High Performance Computing Facility

Operational Assessment Report

for the
Argonne Leadership Computing Facility (ALCF)

August 31, 2009

Pete Beckman, ALCF Director

Executive Summary

The High Performance Computing Facility Operational Assessment (HPCOA) is an Office of Science (SC) programmatic management tool for evaluating the High Performance Computing Facilities' (HPC) plans for providing high-performance computing and network resources as well as support to the scientific user base. Relevant information from the HPCOA is used to respond to the succeeding year's annual Operational Analysis data call for major Information Technology operations from the Office of Management and Budget (OMB) and the Office of the Chief Information Officer (OCIO).

The guidance for the FY '09 Operational Assessment indicated that facilities would provide:

- Responses to recommendations from the 2008 OA review,
- Performance data against the previous year's baseline plan,
- Performance results and projections for the next year.

There were eight recommendations from the previous year's Operational Assessment Review (OAR). The next section addresses those recommendations.

For performance against our baseline plan, the guidance directed us to respond to seven different metrics:

1. Customer Results
2. Business Results
3. Strategic Results
4. Financial Performance
5. Innovation
6. Risk Management
7. Cyber Security

Each of these metrics is addressed, in order, following the recommendations. Finally, we report on our results and projections for FY 2010.

2008 Recommendations and Their Resolution

1. Hire more "Catalysts," since the current ones are already overburdened.
    a. In the period 1 August 2008 through 31 July 2009, we have hired 2 additional catalysts and have 3 strong candidates we believe will be on board within a month.

2. Make sure that the survey is not too long.
    a. We worked with the Argonne Decision Information Systems Division to develop the questionnaire. It is 12 questions and should take an estimated 10 minutes to complete.

3. Consider tracking the most recurring basic questions from users and post solutions on web-based documentation.
    a. We have built a wiki (https://wiki.alcf.anl.gov) to provide user documentation. The support staff contributes to the wiki, which includes a FAQ section that addresses general questions, compiling and linking, queuing, debugging, running, performance, applications, presentations, and documentation. It also contains other sections on hardware overviews, getting started, data transfer, and our outreach activities.

4. Look at job termination information.
    a. We track and analyze job termination information. A discussion of the results and analysis is provided in the business results section titled "Analysis of Job Termination Information."

5. Continue showing how the Argonne Leadership Computing Facility (ALCF) has enabled the science (beyond just providing cycles).
    a. The section on strategic results is directly modeled after last year's presentation and demonstrates significant science accomplishments and how the ALCF has helped enable those accomplishments.

6. Develop metrics to measure the overall accomplishment of goals 3.1 and 3.2.
    a. Through the INCITE program, the Leadership Computing Facilities provide resources for high-impact research. We report on science impact through publications, presentations, and posters from ALCF projects. The majority of the OAR metrics report on the delivery of the scientific facilities. We report on students and post-docs trained using the facility and ALCF activity in workshops and classes.

7. Carefully consider the impact on the science community of the summer 2009 move and ways to mitigate the impact of that move.

   a. A cost-benefit analysis was conducted; as a result, the decision was made to not move Intrepid to the new building.

8. Define the cost (and expectation) model, working with the program office, for the Catalyst team and determine how to measure the cost effectiveness of adding additional staff if more INCITE projects, or large numbers of potentially impactful discretionary projects, are conducted at the ALCF.
   a. An approximate 2-3 INCITE projects per Catalyst is the target ratio. At this ratio, Catalysts are able to provide proactive support, study the details of the application and its performance on ALCF resources, and help guide new development and needs for the projects. When this ratio is exceeded, the quality of service the LCF provides degrades. Projects need to be prioritized, and work needs to be queued based on these priorities.

Performance Data against the Previous Year's Baseline Plan

Baseline Area 1: Customer Results

Metric: *ALCF problem resolution will be measured on problem reports sent by the users to the LCF problem tracking system. The percentage addressed within three working days will reach 66% in FY09, 73% in FY10 and 80% in FY11. ALCF will track its workshops, tutorials, monthly user teleconferences and application support provided to users and will provide quarterly reports to DOE.*

Customer results cover measures of customer satisfaction with the services provided by the ALCF, the resolution of problems encountered by the users, and overall support to the user base.

Customer Satisfaction

ALCF conducts user surveys throughout the year. An annual survey is used to gauge overall feedback from INCITE users regarding their experience in working with the ALCF. Additional surveys are used as part of ALCF workshops to assess the value delivered by the workshops and to identify areas for improvement. The annual survey instrument was initially designed and reviewed with the help of a qualified external reviewer for coverage, clarity, and applicability to the user base. As part of the annual survey, users are asked to rate topical areas (or questions) on the survey using a Likert scale and are given the opportunity to respond to open-ended questions in a more qualitative manner. Specific metrics addressed include:

- The ALCF's overall score on the user survey falls between above average and excellent
    - Based on the 2008 user survey, ALCF received an overall user satisfaction rating of 4.6 out of a 5.0 scale with 5.0 being excellent or the top score.
- The average score on the user survey is above average or better for each category of questions.
    - The average scores for the various questions ranged from 4.45 to 5.0.
- The questions should also capture the satisfaction of new users added during the year.
    - This was our first year of full production and our first user survey; so all users were new users. In the future, we will differentiate between new and returning users.
- The survey will address all topical areas pertinent to users.
    - The survey included questions regarding workshops, user engagement model, experience with support, access to resources, etc.
- Problem Resolution
    - In the 12-month period through July 2009, the number of tickets opened was 18,558.

