ARGONNE **LEADERSHIP COMPUTING** FACILITY

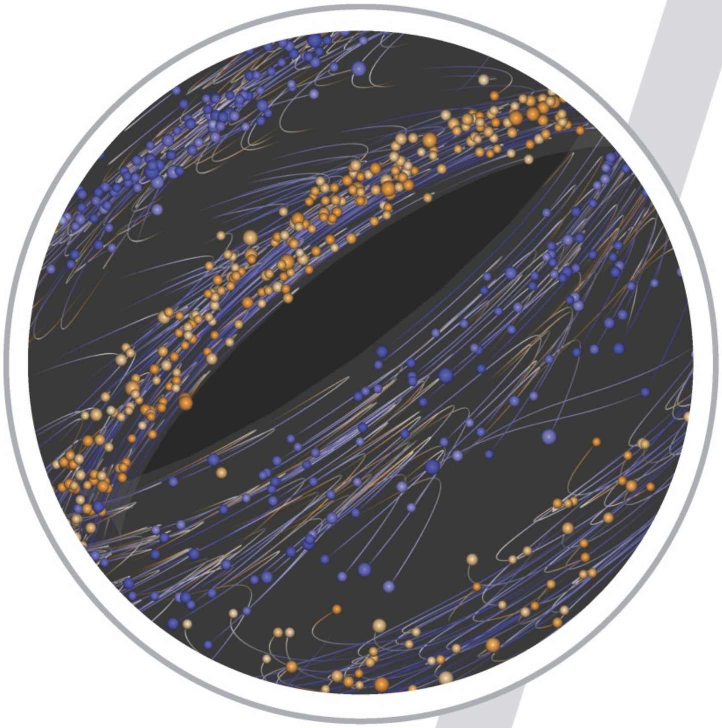

OPERATIONAL ASSESSMENT REPORT

MARCH 2014

On the cover

Particle visualization of a global gyrokinetic particle-in-cell simulation of microturbulence in a tokamak fusion device.

2013 INCITE Project: Kinetic Simulations of Fusion Energy Dynamics at the Extreme Scale
PI: William Tang, Princeton Plasma Physics Laboratory/Princeton University
Image Credit: Chad Jones and Kwan-Liu Ma, University of California, Davis; Stephane Ethier, Princeton Plasma Physics Laboratory

Contents

Executive Summary 1-1

Section 1. User Support Results 1-1

ALCF Response 1-1
2013 ALCF User Survey Results 1-1
Survey Approach 1-2
Likert Scale and Numeric Mapping 1-2
1.1 User Support Metrics 1-3
Day of the Week for Scheduled Maintenance 1-3
1.2 Problem Resolution Metrics 1-4
1.3 User Support, Communication, and Outreach 1-4
1.3.1 User Support 1-4
Phone and Email Support 1-4
Continuous Improvement Efforts in User Support 1-5
Project and User Management Software Efforts 1-6
Configuring Support Software for Mira 1-6
1.3.2 Communications 1-6
Communications through Mailing Lists and Social Media 1-6
Custom Communications and Reports 1-7
Communicating Decommissioning of Intrepid 1-8
Web Documentation Efforts for Mira 1-8
Other Publications 1-9
1.3.3 Outreach Efforts 1-9
User Advisory Council 1-9
Industry Outreach 1-9
2013 Workshops and Webinars 1-9
Conclusion 1-12

Section 2. Business Results 2-1

ALCF Response 2-1
ALCF Resources 2-2
Mira 2-2
2.1 Resource Availability 2-2
2.1.1 Scheduled and 2.1.2 Overall Availability 2-2
Explanation of Significant Availability Losses 2-3
2.1.3 System Mean Time To Interrupt (MTTI) and 2.1.4 System Mean Time To Failure (MTTF) 2-4
ALCF MTTI and MTTF Summary 2-4
2.2 Resource Utilization 2-5
Total System Utilization 2-5
2.3 Capability Utilization 2-7
Intrepid 2-10

Contents (Cont.)

2.1 Resource Availability 2-10
2.1.1 Scheduled and 2.1.2 Overall Availability 2-10
Explanation of Significant Availability Losses on Intrepid 2-10
2.1.3 System Mean Time To Interrupt (MTTI) and 2.1.4 System Mean Time To Failure (MTTF) 2-12
ALCF MTTI and MTTF Summary 2-12
2.2 Resource Utilization 2-13
Total System Utilization 2-13
2.3 Capability Utilization 2-15
ALCF Capability Utilization 2-15
2.4 Financial and Human Resources 2-18
Conclusion 2-18

Section 3. Strategic Results 3-1

ALCF Response 3-1
Community Involvement and Outreach 3-1
3.1 Science Output 3-2
3.2 Scientific Accomplishments 3-2
Kinetic Simulations of Fusion Energy Dynamics 3-3
Thermodynamics of Binding Biomass to Cellulases for Renewable Fuel 3-5
Explosion Research for Safer Building Designs 3-7
Designing Controlled Drug Delivery Systems 3-9
Macromolecular Folding and Aggregation 3-11
Usage of the INCITE and ALCC Hours 3-13
3.3 Allocation of Facility Director's Reserve Computer Time 3-16
Conclusion 3-19

Section 4. Innovation 4-1

ALCF Response 4-1
4.1 Vendor Collaborations to Improve Application Performance 4-1
4.1.1 IBM: Automatic Application Performance Collection 4-1
4.1.2 Pathscale: OpenMP Support in Blue Gene Clang Compilers 4-1
4.2 Operations Improvements 4-2
4.2.1 Machine Time Overlay Graphic 4-2
4.2.2 Continuous Improvement of Operational Assessment Process 4-2
4.2.3 Scheduling 4-3
4.2.3.1 Changing Block Geometries for Machine Packing 4-4
4.2.3.2 Ensemble and Subblock Jobs on Mira 4-4
4.2.4 Networking Improvements 4-4
E600i 10GbE Bandwidth 4-4
DWDM 4-4

Contents (Cont.)

4.3 Business Intelligence Improvements 4-5
4.3.1 Improved Data Management for Facility Data 4-5
4.3.2 Reduced Fragility of Facility Reporting 4-5
4.4 Facility Improvements 4-6
4.5 User Services and Outreach Improvements 4-6
4.5.1 Using Checklist to Communicate with Users 4-6
4.5.2 Using Small-Scale Videoconferences to Replace a Getting Started Workshop 4-7
4.6 Administration Improvements 4-8
Participation in DOE HPC Financing Improvement Team 4-8
Conclusion 4-10

Section 5. Risk Management 5-1

ALCF Response 5-1
5.1 ALCF Risk Management 5-1
Completion of the ALCF-2 Project 5-2
Initiation of the ALCF-3 Project 5-2
ALCF 2013 Risk Management Plan 5-2
Conversion from 3x3 Risk Matrix to a 5x5 Risk Matrix 5-2
Risk Management in Day-to-Day Operations 5-3
Major Risks Tracked for the Review Year 5-4
Risks Encountered in the Review Year and Their Mitigations 5-6
1. Funding/Budget Uncertainties 5-6
Description 5-6
Evaluation 5-6
Management 5-7
2. Staffing Challenges and 3. Staff Retention 5-7
Description 5-7
Evaluation 5-7
Management 5-7
4. Intrepid System Stability Issues 5-8
Description 5-8
Evaluation 5-8
Management 5-8
5. Facility Power Interruptions 5-9
Description 5-9
Evaluation 5-9
Management 5-9
Retired Risks 5-9
New and Recharacterized Risks Since the Last Review 5-10
Projected Major Operating Risks for the Next Year 5-14
Conclusion 5-15

Contents (Cont.)

Section 6. Safety 6-1

ALCF Response 6-1

Section 7. Cyber Security 7-1

ALCF Response 7-1

Section 8: Summary of the Proposed Metric Values for Future OARs 8-1

ALCF Response 8-1
8.1 Overview 8-1
8.2 ALCF 2014 OA Performance Metrics 8-2
8.3 ALCF Proposed 2015 OA Performance Metrics 8-2
8.4 ALCF Reportable Only Metrics (No Targets) 8-3
Conclusion 8-3

Appendix A – Calculations A-1

Scheduled Availability Calculation Details A-1
Overall Availability Calculation Details A-1
ALCF Calculations A-1
MTTI Calculation Details A-2
MTTF Calculation Details A-3
ALCF MTTI/MTTF Calculations A-3

Appendix B – ALCF Director's Discretionary Projects B-1

Intrepid B-1
Mira DD (No ESP) B-8
Mira ESP B-15

Appendix C – Strategic Results Slides C-1

Figures

2.1 Mira Weekly Availability for CY 2013 2-3

2.2 System Utilization over Time by Program 2-6

2.3 Mira Capability Usage over Time 2-9

2.4 Mira Job Usage by Size 2-9

2.5 Intrepid Weekly Availability for CY 2013 2-11

2.6 System Utilization over Time by Program 2-14

2.7 Queue Assignment Based on Job Size and Run Time 2-16

Figures (Cont.)

2.8 Cumulative INCITE Capability Usage during CY 2013 2-17

2.9 Intrepid Jobs Binned by Size 2-17

3.1 Particle visualization of a global gyrokinetic particle-in-cell simulation of microturbulence in a tokamak fusion device from the Princeton Plasma Physics Laboratory project. 3-3

3.2 The absolute binding free energy is defined as the free energy change between a polysaccharide chain (of n monosaccharide units, where n is the chain length required to saturate the GH binding sites) and the enzyme-substrate complex in the catalytically active complex. The team hypothesizes that this quantity, ΔGb°, is directly correlated to processivity and suggests morphology dependence of cellulose attack by enzymes. 3-6

3.3 The flame position is displayed at two instants: before it reaches the second obstacle and after it has interacted with the obstacles. At the final instant, the flame is propagating at speeds of the order of 250 m/s. 3-8

3.4 Water molecules form an ordered structure around the polymer chains. Breaking of this cage acts as a precursor to the collapse. Computational modeling can allow researchers to identify ways of controlling the conformational dynamics and design efficient drug delivery systems. 3-10

3.5 Large-scale ensemble simulations targeted at identifying key effects in the formation of defect states, which can then be eliminated or exploited. This image shows an example of a misfolding. 3-12

3.6 Intrepid INCITE 2013 Allocation Usage 3-14

3.7 Mira INCITE 2013 Allocation Usage 3-15

3.8 Intrepid ALCC 2013 Allocation Usage 3-15

3.9 Mira ALCC 2013 Allocation Usage 3-16

3.10 Intrepid ALCC 2012 Allocation Usage 3-16

3.11 CY 2013 DD Allocation by Standard INCITE Science Domains 3-18

4.1 Machine Time Overlay Graphic 4-2

4.2 Weekly Availability Chart with Annotations 4-3

4.3 Example Message from 2012 INCITE Communications 4-6

4.4 Example Message from 2013 INCITE Communications 4-7

4.5 Arrangement of Video Feeds at ALCF Getting Started Workshop 4-8

4.6 Arrangement of Local Video Conference Room for Getting Started Workshop 4-8

Tables

ES.1 Summary of the Target and Actual Data for the Previous Year (2013) Metrics ES-2
1.1 All 2013 User Support Metrics and Results 1-1
1.2 2013 User Survey Results by Allocation Program 1-2
1.3 2012 and 2013 User Support Metrics 1-3
1.4 2012 and 2013 Problem Resolution Metrics 1-4
1.5 2012 and 2013 Ticket Categorization 1-5
1.6 2013 Primary Communication Channels 1-7
1.7 List of Key Changes and Outcomes Due to Checklist Efforts 1-8
1.8 Publications Designed for Print in 2013 1-9
1.9 2013 Workshops and Webinars 1-10
2.1 Summary of All Metrics Reported in the Business Results Section 2-1
2.2 Availability Results 2-3
2.3 MTTI and MTTF Results 2-5
2.4 System Utilization Results 2-5
2.5 Core-Hours Allocated and Used by Program 2-7
2.6 Capability Results 2-8
2.7 Availability Results 2-10
2.8 MTTI and MTTF Results 2-12
2.9 System Utilization Results 2-13
2.10 Core-Hours Allocated and Used by Program 2-15
2.11 Capability Results 2-15
2.12 Summary of All Metrics Reported in the Business Results Section 2-18
3.1 The total time allocated and used within the discretionary project. The total for Mira includes time for the Early Science Program after the machine went into production. 3-18
5.1 5x5 Risk Assessment Matrix 5-3
5.2 Major Risks Tracked for CY 2013 5-4
5.3 Retired Risks in CY 2013 5-10
5.4 New and Recharacterized Risks from CY 2013 5-11
5.5 Projected Operating Risks for CY 2014 5-14
8.1 Performance Metrics – 2013 Targets, 2013 Actuals, and Agreed-Upon 2014 Targets 8-2

Tables (Cont.)

8.2 Performance Metrics – Agreed-Upon 2014 Targets and Proposed 2015 Targets 8-2
8.3 ALCF Reportable Only Metrics 8-3

This page intentionally left blank.

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 Advanced Scientific Computing Research (ASCR), provides the computational science community with world-class computing capabilities, expertise, and assistance to ensure that every project achieves top performance on its resources.

On April 9 the ALCF moved its new 48-rack IBM Blue Gene/Q system, Mira, into production alongside its legacy system, Intrepid, an IBM Blue Gene/P. Also deployed at this time were two smaller Blue Gene/Q systems, Cetus and Vesta. Cetus is a single Blue Gene/Q rack that shares the same software stack and file systems as Mira and is used for tool and application porting, software testing and optimization, and systems software development. Vesta is a two-rack system used to test new versions of software before they are installed on Mira. Analysis and visualization for Mira data are performed on Tukey, a visualization cluster. Tukey shares the Mira network and parallel file system, enabling direct access to Mira-generated results.

Meanwhile, Intrepid continued to run as a production system, along with its two smaller systems: Challenger and Surveyor. On December 31, Intrepid, an IBM Blue Gene/P, and its two companion systems were retired.

The ALCF has proudly met or exceeded all metrics set for the facility. ALCF delivered 4.79 billion core-hours of compute time in 2013 between January 1 and December 31, with 2.61 billion of those core-hours being used by capability jobs. During that same period, the science done on the machines produced more than 150 publications, all in 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 1,150 DOE facility users and engaged more than 300 active projects from universities, national laboratories, and industry.

The annual Operational 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 2013 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 (2013) Metrics

Area	Metric	2013 Targets	2013 Actuals
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.6/5.0
	User Survey – **Response Rate**	25%	31.7%
	% User Problems Addressed Within Three Working Days	80%	91.9%
Business Results	Intrepid Overall Availability	90%	94.7%
	Intrepid Scheduled Availability	95%	97.0%
	Intrepid INCITE Capability Usage (20% – 100% of system)	40%	64.3%
	Mira Overall Availability	80%	90.6%
	Mira Scheduled Availability	85%	95.5%
	% of INCITE Jobs Run on 16.7% or More of Mira (131,072 – 786,432 cores)	20%	60.7%
	% of INCITE Jobs Run on 33.3% or More of Mira (262,144 – 786,432 cores)	5%	33.3%

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 2013 user survey measured satisfaction, user support, and problem resolution and will direct ALCF where improvements can be made. The sections below document ALCF processes, the effectiveness of those processes, and what improvements to those processes were implemented during CY 2013.

Table 1.1: All 2013 User Support Metrics and Results

		Value for 2012	Target for 2013	Value for 2013
Number Surveyed		920	N/A	1,150
Number of Respondents (Response Rate)		305 (33.2%)	(25%)	364 (31.7%)
Overall Satisfaction	Mean	4.4	3.5	4.5
	Variance	0.5	N/A	0.5
	Std Dev	0.7	N/A	0.7
Problem Resolution	Mean	4.5	3.5	4.6
	Variance	0.4	N/A	0.5
	Std Dev	0.6	N/A	0.7
User Support	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.5
	Std Dev	0.7	N/A	0.7
% User Problems Addressed Within Three Working Days		94.4%	80.0%	91.9%

2013 ALCF User Survey Results

In 2013, 1,150 individual ALCF users met the DOE user definition[1] and were invited to complete a user survey. Of those users, 364 responded. ALCF surpassed all targets for the survey metrics. ALCF measured no statistical change from last year from responses for all users in Overall Satisfaction, Problem Resolution, and User Support.

In Table 1.2, the results are broken down by allocation program. Director's Discretionary users and INCITE users each reported a higher average Overall Satisfaction than ALCC users. This is

[1] The definition of an official user was made consistent across DOE user facilities in 2013 by the program office. The timeframe for an official user is now reported as a fiscal as opposed to calendar year. In addition, staff members and vendors are considered users if they are members of a reviewed project.

noteworthy given the comparison between programs and response rate for ALCC, however considering only 34 ALCC users responded to the survey, the sample is too small to be statistically relevant. User Support and Problem Resolution metrics were similar across all allocation programs.

Table 1.2: 2013 User Survey Results by Allocation Program

2013 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		327	84	411	739	1,150
Number of Respondents		122	34	156	208	364
Response Rate		37.3%	40.5%	38.0%	28.1%	31.7%
Overall Satisfaction	Mean	4.5	4.1	4.4	4.6	4.5
	Variance	0.5	1.1	0.7	0.5	0.5
	Std Dev	0.7	1.1	0.8	0.7	0.7
User Support	Mean	4.5	4.4	4.5	4.6	4.5
	Variance	0.5	0.7	0.6	0.4	0.5
	Std Dev	0.7	0.8	0.7	0.6	0.7
Problem Resolution	Mean	4.5	4.6	4.5	4.6	4.6
	Variance	0.5	0.5	0.6	0.4	0.5
	Std Dev	0.7	0.7	0.8	0.6	0.7
All Questions	Mean	4.5	4.4	4.5	4.6	4.6
	Variance	0.5	0.7	0.6	0.4	0.5
	Std Dev	0.7	0.8	0.8	0.6	0.7

Survey Approach

ALCF hired survey experts from Cvent, a web survey hosting and consulting company, to manage the 2013 survey. The team incorporated lessons learned from last year's survey and internal feedback from various ALCF teams, ALCF leadership, the User Advisory Council, and ASCR. The result was a streamlined survey, improved questions, and a representative user response to the survey.

Likert Scale and Numeric Mapping

All but one Likert Scale question in the ALCF user survey used a six choice scale. This has been a standard for surveys because: 1) it is easy for the users to quickly place the response to a question within a range of options, 2) it can be mapped to a numeric scale and given a certain sample size, be used with a normal distribution. This allowed use of off-the-shelf functions in Excel to determine variance and standard deviation. ALCF maps the Likert Scale in this way:

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

The overall satisfaction question used a 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 2013 User Survey that used the 6-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 2013, ALCF exceeded the Overall Satisfaction and User Support targets.

Table 1.3: 2012 and 2013 User Support Metrics

Survey Area	2012 Target	2012 Actual	2013 Target	2013 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

Day of the Week for Scheduled Maintenance

Last year ALCF asked its users to rank which day of the week was best for scheduled maintenance using the following explanatory sentence: *Given your weekly schedule, if Monday is the best day and Wednesday is the worst for a shutdown, you might select "best" for Monday* *and "worst" for Wednesday and choose neutral for all of the other days.* After reviewing the survey results and the syntax of the question, the survey team decided to use the same coding but remove the explanatory sentence.

The following code was used to generate numerical values for the user responses to this question: Best was coded as +2; Good as +1; Neutral = 0; Bad = -1; and Worst = -2. Summing the resulting re-coded responses for the 364 respondents who answered this question produced the following results:

2012 Results	Monday	Tuesday	Wednesday	Thursday	Friday
Sum of Results	106	-10	-42	-51	-31

2013 Results	Monday	Tuesday	Wednesday	Thursday	Friday
Sum of Results	177	11	-13	-67	-126

Given these two years of responses, ALCF will continue to perform preventive maintenance on Mondays.

1.2 Problem Resolution Metrics

ALCF exceeded the target set for the percentage of problem tickets addressed within three working days. 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 a solution or has found one. The ALCF exceeded the Problem Resolution targets.

Table 1.4: 2012 and 2013 Problem Resolution Metrics

Survey Area	Target for 2012	Value for 2012	Target for 2013	Value for 2013
% User Problems Addressed Within Three Working Days	80.0%	94.4%	80.0%	91.9%
Average of Problem Resolution Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.6/5.0

1.3 User Support, Communication, and Outreach

1.3.1 User Support

Phone and E-mail Support

ALCF answered 5,357 support tickets in 2013, a 41% increase in ticket volume over 2012. The largest number of these tickets involved access and accounts. However, between 2012 and 2013, the percentage of access and accounts tickets dropped from 70% to 56%. The allocations category saw the largest increase in tickets overall, from 39 in 2012 to 581 in 2013.

Table 1.5: 2012 and 2013 Ticket Categorization

Categories	CY 2012	CY 2013
Access	986/(26%)	1,080/(20%)
Accounts	1,679/(44%)	1,934/(36%)
Allocations	39/(1%)	581/(11%)
Applications Software	177/(5%)	306/(6%)
Automated E-mail Responses [2]	239/(6%)	397/(7%)
Data Transfer	24/(1%)	41/(1%)
I/O & Storage	167/(4%)	220/(4%)
Miscellaneous	202/(5%)	190/(4%)
Quota Management	40/(1%)	38/(1%)
System	208/(5%)	559/(10%)
Visualization	41/(1%)	11/(0%)
Total Tickets	**3,802**	**5,357** [3]

This spike in allocation tickets occurred because of a major change to how allocations are now managed. In 2012 a member of the Catalyst team handled the majority of allocations. In 2013 the ALCF User Services team absorbed this task and set up a new process whereby all new, extensions, and renewals also generate a helpdesk ticket, hence the significant increase.

System tickets had the second largest increase after allocations. In 2013, ALCF was supporting both Intrepid and Mira. Looking deeper at the ticketing data, it was found that 26% of these tickets were Intrepid-related and 58% were Mira-related (15% of the tickets were related to visualization clusters and 1% were miscellaneous.) It is possible the significant rise in this category can be attributed to supporting two systems; if so, there will be a decrease in the number of tickets over the next year with the decommissioning of Intrepid. Over 57% of the total increase in tickets was the result of the rise in allocation and systems tickets. The remaining tickets were spread across the categories.

Continuous Improvement Efforts in User Support

In order to accommodate these increases in ticket volume, the User Services team met weekly to consider key support processes. Examples of changes made during continuous improvement meetings:

- Reviewed notifications from various systems and removed notifications that were duplicated or generally ignored.
- Identified a difference between helpdesk and non-helpdesk tasks (i.e., rapid response and real-time work with users v. fostering key processes such as account creation) and changed how User Services approached these different tasks.

2 In previous years, this category was called Bounces. The name was changed at the request of the reviewers.

3 Ticket categorization as of January 2014

- Added communication of the project type between helpdesk and administrator responsible for master user agreement to reduce errors in MUA processing and removed extraneous information sent to the administrator.
- Characterized a problem in Cryptocard token initialization, fixed initialization problem, and redefined debug process for tokens.
- Learned about difference in perception between what ALCF staff believed was the start and end times of account creation and what the user perceived was this time lag. As a result, began to measure account request to login on the system – a perceptual measure of how long it took to provision a user account.
- Discovered that requirements changed on shipping Cryptocard tokens and a new process was added, documented, and reviewed.
- Added step in account renewal process asking users to update all contact information in ALCF databases to improve the quality of information such as affiliation and address for shipping tokens.
- Agreed upon and created ticket handoff process to improve response time to users.
- Reduced complexity in account application web page and made minor changes to decrease the time it takes to apply for an account as a result of direct user feedback.

Project and User Management Software Efforts

Several small but important features were added to the project and user management workflows this year. To assist in managing quotas, a script that synchronizes allocation membership and primary UNIX group was developed. Several ETL (Extract-Transform-Load) scripts were written to aggregate key data for reports to internal staff and to PIs.

An example of such support scripts is a one-way, project-to-UNIX group synchronization. When a member of a group is removed from an allocation, an e-mail is generated to inform the PI that the member has not been removed from the UNIX group. In this way, a PI can then decide whether that project member should continue to have access to the files. It is more often the case that members of a project need access to files, but should no longer run in the allocation.

Configuring Support Software for Mira

ALCF evaluated the Intrepid workflows before applying them to Mira. This included making changes to databases and changing language in automated messaging for users and projects. This project started two months before Mira went online and was interleaved into the existing support for both early users on Mira as well as Intrepid users. Updates included scripted graphical and textual reports for INCITE PIs, notifications sent to project PIs about the termination date of their allocation, and responses from account requests on Mira.

1.3.2 Communications

Communications through Mailing Lists and Social Media

ALCF provided information to users via five electronic communication channels: plain text messages to the system notify mailman lists, HTML-formatted monthly newsletters, intermittent social media postings, the ALCF website and custom tailored e-mail messages via scripts. Users can opt-out of the system notify and newsletter mailing lists.

The targets of these channels are:

- Plain text messages: current users on the systems with accounts
- HTML newsletters: users, customers, members of the scientific communities, students, as well as the general public
- Social media: users, followers of ALCF, collaborators, students, members of the larger scientific communities
- Website: users, collaborators, students, members of the larger scientific communities, as well as the general public
- Custom e-mail message: specific projects, user groups, and individual users (e.g., users with accounts expiring, all ALCC projects)

Table 1.6: 2013 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

This year ALCF changed the rate of release for newsletters from quarterly to monthly and moved the newsletter from a PDF-based format to an HTML formatted e-mail. Every month one or two “science stories” are featured. These stories are either the outcome of research done on ALCF resources or an advancement made by ALCF staff and researchers in the field. Critical deadlines, key events, and significant news stories are also included in the monthly publication. An example of a key deadline was informing users on the deadline for ALCC applications. For the first year of Newsbytes, existing subscribers continued to receive plain text e-mails with the option to change formats and new subscribers received HTML-formatted e-mails.

Custom Communications and Reports

Custom e-mail messages to users include weekly reports sent to PIs of INCITE and ALCC projects; 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.

For the 2013 INCITE and 2013 ALCC project start-up e-mails to PIs and project members, ALCF employed a checklist model to decrease complexity and to increase the number of projects that were up and running on the first day of the allocation. For each engagement, ALCF staff focused on one or two changes that were meant to encourage a faster adoption. The checklists are featured in the innovation section.

Table 1.7: List of Key Changes and Outcomes Due to Checklist Efforts

Year and Program	Change	Outcome
2013 INCITE	Implemented first checklist	PIs told ALCF they liked the concise list and often responded to the prompts in checklist format.
2013 ALCC	Created a second e-mail with a checklist focused on proxy assignment (missed in 2013 INCITE by many PIs)	Increased response rate to the second question and learned that shorter, focused checklists produced better results.

Communicating Decommissioning of Intrepid

This year ALCF started the process of decommissioning Intrepid. A division-wide project, coordination and collaboration were critical to keeping users informed about what actions to take, and when. The effort also involved teaching users about the enhanced data management tools available to them.

Once ALCF decided on a timeline for decommissioning, a messaging plan was developed. The communications team created and launched an informational webpage and hosted two live webinars on the topic in coordination with other ALCF staff. The webinar was recorded and made available for viewing.

Multiple messages were sent to users via system notification lists. ALCF developed targeted e-mails for Intrepid and Surveyor discretionary project PIs instructing them on how to archive, move, or delete data related to their project. All projects were told how to migrate an active project from Intrepid to Mira. Users have until March 31, 2014 to archive, move, or delete data on the file systems serving Surveyor and Intrepid.

Web Documentation Efforts for Mira

In late 2012 ALCF formed a committee devoted to producing documentation to allow users to take full advantage of Mira. This committee had representatives from across the organization, implemented a workflow for writing the documentation, organized a focus group, and worked closely with a set of technical writer consultants from the Communications, Education and Public Affairs (CEPA) Division at Argonne.

The committee set goals, produced a timeline, assessed risks, evaluated content, redefined navigational structure, developed a writing and editorial process, created an editorial schedule and successfully completed this project.

Other Publications

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

Table 1.8: Publications Designed for Print in 2013

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

* Publication developed by ALCF in odd years and by OLCF in even years

1.3.3 Outreach Efforts

User Advisory Council

The seven-member User Advisory Council (UAC) represents all three allocation programs (INCITE, ALCC, and Director's Discretionary.) The UAC provided feedback on the 2013 User Survey such as identifying language that might be difficult for foreign nationals, noting jargon, and beta testing the survey. The UAC also advised ALCF on the following topics:

- Documentation efforts for Mira
- Software support, storage management, and scheduling policies for ALCF systems
- INCITE and ALCC communications
- ALCF's data retention policy
- Messaging strategy for the decommissioning of Intrepid

Industry Outreach

ALCF continued to encourage and support industry use of HPC through focused interactions with companies both large and small. ALCF worked with Argonne's Technology Development and Commercialization (TDC) Division to integrate HPC into Argonne's business-focused outreach. ALCF also has established strong relationships with other Argonne divisions with interests, expertise, and resources in HPC, which has resulted in deeper collaborations with companies such as Shell, GE, and Toyota.

2013 Workshops and Webinars

ALCF conducted workshops and webinars to support the efforts of ALCF users and their project teams.

Table 1.9: 2013 Workshops and Webinars

Event	Description	Date
Getting Started Videoconference	Small-group telepresence event that ran several times over a two-week period for users new to Intrepid and ALCF.	January 2013
Mira Community Conference (MiraCon)	Onsite training event for projects awarded time on Mira for 2013 INCITE.	March 2013
ESP Community Meeting at MiraCon	Ran in tandem with MiraCon. ESP postdocs and users presented key technical challenges and solutions.	March 2013
Mira Performance Boot Camp	Tutorials on scaling and performance tuning codes for projects applying for 2014 INCITE awards.	May 2013
ESP Investigators Meeting	Investigators from each of the 16 ESP projects reviewed their simulation campaigns and scientific discoveries.	May 2013
Thirty Years of Parallel Computing at Argonne	Celebrated Argonne's pivotal role in scientific computing for the last thirty years and assessed the progress of the field.	May 2013
INCITE Proposal Writing Seminar	One of two seminars hosted by LCFs to disseminate INCITE information and provide a forum for questions.	May 2013
Argonne Training Program on Extreme-Scale Computing (ATPESC)	Training program for computational scientists to use current and next-generation leadership computing.	July/August 2013
Decommissioning Webinar	Information about migrating data off Intrepid/Surveyor.	November 2013

Getting Started Videoconference
While a "Getting Started" workshop for Mira was scheduled for March, to meet the needs of a smaller group of new Intrepid users, the ALCF opted to provide "Getting Started" instruction via a series of virtual workshops. By limiting the number of participants to six per session, staff presenters were able to pace the workshops accordingly and to provide individualized attention.

