Argonne NATIONAL LABORATORY

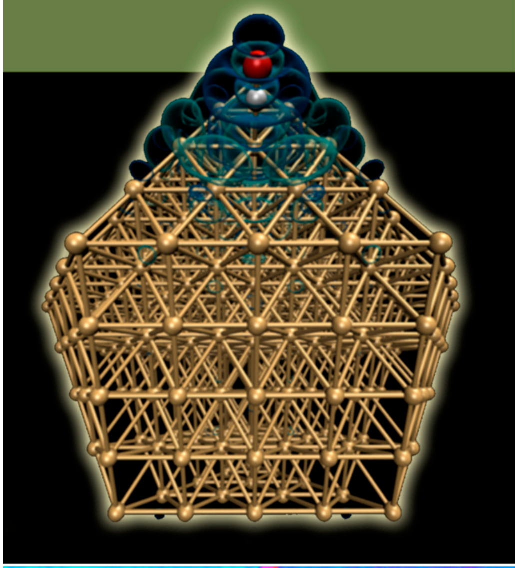

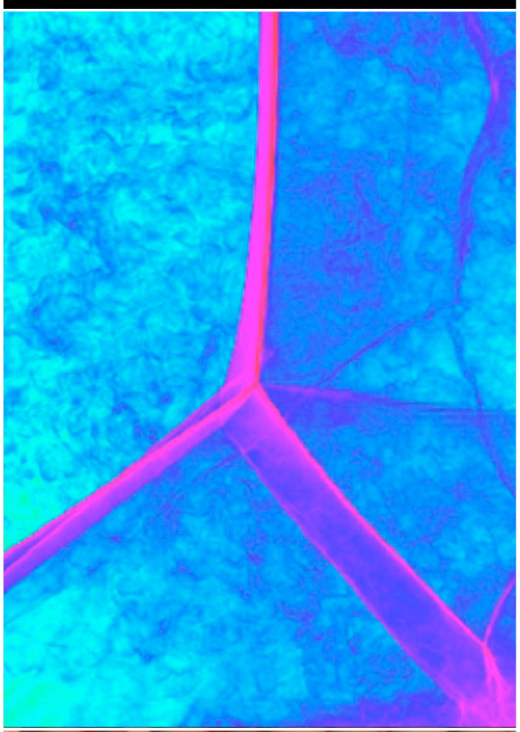

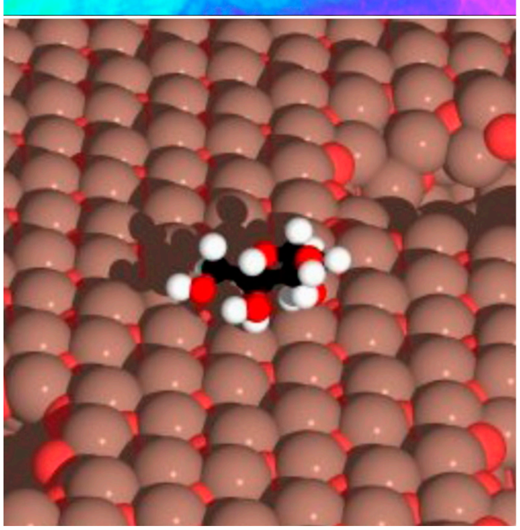

# High-Performance Computing Facility

## Operational Assessment Report

April 2012

Argonne Leadership Computing Facility

## On the cover

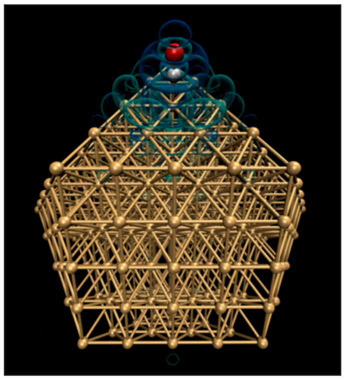

PI: Jeffrey Greeley, Argonne National Laboratory

INCITE Project: Electronic Structure of Gold Nanocluster

The charge density difference of a carbon monoxide molecule – one oxygen atom and one carbon atom — adsorbed (adhered to the surface) on a gold nanoparticle of 309 atoms.

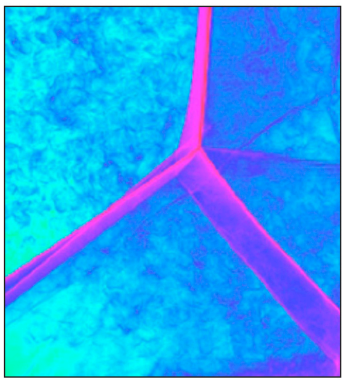

PI: Alexei Khokhlov, The University of Chicago

INCITE Project: Simulations of Deflagration-to-Detonation Transition in Reactive Gases

Temperature in a 3-D Navier Stokes DNS simulation of Mach=3 reflected shock bifurcation in a H-O mixture in a square channel.

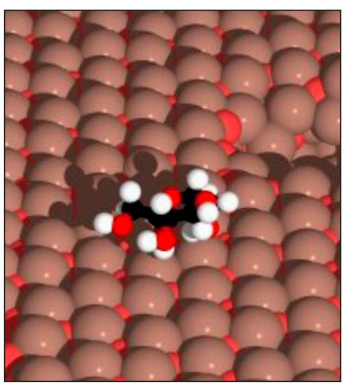

PI: Larry Curtiss, Argonne National Laboratory

INCITE Project: Designing Materials from First Principle Calculations

Thirty-three atom silver cluster that is being studied as a new catalyst for propylene epoxidation.

# Contents

## Contents (Cont.)

# Contents (Cont.)

# Figures

# Tables

*This page intentionally left blank.*

# Executive Summary

The Argonne Leadership Computing Facility (ALCF) is one of two leadership-computing facilities supported by the U.S. Department of Energy. The ALCF provides the computational science community with a world-class computing capability dedicated to breakthrough science and engineering. It began operation in 2006 with its team providing expertise and assistance to support user projects in achieving top performance of applications and maximizing benefits from the use of ALCF resources.

The major resource at the ALCF is Intrepid, a 40-rack IBM Blue Gene/P system, and two development systems named Challenger and Surveyor. Intrepid possesses a peak speed of 557 teraflops, while Challenger and Surveyor are each approximately 14-teraflops systems used for tool and application porting, software testing and optimization, and systems software development. We have currently begun deploying the ALCF's next system named Mira, an IBM Blue Gene/Q system that will surpass 10 petaflops, with the equivalents to Challenger and Surveyor already on site and named Vesta and Cetus.

ALCF supports more than 650 active users (and over 120 active projects) from universities, national laboratories, and industry as of February 2012. More than 28% of our active users in CY2011 have responded to our survey, continuing to give us high marks in overall satisfaction, problem resolution, and user support. We were able to address 90% of our trouble tickets in three business days or under.

ALCF has had an outstanding year, meeting or exceeding all metrics set for the Facility. ALCF has delivered more than 1.2 billion core hours of compute time between January 1, 2011 and December 31, 2011, with more than 680M of those core hours being used by capability jobs (a job using 20% or more of the machine). During that same time period, the science accomplished on the machine produced over 120 publications in CY2011, including nine in *Physical Review Letters* and two in *Nature.*

The annual Operational Assessment Review (OAR) of ALCF by the Advanced Scientific Computing Research (ASCR) program provides the Facility with an opportunity to receive external feedback on ways to improve the ALCF's operation. The review takes into consideration agreed-upon metrics and reports describing the operation of the Facility. The report is organized into six sections. The remaining sections address the 2011 OAR metrics and present User Support Results, Business Results, Strategic Results, Innovation, Risk Management, and a Summary of the Proposed Metric Values for Future OAR's.

| CY 2011 Performance Metrics | | | | | |
|---|---|---|---|---|---|
| **Area** | **Metric** | | **Target** | **Projected** | **Actual** |
| User Results | CY 2011 User Survey | Overall Satisfaction | 3.5/5.0 | - | 4.4/5.0 |
| | | User Support | 3.5/5.0 | - | 4.5/5.0 |
| | | Problem Resolution | 3.5/5.0 | - | 4.5/5.0 |
| | | Response Rate | 25% | 25% | 29.0% |
| | % user problems addressed within three working days | | 80% | 92% | 90.2% |
| Business Results | Intrepid Overall Availability | | 90% | 92% - 94% | 94.0% |
| | Intrepid Scheduled Availability | | 95% | 97% - 99% | 97.8% |
| | Intrepid Capability Usage (old metric) | | 300M | 660M – 685M | 685M |
| | Intrepid Capability Usage (new metric) | | - | 40% | 57.0% |

**Table ES-1: Summary of the Target, Projected, and Actual Data for the Previous Year (2011) Metrics.**

# Section 1. User Support Results

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

## *ALCF Response*

As documented below, ALCF has effective processes for supporting customers, resolving problems, and performing outreach.

## *User Services Overview*

In calendar year 2011, User Services comprised four main areas:

- Managing user communications,
- Coordinating outreach and education efforts,
- Providing a help desk,
- Stewarding key demographic information and surveys.

### User Communications
In the last year, User Services primarily handled email, print, and web messaging to users. Duties included the production of user documentation, various high-profile communications, presentations, and print and presentation materials associated with the call for INCITE proposals. Types of communications managed were weekly emails, monthly and quarterly newsletters, an annual report, an annual science highlight brochure, an industrial outreach brochure, and project tracking reports.

### Outreach Activities
Workshops and facility tours were the primary focus. In addition, User Services also managed industry outreach and lab-wide coordinated interactions with companies such as Dow Chemical. A monthly user call was also part of the outreach activities. User Services conducted individual initial conference call welcoming each new INCITE and ALCC users to the facility. User Services worked closely with the newly formed User Advisory Council (defined later in the document) to provide a sounding board for vetting critical path technology, as well as to represent the variety of users within the ALCF facility.

### Help Desk
ALCF maintains a help desk. User Services is provided via email and phone. If a user is local or visiting Argonne, User Services provides in-person support. User Services also manages issue tracking and resolution in a system called Request Tracker (RT.) In addition, User Services is responsible for escalating issues to subject matter experts within ALCF.

### User Demographics and Surveys
User Services stewards key user data in the USERBASE database. ALCF members and ALCF leadership often request aggregated information from this database through User Services. User Services would also participate in division-wide efforts in aggregating information

from various other data sources. User Services also manages the annual surveys as well as workshop and other ad hoc surveys.

### *User Services Results*

A new website was rolled out in November 2011 to merge the ALCF user documentation and general website under one technology and one place to edit content. In response to dips in the addressed ticket metric from September through November, a ticket review process was revisited and modified within the User Services group. The mission of User Services has changed to become more user-focused – the concept of "user experience" has been introduced and goals are set against this conceptual model. Content creators in User Services have been tasked with developing stories from ALCF users during the arc of ALCC or INCITE projects. This is different from the previous model of being driven by external deadlines such as DOE's ASCR Computing News Roundup newsletter or specific communications pieces such as the science highlights.

### *Operations Support Results*

ALCF Operations developed a project administration tool that allows dissemination of information to users over submission control to very fine granularity. It give the user the information needed to remedy a block on the user's project. ALCF also improved the file system by increasing disk space available and increased reliability of the home directory. Home directories were mounted read-only on the compute nodes so that users were encouraged to use the high-speed, high-capacity scratch file system. Finally, ALCF Operations has improved the resiliency of the Cobalt scheduler during GPFS outages—specifically the scheduler continues to be responsive, and fewer jobs are lost in failures.

### *Application Support Results*

Over the last year, the ALCF team worked closely with many projects to provide code optimization and scaling, compiler support, and job scheduling. Below are a few highlights salient to ALCF's mission.

#### MILC

ALCF staff worked with members of the MILC Collaboration to design a new interface to allow the MILC code to call the LQCD SciDAC libraries that contain important routines used in their simulations with "HISQ" quarks. While looking into optimizing the SciDAC libraries, ALCF staff discovered a new way to calculate the HISQ "fermion force," which combines multiple terms together in a much more efficient manner than the previously optimized version. The new version reduces the number of flops needed by more than a factor of ten. This version also has a slightly better parallel implementation so that the reduction in flops easily translates to a reduction in runtime by more than a factor of ten. The HISQ fermion force is typically responsible for 10% to 50% of the total execution time, depending on the run parameters.

ALCF staff also used a similar strategy to rewrite the existing gauge force routine in the SciDAC libraries. The new version is approximately 2-3 times faster than either the previous SciDAC version or the existing MILC code. The gauge force is usually only about 10% of the total execution time. For completeness, the ALCF staff also added a gauge action routine to the SciDAC libraries to aid the development of new simulation code.

The new routines are now being tested by the MILC Collaboration and should go into their production runs soon. The new code is also being used by the BSM (Beyond Standard Model) Collaboration, who is also currently doing simulations using HISQ quarks on Intrepid. These simulations are currently being conducted using the MILC code; however, the simulation algorithms in MILC were designed for QCD, with two light quarks, while the BSM group is currently performing simulations with eight light quarks. To better support simulations of theories other than QCD, a subgroup of researchers in the USQCD Collaboration have started developing a new, lightweight simulation framework that will take advantage of the existing LQCD SciDAC libraries, while allowing rapid development of new simulation algorithms for QCD and other theories relevant to Beyond Standard Model physics. ALCF staff member James Osborn has been the main developer of this effort and has an initial prototype of the code running on Intrepid. ALCF staff is now beginning to test and tune the simulation algorithms in the new framework to make them more efficient for the eight quark simulations that are in progress.

**QMC**

ALCF staff worked with the main QMCPACK developer, Dr. Jeongnim Kim, to port the QMCPACK code to the Blue Gene/P. QMCPACK is a quantum Monte Carlo code for calculating the electronic structure of molecules and condensed phase systems. QMCPACK is written in C++ and makes extensive use of templates. At the moment, we are encountering a number of errors on the code that appear to be Blue Gene/P-specific. A Project Management Request ticket with IBM was filed in order to investigate potential compiler issues. At the moment, this is still an open issue that will require additional work through the end of CY2011 and beyond.

The QMCPACK code makes heavy use of omp_get_wtime() as a thread-safe timer; ALCF staff found via profiling that more than 10% of the wall time of large jobs was spent in this routine. Both the GNU and IBM XL OpenMP implementations of this function were using gettimeofday(), which took more than 4,000 cycles per call, when it should take less than 100. ALCF staff filed a PMR on this performance bug and described the solution to IBM (reading the timer register without a system call). IBM is integrating it into the company's OpenMP runtime.

**Nek5000**

User Tamay Ozgokomen encountered an I/O problem with the Nek5000 code in which incomplete output data was being written to an output file upon program completion. ALCF staff investigated the problem and found it to be caused by the fact that the output file was not properly closed at program termination, resulting in loss of buffered output. Incorrect closure was due to use of the C exit() routine with the Fortran XL compilers, which does not allow the Fortran runtime system to properly flush I/O buffers. ALCF staff proposed two solutions: use of the Fortran flush_I/O() routine, or use of the XLF exit_(), which were incorporated into the code to resolve the problem.

A second issue encountered with the Nek5000 code involved incorrect results obtained with recent versions of the code when running an input case setup under a previous code version. The issue was debugged by ALCF staff members, who found it to be caused by an inconsistent declaration of 32- and 64-bit integer variables between routines. Variable declarations were corrected to be uniformly 64 bit, and the problem was resolved.

**Exascale Tools Workshop**

ALCF staff presented at the 2011 Exascale Tools Workshop in Annapolis, Maryland on October 13-14, 2011. This talk provided an overview of the current state of performance and correctness tools on ALCF systems, along with an assessment of user needs, and plans and needs for tools on future systems. The presentation emphasized the need for robust, scalable, and easy-to-use tools that better guide users in identifying and resolving application bottlenecks.

**Single Rack Mesh Ensemble**

For some time, ALCF has supported bundling multiple back-end runs into a single Cobalt job to facilitate capability-class ensemble runs. Like normal job runs, these are subject to "passthru" wiring constraints: In order to use the torus network, at most one rack in each odd/even pair may run a 1-rack job; and at most one 4-rack run may be done in each 8-rack row. However, a run may also use mesh communication instead of torus. While communications performance may be reduced for mesh compared to torus, there are no wiring constraints on mesh runs. For example, a capability job of 8 racks can run at most four 1-rack torus jobs within it; whereas it can run a full complement of eight 1-rack mesh jobs. Mesh runs have not previously been supported, but a new version of cobalt-subrun (appropriately named cobalt-subrun-mesh) has been developed to enable them. This made it possible for the Turbulent_Mixing INCITE (user Santhosh Shankar) project to run efficient capability-class jobs to exhaust the project's remaining 2011 INCITE allocation in December.

## *ALCF Support Metrics for 2011*

Tables 1 and 2 show a summary of all user support metrics. Further details are provided in the respective metric sections.

| 2010 to 2011 Metric Comparison | | Value for 2010 | Target for 2011 | Actual for 2011 |
|---|---|---|---|---|
| **Number Surveyed** | | 847 | N/A | 959 |
| **Number of Respondents (response rate)** | | 249 (29.4%) | (25%) | 278 (29.0%) |
| | | | | |
| **Overall Satisfaction** | Mean | 4.4 | 3.5 | 4.4 |
| | Variance | 0.5 | N/A | 0.5 |
| | Std Dev | 0.7 | N/A | 0.7 |
| | | | | |
| **Problem Resolution** | Mean | 4.4 | 3.5 | 4.5 |
| | Variance | 0.5 | N/A | 0.4 |
| | Std Dev | 0.7 | N/A | 0.6 |
| | | | | |
| **User Support** | Mean | 4.4 | 3.5 | 4.5 |
| | Variance | 0.5 | N/A | 0.5 |
| | Std Dev | 0.7 | N/A | 0.7 |
| | | | | |
| **Tickets Addressed in 3 days (business)** | | 82.3% | 80% | 90.2% |

**Table 1: Comparison of Key Metrics between 2010, the Targets, and 2011.**

| 2011 Metrics By Program | | INCITE | ALCC | INCITE + ALCC | Discretionary | All |
|---|---|---|---|---|---|---|
| **Number Surveyed** | | 461 | 50 | 511 | 438 | 959 |
| **Number of Respondents** | | 168 | 19 | 187 | 91 | 278 |
| **Response Rate** | | 36.4% | 38.0% | 36.6% | 20.8% | 29.0% |
| **Overall Satisfaction** | Mean | 4.5 | 3.9 | 4.4 | 4.4 | 4.4 |
| | Variance | 0.4 | 0.7 | 0.5 | 0.6 | 0.5 |
| | Std Dev | 0.7 | 0.8 | 0.7 | 0.8 | 0.7 |
| **User Support** | Mean | 4.5 | 4.4 | 4.5 | 4.4 | 4.5 |
| | Variance | 0.4 | 0.5 | 0.5 | 0.5 | 0.5 |
| | Std Dev | 0.7 | 0.8 | 0.7 | 0.7 | 0.7 |
| **Problem Resolution** | Mean | 4.5 | 4.4 | 4.5 | 4.4 | 4.5 |
| | Variance | 0.4 | 0.5 | 0.4 | 0.4 | 0.4 |
| | Std Dev | 0.6 | 0.7 | 0.7 | 0.6 | 0.6 |
| **All Questions** | Mean | 4.3 | 4.1 | 4.3 | 4.2 | 4.3 |
| | Variance | 0.6 | 0.9 | 0.7 | 0.7 | 0.7 |
| | Std Dev | 0.8 | 0.9 | 0.8 | 0.8 | 0.8 |

**Table 2: Comparison of Key Metrics between Different Allocation Programs.**

## User Survey

In 2010, ALCF sent separate surveys to INCITE/ALCC and Discretionary users. The 2011 ALCF user survey is now a single instrument sent to all active users, PIs, and CoPIs. As users answer questions in the survey, the questions they answer reveal or hide questions later in the survey. For instance, Director's Discretionary allocations do not have a Catalyst assigned to them. Therefore, when the user answers "No" to the question "Did you have a Catalyst assigned to you?" further questions about Catalysts are skipped/hidden from their view. At the recommendation of the DOE Program Manager, ALCF gathered the questions for the User Support metric and the Problem Resolution metric into clearly defined sections.

ALCF leadership, in concurrence with the DOE program manager, chose to wait until 2012 to re-instrument the survey. Changing the questions used to evaluate the OMB metrics requires external oversight from experts in survey instrumentation and analysis. ALCF consulted on-site expertise several years ago and now plans to engage regional experts in survey methods.

ALCF reorganized the questions in the survey and added several new questions. The core set of questions used for metrics last year is largely unchanged. ALCF made one notable change.

*"The amount of communication and support from my Catalyst in 2011 was: (just right/not enough)"*

This question seemed redundant with the other questions in the Catalyst section of the 2011 survey as well as the problem resolution section. It was also part of the metric; however, the way this was added to the computed average was not appropriate. The question was eliminated with the guidance of the DOE program manager.

For more information about how the survey changed, please see the 2011 survey changes in Appendix A.

***Survey Response Results***

| Survey sent to 959 users, PIs, and Co-PIs | 2010 Response Rate (%) | 2011 Response Rate (%) | Target response rate (%) | Total # of responses |
|---|---|---|---|---|
| Response rate | 29.4% | 29.0% | 25% | 278 |

Significant efforts were made to exceed the survey response rate. This was done through a series of 14 email reminders sent over the course of a month. User Services sent specifically tailored messages to PIs. In the final week, Catalysts directly encouraged their collaborators to respond. The ALCF Division Director also reached out to researchers for key projects directly via telephone.

### 1.1 User Support Metrics

| Survey Area | 2010 Target | 2010 Actual | 2011 Target | 2011 Actual |
|---|---|---|---|---|
| **Overall satisfaction rating** | 3.5/5.0 | 4.4/5.0 | 3.5/5.0 | 4.4/5.0 |
| **Average of User Support ratings** | 3.5/5.0 | 4.4/5.0 | 3.5/5.0 | 4.5/5.0 |

Overall Satisfaction and User Support ratings in this metric did not measurably change from 2010. This may be in part due to the previous year's high score of 4.4 for both overall and average ratings. In order to detect trends or changes in user opinion, the questions measured did not change from 2010.

### 1.2. Problem Resolution Metrics

| Survey Area | 2010 Target | 2010 Actual | 2011 Target | 2011 Actual |
|---|---|---|---|---|
| **% of problems addressed in 3 working days** | 80% | 82.3% | 80% | 90.2% |
| **Average of problem resolution ratings** | 3.5/5.0 | 4.4/5.0 | 3.5/5.0 | 4.5/5.0 |

The problem resolution ratings in the survey did not measurably change from 2010. This may be in part due to the previous year's high score of 4.4 or that ALCF put most of its effort in improving the addressed metric. Again, the questions measured did not change from 2010.

The ALCF has a specific process in the problem tracking system for addressed tickets. An ALCF staff member modifies a special field, called addressed, when the definition of addressed is met. A ticket is "addressed" when the following is true: The ticket has been accepted by a staff member, the problem has been identified, the user has received notification of this, and the staff member is working on a solution or has found one.

In 2011 ALCF made it a goal of User Services to achieve this metric. At first, there was an informal process created where ALCF Help Desk staff created nag scripts that would remind ALCF staff members and their supervisors of tickets older than 2 business working days.

This worked initially, but when ALCF experienced a particularly busy time (e.g., the time around Supercomputing 2011), the percentage started to trend down and dipped below 90%. User Services then created a system in which Help Desk staff members would formally visit staff who had not set a ticket to "assigned" after 1 business day. They would see if the ALCF staff member needed assistance or was incorrectly assigned to the ticket in question.

Table 3 shows the difference in categorization of tickets between 2010 and 2011. Ticket categorization within RT, our ticketing system, was implemented only in June 2011. Therefore, the YTD numbers in the OAR 2011 report were an assumption of how tickets were categorized, based on the internal ALCF group member who handled those tickets. After the implementation of ticket categorization within the ticketing system, ALCF staff members were told to retroactively categorize their tickets for CY2011, based on the new system. The current CY2011 categorization of tickets provides more accurate data for ticket categorization.

| Categories | CY2010 | CY2011 |
|---|---|---|
| Access | 1426 | 750 |
| Accounts | 781 | 1397 |
| System | 312 | 312 |
| Misc | 317 | 222 |
| Applications Software | 327 | 217 |
| Visualization & Data Transfer | 168 | 108 |
| I/O & Storage | 91 | 104 |
| Bounces | N/A | 220 |
| **TOTAL TICKETS** | **3422** | **3330** |

**Table 3: Difference in Categorization of Tickets between 2010 and 2011.**

## 1.3 User Support and Outreach

### 1.3.1 Creation of the User Advisory Council

In mid-2011, the ALCF User Advisory Council (UAC) was formed in response to feedback from both the 2010 OAR and at the request of the Division Director. The UAC consists of 8-10 members, roughly distributed by project allocation percentages (60% INCITE, 30% ALCC, 10% Discretionary.) The UAC meets monthly and provides critical feedback to User Services about technical changes, engages in high-level discussions of services and documentation, and acts as a barometer to measure user sentiment.