- The number of tickets addressed in less than three days was 80% of the number of tickets opened. Last year's numbers indicated 78% of the total tickets were resolved in less than three days.

Workshops and User Calls

ALCF conducts monthly user calls with its INCITE users. These calls are scheduled on the last Thursday of each month. The calls typically contain a short presentation, or presentations, on a topic of interest to the users, as well as allot time for user questions. Support staff, including Catalysts, Performance Engineers, User Support Technicians, System Administrators, and Storage Engineers, are available on the call to answer the questions.

The ALCF conducted 5 workshops in the second half of 2008 and the first part of 2009. The first two workshops of each calendar year cover introducing users to the system, software, tools, and ALCF staff. Some workshops focused on having advanced users scale their applications more efficiently on the entire BG/P system. All workshops included presentations and time for hands-on work. The list of workshops and their dates are listed below:

2008

- October 21-22st HPCT Workshop
- October 21st Workshop for SDSU

2009

- February 10-11th INCITE Getting Started Workshop
- March 10-11th Getting Started for Director's Discretionary Projects and Blue Gene Consortium Workshop
- May 27-29th Leap to Petascale Workshop

Baseline Area 2: Business Results

Metric: *Systems in operation more than one year since general availability or the last upgrade will meet 85% scheduled availability and 80% overall availability. 250 million CPU-hours will be delivered to jobs 8,192 nodes or larger.*

Availability

As can be seen in the figure below, the ALCF scheduled and overall availability over the past year has been well above the 75% and 70% respectively required for a first year facility. Overall, for the year, we averaged 98.1% scheduled availability and 93.3% overall availability. We averaged 1.9% unscheduled downtime. There are 5 significant contributors to loss of availability called out in the figure, and they are described below. For clarity, we note that a large spike downward (white) indicates significant unplanned downtime A significant downward spike (blue) indicates significant planned downtime (extended maintenance).