Each 4.5-hour workshop included an overview of ALCF services and resources, technical details on the IBM Blue Gene/P architecture, and guided assistance in porting and tuning applications on Intrepid. The smaller-scale event was an ideal scenario for testing the viability of a virtual workshop and was also used as a training opportunity for ALCF staff seeking to hone their presentation skills. Of the users who responded to a follow-up survey, all said they would participate in a future virtual workshop offered by the ALCF if the content pertained to their needs.

Mira Community Conference

Mira's first users, those involved in scaling representative codes to prototype Mira machines through the ALCF's Early Science Program, presented the results of their efforts over two days of the conference. The Early Science talks were grouped by common themes (i.e., algorithms, threads, communications), and provided attendees with insight and valuable lessons learned to apply to their own projects.

ALCF staff provided additional training sessions on how to take advantage of the latest generation hardware and software available on Mira. In addition, the Blue Gene Consortium convened a special session for system administrators, with attendees from U.S. and European facilities exchanging best-practice processes and solutions based on their respective Blue Gene system experiences.

Mira Performance Boot Camp 2013

The bulk of the four-day event was devoted to hands-on, one-on-one tuning of applications. ALCF experts also presented talks on topics of interest, including an overview of Blue Gene/Q hardware and software architecture, manual code-tuning tips for peak performance, submission of ensemble jobs, and message-passing API on Blue Gene/Q. Tools and debugger developers were on hand to provide information about their tools and individualized assistance to attendees. Accomplishment highlights include the following:

- Researchers using the GAMESS code received assistance redirecting files from GPFS to RAM disk, resulting in a speed-up factor of 13.
- A nuclear physics code that had scaled to only one rack of Blue Gene/Q earlier this year scaled to 32 racks during the workshop.
- Plasma physics researchers at work on quantum lattice gas simulations were able to scale their code from 16 racks to the full 48-rack system.
- Scientists focused on computer-aided drug design received instruction from ALCF staff on the effective use of NAMD and assistance in ensemble mode scripting.
- INCITE researchers in materials science improved the percentage of peak by a factor of two in preparation for a 48-rack run.
- An Early Science Program team compiled and benchmarked a new algorithm up to 32 racks for large-scale DPD-SPH simulations in LAMMPS for use in identifying bottlenecks.

2013 ATPESC

ALCF staff helped support the Argonne Training Program for Extreme-Scale Computing (ATPESC). ALCF provided web development, print design, administrative support, access to and support on Mira, as well as on-site support and expertise on compiling, porting, and tuning codes.

Decommissioning Webinar

A webinar was designed and presented about the Intrepid and Surveyor decommissioning. ALCF provided instruction on understanding the directory structure, using archiving features, taking advantage of data transfer tools, and deleting files on the file systems attached to Intrepid and Surveyor.

Conclusion

As a user facility, ALCF is focused on ensuring the success of all the facility users and customers. During CY 2013 ALCF released documentation about Mira, improved messaging processes, revised the user survey, acted on feedback from the User Advisory Council, and enhanced communication efforts in various e-mail channels. As such, ALCF continues to effectively support customers, resolve problems, and perform outreach.

Section 2. Business Results

Is the Facility maximizing the use of its 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, the 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

	Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM				Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	CY 2012		CY 2013		CY 2012		2013 (Apr 9 – Dec 31)	
	Target	**Actual**	**Target**	**Actual**	**Target**	**Actual**	**Target**	**Actual**
Scheduled Availability	95%	98.5%	95%	97.0%	N/A	N/A	85%	95.5%
Overall Availability	90%	95.9%	90%	94.7%	N/A	N/A	80%	90.6%
System MTTI	N/A	8.55 days	N/A	7.76 days	N/A	N/A	N/A	4.23 days
System MTTF	N/A	21.29 days	N/A	14.91 days	N/A	N/A	N/A	11.29 days
INCITE Usage	736M	803.8M	736M	821.5M	N/A	N/A	2.10B	2.41B
Total Usage	N/A	1.20B	N/A	1.17B	N/A	N/A	N/A	3.61B
System Utilization	N/A	87.6%	N/A	85.5%	N/A	N/A	N/A	79.4%
Intrepid INCITE Capability	40%	63.1%	40%	64.3%	N/A	N/A	N/A	N/A
Mira INCITE Capability A[1]	N/A	N/A	N/A	N/A	N/A	N/A	20%	60.7%
Mira INCITE Capability B[2]	N/A	N/A	N/A	N/A	N/A	N/A	5%	33.3%

[1] Capability A = Jobs using >= 16.7% (8 racks, 131,072 cores) of Mira.

[2] Capability B = Jobs using >= 33.3% (16 racks, 262,144 cores) of Mira.

ALCF Resources

During CY 2013, the ALCF operated two production resources, Mira and Intrepid. Mira, which officially went into production on April 9, is a 48K-node, 768K-core, 10 PF Blue Gene/Q with 768 TB of RAM. Mira mounts a GPFS file system with approximately 20 PB of usable space and has access to the facility-wide HPSS tape archive. Intrepid is the 40K-node, 160K-core, 557 TF Blue Gene/P with 80 TB of RAM. Intrepid mounts a production 4.5 PB GPFS file system, a 0.5 PB PVFS file system (primarily used for performance comparisons) and has access to the facility-wide HPSS system. Intrepid was officially retired from production at the end of CY 2013 and is in the process of being decommissioned. Mira has an associated visualization and analysis cluster called Tukey and Intrepid has one called Eureka. Eureka will continue to run until March 31, 2014 to enable users to complete analysis, transfer and/or archiving of their data and then it will be decommissioned. In the remainder of this section, Mira's metrics and reportable values will be addressed first and Intrepid's second.

Mira

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. ALCF has agreed, with the DOE Program Manager, to metrics of 80% overall availability and 85% scheduled availability, which is consistent with OLCF and NERSC for a machine in its first year of production. Table 2.2 summarizes the availability results:

Table 2.2: Availability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2012		2013 (Apr 9 – Dec 31)	
	Target	Actual	Target	Actual
Scheduled Availability	N/A	N/A	85%	95.5%
Overall Availability	N/A	N/A	80%	90.6%

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, annotated in Figure 2.1.

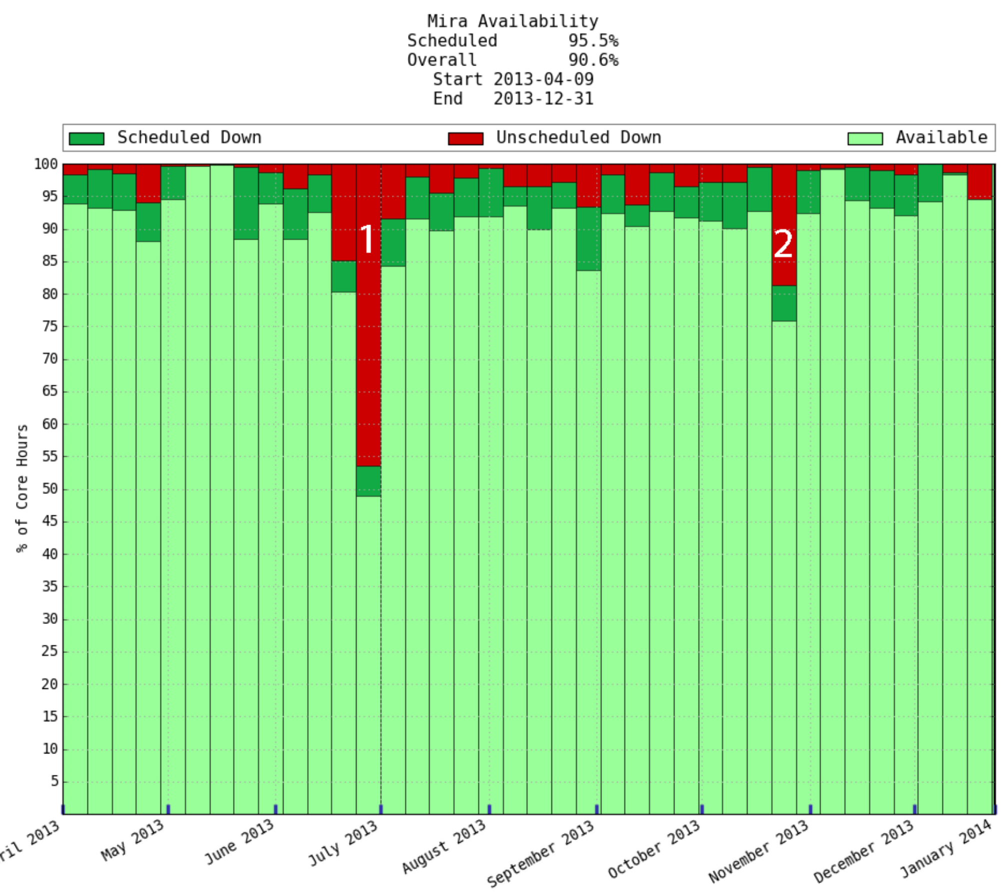

Figure 2.1: Mira Weekly Availability for CY 2013

Graph Description: Each bar represents the average of seven days of core-hour usage. Each bar accounts for all 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 (three weeks): The first of two outages began when a 1200 amp, 480 V breaker that feeds one of the Blue Gene power panels, tripped. This panel powers five of Mira's 48 racks. ALCF staff replaced the breaker, inspected the panel and power supplies in all five racks, and turned the breaker back on. As the system was returning to service, numerous Blue Gene node boards disappeared from the control system – far more than in past instances. (The disappearance of node boards is a known issue currently under investigation.) ALCF returned to service but tripped another breaker of the same type in the same panel five days later. On that occasion, one of the affected racks was a clock rack so all racks downstream were also rendered non-functional. Initially, ALCF swapped power cables to provide power to the clock rack and brought 43 of the 48 racks back into service. ALCF had an electrical contractor come in and megger test all the circuits to ensure there were no ground faults, replaced the breaker, and restored the system to full availability.

Item 2: This outage was due to a third 1200 amp breaker fault. ALCF again had the circuits megger tested, replaced the breaker, and brought the system back into service. ALCF is working with the building trust and the manufacturer (Eaton) to determine the root cause. Eaton has agreed to replace all three breakers free of charge.

ALCF has identified the cause of the node boards disappearing – a bug in the Linux kernel having to do with ARP cache overflow when using Infiniband – which can also prompt the service node to kernel panic and reboot. ALCF has an open bug with Red Hat, and they have provided two patches to try, but have not resolved the issue yet. To work around the situation, ALCF increased the size of the ARP cache and raised the timeouts associated with the Infiniband neighbors to avoid having to refresh the cache. ALCF will continue to work with Red Hat and Mellanox to get an actual fix for this problem.

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

A series of discussions occurred between ALCF, OLCF, and NERSC, and all three sites agreed to a common calculation for System MTTI and System MTTF. MTTI and MTTF are reportable values, but no specific metric 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				
	CY 2012		2013 (Apr 9 – Dec 31)	
	Target	**Actual**	**Target**	**Actual**
System MTTI	N/A	N/A	N/A	4.23 days
System MTTF	N/A	N/A	N/A	11.29 days

Mira is currently on a weekly maintenance schedule. ALCF takes the machine out of service every Monday to perform Blue Gene driver upgrades, hardware replacements, OS upgrades, etc. ALCF plans to move to bi-weekly maintenance in Q1 CY 2014. A weekly maintenance schedule implies a theoretical maximum of seven days for the system MTTI. For a machine in its first year of production, Mira has been relatively stable.

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 that were used (i.e., it is a measure of how busy the system was kept when it was available).

Total System Utilization

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

Utilization is a reportable value, but no specific metric has been set, though 80% or higher is generally considered acceptable for a leadership-class system. Table 2.4 summarizes ALCF utilization results and Figure 2.2 shows system utilization over time by program.

Table 2.4: System Utilization Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2012		2013 (Apr 9 – Dec 31)	
	Target	**Actual**	**Target**	**Actual**
System Utilization	N/A	N/A	N/A	79.4%

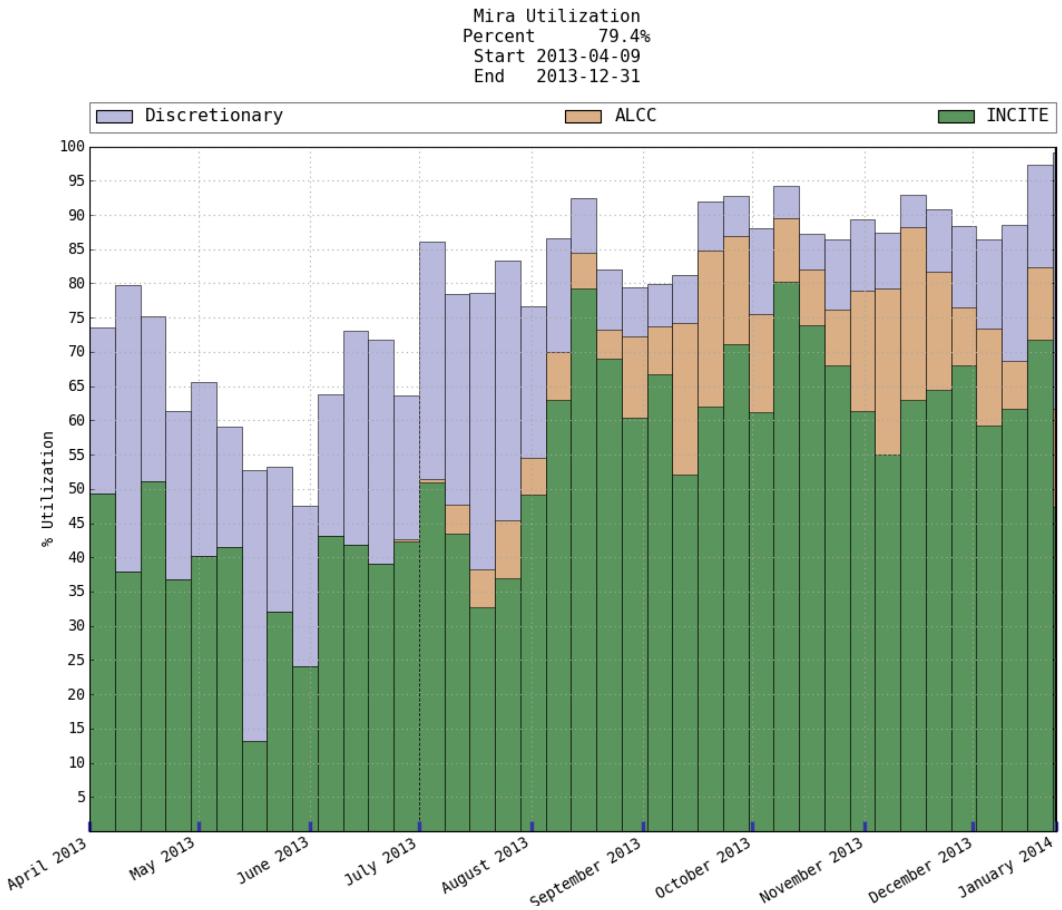

Figure 2.2: System Utilization over Time by Program

Due to ongoing efforts of the Scheduling committee and the Catalyst team, the system utilization for Mira steadily improved during the 2013 production period, achieving 79.4% for the entire period: 86.2% for the period of July 1 to December 31 and 90.1% for the period of October 1 to December 31.

One significant change made to improve the utilization involved changing the default block configuration for the 1K partitions to be a mesh in the D dimension. This enabled the scheduling of 1K jobs adjacent to each other. ALCF also added an additional configuration for the 4K blocks. All dimensions remained tori, but ALCF made the A dimension two hops, which eliminated the need for pass through on the C dimension and allowed ALCF to fully pack a row with 4K jobs. Finally, 32K jobs were made a mesh in the B dimension. This did not affect packing on the machine, but made the jobs more robust. The B dimension connects the rows together. When a 32K job is using B pass through, a failure anywhere on the machine would cause the job to fail. This change made the jobs more resilient to failure and allows staff to do hot maintenance on failed nodes without affecting 32K jobs. Tests showed that these changes had a minimal performance impact for most applications. For those applications that need a torus in all dimensions, they are available by submitting to non-default queues specifically to enable access to the original, full torus block configurations.

Table 2.5 shows how Mira's system hours were allocated and used by the allocation source. Multiplying the theoretical hours by availability and utilization values determines the hours available. Of the hours available, 60% is allocated to the INCITE program, up to 30% is available for ALCC program allocations, and 10% 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 the hours are assigned to each year. The allocated values for the DD allocations appear higher than expected, because they represent a rolling allocation. Since a majority of the DD projects are exploratory investigations, the time allocations are not used in full. The DD allocations are discussed in detail in the Strategic Results section. In CY 2013, the ALCF successfully delivered a total of 3.61 billion core-hours on Mira.

Table 2.5: Core-Hours Allocated and Used by Program

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM							
	CY 2012				2013 (Apr 9 – Dec 31)		
	Allocated		Used		Allocated	Used	
	%	Core-hours	%	Core-hours	Core-hours	Core-hours	%
INCITE	N/A	N/A	N/A	N/A	2.10B	2.41B	66.8%
ALCC	N/A	N/A	N/A	N/A	555.5M	348.4M	9.6%
DD	N/A	N/A	N/A	N/A	824.8M	259.4M	7.2%
ESP*	Multi-year allocation; hours used is production time only				2.04B	591.9M	16.4%
Total	N/A	N/A	N/A	N/A	5.52B	3.61B	100.0%

* As discussed in Strategic Results (Section 3), the Early Science Program (ESP) projects were allowed to continue to run after Mira went production in April. In Table 2.5, the core-hours used by ESP projects are the hours used after Mira was transitioned to production. The ESP allocations in this table reflect the total allocation to these projects, which cover multiple years. These allocation amounts are total for the program and pre-date production.

Summary: For CY 2013, the system usage and utilization values are in line with general expectations. Utilization was down early, but was improved in the last six months 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 Intrepid, capability is defined as using greater than 20% of the machine. However, 20% of Mira would be 9.6 racks, which is not a viable configuration. Hence, the Mira metric was defined in two parts. Part A requires a minimum of 20% of the INCITE core-hours be eight racks or more (16.7%) and Part B requires a minimum of 5% of the INCITE core-hours be 16 racks or more (33.3%). These targets increase to 30% and 10% respectively for CY 2014, and 40% and

10% respectively in CY 2015. Table 2.6 and Figure 2.3 show that ALCF has substantially exceeded these metrics set for INCITE. Although there are not metrics set, data is also provided in the table for ALCC and DD projects as reference, and Figure 2.4 shows the overall distribution of job sizes over time.

Table 2.6: Capability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2012			2013 (Apr 9 – Dec 31)		
Capability Usage	**Total Hours**	**Capability Hours**	**Percent Capability**	**Total Hours**	**Capability Hours**	**Percent Capability**
INCITE A	N/A	N/A	N/A	2.41B	1.47B	60.7%
INCITE B	N/A	N/A	N/A	2.41B	803.1M	33.3%
ALCC A	N/A	N/A	N/A	348.4M	134.0M	38.5%
ALCC B	N/A	N/A	N/A	348.4M	54.9M	15.8%
Director's Discretionary A	N/A	N/A	N/A	851.3M	291.3M	34.2%
Director's Discretionary B	N/A	N/A	N/A	851.3M	190.8M	22.4%
TOTAL A	N/A	N/A	N/A	3.61B	1.89B	52.3%
TOTAL B	N/A	N/A	N/A	3.61B	1.05B	29.0%
ESP Director's Discretionary A*	N/A	N/A	N/A	591.9M	150.3M	25.4%
ESP Director's Discretionary B*	N/A	N/A	N/A	591.9M	95.8M	16.2%

* ESP Director's Discretionary A and B values are additional detail of the Director's Discretionary values.

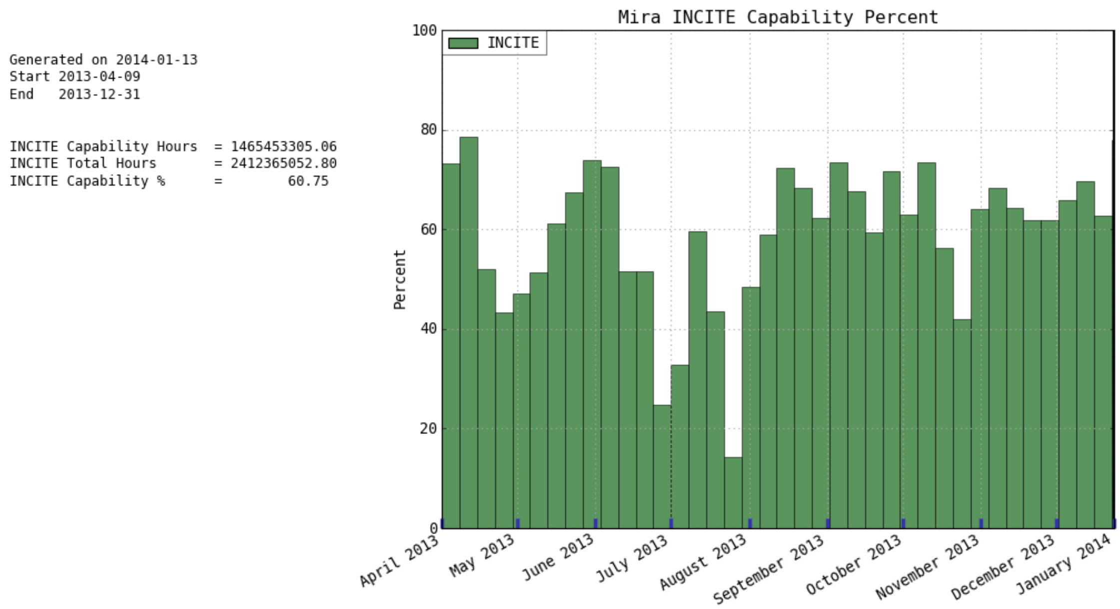

Figure 2.3: Mira Capability Usage over Time

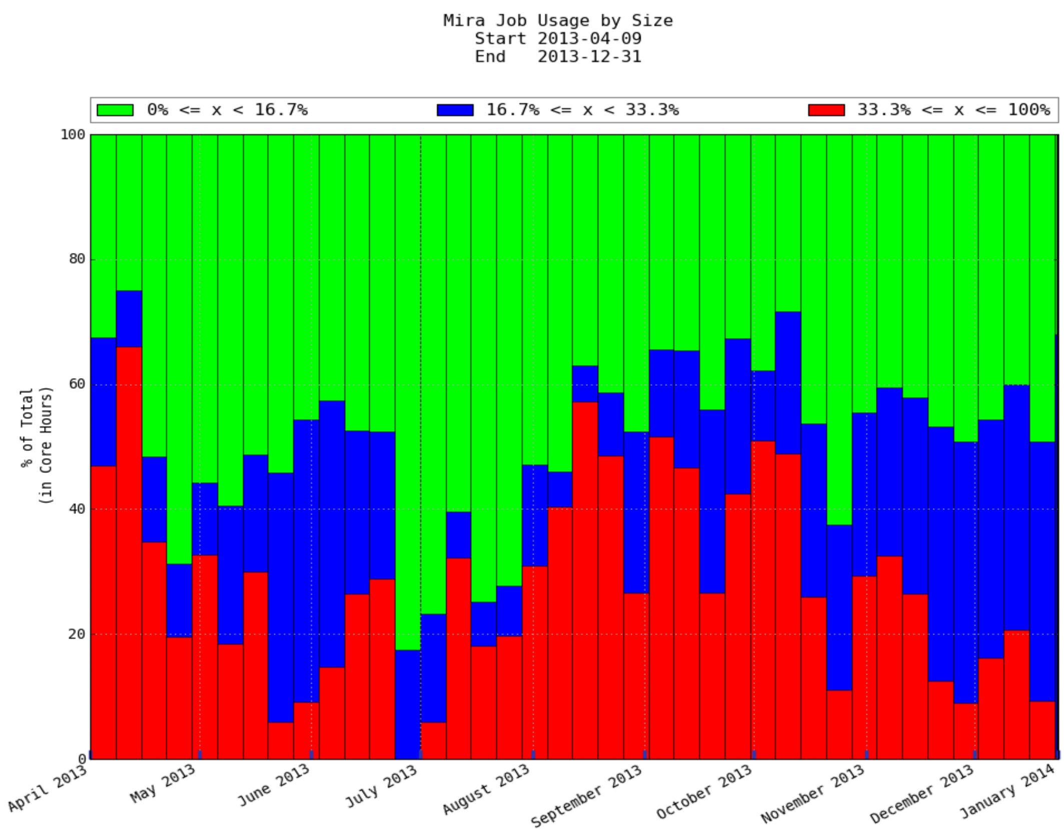

Figure 2.4: Mira Job Usage by Size

Intrepid

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.

2.1.1 Scheduled and 2.1.2 Overall Availability

Intrepid has been in full production since February 2009. ALCF has agreed, with the DOE Program Manager, to metrics of 90% overall availability and 95% scheduled availability, which is consistent with OLCF and NERSC. Table 2.7 summarizes the availability results:

Table 2.7: Availability Results

Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM				
	CY 2012		CY 2013	
	Target	Actual	Target	Actual
Scheduled Availability	95%	98.5%	95%	97.0%
Overall Availability	90%	95.9%	90%	94.7%

Explanation of Significant Availability Losses on Intrepid

The remainder of this section covers significant availability losses, annotated by the numbers in Figure 2.5, and the ALCF's responses to these issues. Details on how the calculations are handled can be found in Appendix A.

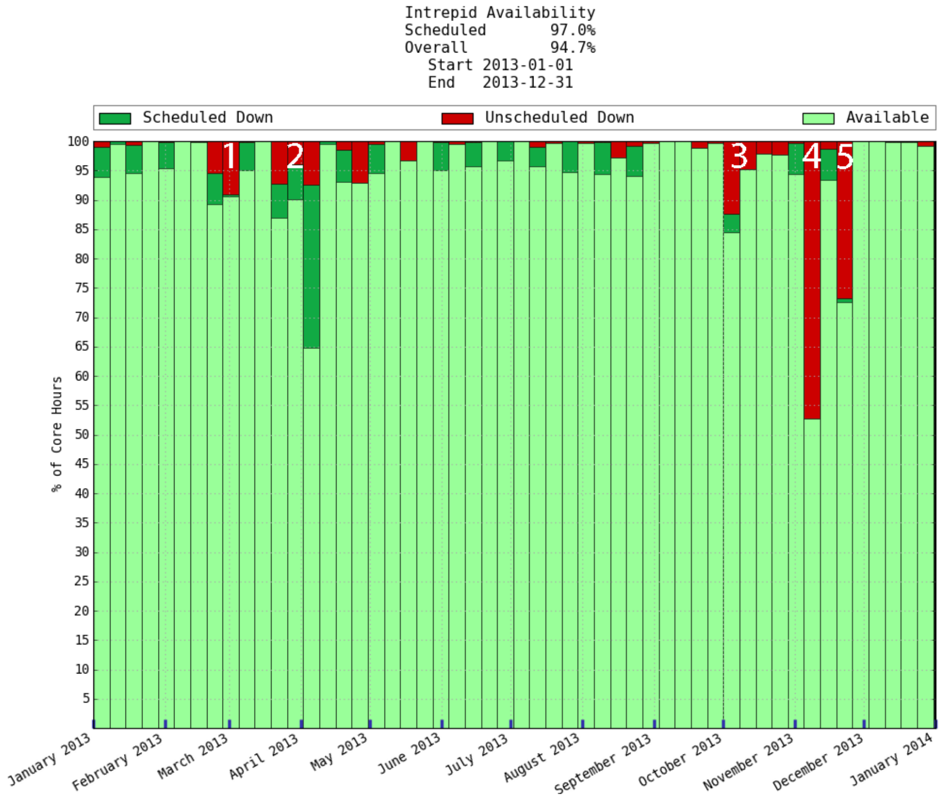

Figure 2.5: Intrepid Weekly Availability for CY 2013

Graph Description: Each bar represents the average of seven days of core-hour usage. Each bar accounts for all 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: This downtime was a combination of several events, two of which were major. ALCF found a race condition in the Cobalt scheduler where a job terminated, but the scheduler missed the notification and therefore did not schedule jobs most of the weekend. There was also another service card failure, as described in the previous two OARs. This causes an electrical short and damages substantial amounts of equipment. IBM will credit ALCF for all the damaged hardware. Additionally, there was scheduled maintenance, a node card replacement, and a SERDES failure.

Item 2: This downtime was due to an offline file system check (fsck) of the primary General Parallel File System (GPFS) (/intrepid-fs0). The downtime was the time it took to reconfigure the system so that /intrepid-fs0 was not mounted and set up special queues for those who could make use of Parallel Virtual File System (PVFS). A likely cause of the corruption was a power outage a few weeks before that caused the ALCF's Uninterruptible Power Supply (UPS) to run dead. Due to Intrepid's scheduled decommissioning, ALCF chose to accept the risk rather than invest in hardware changes. On Mira, a second file system was installed to mitigate such an occurrence. The second file system is expected to be in production use by the end of Q1 2014.

ALCF also developed a script to shut down critical services like GPFS and DB2 if UPS power is used for longer than a configurable window (currently 5 minutes).

Item 3: This was a combination of two issues. First, there was another catastrophic service card failure. This issue has been addressed previously and IBM would replace any needed parts free of charge. Before repairs could be completed, a power failure occurred due to a UPS trip. ALCF had scheduled a vendor to replace the batteries in the UPS. When coming out of UPS bypass, the power tripped. ALCF had made this transition multiple times before and could not determine a root cause for why it tripped out in this particular instance.

Item 4: US-27, the transformer that provides power to the UPS and therefore, everything in the Interim Supercomputing Support Facility (ISSF) other than the Blue Gene racks, shorted to ground and was a total loss. ALCF initially tried to use an Argonne-owned 600 kW generator to restore power. Under normal load, it would have been sufficient, but the additional load required to recharge the UPS batteries ended up being around 900 kW. The next day two 500 kW generators were rented to bring the ISSF back into service.

Item 5: Argonne had a transformer available on site that had been replaced during an upgrade of the Advanced Photon Source (APS) Facility available on site. This was installed as a temporary replacement of US-27 to run the ISSF until operations are completely ceased, which will be some time after April 1, 2014. This outage was required to disconnect and return the rented generators and tie the new transformer into the automatic transfer switch to provide utility power to the UPS again.

