure C.2 Diagram depicting singlet fission. This emerging technology may help improve solar cell efficiency by turning two singlet molecules into two triplet molecules upon the absorption of an incident photon.

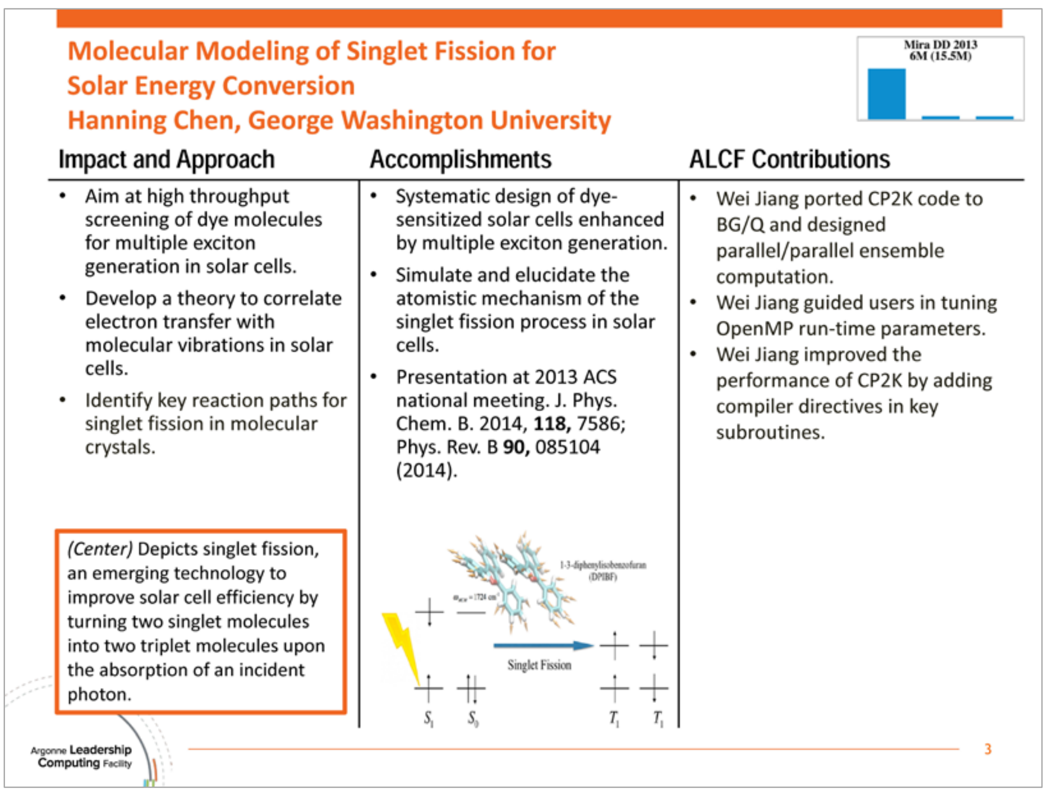

Predictive Materials Modeling for Li-Air Battery Systems

Larry Curtiss, Argonne National Laboratory

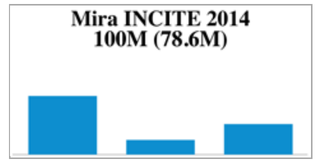

Lithium-air (Li-air) batteries are viewed as a possible game changer for electric vehicles. However, realizing their enormous potential requires developing new materials for electrodes and electrolytes—a challenging scientific problem. Scientists from Argonne and IBM Research are teaming up to use Mira to better understand the physical and chemical mechanisms needed to make Li-air batteries a reality.

With the potential to store up to 10 times the energy of a Li-ion battery of the same weight, Li-air batteries are particularly appealing to researchers because of their theoretical energy density. But developing a viable Li-air battery is a long-range effort that requires scientific breakthroughs in materials design, chemistry, and engineering.

One of the most significant hurdles is finding suitable materials for Li-ion-conducting electrolytes, which enable the transport of ions between the anode and the cathode and promote the diffusion of oxygen from the environment into the electrochemical cell. With this INCITE project, Argonne and IBM Research are conducting ab initio density functional theory (DFT) simulations on Mira to help accelerate efforts to identify novel materials for electrolytes.

In a recent study, IBM researchers focused on the zirconium-containing, garnet-like lithium-lanthanum-oxide known as LLZO ($Li_7La_3Zr_2O_{12}$), a promising material for solid-state electrolytes. Experimental work has determined that LLZO is much better at conducting Li-ions in its cubic phase as compared to its tetragonal phase, but computer simulations were required to investigate and compare the different mechanisms of Li-ion migration in the two phases.

To observe Li-ion migration, the team simulated timescales in nanoseconds rather than picoseconds. Researchers achieved this factor-of-a-thousand improvement by using metadynamics (a method for accelerating rare events, such as mapping the conductivity of a material regardless of its complexity) in their DFT simulations and by taking advantage of Mira's substantial power with full machine runs. Using metadynamics enabled the researchers to obtain, for the first time, the free-energy profile for Li-ion conductivity in LLZO (Figure C.3). One of their key findings was that the presence of vacancies in cubic LLZO is crucial to lowering its activation energy and enhancing Li-ion conductivity.

IMPACT: This project is providing insight into the complexities of the Li-air battery at the molecular level, including an understanding of the microscopic mechanism for high Li-ion conductivity. The results will help inform the design of new materials for Li-air electrolytes and electrodes. If realized, Li-air batteries could enable the widespread deployment of electric vehicles, greatly reducing U.S. dependence on foreign oil.

ALCF Contributions: The team at IBM that carried out these calculations was self-sufficient and required little technical assistance. Staff helped the team increase their job throughput in the Mira queues (Nichols Romero).

Publication: K. Meier, T. Laino, A. Curioni, J. Phys. Chem. C **118**, 6668 (2014).

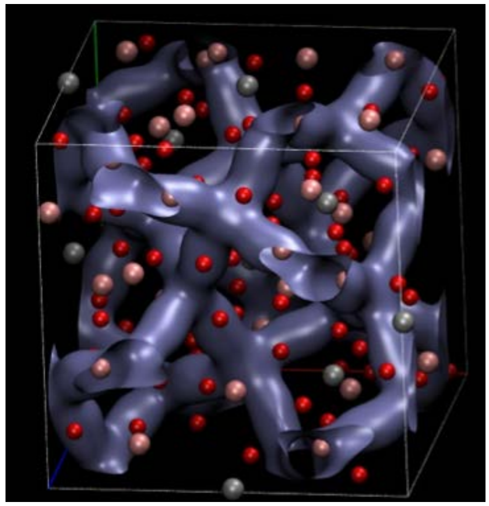

Figure C.3 An image of the LLZO material (200 atoms). The tubes depict iso-surface representations of the free-energy surface of Li-ion conductivity. Within the tubes, Li atoms (pink) can be observed moving inside the material.

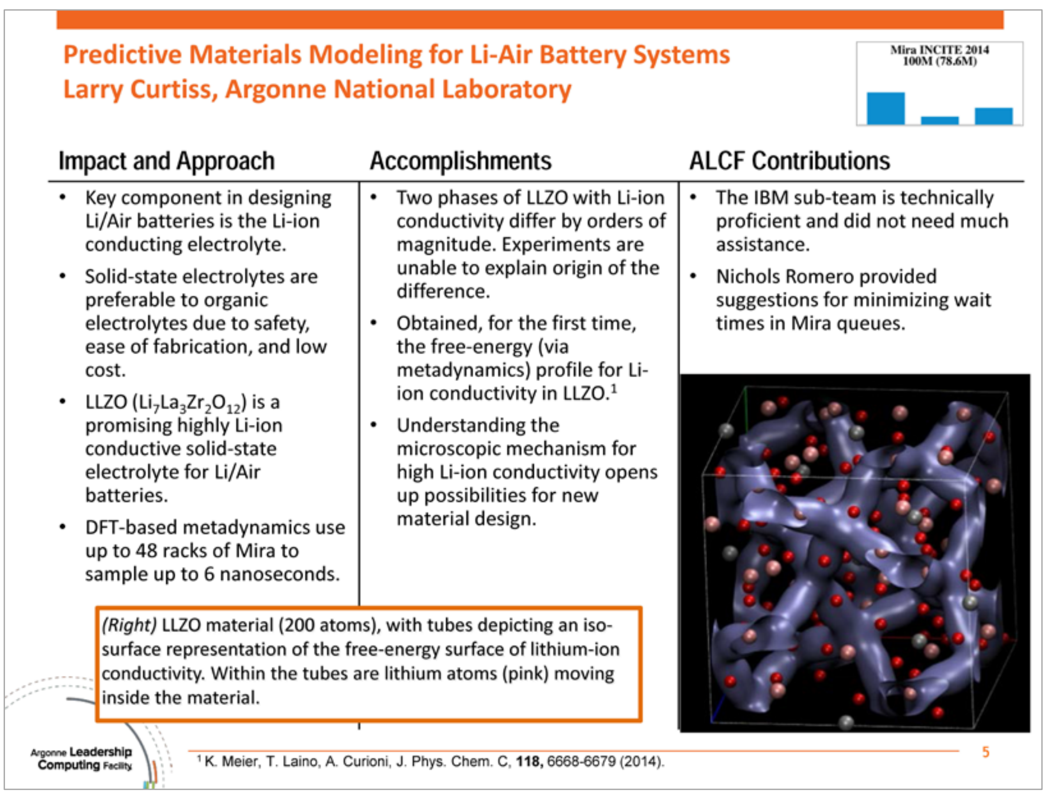

Thermal-Hydraulic Modeling: Cross-Verification, Validation, and Co-Design

Paul Fischer, Argonne National Laboratory

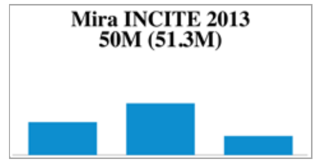

Generating clean, safe nuclear power is essential to meeting the world's growing energy needs, and the use of fast neutron reactors specifically designed to address nuclear waste management concerns is of key interest. Advanced simulation is critical in bringing fast reactor technology to fruition in an economical and timely manner. Paul Fischer's team is performing highly accurate simulations of complex thermo-fluid phenomena in order to characterize the peak material temperature and pressure drops in coolant flow—key safety parameters that feature prominently in fast reactor design. The thermal-hydraulic behavior of multi-pin subassemblies with wire-wrap and grid spacers was identified as a problem of primary interest given that experimental data are available for validation; thus, this problem was chosen as a key area for this multi-year effort.

Fischer's team has investigated the turbulent flow through a 37-pin rod bundle, corresponding to the experiment of Krauss and Meyer. The team has performed the highest-resolution investigation of this flow to date with the Nek5000 code on Mira in order to resolve turbulent transport mechanisms present near the walls of this configuration (Figure C.4). Walker et al. (2013) show that the accurate prediction of spatially varying wall-shear stress requires substantial near-wall resolution for large eddy simulation. Wall-shear stress is an important high-order parameter (a derivative of the velocity) for determining the coolant pressure drop in a nuclear fuel assembly.

In addition, a new benchmark exercise has been initiated as part of an international collaboration with Japan, France, and the United States to compare code predictions with the results of the PLAJEST sodium triple-jet thermal mixing experiment and the WAJECO water-version of the experiment. Initial conditions were generated using Nek5000 for the simulations relating to this experiment (Figure C.5). (A movie of the temperature field that was generated is available upon request.)