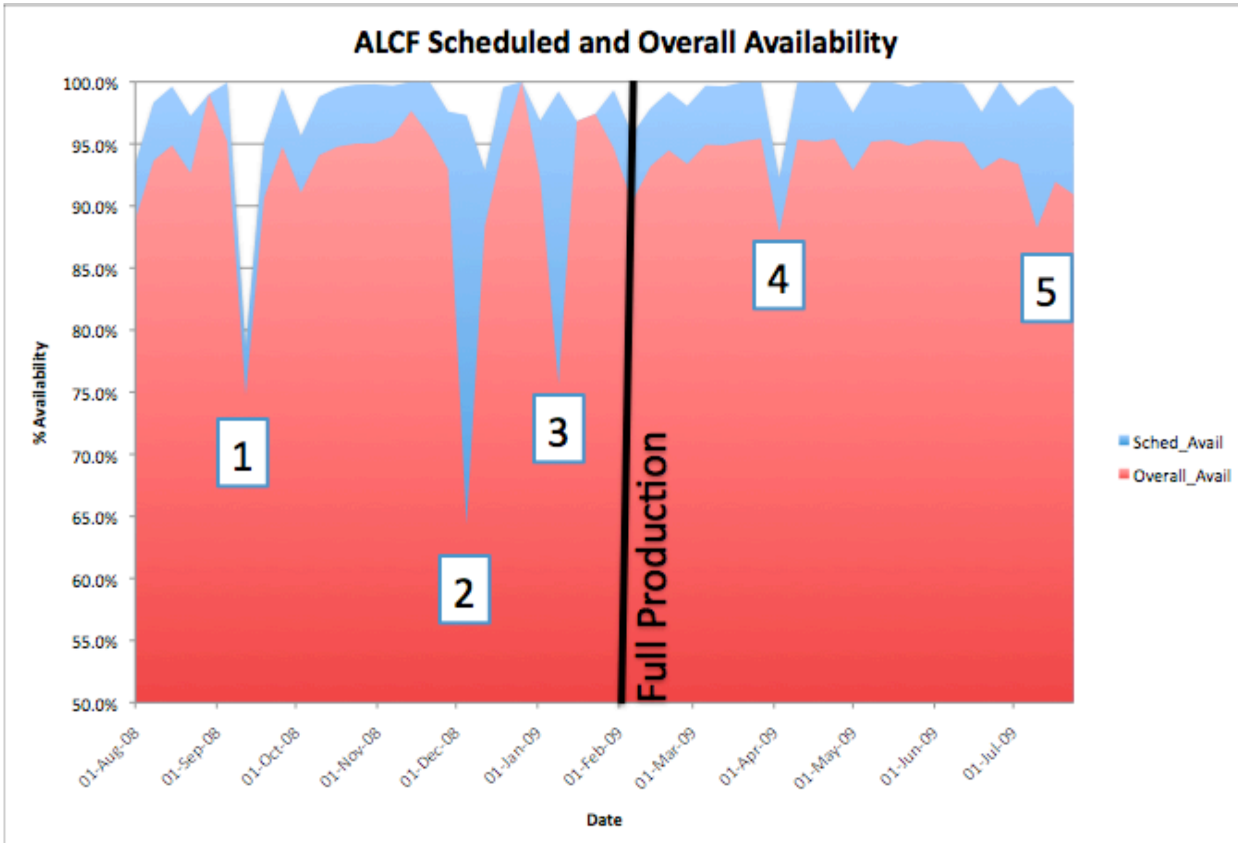

Faulty Bulk Power Modules (BPMs) on the Blue Gene (Items 1 and 2): The significant unscheduled downtime in September was due to an arc flash that took the whole machine down and kept rack R05 down until October 20th. It also accounted for the substantial extra scheduled downtime in December, when IBM replaced all the BPMs free of charge with a re-designed unit. Since then, we have not had an arc flash, and BPM failures have been well within normal limits. We consider this issue closed.

In December, we also had a chiller plant failure that forced an emergency shutdown and the associated problems of a cold restart. We also updated the firmware on nearly 1,000 optical transceivers used in the network switch fabric.

Preparation for Production and General Parallel File System (GPFS)/Myricom Issues (Item 3): We took extended downtime prior to going full production. During this time, we performed many systems tests to tune performance and improve configuration. We also tracked down a difficult-to-find hardware error caused by a bent pin on a high-speed network link.

Other significant work in 2009 included replacing faulty optical transceivers in the Myricom switch network. The ongoing replacement of faulty optical transceivers has been our most significant maintenance issue this year. When an optical transceiver fails, data from the Blue Gene to the storage servers is interrupted. The most common result is a file system-related error. While some transceiver failures over time are expected, the rate of failure from one particular manufacturer, Zarlink, was significantly higher than transceivers from a different manufacturer. The manufacturer has since tracked down the problem to several bad wafers during the manufacturing process. We have replaced approximately 25% of the Zarlink transceivers and are tracking the failure rates over time to predict the likelihood of future port failures from this defect.

Rapid, Continuous Rebooting of Myricom Ports (Item 4): Due to firmware bugs in the Myricom switch network, some ports suffered "reboot storms," which impacted the availability of the file system. A firmware patch corrected this, and we consider this issue closed.

Short-Term Reduction of Overall Availability (Item 5): We planned extended downtime in preparation for a system software upgrade. Since the ALCF has one of the largest BG/P systems in the world, system software updates from IBM cannot be fully verified until they are tested on our platform. We provided IBM with the opportunity to test and debug the next core BG/P software release (driver V1R4). We have successfully run the updated software stack on the full machine several times. We are pleased with the results and believe it will reduce unscheduled downtime when the software is moved into production.

Scheduled Maintenance: Though not called out specifically, the planned Monday maintenance schedule, which usually lasts approximately 8 hours, reduces the availability of Intrepid by 4.7% overall.

Utilization

A brief summary of the status of the ALCF is in order to make this section more clear. We began installation of our Blue Gene/P in late 2007. We follow the basic sequence of order, install, acceptance, early science / transition to operations, and then full operations. We brought intrepid to production in two stages. We brought the first 8 racks up first. These 8 racks went production March 31st, 2008. We then brought the other 32 racks to full production on Feb 2nd of this year. Thus on the graph below, the INCITE hours delivered prior to the point marked full production were delivered on the 8-rack 100TF system, with the remainder of the time being a combination of early science and discretionary time. After all 40 racks reached production (indicated by the black line on the figure below), all early science projects were disabled and you see the steep increase in INCITE usage, with the remainder being discretionary time.

For the reporting period, we averaged 65.3% overall utilization. Our commitment to INCITE was to provide 400 million core-hours. We are currently ahead of schedule, having delivered approximately 253 million core-hours to INCITE since Jan 15th. Extrapolating that rate suggests we will deliver 466 million core-hours to INCITE by the end of the INCITE year.