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

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

ALCF MTTI and MTTF Summary

A series of discussions occurred between ALCF, OLCF, and NERSC, and all three sites agreed to a common calculation for System MTTI and System MTTF. MTTI and MTTF are reportable values, but no specific metric has been set. Table 2.8 summarizes the current MTTI and MTTF values for Intrepid.

Table 2.8: MTTI and MTTF Results

Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM				
	CY 2012		CY 2013	
	Target	Actual	Target	Actual
System MTTI	N/A	8.55 days	N/A	7.76 days
System MTTF	N/A	21.29 days	N/A	14.91 days

Intrepid MTTI and MTTF values are down from last year. This is primarily due to the file system and power problems described above under availability.

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 that were used (i.e., it is a measure of how busy the system was kept).

Total System Utilization

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

Utilization is a reportable value and no specific metric has been set. For leadership-class systems, 80% or higher is generally considered acceptable. System utilization was down slightly from the previous year, but was still well above 80%. Table 2.9 summarizes ALCF utilization results and Figure 2.6 shows system utilization over time by program.

Table 2.9: System Utilization Results

Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM				
	CY 2012		CY 2013	
	Target	Actual	Target	Actual
System Utilization	N/A	87.6%	N/A	85.5%

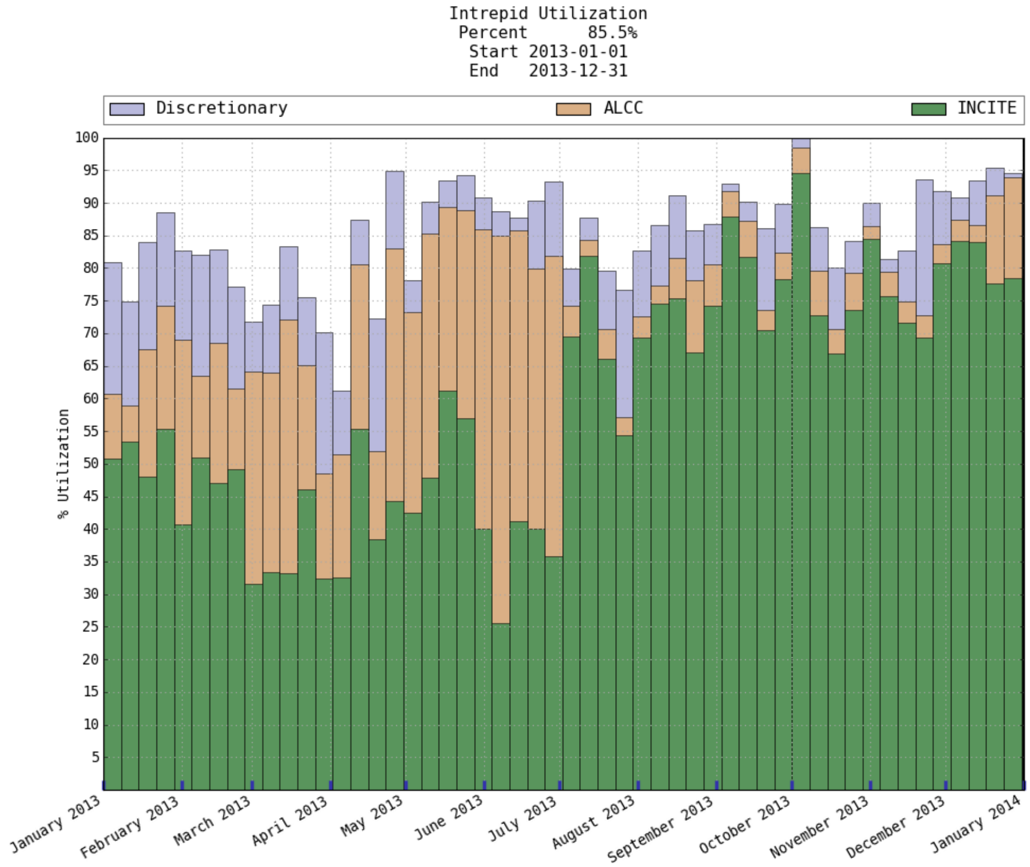

Figure 2.6: System Utilization over Time by Program

Table 2.10 shows how Intrepid's system hours were allocated and used by allocation source. Multiplying the theoretical hours by availability and utilization values determines the hours available. Of the hours available, 60% is allocated to the INCITE program, up to 30% is available for ALCC program allocations, and 10% is available for DD allocations. The ALCC program runs from July through June, so to arrive at allocated values for the calendar year, half the hours are assigned to each year. The allocated values for the DD allocations appear higher than expected, because they represent a rolling allocation. Since a majority of the DD projects are exploratory investigations, the time allocations are not used in full. The DD allocations are discussed in detail in the Strategic Results section. In CY 2013, the ALCF successfully delivered a total of 1.17 billion core-hours across all users, on Intrepid.

Table 2.10: Core-Hours Allocated and Used by Program

Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM							
	CY 2012				CY 2013		
	Allocated		Used		Allocated	Used	
	%	Core-hours	%	Core-hours	Core-hours	Core-hours	%
INCITE	60%	736.0M	66.7%	803.8M	736.0M	821.5M	70.0%
ALCC	30%	216.0M	16.6%	199.6M	225.4M	226.6M	19.3%
DD	10%	433.4M	16.7%	200.6M	488.6M	126.1M	10.7%
Total	100%	1.38B	100.0%	1.20B	1.45B	1.17B	100.0%

Summary: For CY 2013, the system usage and utilization values are in line with general expectations. The calculations for utilization are described in Appendix A.

2.3 Capability Utilization

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

ALCF Capability Utilization

Table 2.11: Capability Results

Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM						
	CY 2012			CY 2013		
Capability Usage	**Total Hours**	**Capability Hours**	**Percent Capability**	**Total Hours**	**Capability Hours**	**Percent Capability**
INCITE*	803.8M	507.4M	63.1%	821.5M	528.5M	64.3%
ALCC	199.6M	154.6M	77.4%	226.6M	140.5M	62.0%
DD	200.6M	133.8M	66.7%	126.1M	50.4M	40.0%
TOTAL	1.20B	795.8M	66.1%	1.17B	719.4M	61.3%

* The 2012 and 2013 metrics are both 40% for INCITE.

ALCF uses the Cobalt job scheduler to automatically schedule user submitted jobs without human intervention. The scheduling policy on Intrepid is designed to support capability usage, machine utilization, and quick job turnaround. Jobs are divided into queues based on their requested node count. Jobs requesting 20% or more of the machine are routed to prod-capability, where they have access to the entire machine. Smaller jobs are split again, based on their requested wall time. Small jobs with a requested wall time over six hours are restricted to running on just 40% of the machine. Small jobs with a shorter runtime are allowed to run anywhere on the machine. The effect of this policy is to ensure that when a capability job is selected to run, 60% of the machine will be available with a maximum eviction time of six hours. It is worth noting that the requested number of nodes does not always equal the allocated number of nodes. This is because the Blue Gene architecture does not allow the allocation of an arbitrary number of nodes. Hardware and policy constraints mean that only specific "quanta" of nodes can be allocated. When a user requests a job size that is not equal to one of these quanta, the next larger size is allocated. A job is not counted as capability unless the requested job size *is* capability. If the request is for something smaller but gets allocated a "capability sized" partition, that job will not count as capability. Requests for odd-sized jobs represent just a small percentage of the core-hours requested over the year.

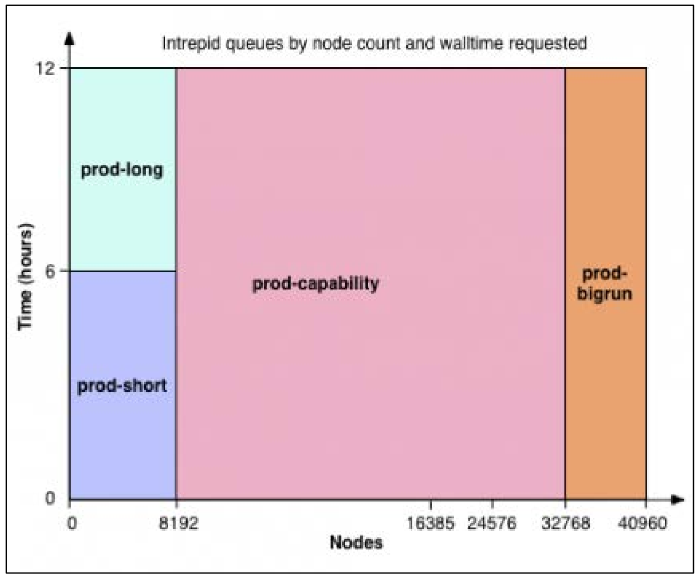

Figure 2.7: Queue Assignment Based on Job Size and Run Time

To improve utilization, ALCF will use any available short job to fill in the sub-partitions that would otherwise be idle while they wait for the partition to drain. This is known as backfill. Another utilization feature used is called "Big Run Monday." Every Monday, job scores are adjusted manually to maximize utilization. In general, a "best pack" is performed, which considers both job size and run time, with job size generally taking precedence. Additionally, relative job order is maintained where possible, no single project is allowed to dominate the time, and packing is limited to a maximum of 48-60 hours. In this way, the machine is not constantly entering a draining/backfill state when one of the large jobs is run. The weak point of the scheduling policy is turnaround time on prod-long (small, long-duration) jobs. ALCF mitigates this through user education about the better queue wait times for prod-short (small, short) jobs.

Summary: ALCF exceeded the target of delivering 40% of the INCITE cycles to capability jobs. In fact, all three categories (INCITE, ALCC, and DD) achieved 40%, as shown in Table 2.11 above.

The remainder of this section presents graphs showing the capability use over time, by INCITE, ALCC, and Discretionary for CY 2013, and then a breakdown of how the total core-hours delivered were distributed across job sizes (Figures 2.8 and 2.9).

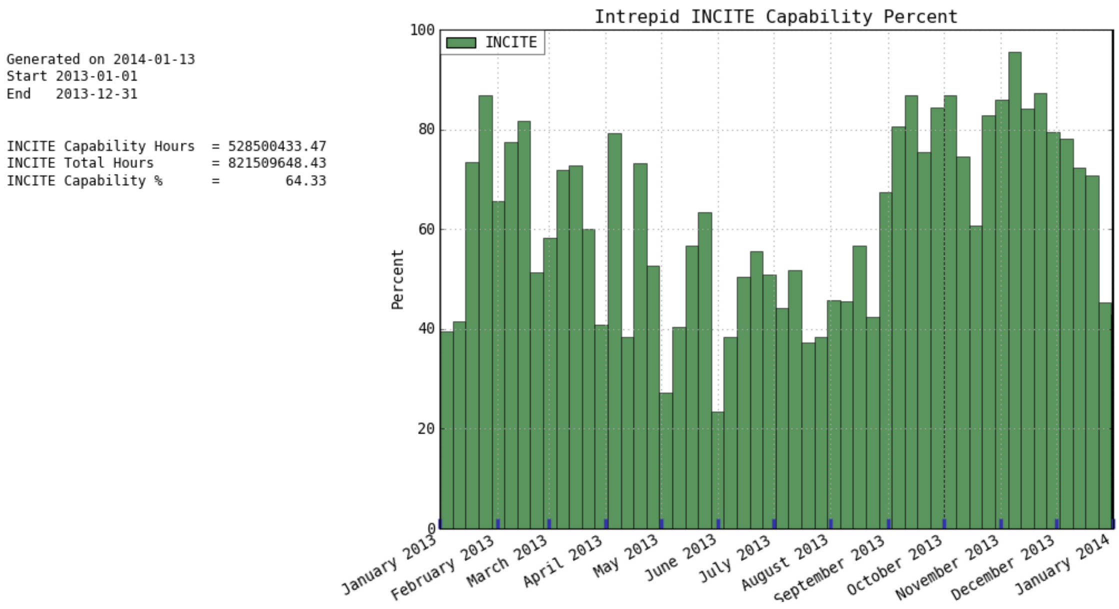

Figure 2.8: Cumulative INCITE Capability Usage during CY 2013

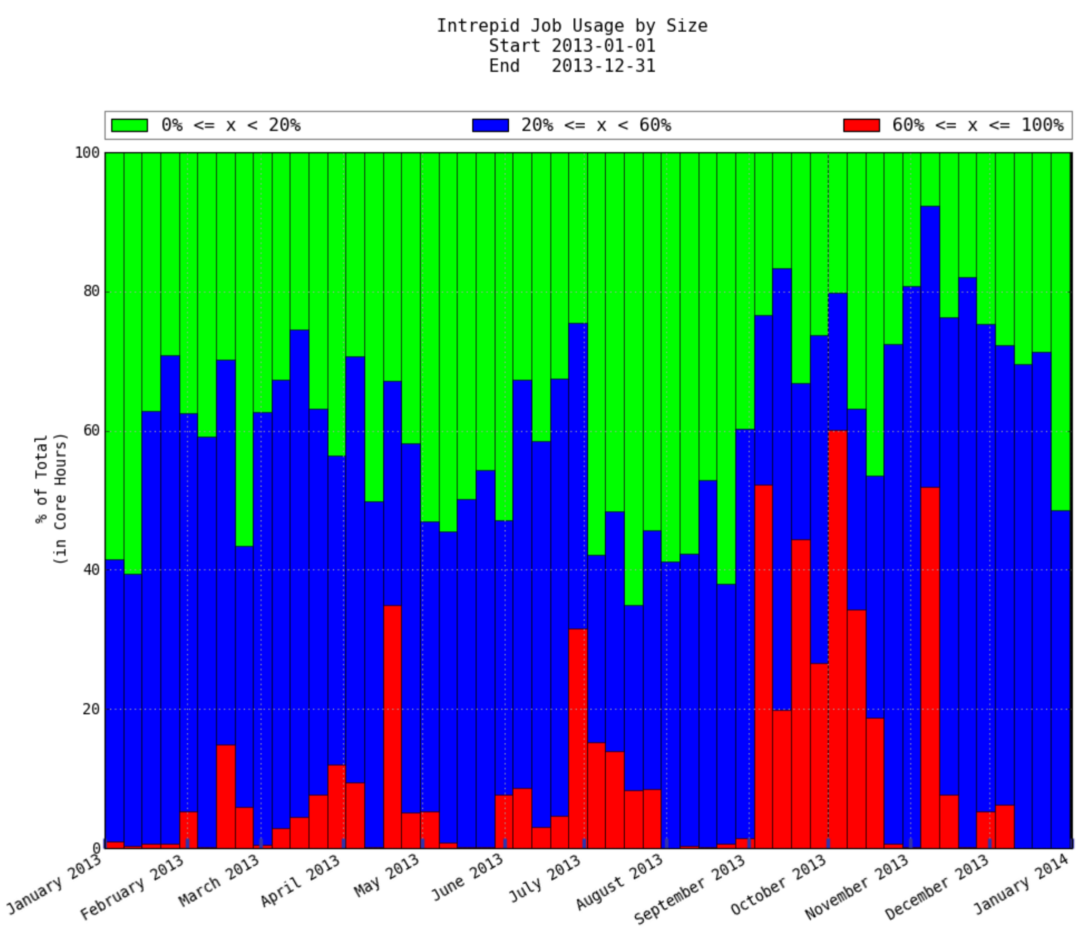

Figure 2.9: Intrepid Jobs Binned by Size

2.4 Financial and Human Resources

Financial data and staffing distribution are covered in detail in the onsite review. A summary is provided here.

- ALCF carefully plans the use of its available resources, maximizing the impact of its staffing distribution through strategic use of matrixed employees and contractors where consistent with its mission.
- Annual and multi-year budgets are reviewed with ASCR and the Argonne Site Office (DOE-ASO) on a regularly scheduled basis. Actual costs are tracked closely against approved budgets, and monthly reports are provided to and discussed with ASCR and DOE-ASO.
- Ongoing communication with ASCR and DOE-ASO includes a weekly conference call and a monthly staffing update. This line of communication allows for a discussion of the best use of ALCF resources. In FY 2013, ALCF-2 equipment purchases, lease down payment, and extra lease payment were coordinated efforts with the ASCR program office.

Conclusion

ALCF is maximizing the use of its HPC systems and other resources, consistent with its mission. ALCF has 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 2.12 below.

Table 2.12: Summary of All Metrics Reported in the Business Results Section

	Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM				Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	CY 2012		CY 2013		CY 2012		2013 (Apr 9 – Dec 31)	
	Target	Actual	Target	Actual	Target	Actual	Target	Actual
Scheduled Availability	95%	98.5%	95%	97.0%	N/A	N/A	85%	95.5%
Overall Availability	90%	95.9%	90%	94.7%	N/A	N/A	80%	90.6%
System MTTI	N/A	8.55 days	N/A	7.76 days	N/A	N/A	N/A	4.23 days
System MTTF	N/A	21.29 days	N/A	14.91 days	N/A	N/A	N/A	11.29 days
INCITE Usage	736M	803.8M	736M	821.5M	N/A	N/A	2.10B	2.41B
Total Usage	N/A	1.20B	N/A	1.17B	N/A	N/A	N/A	3.61B
System Utilization	N/A	87.6%	N/A	85.5%	N/A	N/A	N/A	79.4%

Table 2.12: (Cont.)

	Intrepid (Blue Gene/P) 40K-node, 160K-core, 850 MHz, 80 TB RAM				Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM			
	CY 2012		CY 2013		CY 2012		2013 (Apr 9 – Dec 31)	
	Target	**Actual**	**Target**	**Actual**	**Target**	**Actual**	**Target**	**Actual**
Intrepid INCITE Capability	40%	63.1%	40%	64.3%	N/A	N/A	N/A	N/A
Mira INCITE Capability A[1]	N/A	N/A	N/A	N/A	N/A	N/A	20%	60.7%
Mira INCITE Capability B[2]	N/A	N/A	N/A	N/A	N/A	N/A	5%	33.3%

[1] Capability A = Jobs using >= 16.7% (8 racks, 131,072 cores) of Mira.

[2] Capability B = Jobs using >= 33.3% (16 racks, 262,144 cores) of Mira.

ALCF closely tracks hardware and software failures and their impact on user jobs and metrics. This data is 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 DD projects clearly demonstrate the impact that ALCF is having in supporting scientific breakthroughs. ALCF staff has worked effectively with individual project teams to adapt their simulation codes to run efficiently in an HPC 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.

Community Involvement and Outreach

Early science on Mira was already underway at the start of 2013 and the Early Science Program (ESP) projects continued to run throughout the system's transition to production. Each of the 16 ESP teams contributed a published paper or a technical report on their technical work involving Mira, and presented their scientific and technical findings at two ALCF-sponsored meetings on the topics.

ALCF staff helped organize and run the Argonne Training Program on Extreme-Scale Computing (ATPESC), a two-week training program that hosted 63 students and included lecturers from top institutions in the United States and abroad. Staff also participated in the annual Grace Hopper Celebration of Women in Computing, the 2013 Google Summer of Code, and the Computational Science Graduate Fellowship (CSGF) conference, as well as mentored over a dozen postdocs and students.

The ALCF continues to engage the HPC and science communities. Division staff has memberships in the AAAS, ACM, ACS, APS, OpenMP Architecture Review Board, Open MP Accelerator, Language and Tools forums, SciComp (including president), SPEC-HPG, and the C++ committee. Staff members were on the organizing committees for half a dozen conferences and workshops, including SC'13.

Three members of the facility moved on to faculty positions at universities. Additionally, four ALCF staff members were recognized for their contributions to their respective fields, winning the Thomas Kuhn Paradigm Shift Award, the Adelchi Fabrocini Award, the Thomas Hart Benton Mural Medallion, and the IEEE TCSC Young Achiever in Scalable Computing. Two members of the division participated in the Strategic Laboratory Leadership Program.

The division regularly participated in the larger HPC community through conferences, meetings, and workshops. Staff presented invited talks at universities throughout the country and internationally and traveled to universities, laboratories and business for collaboration meetings. Staff also participated in over 20 domain conferences including SPIRE, ACS, SIAM, ACM, APS, SC'13, and IEEE/IPDPS, with invited talks, paper and poster presentations and tutorials. The invited talks include the Peebles Memorial Lecture in Information Technology; the Hume-Rothery Award Symposium; and the SIAM Conference on Mathematical Aspects of Materials Science. The lectures include the Brazilian School of Molecular Modeling and S212C2M2 2013 (NSF Summer School).

These activities keep the staff engaged with the current state of the art and also provide a key mechanism to inform scientists about the facility. When staff attends conferences or visits universities, they include information about the INCITE and ALCC allocations programs and how ALCF assists researchers to prepare. Researchers are encouraged to apply for a discretionary award to explore the resources and to attend webinars and on-site workshops. The facility provides advice and potential collaboration if more time is needed to develop a competitive INCITE or ALCC proposal. In 2014, all the new INCITE projects had previous access to the facility through a discretionary award.

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.

Nature	PNAS	Physical Review Letters	SC'13	Total 2013 Publications
1	4	9	10	154

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.

Kinetic Simulations of Fusion Energy Dynamics

William Tang, Princeton Plasma Physics Laboratory/Princeton University

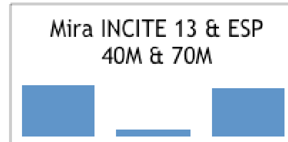

Building the scientific foundations needed to develop fusion power demands high-physics-fidelity predictive simulation capability for magnetically confined fusion energy (MFE) plasmas like ITER. Unavoidable spatial variations in such systems produce microturbulence, which can significantly increase the transport rate of heat, particles, and momentum across the confining magnetic field in tokamak devices. Since the balance between these energy losses and the self-heating rates of the actual fusion reactions will ultimately determine the size and cost of an actual fusion reactor, understanding and possibly controlling the underlying physical processes is key to achieving the efficiency needed to help ensure the practicality of future fusion reactors.

The goal of this project is to gain new physics insights on MFE confinement scaling by making effective use of the ALCF's Mira supercomputer. The knowledge gained will also address the key question of how turbulent transport and associated confinement characteristics scale from present generation devices to the much larger ITER-scale plasmas. This involves the development of modern software capable of using low memory-per-core leadership-class supercomputers to carry out reliable first principles-based simulations of multiscale tokamak plasmas.

The fusion physics challenge is that the key decade-long MFE estimates of confinement scaling with device size (the so-called "Bohm to gyro-Bohm" rollover trend caused by the ion temperature gradient instability) demands much higher resolution to be realistic and reliable. However, results from this project have revealed that this rollover is much more gradual than established earlier (in far lower resolution, shorter duration studies) with magnitude of transport now reduced by a factor of two.

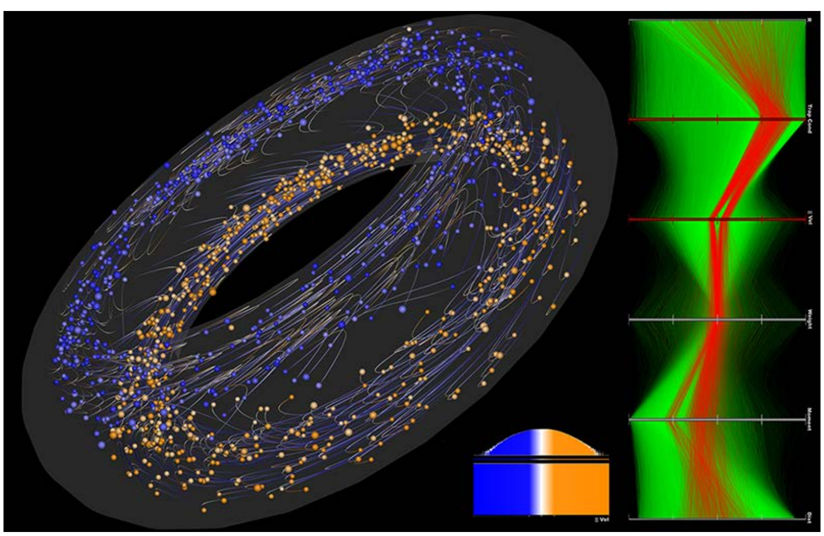

Figure 3.1: Particle visualization of a global gyrokinetic particle-in-cell simulation of microturbulence in a tokamak fusion device from the Princeton Plasma Physics Laboratory project.

IMPACT: Build a predictive foundation for fusion power.

ALCF Contribution: Helped port, benchmark, tune GTC-P C code on Mira, and co-authored an SC'13 technical paper.

Publication: B. Wang, S. Ethier, W. Tang, T. Williams, K. Ibrahim, K. Madduri, S. Williams, L. Oliker, *Kinetic Turbulence Simulations at Extreme Scale on Leadership-Class Systems,* Proceedings of The International Conference for High Performance Computing, Networking, Storage and Analysis (SC'13), November 2013, Denver, CO.

Presentation: W.M. Tang, Supercomputing Grand Challenges & Advances in Fusion Energy Simulations; B. Wang, Gyrokinetic Particle-in-cell Method with Phase-Space Remapping for Long-Duration Simulations of Plasma Turbulence, at the 23rd International Conference on Numerical Simulation of Plasmas, September 14–16, 2013, Beijing, China.

Award: International Data Corporation's 2013 HPC Innovation Excellence Award

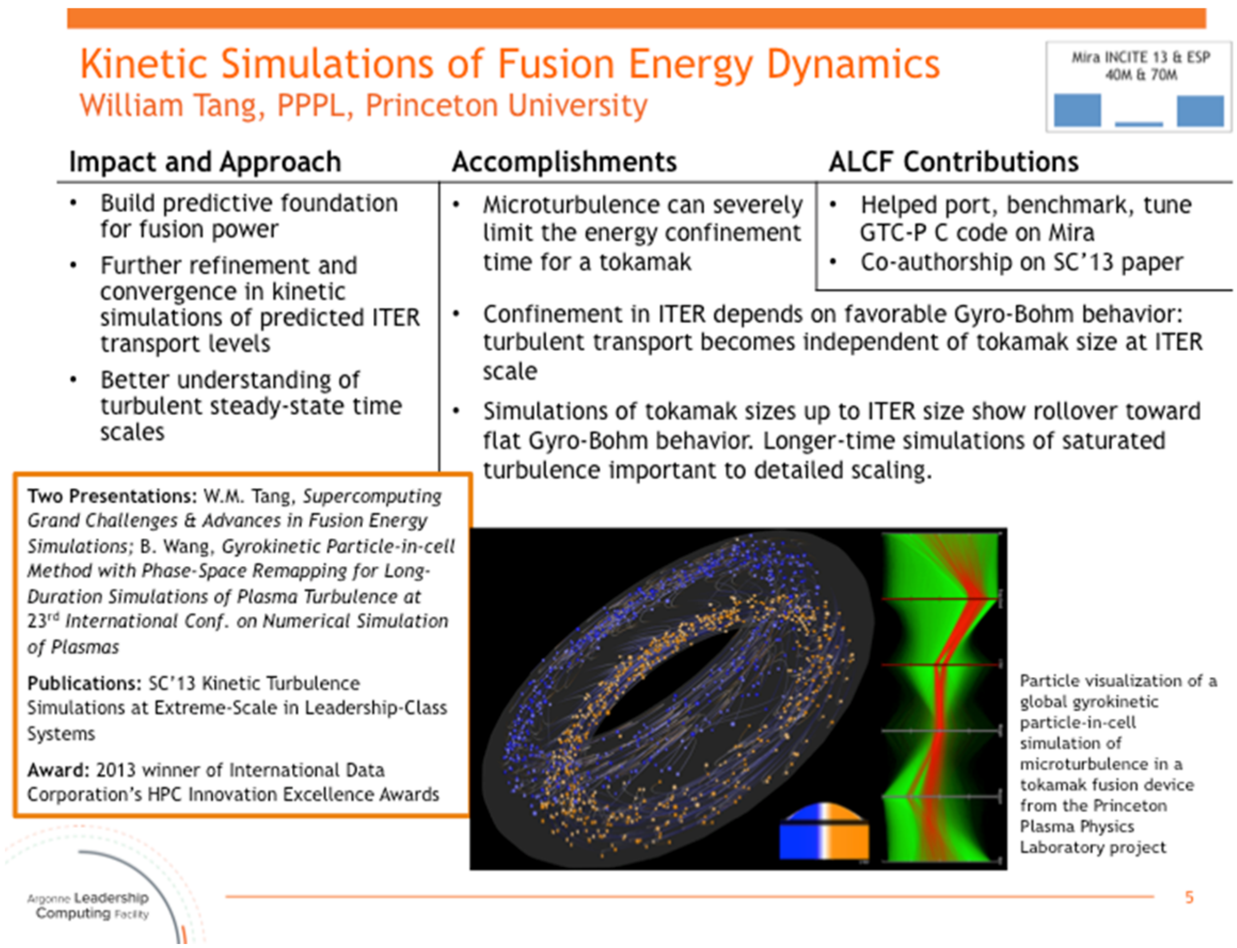

Thermodynamics of Binding Biomass to Cellulases for Renewable Fuel

Michael Crowley, National Renewable Energy Laboratory

The continued development of superior enzymatic catalysts to convert highly recalcitrant biomass into fermentable sugars is key to realizing industrial-scale biofuels production.

Due to long conversion times and the large size of the enzymes and enzyme complexes, conventional experimental design approaches are not suited to screen for and test improved cellulose-degrading enzymes. Understanding the molecular-level steps in cellulose synthesis and conversion will enable improved experimental design approaches to develop superior enzymes and less recalcitrant plant-based feedstocks for liquid biofuels.

For this INCITE project, Michael Crowley and his team performed the first molecular-level theory simulations of an enzyme's ability to catalyze consecutive reactions (processivity) that relate directly to structural features of enzymes. In close collaboration with the ALCF Catalyst team, the researchers ran large-scale free energy perturbation (replica exchange molecular dynamics) simulations on Intrepid and Mira. With the high throughput computations, they obtained the binding free energy information for the most popular enzymes in the biofuel industry.

The research team then developed a mathematical relationship formalizing that the processivity is directly linked to the ability of an enzyme to decrystallize a polymer chain from a crystal, given by the binding free energy of the enzyme to the saccharides. The binding free energies also suggest saccharide morphology is susceptible to enzyme action. Additionally, the research shows the ligand binding free energy is a key parameter in comparing the activity and function of enzymes, and offers a molecular-level basis for developing a theory of enzyme processivity and other enzymes able to hydrolyze linear saccharides.