Since then, the UAC impacted the facility in the following ways:

- Recommended changing the format of INCITE intro calls,
- Asked for clarity in the INCITE application process,
- Recommended eliminating monthly user calls in lieu of subject-specific webinars,
- Recommended changes in the winter workshop,
- Vetted the Challenger configuration and scheduling policies,
- Provided key feedback on the use of IBM and other vendor performance tools,
- Continues to provide recommendations for the ALCF website.

### 1.3.2 ALCF Website Merge and Reorganization

ALCF leadership revamped the organization of the website and documentation in the second half of 2011. The first part of the initial phase of the project was to merge three different content management systems (custom HTML pages, WordPress, and MediaWiki) under the single platform of Drupal. The second part involved training staff, installing critical modules, and refining workflows, such as editing and approval of content to the website. ALCF is now focusing on improving the quality and depth of the web-based documentation.

As part of the initial phase of the website project, ALCF designated key personnel to be responsible for various sections of the website. These designations corresponded with areas of expertise. ALCF managers were placed as the approvers of all published data on the website.

In 2011, ALCF web developers created workflows to enable the human resource (HR) representative in ALCF to be the primary custodian of the staff profiles. The HR staff member is now responsible for collecting and maintaining the photo, contact information, area of expertise, and short biographies for all staff in ALCF.

The primary reason ALCF moved to Drupal is to provide a managed workflow for published content. ALCF now is able to assign each page of the web site to a specific staff member. This ensures that content being published is relevant, and the information being posted is correct.

Another important feature of Drupal is the ability to provide a web presence based upon taxonomies and tagging. The power of the taxonomy is the fluidity in the management of knowledge on the website. Instead of organizing a website using file and folder structures, sections of text within a database are "tagged" with terms (i.e., what Drupal calls "vocabulary") that are then grouped into taxonomies. These taxonomies provide different ways to display data to the end user, depending on context. This allowed ALCF to merge different types of web sites together and share data among them. It also enabled flexible reorganization and redesign without rewriting text or redesigning pages by hand.

The Drupal content management system (cms) has an extensive and continually developing collection of community modules that can be added to the framework. Not only is ALCF collaborating with and contributing to the open source Drupal community, ALCF will also be part of the laboratory-wide efforts in Drupal. Argonne's revamped public website (launching in March 2012) is being developed in Drupal, as are revised websites for several other ANL divisions.

ALCF is evaluating the following key modules:

ShareThis and Facebook Like–These social media modules enable users to "share" stories and documentation in social media like Facebook, Twitter, LinkedIn, email, etc. This connects and grows the ALCF online community.

Media–This module provides ways for ALCF to manage digital assets and build image and video galleries of the users' science visualizations.

Search 404 and Redirect Module–These modules aid visitors in finding information when they are pointed to the ALCF site with outdated links.

Biblio–This publication module provides an easier way of maintaining a publications list for both users and ALCF staff by allowing bulk import of standard publication formats such as BibTex, EndNote, RIS, etc. and providing multiple citation outputs such as Chicago, CSE, and IEEE.

### 1.3.3 2011 Workshops

The most recent workshop in 2011 was a fall Getting Started Workshop. The increase in industry collaboration drove the demand for this workshop. Much like the January 2011 Getting Started Workshop, points covered included an overview of the ALCF, allocation and access information, basics of system architecture, a primer on debugging, and discussion of tools and tips for using them. Part of the workshop time was for users to start working on Intrepid, the ALCF's Blue Gene/P system, with ALCF staff providing hands-on assistance.

For all workshops hosted by ALCF, participants are asked to fill out a survey near the end of the event. For the Getting Started Workshop, ALCF received largely positive feedback on format, topics, material, and resources available at the ALCF. Participants were asked to write about how ALCF could improve workshops. Feedback received was converted into action items for the 2012 Winter Workshop. Here were the most common recommendations:

- More hands-on work [less lecturing],
- Example problems available prior to the workshop and more complex examples to work through,
- Accounts were not working [on the critical first day],
- Subjects were too broad–ALCF should break down this event into small workshops,
- Make presentation materials available prior to the event.

Listed below are all of the workshops hosted at ALCF in 2011. (All workshops except for the fall Getting Started Workshop were reviewed in the August 2011 OAR report.)

- Getting Started Workshop 2011 – January 18
- Productivity Tools for Leadership Science – January 18-19
- Proposal Writing Webinar 2011 – January 24
- Gordon Bell Workshop 2011 – March 8-11
- Leap to Petascale 2011 – June 7-9
- Getting Started Workshop 2011 (Fall) – October 4-5

### 1.3.4 Key Strategic Outreach and Partners

ALCF focused efforts in 2011 on seeking and fostering relationships with industry. ALCF participated in an Argonne-wide strategic effort with Dow Chemical. In addition, the ALCF engaged with the following organizations throughout 2011: Boeing, GE, Caterpillar, John Deere, Eli Lilly, Navistar, UOP/Honeywell, and the Global Midwest Alliance.

To increase reach and distribute focus, ALCF brought in a new lead for User Services. The previous manager was then freed up to increase engagement with industry partners. The new User Services lead is responsible for stewarding the academic partnerships and increasing relationships with computational science-knowledgeable academic and regional institutions worldwide.

### 1.3.5 Streaming Media and Workshop Archiving

ALCF experimented recently with Adobe Connect at a large event. The results were not impressive–bad audio, poor lighting, and content not designed for passive viewing.

The User Advisory Council also commented that efforts would be better spent on improving documentation than developing videos. One specific member felt that working through a 20-40 minute video would not be worth the time investment. They preferred spending that time searching through documentation and working with code snippets or example problems. This follows a rule developed at ALCF around conference calls– people do not want to wait 30 minutes for 30 seconds of information.

### 1.3.6 2012 INCITE Welcome Call Format Change

Based upon feedback from the User Advisory Council, ALCF changed the nature of the initial user conference calls. The call is now science-centric: the participants discuss the story of their research and their upcoming challenges and goals. From this main storyline, participants then dive into relevant user issues that will block the progress of their science. ALCF expects this to have qualitative effects on the newer collaborations and may decrease ramp-up time in the facility.

### 1.3.7 2011 Publications

The latest publication developed in 2011 was an industry-targeted brochure. This promotional flyer discusses the merits of engaging ALCF and promotes the use of the INCITE and ALCC programs in private industry.

The following publications schedule was already discussed in the August 2011 OAR report but is included here to document the breadth of publications ALCF User Services produces. The publications are tailored to address targeted audiences ranging from those who are very knowledgeable about computational science (e.g., researchers, scientists) to those who need just the basics (e.g., Congress, outside media, students who visit the ALCF.)

| Publication | Frequency | When |
|---|---|---|
| ALCF Weekly Update E-mail | Weekly | Thursday |
| Highlights for ASCR Computing News Roundup | Monthly | 1st Week |
| ALCF Newbytes Newsletter | Quarterly | Jan, Apr, Jul, Oct |
| ALCF Fact Sheet | As needed | As needed |
| ALCF Annual Report | Yearly | Mar |
| ALCF Science Highlights | Yearly | Sep |
| ALCF Division Brochure | Yearly | Oct |
| ALCF Calendar | Yearly | Nov |
| ALCF User Packets | As needed | As needed |
| ALCF Industry Brochure | Semi-Annually | Jun, Dec |

## *Conclusions*

As measured by the survey results, ALCF has maintained user perception of the effectiveness of the facility. The number of addressed tickets has risen dramatically in 2011 due to continuous improvement efforts. Outreach activities focused on the commercial sector have increased ALCF industry interaction significantly. Tapping the newly formed User Advisory Council has made changes in how ALCF understands and interacts with users. Finally, ALCF brought new technology to the table to integrate and improve web content and documentation management.

Moving forward, ALCF is committed to continuous improvement. In 2012 ALCF is investing significant time in exploring new methods for assessing the effectiveness of the user experience. ALCF also wants to ask questions on the 2012 survey that expose areas of improvement guided by both the users and the core goals of the organization.

*This page intentionally left blank.*

# Section 2. Business Results

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

### *ALCF Response*

For those measures where there are concrete metrics, availability, INCITE hours delivered, and capability hours delivered, ALCF has exceeded the metrics. For the reportable areas, Mean Time to Interrupt (MTTI), Mean Time to Failure (MTTF), and utilization, ALCF is on par with the other facilities and has demonstrated acceptable performance. (Table 4)

ALCF tracks hardware and software failures closely, plus their impact on the user jobs and metrics, and uses an analysis of this data as significant input into improvement efforts.

| Intrepid (40K-node, 160K-core BG/P) | | | | | | |
|---|---|---|---|---|---|---|
| | CY 2010 | | | CY 2011 | | |
| | Target | Projected | Actual | Target | Projected | Actual |
| **Scheduled Availability** | 85% | 96% – 98% | 97.7% | 95% | 97% – 99% | 97.8% |
| **Overall Availability** | 80% | 91% – 94% | 94.6% | 90% | 92% – 94% | 94.0% |
| **System MTTI** | N/A | See Text | 5.85 days | N/A | 7d +/- 2d | 10.05 days |
| **System MTTF** | N/A | See Text | 10.82 days | N/A | 13d +/- 3d | 17.95 days |
| **INCITE Usage** | 646M | 800M | 829M | 732M | 795M – 820M | 877M |
| **Total Usage** | N/A | N/A | 1.14B | N/A | 1.1B – 1.3B | 1.20B |
| **Capability Usage** | 250M | 375M – 400M | 509M | 300M | 660M – 685M | 685M |
| **System Utilization** | N/A | 78% – 82% | 82.2% | N/A | 81% – 85% | 88.8% |

**Table 4: Summary of All Metrics Reported in the Business Results Section.**

## *2.1 Resource Availability*

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 Summary

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 5 summarizes the availability results:

| Intrepid (40K-node, 160K-core BG/P) | | | | | | |
|---|---|---|---|---|---|---|
| | CY 2010 | | | CY 2011 | | |
| | Target | Projected | Actual | Target | Projected | Actual |
| Scheduled Availability | 85% | 96% – 98% | 97.7% | 95% | 97% – 99% | 97.8% |
| Overall Availability | 80% | 91% – 94% | 94.6% | 90% | 92% – 94% | 94.0% |

**Table 5: Availability Results.**

***Summary:*** For CY2010, ALCF significantly exceeded the targets and met or exceeded the projected values. Based on this and a recommendation from the on-site review held in CY2011, ALCF increased availability targets for CY2011. For CY2011, the facility met projections and exceeded the higher targets. Overall availability was lower in CY2011. There were two reasons for this. First, ALCF lowered the availability projection in last year's OAR due to significant planned scheduled down time power work. Additionally, there was significant unscheduled down time in July 2011. 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 B.

### *Explanation of Significant Availability Losses on Intrepid*

This section provides a brief description of the causes of major losses of availability, annotated in Figure 1.

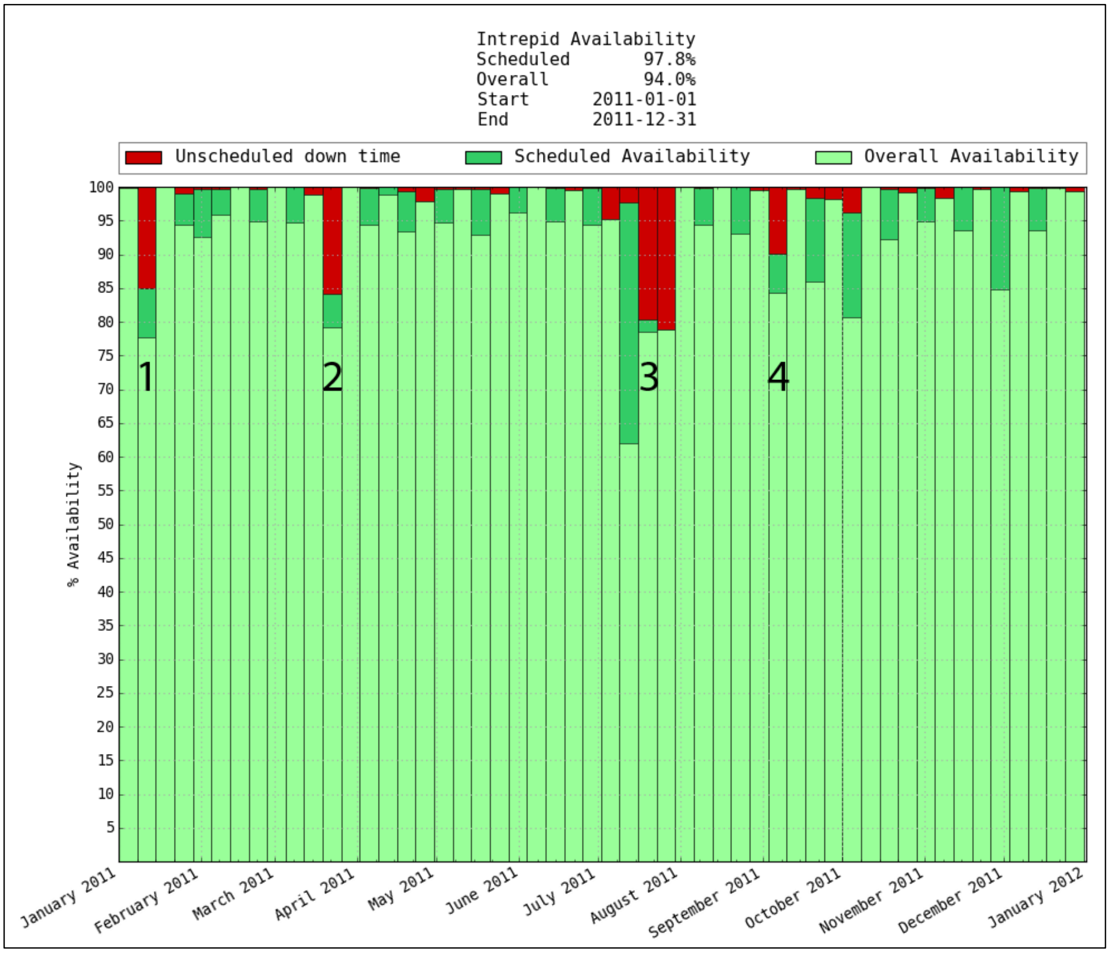


**Figure 1: Weekly Availability for CY2011.**

**Graph Description:** Each bar represents the average of 7 days availability. Each bar accounts for all the time in one of three categories. The light, pale green (lowest part of the bar) is the overall availability for that week. The darker green (middle of the bar) represents scheduled downtime for that week, with the top of that bar being the scheduled availability for that week. The red (at the top), if any, represents unscheduled downtime. As an example, the second bar from the left indicates ALCF had an overall availability of approximately 76%, a scheduled availability of approximately 85%, and approximately 15% unscheduled downtime during the 7-day period from Jan 8, 2011 through Jan 14, 2011. The numeric annotations are the significant losses that will be discussed below.

***Items 1 and 2:*** These items were discussed in last year's OAR report; per this year's guidance, they are not reiterated here. However, note that these issues were addressed and considered closed.

***Item 3:*** As indicated in the graph above, ALCF had very significant down time three out of four weeks during July 2011. This was the result of one significant scheduled outage, and four significant unscheduled outages. Four of the events have been addressed and are considered closed, while the fifth is still being worked and remains open. A detailed description of these five events is given below in chronological order.

***July 6, 2011 Power Outage (ID #2563).*** On July 6 at approximately 8:19 a.m. local time, a voltage disturbance occurred on the ComEd transmission system entering into the Laboratory. This resulted in a voltage loss to a substantial portion of the site. The Facilities Management and Services (FMS) Division responded quickly, and by noon, had safely restored power to all locations. The post-outage investigation revealed an incorrect termination of a single wire on the protective relay from a recent system modification. This particular protective relay should have screened out the disturbance on the ComEd line, but instead, it allowed the bus to be tripped. A Causal Analysis was performed by FMS with the assistance of the Compliance Oversight and Assessment (COA) Division to identify root causes for the outage incident. The causal analysis indicated two general findings:

1) Inadequate design review of proposed system modifications, and
2) Inadequate testing and commissioning for the turnover of the completed system modifications.

Based on these findings, the following recommendations will be incorporated into future system modification activities:

1) Provide an independent, third-party design reviewer for critical and complex utility systems.
2) Provide verification on critical and complex utility systems for changes, modifications, and construction. The verification should validate the installation, testing, commissioning, and training prior to placing that system in operational use (either partial or full use).

Total impact of this event was 8h 32m 2s of unscheduled down time. This issue is considered closed.

***July 9 – 22, 2011 R40 Service Card Catastrophic Power Component Failure (multiple IDs)***. On July 9 at 10:21:53 UTC, the system detected a fatal error. Upon further investigation, the service card was found in the condition shown in the photo in Figure 2. The extended nature of this outage was due to the fact that repeated intermittent link errors occurred over the following weeks. Luckily, the position of this card did not interfere with the 32K or 16K torus partitions. This allowed the machine to be run by disabling all partitions that used R40 link cards, which had minimal scheduling impact. As a result of this, multiple link cards and a midplane also had to be replaced, which essentially involved disassembling half a rack, before normal service could be resumed. This was initially believed to be an anomalous event and so was not pursued with IBM. However, in January 2012 ALCF had a virtually identical failure. ALCF has shipped both service cards to IBM for failure analysis

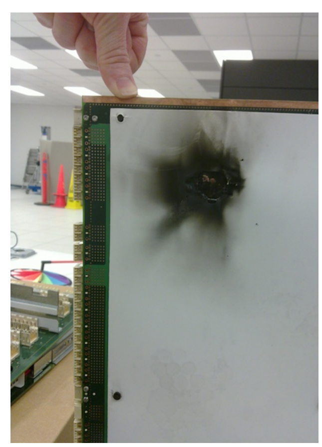

**Figure 2: Picture of the Failed Service Card.**

and is awaiting their report. ALCF continues to monitor the situation, and the issue remains open.

***July 11 – July 14, 2011 Scheduled Power Outage (ID #2561)***. As noted in the August 2011 OAR report, ALCF planned a major outage for July. The primary goal of this work was to tie the various local power measurement readouts at the Power Distribution Units (PDUs), switchgear, and transformers, as well as the Uninterruptable Power Supply (UPS) into the ALCF Building Management System (BMS) so that power consumption could be monitored and trends observed. While it would have been possible with the appropriate arc flash protection to do this work hot, much of it was taking place inside 13.3KV switch gear. It was determined that the safety risk was too great, so the work was performed with the power off. This hardware work was completed successfully, and ALCF is in the process of the software integration into the BMS.

As scheduled power outages do not occur often, the facility took the opportunity to perform other activities as well, in particular, a test of the Emergency Power Off (EPO) system in the data center. As a result of this test, it was discovered that the system was wired per the drawings, and met code, but did not operate as expected. Specifically, the designers had taken the term "Uninterruptable" too literally and had not tied the PDUs that were on the UPS into the EPO. The data center fire suppression system uses water. Although code does not require it, for safety and to reduce the damage in the event of a water release, the fire suppression system activates the EPO before releasing water. As wired, it would power off the Blue Gene racks, but nothing else (disk, servers, switches, tape, etc.). This situation was deemed sufficiently important that ALCF kept the facility down and expedited the work to put the PDUs under the control of the EPO as well. The new configuration was tested and successfully demonstrated a complete EPO of the room prior to the facility being restored to normal operations. The total impact of this was 60 h scheduled downtime and 3h 15m 59s unscheduled downtime. This issue is considered closed.

***July 19 – July 21, 2011 (ID #2757, ID #2699): Database Damage and Network Problems during Challenger Install.*** Prior to July 2011, one of the 40 Intrepid racks was designated as the "prod-devel" rack. The queuing policy was designed to allow short, small debug jobs to make it through quickly. Having such a resource is a necessity. However, on the Blue Gene architecture, it meant that the 40K partition, the 24K partition, and one of the (5) 8K partitions could not be used. To ameliorate this situation, a new rack named Challenger was purchased and designated as the prod-devel, giving back use of the partitions above. During the installation of Challenger, there were two significant issues. The first issue was that an administrative mistake was made, and rather than initializing the new Challenger database, the Intrepid database was initialized, which brought the system to a halt. This resulted in 22h 52m 3s of unscheduled system downtime. While it is impossible to be sure, it appears that it was one of the IBM installers who issued the command. This conclusion was reached because it is ALCF Standard Operating Practice (SOP) that ALCF not run as root, it is IBM's standard practice during an installation, and the command was issued as root. This is significant because under the ALCF SOP, it would have been impossible for this to happen, because ALCF privileges are set so that remote database operations (command issued on Challenger affected Intrepid) are not allowed. Because of this, during the installation of the Blue Gene/Q racks, a requirement has been established that IBM either operate with restricted access, or "shoulder surf" while ALCF executes required commands.

The second issue involved the Myricom network cabling. ALCF scheduled a three-hour maintenance window to make the connections into the production switch infrastructure. It actually took 6h 24m 28s. The difference was due to underestimating the time required, as well as issues with power supplies in the Myricom chassis when bringing the system back up. The total impact for both issues was 29h 16m 31s, of which 3 hours was scheduled. With the mitigations discussed above for database protection, this issue is considered closed.

***July 23 – July 24, 2011 (ID #2721): Power Outage Due to Torrential Rainstorm***. During July 2011, Illinois experienced several severe storms with unusually heavy rainfall. On July 23, this resulted in a power line fault that took out power to a large part of the Lab, including transformer US24, which feeds all power to the ISSF facility, other than the Blue Gene racks. FMS followed the established communications plan and gave ALCF timely notification of the event and provided regular updates during the outage. Once FMS confirmed that power had been restored, ALCF returned the facility to service. Total impact was 36h 28m 51s for the entire facility. This event was considered "force majeure," and no systemic improvements were warranted. This issue is considered closed.

***Item 4: September 4 – 5, 2011 (ID #2810 & #2811) and September 9, 2011 (ID #2812): DDN Hardware Failures Resulting in GPFS Failures.*** This item was the result of two similar incidents occurring five days apart. In both cases, a controller on one of the DDN9900s failed and caused GPFS to go down. In both cases, the Logical Unit Numbers (LUNs) were damaged and rebuilt, but there was no loss of data. These items resulted in a total of impact of 19h 22m 15s. As a result of this and a general increase in failure rates with age, ALCF is negotiating with DDN to increase the onsite stock of spare parts. ALCF will continue to monitor the failure rates and their impact on facility availability.

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

***2012 Operational Assessment Guidance:***

***Description:*** *System Mean Time To Interrupt is the time, on average, to any outage on the system, whether unscheduled or scheduled. This is also known as MTBI (Mean Time Between Interrupts). Mean Time To Failure is similar, but it is the time, on average, to an unscheduled outage on the system*.

### *ALCF MTTI and MTTF Summary*

MTTI and MTTF are reportable values, but no specific metric has been set. Table 6 summarizes the current MTTI and MTTF values.

| Intrepid (40K-node, 160K-core BG/P) | | | | | | |
|---|---|---|---|---|---|---|
| | CY 2010 | | | CY 2011 | | |
| | Target | Projected | Actual | Target | Projected | Actual |
| System MTTI | N/A | See Text | 5.85 days | N/A | 7d +/- 2d | 10.05 days |
| System MTTF | N/A | See Text | 10.82 days | N/A | 13d +/- 3d | 17.95 days |

**Table 6: MTTI and MTTF Results.**

The values reported for CY2010 listed above do not correspond to the values reported in the 2010 OAR report. This is because, for the CY2010 report, ALCF was calculating these values differently than OLCF and NERSC. A series of discussions occurred among the sites, and all three sites agreed to a common calculation for System MTTI and MTTF. The values above were recalculated using that formula. ALCF did make projections in the CY2010 report, but those projections were for the previous calculation, and there is no simple mapping between the two. You will note a significant improvement in the MTTI and MTTF values from CY2010 to CY2011. This is primarily due to resolving the Myricom Transceiver issue. Additionally, the system has matured and stabilized over time, particularly due to our tracking of job interrupts and using the top causes to drive our improvement efforts.

## *2.2 Resource Utilization*

Having addressed availability above, the next item to cover is how the core-hours that were available were used. 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).