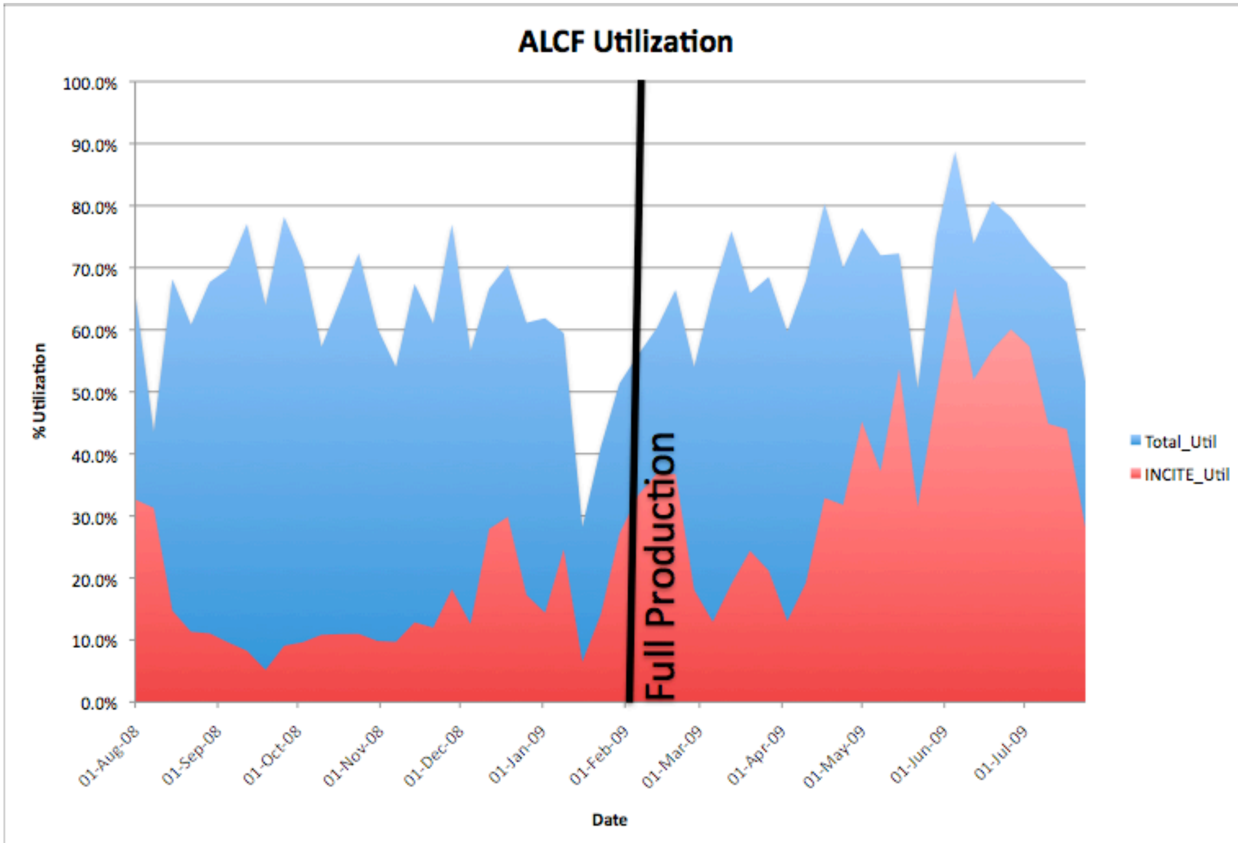

The significant drop in utilization in January 2009 was the result of maintenance, testing, and tuning work to prepare for production.

Improving Utilization

To manage jobs on Intrepid, we use the Cobalt job scheduler, developed here at Argonne. When Intrepid was in pre-production and early production, we made extensive use of job reservations to provide users easy access to very large computing partitions. While reservations are very convenient for users needing the largest partitions, they incur a significant staff overhead and adversely impact utilization. As users became more comfortable with Intrepid and demand increased, we worked with the Cobalt developers to improve the scheduler's algorithms and increase utilization. Cobalt now calculates a "utility score" for each job awaiting execution. This score has three major components:

1. $\left(\frac{QueueWaitTime}{\mathrm{Re}questedWallClockTime}\right)^3$ This means that once a job has waited longer than its requested wall clock time, its priority begins to grow rapidly.
2. Job size: Larger jobs are given a higher utility score.
3. Project priority: ALCF staff can weight projects to help improve turnaround on projects working against short-term deadlines.

Today, because of these improvements, reservations are not often required. Users are able to run jobs using 80% of the machine without any special arrangements. In addition, Cobalt's ability to backfill jobs has been improved by using more sophisticated algorithms. These changes have improved utilization.

In response to user requests, we nearly completed adding support for pre-emption to Cobalt. Scientists who have used their entire CPU allocation would like to still submit jobs if nodes are idle. These "scavenger jobs" would then be pre-empted when a high-priority job arrives in the queue. Sites that have implemented support for scavenger jobs report improvements in utilization.

Capability Usage

The ALCF is required to deliver 250 million core hours to jobs 8K nodes / 32K cores or larger during the year (capability jobs). We delivered over 336 million core-hours to jobs 8k nodes / 32K cores or larger, thus achieving our metric. This equates to 38% of the total cycles being considered capability cycles. The figure below shows total and capability core-hours delivered over time.