Based on the significant simulation progress, the PIs gave two presentations at the 35th Symposium on Biotech for Fuels and Chemicals and the 2013 Gordon Research Conference. They are also drafting a paper "Binding free energy of oligosaccharides to glycoside hydrolases as a metric for processivity on insoluble crystalline substrates" to be submitted to the Journal of the American Chemical Society.

IMPACT: Provided tailored improvements to making biofuels based on the material available.

ALCF Contribution: ALCF staff developed and implemented the replica exchange algorithms in NAMD critical to the workflow for this project.

Publication: C.M. Payne, W. Jiang, M.R. Shirts, M.E. Himmel, M.F. Crowley, G.T. Beckham, *Glycoside Hydrolase Processivity Is Directly Related to Oligosaccharide Binding Free Energy*, Journal of the American Chemical Society, (135)50: 2013.

Presentation: C.M. Payne, W. Jiang, M.E. Himmel, M.F. Crowley, G.T. Beckham, Absolute ligand binding free energy of glycoside hydrolases as a metric for processivity and polysaccharide morphology dependence, at the 35th Symposium on Biotech for Fuels and Chemicals, 2013 Gordon Research Conference, April 29–May 2, Portland, OR.

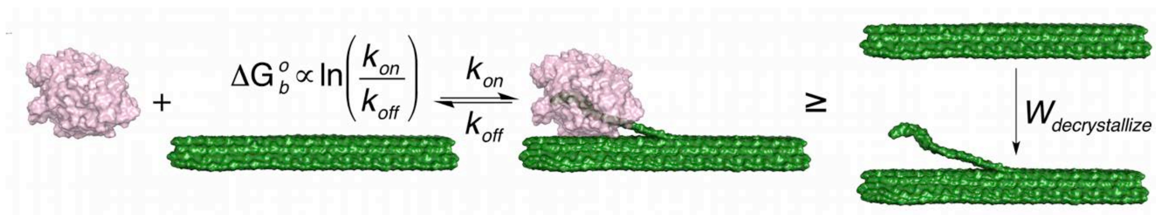

Figure 3.2: The absolute binding free energy is defined as the free energy change between a polysaccharide chain (of n monosaccharide units, where n is the chain length required to saturate the GH binding sites) and the enzyme-substrate complex in the catalytically active complex. The team hypothesizes that this quantity, **ΔGb°**, is directly correlated to processivity and suggests morphology dependence of cellulose attack by enzymes.

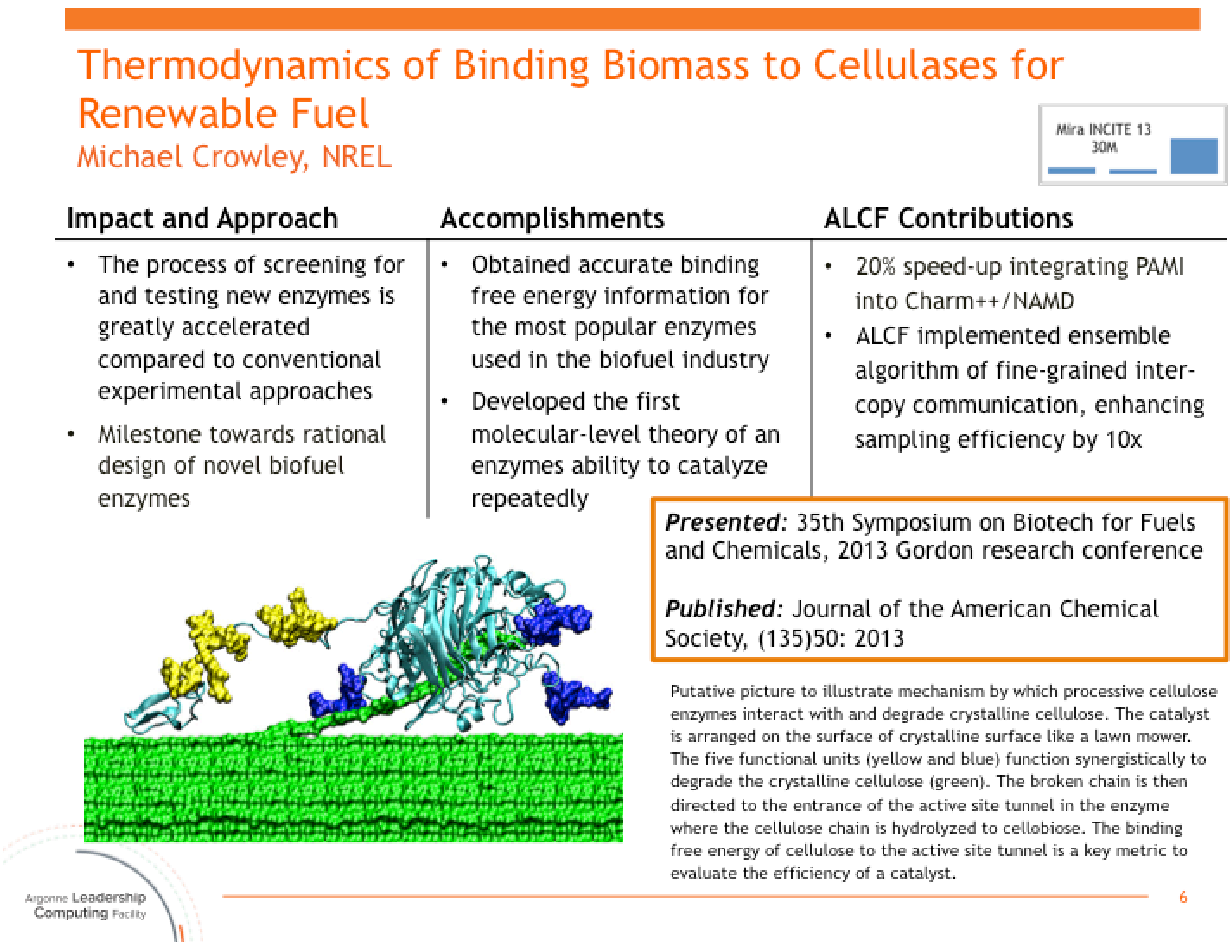

Explosion Research for Safer Building Designs

Thierry Poinsot, CERFACS

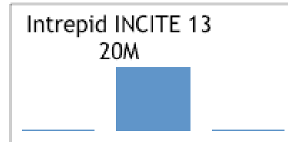

To aid in the design of safer buildings, computations are often used to evaluate the impacts of potential explosions caused by gas leaks. This is especially critical in the oil industry where leaks are common. Such accidents can lead to a financial loss and, more importantly, a loss of lives.

Typical explosion scenarios begin with an ignition of a quiescent flammable mixture inside a confined or semi-confined area. The flame then propagates past the repeated obstacles, generating a strong overpressure that is self-generated by the flow induced by the flame. This overpressure is the parameter of critical interest for safety studies since it can lead to the destruction of an entire facility.

Traditionally, safety-related studies of explosions for industry have been carried out using Unsteady Reynolds Averaged Navier-Stokes (URANS) methods. With this approach most of the physics is parametrized, which means it is not modeled using first principles methods but just mimicked or fitted. However, the emergence of large eddy simulations (LES) on massively parallel computers has significantly improved the ability to precisely simulate fully unsteady flows in complex geometries with much less parametrization and uncertainties. In small-size devices, even direct numerical simulations (DNS) seem feasible today.

With an INCITE allocation at the ALCF, CERFACS researchers are developing high-fidelity LES tools for realistic explosion cases, with sizes on the order of 10 m to 100 m and very large Reynolds numbers. To develop these LES tools and ensure they capture all phases of the flame development, the team's strategy is to begin with a smaller experiment (25 cm) installed in Sydney. The researchers performed the first simulation of an explosion in a building with ALCF resources using a 1-billion-cell mesh. The comparison between LES and experimental results, obtained without ad-hoc parametrization, is very good.

The 2012-2013 INCITE project results have been so interesting that the setup used for the LES validation (an experiment performed in Sydney) was replicated in Norway (GexCon) at two new scales (6 and 24 times larger). This created a unique experimental database where scaling laws can be tested on a range that has never been seen in the turbulent combustion community. Most combustion chambers have similar sizes (10 cm to 50 cm) so that models are rarely tested over a very wide range of Reynolds numbers. With the new database (Australia + Norway experiments), CERFACS now has access to a new range of parameters where the size of the chamber changes from 1 to 24x, everything else being equal (fuel, shape of chamber, pressure). This means that the volume of the chamber changes from 1 to 13,824x, a variation that has never been seen before. This opens the door to a very exciting field of research, not only for LES explosion research but also for modeling turbulent premixed combustion in general.

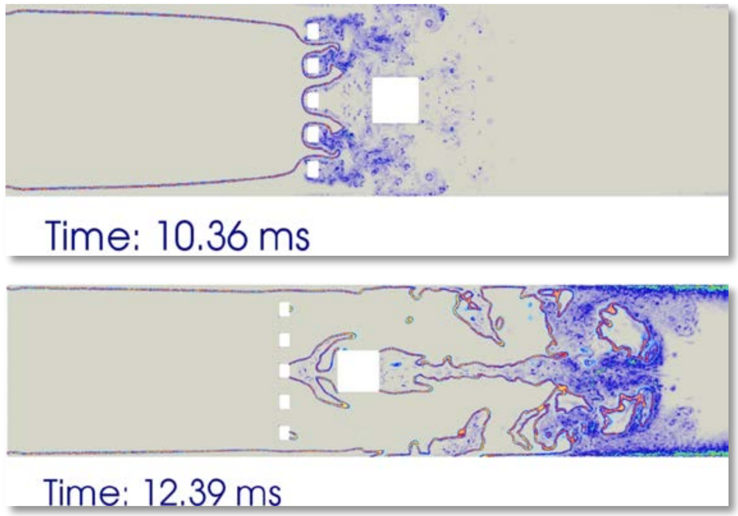

Figure 3.3: The flame position is displayed at two instants: before it reaches the second obstacle and after it has interacted with the obstacles. At the final instant, the flame is propagating at speeds of the order of 250 m/s.

IMPACT: Design of safer buildings by reducing the risk of catastrophic damage from explosion. This project has also driven new experiments.

ALCF Contribution: ALCF staff helped resolve problems related to memory over-allocations on one of the cores of the compute node, as the code transitions from using AVBP in a pure MPI mode to an MPI+OpenMP mode. Resolving this problem was necessary to run the large DNS and LES computations in this project.

Presentations:

- P. Quillatre, O. Vermorel, T. Poinsot, P. Ricoux, *Large Eddy Simulation of Vented Deflagration*, CRCT meeting (French Annual Turbulent Combustion Research meeting), Paris, France, March 2013.
- T. Poinsot, *HPC and combustion*, Invited talk PRACE GENCI meeting, CEA, Paris, France, March 2013.
- G. Staffelbach, *Parametric burner ignition*, Invited Talk Grand Challenge IDRIS, IDRIS, Orsay, France, March 2013.
- G. Staffelbach, *The HPC revolution: Innovation on the Field of Combustion Applications*, BoF Session, PRACE & HPC Services for Industry, in ISC 2013, Leipzig, Germany.

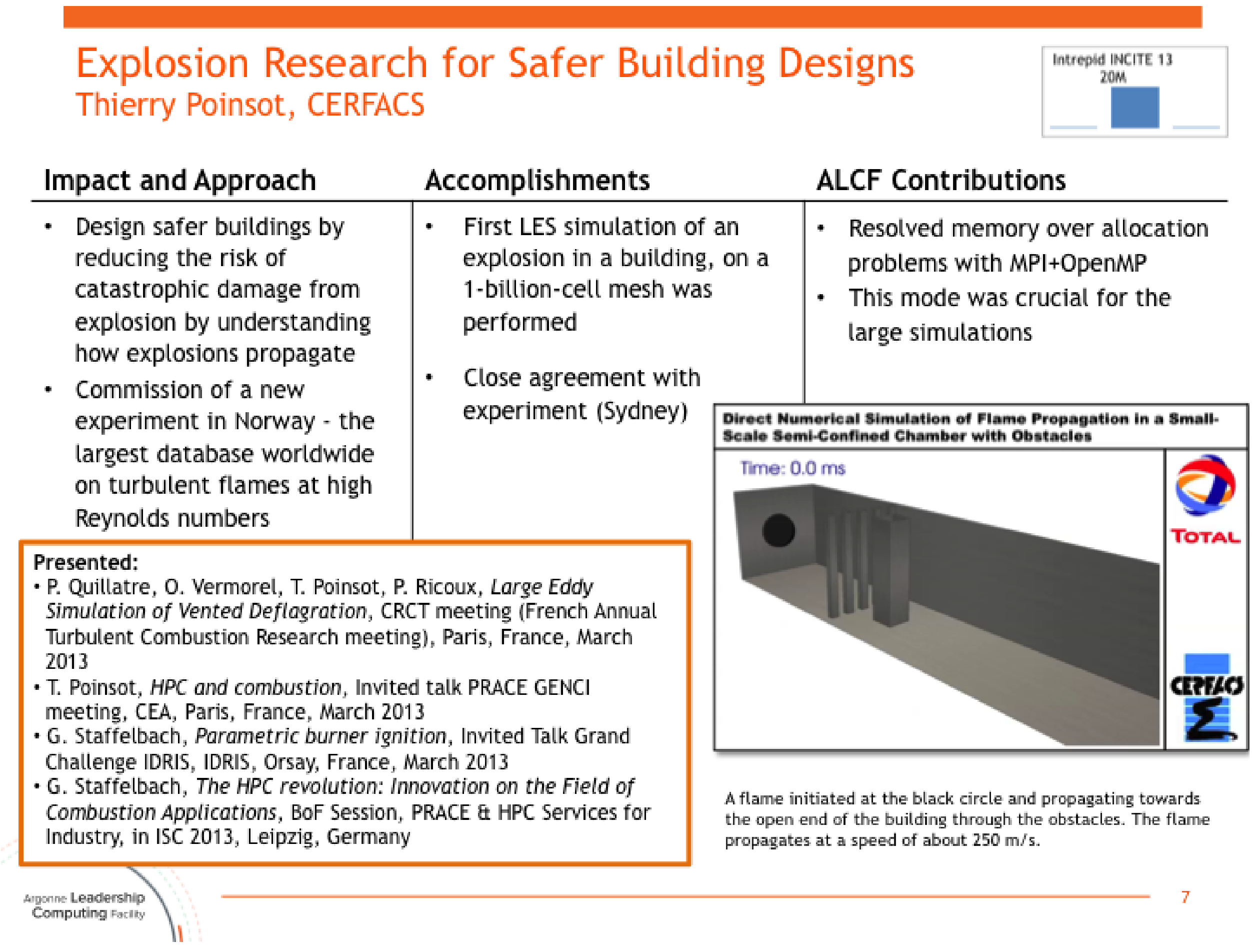

Designing Controlled Drug Delivery Systems

Subramanian Sankaranarayanan, Argonne National Laboratory

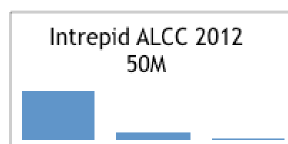

Chemists have been looking at polymers and gels for decades to see how these large chains of molecules respond to external stimuli such as heat, pH, temperature, electric fields, light, and chemical influences. The scientists seek ways to control the polymers' actions and behaviors for a wide range of applications: drug delivery, medical diagnostics, tissue engineering, electrophoresis, and enhanced oil recovery.

Certain smart polymers, for instance, rely on heat, water, and timing to execute their missions in controlled drug delivery systems. The key to employing a polymer is controlling its lower critical solution temperature (LCST). When reached, the LCST triggers a dramatic conformational change in which the polymer chain goes from an expanded coil state to a collapsed, globular one.

A model biopolymer, called thermosensitive poly(n-isopropylacrylamide) – PNIPAM or Poly NIPAM – has drawn considerable attention because its LCST nearly matches human body temperature, which is critical in drug delivery systems. If, for instance, a drug is attached to this polymer and it is given to a person, then when the body temperature goes above 98.6° F, the polymer will collapse from the coil-globule transformation and releases the drug. When body temperature returns to normal, the collapsed globule reforms as a coil.

Sankaranarayanan is studying how polymers respond to temperature, so that when the polymers are systematically heated and cooled across their LCST, the team can observe how their conformations change and respond to their environment, in particular water. His team ran simulations on Blue Gene/P using NAMD, a classical molecular dynamics code, and long trajectories for a number of different PNIPAM structures (3-8 million atoms) were studied. The goal of these simulations is to understand the atomistic origin of the coil-to-globule phase transition in PNIPAM. The team's main finding was that the ordering of water molecules plays a critical role in dictating the polymer conformation. In particular, the surrounding water forms cage-like structures, which drive the coil-to-globule phase transition in PNIPAM. This work has been submitted for publication to Physical Review Letters. Sankaranarayanan has follow-on ALCC and INCITE projects to continue work on the smart polymers on the Blue Gene/Q.

IMPACT: Developed a fundamental understanding of stimuli response of smart polymers and identified the role of solvation in inducing phase transition. Long term, this work will help design controlled drug delivery systems.

ALCF Contribution: ALCF recommended and built memory optimized version of NAMD for large-scale computation, trained postdocs to use NAMD efficiently on Blue Gene/P (load balance parameters), fixed code that parsed input file, and resolved a communication memory issue.

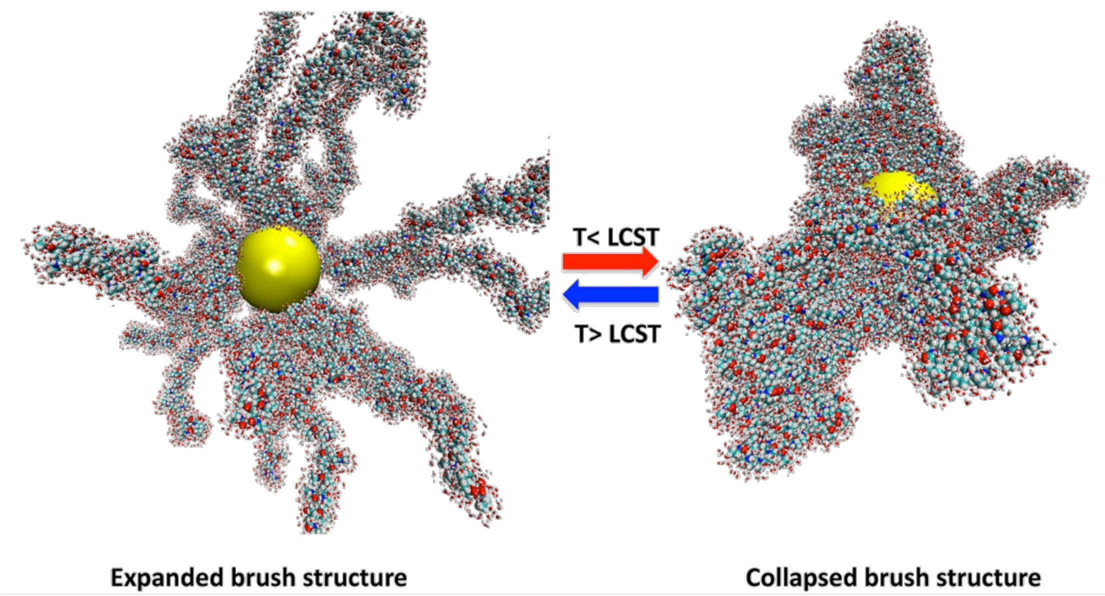

Figure 3.4: Water molecules form an ordered structure around the polymer chains. Breaking of this cage acts as a precursor to the collapse. Computational modeling can allow researchers to identify ways of controlling the conformational dynamics and design efficient drug delivery systems.

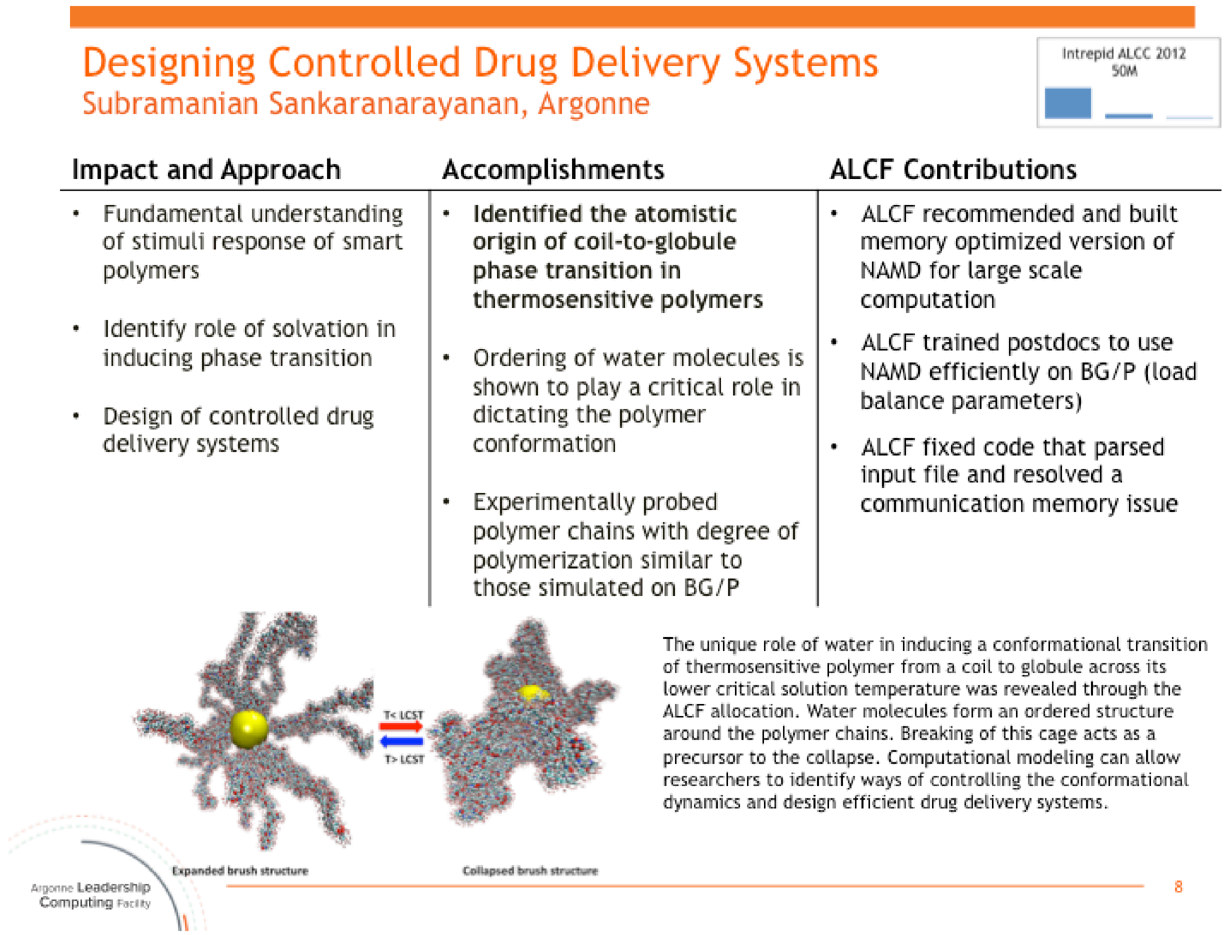

Macromolecular Folding and Aggregation

Juan de Pablo, The University of Chicago

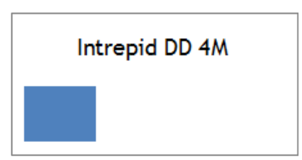

Amyloid fibrils are large structures of misfolded proteins that are strongly associated with various diseases, including Alzheimer's, Parkinson's, and type 2 diabetes. These structures are believed to arise when proteins deviate from their normal 3D structures and instead adopt misfolded states that tend to clump. Understanding how amyloid fibrils develop is critical to controlling these tissue-damaging structures and developing new therapeutic strategies.

To investigate this issue, a multi-institutional team combined experimental and computational research efforts to uncover a complete picture of the early events leading to amyloid formation. As part of the research, scientists from University of Chicago's Institute for Molecular Engineering interpreted experimental spectroscopic data by performing large-scale molecular simulations with Intrepid, the ALCF's Blue Gene/P supercomputer, and University of Chicago computing resources. The results supplied an essential model of the molecular steps involved in the reaction.

Together, the researchers located an entire step that had been missing, and whose absence had been fueling confusion. An earlier study indicated that the intermediate step was likely a floppy loop area formed by proteins, which didn't seem compatible with the tough, damaging fibril as an end result. Researchers believed that the fibrils should come from a rigid structure called a β-sheet. The new data show, however, that both structures occur as the reaction changes over time. Transient rigid β-sheets form, then morph into floppy protein loops, which finally take the form of more β-sheets. The final β-sheets bind together and stack up to form the damaging fibrils.

The researchers will continue to investigate amyloid fibrils at the ALCF as part of a 2014 INCITE allocation. With more data, the research could lead to the design of an inhibitor drug to bind to the offending protein, block the molecule, and halt the pathway's progression.

IMPACT: Target possible treatments for diabetes and other diseases, such as designing an inhibitor to interfere with the harmful pathway.

ALCF Contribution: The ALCF staff helped this project port and optimize their code in preparation of a 2014 INCITE proposal and award.

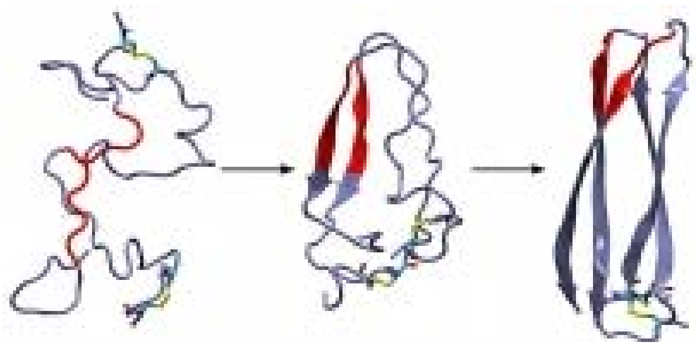

Figure 3.5: Large-scale ensemble simulations targeted at identifying key effects in the formation of defect states, which can then be eliminated or exploited. This image shows an example of a misfolding.

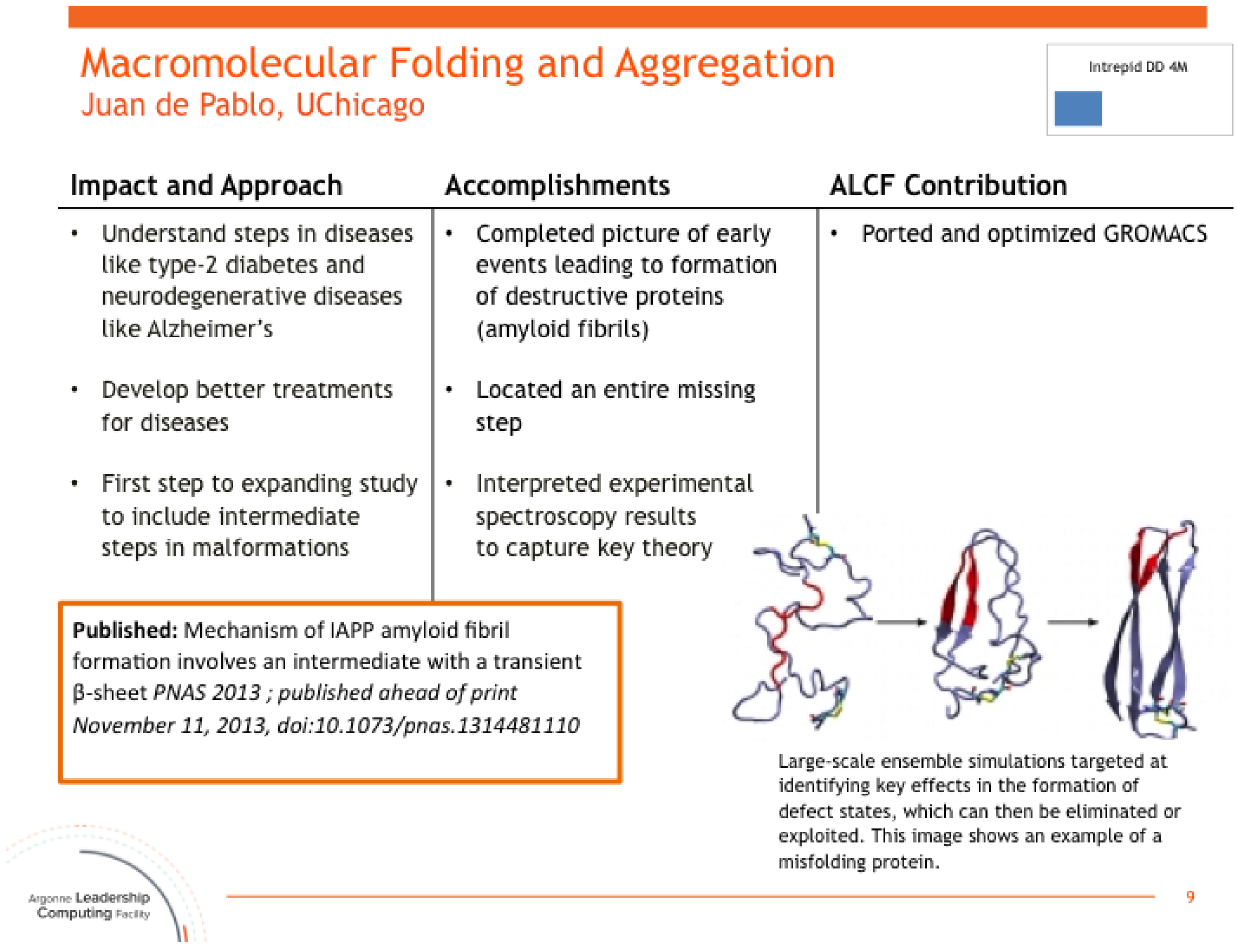

Usage of the INCITE and ALCC Hours

In total, the INCITE program allocated 2.1 billion core-hours on Mira and 736 million on Intrepid. Of the 27 INCITE projects on Mira, 17 of them used their entire allocations (or more). Of the remaining 10 with outstanding allocations, seven projects used more than 50% of their time and of those seven, four projects used more than 75% of their time. Two projects had unexpected scientific challenges using their allocations. The ALCF worked very closely with their teams but the issues were fundamental challenges with their computational models. One project lost staff critical to the project. Of the 27 projects on Intrepid, 21 used 90% or more of their allocation and of those 17 used 100+% of their time. The projects lowest in allocation usage each used 10% or less of their allocation. One had a large technical challenge that the ALCF resolved but it took most of the year. The second project was more focused on their Mira allocation than their Intrepid allocation, and the third was hurt by staffing challenges and never quite took off. A total of 2.41B core-hours were delivered to INCITE on Mira, and 821.5M were delivered to INCITE on Intrepid.