### 2.2.1 System Allocation and Usage

Table 7 shows how Intrepid system hours were allocated and used by the allocation source. Taking the theoretical hours and multiplying them 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 July – June, so to arrive at allocated values for the calendar year, half the hours are arbitrarily assigned to each year. The allocated values for the DD allocations appear higher than expected, because they represent a rolling allocation, typically of a three-month duration with an average of 220M hours allocated at any given time. Since a majority of the DD projects are exploratory investigations, the time allocations are not used in full. The Director's Discretionary allocations are discussed in detail in the Strategic Results section. In CY2011, the ALCF successfully delivered a total of 1.20 billion core-hours across all users, a slight improvement over CY2010.

| Intrepid (40K-node, 160K-core BG/P) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | CY 2010 | | | | CY 2011 | | | |
| | Allocated | | Used | | Allocated | | Used | |
| | % | Core hours | % | Core hours | % | Core hours | % | Core hours |
| **INCITE** | 60 | 646M | 72.6 | 829.0M | 60 | 732M | 73.1 | 876.6M |
| **ALCC** | 30 | 157.9M | 16.5 | 124.7M | 30 | 210.4M | 14.1 | 168.1M |
| **DD** | 10 | 430.4M | 10.9 | 188.4M | 10 | 479.3M | 12.8 | 153.4M |
| **Total** | 100 | 1.23B | 100 | 1.14B | 100 | 1.42B | 100 | 1.20B |

**Table 7: Core-hours Allocated and Used by Program.**

### 2.2.2 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.*

#### *ALCF Utilization Summary*

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 8 summarizes ALCF utilization results and Figure 3 shows system utilization over time by program.

| Intrepid (40K-node, 160K-core BG/P) | | | | | | |
|---|---|---|---|---|---|---|
| | CY 2010 | | | CY 2011 | | |
| | Target | Projected | Actual | Target | Projected | Actual |
| **System Utilization** | N/A | 78% – 82% | 82.2% | N/A | 81% – 85% | 88.8% |

**Table 8: System Utilization Results.**

***Summary***: For CY2011, the system utilization values are in line with general expectations and are an improvement over CY2010. The improvement was greater than expected, which is why ALCF exceeded the projection. This is likely due to the introduction of the Challenger rack and an increase in capability jobs, particularly 8K-node jobs, which pack well on the machine. The graph below shows how utilization varied over the reporting period. The calculations for utilization are in Appendix B.

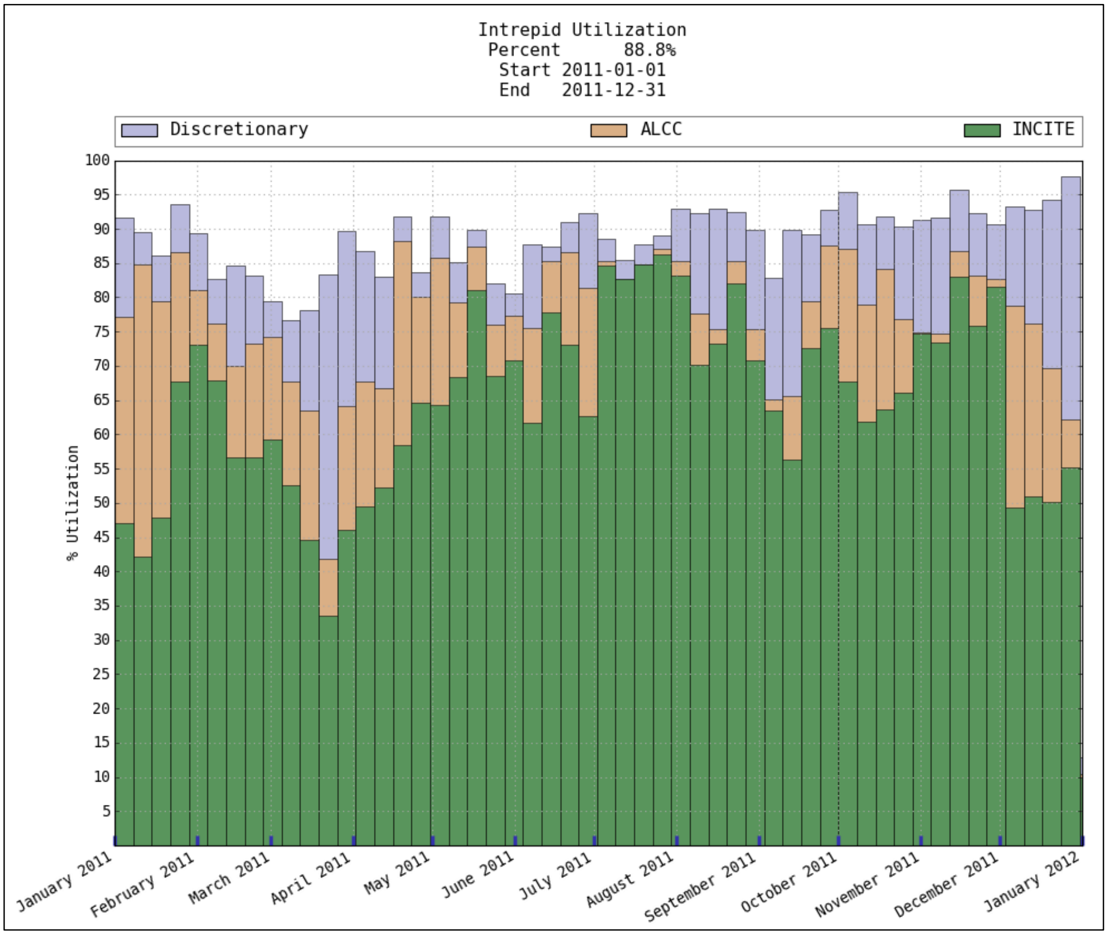


**Figure 3: System Utilization Over Time by Program.**

## *2.3 Capability Utilization*

*2012 Operational Assessment Guidance:*
*Description: The Facility shall describe the agreed definition of capability, the agreed metric, and the operational measures that are taken to support the metric*.

**ALCF Capability Utilization**
For CY2011, ALCF is required to deliver 300 million core hours to capability jobs. The current definition of a capability job is one that requests 20% of the machine or more. For Intrepid, 20% of the machine equates to 8,192 nodes or 32,768 cores. The use of the number of nodes requested, as opposed to allocated, is a change from CY2010 and is a direct result of reviews of the CY2010 OAR report. Note also that for CY2012 the ALCF capability metric will be set as a percent of total core-hours delivered rather than as an absolute number. This is more in line with OLCF and NERSC.

The following Table 9 summarizes the ALCF capability-related results:

| Intrepid (40K-node, 160K-core BG/P) | | | | | | |
|---|---|---|---|---|---|---|
| | CY 2010 | | | CY 2011 | | |
| | Target | Projected | Actual | Target | Projected | Actual |
| **Total Usage** | N/A | N/A | 1.14B | N/A | 1.1B – 1.3B | 1.20B |
| **Capability Usage** | 250M | 375M – 400M | 509M | 300M | 660M – 685M | 685M |

**Table 9: Capability Results.**

The scheduling policy on ALCF 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 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. The introduction of these queues dramatically increased ALCF capability throughput when they were implemented during the 2010 OAR period.

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. Another utilization feature is what is called "Big Run Monday." Every Monday, jobs are organized manually to maximize utilization. In general, a "best pack" is performed, considering both job size and run time, with job size generally taking precedence. Additionally, relative job order is maintained where possible, no one project is allowed to dominate the time, and packing is limited to a maximum of 48-60 hours. (Figure 4) In this way, the machine is not constantly entering a draining/backfill state when one of the large jobs is run.

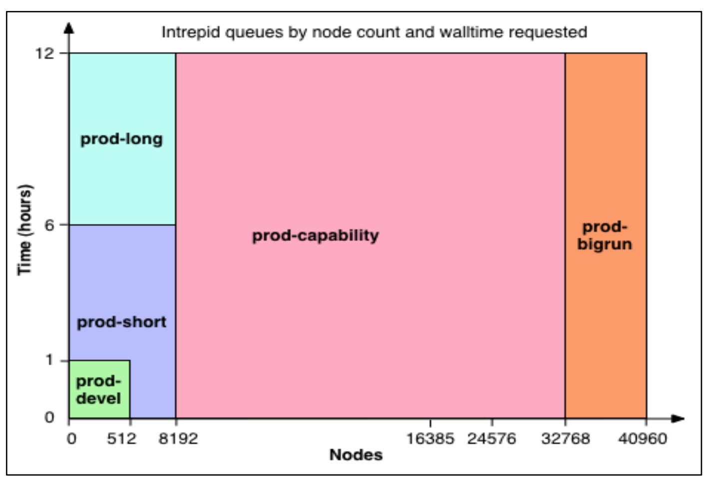


**Figure 4: Diagram Depicting Queue Assignment Based on Job Size and Run Time.**

The current weak point of the scheduling policy is turnaround time on prod-long (small, long-duration) jobs. ALCF tries to mitigate this through user education about the better queue wait times for prod-short (small, short) jobs and is looking at dynamically re-sizing the prod-long percentage of machine in response to queued jobs.

***Summary:*** For CY2010, ALCF exceeded the target for total capability hours delivered, with 75% of the cycles being delivered to INCITE. For CY2011, the facility again exceeded the capability target, with INCITE accounting for nearly 73%. The remainder of this section presents graphs showing the capability use over time, by INCITE, ALCC, and Discretionary for both CY2010 and CY2011 YTD, and then a breakdown of how the total core-hours delivered were distributed across job sizes. (Figures 5 and 6)

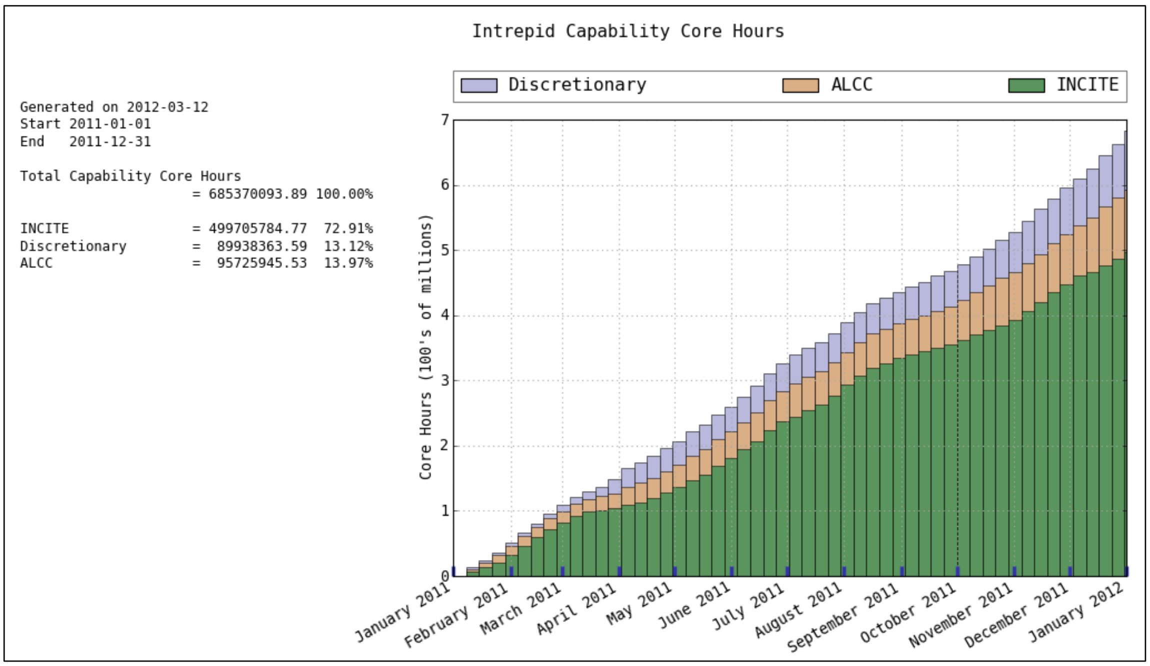


**Figure 5: Cumulative Capability Usage During CY2011.**

Figure 6 shows job distribution by size of run. Note that larger/capability jobs are at the bottom.

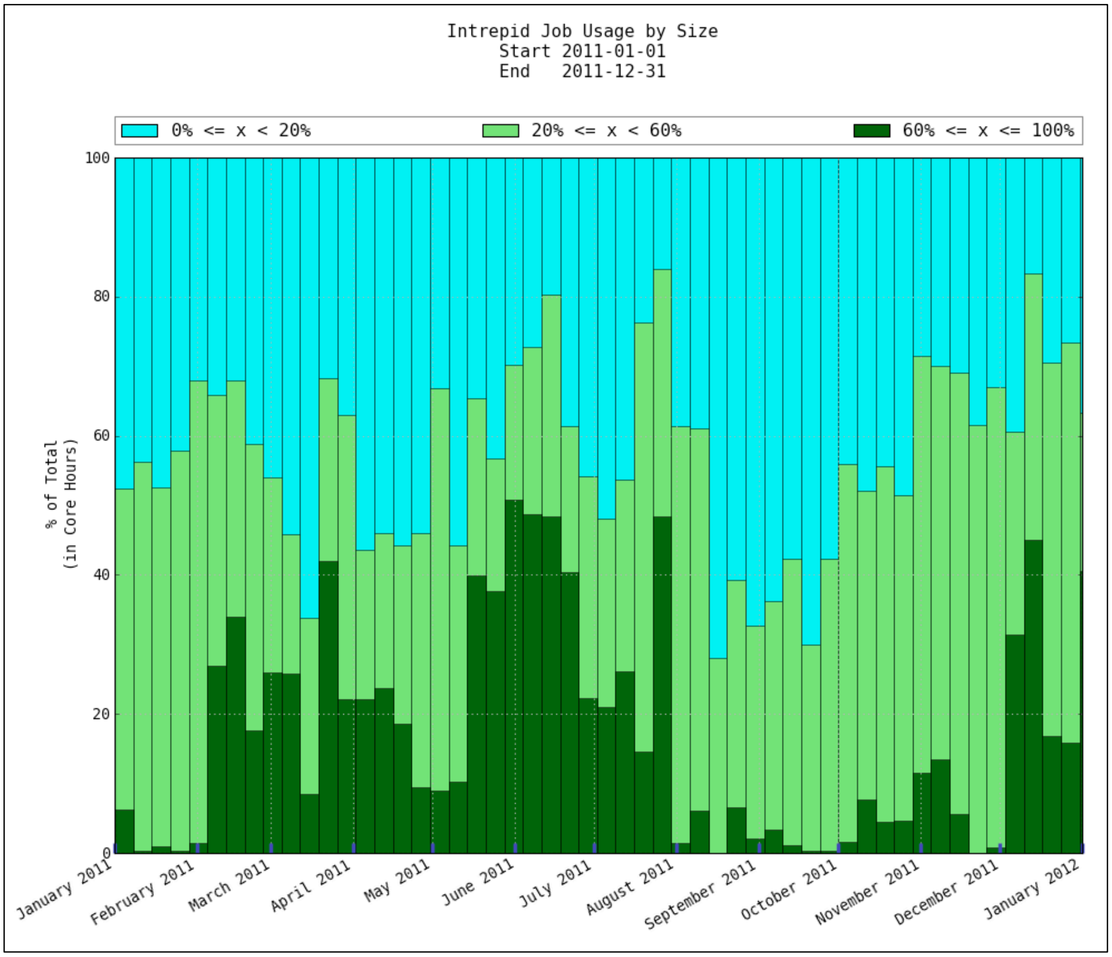


**Figure 6: Intrepid Jobs Binned by Size.**

## *Conclusions*

ALCF is maximizing the use of its HPC systems and other resources, consistent with its mission. For those measures where there are specific metrics (availability, INCITE hours delivered, and capability hours delivered), ALCF exceeded the required metric. For the reportable areas—Mean Time to Interrupt (MTTI), Mean Time to Failure (MTTF), and utilization—ALCF is on par with OLCF and NERSC, and the values reported are reasonable. These measures are summarized in the following Table 10.

| Intrepid (40K-node, 160K-core BG/P) | | | | | | |
|---|---|---|---|---|---|---|
| | CY 2010 | | | CY 2011 | | |
| | Target | Projected | Actual | Target | Projected | Actual |
| **Scheduled Availability** | 85% | 96%-98% | 97.7% | 95% | 97%-99% | 97.8% |
| **Overall Availability** | 80% | 91%-94% | 94.6% | 90% | 92%-94% | 94.0% |
| **System MTTI** | N/A | See Text | 5.85 days | N/A | 7d +/- 2d | 10.05 days |
| **System MTTF** | N/A | See Text | 10.82 days | N/A | 13d +/- 3d | 17.95 days |
| **INCITE Usage** | 646M | 800M | 829M | 732M | 795M – 820M | 877M |
| **Total Usage** | N/A | N/A | 1.14B | N/A | 1.1B – 1.3B | 1.20B |
| **Capability Usage** | 250M | 375M – 400M | 509M | 300M | 660M – 685M | 685M |
| **System Utilization** | N/A | 78% – 82% | 82.2% | N/A | 81% – 85% | 88.8% |

**Table 10: Summary of All Metrics Reported in the Business Results Section.**

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 Departments 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 project members to adapt their project simulation codes to run efficiently in a high performance computing (HPC) environment and has enabled scientific achievements that would not otherwise have been possible.

***2012 Operational Assessment Guidance:*** *In this section, the Facility reports:*
- *Science Output;*
- *Scientific Accomplishments; and*
- *Allocation of Facility Director's Reserve Computer Time (HPC only).*

## *ALCF Science Overview*

The primary science results are highlighted in section 3.1. The ALCF science staff engages the broader science community to better support our existing and future projects. By participating in the communities and understanding where they want to go with their science and application we can better plan how our Facility can support them and help them scale to the leadership class.

2011 was an exciting year as the Early Science Program (ESP) picked up speed. By the end of the year, nine post-docs were working closely with the ALCF and project teams. Time was available on the first Blue Gene/Q hardware, and ALCF is collecting the first promising snapshots of how well applications will fit on Blue Gene/Q. Initial results show many existing codes should be successful very quickly and ALCF staff is quickly building the knowledge base on how to best use the hardware.

The ESP, as well as ongoing collaborations with IBM and Lawrence Livermore National Laboratory (LLNL), has shone a light on ALCF's close relationship with the Mathematics and Computer Science (MCS) Division at Argonne. The research outcomes from this division are frequently adopted in the ALCF's software environment. As ALCF moves toward Mira with a new level of node-level parallelism, the staff will continue to work with the MPI-3 forum to improve interfaces with shared memory models, such as OpenMP. The Asynchronous Dynamic Load-Balancing Library (ADLB) has been adopted by James Vary's INCITE project. ALCF staff collaborates with the developers of Portable, Extensible Toolkit Scientific Computation (PETSc), used by multiple projects, including William Tang's, to ensure that the library is available and performing well on its resources. ALCF and MCS have collaborated to help improve the I/O of several INCITE projects, and that work continues within an INCITE project (PI: Ewing Lusk).

The ALCF has been very active in the HPC community, which is particularly valuable in building the expertise and knowledge base for new hardware over the coming years. ALCF staff members were part of the organizing and program committee for the International Workshop on OpenMP in June 2011. A team member became president of SciComp, the IBM HPC Systems Scientific Computing User Group, and ran a very successful meeting in 2011. Last spring, the ALCF participated in a long program at the Institute of Pure and Applied Mathematics, "Navigating Chemical Compound Space for Materials and Bio Design," which attracted a large, multidisciplinary group of participants. (A staff member organized and chaired the program.) Eleven new collaborations came out of this meeting, which have since translated into one ALCC award, one INCITE proposal, and multiple Director's Discretionary projects that are preparing to be INCITE candidates in future years. Additionally, staff have given talks at the following meetings: DOE-CSGF HPC Workshop, SIAM Conference on Parallel Processing for Scientific Computing, MPI Forum, Swiss National Supercomputing Centre Users' Day, the American Chemical Society National Meeting, the March Meeting of the APS, and the XXIX International Symposium on Lattice Field Theory. ALCF staff has been invited to speak at many institutions, including Berne University, ETHZ, and Technical University Berlin. ALCF staff also has presented seminars at Argonne (Materials Science Division, Laboratory for Advanced Numerical Simulations, MCS Division, onsite ALCF workshops, MCS lecture series for summer students), U.C. Berkeley (Berkeley, CA); the Fritz Haber Institute of the Max Planck Society (Berlin, Germany); Stanford University (Palo Alto, CA); and the Computation Institute (University of Chicago, IL). For SC11, the ALCF staff gave posters and papers, as well as participated in Birds-of-a-Feather sessions.

ALCF staff played a role in preparing the successful multi-institutional proposal for DOE's new Climate Science for a Sustainable Energy Future (CSSEF) program, including attending a writers' workshop at LBNL, writing a white paper on computational and numerical methods, and otherwise working with proposal authors at Argonne and elsewhere.

In addition, the ALCF played roles in various multi-institutional proposals for DOE's Exascale Co-Design Center program, including:

- For *Flash High Energy Density Physics Exascale Co-Design Center*, contributing authorship (two ALCF staff are on the proposal);
- For *Chemistry Exascale Co-Design Center (CECC)*, contributing authorship (two ALCF staff are on the proposal);
- For *A Novel I5 Approach to Exascale Computing and Co-Design for Energy-Related Science,* contributing authorship (ALCF staff member is Argonne PI);
- For *Exascale Center for Earth System Simulation (ExCESS)*, contributing authorship (two ALCF staff are on the proposal).

These strategic activities serve to attract additional, high-impact computational science projects to the ALCF, prepare users for the transition of their applications to the BlueGene/Q, and ensure that the ALCF staff maintains its deep expertise in computational science, thus providing excellent support of its user community on current and future LCF resources.

## 3.1 Science Output

***2012 Operational Assessment Guidance:*** *The Facility tracks and reports the number of refereed publications written annually based on using (at least in part) the Facility's resources. 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.*

#### Publications

Quarterly, the ALCF reports the publications derived from research done at the facility. We maintain a database of publications, and our website presents a bibliography of publications with abstracts and online access to some papers.

During CY2011, 121 new refereed papers were published based on ALCF projects, including nine in *Physical Review Letters* and two in the journal *Nature.*

## 3.2 Scientific Accomplishments

***2012 Operational Assessment Guidance:*** *The Facility highlights a modest number (at least 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. (See ALCF 2011 slides as an example.)*

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

#### ALCF Science Highlights

This section highlights some of the scientific accomplishments achieved using the ALCF resources in the past year.

With the change in the reporting period for the OAR, several 2011 science highlights were covered in the August 2011 report.. Here we summarize some of the highlights that were used in the August 2011 report and then cover some more recent science highlights in greater detail.

- William George with the National Institute of Standards and Technology (NIST) developed a new technique for studying and designing new types of concrete on Intrepid—a rheometer. The $100B concrete industry has a great need to understand the relationships between deformations and stresses and measure them with a rheometer, but no adequate rheometer previously existed for concrete.

- Professor George E. Karniadakis with Brown University is leading a team of scientists using ALCF resources to create highly detailed multiscale models of blood flow within the complex blood vessel networks in the brain. Their main accomplishments to date are: (i) simulation of initial stages of clot formation; (ii) simulation of blood flow, and healthy and diseased red blood cells (RBCs) at different stages of cerebral malaria and sickle cell anemia; (iii) modeling of the glycocalyx layer, which plays an important role in protecting the arterial wall; and (iv) modeling of the microcirculation and distribution of RBCs in Y-shaped bifurcating arteries.

- David Baker's team with the University of Washington developed a new approach for computational analysis of Nuclear Magnetic Resonance (NMR) data that pushes the limits of protein size that can be structurally solved from NMR spectroscopy data. This is a very significant step forward, since many of these larger proteins are not amenable to analysis by X-Ray crystallography, and thus, NMR remains the only way to obtain their structure.

- Donald Truhlar's group with the University of Minnesota studies metallofullerenes important for technology and biological chemistry. These structures can work as molecular electronic switches. From both a fundamental, theoretical point of view and a practical one, it is essential to find the energy minima and saddle points and to map the topography of the seams of conical intersections in these fascinating systems. The energies were calculated, which greatly increases the chances of using these structures as switches.

- Priya Vashista with the University of Southern California studied cavitation erosion, a significant mechanism for long-term component degradation in nuclear power plants. The team completed a 1-billion-atom simulation and improved understanding of the erosion from this process.

- Micheal Smith with Argonne National Laboratory improved capability for modeling nuclear reactor systems by proving a technique for neutronics that scales to Intrepid computing.