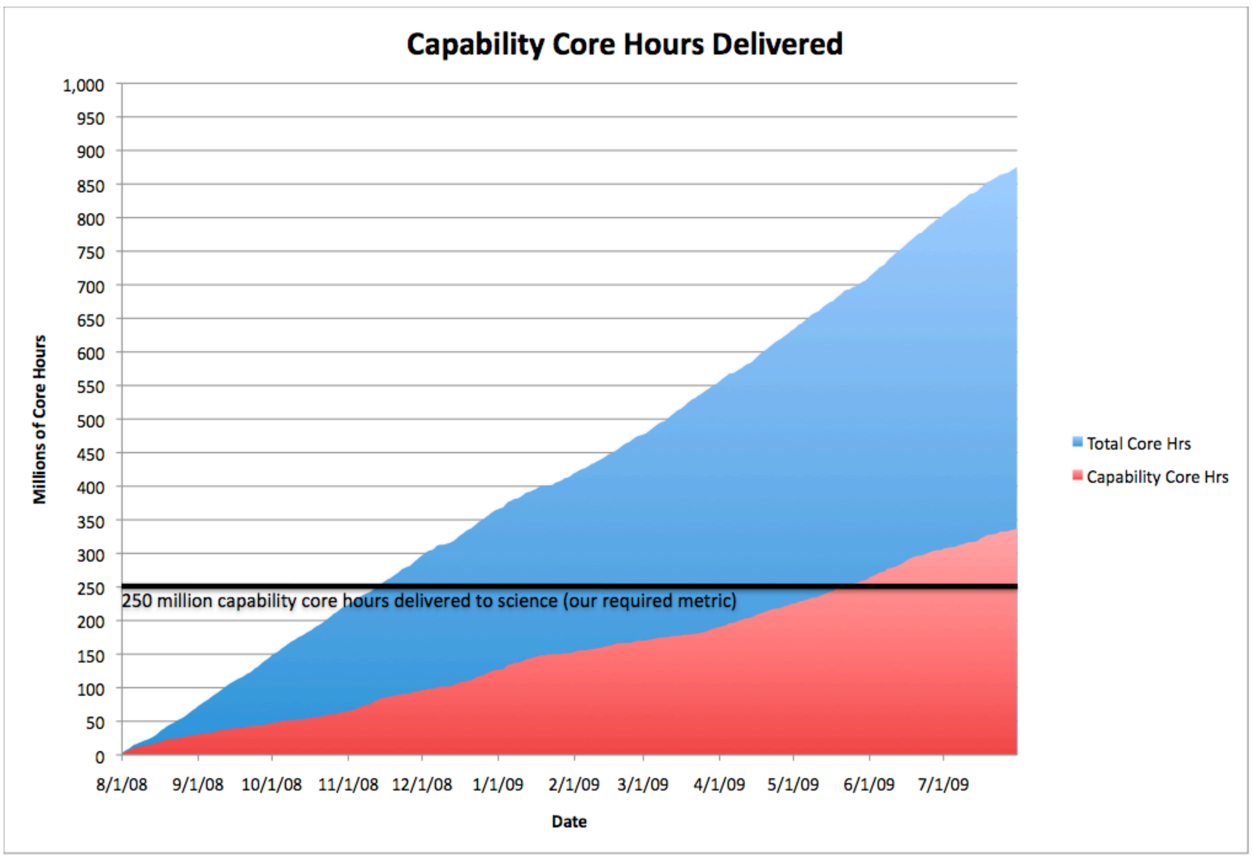

The figure below shows job distribution. Note that larger/capability jobs are at the bottom, and the black line indicates the demarcation between capability and non-capability jobs.