The 2012-2013 Intrepid ALCC allocation year finished on June 30, 2013. The ALCC program allocated nine projects for a total of 307 million core-hours and 271 million core-hours were used. Of these nine, five used 90% or more of their time and two used more than 80%. Finally, one project used 48% and another 5.5%. The very low use on the last project was due to their loss of funding.

The 2013-2014 ALCC year was different than previous years. While ALCC allocations were given on Intrepid for the start of the ALCC year (July 1, 2013), there were very few projects for Intrepid and those ended on December 31, 2013. However, the ALCF and ASCR decided to migrate these projects to Mira to provide them with an opportunity to continue their work and become familiar with the new systems.

The Intrepid ALCC projects had a total of 143 million core-hours and, when Intrepid was shut down halfway through the ALCC year, 34 million core-hours had been used by projects. Two projects had used the entirety of the allocation, and one had used 61% halfway through the allocation year. One project had 14% usage because of delays getting started. They will be able to use Mira for the rest of their work. The final two projects had used no time. One is awaiting experimental results before they can start, and the second had technical difficulties. The ALCF has been assisting the project team, who now believes they will be successful using Mira.

The 14 Mira ALCC projects had 1.11 billion core-hours allocated excluding the Intrepid projects that were migrated. In 2013, these projects used a total of 348.4 million core-hours. Six of the projects have used half of their allocations and are on target for straight-line usage of their time. Four projects have used between 30% and 40% of their time and one is currently running and has used 19%. Three projects have little usage. One found a bug in PETSc as they were ramping up and this has been passed on to the PETSc team for resolution. One project focused on using their Intrepid ALCC time before moving on to Mira. The third project, with no time used, has just resolved a staffing problem and the ALCF is working with them to get started.

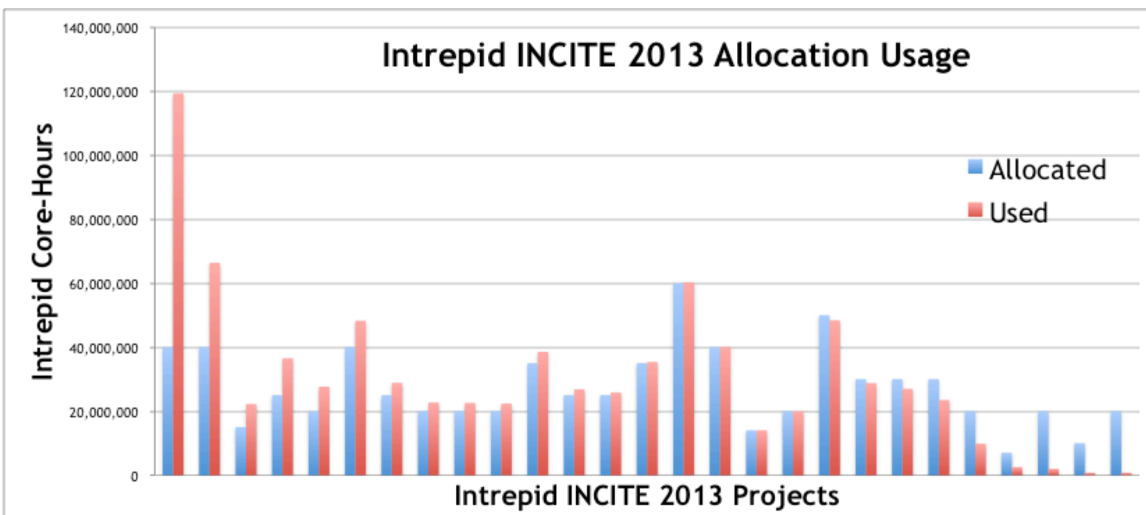

Figure 3.6: Intrepid INCITE 2013 Allocation Usage

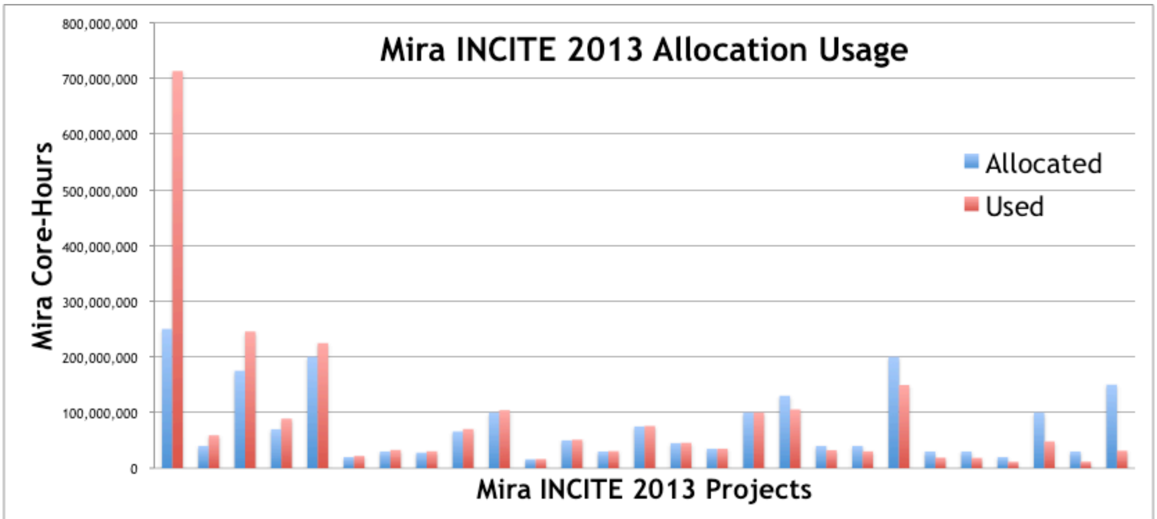

Figure 3.7: Mira INCITE 2013 Allocation Usage

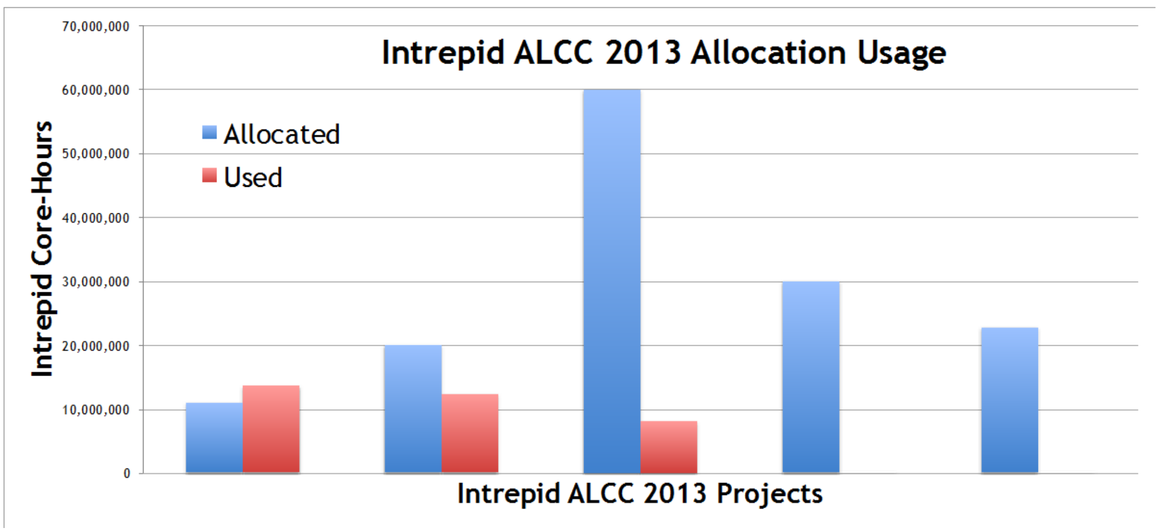

Figure 3.8: Intrepid ALCC 2013 Allocation Usage

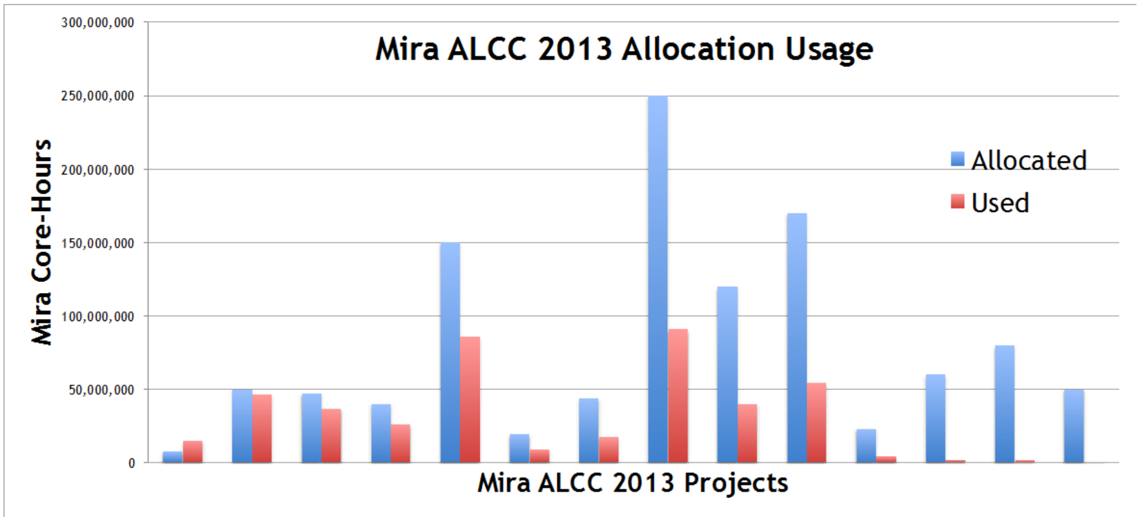

Figure 3.9: Mira ALCC 2013 Allocation Usage

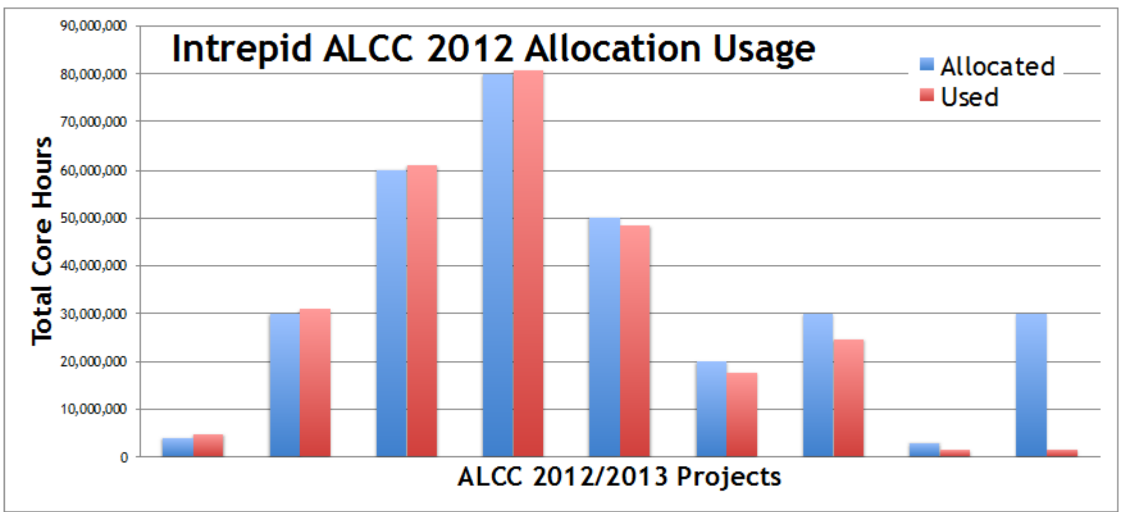

Figure 3.10: Intrepid ALCC 2012 Allocation Usage

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 program, serves the HPC community interested in testing science and applications on leadership-class resources. Projects are allocated in five categories:

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

4) Internal/support
5) Early Science

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., run scaling tests for their computational readiness) or longer-term development and testing. Additionally, in 2013, the ALCF offered limited discretionary time on Intrepid to INCITE projects with a positive balance at the end of 2012. Intrepid experienced an outage in December and the ALCF wanted to support the projects impacted by the outage.

Code support and/or development allocations support teams porting and supporting specific 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 the ALCF. For example, PARTS has supported multiple libraries and software packages. This effort has fueled multiple, successful INCITE proposals and papers.

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

The Internal/Support projects are devoted to supporting the ALCF mission. The ALCF holds back no time for division activities. All activities come out of the discretionary pool. In 2013, this category included all work related to finalizing the punch list for Mira and it included preparations for the request for proposals for ALCF-3. This category regularly includes projects that help staff support the users and maintain the system, e.g., diagnostics and testing tools and applications.

Early science activities for ALCF-2 are also from the discretionary pool. The ALCF extended access for the ALCF-2 early science projects through the end of 2013 by providing discretionary awards. This compensated for a shorter than expected early science period because of a fast acceptance and helped bridge the gap while INCITE projects that were not in the ESP were brought up to speed.

Allocations are requested through the ALCF website and reviewed by the ALCF Director, Director of Science, and Deputy Director of Science at weekly meetings. In 2013, this committee expanded to include representatives from Operations, User Services, and the Catalyst teams. The committee collects additional input from the 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 pool is under great demand and often the requested amount cannot be accommodated.

Table 3.1 shows the number of projects and total time allocated in the DD program in 2013. The nature of the DD pool supports over-allocation, but 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; instead, 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 that is available for their use. Exceptions are made for internal projects that support acceptance of new hardware or support of the users. For example, ALCF maintains a large allocation (100 million core-hours on Mira) called `crit_perf` that is a high-priority project the staff uses when trying to resolve time-sensitive performance problems for ALCF projects. In 2013, these projects used 17.7M core-hours.

Table 3.1: The total time allocated and used within the discretionary project. The total for Mira includes time for the Early Science Program after the machine went into production.

Projects	Intrepid	Mira	ESP
Allocated Core-Hours	488.6M	824.8M	2.04B
Used Core-Hours	126.1M	259.4M	591.9M

Lists of the CY 2013 DD projects, including title, PI, institution, and hours allocated, are provided in Appendix B.

The following charts provide a breakdown of the CY 2013 allocations by the standard INCITE science domains.

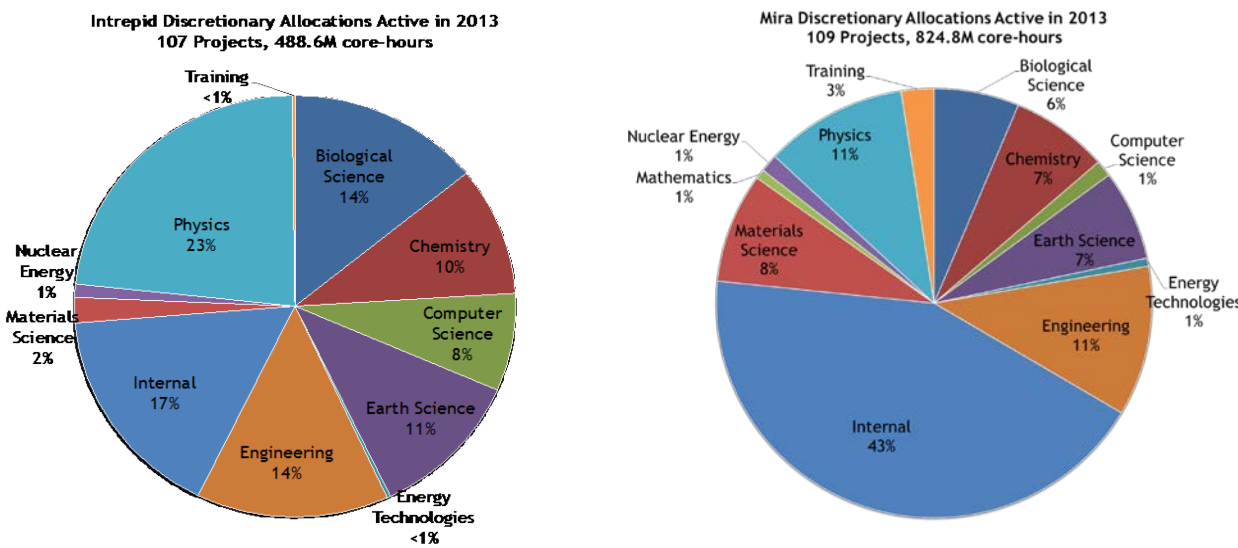

Figure 3.11: CY 2013 DD Allocation by Standard INCITE Science Domains

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 154 papers in the past year. 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 2013: 3.23 billion to INCITE, 575.0 million to ALCC, and 385.5 million to DD. In addition, once Mira was in production the Facility delivered 591.9 million core-hours for Early Science projects in 2013. Of note, the INCITE hours are a significant increase from the 804 million core-hours delivered to INCITE projects in CY 2012. The DD Reserve has been used not only to develop INCITE and ALCC proposals but also to conduct real science of strategic importance and drive development and scaling of key INCITE and ALCC science applications. The excellent ALCF support and the solid, high-performing ALCF resources have enabled the INCITE and ALCC projects to run simulations efficiently on HPC machines and achieve science goals that could not otherwise have been reached.

This page intentionally left blank.

Section 4. Innovation

Have innovations been implemented that have improved Facility operations?

ALCF Response

Listed below are the innovations and best practices carried out at the ALCF during CY 2013.

4.1 Vendor Collaborations to Improve Application Performance

4.1.1 IBM: Automatic Application Performance Collection

The ALCF and IBM have been working together to develop and deploy an automatic application performance collection library to gather MPI and hardware counter performance data for jobs running on Blue Gene/Q systems. As one of the largest and most powerful platforms for scientific computing in the country, Mira is a production resource that runs a wide variety of complex, highly optimized, scientific simulations at extremely large scale. It is critical that any library automatically linked into user software in this environment must be robust and scalable, and its overhead must be low. In addition, this library's presence must be transparent to the user.

After testing and evaluating an IBM prototype automatic performance library on ALCF systems, numerous modifications were introduced to the source code to improve robustness, reduce overhead, and allow interoperation with other data collection libraries used at the ALCF. Additionally, the process for outputting the collected performance data to an ALCF database was re-architected to make it more robust and secure. Scripts were developed to integrate the performance collection library into the default ALCF compilation environment and a beta version of the library is being tested on ALCF's testbed resource, Vesta.

4.1.2 Pathscale: OpenMP Support in Blue Gene Clang Compilers

Partnering with Pathscale, the producers of the EKOPath compiler suite, ALCF brought an OpenMP capability to the bgclang compiler. ALCF staff created a customized OpenMP runtime library for the Blue Gene/Q derived from the recently released Intel runtime, and provided bgclang with full support for the OpenMP 3.1 standard. ALCF now offers its users a compiler supporting both the latest C++ language standards and also multicore parallelism via OpenMP directives. Because exploiting multicore parallelism is paramount to Blue Gene/Q users, the availability of OpenMP in bgclang has transformed bgclang from an interesting experiment to a compiler that can be used by a wide array of HPC applications.

By bringing OpenMP support to bgclang this partnership has also impacted the wider LLVM/Clang community. The work done by Pathscale has aided the process of bringing Intel's OpenMP contributions, which were based on an older version of Clang compiler, into the upstream Clang project. In addition, because Pathscale will be able to fold this work into its own compiler products, Pathscale's customers, which include several DOE laboratories, will benefit from enhanced OpenMP support in future Pathscale compiler releases.

4.2 Operations Improvements

4.2.1 Machine Time Overlay Graphic

ALCF has developed a new graphical tool that allows staff to assess the operational state of the machine over a period of time with a fair degree of detail. It is currently a prototype, but has already identified a number of operational issues including a bug in the scheduler, a user who was running a large number of one-minute jobs that were all failing, and a situation where the max hold setting on the queue was hurting utilization. Currently, the graphic is static as a .PNG file. One of ALCF's future goals is to make this a scalable vector graphics (SVG) file, which will allow the text to scale when zooming in, and natively supports layers so information can be turned on and off as needed to reduce clutter in the graph. An example is shown below:

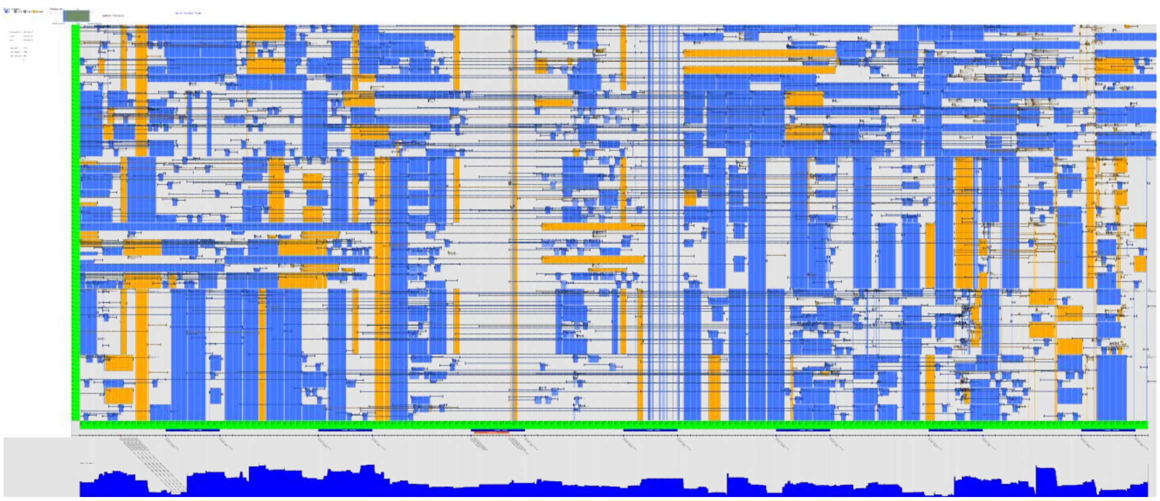

Figure 4.1: Machine Time Overlay Graphic

The Y-axis represents location within the machine. Scheduling is done at the midplane level (there are two midplanes per rack) and there are 96 midplanes in Mira. The bottom represents R00-M0 and the top represents R2F-M1. The X-axis is time. The colors are definable. Red is an outage, either scheduled or unscheduled. Blue is a job that had a zero exit code. Two particular problems are being tracked that have characteristic exit codes; those are two shades of green. Orange is all other jobs. Grey is no job running at all. Though it is not discernable at this scale, there is job data (e.g., job id, number of nodes) embedded in the color block representing that job. The bars indicate the time the job was selected to run on a particular resource to the time it was scheduled to end.

4.2.2 Continuous Improvement of Operational Assessment Process

ALCF has reported on its operational assessment process in the past. In weekly meetings, the fate of every job that ran on the system is assessed. It is assumed that job failures occur due to system issues unless it can be proved otherwise. ALCF continues to improve this process. The Machine Time Overlay discussed above is a tool developed to assist with investigating failures. Additionally, a full suite of reports and graphics is used to help assess the state of the system and then annotated with interrupt identifiers. ALCF also has an internal wiki where details of each of these events are kept. Essentially, every core second on the machine is accounted for and classified as available or not, with an explanation for unavailable time and available time, utilized or not, and a determination of the fate of the process utilizing the time. The figure below shows an example of an annotated availability graphic:

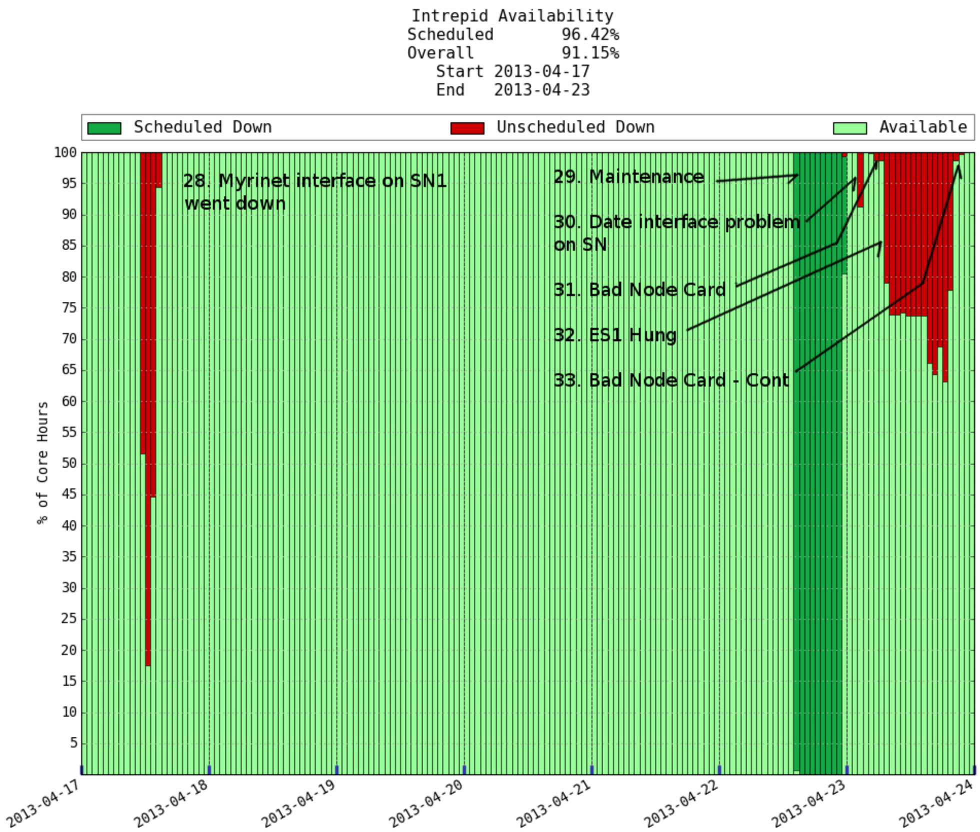

Figure 4.2: Weekly Availability Chart with Annotations

4.2.3 Scheduling

ALCF has a scheduling committee composed of members from the various teams that support the facility's user base. This committee is dedicated to assessing and improving scheduling related issues on Mira and its testing and development systems with the goal of improving job throughput. The machine time overlay graphic has proven to be a valuable tool for analyzing utilization on these resources and for identifying areas for enhancement.

In the second quarter, ALCF identified two opportunities for improving job throughput on Mira. Due to network wiring constraints, one-rack jobs blocked the ability of jobs to run on the adjacent rack. A similar issue was identified for four-rack jobs. In the limit of a pure workload consisting of one-rack and four-rack jobs, 50% of Mira would be idle. The following two modifications to the Mira partitions resolved this issue:

1. Adjusting the D-dimension of the Blue Gene/Q network for one-rack partition from torus to mesh.

2. Creating four-rack partitions with two different shapes.

With these two modifications, one hundred percent utilization can be achieved, in principle, in the limit of pure workload of one-rack and four-rack jobs. Before implementing these changes, the impact to the user was assessed by running tests via the Harness test environment. Analysis showed that there was minimal impact to users, and thus these partition modifications were implemented during a regularly scheduled maintenance cycle.

4.2.3.1 Changing Block Geometries for Machine Packing

Early in a machine's life, smaller job sizes, and in particular those using 1,024 nodes and 4,096 nodes, are more heavily run as users adjust their codes and processes to the new machine. Due to restrictions in Mira's wiring, however, running full-torus, homogenous geometry versions of these blocks meant that the entire machine could not be packed with either of these job sizes, and that at a minimum, other smaller sizes would be needed to fill in. ALCF adjusted the geometry of 4,096 node blocks allowing this type of job to pack the entire machine, while minimizing impacts on capability-sized jobs (8,192 nodes or greater) and preserving a full-torus on the block. For the 1,024 node jobs, ALCF disabled the torus in the one dimension that made the wiring problematic, something that was new to the Blue Gene/Q. This resulted in a minimal performance impact, while allowing these jobs to pack across the entire machine. This change increases the maximum usage for both of these sizes from fifty percent of Mira to one hundred percent.

4.2.3.2 Ensemble and Subblock Jobs on Mira

To exploit new functionality available on the Blue Gene/Q for ensemble jobs, users were provided with new interfaces that allow them a very high degree of flexibility when subdividing their allocated resources, thereby facilitating more efficient bundling of user workloads. This resulted in better throughput for the user, and enabled ALCF to better reduce draining costs. This allows users to set up multiple sub-midplane jobs on the same midplane, down to a single node. ALCF has also worked to educate users on the use of these features through interactions with their Catalysts as well as presentations in its workshops.

4.2.4 Networking Improvements

E600i 10GbE Bandwidth

The existing Force10 E600i distribution switch at the ISSF was expanded with two 10 GbE expansion modules in preparation for Mira's production state and to support connectivity between the Theory and Computing Sciences (TCS) Building and ISSF. This expansion provided 10 GbE ports for both HPSS and data transfer nodes (DTNs), with the latter being used to move large amounts of user project data to the new Mira environment. This expansion also addressed the over-subscription problems on the older E600i chassis that would have greatly diminished the capability of the switch to move data.

DWDM

To facilitate user data migration and to provide for disaster preparedness, the ALCF installed dedicated DWDM connectivity between the two supercomputing facilities at the TCS and the ISSF. Notably, this technology allowed for the economical use of the remaining pair of dark fiber strands connecting the two sites. This DWDM connectivity currently multiplexes 12x8 Gbps of Fibre Channel waves and 4x10 Gbps Ethernet waves. The Ethernet is available for both administrative and user data, allowing users access to the new GridFTP DTNs at the TCS for data transfer. The Fibre Channel is being used to bridge the critical storage fabrics between the two sites for geographical redundancy.

4.3 Business Intelligence Improvements

4.3.1 Improved Data Management for Facility Data

Data is generated throughout the facility, for example job data, error data, usage data, and user and project data. The ALCF has different teams and groups that are responsible for the generation, storage, and management of different pieces of data. Each of these teams were developing and generating their own reports, pulling the data directly from their individual sources. This led to confusion as to what data was available at the facility, where that data was located, and how it was being generated, i.e., if there were calculations, they weren't transparent to others at the facility.

To address this challenge, the ALCF has implemented a data warehouse to provide a single, canonical source of facility data that will be the source for all production reports. The ultimate goal is to have all data pushed to this warehouse from all sources. This will provide greater transparency to management and staff as to what data is available and how it was generated. This has resulted in more consistent reporting, reflecting a single view of the data. In addition, this change has brought to light data that was previously not utilized by management to do broader analysis.