Science achievements not featured in the previous OAR report follow. Included for each highlight is a summary graphic stating the type of project (INCITE or ALCC,); the core hours allocated, followed by the actual core hours used; and a bar chart showing the breakdown of machine usage into three categories: 1) hours using less than 20% of Intrepid, 2) hours using from 20% to less than 60% of the machine, and 3) hours using 60% or greater of the machine.

## *Reducing Toxic Gas through Metal Catalysis*

*INCITE: Jeffrey Greeley, Argonne National Laboratory*

*Allocated/Used: 15M/16M*

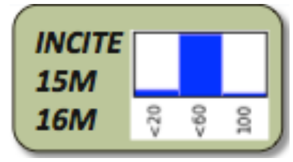


Carbon monoxide is a toxic gas that is produced in car exhaust as well as in a number of other industrial processes. A catalytic agent can aid the conversion of carbon monoxide to carbon dioxide. Nanoparticles of precious metals such as gold, platinum, and rhodium are of particular interest for the conversion of carbon monoxide to carbon dioxide. Quantum mechanical calculations of adsorption energies and charge transfer effects are indicators of catalytic properties and are very complex and time-intensive calculations. Due to the availability of computational resources made possible through the Department of Energy (DOE) INCITE program, it has been possible to study the catalytic properties of nanoparticles with over 10,000 valence electrons at the quantum mechanical level of theory.

Nanoparticles interact differently from bulk materials. Determining the size of the particles—the number of atoms—to use in catalysts for optimal results is difficult to do with physical experiments. Large-scale computer simulation enables the study of the interactions—and therefore the catalytic properties—for many configurations. A number of factors were analyzed in this study; these charge density plots are just one of many. The adsorption energies are the clearest indicator of finite-size effects. GPAW, a $O(N^3)$, real space, grid-based Density Functional Theory (DFT) code, was used for the modeling.

Adsorption energies for the platinum nanoparticles converge much more quickly (as a function of the number of atoms) to their bulk values than for gold nanoparticles (the latter was completed in CY2010). For example, gold nanoparticles adsorption energies converge to bulk values at about 561 atoms, while platinum nanoparticles converge to bulk values at 147 atoms. The adsorption energies calculated were for the adsorption of CO and O on these nanoparticles.

**IMPACT**: Enable design of improved catalytic systems (like catalytic converters) for a wide range of industrial uses, for example, electrocatalysis for fuel cells and metal catalysis of CO to $CO_2$.

**ALCF Contributions**: ALCF staff collaborated with the project to resolve difficulties converging geometry optimizations. ALCF staff reduced the memory footprint, added a new layer of parallelization to the code, and identified a very difficult ScaLAPACK memory bug. They assisted the project team with scaling from 512 cores to 131,072 cores. They also added parallel I/O, achieving 40% of I/O peak on runs of 32, 678 cores.

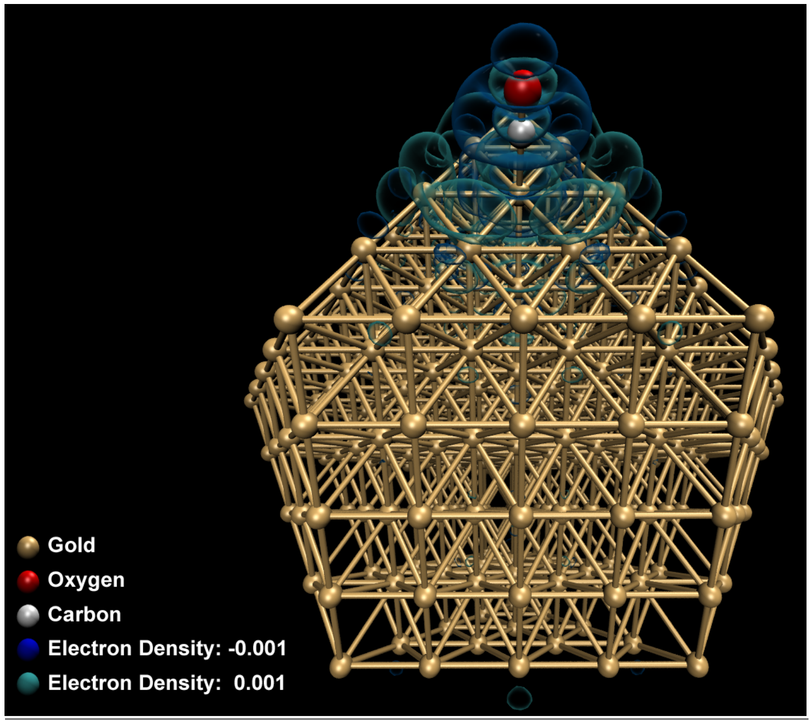


**Figure 7: The charge density difference of a carbon monoxide molecule—one oxygen atom and one carbon atom—adsorbed (adhered to the surface) on a gold nanoparticle of 309 atoms.**

## *Boosting Fuel Economy through Cutting-Edge Computational Physics*

*INCITE: Robert Moser, University of Texas at Austin*
*Allocated/Used: 40M/41M*

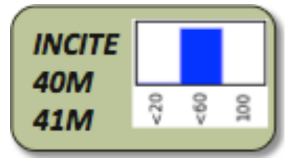


United Airlines, the world's largest airline, consumes one of every 350 gallons of global oil production. In 2010, the airline's fuel costs reached $9.6 billion. To improve jet fuel usage efficiency and decrease the impact of soaring fuel costs, scientists need a thorough understanding of the physics at work in turbulent flows, including drag.

Before it can become airborne, an aircraft must power through the effects of drag, burning costly fuel in the process. A team of researchers led by Robert Moser from the University of Texas is using the ALCF to shed light on the physics that impact fuel economy in jets and other systems affected by drag in turbulent flows. With commercial aircraft flying at near supersonic speeds, the calculations in these studies require leadership-class supercomputing resources in order to gather the required turbulence statistics.

In 2011 the team completed the highest-ever Reynolds number simulations for spatially evolving incompressible turbulent boundary layers using direct numerical simulation (DNS). The simulations show excellent agreement with experimental datasets in literature and other simulations. The team found that turbulent boundary layer fluctuations are higher than those in channel flows. This was one of the open questions in the literature: whether the difference between boundary layer and channel flows were Reynolds number-dependent or not. The lack of high Reynolds number boundary layer simulations had made it impossible to answer this particular question until now.

**IMPACT**: Improve the accuracy of modeling turbulent phenomena in boundary layers, enabling vehicle designs that reduce energy loss due to drag and increase efficiency.

**ALCF Contributions**: The ALCF ported and conducted initial performance optimizations. These drove further optimizations by Moser's team. The ALCF team is also in the process of testing the performance and scalability of the code for Mira.

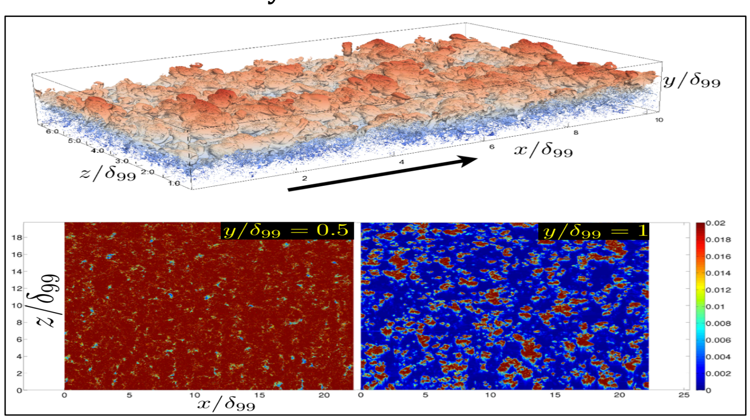


**Figure 8: This shows the turbulent/non-turbulent interface at the boundary layer; turbulent zones are red, and non-turbulent zones are blue. Rotational flow is a key difference between boundary layer flow and internal flows (in channels and pipes).**

## *Minimizing the Acoustic Signature of Jet Engines and Wind Turbines*

*INCITE: Umesh Paliath, G.E. Global Research, Niskayuna, NY*
*Allocated/Used: 20M/18M*

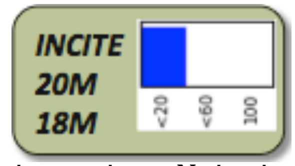


Umesh Paliath of G.E. Global Research is carrying out large-eddy simulations (LES) on jet nozzles to understand the noise generation in jet engines. Noise is an indicator of engine inefficiency, and understanding the source of noise in these engines facilitates modifications that increase fuel efficiency in existing engines. More importantly, it guides design of improved, lighter engines that have more significant fuel savings and reduce $CO_2$ emissions. Additionally, decreased noise will reduce ear damage, a significant health issue for those who work near jet engines.

In 2011 simulations were completed on a proof-of-concept problem to assess the capability of the LES approach in understanding the jet–flap interaction effect. Paliath and his team demonstrated that a LES-based approach can be extended from predicting jet noise for an isolated configuration to predicting the acoustics for an installed configuration. The research team also investigated sound generation from wind turbine blades using LES to understand and predict the broadband noise that arises from the interaction of turbulence with the airfoil trailing edge.

**IMPACT**: Reduce the noise generation of jet engines and wind turbines, thereby improving their efficiency and reducing health and environmental impact of the noise.

**ALCF Contributions**: The ALCF ported and performed performance optimizations of the code CharLES. The team is tightly engaged getting this and a preferred software package ready for Mira.

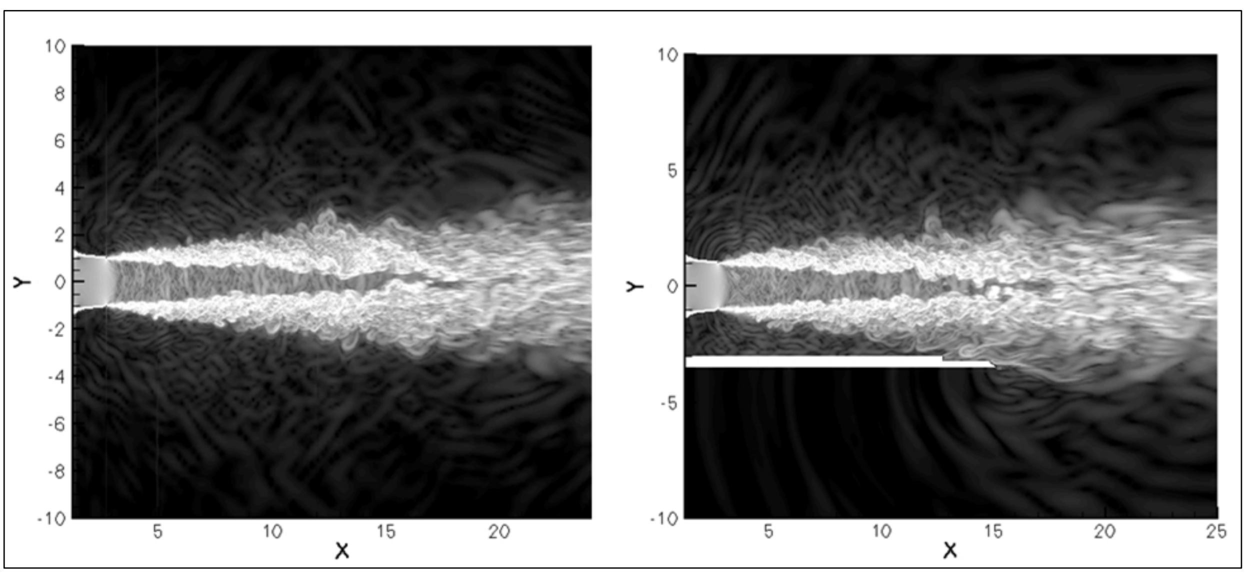


**Figure 9: Density gradient contours for simulation of conic nozzle with and without the presence of flat plane. The density gradient is a qualitative picture of noise.**

## *Simulations of Deflagration-to-Detonation Transition in Reactive Gases*

***INCITE: Alexei Khokhlov, University of Chicago***
***Allocated/Used: 18M/19M***

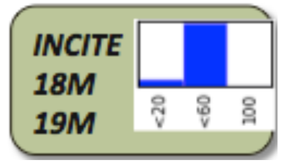


Hydrogen is the most abundant element in the universe. It is an environmentally friendly, clean fuel that has the potential to reduce the nation's dependence on foreign oil, improve the environment, and boost our economy. The challenge with hydrogen fuel is bringing it safely into our everyday lives. This fuel is very energetic, and under certain conditions, a hydrogen-oxygen mixture can react violently and detonate.

The process of transition from slow burning to a detonation is called a deflagration-to-detonation transition or DDT. Predicting DDT in various combustion settings remains an outstanding combustion theory problem. Led by Alexei Khokhlov with The University of Chicago, the High Speed Combustion and Detonation (HSCD) project uses ALCF resources to perform first-principles, reactive flow Navier-Stokes fluid dynamic simulations of DDT. The team completed simulations of a 7 micron resolution shock tube in $CO_2$ with heat conduction and isothermal walls and found excellent agreement with experimental results. This step with $CO_2$ is key to using the approach for hydrogen. The next suite of simulations will address more complex physics.

**IMPACT:** Engineering insight required to make hydrogen fuel a viable energy alternative.

**ALCF Contributions:** ALCF staff improved the visualization capabilities of the code by adding 3-D VisIt output to the existing 2-D slice output. They also rewrote the I/O routines to speed up checkpointing by 10x. ALCF staff added OpenMP threading to some of the models, resulting in a 3x application speedup. Finally, ALCF and the project are rewriting the rebalance algorithm to reduce inefficiencies that appear at large scale.

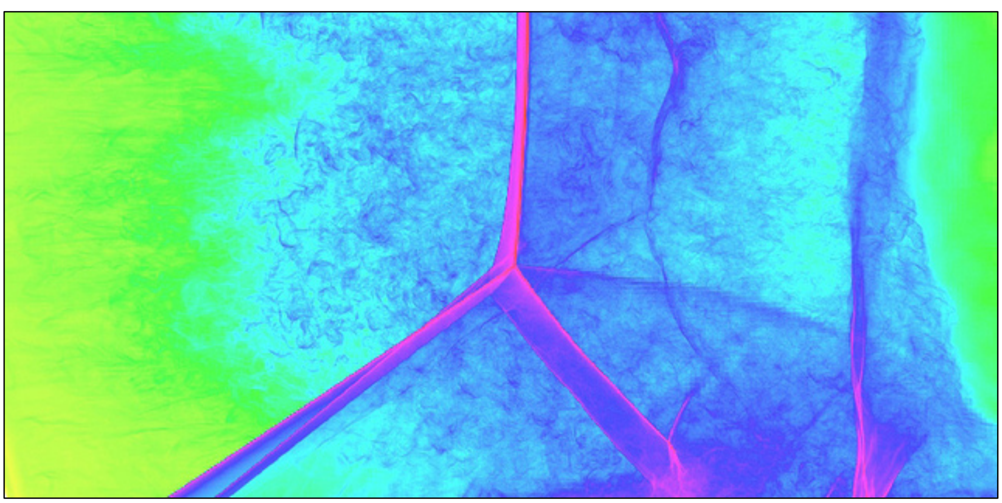

**Figure 10: Temperature in a 3-D Navier Stokes DNS simulation of Mach=3 reflected shock bifurcation in a H-O mixture in a square channel.**

## *Laser-Plasma Interactions for the National Ignition Facility*

***INCITE: Denise Hinkel, Lawrence Livermore National Laboratory***
***Allocated/Used: 50M/46M***

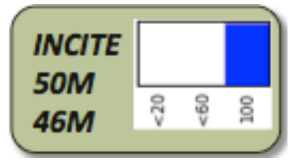


Lawrence Livermore National Laboratory (LLNL) has been tasked with achieving ignition at the National Ignition Facility (NIF), where scientists could study the phenomena of fusion applicable to alternative energy, materials, astrophysics and nuclear science. NIF focuses 192 intense laser beams at the target to create the extreme environment needed for these studies. An important aspect of focusing 192 lasers is minimizing backscatter; otherwise, the facility loses input energy needed to drive the process. Backscatter can also alter implosion symmetry as well as preheat the ignition capsule via generation of hot electrons.

Recent experimental results from NIF show that backscatter occurs in the laser target where quads of laser beams are overlapped. The goal of simulations on Intrepid is to quantify how overlapping beam quads impact backscatter. The team has found that overlapping quads in the region of backscatter can share a reflected light wave, which acts to increase reflectivity. Additionally, the team determined that spatial non-uniformity across the quads, as happens with overlapping beams, also increases reflectivity. Both effects are important to include when simulating beam propagation and backscatter in NIF. Besides already yielding important design information, these results show that expanding the simulations to more beams is critical to improving target design. These calculations were performed using the pF3D code to simulate both radiation and hydrodynamics on 80% of Intrepid for three weeks.

**IMPACT**: Improved design of physical target for future NIF experiments. These experiments work to bring sustainable, green energy production to the world.

**ALCF Contributions**: The ALCF staff helped the team debug a few critical issues and coordinate large campaigns.

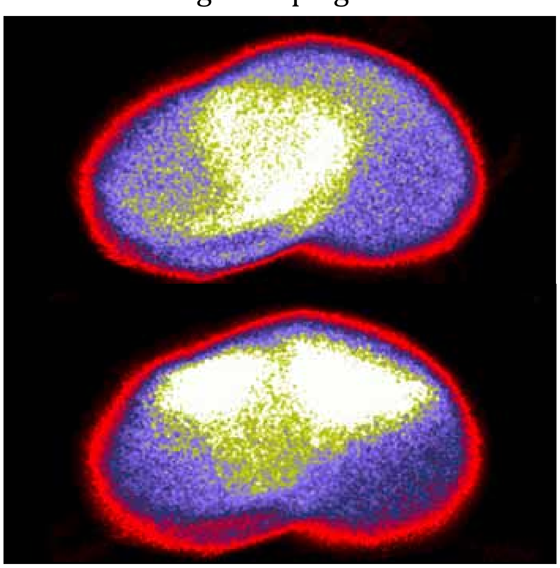

**Figure 11: Input to a pF3D simulation of two NIF laser quads propagating through an ignition target. Here, power transferred from other quads of laser beams is distributed uniformly across the laser beams. Two quads overlap in the simulated region. This enhances reflectivity through a shared reflected light wave.**

**Figure 12: Laser input to a pF3D simulation of two NIF laser quads propagating through an ignition target. Here, power transferred from other quads provides a spatially non-uniform distribution of power across the beams. The bright "triangle" in the upper region of each laser quad drives high levels of reflectivity within each quad. The overlap of the two quads drives reflectivity through a shared reflected light wave.**

## *Designing Materials from First Principles Calculations*

*ALCC: Larry Curtiss, Argonne National Laboratory*
*Allocated/Used: 20M/21M*

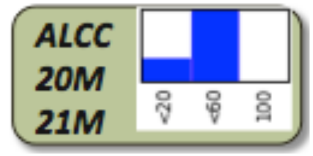


The design and discovery of new materials are crucial to our energy future. Electronic structure calculations will play a crucial role in the design of breakthrough materials in catalysis and electrical energy storage. These first principles calculations inform the fundamental understanding of the materials needed to design these new materials.

Larry Curtiss and his team are exploring ways to generate new biofuels through use of new and efficient catalysts. This work is in support of the Institute for Atomic-efficient Chemical Transformation (IACT), an Energy Frontiers Research Center (EFRC). Its primary focus is to study the transformation of biofuels with efficient catalysts. The researchers completed a complicated transformation that optimized the structure of $ZrO_2$ nanobowls in dry and hydroxylated alumina. These nanobowls are used to study fructose and conversion.

Another problem Curtiss' team is studying reducing the significant chlorinated (toxic) waste generated by production of the important industrial chemical propylene. The team completed a study of silver, $Ag_{377}$, which has been key to understanding new experimental results. Subnanometer Ag clusters may provide pathways to overcome this waste problem.

**IMPACT**: Provide the fundamental understanding and predictions needed to design new materials for catalysis and energy storage.

**ALCF Contributions**: ALCF staff spends significant amounts of time training new group members in how to configure and use the applications GPAW and NWChem. The ALCF spent significant effort on the performance and scalability of these codes over the past two years.

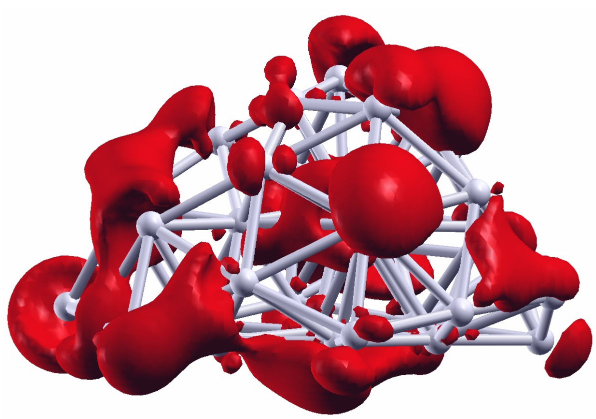

**Figure 13: Thirty-three atom silver cluster that is being studied as a new catalyst for propylene epoxidation .**

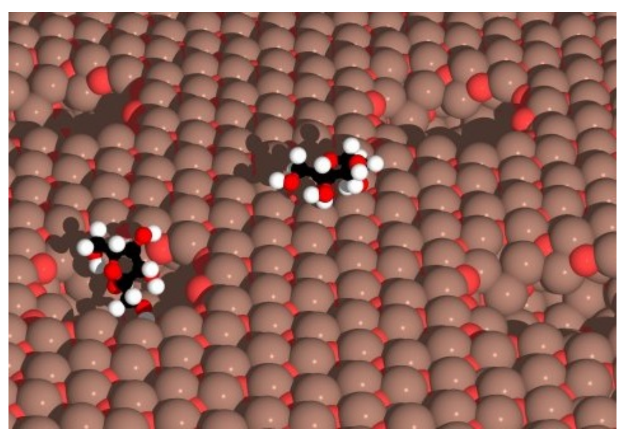

**Figure 14: Nanobowl in ZrO2 surface (1nm diam.) that is being studied for biomass conversion to fuels.**

## *Deepening the Understanding Between Quarks and Gluons Guiding Experiments*

*Paul Mackenzie, Fermilab*
*Allocated/Used: 50M/182M*

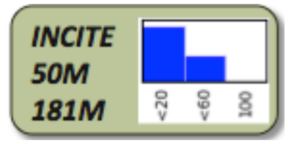


Quantum Chromodynamics (QCD) research plays a key role in the ongoing efforts to develop a unified theory of the fundamental forces of nature. While the behavior of atomic particles such as protons and neutrons is well understood, less is known about the interactions of subatomic particles like quarks and gluons, which compose the atomic particles. Paul Mackenzie from Fermilab leads the US-QCD INCITE project that studies these interactions. Not only does this work directly impact the understanding of physics, but also the research done on ALCF resources provides crucial high-precision lattice QCD calculations needed for new or in-progress experiments and for analyzing results from completed experiments.

Brookhaven National Laboratory's Relativistic Heavy Ion Collider used results from US-QCD computation to firmly constrain heavy-ion collision models for the first time. The papers that detail these results are the first- and second-most cited papers in LQCD. Fermilab has used calculations by US-QCD members combined with experimental results, allowing many of the fundamental parameters of the Standard Model to be determined more accurately than ever before.

**IMPACT**: Deliver essential theoretical results to experimental programs such as RHIC, and Fermilab.

**ALCF Contributions**: ALCF staff discovered a new, much more efficient way to calculate the HISQ "fermion force." This effort reduced the total FLOP requirement by 10 times and improved parallel efficiency, which saves about 10%-50% of total runtime. Similar improvements were made for the gauge force routines. About 10% of the runtime is now 2-3 times faster in the MILC code. These changes are captured in the LQCD SciDAC libraries.