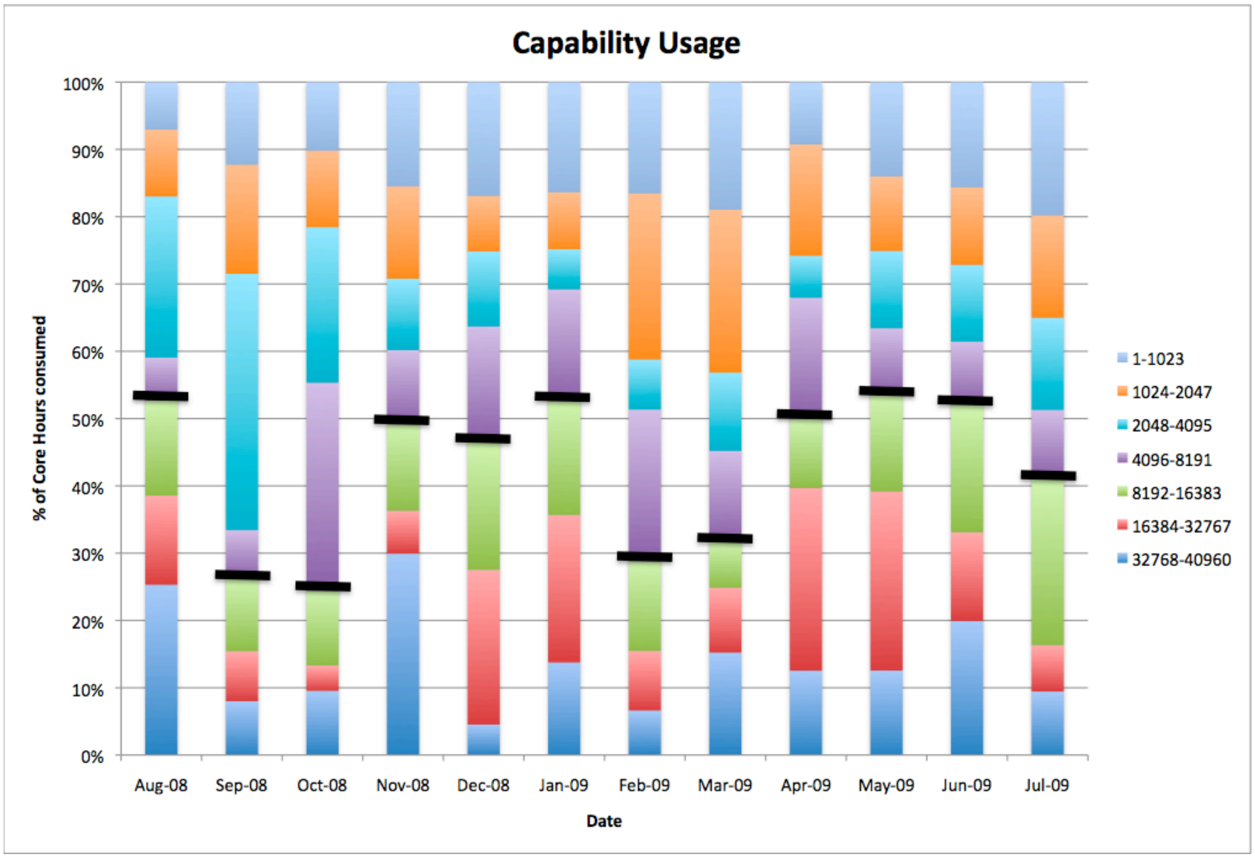

Job Throughput

The expansion factor[1] by job size is shown in the figure below. As noted in the discussion about the Cobalt job scheduler, the utility score begins to rapidly increase once the job has waited longer than its requested wall-clock time (an expansion factor of 2). This tends to automatically keep the expansion factor under control. Two issues negatively impact our expansion factor. First, this data is not filtered to remove jobs that failed prior to successful completion. A large 32-rack job that fails almost immediately will substantially skew the expansion factor. A similar effect occurs with scaling runs. These tend to be large jobs, that naturally have longer wait times, that may run to successful completion in 15-30 minutes, again, causing a skew in the results. Also note that jobs less than 512 nodes run from our prod-devel (production development) queue. These jobs generally are intentionally short, and by policy, constrained to no more than 1 hour.

---

[1] The expansion factor is defined to be $QueueWaitTime + JobRunTime / JobRunTime$

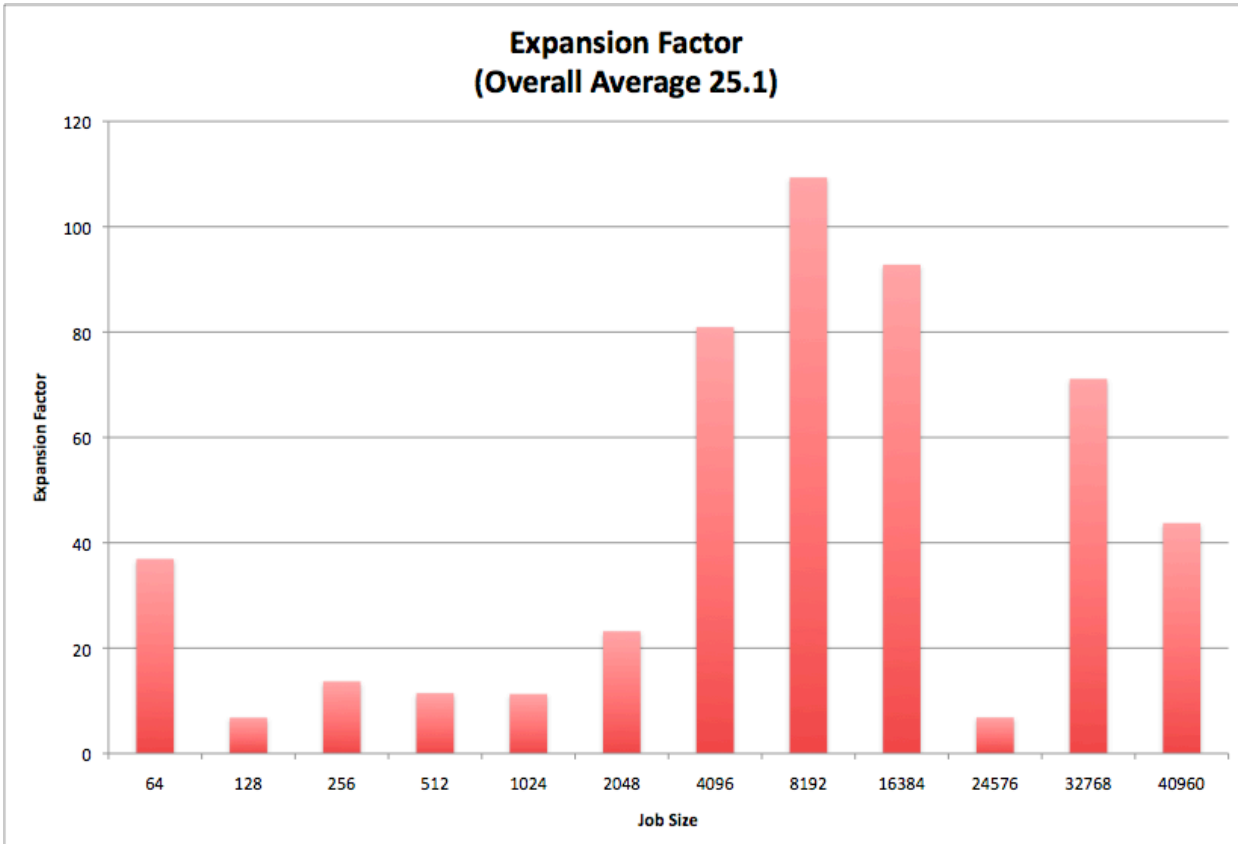

We do not typically track the expansion factor. However, we do hold monthly meetings to discuss scheduler issues. During these meetings we review the unit-less wait times (queue wait time/requested wall clock time) by queue, INCITE vs. non-INCITE, job size, etc. We also elicit feedback from our users and have made modifications to scheduler policy and the cost functions based on this feedback.