In addition to performing high-fidelity calculations of turbulent flow for nuclear reactor assemblies, Fischer's team is evaluating algorithms for the future exascale simulation of nuclear reactors under the CESAR program (Center for Exascale Simulation of Advanced Reactors). Re-casting today's classic bulk-synchronous algorithms into more asynchronous, coarse-grained parallelism is a critical area of research for next-generation computing. The team proposed a robust class of parallel algorithms for explicit, stencil-based discretizations in the presence of non-uniformities in processor execution time that are likely to become more prominent at exascale. These algorithms were tested in the presence of simulated processor interruptions, and their timing was compared to classical algorithms on a simple model problem. The results showed that factors of up to three in speedup can be obtained for a broad range of parameter values.

IMPACT: High-fidelity simulation allows the accurate prediction of complex thermal-hydraulic parameters, which can improve fast neutron reactor safety. Also, the initial investigation of future exascale algorithms will help transition today's codes to future architectures.

ALCF Contributions: Staff helped the project team move data to tape storage in preparation for the decommissioning of Intrepid (Emily Shemon). ALCF also set up the reservations and special queues needed to complete work in time for an international conference (ALCF Scheduling Committee).

Publications:

- J. Walker, E. Merzari, A. Obabko, P. Fischer, and A. Siegel, "Accurate Prediction of the Wall Shear Stress in Rod Bundles with the Spectral Element Method at High Reynolds Numbers," *International Journal of Heat and Fluid Flow,* Volume 50, December 2014. 287-299.
- A. Hammouda, A. R. Siegel, and S. F. Siegel, "Noise-Tolerant Explicit Stencil Computations for Nonuniform Process Execution Rates," *ACM Transactions on Parallel Computing*, Accepted November 2014.

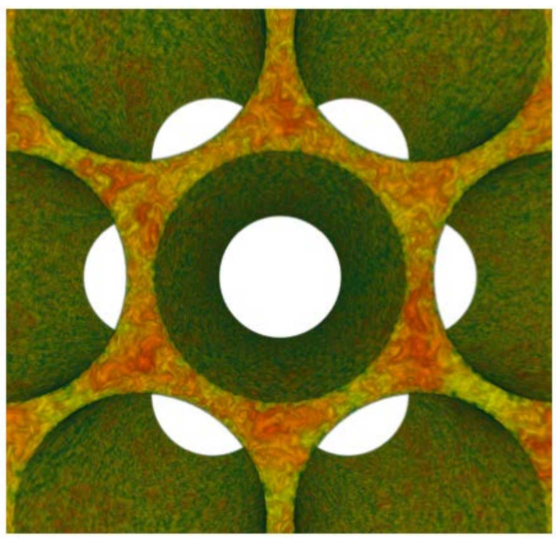

Figure C.4 Nek5000 velocity results for 37-pin rod bundle of Krauss and Meyer experiment. Results depict the range of scales (and hence the large number of degrees of freedom) needed to resolve small features and accurately calculate wall-shear stress in large eddy simulations.

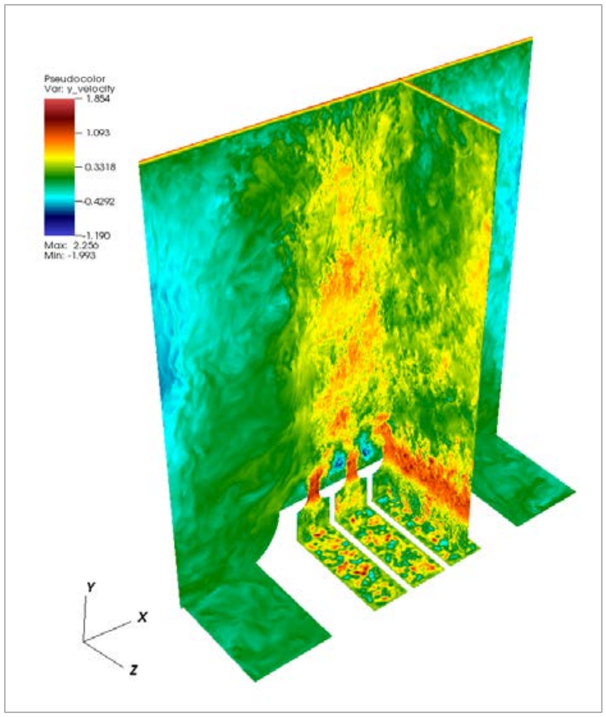

Figure C.5 Vertical velocity calculated by Nek5000 in the triple-jet thermal mixing experimental benchmark WAJECO.

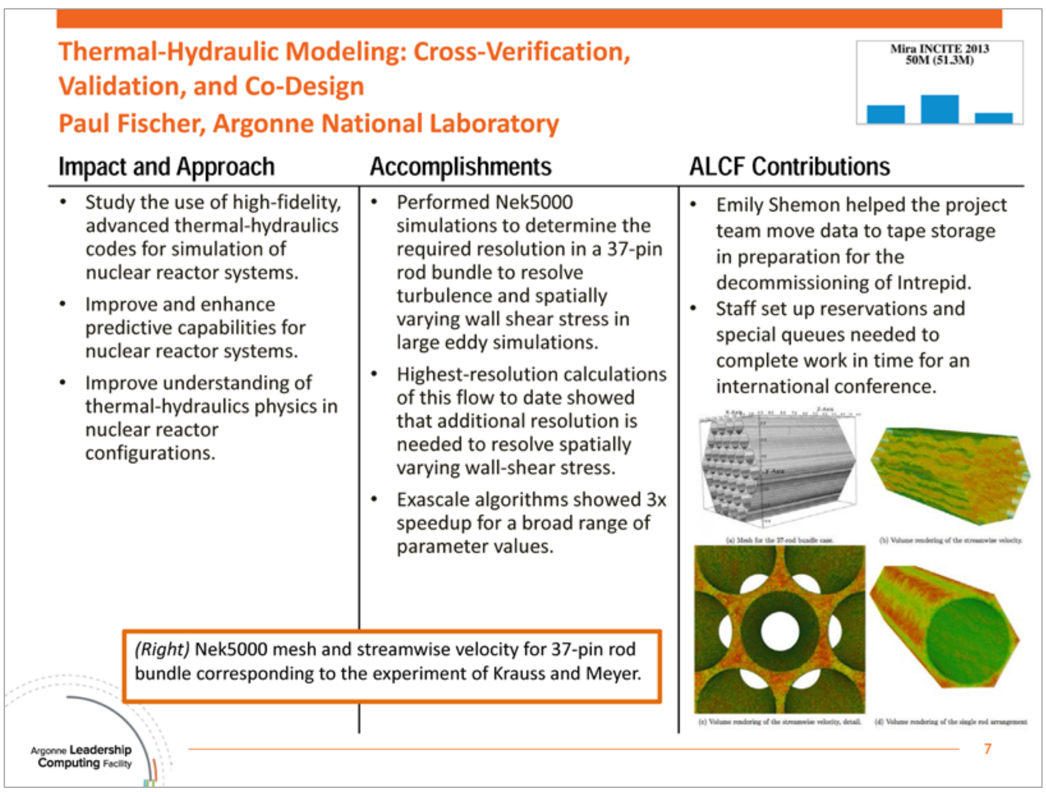

Thermal-Hydraulic Modeling: Cross-Verification, Validation, and Co-Design

Paul Fischer, Argonne National Laboratory

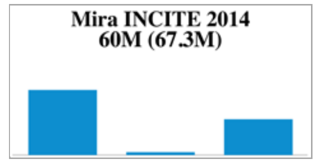

The 2011 accident at Japan's Fukushima Daiichi nuclear power plant involved a series of several events ultimately leading to radioactive release and the disabling of several reactors. In fact, the convection and mixing flow in the containment facility played an important role in the Fukushima accident as buoyant hydrogen gas mixed with oxygen and detonated, resulting in significant destruction and radioactive pollution. Following the Fukushima disaster, the Nuclear Energy Agency within the Organization for Economic Co-operation and Development chose the PANDA facility for a 2014 benchmark exercise where predictive capabilities of computational fluid dynamics (CFD) tools are tested for multispecies convection under a notorious transition regime from turbulent to laminar flow and from forced to natural convection. The accurate prediction of these phenomena will increase understanding of reactor behavior during accidents and help with the design of safer, more efficient reactors for a carbon-free energy option.

Designed to study buoyancy, convection, and mixing flows, the PANDA experiment facility is a multi-compartment, large-scale thermal-hydraulics test rig located at the Paul Scherrer Institute in Switzerland. Initially, the experiment vessel contains a well-defined gas mixture of helium/air at the top, and air below with measured stratification. A jet of air and buoyant helium is injected into the vessel, eventually eroding the initial stratification layer (Figure C.6) and posing significant challenges for the accurate prediction of flow evolution.

Scientists used the Nek5000 code to conduct a series of numerical experiments to optimize computational cost and accuracy in this type of flow prediction. Insights were made into simulating this difficult flow regime with Nek5000, including discovery of the optimal meshing, and changing the boundary conditions at the inlet. The team tested varying approximations for modeling buoyancy and diffusion effects to determine the effect on the flow.

IMPACT: The Nek5000 code has now been studied under flow regimes previously not tested, which are important for nuclear reactor containment accident scenarios. Researchers gained important insights that will optimize computational cost and accuracy in this type of flow prediction. Eventually the results will be compared with the experimental benchmark data to validate Nek5000 in these conditions.

ALCF Contributions: This INCITE team is well-established on Blue Gene systems and required little technical assistance. ALCF helped them simulate extra-long jobs on Mira to reach longer times in the simulation (ALCF Scheduling Committee). ALCF also provided assistance with the visualization cluster Tukey.

Publications:

- A. Tomboulides, S. M. Aithal, P. Fischer, E. Merzari, and A. Obabko, "A Novel Variant of the K-Omega uRANS Model for Spectral Element Methods—Implementation, Verification and Validation in Nek5000," *Proceedings of the ASME 2014 4th Joint U.S.-European Fluids Engineering Division Summer Meeting & 11th International Conference on Nanochannels, Microchannels, and Minichannels,* FEDSM2014-21926, August 3–7, 2014, Chicago, Illinois.
- A. Hammouda, A. R. Siegel, and S. F. Siegel, "Noise-Tolerant Explicit Stencil Computations for Nonuniform Process Execution Rates," *ACM Transactions on Parallel Computing*, Accepted November 2014.

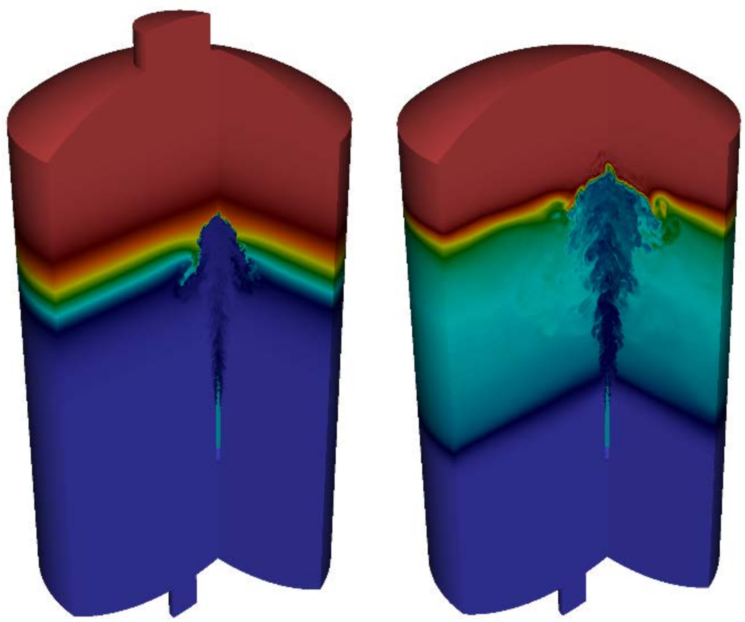

Figure C.6 This visualization depicts the helium mass fraction (red=high, dark blue=low) of the mixture inside the PANDA containment at early time *(left)* and late time *(right).* At late times, the helium has mixed with the lower air-rich layer significantly and has also completely penetrated the air-helium stratification layer.