All data necessary for the current production reports is now in the warehouse. ALCF is in the process of revising all production reports to only use data from the warehouse.

4.3.2 Reduced Fragility of Facility Reporting

The ALCF was initially stood up with a lean staff and developed the scripts to generate reports to stakeholders on demand and in an *ad hoc* fashion. As a result, business policies were embedded in those report-generating scripts and these policies were implemented and interpreted differently by different scriptwriters. That could result in reports against the same dataset generating different values. To resolve this issue the ALCF now identifies, applies, and manages all business policy decisions in a single place: the Extract-Transform-Load (ETL) for the data warehouse. This exposes these policies to all stakeholders and allows the ALCF to periodically review and evaluate them. This protects data and reports from unintentional changes and repercussions. All of this results in more consistent reports.

ALCF is in the process of removing all business policy from existing report-generating scripts. Going forward, all ALCF production reports will use data from the warehouse, incorporating business policies from the ETL.

4.4 Facility Improvements

ALCF improved the system that directly cools Mira, Cetus, and Vesta by installing a second chilled-water process loop, which became operational in July 2013. The new loop was installed and commissioned without affecting the users' access to the systems. The original loop stretches from the Argonne Chilled Water Plant building to the TCS building. The new loop has an added connection to a separate, alternative Argonne campus cooling loop to provide redundancy. Additionally, there are built-in options to extend the second loop for additional water-cooled racks in the future. Also, the new loop improves the water quality to the Blue Gene/Q racks by utilizing a finer filtration (100 micron) system and a pure water (reverse osmosis) input system.

4.5 User Services and Outreach Improvements

4.5.1 Using Checklist to Communicate with Users

ALCF User Services implemented a best practice similar to many of the web resellers and parcel tracking systems. Instead of sending e-mails to PIs and users with extensive prose, messages were broken down into checklists — an idea inspired by Atul Gawande's book, *The Checklist Manifesto*.

Figure 4.3 is an example of an e-mail sent to PIs during the startup of INCITE 2012. The format is paragraphs of prose separated by major titles.

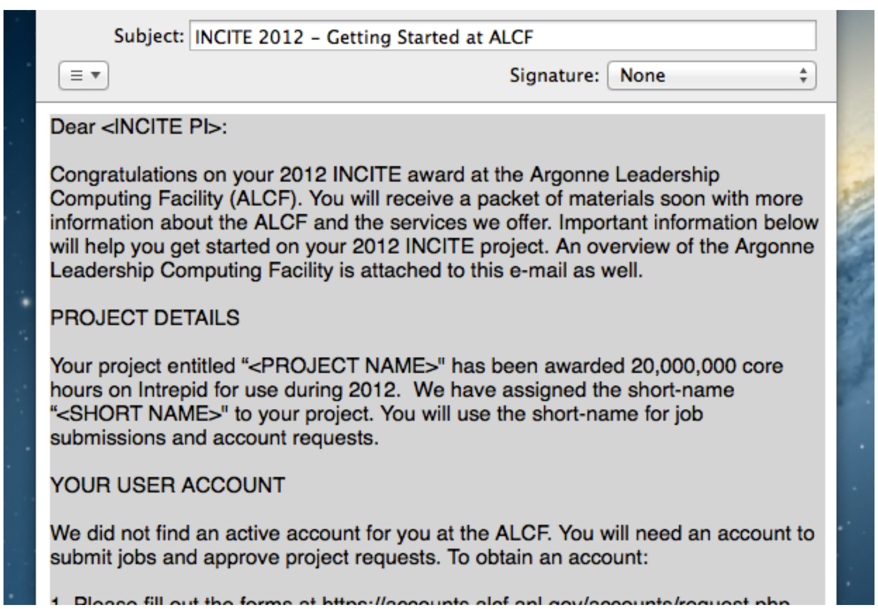

Subject: INCITE 2012 – Getting Started at ALCF

Signature: None

Dear <INCITE PI>:

Congratulations on your 2012 INCITE award at the Argonne Leadership Computing Facility (ALCF). You will receive a packet of materials soon with more information about the ALCF and the services we offer. Important information below will help you get started on your 2012 INCITE project. An overview of the Argonne Leadership Computing Facility is attached to this e-mail as well.

PROJECT DETAILS

Your project entitled "<PROJECT NAME>" has been awarded 20,000,000 core hours on Intrepid for use during 2012. We have assigned the short-name "<SHORT NAME>" to your project. You will use the short-name for job submissions and account requests.

YOUR USER ACCOUNT

We did not find an active account for you at the ALCF. You will need an account to submit jobs and approve project requests. To obtain an account:

Figure 4.3: Example Message from 2012 INCITE Communications

In Figure 4.4, ALCF implemented a checklist. It draws attention to the most salient things that a PI needs to do, and provides positive feedback on what they have already completed.

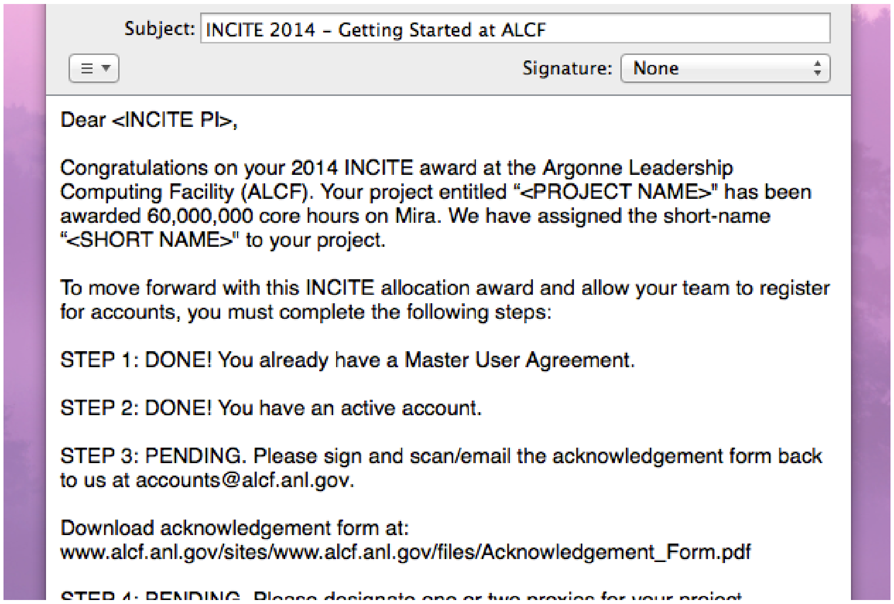

Subject: INCITE 2014 – Getting Started at ALCF

Signature: None

Dear <INCITE PI>,

Congratulations on your 2014 INCITE award at the Argonne Leadership Computing Facility (ALCF). Your project entitled "<PROJECT NAME>" has been awarded 60,000,000 core hours on Mira. We have assigned the short-name "<SHORT NAME>" to your project.

To move forward with this INCITE allocation award and allow your team to register for accounts, you must complete the following steps:

STEP 1: DONE! You already have a Master User Agreement.

STEP 2: DONE! You have an active account.

STEP 3: PENDING. Please sign and scan/email the acknowledgement form back to us at accounts@alcf.anl.gov.

Download acknowledgement form at:
www.alcf.anl.gov/sites/www.alcf.anl.gov/files/Acknowledgement_Form.pdf

STEP 4: PENDING. Please designate one or two proxies for your project

Figure 4.4: Example Message from 2013 INCITE Communications

ALCF received informal feedback from users about this change, which was positive. While there is not a direct measure, User Services staff noted an increase in response rate to the key e-mails where this technique was employed.

4.5.2 Using Small-Scale Videoconferences to Replace a Getting Started Workshop

Inspired by the extensive research done by the Access Grid community of the late 1990s, User Services outreach staff used beta software from ESnet to conduct small-group workshops. The idea was to create a virtual environment that mimicked a small group working around a table.

This environment initially surprised users. They were expecting a unidirectional webinar with little to no interaction. However, when ALCF explained how this would work, users interacted well with the Catalyst and User Services presenter.

Figure 4.5 shows how the videoconference was set up using side-by-side HD displays. On one display, the presenter appears in one video stream on the left side, surrounded by the other participant video streams, each appearing as a separate window. The other HD display provides the presenter with a view of the shared slide presentation.

Figure 4.5: Arrangement of Video Feeds at ALCF Getting Started Workshop

Figure 4.6 shows the arrangement of the presenter's space. ALCF User Services and Catalyst staff communicated with the users via camera or scan-converted, PowerPoint video feed.

Figure 4.6: Arrangement of Local Video Conference Room for Getting Started Workshop

The Catalyst could share a PowerPoint, do a live demo, or ask the participants to share their individual screen on the main screen in order to "group troubleshoot" an issue. Audio was open and available during the entire meeting.

4.6 Administration Improvements

Participation in DOE HPC Financing Improvement Team

In conjunction with participants from ASCR, ALCF, NERSC, and OLCF, ALCF remains an ongoing coordinator and participant in the Financing Improvement Team best practices group. All participated and shared information in order to improve leasing and negotiating practices with the end goal of obtaining the very best rates possible for all the facilities. General areas of interest included lease options, requests for proposals, evaluation templates, processes, and best practices.

Within the area of leases, a number of options continued to be discussed and evaluated. These areas include the following:

- Buyout option – the ability to pay off the lease principal at any point in time to save subsequent interest. A mechanism for pricing the option was developed and used for evaluating different lease options for ALCF-2.
- Buydown option – the ability to pay ahead or buy down lease principal at any point in time to save interest without having to renegotiate the lease.
- Discount point payment – the ability to pay a fee to be traded for a lower interest rate on a lease. A mechanism for pricing the option was developed and used for evaluating different lease options for ALCF-2.
- Flexible payment schedule – the ability to change payment schedules and amounts depending on revised funding profiles. This could include the ability to skip a lease payment or two without defaulting on the lease.
- Adding on to the lease – the ability to add equipment to an existing lease so a new agreement need not be negotiated.

Templates for evaluating and comparing different leases were enhanced and refined. One example is the template providing the capability to measure the cost and benefit of paying an upfront fee to reduce the interest rate paid on the lease. Vendor lists continue to be updated and shared with a composite list provided to members of the group.
Shared best practices include specifics such as:

- Saving time and lining up competitors through pre-approval by setting up a "master lease" prior to the specific lease proposal.
- Separating equipment purchases from the financing (which is competitively bid).
- Obtaining a financing package from the equipment provider to serve as a competitive starting point to negotiate lease rates (provides lower risk since a maximum rate is locked at this point).
- Using consistent methods and market data for evaluating interest rate ranges and associated risks.
- Shared analysis tools for pricing and comparing diverse proposals.

A prescient best practice arrived at in 2012 involved discussion of pricing for various lease options. After considering many methods and much discussion, the group opted for lease bids and pricing based on Treasury rates plus a spread. An alternative option using Libor (London Interbank Offer Rate) was ruled out on the basis that the index was self-reported by banks and could be prone to unpredictable swings as well as potential manipulation. Subsequent to the group's decision, it was discovered that banks have manipulated the rate resulting in indictments, sanctions, penalties, and lawsuits.

Argonne also broke new ground by providing a detailed financial information package to prospective lenders comprised of many materials often found in due diligence processes. The package provided information not available in the past in the hope it would increase confidence in ALCF as a borrower and result in lower lease rates. The lease rate for ALCF-2 came in lower than anticipated; it is believed this is due at least in part to the financial information package.

During the ALCF-2 financing process, the team evaluated options from both the hardware supplier and multiple independent financiers. The ALCF-2 financing team discovered and exploited cost saving synergies available from the hardware supplier.

Some white papers are currently being drafted, including:

- A method for estimating future interest rates based on market data.
- An assessment of potential changes in lease rules.

The Financing Improvement Team best practices group continues to share information and seeks to continually improve practices among the participants.

Conclusion

ALCF has identified innovations and best practices that have improved operations in several areas the past year.

- The center has found ways to help application developers improve application performance through automatic profiling of applications and by supporting the OpenMP 3.1 programming model in the open source CLANG compilers.
- On the operations side, ALCF had identified a number of improvements including development of a time overlay graphic that provides the state of the machine at any given time. Study of this static graph has helped identify cause of machine failures and more importantly improve scheduling algorithms. Networking improvements have been made to improve data migration between Intrepid (housed at ISSF) and Mira (housed in the TCS building).
- In the Business Intelligence area ALCF has implemented a data warehouse to provide a single canonical source of data for business reports and with all business policies implemented through the ETL in the data warehouse to ensure consistent reports.
- ALCF added a second chilled-water process loop.
- ALCF has improved user support through use of a checklist to communicate with users and adding small-scale videoconferences to onsite workshops.
- Collaboration with other facilities, including participation in a DOE-sponsored HPC Financing Improvement Team.

Section 5. Risk Management

Is the Facility effectively managing risk?

ALCF Response

The 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. In preparation for the start of the ALCF-3 project, the RMP was significantly revised to create a document based on the Argonne Office of Project Management (OPM) Risk Management Plan. This allows ALCF to leverage the expertise OPM has developed in risk management; it also results in an ALCF Risk Management Plan that the Risk Manager, who is matrixed from OPM, can work with using a process similar to what is used with RMPs in other Argonne divisions. Risk management is a part of the ALCF culture, and the RMP processes have been incorporated into both its normal operations and all projects, such as the newly completed ALCF-2 project, the acquisition of a 10 PF Blue Gene/Q system called Mira, as well as the recently launched ALCF-3 project. 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. The ALCF currently has 45 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 their mitigations described in more detail, along with new and retired risks, as well as the major risks that will be tracked in 2014.

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

The ALCF uses the documented risk management processes, first implemented in June 2006 and outlined in its RMP, for both operations and project risk management. The ALCF reviews and updates its RMP annually, and it is also updated during the year if changes at the ALCF necessitate an update (e.g., changes to Argonne risk management), as well as to incorporate new risk management techniques when 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 the 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.

The ALCF currently has 45 open risks in its operations risk register. These risks include general facility risks, such as funding uncertainties, staffing issues, and safety concerns, as well as specific risks such as system component failures, availability of resources, and cost of electricity. Risk mitigation costs on the project side are developed using a bottom-up cost analysis, then used to set the contingency pool utilizing the OPRA integration with Primavera. On the operations side, the costs are estimated by subject matter experts and used to inform management reserves.

Completion of the ALCF-2 Project

ALCF keeps project risks and operation risks in separate registers within the same tool. This makes it possible to build reports easily, as well as manage similar risks that have different impacts between operations and the project. With the CD-4 approval of the ALCF-2 project in January 2013, four risks were transferred from the project risk register to the operations risk register. These risks were updated to reflect operational impacts during the ALCF-2 Transition to Operations period and were opened in the register during CY 2013.

Initiation of the ALCF-3 Project

A new project, ALCF-3, to procure and deploy the next ALCF supercomputer, was initiated in 2013. A project risk register was implemented and a set of detailed risks was created. These risks are not included in the risk numbers covered in this document and are not discussed further.

ALCF 2013 Risk Management Plan

In 2013, the ALCF RMP underwent significant revisions to create a document based on the Argonne Office of Project Management (OPM) RMP, tailored for ALCF and ALCF upgrade project use. This revised RMP went into effect in August. Prior to development of the OPM RMP, ALCF utilized an ALCF-developed RMP that was used for ALCF-1, ALCF-2, and ALCF operations into 2013. The August 2013 RMP retains the major elements of earlier ALCF RMP's and adds information on specific risk management tools that are used by ALCF, organizational structure and responsibilities, related project documents, and significant detail about qualitative and quantitative risk evaluation. Risk reviews for both project evolution and steady-state operations are also described.

Conversion from 3x3 Risk Matrix to a 5x5 Risk Matrix

As part of revising the risk management processes, ALCF moved to the Argonne site-wide risk matrix, shown in Table 5.1. ALCF management, staff, and project teams use it during the risk identification and risk scoring steps, and it is modeled in the OPRA tool for use while analyzing risks. This is a refinement of the 3x3 matrix that had been used prior to adoption of the 2013 RMP and is consistent with the Argonne site-wide RMP. Risks in the ALCF risk register that had been scored with the 3x3 matrix have been rescored so as to make them consistent with the 5x5 matrix while keeping the level of risk (e.g., high, low, moderate) the same. This rescoring has not been counted as a recharacterization of a risk.

Table 5.1: 5x5 Risk Assessment Matrix

CONSEQUENCES				INCREASING PROBABILITY →				
Technical Scope	Schedule	Cost		Risk Event Almost Never Occurs <10%	Risk Event Rarely Occurs 10% - 25%	Risk Event Occurs On Occasion 26% - 74%	Risk Event Occurs Often 75% - 90%	Risk Event Almost Always Occurs >90%
				VL	L	M	H	VH
Very High Impact	> 3 months	> 2 million	VH	5	10	15	20	25
High Impact	2 - 3 months	1.0 - 2.0 million	H	4	8	12	16	20
Moderate Impact	1 - 2 months	0.5 - 1.0 million	M	3	6	9	12	15
Low Impact	0.5 - 1 month	0.1 - 0.5 million	L	2	4	6	8	10
Very Low Impact	< 0.5 months	< 0.1 million	VL	1	2	3	4	5

INCREASING IMPACT ↑

Risk Management in Day-to-Day Operations

Beyond the formal monthly and individual risk meetings, the 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, the risks are not added to the registry. Otherwise, if new significant risks are identified during the activity planning, they 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 (WPC) and Job Hazard Questionnaires (JHQ). Both of these are used to manage risks for activities where safety is a potential concern. JHQ are used for all staff and all contractors, and cover all work, both routine and non-routine. WPC 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 for 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 proposal risks (risk of too few proposals, risk of a lack of diversity across science domains, 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.

Major Risks Tracked for the Review Year

Since Q4 of FY 2010, the ALCF has had several eventful years. This was due to 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, the ALCF was monitoring, and continues to monitor, a large number of major risks for the facility. These risks are described below in Table 5.2. All risk ratings shown are post-mitigation ratings. There were 19 major operations risks tracked for CY 2013, one with a risk rating of *High*, and 18 with a risk rating of *Moderate.* Of these 19 risks, five were encountered and managed, and six were retired with Mira beginning production operations and Intrepid operations being terminated at the end of 2013. The risks are color coded in the following way to assist with reading the table:

- Red risks were encountered and remain moderate or high risks (with the exception of risk 1060, which is green because it was also retired).
- Orange risks were not encountered but remain moderate or high risks.
- Green risks are now retired.

Table 5.2: Major Risks Tracked for CY 2013

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 for both this year and the out years as the facility begins 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.
1049	Staff retention	Yes	Mod	Between the salary freeze and budget concerns at Argonne and the growth in high paying industry jobs for system administrators and programmers with HPC expertise, the ALCF has lost several staff members during CY 2013. This remains a concern.

Table 5.2: (Cont.)

ID	Title	Encountered	Rating	Notes
31	Facility power interruptions	Yes	Mod	A transformer failure at the ISSF along with UPS battery issue caused outages for Intrepid. Multiple breaker failures in the TCS data center caused Mira outages. Intrepid is decommissioned but root cause of breaker failures for TCS has not been determined. This remains a concern.
1018	INCITE users are not provided adequate support by ALCF	No	Mod	INCITE teams are solicited periodically for feedback on the support they receive. ALCF staff manage INCITE teams expectations for support.
1050	Insufficient disk space to support science needs	No	Mod	Addressed through quotas, purchase of additional disks, and moving disk data to tape as needed.
1054	Catastrophic failure of home file system	No	Mod	Mitigate the potential impact by using several methods of backing up and duplicating files.
1056	System stability issues due to upgrades	No	Mod	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.
1067	Total project count grows beyond staffing levels	No	Mod	This is an ongoing risk that is managed by careful project selection and educating projects on realistic support expectations.
1076	If the ISSF is decommissioned, ALCF will not have an appropriate location to host disaster recovery equipment.	No	Mod	Argonne management has committed to providing ALCF with space if ISSF is decommissioned. Alternative locations both onsite and offsite are being explored.
1085	Diagnostic suite and utilities fail to detect hardware problems	No	Mod	Addressing through elevating the issue with the vendor to accelerate a vendor response.
1091	Injury to workers, overall safety of the division	No	Mod	Promote safety culture at all levels of the division.
1099	INCITE does not use all allocated core-hours	No	Mod	Generally occurs with only a few projects. Addressed through catalyst/project communications.
1052	The BG/P (Intrepid) service node Is a single point of failure	No	Retired	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.
1060	Intrepid system stability issues	Yes	Retired	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.
1061	System performance issues	No	Retired	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.

Table 5.2: (Cont.)

ID	Title	Encountered	Rating	Notes
1071	Priority 1 and 2 user documentation unavailable for Mira production	No	Retired	Documentation was completed within the planned schedule, ending the risk.
1075	Not enough Early Science time	No	Retired	Mira was accepted early enough to provide three months of dedicated Early Science access. In addition, Early Science projects were allowed to run through the end of CY 2013. Risk was closed with the end of Early Science time on Dec. 31, 2013.
1098	Unable to recruit enough projects for Mira production in 2014	No	Retired	Demand for INCITE 2014 hours remained high, all available INCITE 2014 hours have been allocated, and this risk has been retired.

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.

1. Funding/Budget Uncertainties

1059: Funding/Budget uncertainties	
Risk Owner	Michael Papka
Probability	High
Impact	Schedule: Low; Cost: Negligible; Tech.: High
Risk Rating	High
Primary Management Strategies	Implement austerity measures. Work closely with DOE sponsors to manage expectations and scope.
Triggers	Budget information from DOE; information from DOE indicating a likely extended CR

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 the ALCF from receiving planned funding. These scenarios could result in the inability to pay leases, contracts, and staff, as well as the inability to deploy future machines.

Evaluation

During the past year, the Funding/Budget uncertainties risk was the ALCF’s highest risk, and it was also one of the risks that was encountered. The ALCF was supposed to be in a growth phase. A reduction of incoming funds carried a high impact.

Management

In conjunction with the DOE-ASCR Budget Deep Dive, ALCF prepared for a full-year CR and reduced budget scenarios. To assure 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.

Funds in FY 2014 were reduced from the original plan of record, and the ALCF continues to closely monitor budget information for FY 2015 and beyond in case of a further reduction in funds from the plan of record. Austerity measures remain in place and may be increased depending on the budget through FY 2014.

2. Staffing Challenges and 3. Staff Retention

25: Staffing challenges; 1049 Staff retention	
Risk Owner	Michael Papka
Probability	Moderate
Impact	Schedule: Mod; Cost: Negligible; Tech.: Mod
Risk Rating	Moderate
Primary Management Strategies	Prioritized staffing needs, re-planned work, and re-tasked staff.
Triggers	Staff resignation; lack of response to job postings; staff showing signs of burnout

Description

ALCF was, and remains, in a time of necessary growth, as it staffs up to operate Mira and prepare for ALCF-3. An aggressive staff ramp up was originally planned for FY 2010 through FY 2012. This was extended due to budget reductions. ALCF risk evaluation identified two key risks associated with this ramp up, and both occurred in FY 2013 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. The risks are:

- 25: Challenges for hiring new staff.
- 1049: Unable to retain staff due to heavy workloads and the increased demand for staff with compute expertise.

Evaluation

The recovering economy is opening more jobs in the industry for ALCF staff. In the past year, as a result, three ALCF staff left for higher paying jobs in industry. Ten new staff were added during CY 2013, for a net gain of +7 ALCF staff for the year. Thus the hiring rate has been able to keep pace with the attrition and add priority new hires, but because the facility was already significantly understaffed, the existing staff continues to be overworked.

Management

It can be very challenging to hire experienced HPC staff. Because of this, the ALCF risk management team started several years ago preparing to execute mitigations prior to these risks occurring. When they occurred, the ALCF was able to successfully continue supporting existing projects and continue to deploy the new system even while understaffed. The ALCF has continued to use mitigations to manage this risk over the past year. Staff hires were prioritized, and once the austerity measures were implemented, top priority staff (replacement staff and key staff for the ALCF-2 project) were hired first. The ALCF continues to re-plan work as possible, delaying both planned improvements and lower priority work. Staff have continued to be re-tasked, dropping lower priority tasks, expanding job descriptions, and, where possible, sharing staff with other divisions.

With careful and judicious management of this risk, the ALCF has successfully run the facility, and completed the deployment of Mira and the two Test & Development systems, Cetus and Vesta. The facility continues to remain understaffed. These risks will remain a concern for the next few years, and the ALCF will continue to carefully manage them.

4. Intrepid System Stability Issues

1060: Intrepid system stability issues	
Risk Owner	William Allcock
Probability	Moderate
Impact	Schedule: Mod; Cost: Mod; Tech: Low
Risk Rating	Moderate
Primary Management Strategies	Automatic monitoring for fault tracking. Weekly analysis of small issues to determine whether larger ones are developing. Work closely with vendors when problems appear. Dynamic mitigation development.
Triggers	Higher than expected hardware failure rates; catastrophic failures; early indicators of an upcoming significant failure; planned activities that increase the probability of circumstances known to trigger a critical failure.

Description

Risk 1060 is a general risk that is used to track the health of the system and manage stability problems as they are encountered. One such problem was encountered late in CY 2013 when a catastrophic service card failure occurred. With the decommissioning of Intrepid at the end of 2013, this risk has been retired.

Evaluation

As described in the CY 2012 ALCF OAR, working with IBM, the facility discovered a design flaw in the Blue Gene/P service card that could cause a catastrophic electrical failure (burn a hole through the card, short to ground, and send a power surge through a mid plane) in certain "idle" states, which often occur during maintenance. Because the Blue Gene/P family was planned for End-Of-Life on Dec. 31, 2013, IBM chose not to redesign the card and instead agreed to replace any hardware that was damaged during the events. See the CY 2012 OAR for more details on the IBM failure analysis.

Management

As discussed in the CY 2012 ALCF OAR, IBM and ALCF reached an agreement by which IBM gave ALCF credit for the purchase of future Blue Gene/P spare parts equal to the value of the damaged parts resulting from any Blue Gene/P service card failures. See the CY 2012 OAR for further details.

5. Facility Power Interruptions

31: Facility power interruptions	
Risk Owner	William Allcock
Probability	Moderate
Impact	Schedule: Mod; Cost: Low; Tech: Low
Risk Rating	Moderate
Primary Management Strategies	Participation in Data Center Management Group by the ALCF Operations manager. ALCF pays part of the cost of an Argonne Data Center liaison. Initiate root cause analysis on any failures or degradation; remediate accordingly.
Triggers	Electrical failure; multiple power quality events.

Description

Risk 31 is a general risk that is used to monitor the facility power and manage problems as they are encountered. Problems were encountered with the facility power at both ISSF and the TCS data center in CY 2013.

Evaluation

The transformer that provides power to the ISSF UPS, and therefore everything but the Blue Gene racks, shorted to ground and was a total loss. This resulted in multiple outages, one for the failure and others as ALCF and FMS implemented temporary fixes and later replaced the transformer.

In the fall of 2013, ALCF had multiple breakers powering some of Mira's Blue Gene/Q racks fail, resulting in a total of three outages. These Mira outages represented about 20% of its CY 2013 downtime.

Management

Working with TCS building owners and building managers as well as the breaker manufacturer, ALCF has sent the faulty breakers to the manufacturer to determine the root cause. ALCF is awaiting the report. An electrical contractor was brought in following the second and third incidents to ensure there were no ground faults. The manufacturer has agreed to replace all three breakers at no charge, but it is important to determine root cause of the failures in order to ensure that the remaining breakers, which are from the same manufacturer, don't fail.

For Intrepid, because it has been decommissioned, no further action will be taken regarding the power failures at ISSF.

Retired Risks

The eight risks in Table 5.3 below are risks that have been retired in the past year. Some were retired because the threat has been managed and/or no longer exists. One was a new short-term risk that was developed, monitored, and tracked, but was retired within the same year. The risk colored light green is the risk that was created and then retired in CY 2013.

Table 5.3: Retired Risks in CY 2013

ID	Title	Rating	Management Strategies	Notes
1052	The BG/P (Intrepid) service node Is a single point of failure	Mod	Upgrade service contract put standby service node options in place.	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.
1060	Intrepid system stability issues	Mod	Monitor and analyze system failure data, perform regression testing and baseline testing.	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.
1061	Intrepid system performance issues	Mod	Monitor and analyze system failure data, perform regression testing and baseline testing.	With the decommissioning of Intrepid at the end of CY 2013, this risk has been retired.
1067	Total project count grows with Mira and Intrepid beyond staffing levels	Low	Encourage quality large projects.	The staff has been increased and the plan is to keep project counts within current numbers, this risk has been retired.
1070	Installation of 2nd cooling loop could disrupt operations	Low	Implement process designed to avoid disrupting current operations, require subcontractors follow Argonne safety and LOTO standards.	The second loop was completed in the summer of 2013 and has been used successfully since, ending the risk.
1071	Priority 1 and 2 user documentation unavailable for Mira Production	Mod	Use contingency to hire contractors and add contingency to schedule to assure success.	Documentation was completed within the planned schedule, ending the risk.
1075	Not enough Early Science time	Mod	Mira was accepted early enough to provide three months of dedicated Early Science access. In addition, Early Science projects were allowed to run through Dec. 31, 2013.	Risk was closed with the end of Early Science time on Dec. 31, 2013.
1098	Unable to recruit enough projects for Mira production in 2014	Mod	Recruit new users earlier.	Demand for INCITE 2014 hours remained high, all available INCITE 2014 hours have been allocated, and this risk has been retired.

New and Recharacterized Risks Since the Last Review

The ALCF risk culture leads to the identification of new risks and the recharacterization of existing risks on a regular basis. In the past year, 15 new risks were identified, developed, and monitored and tracked as part of the Steady State risk register. One of them, 1098 (shown in green in Table 5.4 below), was created and subsequently retired in CY 2013. When Mira began production operation in April, four existing project risks were moved to the steady state risk register in February, and a number of risks underwent minor changes consistent with the planned shutdown of Intrepid at the end of CY 2013 and of Intrepid's file system and visualization hardware in the first half of CY 2014.