Analysis of Job Termination Information

From August 1, 2008 through July 31, 2009, there were 128,167 jobs submitted to Intrepid. Of these, 81,607 of them, or about 63.37%, ran to successful completion, leaving 46,560 (36.44%) of them that did not complete successfully.

Of the jobs that failed to complete, 41,731 (89.6%) of them were due to user error, and 4,829 (10.4%) due to system failures. The figure below shows a Pareto chart of the failure causes. Over 60% of the user error failures were a combination of jobs running beyond the requested time limit, users deleting the job from the queue before it ran, or users killing the job during the run.

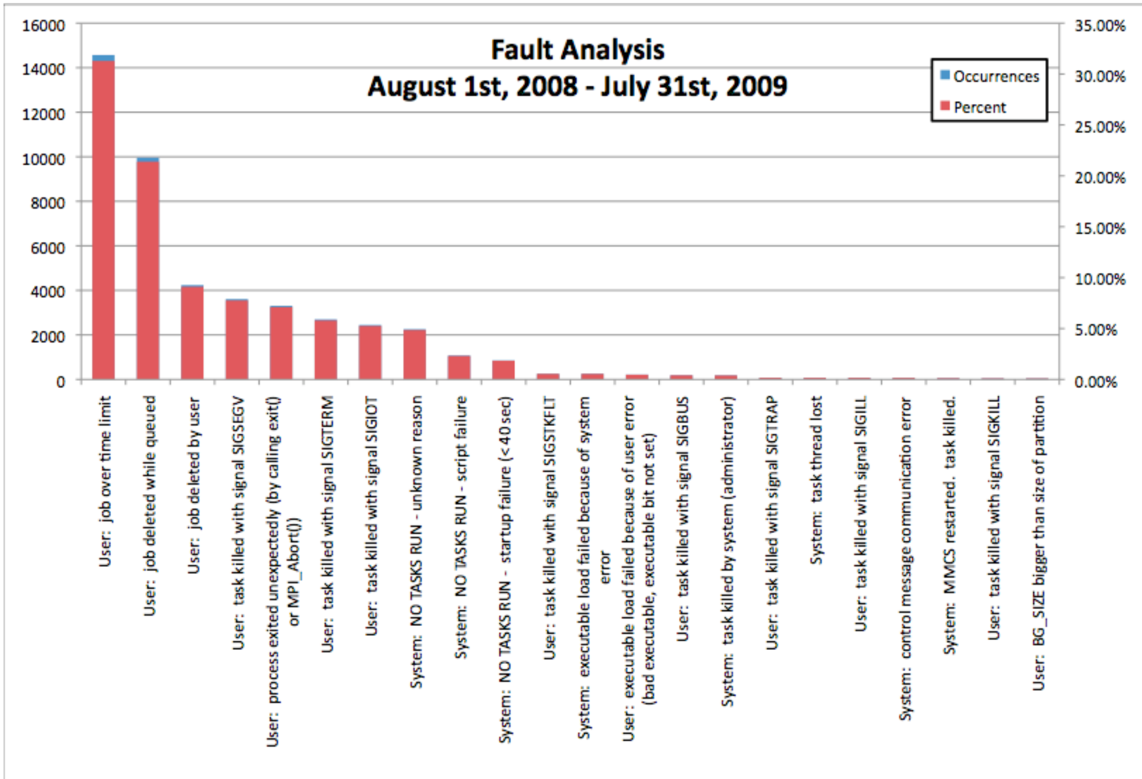

Of the system failures, 90% of them were boot failures (error messages with NO TASK RUN in the text). We are working on procedures that will allow us to correlate external issues with this data, but they are not fully in place at this time, so an exact analysis of what caused the boot failures is not possible. However, we do know that 23% of the boot failures come from one of two causes:

1. GPFS file system problems caused by failing Myricom optical transceivers, discussed above in the availability section.

2. NFS mount failures of the /bgsys file system, which is required for booting. We recently completed work that we believe will greatly reduce this issue. This included updating a 10GigE driver, reducing the NFS mounting concurrency, and moving the NFS server off the service node and onto a dedicated NFS server. We do not yet have sufficient data to consider this closed, but early results look promising.

Baseline Area 3: Strategic Results

Metric: *The LCF will track the science output and accomplishments for each project including milestone reports, presentations, publications, journal covers, and awards and will provide quarterly reports on the project results. The LCF will also track technology accomplishments, such as development of reusable code that results in a new tool for its discipline and new algorithm design ideas or programming methodologies.*

The ALCF tracks the scientific impact of research performed on ALCF resources. There have been over 23 publications and 20 presentations of work related to the Leadership Computing Facility, including an outstanding paper award at EuroPar for work developing a load-balancing library for Steve Pieper's INCITE project. The Lattice QCD project has generated public datasets used internationally. ALCF projects had significant presence at the 2009 SciDAC meeting, including two of the ten best scientific visualizations. At current count, discretionary projects have over 15 publications. One project from Indiana University led by Guoping Zhang is studying ultrafast demagnetization of high-temperature superconductors and has generated publications in *Nature Physics, PRB, PRL, JAP*, and received an invitation from the *New Journal of Physics*.