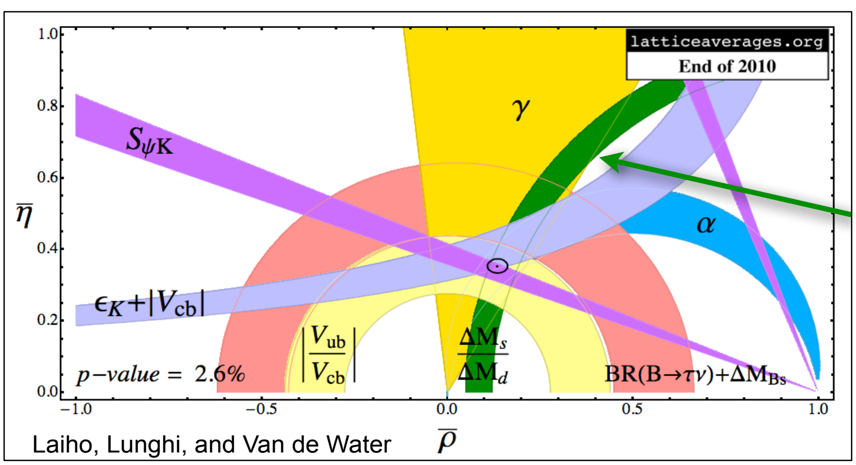


**Figure 15: The current summary of Lattice QCD averages from Laiho, Lunghi, & Van de Water, Phys.Rev.D81:034503,2010.**

## *Understand the Impact of Sickle Cell Anemia on Blood Flow*

*INCITE: George E. Karniadakis, Brown University*

*Allocated/Used: 50M/44M*

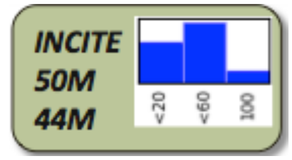


Sickle cell (SS) anemia was the first disorder to be identified as a molecular disease back in 1949. This chronic inflammatory disease is the most common genetic disease among African Americans, with an 8% incidence of the trait among this population. In the U.S., 72,000 individuals suffer from sickle cell anemia. George E. Karniadakis from Brown University is leading a team of scientists using ALCF resources to study sickle red blood cells (SS-RBCs) and their impact on blood rheology properties, specifically, changes of blood viscosity in the presence of various sickle blood cell shapes (sickle, granular, and elongated) and blood flow resistance with/without adhesive dynamics between the RBCs and the vessel walls (Figure 16). These simulations are also used to map exactly how red blood cells move through the brain, allowing scientists and cardiologists to conduct virtual experiments to study cerebral blood flow. The results of these simulations, performed normally on up to 32,768 cores of Intrepid, are published in *Biophysical Journal*.

In addition to work on sickle cell anemia, other major accomplishments of the Karniadakis team include: (i) simulation of initial stages of clot formation; (ii) simulation of blood flow, and healthy and diseased RBCs at different stages of cerebral malaria; (iii) modeling of the glycocalyx layer, which plays an important role in protecting the arterial wall; (iv) modeling of the microcirculation and distribution of RBCs in Y-shaped bifurcating arteries; (v) fluid-structure interaction simulations of arterial network; and (vi) simulations of blood thrombus formation.

**IMPACT**: Improved diagnosis and treatment for patients with blood flow diseases.

**ALCF Contribution**: ALCF staff provided assistance to Karniadakis' team with job scheduling, reservations, and general guidance of the systems to use them more efficiently. ALCF staff worked another aspect of the project on a Gordon Bell submission. This specific subproject benefited from porting and basic performance improvements.

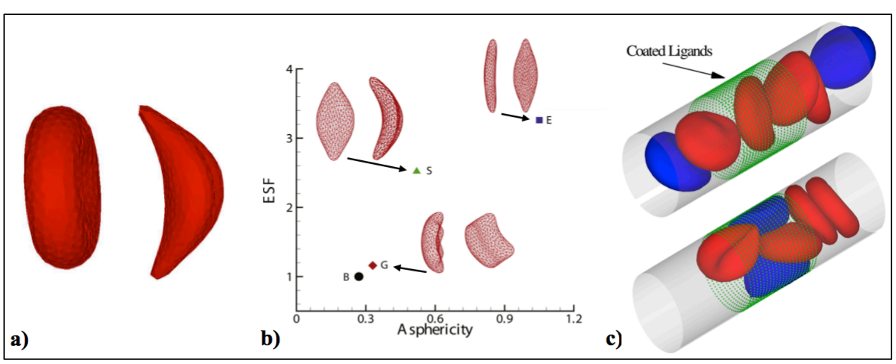


**Figure 16: a) and b) show the differences in a sickle-shaped blood cell, including how it deforms through flow; c) Sickle blood flow with adhesive dynamics. The**

**blue cells represent the active group of sickle cells that exhibit adhesive interaction with the coated ligands. The red cells represent the nonactive group of cells. Top: A snapshot showing an active group of cells flowing into the region coated with ligands. Bottom: A snapshot of the sickle red blood cells in the local occlusion state.**

## INCITE and ALCC Project Usage on Intrepid

### *CY2011 INCITE Usage*

During CY2011, 877 million core-hours were delivered to INCITE projects, with 732 million core-hours allocated to the projects. Of the 30 projects, 15 used all of their allocations, and 23 used over 90% of their allocations. A total of 44 million core-hours were unused by projects (6% of the total INCITE allocations). In one ALCF INCITE project, the PI left his position, so there was almost no usage. Two other projects accomplished their work in much less time than expected (in one case because ALCF staff sped up their application). Projects are allowed to use more than their allocation in backfill mode. When in backfill, a job can run only if no job belonging to a non-negative project can run. Projects that use significantly more than their allocations often have very fluid workflows that can run at many scales and wall times. Figure 17 shows the core-hours allocated to, and used by, each INCITE project during 2011.

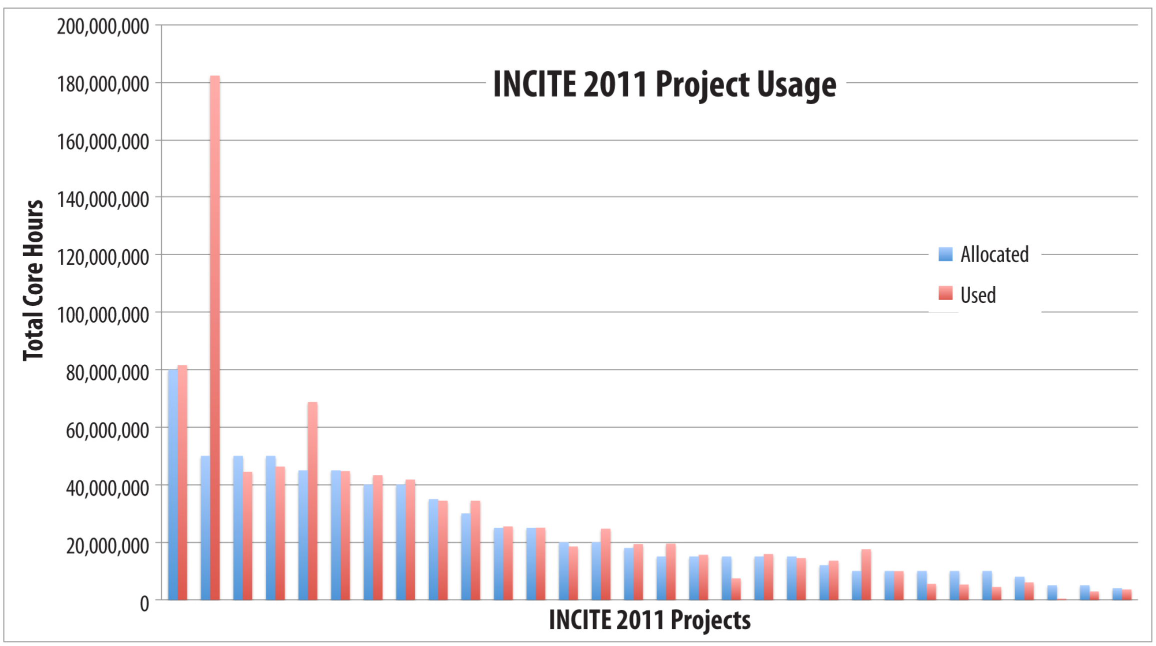


**Figure 17: For the 30 CY2011 INCITE projects, core-hours allocated (blue) and used (red).**

### *ALCC Usage*

The 2010/2011 ALCC projects started on July 1, 2010, with allocations lasting one year. During this time period, 223 million core-hours were delivered to the projects, out of the 316 million core-hours allocated. There were 51 million core-hours allocated that remained unused, and 40.1 million core-hours were pulled back from projects that failed to use at least 50% of the project by February 2011. Three projects were able to use all of their allocations, with 80 million core-hours allocated and 86 million core-hours delivered to these projects. Figure 18 shows the core-hours allocated to and used by each ALCC project.

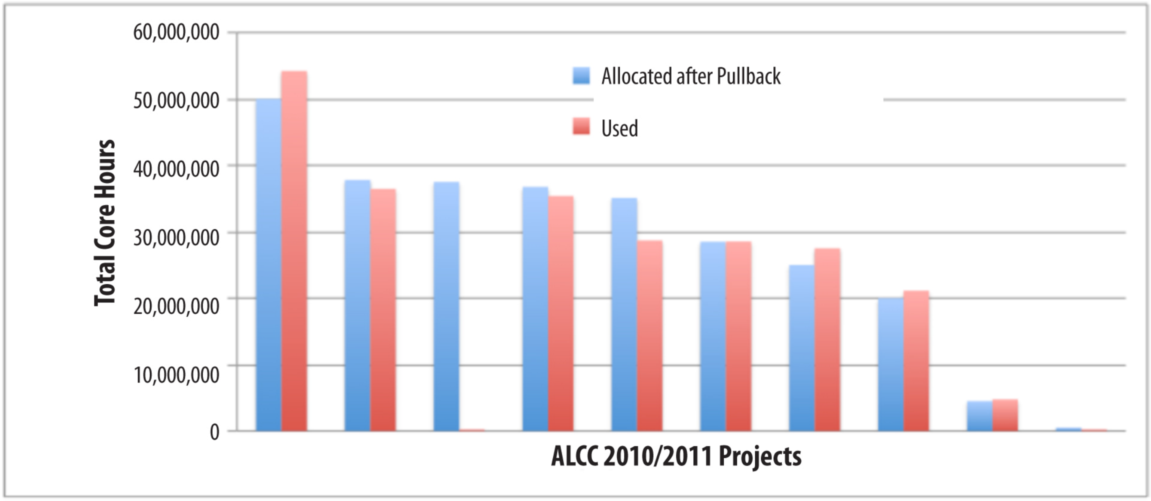


**Figure 18: For the ten 2010/2011 ALCC projects, core-hours allocated (blue) and used (red).**

The 2011/2012 ALCC projects started on July 1, 2011, with allocations lasting one year. Halfway through the year, projects have used 67 million core-hours of 185 million core-hours allocated; 36.5% of the time has been used (Figure 19).

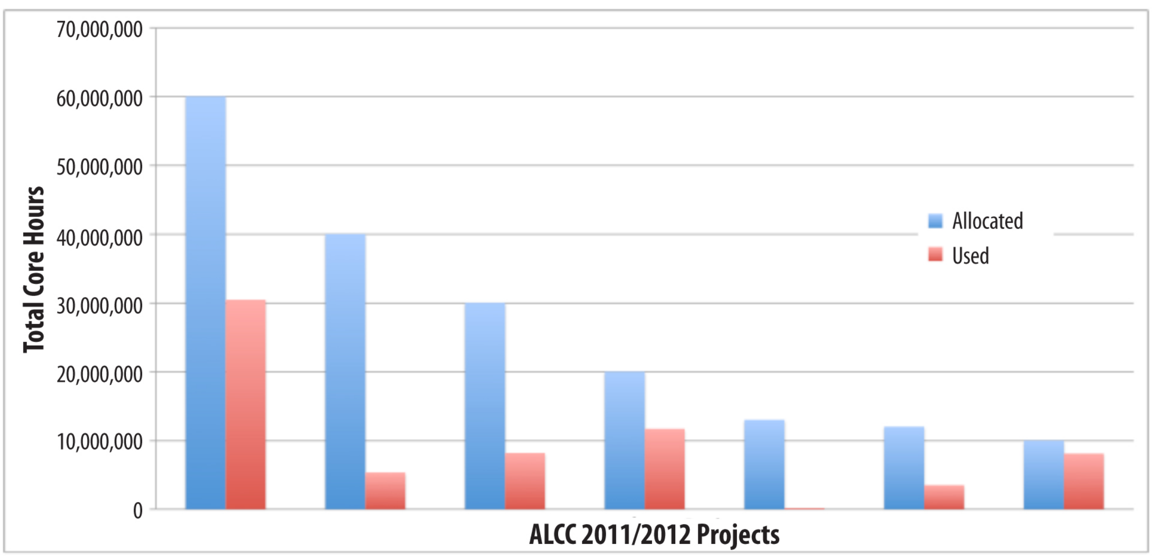


**Figure 19: For the seven 2011/2012 ALCC projects, core-hours allocated (blue) and used (red).**

## *3.3 Allocation of Facility Director's Reserve Computer Time*

***2012 Operational Assessment Guidance:*** *The Facility describes how the Director's Reserve is allocated and lists the awarded projects, showing the PI's name, organization, hours awarded, and project title.*

**ALCF Director's Reserve Program**

The Director's Reserve, or Director's Discretionary (DD) program, serves the HPC community interested in testing science and applications on leadership-class resources. Projects are allocated in five categories: 1) INCITE or ALCC preparation, 2) code support and/or development, 3) strategic science, 4) internal/support, and 5) early science.

"INCITE and ALCC preparation" allocations are offered for projects that are preparing proposals. These projects can involve short-term preparation (e.g., run scaling tests for their computational readiness) or longer-term development and testing. "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, NWChem and its required libraries now perform very well on Intrepid due to work in the DD pool. This effort has fueled multiple, successful INCITE proposals. Computer science is a large fraction of this category covering areas including software tools, benchmarks, performance and applied mathematics. 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.

In CY2011, we established another category for projects—the Early Science Program (ESP). These projects were selected to collaborate with the ALCF on preparing their codes to use ALCF's next-generation Blue Gene/Q, Mira. To support these projects, ALCF has provided modest DD allocations for basic testing.

Allocations are requested through the ALCF website and reviewed by the ALCF's Director, Director of Science, and Deputy Director of Science. Input is collected from the ALCF staff where appropriate. Allocations are reviewed on their readiness to use the resources and are awarded time on a quarterly basis. The DD pool is under great demand and often the requested amount cannot be accommodated.

In 2011, we had a total of 112 DD projects. Of these, 11 are labeled "Internal/Support," as these projects are used by the ALCF to support the facility. For example, the Operations team uses this time to run their hardware diagnostics.

The Director's Reserve allocated 480M core-hours and used 154M core-hours from January 1, 2011 to December 31, 2011. The nature of the DD pool supports over-allocation, but it should be noted that 480M core-hours does not represent open allocations for the entire calendar year. A project might have a 1M core-hour allocation that only persists for three months, but that 1M core-hour allocation is counted entirely in the 480M core-hour number. Project PIs are experimenting with their ability to scale their applications or the scale of their science. The 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. Lists of the CY2010 and CY2011 DD projects, including title, PI, institution, and hours allocated, are provided in Appendix B.

To provide an idea of the distribution of the allocations, Figure 20 provides a breakdown of the CY2011 usage by the standard INCITE science domains.

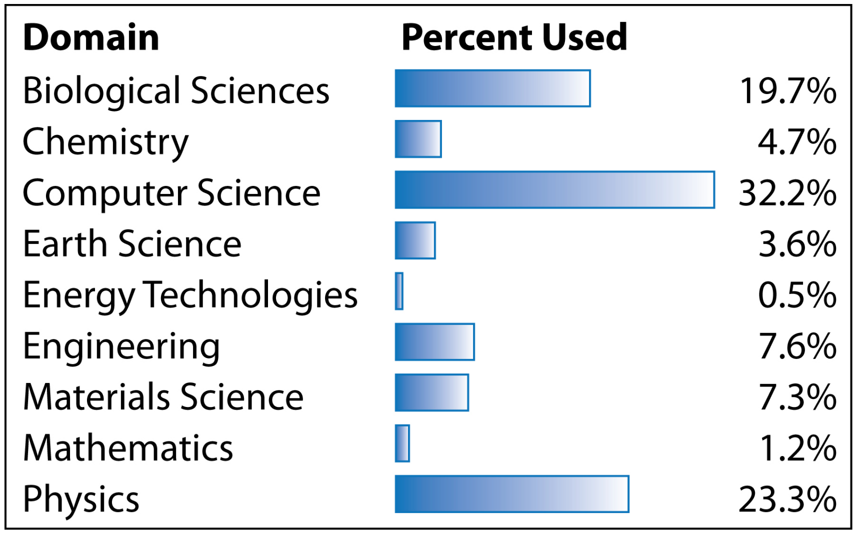


**Figure 20: CY2011 Discretionary time used by domain (of 154M used).**

## *Conclusions*

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 103 papers in the past 15 months. 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.

The ALCF delivered 877 million core-hours to INCITE projects in CY2011, far exceeding the 732 million core-hours committed. The Director's 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 the Facility's operations?***

## *ALCF Response*

Listed below are the innovations and best practices carried out at the ALCF during July 1 to December 31, 2011. The accomplishments for the first six months of CY2011 are listed in the August 2011 OAR report and are not repeated here.

## *4.1 Application Support Improvements*

**Improving Debugging at Large Scale: Collaboration with Allinea**
One of the ongoing challenges for large-scale computing systems such as Intrepid is debugging. Initially, debuggers were not usable at more than 4,096 cores on the Blue Gene/P (BG/P), causing users to revert to using printf to find and debug problems. In an effort to improve the situation for the users, the ALCF staff and Allinea formed a team to scale the Allinea debugger (Distributed Debugging Tool, DDT) to BG/P beyond 4K cores on Intrepid.

In the first phase of the project, reported in OAR 2011, investigation revealed bottleneck issues within the I/O nodes of the BG/P architecture. DDT creates a GNU DeBugger (GDB) server process on an I/O node of Intrepid associated with each Message Passing Interface (MPI) rank. Because 64 compute nodes on Intrepid share a single I/O node, there are significant memory issues on the I/O node and poor debugger performance. The ALCF/Allinea team greatly reduced the memory issues by modifying the debugger to allow all server processes on the same I/O node to share a single set of memory-mapped program data. Implementation of this solution resulted in being able to debug a job running across 32,678 cores (20% of Intrepid).

In the second phase of the project, the team will further address the I/O node bottleneck. While memory use has been brought under control, the overall number of processes contending for I/O node resources can be very high. As mentioned previously, there is one GDB server process per MPI rank. There is also one DDT daemon process per MPI rank. Therefore, running a BG/P program in Virtual Node (VN) mode (4 MPI ranks/node) loads each I/O node with over 512 processes. The work in process is twofold. First, the number of required debugger processes per I/O node will be reduced from 256 to arbitrarily few. This will be done by multiplexing control of multiple MPI ranks by each debugger process. Second, the number of required daemons per I/O node will be reduced from 256 to arbitrarily few, in a similar way. In practice, the most efficient number of servers and daemons may be greater than 1. It is anticipated the improvements made by this project will not only be useful on BG/P but also apply to directly addressing similar scaling issues on BG/Q.

**Partnerships with Vendors and External Organizations**

*ALCF Leadership Roles in User Groups and Standards Organizations*. ALCF serves in leadership roles within IBM scientific user groups and standard organizations that define standards for next-generation programming languages, programming models, and benchmarking. ALCF staff member Ray Loy is President of ScicomP. ALCF staff member Kalyan Kumaran is the long-time Chair of the SPEC High Performance Group. ALCF staff members also actively participate in the working groups of OpenMP and MPI standard bodies to drive standards in these areas.

*Automatic Performance Data Collection on Applications*. ALCF staff collaborated with ParaTools to develop an application performance data collection library for the Blue Gene/P that can be transparently enabled systemwide to collect application performance data from the on-chip hardware performance counters and MPI call information from the PMPI interface. This library enables data collection for applications with no user involvement and is currently being tested by ALCF staff.

*#1 Result in Graph 500 benchmark.* ALCF staff collaborated with Prof. Andrew Lumsdaine and Jeremiah Willcock of Indiana University, and Fabrizio Petrini of IBM to submit improved numbers on the Graph 500 benchmark for Intrepid on 131,072 cores. In the latest list released in November 2011, Intrepid is ranked #5. IBM also submitted a result on BlueGene/Q, ALCF's next-generation system, on 65,536 cores. This result is ranked #1.

### *Improvements for the Future*

*Code Optimization on Next-Generation Systems.* ALCF is working with IBM to develop a knowledge base of porting and optimization best practices for BG/Q in the form of presentations and Red Books. ALCF staff presented some of this material at the BG/Q Birds of a Feather at SC11, jointly hosted with LLNL and IBM.

*Porting Performance Tools to Blue Gene/Q.* ALCF, in collaboration with a number of tools researchers and IBM, has ported, or is in the process of porting, tools to BG/Q. The tools and teams include PAPI (UTK), HPCToolkit (Rice University), Open|Speedshop (Krell Institute), Scalasca (Juelich), and TAU (ParaTools). These tools will be made available on the BG/Q Test & Development (T&D) systems for ALCF Early Science users.

## *4.2 User Services Improvements*

**Improvements in Communication**

ALCF integrated content and deployed a new content management system (CMS) using the Drupal framework. This has enabled ALCF to unify the user documentation and web communications.

Based upon feedback from the User Advisory Council, ALCF changed the nature of the initial user conference calls. The call is now science-centric, with the participants discussing the story of their research and their upcoming challenges and goals. From this main storyline, the dicussion then dives into relevant user issues that will block the progress of their science.

## 4.3 Operations Improvements

**Improvements in Storage**

ALCF has a major storage improvement project under way to better serve the users' space and performance requirements. The project has four major components to it, which were listed in the August 2011 OAR report. One of the components was completed during the August 2011 OAR review period. Since then, two more have been completed and are listed below:

- *Space Expansion and Resiliency Improvements for the Home File System.* Intrepid was initially brought into production with only 8 racks of compute nodes and approximately 0.5PB of parallel "scratch" storage. It was subsequently upgraded to 40 racks with an upgrade to approximately 5PB of usable storage. The original storage was not integrated into the new storage system, because it was an older, slower generation of hardware, so it was sitting idle. A disaster recovery evaluation of the facility produced an estimate that if a catastrophic loss of the home file system occurred, it would take approximately six weeks to completely restore the data. To mitigate this, ALCF migrated the home file system to the original "scratch" storage system, which used hardware identical to the home file system. By doing so, it was possible to both increase the capacity of the home file system and reduce the probability of a catastrophic loss of the home file system. By enabling data redundancy, an entire Data Direct Network storage array could be lost and the home file system would not be lost. In addition, the number of home file servers increased from 8 to 16, doubling the maximum server throughput.

- *Metadata Performance Improvements for the New Home File System.* During large compiles, tar operations, etc., the home file system can exhibit slow or sluggish behavior. An analysis determined that the bottleneck was the file system metadata performance. To overcome this, during the migration of the home file system to the older hardware as described above, six of the eight file servers that had been acting as both metadata and file servers were converted to dedicated metadata servers with Fusion-IO solid state storage cards for metadata storage. ALCF performed simple "before and after" tests of operations and recorded total time for the operation. Tests involving pure file I/O, such as tar, showed up to a 68% improvement, while more mixed tests, such as compilation, showed about a 10% improvement.

## 4.4 Infrastructure Improvements

The Argonne National Laboratory TCS Data Center has a 48" raised floor assembly utilizing a pedestal and stringer substructure. The IBM BG/Q racks are some of the densest computer racks in the world, and weigh significantly more than any other racks in the data center. Third-party structural analysis and manufacturer's load testing concluded that the data center floor pedestals and floor tile required upgrades to meet the requirements of the increased static load, as well as the rolling load, given the required tile cutouts, for the racks.

The ALCF-2 project upgraded approximately 60% of the structural pedestals in the 25,000-sq-ft data center. In addition, of that 60% of floor area, the ALCF-2 project will upgrade approximately 650 2'x2' floor tiles. As of the date of this report, the data center has three different load capacity conditions: 1) no upgrades, 2) upgraded pedestals only, and 3) upgraded pedestals and upgraded floor tiles. For safety reasons, it is important to delineate the various structural capacities of the floor assemblies in such a way that the differences are easily determined with a quick glance. To accomplish this delineation, colored tile will be utilized to provide a visual indicator and reference of the load capacity of any given floor location. Color legends to be posted around the perimeter of the room will provide further details of the load capacity of each area.


**Figure 21: The two BG/Q racks rest on the new floor tiles. The new floor tiles have a darker shade compared to the initial tiles.**

A clear line of distinctively colored tiles will delineate the area where "old" pedestals are installed vs. the area where the upgraded pedestals are installed. In addition, a separate and distinctive color tile will demarcate the upgraded floor tiles. This is done to help ensure recognition and visual reference for all users of the multi-divisional shared space as equipment is delivered, moved, and installed around the data center. Currently, only a portion of the floor tiles have been upgraded, just enough to safely install the two T&D racks and test the visibility of the colored tiles (Figure 21). The color of these initial floor tiles was too light and not visible enough when viewed from a distance. The remainder of the upgraded tiles was ordered in the darkest color available and will be installed before the remainder of the BG/Q racks is delivered.