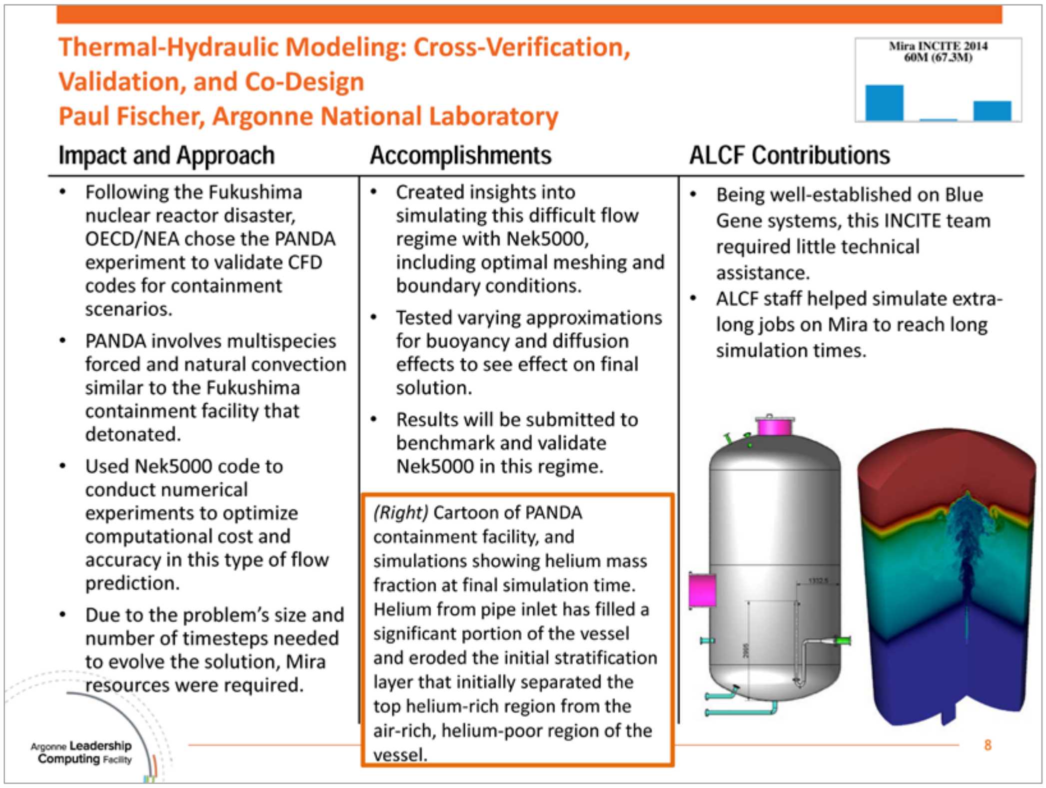

Vibrational and Optical Spectroscopy of Electrolyte/Solid Interfaces

Giulia Galli, University of Chicago

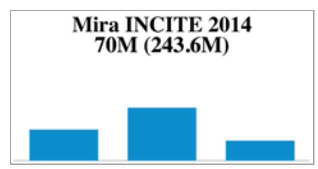

Photoelectrochemical (PEC) energy conversion is a promising approach to hydrogen production that uses solar energy to split water into hydrogen and oxygen gases. To help accelerate research and development efforts in this area, scientists from the University of Chicago and the University of California, Davis, are using Mira to build novel computational tools and generate knowledge that can be used to better understand PEC processes at the microscopic scale.

One of the key challenges to developing scalable and commercially viable PEC cells is identifying stable and efficient photoelectrode materials. For this INCITE project, the research team is carrying out large-scale simulations at ALCF to model the physical and chemical processes occurring at the interface between solid photoelectrodes and electrolytes. The simulation results can be used to create design rules that help predict the optimal photoelectrode materials to use for PEC water splitting.

To enable these studies, researchers developed a set of simulation codes to probe and predict vibrational and electronic properties of the solid-liquid interfaces. In particular, they substantially improved the performance, scaling, and parallelization of Qbox (a first-principles, molecular dynamics code developed at the University of California, Davis) on Mira. The Qbox development team also added modules to compute Raman and surface-sensitive sum-frequency generation spectra, including hybrid density functionals.

The use of hybrid functionals allowed the researchers to calculate the electronic properties of electrolytes, a challenging task for ab initio simulations. This computational approach, applied to both sodium chloride (Figure C.7) and sulfuric acid in water, takes full advantage of the massively parallel capabilities of Mira.

The team also developed algorithms and codes to enable first-principles calculations of the electronic properties of complex surfaces and interfaces within many-body perturbation theory at the GW level, enabling GW calculations of unprecedented size for slabs containing solid-liquid interfaces. Although this INCITE project is focused on PEC, the methodologies and results stemming from this work are relevant to other energy-related research efforts, including electrical energy storage and solar cells.

IMPACT: This project is providing knowledge and computational tools to interpret and inform ongoing experiments on hydrogen production from water. The results can be used to establish design rules to predict the materials that will be optimally suited to reducing and oxidizing water, which could ultimately help accelerate the development of technologies for sustainable and clean hydrogen production.

ALCF Contributions: Prof. Galli's team worked with staff to leverage parallel I/O libraries (Venkatram Vishwanath), port their codes to BG/Q, and improved the performance of the Qbox and Quantum Espresso codes (Alvaro Vazquez-Mayagoitia).

Publication: A. P. Gaiduk, C. Zhang, F. Gygi, and G. Galli, Chem. Phys. Lett. **604**, 89 (2014).

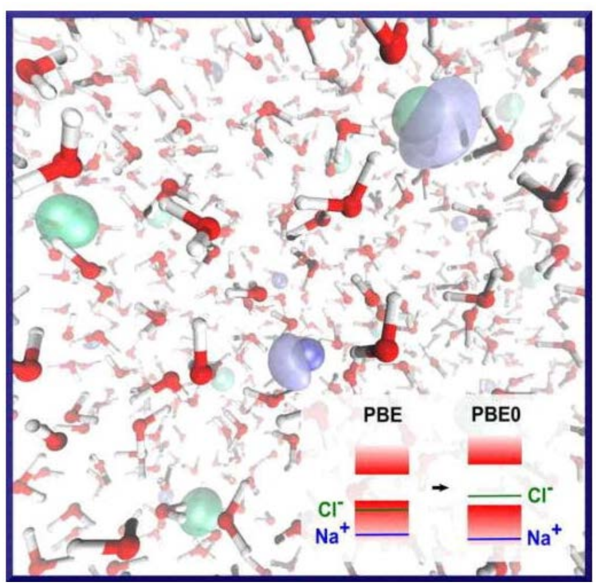

Figure C.7 Snapshot of a 1-M solution of NaCl obtained from an ab initio molecular dynamics simulation with a hybrid (PBE0) functional. Red, white, purple, and green spheres represent oxygen, hydrogen, sodium, and chlorine atoms, respectively. The electronic properties of $[NaCl]_{aq}$, computed with PBE0 and semi-local (PBE) functionals, are shown in the right corner of the figure.

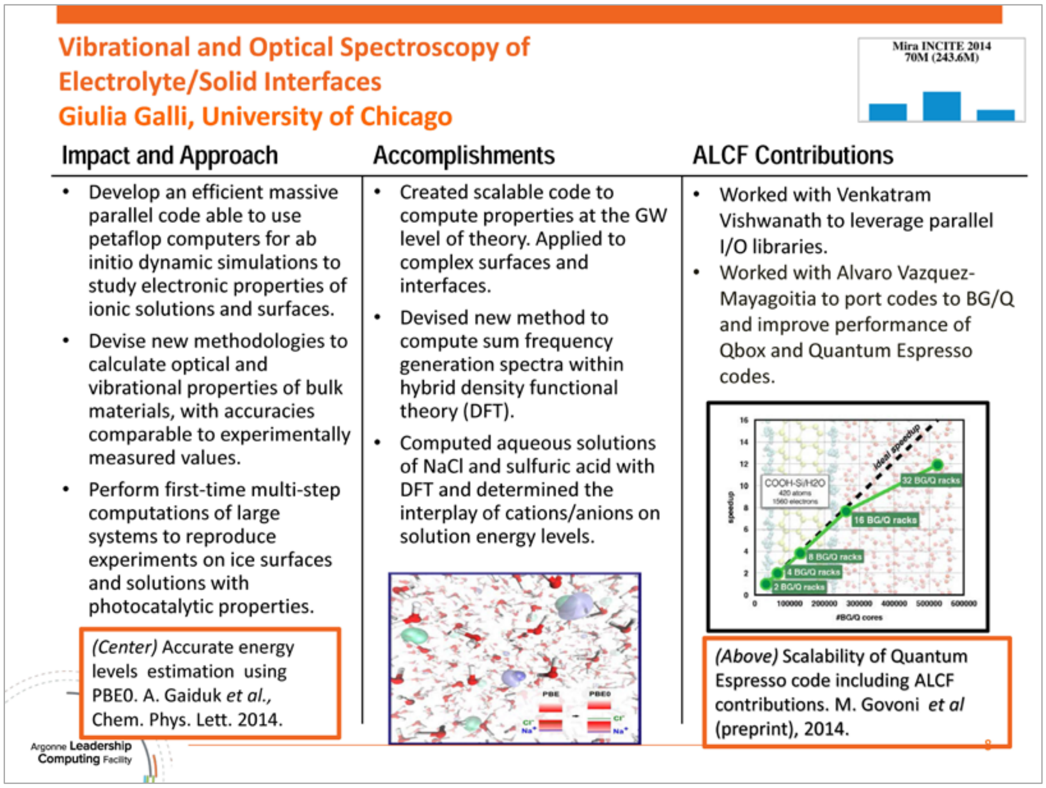

High-Fidelity Simulation of Complex Suspension Flow for Practical Rheometry

William George, National Institute of Standards and Technology

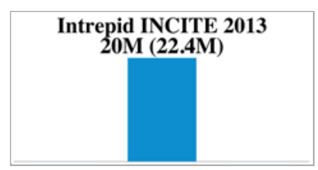

Concrete begins as a thick, pasty fluid containing innumerable particles in suspension that, ideally, can flow into a space of nearly any shape, where it hardens into a durable, rock-like state. Its initial flexibility, combined with its eventual strength, has made it the material of choice for building everything from the ancient Roman Coliseum to the foundations of countless modern bridges and skyscrapers.

But just because concrete is the most widely used building material in human history doesn't mean it can't be improved. A study conducted by a team of scientists from the National Institute of Standards and Technology (NIST), the University of Strasbourg, and Sika Corporation that used Intrepid has led to a new way to predict concrete's flow properties from simple measurements. By circumventing the need for complex and expensive testing, this new capability will speed the development of new concrete mixtures that not only flow better, but are friendlier to the environment.