Table 5.4: New and Recharacterized Risks from CY 2013

ID	Title	Rating	Management Strategies	Notes
31	Facility power interruptions	Mod	Participation in Data Center Management Group by the ALCF Operations Director. ALCF covers the cost of a Argonne data center liaison. Initiate root cause analysis on any failures or degradation; remediate accordingly.	Risk score was lowered initially because there is no expected long-term performance impact. It was later increased due to the Mira breaker failures and the lack of a determination of the root cause.
129	Storage bandwidth insufficient to meet Science/Analytics needs	Low	Consider adding more storage couplets, installing a “burst buffer.”	Risk lowered once alternatives were identified.
990	Electric cost could increase beyond the budget	Low	Lock in long-term contract.	Recharacterized when the electric rate was fixed through a two-year contract with the supplier.
991	Electric cost could decrease from budget	Low		Recharacterized when the electric rate was fixed through a two-year contract with the supplier.
1047	Compiler bugs do not get fixed	Low	Manually fix code with a workaround; use open source compilers; support development of open source compilers.	Recharacterized. An unavoidable issue with using an HPC machine where it is not in the vendor’s mainstream product line.
1056	System stability issues due to upgrades	Mod	Perform upgrades on non-critical systems first when feasible, have a rollback plan in place, monitor performance, team with vendor.	Risk was recharacterized to include Mira.
1060	Intrepid system stability issues	Mod	Monitor and analyze system failure data, perform regression testing and baseline testing.	Risk was recharacterized when it was concluded that the development of management strategies should result in lowering the technical impact score from moderate (3) to low (2).
1061	Intrepid system performance issues	Mod	Monitor and analyze system failure data, perform regression testing and baseline testing.	Risk was recharacterized when it was concluded that the development of management strategies should result in lowering the technical impact score from moderate (3) to low (2).

Table 5.4: (Cont.)

ID	Title	Rating	Management Strategies	Notes
1068	Rack coolant monitor fails to function	Low	Consider increasing coolant flow; contact other users for lessons learned.	Risk was recharacterized when it was concluded that the development of management strategies should result in lowering the technical impact score from moderate (3) to low (2).
1075	Not enough Early Science Program project time	Retired	Work with projects to effectively utilize the system to accomplish science goals as quickly as possible; balance production date with time for ESP work.	Risk was recharacterized to an anticipated end date of Dec. 31, 2013.
1076	If the Interim Supercomputing Support Facility (ISSF), where ALCF-1 resources are currently housed, is decommissioned, ALCF may not have an appropriate facility to host its disaster recovery resources	Mod	Raise concerns with upper lab management (completed); 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.	Risk was recharacterized as strategies to provide alternative disaster recovery resources evolved.
1077	Cost of decommissioning and removing Intrepid could exceed expectations	Low	Approach other sites to see if they would like to take part of Intrepid (completed); explore non-standard options; use management reserves.	Risk recharacterized as removal options were developed. Brookhaven, Rice, and NC State are all taking some racks. Argonne will take care of cutting up and disposing of the rest of the Intrepid racks. The parts will go to a company that recovers usable materials from computer components.
1078	Mira system stability issues	Low	Conduct regression testing, monitor and analyze system failure data.	New risk; was also recharacterized as experience with operating Mira led to lowered probability, lowered cost score and lowered schedule impact.
1079	Mira compute performance issues	Low	Conduct regression testing, monitor and analyze system failure data.	New risk; was also recharacterized as experience with operating Mira led to lowering cost score from high to low and technical impact score from moderate to low.

Table 5.4: (Cont.)

ID	Title	Rating	Management Strategies	Notes
1080	Risk of operating in a shared data center	Low	Coordinate with other divisions located in the Data Center; coordinate with TCS building management; plan with all stakeholders for capacity expansion and actions for addressing a major hardware failure.	New risk due to beginning of operation of Mira in a shared data center.
1081	InfiniBand hardware Is not functional/stable	Low	Test on Vesta when feasible; roll back changes when situation is encountered; purchase backup hardware; work with vendor.	New risk. Risk also recharacterized with probability lowered from moderate to low due to implementing management strategies.
1082	File system metadata performance is inadequate	Low	Help users modify code to not use file per process approach; provide alternative technologies; install additional hardware	To date, metadata performance has been better than needed, risk therefore recharacterized to low probability.
1083	InfiniBand hardware does not meet performance targets	Low	Accept lower performance; add hardware; work with vendor to tune performance.	New risk to address possibility of hardware performance not meeting targets.
1084	Rate of hardware failure exceeds projections	Low	Work with vendor to resolve issues.	Risk brought over from ALCF-2's project risk as a new risk in steady state
1085	Diagnostic suite and utilities fail to detect hardware problems	Mod	Work with vendor to resolve issues. Improve diagnostic suite.	Risk brought over from ALCF-2's project risk as a new risk in steady state. Recharacterized as low since the diagnostic suite has been improved and is now detecting all known hardware problems.
1087	Data analysis servers do not have sufficient memory or memory is too expensive	Low	Include SSD in design study as potential stop gap solution	Risk brought over from ALCF-2's project risk as a new risk in steady state
1090	Overhead rates could increase	Low	Work with lab management	New risk
1091	Injury to workers/overall safety of the division	Mod	Promote safety culture at all levels of the division. Follow Argonne ISM.	Risk brought over from ALCF-2's project risk as a new risk in steady state
1096	Mira file system single point of failure	Low	Purchase a second file system.	New risk created when Mira entered into production mode. Situation resulted from the need to alter the initial file system plans due to hardware limitations, resulting in only one scratch file system.

Table 5.4: (Cont.)

ID	Title	Rating	Management Strategies	Notes
1098	Unable to recruit enough projects for Mira production in 2014	Retired	Recruit new users earlier.	New risk created when Mira entered into production mode. Demand for INCITE 2014 hours remained high, all available INCITE 2014 hours have been allocated, and this risk has been retired.
1099	INCITE does 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 INCITE workflows.	New risk that will remain as an ongoing risk.
1101	Facility that provides cooling is insufficient	Low	Ongoing communications with FMS (provide electrical power to the building) and TCS building management.	New risk created when Mira entered into production mode.

Projected Major Operating Risks for the Next Year

The current top operating risks projected for CY 2014 are listed below in Table 5.5, along with the current risk rating and management strategies for the risk.

Table 5.5: Projected Operating Risks for CY 2014

ID	Title	Rating	Management Strategies
1059	Funding/Budget uncertainties	High	Careful planning, in conjunction with program office, for handling CR, leasing costs, and hires; maintain austerity measures; forward pay to reduce overall leasing costs.
25	Staffing challenges	Mod	Continue to re-plan work, multi-purpose existing staff, and share staff with other divisions.
1049	Staff retention	Mod	Use non-salary based compensation (bonuses, office improvements, work from home); cross-train team members.
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. Initiate root cause analysis on any failures, or degradation, remediate accordingly. Push to obtain a root cause analysis of breaker failures.
1091	Injury to workers/overall safety of the division	Mod	Promote safety culture at all levels of the division. Follow Argonne ISM

Conclusion

The ALCF uses a proven risk management strategy that is documented in its RMP. This document is regularly reviewed and updated to reflect the dynamic nature of risk management, as well as new lessons learned and best practices captured from other facilities. Risk management is a part of the ALCF 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, and informal ones are used for smaller activities within the ALCF. Risks are monitored and tracked using the commercial risk management tool OPRA. Over the past year, eight risks were retired, 15 new risks were added, and many risks were updated. Beyond this, many tools are used to manage risks at the ALCF, particularly in the area of safety. The ALCF's effective risk management plan has contributed to the successful management of a number of significant risks that were 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

The ALCF has an exemplary safety record. Since the division's inception in 2006, the ALCF has never experienced a lost time incident. The facility had one minor first aid case in CY 2013.

ALCF employs appropriate work planning and control principles. A formal "skill of the worker" document is used for routine tasks. Formal specific procedures are in place for more complex tasks, such as changing out the Blue Gene/Q power supplies (thermal hazard) and node boards (very mild chemical hazard due to water treatment chemicals, weight, and potential damage to hardware), as well as medium-voltage electrical maintenance. The facility performs hazard analysis and creates work planning and control documents for emergency work or when there is an unexpected change to previously planned work. For example, when the power transformer that supplied power to non-Blue Gene/P equipment in Building 369 (ISSF) failed, ALCF successfully executed the installation of a pair of 500 kW generators, a project that involved multiple divisions within Argonne and an external vendor. All paperwork was filled out, individuals were contacted directly to expedite approvals, the vendors received the appropriate safety training, and there were no safety issues of any kind.

This page intentionally left blank.

Section 7. Cyber Security

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

ALCF Response

The Argonne Authority to Operate, which includes the ALCF, is included below. During CY 2013, the ALCF had one reportable cyber security incident. Due to the failure of a security device at Argonne, the authentication server that the ALCF relies on suffered a minor denial-of-service outage that lasted approximately 15 minutes. During that time, users were unable to log into Mira. Existing connections were unaffected. A secondary system at Argonne blocked the attacker and the authentication system returned to normal operation shortly thereafter. The failed piece of equipment was subsequently replaced.

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

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

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

0212-022

Dear Dr. Isaacs:

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

The ALCF and the DOE have agreed to the 2014 metrics and targets as proposed in the April 2013 OAR report. These metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2015, the proposed metrics and targets for the current production resources remain the same as for 2014, with the exception of the Mira capability metric. This metric ultimately provides a similar 20% of the system capability threshold, while accommodating the new partition sizes of Mira.

The facility should provide a summary table of the metrics and targets agreed upon for the review of Calendar Year 2014 and include the target and actual values of similar metrics used for 2013 for comparison. The facility should also provide metrics and targets under consideration for CY 2015. 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 2014 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2015, the proposed metrics and targets for the current production resource, Mira, will remain the same as for 2014, with the exception of the capability metric, which has a proposed increased target. The 2014 metrics are covered in Section 8.2 and the 2015 metrics are covered in Section 8.3. Unlike the other facilities, ALCF uses a two-tier capability target for Mira because of partition size restrictions. In 2014, the percentage of INCITE jobs running on 16.7% or larger of the system (131,072 or more cores) will be at least 30%, and 10% or more of the INCITE jobs will be on 33.3% of the system or larger (245,761 or more cores). For 2015, ALCF proposes a percentage of INCITE jobs running on 16.7% or larger of the system (131,072 or more cores) of at least 40%, and 10% or more of the INCITE jobs running on 33.3% of the system (245,761 cores or larger).

8.2 ALCF 2014 OA Performance Metrics

The OA performance metrics, 2013 targets and actuals, and agreed upon 2014 targets are presented in Table 8.1.

Table 8.1: Performance Metrics – 2013 Targets, 2013 Actuals, and Agreed-Upon 2014 Targets

Area	Metric	2013 Targets	2013 Actuals	2014 Targets
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.6/5.0	3.5/5.0
	User Survey – **Response Rate**	25%	31.7%	25%
	% User Problems Addressed Within Three Working Days	80%	91.9%	80%
Business Results	Intrepid Overall Availability	90%	94.7%	-
	Intrepid Scheduled Availability	95%	97.0%	-
	Intrepid INCITE Capability Usage (20% – 100% of system)	40%	64.3%	-
	Mira Overall Availability	80%	90.6%	90%
	Mira Scheduled Availability	85%	95.5%	95%
	% of INCITE Jobs Run on 16.7% or More of Mira (131,072 – 786,432 cores)	20%	60.7%	30%
	% of INCITE Jobs Run on 33.3% or More of Mira (262,144 – 786,432 cores)	5%	33.3%	10%

8.3 ALCF Proposed 2015 OA Performance Metrics

The OA performance metrics, agreed upon 2014 targets, and 2015 proposed targets are shown in Table 8.2.

Table 8.2: Performance Metrics – Agreed-Upon 2014 Targets and Proposed 2015 Targets

Area	Metric	2014 Targets	Proposed 2015 Targets
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	3.5/5.0
	User Survey – **Response Rate**	25%	25%
	% User Problems Addressed Within Three Working Days	80%	80%
Business Results	Mira Overall Availability	90%	90%
	Mira Scheduled Availability	95%	95%
	% of INCITE Jobs Run on 16.7% or More of Mira (131,072 – 786,432 cores)	30%	40%
	% of INCITE 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, # 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
Risk Management	Report on major risks

Conclusion

The agreed-upon 2014 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2015, the proposed metrics and targets for the current production resources remain the same as for 2014, with the exception of an increased capability target. Achieving the agreed-upon 2014 and the proposed 2015 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

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, 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.

Formula:

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

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:

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

ALCF Calculations

A simple example should make this clear. If on a particular day, there were 14 hours of scheduled maintenance, and two hours where the machine was down due to unexpected failures, there were 8 hours actually available (24-14-2 = 8) resulting in 33.3% overall availability (8/24). There were 10 hours scheduled to be available (24-14 = 10), but there were actually 8 hours available, resulting in 80% scheduled availability (8/10).

To implement the above, ALCF tracks availability at the core-second level. The Blue Gene architecture allows an individual node cards, 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 512 nodes (8,192 cores) are allowed to run, resulting in 512 nodes being 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 only one production file system for Mira. Therefore, if it is down, the entire machine is considered to be down. There is an exception to the above. 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 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 the job that occurred during the downtime. This is done to prevent reporting greater than 100% utilization.

To produce the actual numbers, ALCF takes the downtime data and 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, the overall availability and the scheduled availability daily values are arithmetically averaged over 7-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 Jan 1 – Jan 7, the second data point is the average of Jan 8 – Jan 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 (4 or more) of the data points, a final data point is calculated from those values and plotted. If not (3 or fewer), those values are included in the previous data point, which becomes an average of between 8 and 11 data points. This is to avoid significant deviations of the last point due to a small average.

MTTI Calculation Details

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

Formula:

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

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

MTTF Calculation Details

Time, on average, to an unscheduled outage on the system.

Formula:

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

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

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 this data is 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. Projects have attributes associated with them (INCITE, Discretionary, type of science, etc.). 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. Jobs that cross day boundaries have the hours 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:

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

The daily values are then averaged as described in the availability section and plotted. The darker black vertical line marks the calendar year boundary.

Capability Calculation Detail: The data for everything except the job usage by size graph is the sum of the core-hours for qualifying jobs, with the plots showing daily values. Each bar in the job usage by size graph covers one week of data. The data is summed by type and then divided by the total for the week to determine the percentage.

This page intentionally left blank.

Appendix B – ALCF Director's Discretionary Projects

January 1, 2013 – December 31, 2013

Intrepid

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
AbInitioC12_esp	Argonne National Laboratory	Steven C. Pieper	Ab-initio Reaction Calculations for Carbon-12	Physics	7,500,000
ALCF-LANL_withGO	Los Alamos National Laboratory	Susan Kurien	Connect Data at ALCF to LANL Using Globus Online	Engineering	10
Allinea	Allinea Software	Ray Loy	Improved Debugging Memory Usage for BG/P	Internal	500,000
alpha-nek	University of Chicago, Fermilab	Maxwell Hutchinson	DNS of Multi-mode Rayleigh-Taylor Instability	Engineering	10,000,000
AR-1-2012	Toyota Technological Institute at Chicago	Yang Shen	Molecular Dynamics Simulations for Exploring Androgen Receptor Antagonism, Drug-Resistant Mutations, and Antagonist Design	Biological Sciences	2,000,000
ATLASP	Argonne National Laboratory	Thomas J. LeCompte	Grid-Enabling High Performance Computing for ATLAS	Physics	5,000,000
Autoignition_esp	Eidgenössische Technische Hochschule Zürich (ETH Zurich)	Christos Frouzakis	Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow	Chemistry	5,000,000
BackFlexMed	The University of North Carolina at Chapel Hill	Nikolay V. Dokholyan	Incorporating Backbone Flexibility in MedusaDock	Biological Sciences	5,000,000
BGQtools_esp	Argonne National Laboratory	Kalyan Kumaran	Enabling Petascale Science on BG/Q: Tools, Libraries, Programming Models, & Other System Software	Computer Science	4,000,000
BifurcationALCF	University of Canterbury	Timothy David	Simulation of Cell Coupling in Arterial Bifurcation	Biological Sciences	1,100,000
Bulk_Properties_esp	Iowa State University	Mark Gordon	High Accuracy Predictions of the Bulk Properties of Water	Chemistry	2,000,000
Catalyst	Argonne National Laboratory	Katherine Riley	Catalyst	Internal	50,000,000
CFDAnisotropic_esp	University of Colorado-Boulder	Kenneth Jansen	Petascale, Adaptive CFD	Engineering	2,000,000
CharmRTS	University of Illinois at Urbana-Champaign	Laxmikant V. Kale	Charm++ and its Applications	Computer Science	1,000,000
ChemMechOxides	King's College London/Argonne National Laboratory	Alessandro De Vita, James Kermode	Chemo-mechanical Properties of Complex Oxide Interfaces	Materials Science	500,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
CNL_WSU	Washington State University	Soumik Banerjee	Molecular Modeling of Nanoscale Transport Pertinent to Energy Conversion and Storage	Energy Technologies	250,000
CobaltDevel	Argonne National Laboratory	Narayan Desai	Cobalt Development	Internal	1
CONVERGE-BGP-LDRD	Argonne National Laboratory	Marta Garcia Martinez	Improvement I/O on CFD code CONVERGE	Engineering	2,000,000
critical_perf	Argonne National Laboratory	Kalyan Kumaran	Critical Debugging Project	Internal	5,000,000
Crystal_Eng_disc	The University of Texas at Austin/Fritz-Haber-Institut der Max-Planck-Gesellschaft/Argonne National Laboratory	J. R. Chelikowsky, N. Marom, A. Tkatchenko, L. Ghiringhelli, O. A. von Lilienfeld, J. Hammond	Toward Crystal Engineering From First Principles	Materials Science	2,000,000
DarkUniverse_esp	Los Alamos National Laboratory	Salman Habib	Cosmic Structure Probes of the Dark Universe	Physics	16,000,000
DNS-hot_mixing_layer	Conservatoire National des Arts et Metiers	Francesco Grasso	DNS Study of a Spatially Developing Compressible Mixing Layer with Non-unity Density and Temperature Ratios	Engineering	2,500,000
DNS-velvor	Argonne National Laboratory/University of Washington/The University of Texas at Austin/Universidad Carlos III de Madrid	Ammar Abdilghanie	Development of a Novel Methodology for Petascale simulation of Low Mach Combustion	Engineering	400,000
Ducts	Illinois Institute of Technology/Argonne National Laboratory/KTH Royal Institute of Technology	Hassan Nagib	Aspect Ratio Effects in Turbulent Duct Flows Studied Through Direct Numerical Simulation	Engineering	1,500,000
Epigenetics	The University of Utah	Nadeem Vellore	Nanoscale Biomolecular Simulation for Understanding of Peptide Recognition by Epigenetic Protein	Biological Sciences	300,000
ExM	Argonne National Laboratory	Michael Wilde	Systems Support for Extreme-Scale Many-Task Applications	Computer Science	5,000,000
Femtomagnetism	Indiana State University	Guoping Zhang	First-Principles Calculation of Laser Induced Ultrafast Magnetism	Physics	3,000,000
FusionFS	Illinois Institute of Technology	Ioan Raicu	Distributed File Systems for Exascale Computing	Computer Science	500,000
GFDL_esp	Geophysical Fluid Dynamics Laboratory	Venkatramani Balaji, Tim Williams, Chris Kerr	Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model	Earth Science	7,500,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
GroundMotion_esp	University of Southern California	Thomas Jordan	Using Multi-scale Dynamic Rupture Models to Improve Ground Motion Estimates	Earth Science	7,500,000
HSCD_esp	The University of Chicago	Alexei Khokhlov	High-Speed Combustion and Detonation (HSCD)	Chemistry	5,000,000
HEDhydro_RMRT	University of Michigan	Eric Johnsen	High-Fidelity, High-Energy-Density Hydrodynamics Simulations of Shocks Interacting with Material Discontinuities	Engineering	250,000
HighReyTurb	The University of Texas at Austin	Robert Moser	Direct Numerical Simulations of High Reynolds Number Turbulent Channel Flow	Engineering	10,000,000
i-SNSDCs	University of California, Los Angeles	Anastassia N. Alexandrova	Multi-Scale Modeling of Catalytic Interfaces Based on 2-D Sub-nano Surface-Deposited Clusters	Chemistry	2,000,000
I12_BloodFlow	Brown University	George Karniadakis	Multiscale Blood Flow Simulations	Biological Sciences	24,632,149
I12_FaultOblivious	Lawrence Livermore National Laboratory	Maya Gokhale	Fault-Oblivious Exascale Computing Environment	Computer Science	2,837,111
I12_LaserPlasmaSim	Lawrence Livermore National Laboratory	Denise Hinkel	Simulations of Laser-Plasma Interactions in Targets for the National Ignition Facility and Beyond	Physics	11,002,552
I12_LES_Turbines	Centre Europeen de Recherche et de Formation Avancee en Calcul Scientifique (CERFACS)	Thierry Poinsot	Large Eddy Simulation of Two Phase Flow Combustion in Gas Turbines	Chemistry	4,909,488
I12_NucStructReact	Iowa State University	James Vary	Nuclear Structure and Nuclear Reactions	Physics	1,490,805
I12_PEACEndStation	Lawrence Livermore National Laboratory	Leonid Oliker, Todd Gamblin	Performance Evaluation and Analysis Consortium End Station	Computer Science	3,228,834
I12_PESurfChem	University of Minnesota	Donald Truhlar, Osanna Tishchenko	Potential Energy Surfaces for Simulating Complex Chemical Processes	Chemistry	10,908,660
I12_SSSPP	Argonne National Laboratory	Ewing Lusk, Rob Latham	Scalable System Software for Performance and Productivity	Computer Science	26,094
I12_StochasticConver	The State University of New York at Stony Brook	James Glimm	Stochastic (w*) Convergence for Turbulent Combustion	Engineering	15,734,380
I12_VaporDroplets	University of California, Irvine	Said Elghobashi	Direct Simulation of Fully Resolved Vaporizing Droplets in a Turbulent Flow	Engineering	14,618,179
IBM-performance	Argonne National Laboratory	Kalyan Kumaran	BG/P Performance Runs carried out by IBM employees	Internal	500,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
IOFSL	Argonne National Laboratory	Rob Ross	I/O Forwarding Scalability Layer	Computer Science	1,500,000
LatticeQCD_esp	Fermi National Accelerator Laboratory	Paul Mackenzie	Lattice Quantum Chromodynamics	Physics	7,500,000
LES_Turbines_DD	Centre Europeen de Recherche et de Formation Avancee en Calcul Scientifique (CERFACS)	Thierry Poinsot	Large Eddy Simulation of Two Phase Flow Combustion in Gas Turbines	Chemistry	1,700,000
Licrys	The University of Chicago/Argonne National Laboratory	Juan J. de Pablo	Liquid Crystal Based Functional Materials	Materials Science	2,000,000
LLVM	Argonne National Laboratory	Hal Finkel	LLVM Compiler Tools for the BG/Q	Internal	50,000
LowPriority	Argonne National Laboratory	Katherine Riley	LowPriority	Internal	5,000,000
LPIbenchmark	Lawrence Livermore National Laboratory	Steven Langer	Benchmarking LPI Simulation Performance on Blue Gene Systems	Physics	1,500,000
MADNESS_MPQC_esp	Brookhaven National Laboratory	Robert Harrison	Accurate Numerical Simulations of Chemical Phenomena Involved in Energy Production and Storage with MADNESS and MPQC	Chemistry	7,500,000
Mat_Design_esp	Argonne National Laboratory	Larry Curtiss	Materials Design and Discovery: Catalysis and Energy Storage	Materials Science	3,000,000
Maintenance	Argonne National Laboratory	Cheetah Goletz	LCF Operations System Maintenance	Internal	10,000,000
MCA	Pacific Northwest National Laboratory	Yousu Chen	Massive Contingency Analysis	Energy Technologies	1,000,000
MembraneChannels	University of California-San Diego	Igor Tsigelny	Simulation and Modeling of Membranes Interactions with Unstructured Proteins and Computational Design of Membrane Channels for Absorption of Specified Ions	Biological Sciences	1,000,000
MiningSysLogs	Argonne National Laboratory	Michael E. Papka	Data Mining ALCF System Logs to Better Understand System Faults, Usage, and Application Characteristics	Computer Science	100,000
MiraCon2013	Argonne National Laboratory	Richard Coffey	Project for Mira Conference 2013	Training	1,000,000
MLearning4CCS	Argonne National Laboratory	O. Anatole von Lilienfeld	Machine Learning for the Exploration of Chemical Compound Space	Chemistry	5,000,000
MOAB-scalability	Argonne National Laboratory	Tim Tautges	MOAB Debugging & Scalability	Computer Science	200,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
MRCCSN	The University of Chicago/Fermi National Accelerator Laboratory	Sean M. Couch	3D Simulations of Magnetorotational Core-Collapse Supernovae	Physics	35,000,000
MRDFT	Argonne National Laboratory	Alvaro Vazquez-Mayagoitia	Wavelet Basis Set in Density Functional Theory Methods for Photoelectric Materials	Materials Science	900,000
mri_athena	University of Rochester	Eric Blackman	Shearing Box Simulations of Accretion Disks	Physics	300,000
MTCScienceApps	Argonne National Laboratory	Michael Wilde	Many Task Computing Science Applications	Computer Science	5,000,000
MultiscaleMolSim_esp	The University of Chicago	Gregory Voth	Multiscale Molecular Simulations at the Petascale	Biological Sciences	17,500,000
NAMD_esp	The University of Chicago	Benoit Roux	NAMD - The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based on a Polarizable Force Field	Biological Sciences	7,500,000
NekCEM	Argonne National Laboratory	Misun Min	Electromagnetics	Physics	1,500,000
NekLBM	Argonne National Laboratory	Misun Min	Lattice Boltzmann Simulations for Fluids	Engineering	1,300,000
NEURO_EVOLVE	Indian Institute of Science/University College London	Biswa Sengupta	The Role of Constraints in the Design of the Nervous System	Biological Sciences	80,000
neuronq	Yale University	Michael Hines	High Performance Neuron simulations on the Blue Gene/Q	Biological Sciences	1,000,000
Neutronics_MAS	Argonne National Laboratory	Micheal A. Smith	Neutronics Simulation of FFTF and EBR-II	Nuclear Energy	5,000,000
NRCM_DD	Argonne National Laboratory	V. Rao Kotamarthi	Dynamic Downscaling of Climate Models	Earth Science	36,300,000
NuMMo_GC	The Dow Chemical Company/Northwestern University	Dee Dickerson	Coarse-Grain Model Development for Photoresist Polymers	Materials Science	500,000
OpenFOAM-ALCF	Argonne National Laboratory	Ramesh Balakrishnan	OpenFOAM Based Computational Fluid Dynamics Simulations at the Argonne Leadership Computing Facility	Engineering	1,000,000
Operations	Argonne National Laboratory	Bill Allcock	Systems Administration Tasks	Internal	0
Operations_Test	Argonne National Laboratory	Cheetah Goletz	Operations Infrastructure Testing	Internal	500,000
Oxygen_defects	Argonne National Laboratory	Olle Heinonen	First-Principle Investigations of Oxygen Defects in Metal-Oxide-Metal Heterostructures	Materials Science	100,000
PARTS	Argonne National Laboratory	Jeff Hammond	Parallel Run-Time Systems	Computer Science	5,500,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
PBGL	Indiana University	Andrew Lumsdaine	Parallel Boost Graph Library	Computer Science	100,000
Pepag	The University of Chicago/Argonne National Laboratory	Juan J. de Pablo	Macromolecular Folding and Aggregation	Biological Sciences	2,000,000
Performance	Argonne National Laboratory	Kalyan Kumaran	Performance	Internal	5,000,000
PlasmaMicroturb_esp	Princeton Plasma Physics Laboratory	William Tang	Global Simulation of Plasma Microturbulence at the Petascale & Beyond	Physics	7,500,000
prec_sense	Argonne National Laboratory	Laura Zamboni	Sensitivity and Uncertainty of Precipitation of the GFDL High Resolution Model	Earth Science	4,225,000
QLG-Turbulence	The College of William & Mary	George Vahala	Quantum Lattice Algorithm for Quantum Turbulence	Engineering	100,000
QMC_MEP	University of L'Aquila (Univaq)	Leonardo Guidoni	Minimum Energy Path with Quantum Monte Carlo	Chemistry	1,000,000
radix	Argonne National Laboratory	Kamil Iskra	MCS Radix Team	Computer Science	2,000,000
Repast	Argonne National Laboratory	Michael North	Exascale Agent-Based Modeling System	Computer Science	750,000
rlins	Universidade Federal de Pernambuco	Roberto D. Lins	A Microscopic Perspective on Outer Membrane Remodeling and Antimicrobial Peptide Resistance	Biological Sciences	1,000,000
ScalableML	General Electric Company	Dongryeol Lee	Distributed Optimization for Scalable Machine Learning	Computer Science	60,000
ScaleAnalysisTurb	Los Alamos National Laboratory	Susan Kurien	Scale-Analysis of Data from High-Resolution Simulations of Boussinesq Turbulence	Engineering	490,000
SDAV	Argonne National Laboratory	Michael E. Papka	SciDAC Scalable Data Management Analysis and Visualization	Computer Science	4,800,000
SecCG	California Institute of Technology	Thomas F. Miller, III	Coarse-Grained Simulations of Sec-Facilitated Protein Translocation and Membrane Integration	Biological Sciences	6,000,000
Shell_Cat	Shell International E&P, Inc.	Leonardo Spanu	Investigation of Catalytic Properties of Nanoclusters	Chemistry	2,000,000
SiliconeRubberAlt	The Dow Chemical Company	Dee Dickerson	Novel Strategies for the Design and Preparation of Liquid Silicone Rubber Alternatives	Chemistry	501,200
Spray_and_soot	Stanford University	Colleen Kaul	DNS of Spray-Soot-Turbulence-Chemistry Interactions	Chemistry	750,000
TestHarness	Argonne National Laboratory	Eric Pershey	TestHarness	Internal	2,000,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
Tools	Argonne National Laboratory	Scott Parker	ALCF Performance Tools	Internal	1,000,000
TurbChannelFlow_esp	The University of Texas at Austin	Robert Moser	Petascale Direct Numerical Simulations of Turbulent Channel Flow	Engineering	5,000,000
TurbNuclComb_esp	The University of Chicago	Don Lamb, Carlo Graziani	Petascale Simulations of Turbulent Nuclear Combustion	Physics	15,000,000
turbulentRBC	Occidental College	Janet Scheel	DNS Simulations of Turbulent Convection	Physics	600,000
ucl_qmc	University College London	Dario Alfe	Quantum Monte Carlo Methods for Solids and Liquids	Materials Science	500,000
User_Services	Argonne National Laboratory	Richard Coffey	User Services	Internal	0
USO_InterimDD	Argonne National Laboratory	Richard Coffey	User Services Organization Interim Project	Internal	10,000
UVa-laser-multilayer	University of Virginia	Leonid V. Zhigilei	Atomistic Simulation of Laser Processing of Metal Multilayers	Engineering	1,000,000
visualization	Argonne National Laboratory	Michael E. Papka	Visualization and Analysis Research and Development for ALCF	Internal	1,000,000
Wall_Turb_DNS	National Aeronautics and Space Administration	Gary N. Coleman	DNS of Wall-Bounded Turbulence	Engineering	1,000,000
				Total Intrepid DD	488,604,463