## *4.5 Participation in the DOE HPC Best Practices Workshop*

ALCF participates in the DOE HPC Best Practices workshop every year through chairing sessions and breakouts, note taking, and participating in breakouts. Most recently, ALCF staff members participated in the 5th DOE Workshop on HPC Best Practices: File Systems and Archives, held September 2011 in San Francisco, California. ALCF staff chaired the "Administration of Storage Systems" breakout session. Staff members also attended the "Business of Storage Systems" and "Reliability and Availability of Storage Systems." As part of the workshop, the ALCF submitted a position paper for use at the workshop. The ALCF is considering implementing two of the best practices discussed at the meeting: Multiple Scratch File Systems and a Hot Spare Test Bed.

## 4.6 Participation in DOE HPC Financing Improvement Team

ALCF staff was a principal coordinator and participant in the Financing Improvement Team best practices meetings. Meetings included participants from ASCR, ALCF, NERSC and OLCF. All participated and shared information to improve leasing practices, with the end goal of obtaining the very best rates possible for all facilities. General areas of discussion included lease options, requests for proposals, evaluation templates, processes, and best practices.

With lease options, a number of options were discussed and evaluated, including:

- Buyout option – the ability to pay off the lease principal at a/any point in time and save subsequent interest.
- Buydown option – the ability to pay ahead or buy down lease principal at a/any point in time to save interest without having to renegotiate the lease.
- 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.
- Discount point payment – the ability to pay a fee to be traded for a lower interest rate on a lease.
- Adding on to the lease – the ability to add equipment to an existing lease so a new agreement need not be negotiated.

The group also developed and shared templates for evaluating and comparing different leases. For example, one template provides the capability to measure the cost and benefit of paying an upfront fee to reduce the interest rate paid on the lease. Vendor lists were also shared, and a composite list will be forthcoming.

The group shared best practices 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 mechanisms and market data for evaluating interest rate ranges and associated risks.

The group continues to share information and seeks to continually improve practices beyond this first set of meetings.

## *Conclusions*

ALCF has performed innovations and improved operations in many areas in the past year.

- The center has improved the user experience with system and application software improvements, such as scaling the DDT debugger, and through partnerships with vendors and other organizations.
- ALCF has improved communication with their users through a new and easy- to-use website based on Drupal and making the initial calls with user projects more “science-centric.”
- Systems operations have been improved through implementation of a storage plan to improve performance, resiliency, and space.
- Facility efficiency has been expanded through improvements in safety, by using different-colored tiles in the machine room to handle the denser BG/Q racks.
- Finally, collaborations with other facilities included participating 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. 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 ALCF-2 project, the acquisition of a 10-PF Blue Gene/Q system called Mira. 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 Argonne Risk Manager. All risk ratings in this report are post-mitigation ratings. The ALCF currently has 36 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. 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 2012. Risks covered in depth in the 2011 OA Report will not be covered in detail.

***2012 Operational Assessment Guidance:*** *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 re-characterized 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. The RMP is also updated during the year if changes at the ALCF necessitate an update (e.g., personnel changes), 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 addressed in past reports and will not be discussed further here. Risks are tracked in a risk register using the commercial management tool, PertMaster®, which integrates with the Primavera project management tool used to manage all large ALCF projects.

The ALCF currently has 36 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 with a bottoms-up cost analysis, then used to set the contingency pool utilizing the PertMaster integration with Primavera. On the Operations side, the costs are estimated by subject matter experts and used to inform management reserves. One of the benefits of the new deep dives, discussed below, has been the dissemination and discussion of the basis of cost estimations, allowing the team to refine the mitigation costs within the risk register.

**Improvements in Risk Management since July 2011**

Triggers are risk symptoms or warning signs that are used by the risk owners in order to anticipate a risk before it becomes an issue. In the past, the ALCF risk owners had developed causes for the risks, but these proved less than useful as triggers. As a result of suggestions made during last year's onsite OAR, ALCF staff have been developing more focused triggers for the tracked risks. Currently, the Argonne Risk Manager is using the triggers as reminders for the risk owners, while the risk management team explores other ways to utilize them effectively.

Risk owners meet individually with the Argonne Risk Manager each month to update their risks. In addition, the risk management team, which includes all risk owners and the Argonne Risk Manager, as well as the ALCF Risk Officer, meets monthly to review changes from the previous month and to discuss any concerns not addressed within the current risk register. Each month, a small number of high or moderate risks are selected for a deep dive at the monthly meeting. During the deep dive, the risk owner presents all details on the selected risk, including:

- Full description of the risk,
- Basis for probability,
- Basis for impact levels,
- Triggers and any time periods when they might occur,
- Actions planned if risk occurs,
- Mitigations proposed or in progress, costs to implement, and
- Status of in-progress mitigations.

The Risk Manager captures the data and the discussion, and uses it to improve the risk data in the risk register. The risk deep dive is a new activity that is already proving valuable for helping ALCF staff to gain deeper understanding of the risks and their potential impact on the facility.

**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 monthly 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 reviews the planned actions and evaluates them 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.

Another example of a tool used for operational risk management is the OARTool. This tool, and its backend database, assists the team with identifying potential risks and triggers, developing possible mitigations, developing data for calculating impacts and probabilities, and tracking outcomes of the risk management itself. The OARTool is used for risk management on a weekly basis by the Operations team. For details on the tool and its database, please see past OAR reports. Beyond the operations of the machine, risk management is used in such diverse ways as evaluating and managing INCITE proposal risks (e.g., too few proposals, lack of diversity across science domains, too few capability proposals, etc.), safety risks in staff offices, ALCF-2 leasing risks (including the opportunity risk that interest rates could be lower than planned), support risks, etc.

**Major Risks Tracked for the Review Year**

Beginning in Q4 of FY2010, the ALCF expected to have several eventful years with Mira, scheduled to be deployed in FY2012, and the planned growth of the ALCF staff and budget 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. All risk ratings shown are post-mitigation ratings. There were twenty major operation risks tracked for the current year: six with a risk rating of *High*, thirteen with a risk rating of *Moderate*, and one that was reduced to *Low* in spring 2011. Of these twenty risks, three of the high risks were reduced to moderate, one high risk was retired, and seven of the moderate risks have been either retired or reduced to low. Ten of these risks were encountered in the past year. See the section titled "Projected Major Operating Risks for the Next Year" for current risk ratings. 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.
- Orange risks were not encountered but remain moderate or high risks.
- Purple risks have been reduced to low and were fully covered in the 2011 OAR.
- Green risks are now retired and were fully covered in the 2011 OAR.

| ID | Title | Encountered | Rating | Notes |
|---|---|---|---|---|
| 1059 | Funding is delayed or reduced (FY2011) | Yes | High | Uncertainty, a long CR, and facility need for growth to prepare for Mira combined to make this a significant challenge in the past year. This risk remains a major concern for both FY2012 and FY2013. |
| 25 | Money insufficient to hire to plan | Yes | High | Reduced rating for 2012. This remains a concern as we bring up Mira. |
| 26 | Unable to recruit qualified staff within required time | Yes | High | Reduced rating for 2012. This remains a concern as we bring up Mira. |
| 1049 | Staff Retention | Yes | Mod | Salary freeze remains in effect. This remains a concern. |
| 1050 | Insufficient disk space to support science needs | Yes | High | Mitigations completed in fall 2011 reduced rating. This remains a concern due to reduced funding, restricting ability to mitigate. |
| 995 | Interest rates on leases could be higher | No | High | Reduced rating for 2012, trigger for this risk should occur this FY. |
| 1060 | System stability problems | No | Mod | Remains a concern as Intrepid ages. |
| 1056 | System stability issues due to upgrades | No | Mod | Remains a concern as Intrepid ages. |
| 1061 | System performance issues | No | Mod | Remains a concern as Intrepid ages. |
| 1062 | Facility problems with ALCF-2 deployment | No | Mod | Remains a concern; triggers for this risk will occur in 2012. |
| 994 | Difficulties obtaining financing for ALCF-2 | No | Mod | Remains a concern; trigger for this risk will occur in 2012. |
| 1063 | IBM may back out of Mira contract | No | Mod | New risk for 2011. While probability is greatly decreased, impact remains high. |
| 990 | Electrical costs could be higher than planned | No | Low | Reduced rating to low with the signing of 2-year contract. *Fully covered in 2011 OAR. |
| 30 | Interruptions to the facility cooling | Yes | Mod | Mitigations completed fall 2011. Reduced rating for 2012. *Fully covered in 2011 OAR. |
| 31 | Interruptions to the facility power | Yes | Mod | Mitigations completed in fall 2011 reduced probability but not overall rating. This remains a concern due to aging Argonne infrastructure. *Fully covered in 2011 OAR. |
| 1051 | Component failure leads to cascading failure | Yes | Mod | Mitigations completed in fall 2011. Reduced rating for 2012. *Fully covered in 2011 OAR. |
| 1052 | Service node single point of failure | Yes | Mod | Mitigations covered in the 2011 OAR completed. Reduced rating for 2012. *Fully covered in 2011 OAR. |
| 1054 | Catastrophic failure of home FS | No | Mod | Mitigations completed in fall 2011 reduced impact. Reduced rating for 2012. *Fully covered in 2011 OAR. |
| 992 | BG/Q memory costs could be higher than planned | No | High | Retired in summer 2011. *Fully covered in 2011 OAR. |
| 1055 | Difficulty getting Myricom support in out years | Yes | Mod | Retired summer 2011. *Fully covered in 2011 OAR. |

**Table 11: Major Risks Tracked for CY2011.**

* Risk marked fully covered in the 2011 OAR will not be covered in the following sections on risks managed during 2011.

**Risks Encountered in the Review Year and Their Mitigations**

The top risks encountered in the last 12 months (except those fully covered in the 2011 OAR) are discussed below, along with the risk owners, their probability and impacts, a description of the actual problems that occurred, and the management of the risks. Some risks covered in detail in the 2011 OAR will not be discussed further in this section (see Table 11). For others, the level of detail is reduced, and a discussion of changes is provided. Risks not covered in the 2011 OAR are discussed in detail.

### *1. Funding/Budget Uncertainties (Updated)*

| 1059: Funding/Budget Uncertainties | |
|---|---|
| Risk Owner | Mike Papka |
| Probability | High |
| Impact | Schedule: Low; Cost: Low; Perf: High |
| Risk Rating | High |
| Primary Management Strategies | Implement austerity measures. Work closely with DOE sponsors to manage expectations and scope. |

#### 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, funding uncertainties ranked as 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 and also had a large lease payment scheduled for FY2011. A reduction, or delay, of incoming funds carried a high impact.

#### Management (Updated)

In conjunction with the DOE ASCR Budget Deep Dive, the ALCF prepared full-year CR and reduced budget scenarios. With DOE's concurrence, the ALCF immediately implemented the reduced budget scenario to provide maximum flexibility for the coming fiscal year. As a result of the ALCF risk management, the issue of the large lease payment due early in FY2011 was made top priority with enough time to take action to address it before the risk would occur. The lease payment was made on time and without difficulties. Once the CR ended and the budget was restored, austerity measures were eased. Hiring remained slow to manage a possible FY2012 reduction in funds.

Update: The increased funds available at the end of the fiscal year, a result of the austerity measures and the restoration of the full budget late in the year, were used to reduce the amount of equipment on the ALCF-2 leases. This will reduce the funding risk in the out years. FY2012 and FY2013 budget information indicates a likely reduction in funds from the plan of record of approximately 5-10%. Austerity measures remain in place and may be increased depending on the budget through FY2013.

### *2, 3, 4. Staffing Challenges (Updated)*

| 25, 26, 1049: Staffing Challenges | |
|---|---|
| Risk Owner | Mike Papka |
| Probability | High |
| Impact | Schedule: High; Cost: Low; Perf: High |
| Risk Rating | High |
| Primary Management Strategies | Prioritized staffing needs; re-planned work; re-tasked staff. |

#### Description (Updated)

The ALCF was, and remains, in a time of necessary growth, as it staffs up in preparation for deployment of Mira. An aggressive staff ramp was planned for FY2010 through FY2012. ALCF risk evaluation identified two key risks associated with this ramp up, and both occurred in FY2011 as a result of the budget and funding uncertainties. A third risk, the risk of losing staff due to salary freezes and budget uncertainties, had the potential to impact the planned ramp up as well. The risks have been combined into this discussion, as they are related. The risks are:

- 25: Funds unavailable to hire to plan
- 26: Unable to recruit qualified staff within required timeframes
- 1049: Unable to retain staff due to salary freeze, funding uncertainties, and heavy workloads

#### Evaluation

For the past few years, the ALCF has been working hard to increase recruiting efforts, aware of the greater needs for this time period. These intense recruiting efforts had built a full pipeline of potential candidates by fall 2010. Unfortunately, the funding uncertainties led to hiring delays and freezes, which interrupted this pipeline and prevented the ALCF from taking advantage of it. Potential candidates had also expressed concern about the funding instabilities, as well as the lab-wide salary freeze. In addition, because the facility was understaffed, the existing staff was overworked.

#### Management (Updated)

It can be very challenging to hire experienced HPC staff. Because of this, the ALCF risk management team had started arranging to execute mitigations prior to these risks occurring. When they happened, the ALCF was able to successfully continue supporting existing projects and preparing to deploy the new system, even while understaffed. 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 has re-planned work as possible, delaying planned improvements and lower-priority work. In choosing the integrated IBM storage solution, the amount of effort required to prepare for the ALCF-2 system was greatly reduced, allowing a delay in hiring a storage engineer. Staff has also been re-tasked, dropping lower-priority tasks, expanding job descriptions, and, where possible, sharing staff with other divisions.

Update: With careful and judicious management of this risk, the ALCF has successfully run the facility, prepared for the deployment of Mira, and taken delivery of the first two Test and Development racks of Blue Gene/Q. With DOE concurrence, strategic planned hires (primarily in Operations and the Applications group) were completed. At the time of this report, the facility remains below the planned staff levels by 11%. These risks will remain a concern for the next few years, and the ALCF will continue to carefully manage them.

### 5. *Insufficient Disk Space (New)*

| 1050: Insufficient Disk Space to Support Science Needs | |
|---|---|
| Risk Owner | William Allcock |
| Probability | High |
| Impact | Schedule: Low; Cost: High; Perf: High |
| Risk Rating | High |
| Primary Management Strategies | Implement quotas; Monitor; Incorporate unused capacity from original 100T system. Increase capacity using management reserves. |

#### Description

Over time, the scratch and home file systems have become full. When these file systems are full, users are unable to write output from their science runs, causing their data to be lost and/or their jobs to fail. While some files can be archived or deleted, freeing space to allow user jobs to continue, it can be challenging to motivate the users to do the clean up.

#### Evaluation

The ALCF file systems have been running close to full for over a year. The probability for this risk was raised to High in spring 2011 as the file system reached a critically full point, and requests to users to reduce usage ceased to have a measurable impact.

#### Management

A number of mitigation activities were undertaken to address this. Additional low-impact monitoring was put into place and frequent reports sent to key staff. The file systems for Intrepid were reworked to incorporate repurposed hardware from the original 100T, 8-

rack BG/P system. Hands-on assistance was provided to the users who agreed to move large amounts of their data to the archive. Once free space was increased to a reasonable level again, quotas were implemented where possible. Issues with the layout of the underlying disk storage structure currently prevent implementing the planned quotas for the large scratch file system. A workaround was developed but is on hold as the file system became full again. The disk storage hardware has room for additional drives, but the cost, because of the flooding in Thailand, is too high to implement increasing capability at this time. This will be re-evaluated when prices recover.

**Retired Risks**

The risks in the following table have been retired in the past year. Some were retired because the threat has been managed and/or no longer exists. Others were very specific short-term risks that have been reduced in impact such that they have been combined into a more generic long-term risk. Risks retired on or before the 2011 OAR are marked as such with *2011 OAR.

| ID | Title | Rating | Management Strategies | Notes |
|---|---|---|---|---|
| 1048 | Staff impacts to Intrepid due to Magellan support | Low | Hired specific personnel to administrate the Magellan cluster. Management decided Intrepid production environment was higher priority. Matrix staff from MCS as needed. | Retired upon completion of the Magellan project December 31, 2011 |
| 142 | Funding is reduced in FY2012 | High | See risk ID #1059 | Decision to manage just one funding risk. Specific FY risk was retired. |
| 8 | Poor HDF5 performance | Mod | Initial issues resolved; combined into general performance risk. | *2011 OAR |
| 137 | Poor archival storage performance | High | Initial issues resolved; combined into general performance risk. | *2011 OAR |
| 140 | Excessive load on service node during large boots | High | Implemented NAS Filer solution; additional details provided in the Innovations section. | *2011 OAR |
| 141 | High spare parts consumption | High | Issue resolved; continuing to track parts usage. | *2011 OAR |
| 139 | DDN stability issues | Mod | Initial issues resolved, combined into general stability risk. | *2011 OAR |
| 979 | Network stability problems | Mod | Initial issues resolved, combined into general stability risk. | *2011 OAR |
| 1055 | Difficulties getting Myricom support in the out years | Mod | Contract awarded that extends maintenance until June 2014. | *2011 OAR |
| 992 | BG/Q memory costs could be higher than planned | High | Memory cost locked with IBM on July 27, 2011. No additional cost to Argonne. | *2011 OAR |
| 993 | BG/Q memory costs could drop below estimates | Low | Memory cost locked with IBM on July 27, 2011 was for the contracted amount. No additional benefit to Argonne. | *2011 OAR |

**New and Re-characterized Risks Since the Last Review**

The ALCF risk culture leads to new risks being identified and risks being re-characterized on a regular basis. In the past year, 12 new risks have been identified, developed, and are now being monitored and tracked as part of the Steady State risk register. Most active risks in the risk register have been re-characterized, adding triggers, adjusting them to accommodate changes, updating mitigations, etc. in the past year. For the new risks (shown in blue in the table below), per the guidance, only those added since the last review are shown and discussed. The previous new risks can be viewed in the 2011 OAR. Because so many risks have been re-characterized in the past year, only those that had the risk rating change, or other significant change are covered (shown in purple in the table).

| ID | Title | Rating | Management Strategies | Notes |
|---|---|---|---|---|
| 1050 | Insufficient disk space. | High | Implement quotas; use monitoring scripts; increase capacity. | New risk in 2011, re-characterized after 2011 OAR. Initial rating high, lowered to moderate after several mitigations completed. |
| 1063 | IBM may back out of the Mira contract. | Mod | Run Intrepid and existing hardware for longer while locating and finding another system; verify sufficient management reserves exist to increase power/cooling, or reduce scope; utilize new infrastructure hardware with new system. | New risk in 2011, added after IBM cancelled the Blue Waters contract. Initial rating was moderate, but it has been reduced to low for 2012, based on the delivery of a significant portion of the LLNL racks and both of the ALCF T&D racks. |
| 1064 | I/O libraries are inadequate for user needs. | Low | Work with developers to stabilize, add features, and/or increase performance. Use management reserves to fund work on the libraries. | New risk in 2011, added for Mira. Triggers for this risk occur in 2012. |
| 1050 | Insufficient disk space to support science needs. | Mod | See discussion in Major Risks Encountered in 2011. | Risk rating reduced from High to Moderate after completion of mitigations. This remains a concern, and will until additional capacity can be added, or until Intrepid is de-commissioned. |
| 25 | Money insufficient to hire to plan. | Mod | See discussion in Major Risks Encountered in 2011. | Risk rating reduced from High to Moderate, based on full funding and hires managed during 2011. This risk remains a concern until ALCF is fully staffed. |
| 995 | Interest rates on leases could be higher than planned. | Mod | Competitive procurement; reduce scope; inform management reserves; collaborate with other DOE facilities. | Risk rating reduced from High to Moderate based on 1) reduced uncertainty as the time for the lease gets closer and 2) the results of the collaboration with OLCF and NERSC. This will remain a concern until the lease rate is locked. |
| 26 | Unable to recruit qualified staff within required timeframes. | Low | See discussion in Major Risks Encountered in 2011. | Risk rating reduced from High to Low, based on full funding and hires managed during 2011. With the current levels and quality of staff, we can now afford to higher less-skilled staff and train them. |
| 30, 31 | Interruption to facility cooling. | Low | Covered fully in the 2011 OAR. | Risk rating reduced from Moderate to Low after the completion of the mitigations. |
| 1051 | Component failure leads to cascading failure. | Low | Covered fully in the 2011 OAR. | Risk rating reduced from Moderate to Low after completion of the mitigations. |
| 1052 | Service node single point of failure. | Low | Covered fully in the 2011 OAR. | Risk rating reduced from Moderate to Low after completion of the mitigations. |
| 1054 | Catastrophic failure of home FS | Low | Covered fully in the 2011 OAR. | Risk rating reduced from Moderate to Low after completion of the mitigations. |

**Projected Major Operating Risks for the Next Year**

The top operating risks projected for the next year are listed below, along with the current risk rating and management strategies for the risk.

| ID | Title | Rating | Management Strategies |
|---|---|---|---|
| 1059 | Funding 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 | Look into non-salary-based compensation (bonuses, office improvements, work from home); cross-train team members. |
| 995 | Interest rates could be higher than planned | Mod | Competitive procurement of financing; reduce scope; budget reserves based on interest rate models. |
| 994 | Problems obtaining financing for ALCF-2 | Mod | Lock financing as early as possible; use relationship with IBM; delay purchase until financing is available. |
| 31 | Interruptions to the facility provided power | Mod | Continued high-quality communications with Facilities Management and Services (FMS); Lab-wide improvements for the aging Argonne infrastructure are under way. We will continue to monitor FMS efforts and provide data for improvement choices. |
| 1056 | System stability issues due to upgrades | Mod | Work with vendor closely on planning; develop rollback plans; deep tests on T&D systems; increase monitoring and notification. |
| 1060 | System stability problems | Mod | Initial issues resolved, combined into general stability risk. |
| 1061 | System performance issues | Mod | Baseline; regression testing; monitoring. |
| 1062 | Facility problems with ALCF-2 deployment | Mod | Monitor LLNL deployment; vendor facility visits; test early; independent analysis; prepare safety and contingency plans. |
| 1063 | IBM may back out of the Mira contract | Mod | Run Intrepid and existing hardware for longer while locating and finding another system; verify sufficient management reserves exist to increase power/cooling, or reduce scope; utilize new infrastructure hardware with new system. |

### *Conclusions*

The ALCF uses a proven risk management strategy that is documented in its Risk Management Plan (RMP). The document is reviewed and updated regularly 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 for all staff, from senior management through summer students, and is used both formally and informally. A formal risk assessment is performed for every major activity; informal ones are used for smaller activities within the ALCF. Risks are monitored and tracked using the commercial risk management tool PertMaster™. This risk register is a living document, and in the past year, eleven risks have been retired, twelve new ones added, and every active risk updated. Beyond the risk register, 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: 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 2012 metrics and targets as proposed in the August 2011 OAR report. These metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2013, the proposed metrics and targets for the current production resources remain the same as for 2012. For the new BG/Q system, Mira, expected to go into production some time in 2013, its proposed targets will be the standard first year of operation targets, with the exception of the capability metric, which has been modified to be a two-tier metric. This new metric ultimately provides a similar 20% of the system capability threshold, while accommodating the new partition sizes of Mira.

***2012 Operational Assessment Guidance:*** *The Facility should provide a summary table of the metrics and targets agreed upon for the review of Calendar Year 2012 and include the target and actual values of similar metrics used for 2011 for comparison. The Facility should also provide metrics and targets under consideration for the review of CY 2013. 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.*

## *6.1 Overview*

The 2012 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2013, the proposed metrics and targets for the current production resources remain the same as for 2012 and are covered in Section 6.2. For the new BG/Q system, Mira, expected to go into production some time in 2013, its proposed targets will be the standard first year of operation targets, with the exception of the capability metric. Because of the architecture and the configuration of Mira, the current job capability threshold of 20% of the system doesn't fit the partition sizes that are available. As a result, ALCF, with DOE concurrence, proposes a two-tier capability target, with 15% of the INCITE jobs running on 16.7% to 33.2% of the system (131,072 to 245,760 cores) and 5% of the INCITE jobs running on 33.3% of the system (245,761cores or larger). A detailed discussion of this new metric is provided in Section 6.3. The summed metric is 20% of the INCITE jobs will be run on 20% of the machine or larger. The ALCF expects to increase the % of jobs in later years as the users become comfortable on the new architecture.