While it's a simple goal to describe, accomplishing it demanded some complex math and physics and, at the same time, an enormous amount of computer power to study how all the particles and fluid react as they are mixed. Making use of 32,768 cores on Intrepid, William George's team simulated how a suspension would change if one or more parameters varied—the number of suspended particles, for example, or their size (Figure C.8). Results from these simulations will be used to help interpret simulations using a vane rheometer with the ultimate goal to help calibrate rheometers used in experiments.

Suspensions have a remarkable property: Plotting two of its parameters—viscosity versus shear rate—always generates the same shaped curve as plotting them only for the suspending fluid alone without added particles. This is true no matter what fluid is used. The curve just sits on a different place on the X-Y axis, as though someone had pushed it upwards or off to the side without otherwise altering it.

What the team unexpectedly found was the amount that the curves had to be shifted could be predicted based on the microscopic shear rates that existed between neighboring particles. Experiments at the University of Strasbourg confirmed the simulated results, which allowed the team to come up with a general theory of suspensions' properties.

These results should help accelerate the design of a new generation of high-performance and eco-friendly cement-based materials by reducing time and costs associated with R&D. Scientists are looking at several materials that could replace traditional cement ingredients, such as slag and fly ash, two industrial waste products that ordinarily require disposal. While it is not yet known whether these materials will be ideal candidates, the NIST team's research could eventually help identify a suitable candidate.

The description of this science accomplishment is, in part, based on a NIST press release from January 21, 2015: http://www.nist.gov/itl/math/concrete-012115.cfm.

IMPACT: Accelerate the design of a new generation of high-performance and eco-friendly cement-based materials by reducing time and costs associated with R&D. Scientists are looking at several materials that could replace traditional cement ingredients, such as slag and fly ash, two industrial waste products that ordinarily require disposal.

ALCF Contributions: The team was self-sufficient for code development. Staff provided assistance and guidelines to improve turnaround time on Intrepid to achieve their milestones (Marta García).

Publication: M. Liard, N. S. Martys, W. L. George, D. Lootens, and P. Hébraud, Journal of Rheology, **58**, 1993 (2014).

Presentation: N. S. Martys, D. Lootens, W. L. George, P. Hébraud, and M. Liard, "Universal Scaling of Microscopic and Macroscopic Behavior in Spherical Non-Colloidal Suspensions with a Non-Newtonian Fluid Matrix," Society of Rheology 85th Annual Meeting, October 16, 2013, Montreal, Canada.

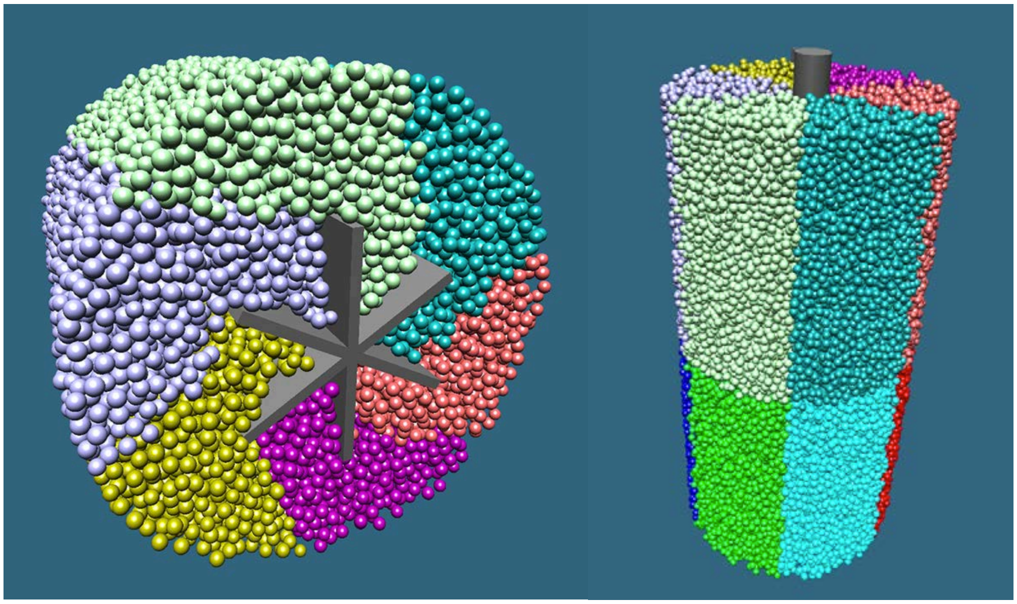

Figure C.8 These simulation images show suspended particles in a 6-blade rheometer for NIST's proposed mortar Standard Reference Materials (SRM). The spheres, which are color-coded by their starting location in the rheometer, are suspended in a cement paste with properties derived from NIST's cement paste SRM.

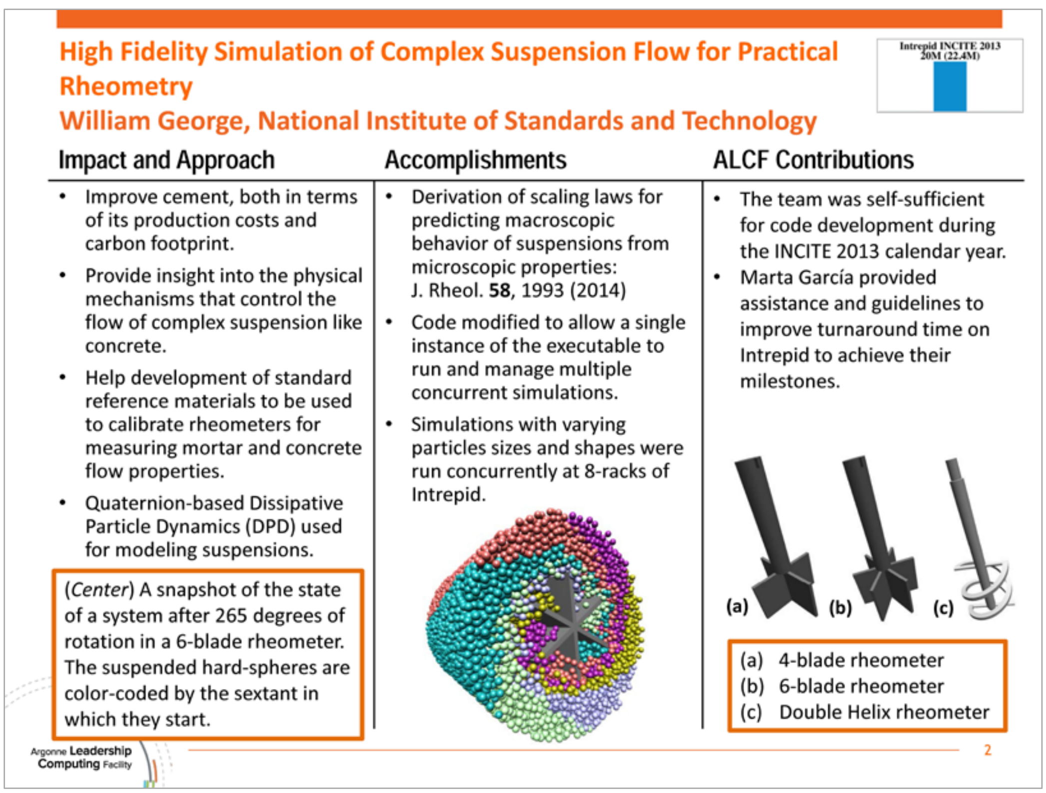

Does a Turbulent Duct Flow Ever Become Two-Dimensional?

Hassan Nagib, Illinois Institute of Technology (IIT)

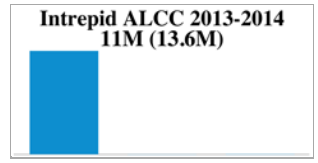

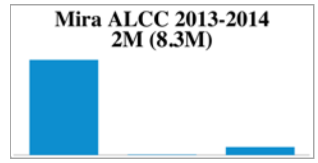

Numerous real-world applications involve the flow of fluid through a rectangular duct: urban drainage systems, ventilation systems, and combustion engines, to name a few. Globally, about 10-20 percent of energy consumption is used to overcome turbulent drag in one way or another. A deeper understanding of flow physics in rectangular ducts could lead to the discovery of scientific methods to reduce drag and skin friction, thereby also reducing global energy consumption.

Traditional computational studies of turbulent flow within a duct have always assumed that the duct can be treated as infinitely wide (2D) if the duct is wide enough. However, recent experiments by Nagib and co-investigator Ricardo Vinuesa showed that this 2D assumption is not representative of the physical flow for many realistic flow regimes. Using the computational resources at ALCF, Nagib and Vinuesa are studying the behavior of turbulent flow through rectangular ducts for an increased number of aspect ratios and more realistic Reynolds numbers using the highly scalable Nek5000 code on Intrepid (BG/P) and Mira (BG/Q) in order to better understand the results of experimental measurements carried out at IIT.

Several duct aspect ratios (ratios of width/length) were studied; the analysis of their computational results led to some interesting physical conclusions. At higher aspect ratios (ranging from 12.8 to 48), the calculated skin friction decreases as the aspect ratio increases up to 24, where it then appears to reach an asymptotic value; this result is consistent with experimental data. However, when the aspect ratio is increased from 1 to 3, the friction actually increases rather than decreases (Figures C.9 and C.10). Close inspection of the computational flow determined that the increased friction results from a complex relationship between the side-wall boundary layer thickness and secondary motions in the flow. These secondary motions significantly affect the overall flow physics but have not been reported in the literature because they are extremely difficult to detect in the laboratory. The secondary vortices and side-wall boundary layers have been studied in detail, as well as their evolution with Reynolds number and aspect ratio. Understanding these 3D effects through computation leads to improved computational models and better fundamental understanding of the flow physics.

IMPACT: The data produced in this work provides justification to redirect efforts of the computational fluid dynamics community from studying flows that have no counterparts in real life, such as the z-periodic channel, to more relevant wall-bounded and ducted flows. The deeper understanding gained in such simulations will eventually be used to improve prediction, design, and control of ducted flows and other flow geometries. In the meantime, the computational models helped show the existence of important fluid flow physics that have been postulated but are difficult to observe in the laboratory.

ALCF Contributions: Previous efforts by ALCF have led to the porting and efficient use of Nek5000 of Blue Gene/P machines (Scott Parker). By the time this project was running, the Nek5000 code was already stable and optimized for this platform. Staff provided assistance with visualization/data analysis using VisIt.

Publication: R. Vinuesa, E. Bartrons, D. Chiu, J.-D. Ruedi, P. Schlatter, A. Obabko, and H. M. Nagib, "On Minimum Aspect Ratio for Experimental Duct-flow Facilities," Progress in Wall Turbulence: Understanding and Modelling, June 18–20, 2014, Lille, France.

Presentation: H. M. Nagib and R. Vinuesa, 2013, "The Myth of Laboratory Two-Dimensional Wall-Bounded Turbulent Flows," High Reynolds Number Boundary Layer Turbulence: Integrating Descriptions of Statistical Structure, Scaling, and Dynamical Evolution Workshop, November 22, 2013, University of New Hampshire, Durham, New Hampshire.