January 1, 2013 – December 31, 2013
Mira DD (No ESP)

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
Accretion	The University of Chicago	Fausto Cattaneo	Global Simulations of Accretion Discs	Physics	5,914,443
ALCF_GEGR_NOISE	Argonne National Laboratory	Ramesh Balakrishnan	Porting and Tuning CFD codes for Aero-acoustics Computations	Engineering	5,000,000
ALCF_Getting_Started	Argonne National Laboratory	Chel Heinzel	ALCF Getting Started	Computer Science	10,000,000
ALCF_SANDIA_RAPTOR	Sandia National Laboratories	Joseph C. Oefelein	Scaling the RAPTOR Combustion LES Code on Mira	Engineering	5,000,000
ALCFCERFACS_CFD	Argonne National Laboratory/Centre Europeen de Recherche et de Formation Avancee 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
Allinea	Allinea Software	Ray Loy	Improved Debugging Memory Usage for BG/P	Internal	2,000,000
AR-1-2012	Toyota Technological Institute at Chicago	Yang Shen	Molecular Dynamics Simulations for Exploring Androgen Receptor Antagonism, Drug-Resistant Mutations, and Antagonist Design	Biological Sciences	1,000,000
ATPESC2013	Argonne National Laboratory	Paul Messina	Argonne Training Program on Extreme Scale Computing	Computer Science	10,000,000
BIG_MAC	Northwestern University	Monica Olvera de la Cruz	Effective Interactions in Coulombic Systems with Highly Disparate Particle Sizes	Physics	3,145,728
BloodFlowBGQ	Brown University	George Karniadakis	Multiscale Blood Flow Simulations	Mathematics	3,000,000
Camellia	Argonne National Laboratory	Nathan Roberts	Camellia for Discontinuous Petrov-Galerkin Simulations of Incompressible Flow	Physics	5,000,000
Catalyst	Argonne National Laboratory	Katherine Riley	Catalyst	Internal	10,000,000
charmm_zmod	Harvard University/Harvard Medical School	Robert J. Petrella	Highly Parallel Macromolecular Conformational Searches and Energy Evaluations with the CHARMM Program	Biological Sciences	10,000,000
ChemiCAD	The University of Chicago	Rick Stevens	Parallel Multi-objective Optimization Algorithm for de Novo Ligand Design	Biological Sciences	1,000,000
ChemMechOxides	King's College London/Argonne National Laboratory	Alessandro De Vita, James Kermode	Chemo-mechanical Properties of Complex Oxide Interfaces	Materials Science	1,500,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
Chombo_Crunch_DD	Lawrence Berkeley National Laboratory	David Trebotitch	Advanced Simulation of Subsurface Flow and Reactive Transport Processes Associated with Carbon Sequestration	Chemistry	10,000,000
ClimateUncertainty1	Argonne National Laboratory	Ian Foster	Investigation of initial condition uncertainty in climate models	Earth Science	5,000,000
CMT	Argonne National Laboratory	Scott Parker	Compressible Multiphase Turbulence	Engineering	12,000,000
CobaltDevel	Argonne National Laboratory	Narayan Desai	Cobalt Development	Computer Science	10,000,000
CombStab_CEF_DD13	Cascade Technologies, Inc.	Lee Shunn	Combustion Stability in Complex Engineering Flows	Engineering	10,000,000
CompBIO	Argonne National Laboratory/The University of Chicago	Rick Stevens	Multiscale Simulations in Biology: Evolution and Ecology of Microbes	Biological Sciences	1,000,000
ComplexCrystalSearch	Argonne National Laboratory	Carolyn Phillips	Design Space Searching for Complex Crystals	Materials Science	28,151
Concerted_Flows	Argonne National Laboratory	Venkatram Vishwanath	Topology-Aware Data Movement and Flow Coordination for Supercomputing	Computer Science	1,000,000
Coral_rfp_prep	Argonne National Laboratory	Kalyan Kumaran, Vitali Morozov	CORAL RFP Preparatory work	Internal	5,500,000
critical_perf	Argonne National Laboratory	Kalyan Kumaran	Critical Debugging project	Internal	100,000,000
CROC	Fermilab, University of Chicago	Nickolay Y. Gnedin	Cosmic Reionization On Computers	Physics	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	1,000,000
CSM_PE_LBM	Colorado School of Mines	Feng Xiao	Lattice Boltzmann Simulations in Complex Porous Medium Geometries	Earth Science	5,000,000
DE-NE0000533	Massachusetts Institute of Technology	Kan-Ju Lin	Radiation Tolerance and Thermo-mechanical Properties of Amorphous SiCO glasses and SiCO/Fe composites	Materials Science	1,024,000
DiffLight_APS	Argonne National Laboratory	Michael Borland	Diffraction-Limited Light Source for APS Upgrade	Physics	10,000,000
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

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
duanlianx	Missouri University of Science and Technology	Lian Duan	Direct Numerical Simulations of laminar-Turbulent Transition in Swept-Wing Boundary Layers	Engineering	2,000,000
ExM	Argonne National Laboratory	Michael Wilde	Systems Support for Extreme-Scale Many-Task Applications	Computer Science	5,000,000
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	Benson Muite	Large and Accurate Numerical Solutions of Partial Differential Equations	Mathematics	3,000,000
FutureLight_DD13	Lawrence Berkeley National Laboratory	Robert Ryne	Advanced Accelerator Simulations for Future Light Sources	Physics	10,000,000
gamra	Earth Observatory of Singapore	Sylvain Barbot	Aerodynamics with Adaptive Mesh Refinement	Earth Science	2,000,000
GE_CTS	General Electric Company	Anne L. Dord	Highly Resolved LES of a GE 3 Cup Combustor System	Engineering	10,000,000
GEA_Noise_Combustion	General Electric Company	Trevor Tersmette, Jason Lee, Erin Lariviere	GE Aviation Engine Simulations	Engineering	7,000,000
HiFi_Carlo_DD13	Massachusetts Institute of Technology	Benoit Forget	High-Fidelity Simulation of Small Modular Reactors with Monte Carlo	Nuclear Energy	10,000,000
IBM_CORAL	IBM Watson Research Center	Bob Walkup	IBM Coral Benchmarks	Internal	6,553,600
IBM_graph500	Argonne National Laboratory/IBM Corporation	Fabrizio Petrini	June 2012 Graph 500 Submission	Computer Science	3,200,000
IBM-performance	Argonne National Laboratory	Kalyan Kumaran	BG/P Performance Runs carried out by IBM employees	Internal	1,100,000
IBM-Software-Testing	Argonne National Laboratory	Susan Coghlan	IBM Testing of Early Release Software for the BG/Q	Internal	50,000,000
icee	LBNL	John Wu	ICEE	Computer Science	500,000
InterX	InterX Inc.	Boris Fain	A Comprehensive Survey of Molecular Interactions at Chemical Accuracy	Biological Sciences	3,000,000
LAMMPSopt	IBM	Paul Coffman	LAMMPS Performance Optimization	Materials Science	2,000,000
Large_Scale_Sim	University of California-San Diego	Rommie E. Amaro	Simulation of Large Scale Biomolecular Systems	Biological Sciences	20,000,000
Licrys	The University of Chicago/Argonne National Laboratory	Juan J. de Pablo	Liquid Crystal Based Functional Materials	Materials Science	2,000,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
Maintenance	Argonne National Laboratory	Cheetah Goletz	LCF Operations System Maintenance	Internal	10,000,000
MCMD_Nano	Argonne National Laboratory	Phay J. Ho	Radiation Damage in Nanoparticles by Intense X-ray Interactions	Materials Science	2,000,000
Meso_CCS_DD13	Centre Europeen de Recherche et de Formation Avancee en Calcul Scientifique (CERFACS)	Roberto Paoli	Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus	Earth Science	15,000,000
MiraBootCamp2013	Argonne National Laboratory	Kalyan Kumaran	Project for Mira Boot Camp 2013	Internal	25,000,000
MiraPunchList	Argonne National Laboratory	Susan Coghlan	Mira Punch List	Internal	10,000,000
ML4TMO	Argonne National Laboratory/Columbia University	Alejandro Lopez-Bezanilla	Machine Learning for Transition Metal Oxides	Physics	5,000,000
MMF-IPHOC	Science Systems and Applications, Inc./ NASA	Anning Cheng	Climate Sensitivity Experiments Using a Multi-scale Modeling Framework with a Higher-order Turbulence Closure in its CRM	Earth Science	3,595,200
MOAB-scalability	Argonne National Laboratory	Tim Tautges	MOAB debugging & scalability	Computer Science	200,000
ModSS	RWTH Aachen University	Ahmed E. Ismail	Molecular Dynamics Simulations of Superspreaders	Materials Science	5,000,000
moose	Idaho National Laboratory	Derek Gaston	Full Core Reactor Modeling with MOOSE	Nuclear Energy	1,000,000
MRCCSN	The University of Chicago/Fermi National Accelerator Laboratory	Sean M. Couch	3D Simulations of Magnetorotational Core-Collapse Supernovae	Physics	30,000,000
MSMClimate_DD13	The National Center for Atmospheric Research	Greg Holland	Multi-Scale Modeling of Climate and High-Impact Weather	Earth Science	20,000,000
NanoCrysTa_DD13	Los Alamos National Laboratory	Timothy Germann	Multi-billion-atom MD Studies of the Mechanical Response of Nanocrystalline Ta	Materials Science	10,000,000
NekCEM	Argonne National Laboratory	Misun Min	Electromagnetics	Physics	1,000,000
NekLBM	Argonne National Laboratory	Misun Min	Lattice Boltzmann 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
NMGC-Mira-2013	University of Minnesota	J. Ilja Siepmann	Computations for the Development of the Nanoporous Materials Genome	Materials Science	8,000,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
NRCM_DD	Argonne National Laboratory	V. Rao Kotamarthi	Dynamic Downscaling of Climate Models	Earth Science	500,000
NuStab_DD13	Northwestern University	George Schatz	Computational Studies of Nucleosome Stability	Biological Sciences	10,000,000
OpenFOAM-ALCF	Argonne National Laboratory	Ramesh Balakrishnan	OpenFOAM based Computational Fluid Dynamics Simulations at the Argonne Leadership Computing Facility	Engineering	1,000,000
Operations	Argonne National Laboratory	Bill Allcock	Systems administration tasks	Internal	10,000,000
Operations_Test	Argonne National Laboratory	Cheetah Goletz	Operations infrastructure testing	Internal	10,000,000
OPV	Currently University of Chicago Ph. D. student, project will be continued after graduation at an institution TBD.	Kenley Pelzer	Mesoscale modeling of charge transport in organic photovoltaics	Materials Science	1,000,000
parallelQD	Texas Tech University	Bill Poirier	Massively Parallel Quantum Dynamics	Chemistry	4,000,000
ParamStudyBouss	Los Alamos National Laboratory	Susan Kurien	Parameter Studies of Rotating Stratified Boussinesq Flows	Engineering	2,250,000
Performance	Argonne National Laboratory	Kalyan Kumaran	Performance	Internal	10,000,000
PESurfChem	University of Minnesota	Donald Truhlar, Osanna Tishchenko	Potential Energy Surfaces for Simulating Complex Chemical Processes	Chemistry	4,915,200
PHASTA_NCSU	North Carolina State University	Igor A. Bolotnov	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	2,500,000
Poly_Self_Assembly	Argonne National Laboratory	Derrick Mancini	Meso-scale Modeling of Self-assembly of Polymer Grafted Nanoparticles	Materials Science	12,000,000
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 Computational Experiments of Atmospheric Turbulence-Driven Wind Turbine Aerodynamics and Wakes	Engineering	2,000,000
QBOX	University of California, Davis	Francois Gygi	QBOX	Materials Science	10,000,000
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	2,800,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
QMC_Bio_Catalysis	Argonne National Laboratory	William Parker	Quantum Monte Carlo Simulations of Biochemical and Catalysis-Related Systems	Chemistry	10,000,000
qmcsc	University of Illinois at Urbana-Champaign	Lucas K. Wagner	First Principles Quantum Monte Carlo for Superconducting Materials	Physics	500,000
Repast	Argonne National Laboratory	Michael North	Exascale Agent-Based Modeling System	Computer Science	100,000
replisome-model	University of Illinois at Urbana-Champaign	Aleksii Aksimentiev	Towards a Model of the Replisome	Biological Sciences	4,104,000
rlins	Universidade Federal de Pernambuco	Roberto D. Lins	A Microscopic Perspective on Outer Membrane Remodeling and Antimicrobial Peptide Resistance	Biological Sciences	2,000,000
rtflames	Northwestern University	Elizabeth P. Hicks	DNS Simulations of Turbulent Rayleigh-Taylor Unstable Flames using Nek5000	Physics	5,000,000
SC13_prep	Argonne National Laboratory	Richard Coffey	Scaling and Testing Runs for Supercomputing 2013	Internal	10,000,000
SDAV	Argonne National Laboratory	Michael E. Papka	SciDAC Scalable Data Management Analysis and Visualization	Computer Science	250,000
shan-2013-shock	Sandia National Laboratories, New Mexico	Tzu-Ray Shan	Nanostructure-enhanced Detonation in Energetic Materials	Materials Science	2,600,000
Singlet_Fission	The George Washington University	Hanning Chen	Probing the Free Energy Surface of Spin Separation in Singlet Fission	Chemistry	6,000,000
Solitons	Argonne National Laboratory/University of British Columbia	Benjamin Gutierrez	Solitons in Field Theory and Applications	Physics	1,000,000
SteamGenerator	The University of Utah	Jeremy Thornock, Jennifer Spinti	Assessing Steam Generator Readiness for Compliance with Future NOx and CO2 Regulation	Engineering	2,000,000
SuspRheometry_DD13	National Institute of Standards and Technology	William George	High Fidelity Simulation of Complex Suspension Flow for Practical Rheometry	Materials Science	10,000,000
TACOMA	General Electric Company	Brian E. Mitchell	TACOMA Porting and Scaling Study	Engineering	1,000,000
TECA_Scaling	Argonne National Laboratory	Venkatram Vishwanath	Scaling the Toolkit for Extreme Climate Analysis on Blue Gene/Q systems	Earth Science	5,000,000
Tools	Argonne National Laboratory	Scott Parker	ALCF Performance Tools	Internal	1,000,000
TotalView	Rogue Wave Software, Inc.	Peter Thompson, Ray Loy	TotalView Debugger on Blue Gene/P	Internal	1,000,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
turbulentRBC	Occidental College	Janet Scheel	DNS Simulations of Turbulent Convection	Physics	400,000
UdSaeroacoustics	Universite de Sherbrooke	Marlene Sanjose	LES of Turbulent Jet Noise	Engineering	1,572,864
Uintah_DD13	University of Utah	Martin Berzins	Porting Uintah to the Blue Gene/Q Architecture	Engineering	15,000,000
umn_crackle	University of Minnesota	Joseph Nichols	Large Eddy Simulation of Crackling Supersonic Jets	Engineering	5,000,000
urban_street_canyon	Illinois Institute of Technology	Dietmar Rempfer	Characterization and Low-Dimensional Modeling of Urban Fluid Flow	Engineering	2,000,000
Vestas_Park_LES	Vestas Wind Systems A/S	Gregory Oxley	LES Investigation of Stability Enhanced Wake Losses on Large Wind Parks	Energy Technologies	5,000,000
visualization	Argonne National Laboratory	Michael E. Papka	Visualization and Analysis Research and Development for ALCF	Internal	83,333
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
				Total Mira DD (no ESP)	824,755,911

Mira ESP

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
AbInitioC12_esp	Argonne National Laboratory	Steven C. Pieper	Ab-initio Reaction Calculations for Carbon-12	Physics	110,000,000
Autoignition_esp	Eidgenössische Technische Hochschule Zürich (ETH Zurich)	Christos Frouzakis	Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow	Chemistry	150,000,000
BGQtools_esp	Argonne National Laboratory	Kalyan Kumaran	Enabling Petascale Science on BG/Q: Tools, Libraries, Programming Models, & Other System Software	Computer Science	20,000,000
Bulk_Properties_esp	Iowa State University	Mark Gordon	High Accuracy Predictions of the Bulk Properties of Water	Chemistry	150,000,000
CFDAnisotropic_esp	University of Colorado-Boulder	Kenneth Jansen	Petascale, Adaptive CFD	Engineering	150,000,000
DarkUniverse_esp	Los Alamos National Laboratory (LANL)	Salman Habib	Cosmic Structure Probes of the Dark Universe	Physics	150,000,000
GFDL_esp	Geophysical Fluid Dynamics Laboratory	Venkatramani Balaji, Tim Williams, Chris Kerr	Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model	Earth Science	150,000,000
GroundMotion_esp	University of Southern California	Thomas Jordan	Using Multi-scale Dynamic Rupture Models to Improve Ground Motion Estimates	Earth Science	150,000,000
HSCD_esp	The University of Chicago	Alexei Khokhlov	High-Speed Combustion and Detonation (HSCD)	Chemistry	150,000,000
LatticeQCD_esp	Fermi National Accelerator Laboratory	Paul Mackenzie	Lattice Quantum Chromodynamics	Physics	150,000,000
MADNESS_MPQC_esp	Brookhaven National Laboratory	Robert Harrison	Accurate Numerical Simulations of Chemical Phenomena Involved in Energy Production and Storage with MADNESS and MPQC	Chemistry	150,000,000
Mat_Design_esp	Argonne National Laboratory	Larry Curtiss	Materials Design and Discovery: Catalysis and Energy Storage	Materials Science	50,000,000
MultiscaleMolSim_esp	The University of Chicago	Gregory Voth	Multiscale Molecular Simulations at the Petascale	Biological Sciences	150,000,000
NAMD_esp	The University of Chicago	Benoit Roux	NAMD - The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based on a Polarizable Force Field	Biological Sciences	80,000,000
PlasmaMicroturb_esp	Princeton Plasma Physics Laboratory	William Tang	Global Simulation of Plasma Microturbulence at the Petascale & Beyond	Physics	70,000,000
TurbChannelFlow_esp	The University of Texas at Austin	Robert Moser	Petascale Direct Numerical Simulations of Turbulent Channel Flow	Engineering	60,000,000

Project Name	PI Institution	PI Name	Short Title	Science Field (Short)	Allocation Amount
TurbNuclComb_esp	The University of Chicago	Don Lamb, Carlo Graziani	Petascale Simulations of Turbulent Nuclear Combustion	Physics	150,000,000
				Total Mira DD ESP	2,040,000,000

Appendix C – Strategic Results Slides

Better Cancer Drugs with Quantum Chemistry

Anatole von Lilienfeld and Larry Curtiss, Argonne

Impact and Approach

- Improve the effectiveness of ellipticine – a promising drug for treatment of uterine cancer
- Determine why the drug only works intermittently
- Nature of the binding requires many-body quantum theory. Quantum Monte Carlo (QMC) is the only scalable approach.

Accomplishments

- Characterized the binding energies, providing critical input for improved modeling
- Proved that many molecules thought too complicated or too large for quantum chemistry can be tackled with QMC
- **First time simulations of this type have been completed**

Sped-up QMCPACK 2.68x time-to-solution

ALCF Contributions

- **2.7x speed-up was driven heavily by good use of the quad-vector FP unit. Reducing time-to-solution makes an accurate, but very expensive technique tenable.**

Presented: Quantum Monte Carlo in the Apuan Alps VIII

Publication: Submitted to Physical Review Letters.

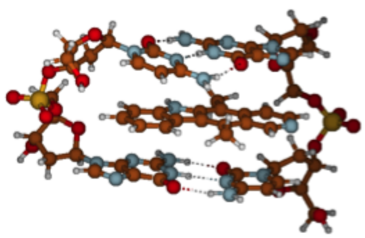

Ellipticine is in the middle binding to the DNA of the cancer cells.

Argonne Leadership Computing Facility

18

Exploring Biophysical Mechanism for Intracellular Messaging

Benoit Roux, University of Chicago

Impact and Approach

- Assist drug design for neurodegenerative diseases such as Parkinson's disease
- Advanced workflow using replica exchange molecular dynamics, visualization and analysis

Accomplishments

- Elucidated calcium ion pumping mechanism by simulating conformational transition pathway
- Overcame time-scale bottleneck of biomolecular machine simulation by harnessing massively parallel resource of Mira

ALCF Contributions

- 20% speed-up integrating PAMI into Charm++/NAMD
- Implemented NAMD API to carry out replica exchange
- Optimized run-time environment variables for NAMD and debugging inter-copy synchronization of NAMD ensemble run

Presented: National Meeting of American Chemical Society , Sep. 2013

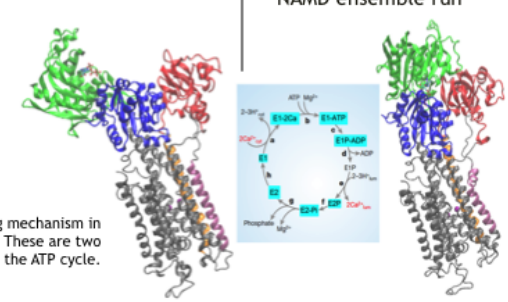

Calcium ion pumping mechanism in two conformations. These are two end states in the ATP cycle.

Argonne Leadership Computing Facility

19

Decreasing Jet Engine Noise

Umesh Paliath, G.E. Global Research

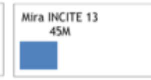

Impact and Approach

- Reduce the noise from jet engines, improving safety and efficiency
- Created a numerical wind tunnel that provides detailed large datasets that are enabling improved understanding of jet mixing noise generation mechanism

Accomplishments

- GE has demonstrated the ability of the LES approach to predict the acoustic signature from a canonical nozzle to complex exhaust nozzles with noise control devices such as chevrons/lobed mixers

ALCF Contribution

- Findings from this research will significantly improve the prediction and design capabilities for next-generation aircraft engines, both from demonstrating the viability of LES as a characterization tool and as a source of physics guidance.

GE is leveraging the LES +HPC tool as a numerical wind tunnel to provide insight into the jet noise generation mechanism from turbulent free shear layer flow.

Award: 2013 winner of International Data Corporation's HPC Innovation Excellence Awards

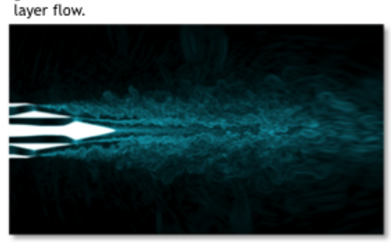

20

Improve Wind Turbine Design

Giridhar Jothiprasad, G.E. Global Research

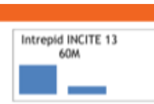

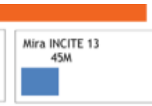

Impact and Approach

- Reduce the noise from jet engines, improving safety and efficiency
- Created a numerical wind tunnel providing detailed large data sets that are enabling improved understanding of jet mixing noise generation mechanism

Accomplishments

- Predicting separated flow aerodynamics at a near-stall flow angle
- Improved LES Wall-models to enable computational investigation of large span noise-reduction concepts

ALCF Contribution

- Onsite ALCF machines allow for post-processing of very large, time accurate full 3-D data sets to provide insight into flow physics

Award: 2013 winner of International Data Corporation's HPC Innovation Excellence Awards

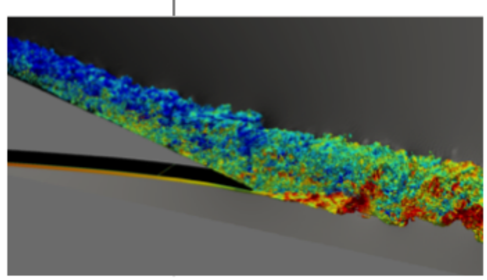

LES of the DU96 airfoil at Re = 1.5 million and flow angle of 10.3 degrees. Flow separating from the airfoil visualized by Q-criterion iso-surfaces colored by the vertical velocity

Protons and Neutrons Get Social in the Nucleus

Steven Pieper, Argonne

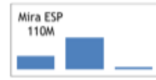

Impact and Approach

- Sets the stage for neutrino-nucleus calculations relevant for neutrino detector calibration and supernovae explosion mechanism
- Predict the results of a recent experiment at Jefferson Lab, which are not yet released

Accomplishments

- Computed charge form factor and sum rules for the electromagnetic response of ^{12}C
- The two-body meson exchange currents are important for agreement with existing data

ALCF Contributions

- Implemented physics routines key to calculation
- Optimized to 32,000 processes with an efficiency above 95%
- Optimized best MPI/OpenMP balance

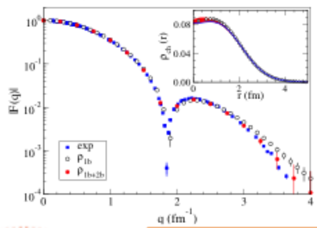

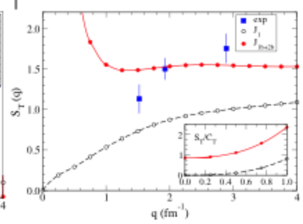

(left) Longitudinal form factor: Two body terms in the density operator bring theoretical prediction closer to experimental data in the high-momentum transfer tail (q).

(right) The transverse sum rule. Two-body contribution is large for all momentum transfers. Good agreement with experiment.

A. Lovato, S. Gandolfi, Ralph Butler, J. Carlson, Ewing Lusk, Steven C. Pieper, R. Schiavilla, "Charge Form Factor and Sum Rules of Electromagnetic Response Functions in 12C," Physical Review Letters, August 2013, vol. 111, no. 9, The American Physical Society, 2013 10.1103/PhysRevLett.111.092501

Argonne Leadership Computing Facility

22

Visualizing Adaptive Grids

Michael E. Papka, Argonne

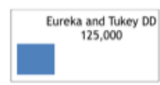

Impact and Approach

- Scalable mechanisms to visualize Adaptive Mesh Refinement (AMR) simulations are critical in order to gain insights from the large-scale datasets being generated by these simulations
- Enable visualizations of the scale expected from future supercomputers
- Special subdividing of data, static load balancing, and improved use of data locality

Best Paper Award: N. Leaf, V. Vishwanath, J. Insley, M. Hereld, M. E. Papka, K.L. Ma, *"Efficient Parallel Volume Rendering of Large-Scale Adaptive Mesh Refinement Data"*, In the Proceedings of the IEEE Symposium on Large Data Analysis and Visualization (LDAV), Atlanta, Georgia, USA, October 2013

Accomplishments

- Developed a new approach to AMR visualization that scales comparable to regular grid renderers
- Identified optimal decomposition and views

ALCF Contribution

- The team developed the algorithms for AMR volume rendering and the parallel implementation for GPU-based clusters
- This is joint project with University of California at Davis, University of Chicago and Argonne National Laboratory

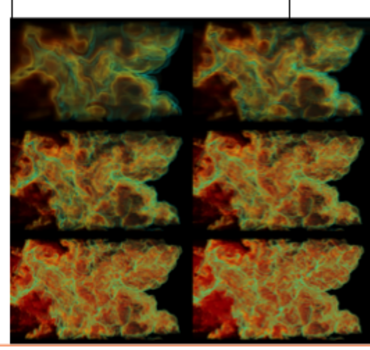

An analytical estimate of the flame front in a Rayleigh-Taylor Flame and the effect of additional refinement levels. The green-blue color highlights the flame surface, and the yellow-orange-red transition shows the fuel-ash mixture, which the flame front leaves behind.

Argonne Leadership Computing Facility

23

Lattice QCD
Paul Mackenzie, FNAL

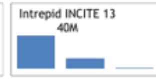

Impact and Approach

- Provide crucial theoretical inputs for searches of new physics beyond the Standard Model of particle physics
- Determine a number of underlying parameters of the Standard Model
- Importance sampling of QCD vacuum configurations discretized on a 4-D space-time lattice

Accomplishments

- MILC subproject has been generating new high- resolution QCD configurations since 2010
- Mira has enabled these configurations to reach physical quark masses
- First physics results published in April contain most precise value of ratio of kaon to pion decay constants to date

ALCF Contributions

- ALCF catalyst and ESP postdoc optimized linear solver using QPX intrinsics, low-level SPI communications and optimal mappings
- ALCF catalyst implemented new algorithm for fermion force routine, giving 7x speedup for that routine

Total improvement in time-to-solution is about 2x

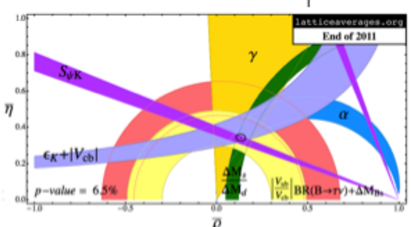

Global fit of CKM matrix elements combining experimental and lattice QCD results. (Laiho, Lunghi, & Van de Water, Phys. Rev. D81 034503, 2010)

MILC Collaboration, Phys. Rev. Lett. 110, 172003 (2013).

Open Source Compiling for Supercomputers
Hal Finkel, Argonne

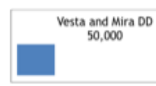

Impact and Approach

- Provides a development environment on BG/Q systems similar to the environment on standard laptops
- Influences the direction of C++ standards and implementations for HPC
- In-house compiler expertise for better leverage in the compiler community

Accomplishments

- Released a LLVM/Clang version with advanced, modern features for the BG/Q
- Implemented QPX auto-vectorization (including C++)
- Full OpenMP 3.1 support
- Address sanitizer
- Advanced MPI-specific warnings

ALCF Contributions

- ALCF develops and maintains the bgclang
- ALCF partnered with Pathscale to update Intel's open source run time library to latest CLANG and OpenMP versions

Select bgclang Users

- GROMACS (de Pablo, INCITE 2014)
- HACC (Habib, INCITE 2013-2014)
- Elemental (library used by 1 INCITE project)
  - (The Elemental library with bgclang has already been used for a 2013 paper in the SIAM Journal on Scientific Computing)
- Rosetta (Baker, INCITE 2013)
- OpenFOAM