## 6.2 ALCF 2012 OA Performance Metrics

The OA performance metrics, 2011 targets and actuals, and agreed upon 2012 targets are presented in Table 12.

| Area | | Metric | 2011 Targets | 2011 Actuals | 2012 Targets |
|---|---|---|---|---|---|
| User Results | CY2011 User Survey | **Overall Satisfaction** | 3.5/5.0 | 4.4/5.0 | 3.5/5.0 |
| | | **User Support** | 3.5/5.0 | 4.5/5.0 | 3.5/5.0 |
| | | **Problem Resolution** | 3.5/5.0 | 4.5/5.0 | 3.5/5.0 |
| | | **Response Rate** | 25% | 29.0% | 25% |
| | | % user problems addressed within three working days | 80% | 90.2% | 80% |
| Business Results | | Intrepid Overall Availability | 90% | 94.0% | 90% |
| | | Intrepid Scheduled Availability | 95% | 97.8% | 95% |
| | | Intrepid Capability Usage (Old metric) | 300M | 685M | – |
| | | Intrepid Capability Usage (New metric) | - | 57.0% | 40% |

**Table 12: ALCF Performance Metrics – 2011 Targets, 2011 Actuals, and Agreed Upon 2012 Targets.**

## 6.3 ALCF Proposed 2013 OA Performance Metrics

The OA performance metrics, agreed upon 2012 targets, and 2013 proposed targets are shown in Table 13.

### 6.3.1 Mira Capability Metric

Mira's configuration consists of 48 racks (786,432 cores), therefore 20% of the system is 9.6 racks, which is not a practical size partition for the Blue Gene/Q architecture. Of all the possible partitions on Mira, the closest feasible sizes are 8, 9, 12, and 16 racks. While very close in size, the 9-rack partition would not be a good match with the workflow of our projects and would cause significant drain costs. The 8-rack partition is an attractive target size because it is the same number of cores (131,072) as Intrepid's 32-rack partition (though Mira jobs are likely to have more threads). Of the feasible partitions greater than 8 racks, 16-rack partitions (262,144 cores) provide a power-of-two growth that is appropriate for many applications in the expected Mira workload. Twelve-rack partitions (196,608 cores) are feasible and it is believed that some projects may use this size. Because of these complications, a modified two-tier capability target is proposed for Mira. This target will be that 20% of the core-hours used by INCITE projects will be generated by capability jobs of two sizes: 15% of the INCITE core-hours delivered by jobs that use 131, 072 cores (8 racks) or 196,608 cores (12 racks) and 5% of the core-hours generated by jobs that use 262,144 cores or greater (16 or more racks). This two-tier definition of “capability use” ensures that capability jobs will use an average of 10 racks, which is slightly over 20% of the entire Mira configuration.

The 10-rack average use results from the requirement that ¾ of the 20% of INCITE core-hours (15%) come from jobs that use a minimum of 8 racks and ¼ of the 20% (5%) come from jobs that use a minimum of 16 racks: ¾*8+ ¼*16=10.

| Area | | Metric | 2012 Targets | 2013 Targets |
|---|---|---|---|---|
| User Results | User Survey | **Overall Satisfaction** | 3.5/5.0 | 3.5/5.0 |
| | | **User Support** | 3.5/5.0 | 3.5/5.0 |
| | | **Problem Resolution** | 3.5/5.0 | 3.5/5.0 |
| | | **Response Rate** | 25% | 25% |
| | | % user problems addressed within three working days | 80% | 80% |
| Business Results | | Intrepid Overall Availability | 90% | 90% |
| | | Intrepid Scheduled Availability | 95% | 95% |
| | | Intrepid Capability Usage (20% of system) | 40% | 40% |
| | | Mira Overall Availability | - | 80% |
| | | Mira Scheduled Availability | - | 85% |
| | | Mira Capability Usage A (16.7% of system) | - | 15% |
| | | Mira Capability Usage B (33.3% of system) | - | 5% |

**Table 13: ALCF Performance Metrics – 2011 Targets, Agreed Upon 2012 Targets, and Proposed 2013 Targets.**

## *Conclusions*

The agreed upon 2012 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2013, the proposed metrics and targets for the current production resources remain the same as for 2011. For the new BG/Q system, Mira, the proposed targets will be the standard first year of operation targets, consistent with metrics and targets used at other facilities, with the exception of the capability metric. Because of the architecture and the configuration of Mira, the current job capability threshold of 20% of the system doesn’t fit the partition sizes that are available and a two-tier capability metric is proposed that provides 20% of the INCITE jobs will run on an average of 20% of the system. Achieving the agreed upon 2012 and the proposed 2013 targets will demonstrate 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 – ALCF 2011 User Survey Analysis

Richard Coffey
ALCF User Experience Team
January 20, 2012

## *Overview of Changes for the ALCF 2011 User Survey*

**Summary**

ALCF practices continuous process improvement throughout the facility. The user survey ties directly to OMB and DOE metrics the leadership uses to measure the performance of the ALCF. This year the organization of the user survey was improved and questions added to increase understanding of ALCF users. In 2012, the fidelity of this critical instrument will be improved with the help of our various communities and external expertise in statistics and survey management.

**Reorganization of the Survey**

Last year's survey 2010 improved upon the 2009 survey by asking questions focused on providing greater insight into the particular pain points of the users. As a first step to categorizing by user allocation categories (i.e., INCITE, ALCC, Discretionary), separate surveys were provided for the different categories. This proved slightly awkward when collating the surveys, so the 2011 survey is now a single instrument that was provided to all active users. Logic in the survey was used to eliminate questions that do not pertain to a user segment (e.g., Director's Discretionary allocations do not have a Catalyst assigned to them.) At the recommendation of the ALCF DOE Program Manager, the questions for the user support and problem resolution metrics were broken into clearly defined sections.

**Looking Ahead: The ALCF 2012 User Survey**

ALCF is working in 2012 to properly instrument the survey. ALCF is in the process of engaging survey consultants from well-respected, regional institutions. The consultant will help improve the survey via a process similar to the following: First, derive a core set of questions from the ALCF's primary goals and objectives. Next, conduct a focus group – current plans are to engage the User Advisory Council as that focus group. The focus group will draw out themes in the qualitative results. ALCF will then design a set of questions that drive those surveyed to provide focused insight into the core questions. The consultant will ensure that these questions do not violate any rules of survey analysis as well as assist in the final analysis.

**Items That Did Not Change in the 2011 Survey**

The survey population in 2011 continued to be the active user accounts in the year the survey covers. All users with active accounts in December 2010, and all PIs with active projects, were sent a survey. In addition, questions that used the Likert-type scale used in last year's OMB results remained untouched. Finally, ALCF used Survey Monkey to collect and analyze the results of the user survey.

## *Specific Changes in the 2011 Survey*

**One Survey for All Groups**

As stated above, ALCF sent out a single survey. This reduced the complexity of managing the instruments and ensured it was easy to compare the responses from the different types of allocations.

**Clearly Defined Sections for Critical Metrics**

ALCF leadership and the ALCF program manager agreed that grouping the user support and problem resolution questions in sections clarified the basis for the questions in those sections.

**Problem Resolution is a Subset of the User Support Metric**

ALCF leadership and the ALCF program manager agreed that problem resolution is a subset of the user support metric. This was not the case in 2010 so problem resolution was not added to the user support metric.

## *Questions Added in the 2011 Survey*

***"Please briefly describe the science goal of your ALCF allocation."***

This information can be obtained from the user's original proposal. However, asking them this question one more time measures if their science goal changed in their mind. Their answer to this question sets up the next question: *"Were you able to achieve the science goals for your project?"*

***"The Catalyst understood the deadlines and constraints of my project."***

***"The Catalyst understood the core scientific questions driving my research."***

***"The Catalyst helped me work through one or more critical performance issues in my code."***

***"The Catalyst helped me approach my scientific problem in a fundamentally different way."***

These four questions were added as a result of discussions with the User Advisory Committee. The committee has underscored over and over the positive effect of a Catalyst on an allocation because of these reasons. The Catalysts and Performance Engineers improve the performance of a project because they understand the science, the culture, and the incentives of the domain and its interface to computational science.

***"The ALCF staff resolved my problems within a reasonable amount of time."***

***"The ALCF staff provided additional follow-up and materials."***

ALCF added these two questions because they are the "end caps" of problem resolution: timeliness and going above and beyond; essentially the difference between user service and user experience. Measures user perception of not just the quality, but the speed at which problems are solved.

***"ALCF provided sufficient software infrastructure to install, compile, test, and run the tools and libraries I needed."***
ALCF did not distinguish between ALCF libraries and tools and the ease of getting a user's own tools and libraries installed on the Blue Gene/P.

***"ALCF staff enabled me to utilize a new tool or library that improved the performance of my code."***

***"By the end of the workshop, ALCF staff understood my science."***

***"By the end of the workshop, ALCF staff understood my computational bottlenecks."***

Measures the primary reasons for running these workshops. These three questions clearly embody why ALCF has conducted the Getting Started and Leap to Petascale Workshops.

### *Questions Deleted in the 2011 Survey*

***"The amount of communication and support from my Catalyst in 2010 was: (just right/not enough)"***
This question seemed redundant with the other questions in the Catalyst section of the 2011 survey as well as the problem resolution section. It was also part of the metric.

***"Ideally, how often would you like to communicate with your Catalyst?"***
These two questions did not provide much added insight in 2010. Similar questions could be used in the 2012 report with some professional guidance on how to make these meaningful. These questions were not part of the metric calculations.

***"Did you participate in one or more of the monthly ALCF User Calls in 2010?"***
Since ALCF has decided to eliminate the ALCF User Call based upon low attendance and recommendations from the User Advisory Committee, it was decided to eliminate this question from the 2011 survey. This question was not part of the metric calculations.

***"What other resources could we provide that would be helpful to you now or in the future?"***
In an effort to trim the length of the 2011 survey and reduce the complexity of the analysis, this question was eliminated.

### *Other Changes to the 2011 Survey*

***An estimated time to completion and types of questions description added to the invitation***
Not shown on the survey itself, the survey invitation email stated the estimated time to completion and the types of questions on the survey. This sets expectations for the end user and attempts to increase the number of respondents.

***"Were you able to achieve…" question has an addendum to accommodate ALCC, INCITE, and Director's Discretionary allocations***
Previously the different allocations were polled separately. This question allows for answers from all three types of allocations.

### *The number of comment sections were reduced*

In an effort to increase the number of respondents, the comment sections were trimmed so people don't feel obligated to fill in so many comments.

### *Indications of the end of a section, a progress bar and a percentage complete have been added*

By providing signposts, the user is relieved of the anticipation of yet another question.

### *Descriptions of the purpose of the section have been added*

When users have a better sense of the purpose of a section, they tend to give answers to the questions as intended. Descriptions of the purpose help with this.

*This page intentionally left blank.*

# Appendix B – Calculations

## *Scheduled Availability Calculation Details*

***2012 Operational Assessment Guidance:***
***Description***: *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*

***2012 Operational Assessment Guidance:***
***Description***: *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. Some hardware platforms have to shut down in order to do hardware maintenance, so the availability tends to be binary, either the whole system is available or it is not. This is not the case with Blue Gene. The Blue Gene is capable of taking individual node cards (32 nodes) off line for maintenance, while the rest of the machine continues to run. However, ALCF also takes into account scheduling policy. By policy, jobs smaller than 512 nodes (2,048 cores) are not allowed to run, and that is the smallest number of nodes that will be allocated. Therefore, if a single node were to fail for exactly one hour, it would be recorded as 2,048 cores * 3600 seconds = 7,372,800 core-seconds of down time. ALCF has only one production file system. 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*

***2012 Operational Assessment Guidance:***
***Description:*** *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*

***2012 Operational Assessment Guidance:***
***Description:*** *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 in the PBS format. Each night at midnight, a script runs and processes the day's records and imports the data into the internal accounting package, called clusterbank. Clusterbank 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 clusterbank 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:

$$Utilization = \frac{CoreHours\ consumed}{Total\ core\ Hours\ that\ were\ available} * 100$$

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***: There is little calculation involved with the capability numbers. The data for everything except the job usage by size graph is simply 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 C – ALCF Director's Discretionary Projects

## January 1, 2011 – December 31, 2011

**ALCF Discretionary Projects**
January 1, 2011 - Decemeber 31, 2011 active allocations

| Project Name | Project Title | PI | Institution | Allocated | Charged | Field | Category |
|---|---|---|---|---|---|---|---|
| CompBIO | Multiscale Simulations in Biology: Evolution and Ecology of Microbes | Rick Stevens | Argonne National Labor | 400,000 | 0 | Biological Sciences | Code Development/Support |
| fmobench | FMO benchmarking and tuning | Yury Alekseev | Argonne Leadership Cor | 500,000 | 726,467 | Biological Sciences | Code Development/Support |
| UNRES-DEV | Optimization for memory and development of the UNRES software for coarse-grained molecular dynamics simulations of proteins | Harold A. Scheraga | Cornell University | 200,000 | 0 | Biological Sciences | Code Development/Support |
| VLP123 | Computer-aided design of vaccine nanoparticles | Peter Ortoleva | Indiana University | 2,250,000 | 556,153 | Biological Sciences | Code Development/Support |
| neurosim | Neural Network Simulations | Mark Hereld | University of Chicago | 50,000 | 0 | Biological Sciences | Code Development/Support |
| IIT-scott-lipids | Atomistic and Coarse-grained molecular dynamics simulations of model biological membranes | H. Larry Scott | Illinois Institute of Techı | 2,867,200 | 0 | Biological Sciences | Code Development/Support |
| HTCScienceApps | HTC-Linux Loosely-coupled science applications | Michael Wilde | Argonne National Labor | 3,000,000 | 1,253,738 | Biological Sciences | Code Development/Support |
| Voth-Discretionary | MULTISCALE MOLECULAR SIMULATIONS | Gregory A. Voth | Argonne National Labor | 5,000,000 | 842,225 | Biological Sciences | ESP |
| MultiscaleMolSim_e | Multiscale Molecular Simulations at the Petascale | Gregory Voth | University of Chicago | 25,000,000 | 9,665,956 | Biological Sciences | ESP |
| NAMD_esp | NAMD - The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based on a Polarizable Force Field | Benoit Roux | University of Chicago | 25,000,000 | 16,357,686 | Biological Sciences | ESP |
| pocket_druggability | Predicting druggability of protein interaction sites by pocket optimization | John Karanicolas | University of Kansas | 500,000 | 0 | Biological Sciences | INCITE Prep |
| MUPHY | Multiscale Hemodynamics | Amanda Peters | Harvard University | 500,000 | 0 | Biological Sciences | INCITE Prep |
| IBM-performance | BG/P performance runs carried out by IBM employees | Kalyan Kumaran | Argonne Leadership Cor | 500,000 | 204,690 | Biological Sciences | Internal/Support |
| AlchemicalDrugDesi | Alchemical Drug Design by Ensemble Monte Carlo Molecular Dynamics Simulation | Michel A. Cuendet | New York University | 1,000,000 | 0 | Biological Sciences | Science |
| nrnmultisend | Load balance spike communication in large scale neural networks | Michael Hines | Yale University | 970,000 | 612,544 | Biological Sciences | Science |
| Nanofluid | Mesoscale simulation of nanoflow and surface properties of amorphous materials | Zeke Insepov | Argonne National Labor | 500,000 | 0 | Biological Sciences | Science |
| fmo4bgp | Fragment Molecular Orbital Calculations on BlueGene/P | Graham Fletcher | Argonne Leadership Cor | 3,000,000 | 2,341,473 | Chemistry | Code Development/Support |
| MLearning4CCS | Machine Learning for the exploration of chemical compound space | O. Anatole von Lilienfeld | Argonne Leadership Cor | 1,000,000 | 0 | Chemistry | Code Development/Support |
| nwchem4bgp | Testing and tuning NWChem for BlueGene/P and studies of nonlinear optical properties of conjugated chromophores | Jeff Hammond | University of Chicago | 16,000,000 | 927,622 | Chemistry | Code Development/Support |
| Bulk_Properties_esp | High Accuracy Predictions of the Bulk Properties of Water | Mark Gordon | Iowa State University | 14,000,000 | 12,276 | Chemistry | ESP |
| HSCD_esp | High-Speed Combustion and Detonation (HSCD) | Alexei Khokhlov | University of Chicago | 20,000,000 | 2,964,549 | Chemistry | ESP |
| MADNESS_MPQC_e | Accurate Numerical Simulations Of Chemical Phenomena Involved in Energy Production and Storage with MADNESS and MPQC | Robert Harrison | Oak Ridge National Labo | 15,000,000 | 729,941 | Chemistry | ESP |
| ACESIII | Port ACES III and SIAL | Erik Deumens | University of Florida | 1,500,000 | 377 | Chemistry | INCITE Prep |
| eon | Long time simulations of catalytic reaction dynamics | Graeme Henkelman | University of Texas at A | 1,000,000 | 204,866 | Chemistry | INCITE Prep |
| HSCD | DNS simulations of high-speed combustion and detonations | Alexei M. Khokhlov | University of Chicago | 300,000 | 0 | Chemistry | INCITE Prep |
| i-SNSDCs | Multi-scale modeling of catalytic interfaces based on 2-D sub-nano surface-deposited clusters | Anastassia N. Alexandrova | University of California, | 1,000,000 | 0 | Chemistry | INCITE Prep |
| Allinea | Improved debugging memory usage for BG/P | Kalyan Kumaran | Allinea Software | 650,000 | 278,890 | Computer Science | Code Development/Support |
| BloodFlowGB | Preparation for Gordon Bell Competition | Vitali Morozov | Argonne Leadership Cor | 15,000,000 | 13,263,662 | Computer Science | Code Development/Support |
| CharmRTS | Charm++ and its applications | Laxmikant V. Kale | University of Illinois at U | 4,000,000 | 866,374 | Computer Science | Code Development/Support |
| CScADS | Center for Scalable Application Development Software | Tim Williams | Argonne Leadership Cor | 200,000 | 0 | Computer Science | Code Development/Support |
| Fathom | Fathom: geometry, mesh generation, and related technologies | Tim Tautges | Argonne National Labor | 1,100,000 | 15,066 | Computer Science | Code Development/Support |
| FusionFS | Distributed File Systems for Exascale Computing | Ioan Raicu | Illinois Institute of Techı | 50,000 | 0 | Computer Science | Code Development/Support |
| graph500 | Graph500 benchmark run on Intrepid | Andrew Lumsdaine | Indiana University | 2,250,000 | 3,193,608 | Computer Science | Code Development/Support |
| HDF5 | Daily Testing of HDF5 | Quincey Koziol | The HDF Group | 200,000 | 0 | Computer Science | Code Development/Support |
| IOFSL | I/O Forwarding Scalability Layer | Rob Ross | Argonne National Labor | 1,500,000 | 727,143 | Computer Science | Code Development/Support |
| MTCScienceApps | Many Task Computing Science Applications | Michael Wilde | Argonne National Labor | 1,500,000 | 86,669 | Computer Science | Code Development/Support |
| OnSite_Workshop | Project for ALCF Workshops Onsite | David Martin | Argonne National Labor | 20,000 | 2,271 | Computer Science | Code Development/Support |
| Orio | Orio: Annotation-Based Performance Tuning of C and Fortran Applications | Boyana Norris | Argonne National Labor | 100,000 | 11,712 | Computer Science | Code Development/Support |
| PARTS | PArallel Run-Time Systems | Jeff Hammond | Argonne National Labor | 7,000,000 | 1,735,150 | Computer Science | Code Development/Support |
| PBGL | Parallel Boost Graph Library | Andrew Lumsdaine | Indiana University | 250,000 | 0 | Computer Science | Code Development/Support |
| PetaIO | Understanding I/O on Petascale Platforms | Robert Ross | Argonne National Labor | 250,000 | 17,220 | Computer Science | Code Development/Support |
| Repast | Exascale Agent-based Modeling System | Michael North | Argonne National Labor | 700,000 | 92,987 | Computer Science | Code Development/Support |
| Spiral_OnBGP | Generation of a Multicore FFT library for BlueGene/P | Franz Franchetti | Carnegie Mellon Univer | 50,000 | 0 | Computer Science | Code Development/Support |
| TAU | Parallel Performance Evaluation Using the TAU Performance System | Sameer Shende | ParaTools, Inc. | 500,000 | 2,309 | Computer Science | Code Development/Support |
| TotalView | TotalView Debugger on Blue Gene P | Peter Thompson | TotalView Technologies | 1,250,000 | 1,946,413 | Computer Science | Code Development/Support |
| UltraVis | UltraVis Institute Research for Extreme Scale Visualization | Robert Ross | Ultrascale Visualization | 1,300,000 | 811,646 | Computer Science | Code Development/Support |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| ZeptoOS | ZeptoOS project | Pete Beckman | Argonne National Labor | 1,500,000 | 306,930 | Computer Science | Code Development/Support |
| BGQtools_esp | Enabling Petascale Science on BG/Q: Tools, Libraries, Programming Models, & Other System Software | Kalyan Kumaran | Argonne Leadership Coı | 8,000,000 | 509,868 | Computer Science | ESP |
| Catalyst | Catalyst | Katherine Riley | Argonne Leadership Coı | 5,000,000 | 13,102,962 | Computer Science | Internal/Support |
| CobaltDevel | Cobalt Development | Narayan Desai | Argonne Leadership Coı | 1 | 68,260 | Computer Science | Internal/Support |
| disk_study | Intrepid disk I/O analysis | Justin Binns | Argonne Leadership Coı | 750,000 | 252,778 | Computer Science | Internal/Support |
| DynamicDevel | Development and Improvement of Dynamic Programming and Application Environments at Scale | William Scullin | Argonne Leadership Coı | 500,000 | 522,590 | Computer Science | Internal/Support |
| LowPriority | LowPriority | Katherine Riley | Argonne Leadership Coı | 5,000,000 | 0 | Computer Science | Internal/Support |
| Operations | Systems administration tasks | Susan Coghlan | Argonne Leadership Coı | 0 | 5,574,885 | Computer Science | Internal/Support |
| Operations_Test | Operations Infrastructure Testing | Cheetah Goletz | Argonne National Labor | 500,000 | 118 | Computer Science | Internal/Support |
| Performance | Performance | Kalyan Kumaran | Argonne Leadership Coı | 5,000,000 | 5,062,944 | Computer Science | Internal/Support |
| TestHarness | TestHarness | Eric Pershey | Argonne Leadership Coı | 2,000,000 | 0 | Computer Science | Internal/Support |
| visualization | Visualization and Analysis Research and Development for ALCF | Michael E. Papka | Argonne Leadership Coı | 1,200,000 | 640,818 | Computer Science | Internal/Support |
| nu-xanalytics | Scalable data analytics and I/O | Alok Choudhary | Northwestern Universit | 250,000 | 283,811 | Computer Science | Science |
| GFDL_esp | Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model | Venkatramani Balaji | GFDL | 15,000,000 | 0 | Earth Science | ESP |
| GroundMotion_esp | Using Multi-scale Dynamic Rupture Models to Improve Ground Motion Estimates | Thomas Jordan | University of Southern ( | 25,000,000 | 0 | Earth Science | ESP |
| climuq | Uncertainty Quantification of Climate Sensitivity | Richard Klein | Lawrence Livermore Na | 200,000 | 0 | Earth Science | INCITE Prep |
| uqsci | Climate Modeling Uncertainty Quantification | Richard Klein | Lawrence Livermore Na | 300,000 | 71,382 | Earth Science | INCITE Prep |
| MidwestClimate | High-resolution Simulation of Midwestern Climate Using NRCM | J.W. Larson | Argonne National Labor | 100,000 | 0 | Earth Science | Science |
| prec_sense | Sensitivity and uncertainty of precipitation of the GFDL high resolution model | Laura Zamboni | Argonne National Labor | 100,000 | 111,402 | Earth Science | Science |
| theEyeOfTheHurrica | Large Scale Hurricane Simulations | John Michalakes | Argonne National Labor | 17,010,000 | 5,214,409 | Earth Science | Science |
| HohokamWaterMgr | Hohokam Water Management | John T. Murphy | Argonne National Labor | 125,000 | 77,975 | Earth Science | Science |
| Oil | Oil Transport Analysis | Paul Fischer / Tamay Ozgokn | University of Miami | 3,000,000 | 27,584 | Earth Science | Science |
| Heron | Wind turbine array fluid dynamic and aero-elastic simuations | John Michalakes | NREL | 2,000,000 | 3,640 | Energy Technologies | Science |
| RERTR | RERTR_Project | Z.Insepov | Argonne National Labor | 2,000,000 | 714,886 | Energy Technologies | Science |
| PHASTA_GB | PHASTA-GLEAN Gordon Bell submission | Ray Loy | University of Colorado | 3,500,000 | 8,010,306 | Engineering | Code Development/Support |
| Autoignition_esp | Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow | Christos Frouzakis | Swiss Fed. Inst. Tech. | 10,000,000 | 715,923 | Engineering | ESP |
| CFDAnisotropic_esp | Petascale, Adaptive CFD | Kenneth Jansen | University of Colorado | 14,000,000 | 309,039 | Engineering | ESP |
| TurbChannelFlow_e | Petascale Direct Numerical Simulations of Turbulent Channel Flow | Robert Moser | University of Texas at A | 10,000,000 | 0 | Engineering | ESP |
| DNS-TURB-DROPLET | The interactions between vaporizing liquid droplets and a turbulent flow: Fully resolved DNS | Said Elghobashi | University of California, | 1,200,000 | 1,289,298 | Engineering | INCITE Prep |
| INCITE2010_TELEM/ | Testing TELEMAC on Intrepid | Charles Moulinec | STFC Daresbury Laborat | 340,000 | 141,651 | Engineering | INCITE Prep |
| ParaFlow | Multiphase Mixing Simulations using the Implicit Lattice Kinetics Method | David R. Rector | Pacific Northwest Natio | 200,000 | 186,590 | Engineering | INCITE Prep |
| RTInstability | Large Eddy Simulation of turbulent mixing by Rayleigh-Taylor instability | James Glimm | SUNY at Stony Brook | 500,000 | 901,907 | Engineering | INCITE Prep |
| MFiX | Achieving Scalable, Clean & Efficient Coal Gasifier Designs with High Performance Computing | Aytekin Gel | DOE | 200,000 | 0 | Engineering | INCITE Prep |
| NekLBM | Lattice Boltzamann Simulations for Fluids | Misun Min | Argonne National Labor | 1,100,000 | 87,746 | Engineering | Science |
| SMM | Development and application of self-consistent multiscale method (SMM) | Suvranu De, DSc | Rensselaer Polytechnic | 100,000 | 1,265 | Engineering | Science |
| Gpaw | Optimization of real-space, PAW-based DFT codes for use on the BlueGene/P | Nichols Romero | Argonne Leadership Coı | 2,000,000 | 1,300,460 | Materials Science | Code Development/Support |
| Oxygen_defects | First-principle investigations of oxygen defects in metal-oxide-metal heterostructures | Olle Heinonen (MSD) | Argonne National Labor | 500,000 | 0 | Materials Science | Code Development/Support |
| Mat_Design_esp | Materials Design and Discovery: Catalysis and Energy Storage | Larry Curtiss | Argonne National Labor | 11,000,000 | 6,394,078 | Materials Science | ESP |
| qe_parallel_ph | Phonon Parallelization in Quantum ESPRESSO | William Parker | Argonne National Labor | 500,000 | 15,818 | Materials Science | INCITE Prep |
| ucl_qmc | Quantum Monte Carlo methods for solids and liquids | Dario Alfe | University College Lond | 500,000 | 6,227 | Materials Science | INCITE Prep |
| VDW-CMAT | Van der Waals Interactions in Complex Materials | Alexandre Tkatchenko | Fritz Haber Institute of t | 1,000,000 | 1,003,722 | Materials Science | INCITE Prep |
| MolSaltMix | Design of low melting heat-transfer-fluids via adaptive screening in composition space based on alchemical interpolations between replicas. | Sai Jayaraman | Sandia National Laborat | 1,700,000 | 2,019,022 | Materials Science | INCITE Prep |
| CACS | Hierarchical Petascale Simulation Framework for Stress Corrosion Cracking | Priya Vashishta | University of Southern ( | 3,000,000 | 0 | Materials Science | Science |
| MatEne_Sim | Algorithms, codes and data analysis in quantum simulations of materials for energy applications | Giulia Galli | University of California, | 5,000,000 | 321,411 | Materials Science | Science |
| materialsdesign | Materials Design from First Principles Calculations | Larry Curtiss; Nick Romero | Argonne National Labor | 5,000,000 | 0 | Materials Science | Science |
| PNIPAM_MD_Simul | Large scale atomistic simulations to evaluate dynamics of conformational transitions in thermo-sensitive polymers | Subramanian K.R.S. Sankarar | Argonne National Labor | 400,000 | 77,668 | Materials Science | Science |
| StochProg | Stochastic Programming And Applications for Energy Systems | Mihai Anitescu | Argonne National Labor | 2,750,000 | 1,812,723 | Mathematics | Code Development/Support |
| accel_code_perf | Performance Studies of the IMPACT-Z Parallel Particle-In-Cell Accelerator Modeling Code | Ji Qiang | Lawrence Berkeley Nati | 200,000 | 52,015 | Physics | Code Development/Support |
| AbInitioC12_esp | Ab-initio Reaction Calculations for Carbon-12 | Steven C Pieper | Argonne National Labor | 15,000,000 | 0 | Physics | ESP |
| DarkUniverse_esp | Cosmic Structure Probes of the Dark Universe | Salman Habib | Los Alamos National Lal | 22,000,000 | 2,008,005 | Physics | ESP |
| TurbNuclComb_esp | Petascale Simulations of Turbulent Nuclear Combustion | Don Lamb | University of Chicago | 20,000,000 | 1,083,662 | Physics | ESP |
| LatticeQCD_esp | Lattice Quantum Chromodynamics | Paul Mackenzie | Fermilab | 25,000,000 | 22,989,271 | Physics | ESP |
| PlasmaMicroturb_e | Global Simulation of Plasma Microturbulence at the Petascale & Beyond | William Tang | PPPL | 25,000,000 | 19,729 | Physics | ESP |
| MF_ADD | Accretion Disk Dynamics: the multifluid regime | Dr Turlough Downes | Dublin City University | 462,000 | 62,106 | Physics | INCITE Prep |
| sncore | 3D Core-Collapse Supernova Simulations | Adam Burrows | Princeton University | 5,000,000 | 0 | Physics | INCITE Prep |
| Femtomagnetism | First-principles calculation of laser induced ultrafast magnetism | Guoping Zhang | Indiana State University | 5,250,000 | 7,659,926 | Physics | INCITE Prep |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| HEIGHTS-3D | Advanced 2D and 3D Laser-Produced and Electrically Driven Plasma Processes Modeling for EUV Lithography and other advanced plasma applicactions | Ahmed Hassanein | Argonne National Labor | 300,000 | 0 | Physics | INCITE Prep |
| Reflection_MHDTurl | Reflection-Driven Magnetohydrodynamic Turbulence | Jean C Perez | University of New Hamp | 300,000 | 227,391 | Physics | INCITE Prep |
| CRYSTALLISATION | Massively parallel molecular simulation studies of nano-scale crystal formation | Niall English | University College Dubli | 1,000,000 | 0 | Physics | INCITE Prep |
| qmc | QMC development on BG | Jeongnim Kim | University of Illinois at U | 995,000 | 79,450 | Physics | INCITE Prep |
| B_I_MHD_Turb | Balanced and Imbalanced MHD Turbulence | Andrey Beresnyak | Los Alamos National Lab | 500,000 | 473,550 | Physics | Science |
| ComPASSatBGP | Community Petascale Project for Accelerator Science and Simulation Codes on BG/P | Panagiotis Spentzouris | Fermilab | 1,000,000 | 257,691 | Physics | Science |
| Cosmology | Computational Cosmology | Stephen Kuhlmann | Argonne National Labor | 2,400,000 | 490,515 | Physics | Science |
| CTJSNIA | Research Into the Systematics of Type Ia Supernovae | Alan Calder | Stony Brook University | 120,000 | 0 | Physics | Science |
| LSBDO | Large scale beam dynamics optimization for more efficient operation of large user facilities | P. Ostroumov | Argonne National Labor | 5,000,000 | 13,012 | Physics | Science |
| NekCEM | Electromagnetics | Misun Min | Argonne National Labor | 1,800,000 | 281,293 | Physics | Science |
| | | | | 479,259,201 | 153,291,394 | | |