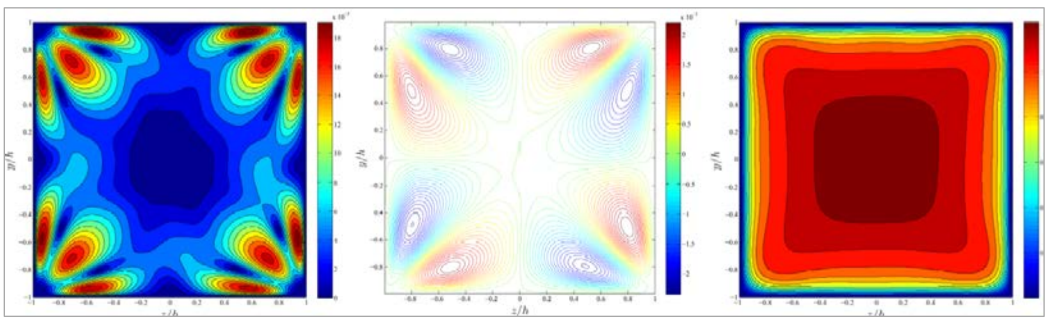

Figure C.9 In-plane velocity, streamlines, and streamwise velocity of the aspect ratio 1 case.

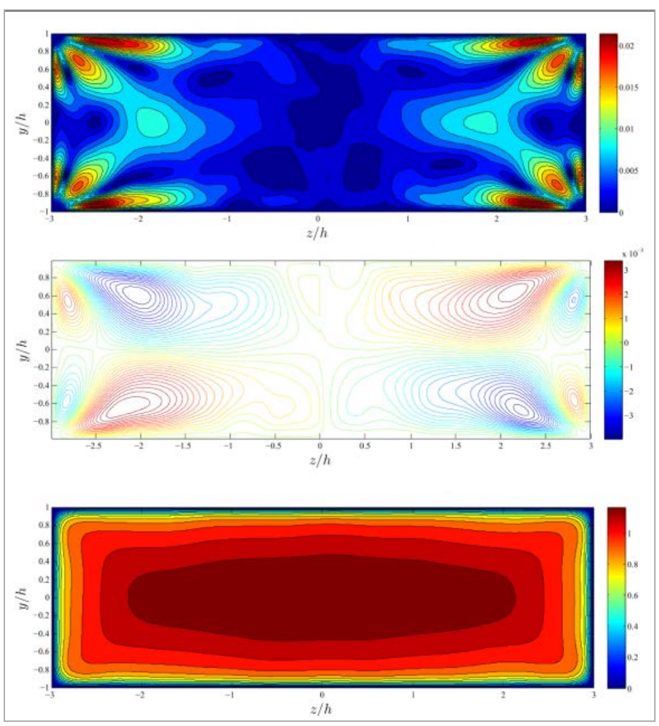

Figure C.10 In-plane velocity, streamlines, and streamwise velocity of the aspect ratio 3 case.

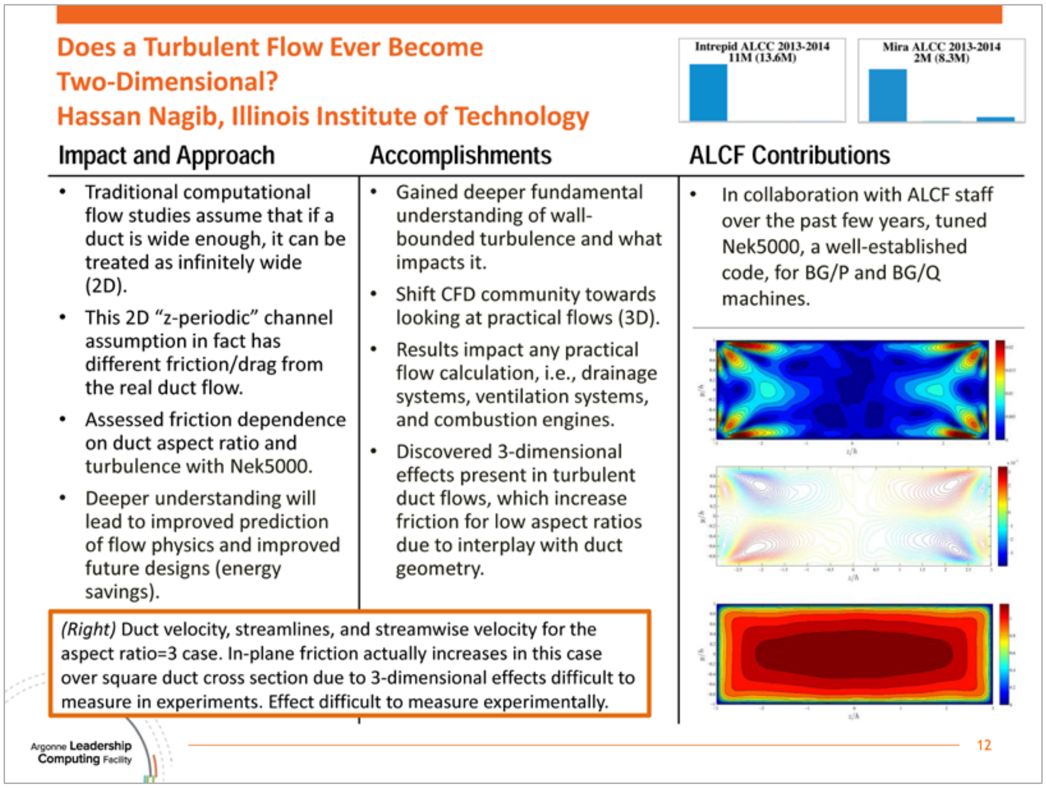

Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus Simulations

Roberto Paoli, CERFACS

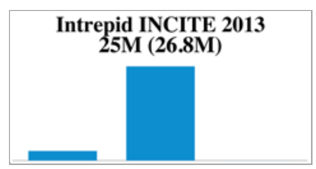

Contrails, the visible white lines in the sky left behind an airplane, are ice clouds made by water exhaust from the aircraft's engine. Contrails can spread to form cirrus clouds that affect the Earth's radiation budget. As the volume of air travel continues to grow, the impact of more contrail cirrus becomes an increasing concern for scientists and policymakers.

A research team from CERFACS is using ALCF computing resources to fine-tune numerical models that will allow for a more accurate understanding of the impact of contrail cirrus on global climate. With a 2012 INCITE award, the researchers used high-resolution large eddy simulations (LESs) to characterize the mechanisms that control the transition from the contrail stage to the young cirrus stage as a function of the wake age. These simulations need fine grids to capture the smallest flow structures and large domains to accommodate the contrail expansion. Hence, they can be run only on large supercomputing systems. The research team observed that atmospheric turbulence is the main driver of contrail evolution after vortex breakup; whereas, at later stages, radiative transfer and sedimentation affect the transition by controlling the plume's horizontal and the vertical spreading.

With their 2013 INCITE allocation, the team examined contrail evolution up to one hour after emission. One important finding of their work is that radiative transfer is the key mechanism controlling the global ice characteristics and the vertical extension of the contrail, and is predominant on the atmospheric turbulence on a time scale of one hour after the contrail has formed (see Figure C.11). This effect is magnified during the day by direct heating mostly at the top of the contrail in the visible band that adds to the heating at the bottom in the infrared band from the Earth's surface. The team is now using a Director's Discretionary allocation at ALCF to carry out a detailed analysis of the microphysical and optical properties of contrails, and to propose parameterizations of these quantities for global and climate models.

IMPACT: By understanding and characterizing the physical mechanisms that control the formation of contrail cirrus, this INCITE research led to more complete and effective high-resolution atmospheric LES models for the development of parameterizations of aircraft perturbations in next-generation global and climate models.

ALCF Contributions: Staff proposed improvements to optimize the I/O memory management and the use of the global arrays library to increase I/O scalability (Marta Garcia, Jeff Hammond, Vitali Morozov). This work helped improve the code performance on Mira. Staff also assisted the team by allowing temporary access to dedicated queues when long runs were necessary (ALCF Scheduling Committee).

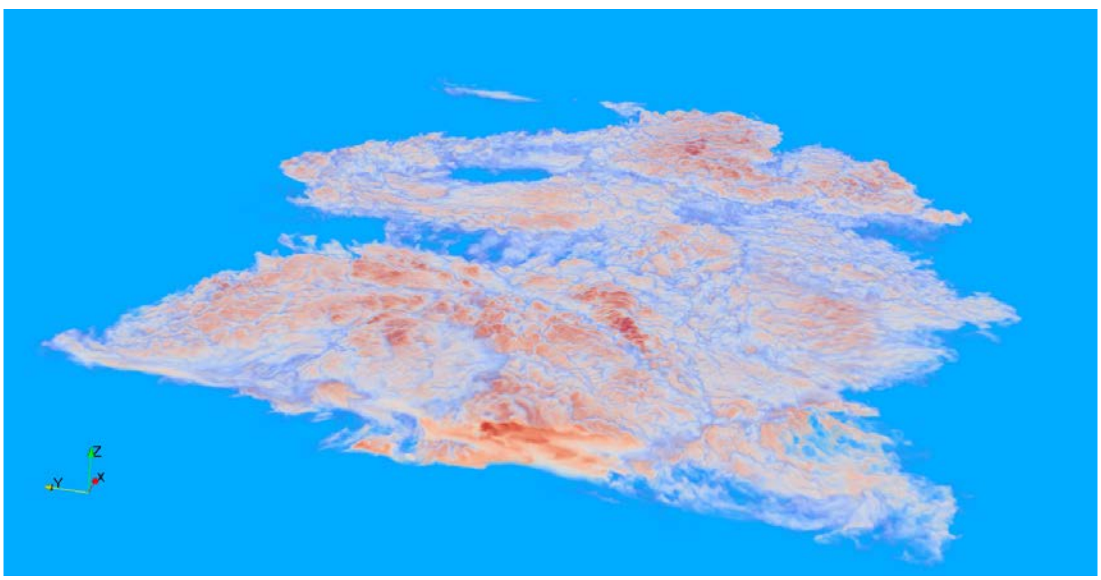

Figure C.11 Three-dimensional snapshot of ice concentration in a one-hour-old contrail. Atmospheric turbulence tends to shear the contrail whereas the radiative transfer creates additional vertical circulation inside the cloud, leading to the formation of packets of large ice concentration (red areas).

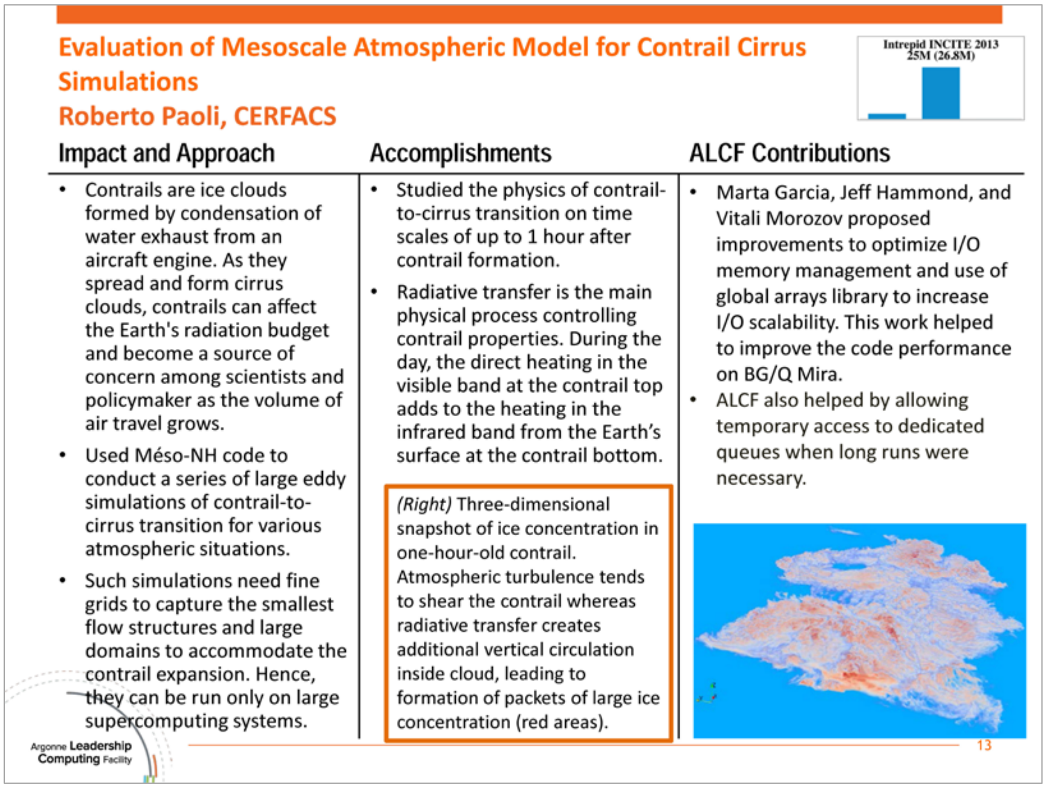

Petascale Simulations of Inhomogeneous Alfvén Turbulence in the Solar Wind

Jean C. Perez, University of New Hampshire

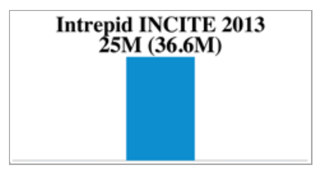

The origin of the solar wind and heating of the solar corona are two of the most important unsolved problems in heliophysics. This project aims to advance understanding of Alfvén wave turbulence in the inner heliosphere and the role that it plays in the heating and acceleration of the solar wind.

In the solar wind acceleration region, background solar wind properties—including density, solar wind outflow velocity, and magnetic field strength—are highly inhomogeneous. These inhomogeneities lead to wave reflections that trigger a turbulent cascade, called “reflection-driven” turbulence (Figure C.12).

Using the inhomogeneous magnetohydrodynamics (MHD) code *REFLECT* and the petascale power of Mira, PI and code developer Jean Perez is able to resolve both the inhomogeneity of background profiles and the small-scale structures that result from the turbulent dynamics. *REFLECT* simulations have shown that reflection-driven turbulence can be compared directly with observations from the Helios space probe. While there has been significant debate on the observed spectrum’s origins, these simulations show evidence that reflection-driven turbulence can lead to a flat spectrum by a local mechanism.

IMPACT: An unprecedented simulation on an 8.6 billion-point mesh has given the first observation of a 1/*k* spectrum of Alfvén waves in a direct numerical simulation of MHD turbulence, similar to spectrum measurements made *in situ* by the Helios mission.

ALCF Contributions: After some previous ALCF work on optimizing the code, Perez's projects have been quite self-sufficient (Tim Williams). Staff has worked with the project on visualization, helping them to use Paraview on Tukey and advising on filters to apply to bring out desired results in visualization (Joe Insley).

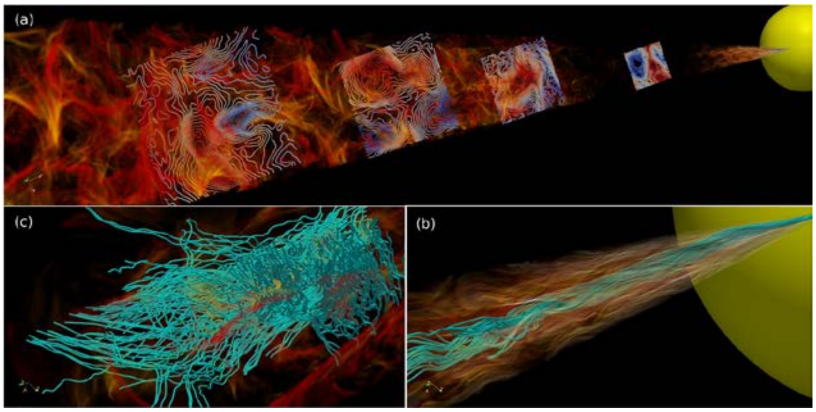

Figure C.12 (a) Large-scale Alfvén waves, shown as contour lines at selected radii, transport energy from the solar surface into the solar atmosphere along a coronal hole. The waves become turbulent and heat the ambient plasma in small-scale turbulent structures, shown in the orange gray scale. (b) Magnetic field line rendering near the sun. (c) Turbulent magnetic field lines above 60 solar radii, covering the lowest orbit of the Helios mission.

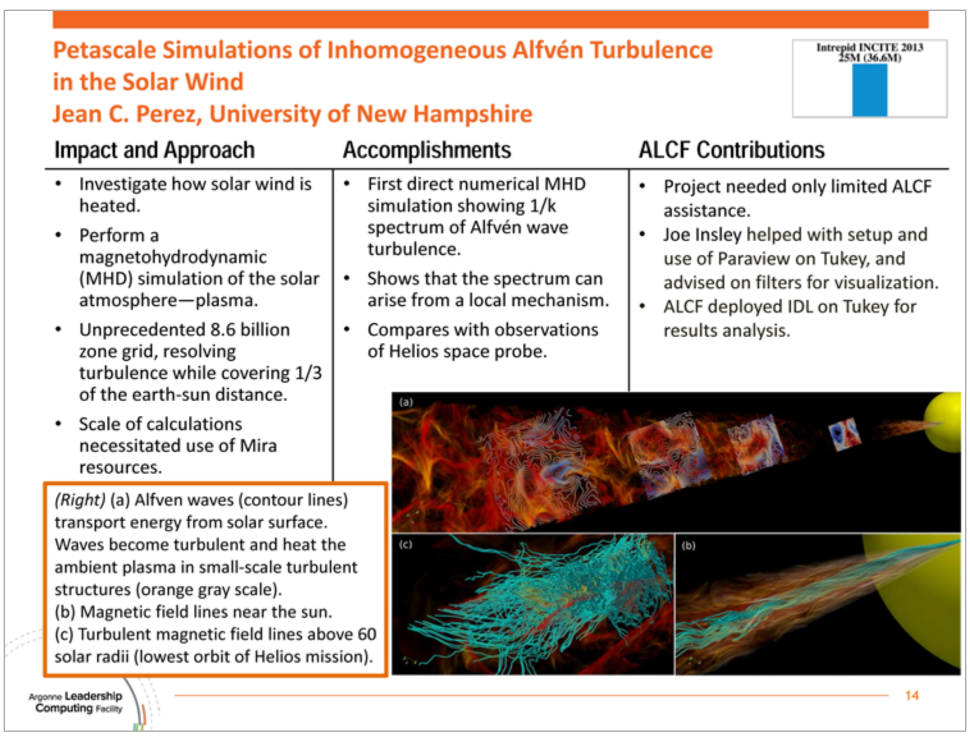

Petascale Simulations of Stress Corrosion Cracking

Priya Vashishta, University of Southern California

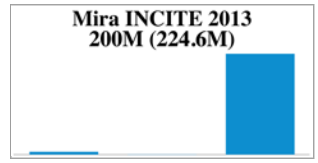

Producing hydrogen from aluminum-water reactions has potential for clean energy applications, including on-board fuel production for hydrogen-powered vehicles. The process has been limited by its low reaction rate and poor yields, but recent experiments have suggested that alloying aluminum with lithium could help overcome these issues. A research team from the University of Southern California (USC) used the supercomputing capabilities of Mira to shed light on the atomistic mechanisms of this promising approach.

In the final year of this three-year INCITE allocation, USC researchers performed quantum molecular dynamics (QMD) simulations to study the atomistic mechanisms of the reaction of a lithium-aluminum (Li-Al) alloy particle with water for rapid, high-yield hydrogen production. The team previously used the allocation to investigate fundamental mechanisms of stress corrosion cracking, nanoindentation of amorphous silica in the presence of water, and impurity segregation-induced embrittlement of metallic alloys.

For the hydrogen production study, researchers performed a 16,611-atom QMD simulation that revealed that alloying aluminum particles with lithium results in orders-of-magnitude acceleration of the reaction rate as well as higher yield. Their simulations also helped identify the key mechanisms underlying this rapid, high-yield reaction. The mechanisms included having efficient charge pathways in the aluminum atoms that collectively act as a "superanion" and the dissolution of lithium atoms into water to produce a corrosive basic solution that prevents the formation of a reaction-stopping passive oxide layer on the particle surface.

Scalability is another major obstacle to using aluminum particles for hydrogen production. Namely, the high reactivity of aluminum particles cannot be sustained for larger particles that are commercially mass-produced. To investigate the scalability of the process, the researchers compared simulations involving $Li_{135}Al_{135}$ and $Li_{441}Al_{441}$ in water (a total of 4,836 and 16,611 atoms, respectively) (Figure C.13). They found these surfaces were equally reactive regardless of the surface curvature, indicating that alloy design has the potential to scale up to industrially relevant particle sizes.

IMPACT: These simulations provide a microscopic understanding of how hydrogen is produced from water reacting with lithium-aluminum alloy particles while validating its potential for industrial scalability. Results from this INCITE project can inform design principles for rapid, high-yield hydrogen production for clean energy applications such as in hydrogen-powered vehicles.

ALCF Contributions: Based on work started at Mira Performance Boot Camp, the FLOP rate was doubled with the help of staff (Nichols Romero) and IBM (Robert Walkup). Early on in the project, an MPI hang was encountered and resolved with ALCF assistance (Vitali Morozov).

Publication: K. Shimamura, F. Shimojo, R. K. Kalia, A. Nakano, K. Nomura, and P. Vashishta, Nano Lett. **14(7)**, 4090 (2014).

Presentation: "Quantum Molecular Dynamics Simulations of Hydrogen-on-Demand on Blue Gene/Q," IBM HPC Systems Scientific Computing User Group Meeting, ScicomP 2014, May 28, 2014, Chicago, Illinois.

Acknowledgment: Dr. James Davenport, Program Manager of Theoretical Condensed Matter Physics, Division of Materials Science and Engineering, BES (Grant Number DE-FG02-04ER46130) supported this research. Computing resources for this DOE-supported research were provided by the INCITE program.

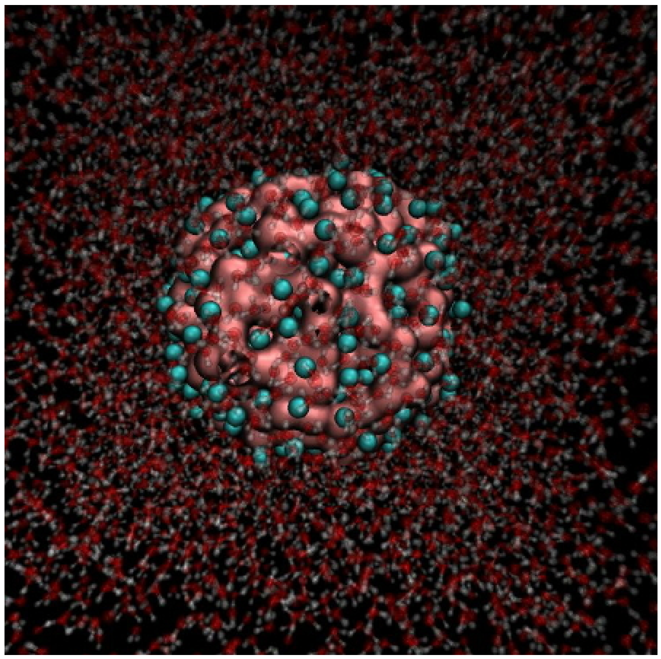

Figure C.13 Hydrogen production from water using a $Li_{441}Al_{441}$ particle. The valance electron density represented by the silver isosurface is centered around Al atoms, whereas some of the Li atoms represented by red spheres are dissolved in water. Produced hydrogen molecules are represented by green surfaces. For clarity of presentation, water molecules are not shown.