## Category Definitions

| | |
|---|---|
| **INCITE Prep** | Preparation for an INCITE proposal. These can be short term prep (needing to run some scaling tests) to longer term development and testing. |
| **Code Development /Support** | are supporting tools and abilities on ALCF resources and development needed for competitions such as the Gorodon Bell. |
| **Science** | encouraged to work toward INCITE or ALCC but these projects might need a much longer lead time to get to that level. |
| **Internal/Support** | These are ALCF projects used to support the machine or users. |
| **ESP** | Early Science Program time (applications preparation for transition to BlueGene/Q). |

*This page intentionally left blank.*

# Appendix D – Strategic Results Slides

4/10/12

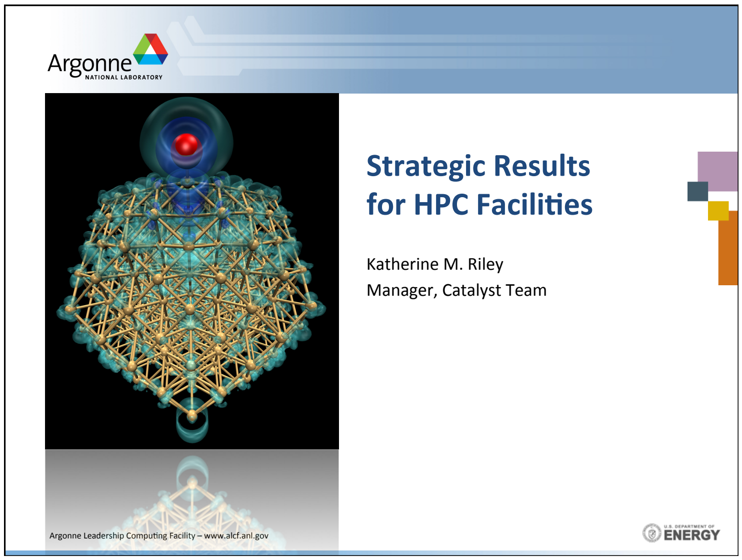


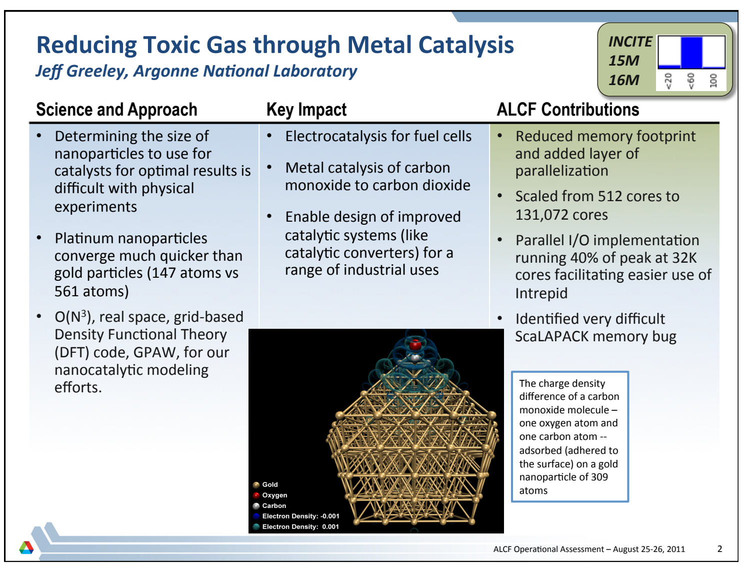


ALCF OAR 2011

1

## Boosting Fuel Economy through Cutting-Edge Computational Physics

INCITE
40M
41M

Robert Moser, University of Texas

**Science and Approach**

- Highest ever Reynolds number simulations for spatially evolving incompressible turbulent boundary layers.
- The team found that turbulent boundary layer fluctuations are higher than those in channel flows. A ***key*** question in turbulent literature.
- Direct numerical simulation (DNS) of high Reynolds number turbulent flows in a boundary layer
- Gathering the required turbulence statistics in analysis requires significant computation

**Key Impact**

- Boundary layers are central to the energy losses in transportation
- Understanding these flows can improve the design and efficiency of vehicles. These designs are impeded by the lack of accurate models of the turbulent phenomena

**ALCF Contributions**

- Initial port and performance optimizations for Intrepid which drove further optimizations by Moser's team.
- Testing threading performance for Mira

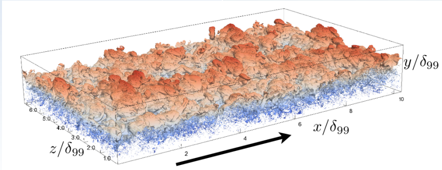


This shows the turbulent/non-turbulent interface at the boundary layer; red are turbulent zones and blue are non-turbulent zones. Rotational flow is a key difference between boundary layer flow and internal flows (in channels and pipes).



## Minimizing the Acoustic Signature of Jet Engines and Wind Turbines

INCITE
20M
18M

Umesh Paliath, G.E. Global Research, Niskayuna, USA

**Science and Approach**

- Carry out large-eddy simulations on jet nozzles to understand noise generation
- Completed a proof-of-concept problem that proved the LES approach captures the acoustics

**Key Impact**

- Design improved, lighter engines with more fuel savings and fewer $CO_2$ emissions
- Reduce ear damage for these working near engines
- Decrease noise impact and improve design of wind turbines

**ALCF Contributions**

- The ALCF ported and performed performance optimizations of the code CharLES.
- The team is tightly engaged getting this and a preferred software package ready for Mira.

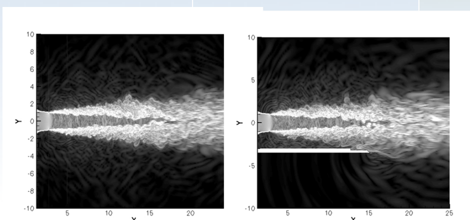


**Density gradient contours for simulation of conic nozzle with and without the presence of flat plane. The density gradient is a qualitative picture of noise.**

## Deflagration-to-Detonation Transition in Reactive Gases

Alexei Khoklov, University of Chicago

INCITE
18M
19M

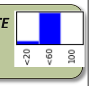


**Science and Approach**

- First principle, reactive flow direct numerical simulations
- 7 micron resolution shock tube in $CO_2$ with heat conduction and isothermal walls has excellent agreement with experiment
- More complicated physics for next suite of simulations

**Key Impact**

- Design safe systems for future hydrogen energy systems by gaining insight into the physical mechanisms of the burning and detonation of hydrogen-oxygen mixtures

**ALCF Contributions**

- Revealed a strong load imbalance and presented a recommended solution. Collaborating on implementation.
- Visualization team helped visualize data in VISIT
- To prepare for initial INCITE award, ALCF: Implemented OpenMP allowing use of all cores on the node giving 3x speedup and rewrote I/O.

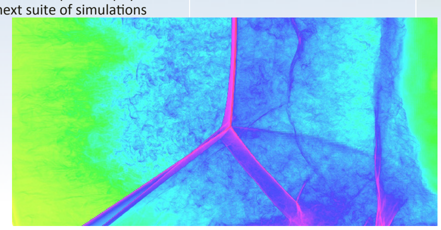

Temperature in a three-dimensional Navier-Stokes first-principles direct numerical simulation of a Mach=3 reflected shock bifurcation in a hydrogen-oxygen mixture in a square channel.

5

ALCF Operational Assessment – August 25-26, 2011

## Laser-Plasma Interactions for the National Ignition Facility

Denise Hinkel, LLNL

INCITE
50M
46M

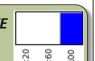


**Science and Approach**

- Using code pF3D simulating both radiation and hydrodynamics
- A large-scale simulation of laser-plasma interaction in the National Ignition Facility (NIF) looking at non-uniform cross-beam energy transfer
- Double quad-beam simulations have shown an increase of reflectivity out of the target
- Expanding to more beams is critical to improving target design

**Key Impact**

- Impacted design of target for future National Ignition Facility experiments.
- Fusion energy promises to bring sustainable, green energy production to the world

**ALCF Contributions**

- Debugging a few critical issues
- Coordinated large campaigns

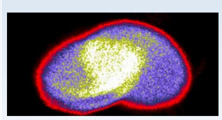

Power transferred from other quads of laser beams is distributed uniformly across the laser beams. Two quads overlap in the simulated region. This enhances reflectivity through a shared reflected light wave.

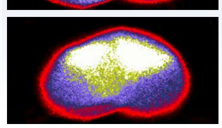

Power transferred from other quads provides a spatially non-uniform distribution of power across the beams. The bright "triangle" in the upper region of each laser quad drives high levels of reflectivity within each quad. The overlap of the two quads drives reflectivity through a shared reflected light wave.

6

ALCF Operational Assessment – August 25-26, 2011

## Materials Design from First Principles

Larry A. Curtiss, Argonne National Laboratory

ALCC
20M
21M

**Science and Approach**

Reduce the significant chlorinated waste from producing propylene

- Subnanometer Ag clusters may provide pathways to overcome this waste problem.
- Structure of key molecule optimized and study completed on part of the process

Study the transformation of biofuels with efficient catalysts

- Structure of $ZrO_2$ nanobowls in dry and hydroxylated alumina have been optimized
- In support of the Institute for Atomic-efficient Chemical Transformation (IACT), an Energy Frontiers Research Center (EFRC)

**Key Impact**

- Provide the fundamental understanding and predictions needed to design new materials for catalysis and energy storage

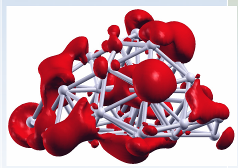

Thirty-three atom silver cluster that is being studied as a new catalyst for propylene epoxidation.

**ALCF Contributions**

- Significant time was spent mentoring 4 post-docs on using GPAW and how to use Intrepid
- Building off of significant effort on the performance and scalability of GPAW

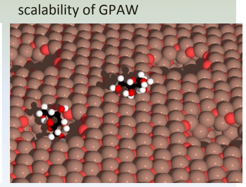

Nanobowl in ZrO2 surface (1nm diam.) that is being studied for biomass conversion to fuels

7

ALCF Operational Assessment – August 25-26, 2011

## Lattice QCD

Paul Mackenzie, Fermilab

INCITE
50M
181M

**Science and Approach**

- Provide crucial high precision lattice QCD calculations needed for new or in-progress experiments and analyze results from experiments
- Used in many high energy and nuclear physics research projects

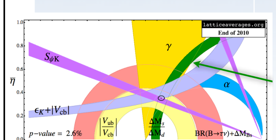

Laiho, Lunghi, and Van de Water

**Key Impact**

Lattice QCD has delivered essential results to experimental programs

- RHIC - Helped to firmly constrain heavy-ion collision models for the first time. First and second most cited papers in LQCD since 2006
- Fermilab – Calculations by USQCD members combined with experimental results have allowed many of the fundamental parameters of the Standard Model to be determined more accurately than ever before.

**ALCF Contributions**

Discovered a new, much more efficient, way to calculate the HISQ "fermion force".

- Reduces FLOPS 10x
- Improves parallel efficiency
- Represents 10-50% of total time

Similar improvements were made for the gauge force routines.

- ~10% of runtime and is now 2-3x faster in the MILC code

Changes are captured in the LQCD SciDAC libraries.

8

ALCF Operational Assessment – August 25-26, 2011

## Understand the Impact of Sickle Cell Anemia on Blood Flow

George E. Karniadakis, Brown University

INCITE
50M
44M

| Science and Approach | Key Impact | ALCF Contributions |
|---|---|---|
| • Mapping exactly how **red blood cells** move through the brain and study SS and CM diseases.<br>• Understand blood viscosity and blood flow resistance<br>• Using **computers**, **scientists** and **cardiologists** will conduct virtual experiments to study cerebral blood flow and improve **diagnosis** and **treatment**. | • **Sickle cell anemia** (a chronic inflammatory disease) affects **72,000** individuals in USA.<br>• Improve potential treatments | • Scheduled significant projects campaigns, job scheduling and interacting with the system<br>• ALCF staff worked another aspect of the project on a Gordon Bell submission. This specific subproject benefited from porting and basic performance improvements. |

*Biconcave and sickle shape of the red blood cells (RBC)*

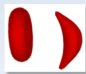

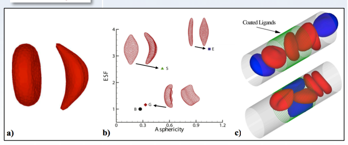


*Sickle shaped blood cells and how they deform in blood flow*

9

ALCF Operational Assessment – August 25-26, 2011

*This page intentionally left blank.*

# Appendix E – Report on INCITE Activities

In early 2011, DOE initiated a review of the INCITE program to assess the processes that the ALCF and OLCF use to solicit, review, recommend, and document proposal actions and monitor active project[s] and evaluate their INCITE portfolio. The six-member panel of national and international experts met in June with the INCITE manager and OLCF and ALCF senior management. There were no negative findings. The panel judged that the program has addressed the 2008 Committee of Visitors recommendations from the previous review of INCITE and had few additional suggestions. The INCITE manager and center directors were complimented for their effective management of the program.

A total of approximately 1.7 billion processor hours were allocated to 57 INCITE projects in CY 2011. (930 billion hours on OLCF's Cray XT Jaguar were awarded to 32 projects and 732 billion hours on ALCF's IBM Blue Gene/P were awarded to 30 projects; several projects received awards of time at both centers). The scientific peer review was carried out with nine panels of experts, with nearly 70 reviewers in total. INCITE is open to researchers from around the world and the panels reflect this: 15% of the reviewers were from outside of the United States.

The 2012 INCITE Call for Proposals (CFP) yielded a total of 119 submissions. These submissions underwent computational readiness and scientific review. The demand for time on the leadership systems continued to be high. In the 2012 CFP, INCITE received requests for 5 billion hours of time, nearly three times greater than the combined OLCF and ALCF hours available for allocation.

Peer review represents a best practice for the assessment of programmatic efficacy as well as for the identification of high-impact research activities. For INCITE, not only are the proposals peer reviewed, we also ask the scientific panels to provide INCITE management with feedback about the quality of the submittals and the operation of the program. To gauge the quality of the proposals received, the panel reviewers are asked to rate their response to the statement "INCITE proposals discussed in the panel represent some of the most cutting-edge computational work in the field." On a scale from 1 (strongly disagree) to 5 (strongly agree), the reviewers in the 2010 and 2011 CFP strongly agreed, with average ratings of 4.51 and 4.52, respectively. 94% of the attending panel reviewers last year responded. See Table E-1 below for the survey questions and average responses.

| | 2010 INCITE CFP Avg Score | 2011 INCITE CFP Avg Score |
|---|---|---|
| INCITE proposals discussed in the panel represent some of the most cutting-edge computational work in the field. | 4.51 | 4.52 |
| The proposals were comprehensive and of appropriate length given the award amount requested. | 3.89 | 4.15 |
| Please rate your overall satisfaction with the 2010 [2011] INCITE Science Panel review process (where 1 is "very dissatisfied" and 5 is "very satisfied.") | 4.67 | 4.79 |

**Table E-1: Results of survey of INCITE scientific peer-reviewers at the annual panel review meeting. Average scores are based on ratings between 1 ("strongly disagree") to 5 ("strongly agree").**

Refinements to the program policies and procedures were introduced in April 2010 for the 2011 CFP; see the 2010 OAR for details. These changes resulted in an improvement in the panel rating for the second survey statement, "The proposals were comprehensive…" with an increase in average rating from 3.89 to 4.15. Additional changes were introduced in April 2011 for the 2012 CFP. For example, the program revised the renewal proposal form (the new submittal form was previously redone) and emphasized the authors' achievements to date. After the 2012 CFP ended, the authors were invited to respond to a short survey asking for input about the proposal form and templates. Nearly 20% of the authors responded and expressed satisfaction with the INCITE proposal form. Several suggested modifications that will be incorporated into the 2013 INCITE CFP. Some comments are provided below.

*"Templates were great, wish other programs such as Teragrid, GENCI or PRACE provided these."*

*"I really think the increased emphasis on results for renewals is a good change. Previous years it seemed like the important thing was how many jobs were run and at what size for each objective, and not so much what you get out of the simulations. Since obtaining science results is the ultimate objective, this change is appropriate, and prevents users spending time collecting statistics that are not particularly enlightening themselves when it comes to science results."*

Authors also provided these suggestions for future consideration:

*"I would like to see in the proposal the section devoted to a position of the proposed project as compared with the existing 'state of the art' in the field of the proposal."*

*"I had trouble figuring out how the best way to report some of our Computing Resource Allocations. They did not follow a fiscal year pattern and the webpage only allowed one to enter fiscal years. Maybe having the option to give start and end date would help."*

**Argonne Leadership Computing Facility**
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

www.anl.gov