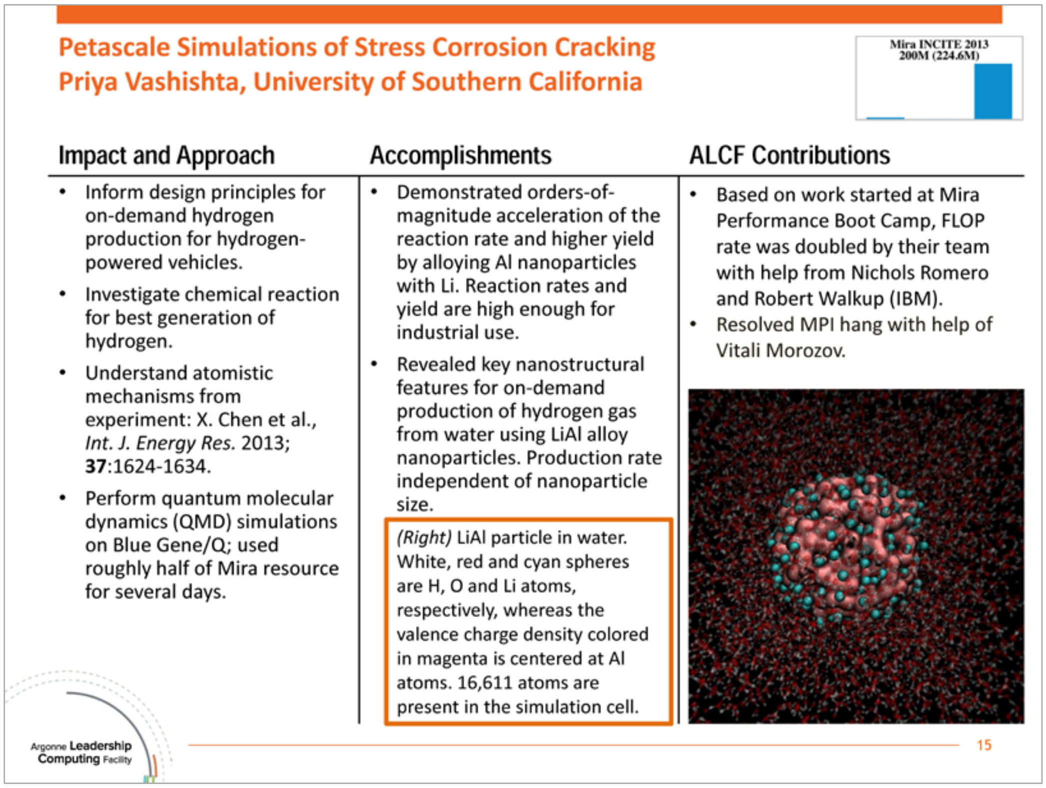

Correlated Electrons in Photoactive and Superconducting Materials

Lucas K. Wagner, University of Illinois at Urbana-Champaign

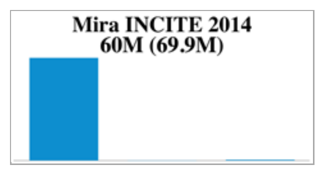

Superconducting materials are currently used to generate high magnetic fields for applications ranging from medical magnetic resonance imaging to the Large Hadron Collider at CERN. But their need for very low operating temperatures prevents superconducting materials from being used everywhere. Although the existence of high-temperature superconductors has been known for years, the microscopic origins of their properties have remained largely elusive.

There are two categories of superconductors:

- The conventional type which are well understood with properties arising from the electron-phonon coupling interaction.
- The unconventional type which are not well understood with properties attributed to electron-electron correlations.

The first unconventional superconductors were discovered about 30 years ago. Collectively called the cuprates (containing copper, oxygen, and other species of atoms), they have unusually high superconducting transition temperatures (up to ~150K). Interest in the cuprates, once high, has waned. Although some regarded cuprates as unique, another family of unconventional superconductors was discovered in 2008—this one based on iron instead of copper—and has re-awakened interest in high-temperature superconductors.

Lucas Wagner's team uses quantum Monte Carlo (QMC) methods to understand the microscopic mechanism that gives rise to high-temperature superconductivity. The advent of massively parallel supercomputers has brought the QMC method to the forefront and permitted calculations previously accessible only by computationally cheaper mean-field methods, such as density functional theory (DFT). Wagner and his team use QWALK, a QMC code they developed to perform diffusion Monte Carlo calculations on iron selenide (FeSe), a known high-temperature superconductor.

In experiments, FeSe starts in a magnetic state and becomes superconducting as pressure is applied; with more pressure, superconductivity first vanishes but then re-emerges with additional pressure. The difference between the energies in the checkerboard spin ordering versus the other two is a measure of the electron-spin interaction, which is increasing under pressure (Figures C.14 and C.15). This research, it is hoped, will lead to a better understanding of how superconductivity emerges in these materials. An INCITE award allowed Wagner's team to study the magnetic states of FeSe at varying pressure. Their calculations add to the existing body of evidence supporting the notion that high-temperature superconductivity is magnetic in origin.

IMPACT: Calculations on iron selenide (FeSe) represent the first application of this level of theory (QMC) to these materials, adding to evidence of the magnetic origin of superconductivity.

ALCF Contributions: Staff helped with project set-up and solved technical issues (Anouar Benali).

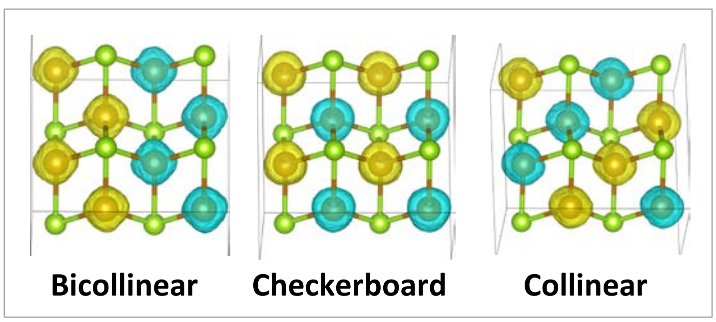

Figure C.14 Magnetic states of FeSe. Yellow (green) translucent spheres indicate a net spin of up (down) on the iron atoms. Small green spheres are selenium atoms.

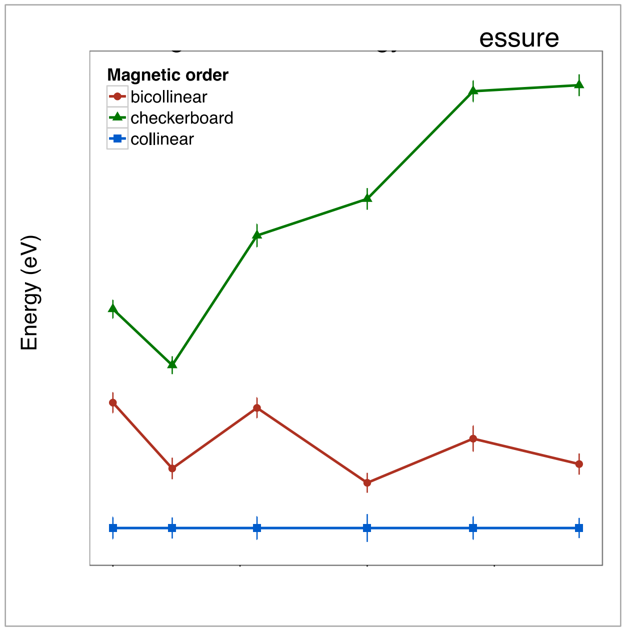

Figure C.15 Magnetic states energy vs. pressure in FeSe.

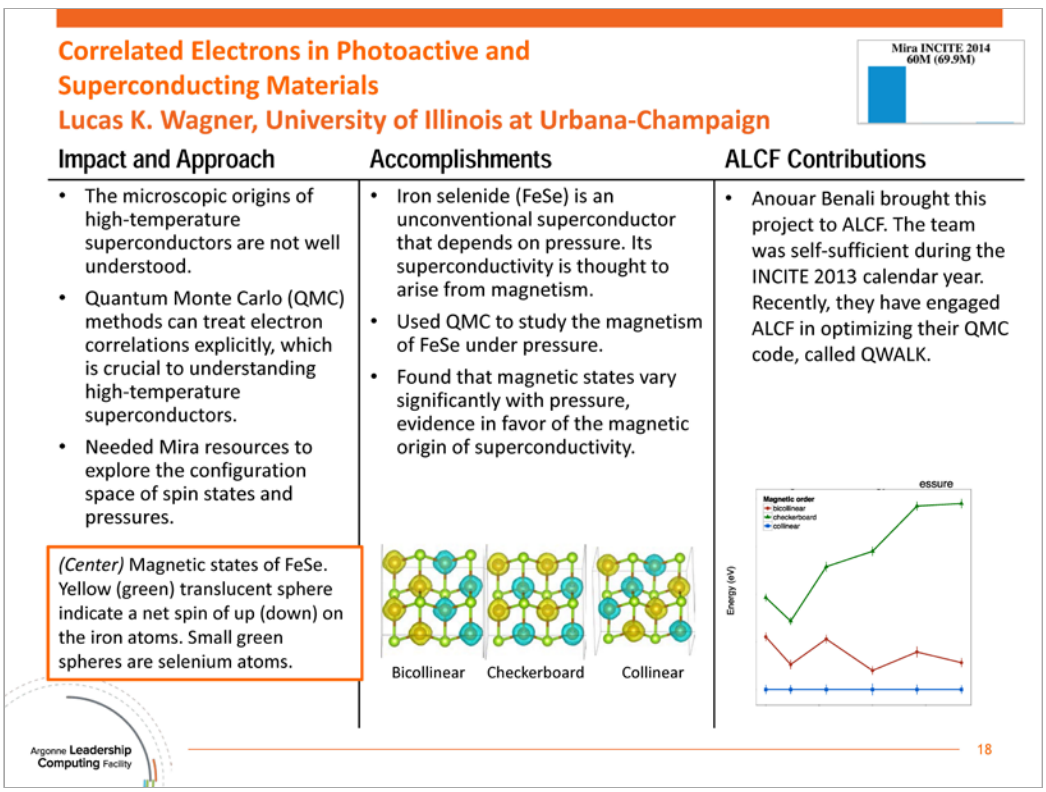

Understanding Helium Plasma-Mediated Tungsten Surface Response That Controls Plasma-Facing Component Performance and Lifetime

Brian D. Wirth, University of Tennessee, Knoxville

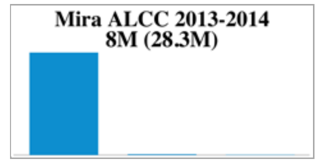

Providing energy from terrestrial nuclear fusion reactions is among the great scientific and engineering challenges of our time. One possible route to fusion energy is magnetically confined plasma reactors. This project studies the reaction of helium plasma on specific surface structures of tungsten, the proposed material for large tokamak divertors, the device within a tokamak that allows the removal of fusion reaction waste by-products.

Tungsten is currently the material of choice for the divertors in prototype magnetic-confinement nuclear fusion reactors, such as the super tokamak International Thermonuclear Experimental Reactor (ITER). Research used the LAMMPS code to perform large-scale molecular dynamics (MD) simulations to develop an atomistic understanding of the helium bubble formations on tungsten (Figure C.16). LAMMPS calculations that ran on Mira provided a detailed microscopic picture of helium plasma structures on the surface of tungsten.

Simulation results have focused on implanting helium gas atoms on tungsten surfaces. The most current and significant scientific results are associated with how the trap mutation processes are modified in the presence of a free surface. There is a marked difference in the helium retained between different orientations/layers of surfaces. This suggests the existence of a mechanism that prevents helium from leaving some specific orientations, which does not exist for other surface orientations and which is quite strong.

MD simulations have shown the presence of highly inhomogeneous helium layering along different surface layers, and such clustering may substantially influence the sputtering and erosion behavior of tungsten.

IMPACT: The results of this work may indicate that some surface orientations are preferable to others for the exposed surfaces of plasma-facing components and, thus, provide a greater physical understanding and predictive modeling capability for materials design of the ITER divertor to survive the incredibly extreme conditions of a fusion power plant.

ALCF Contributions: Staff guided users in tuning OpenMP run-time parameters, improving throughput on Mira queues and obtaining strong-scaling benchmark data (Wei Jiang).

Publication: L. Hu, K. D. Hammond, B. D. Wirth, and D. Maroudas, “Interactions of mobile helium clusters with surfaces and grain boundaries of plasma-exposed tungsten,” Journal of Applied Physics **115**, 173512 (2014).

Presentations:

- B. D. Wirth, T. Faney, K. Hammond, N. Juslin, F. Sefta, D. Xu, and the SciDAC-PSI team, "Modeling Plasma Surface Interactions in Tungsten through High Performance Computing," Oxford University Workshop on Modelling Radiation Damage and Its Effects on Materials, September 26, 2013, Oxford, England.
- B. D. Wirth, "Modeling Plasma Surface Interactions in Tungsten through High Performance Computing," 11th International Symposium on Fusion Nuclear Technology, September 16, 2013, Barcelona, Spain.
- B. D. Wirth, "Plasma Surface Interactions Involving He," International Workshop in Plasma Materials Interaction Facilities, Oak Ridge National Laboratory, September 11, 2013, Oak Ridge, Tennessee.

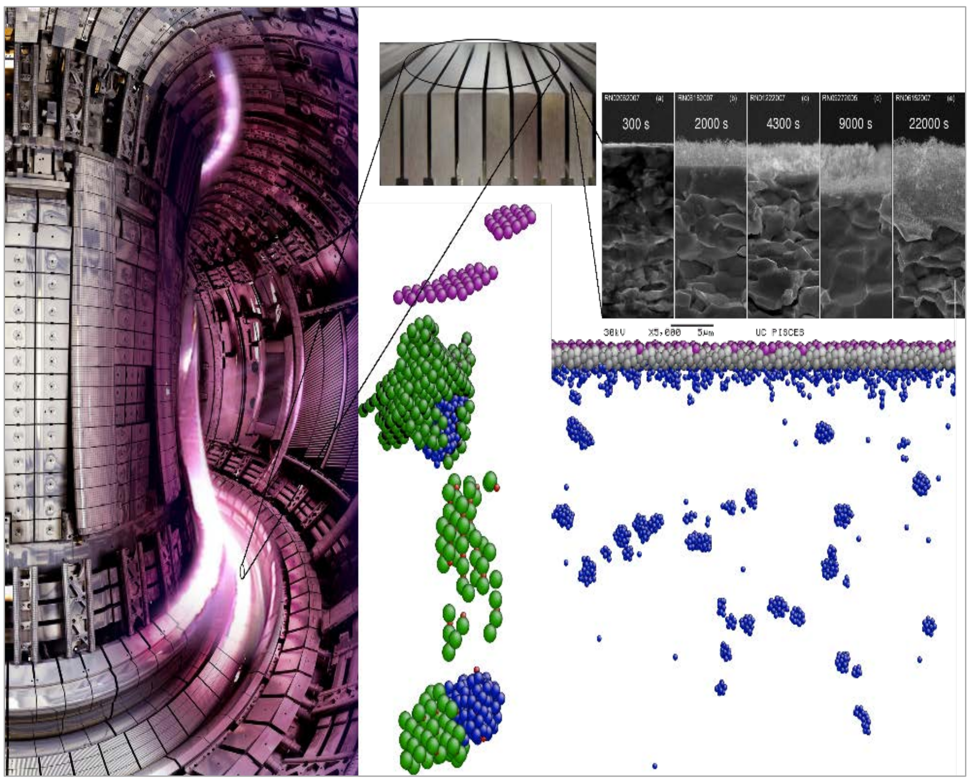

Figure C.16 Composite image showing internal surfaces of the tokamak fusion reactor and the castellated tiles of the divertor region, which faces severe plasma surface interactions from high fluxes of low-energy helium and hydrogen plasmas. *Top right inset:* Tungsten surface morphology modification observed in linear plasma device exposure to low-energy He plasma. It is now recognized through atomistic, molecular dynamics simulations that growing helium atom clusters (blue spheres) punch dislocation loops (green spheres) during bubble formation and growth that produce adatom islands (purple spheres) that drive the initial surface roughening, leading to nano-scale fuzz formation.

Understanding Helium Plasma-Mediated Tungsten Surface Response That Controls Plasma-Facing Component Performance and Lifetime
Brian Wirth, University of Tennessee

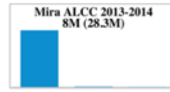

Impact and Approach	Accomplishments	ALCF Contributions
• Correlating surface structure of divertor material to sub-surface gas clustering and bubble formation in tungsten exposed to He plasma. • Input to large-scale simulations for divertor materials on length and time scales relevant to fusion reactors.	• Improved knowledge of plasma surface interactions and the materials engineering design of component systems under extreme conditions of a fusion power plant. • Secured three invited talks. • *Journal of Applied Physics* **115** (2014) 17351.	• Wei Jiang guided users in tuning OpenMP run-time parameter, improving throughput on Mira queues, and obtaining strong-scaling.

(Right) Composite image showing internal surfaces of tokamak fusion reactor and castellated tiles of the divertor region, which faces severe plasma surface interactions from high fluxes of low-energy helium and hydrogen plasmas. *Top right inset:* Tungsten surface morphology modiciation observed in linear plasma device exposure to low-energy He plasma. It is now recognized through atomistic, molecular dynamics simulations that growing helium atom clusters (blue spheres) punch dislocation loops (green spheres) during bubble formation and growth that produce adatom islands (purple spheres) which drive the initial surface roughening leading to nano-scale fuzz formation.

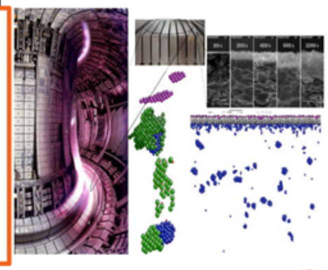

Argonne Leadership Computing Facility

17

This page intentionally left blank.

Appendix D – The Impact of Checklists

Ira Goldberg, Richard Coffey, and Jini Ramprakash
Leadership Computing Facility, Argonne National Laboratory

Executive Summary
A checklist is a framework that individuals or teams use to confirm completion of the most important steps in a process. It has the greatest impact when employed in conjunction with complex repeatable tasks. The ALCF uses checklists for responding to helpdesk tickets, in project management, in event planning, to improve INCITE project efficacy, and to analyze cryptocard failures. All are employed as a means to improve quality and to reduce defects in a process.

What Is a Checklist?
A checklist is an informational job aid. It reduces failure by compensating for potential limits of human memory and attention. It also helps to ensure that tasks are executed consistently and completely. A checklist provides a framework for tracking desirable job performance traits as well as failures (defects) in order to achieve goals. It can be a tangible mechanism to train a team or to elicit recommendations from a group.

A checklist can range from the very basic, like a shopping list, to the highly complex, like all of the ordered steps for constructing a skyscraper. In checklist taxonomy, a "task-focused" checklist specifies tasks that must be completed to successfully complete a project (like the Boeing example cited below), and a "communication" checklist is designed to assure information exchange and problem resolution among different specialists in cases where solutions are not yet known. The latter process is designed to produce more complete understanding and improved problem resolution as glitches arise.

Why Use Checklists?
Simply put, appropriate use of checklists reduces errors/defects and provides better results than without them. Time spent on implementing a checklist is an investment. By assuring a standard procedure, the checklist reduces the potential of missing a crucial step in a process, and can thereby save time, money, and even lives.

Organizations can benefit from process improvement by including the use of checklists. A high positive impact may occur when used in conjunction with what we might call "Pareto Principle triage," where groups first identify the few types of "defects" that create the largest number of significant problems, and then focus on ways to lessen the chance of those "defects" from occurring. For an organization like ALCF, the method can be applied to most tasks. Examples include administrative interactions with DOE, help desk tickets, machine maintenance, catalyst support for scientists, and purchasing.

How Checklists Are Used

Checklists are used to monitor time-saving practices, value-added activities, and defects. The knowledge gained is important to attain systemic improvement. As any process becomes more complex, the probability of human error increases. Checklists are used to break complex tasks into simpler, more manageable, and less error-prone components. Some examples are shown below.

Examples

In 1935, an experienced pilot named Major Ployer P. Hill conducted a test flight of the Boeing Model 299 (the "flying fortress"), during which the plane stalled, crashed, and exploded. Why did it happen? The new airplane was substantially more complicated than Hill's previous aircraft and one critical step (releasing a new locking mechanism) was missed. After analyzing the incident thoroughly, a group of test pilots came up with a simple checklist process covering takeoff, flight, landing, and taxiing that was short enough to fit on a small card. The principles underlying the checklist are still used by pilots today.

In 2001, Peter Pronovost, a physician at Johns Hopkins Hospital, documented a specific process to avoid infections when putting a central line into a patient. The process was converted into a checklist with a 15-month estimated impact of 43 fewer infections, $2 million in lowered costs, and 8 lives saved. In 2004, the program was expanded to hospitals in Michigan. Within the first 3 months, infection rates dropped by 66% and over the first 18 months, savings were estimated at $175 million. The success of this approach was due in large part to management participation. Staff teams pointed out a lack of infection-reducing soap and sterile drapes. Without management involvement, the supplies would not have become available and progress would not have been made.

Conclusion

Checklists can be useful tools through which individuals and groups can focus on process improvement and improve outcomes. In order to be effective, they must be short, easily understood, and straightforward to implement (i.e., simple to use).

References

The Checklist Manifesto, Atul Gawande, Metropolitan Books, 2009.

Quality is Personal: A Foundation for Total Quality Management, Harry Roberts and Bernard F. Sergesketter, The Free Press, 1993.

"Using Personal Checklists to Facilitate Total Quality Management," Harry Roberts (University of Chicago), http://www.chicagobooth.edu/faculty/selected-paper-series.