Continuing its success in 2008, the Director's Discretionary time played a direct role in 41 new INCITE proposals in 2010 through both outreach/education and more independent projects. Research highlights include Priya Vashista's project (see slides) and work with Paul Fischer on auto-ignition in turbulent flows:

http://www.sciencedaily.com/releases/2009/07/090725203431.htm.

The ALCF conducted three workshops, "INCITE Getting Started/Introduction to Blue Gene/P," "Leap to Petascale," and "INCITE Proposal Writing Workshop." At SC08, we ran a Birds of Feather (BOF) session to discuss the challenges and successes using Blue Gene/P, and we participated in the Blue Gene Consortium. Additionally, staff have given talks at the following meetings: SciDAC's CScADS summer workshops, SPEC HPG, NERSC HDF5 Workshop, Workshop on Climate Models on 100K-1M Cores (ANL, Jan 2009), and the 14th Annual CCSM Workshop.

2008-2009 INCITE project teams made abundant progress on code development. Christopher Mundy's project had a major breakthrough in the last year with the implementation in CP2K of a robust, accurate, and efficient implementation of Hartree-Fock exchange. HFX is an essential ingredient for hybrid functionals that provide accuracy for problems where local (GGA) DFT fails. Igor Tsigelny's team created an application, MPMD, which has four different physics components working together to further the scale of their research into Parkinson's disease. William Tang's project added a radial decomposition to GTC-P, dramatically improving performance and scalability. David Baker is developing a new ability in Rosetta to make structure projections based on sparse experimental data. In consultation with the ALCF, Susan Kuriens' project modified their application, DNS3D, to use p3dfft, resulting in significantly larger scaling. Finally, the ALCF played a key role in the development of MILC, FLASH and CCSM.

Both the Lattice QCD projects and the FLASH projects have generated public datasets for the larger computational community. The lattice configurations from Lattice QCD are used internationally.

2009 INCITE Allocations at ALCF

INCITE
Advancing America's Science and Industrial Competitiveness

28 projects

400 M CPU Hours

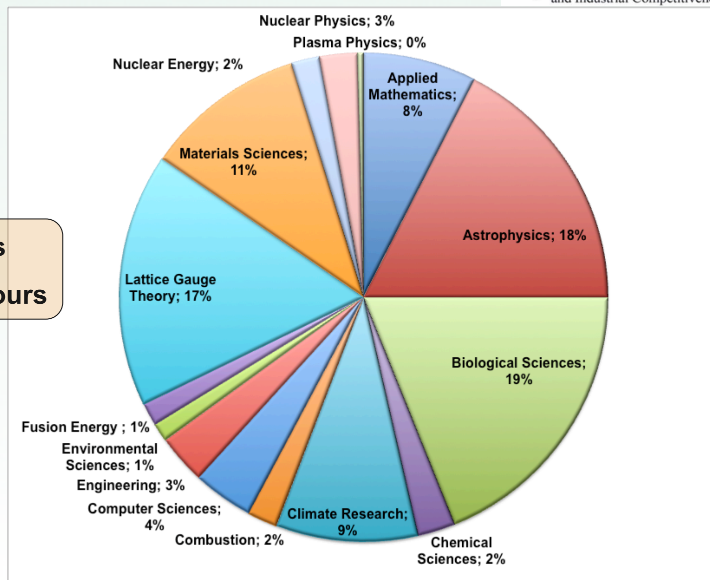

Argonne National Laboratory | ALCF Stock Science Slides June, 10, 2009 | 1

2008-2009 ALCF Discretionary Allocations

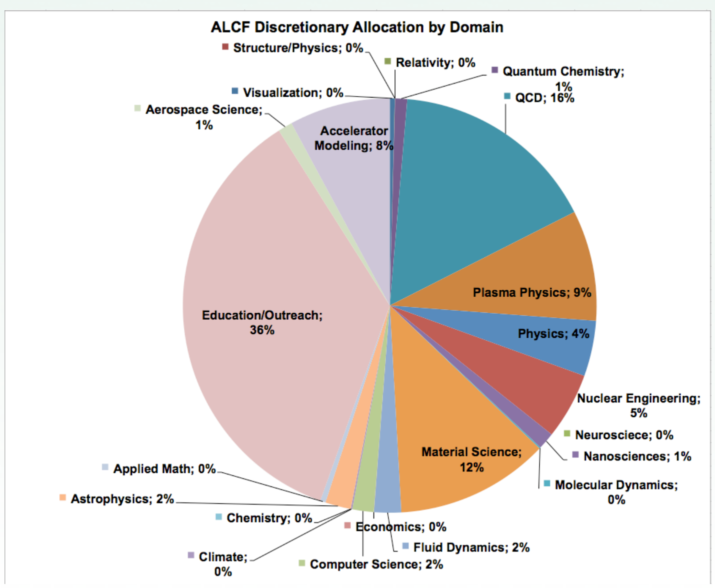

Argonne National Laboratory | ALCF Stock Science Slides June, 10, 2009 | 2

Better Reactors Faster at Scale

Paul Fischer ANL

Pressure distribution of coolant flow

Breakthrough computation of thermal-hydraulics with Nek5000

2.95M spectral elements

1 billion grid points

- Full physical configuration of 217 wire-wrapped fuel pins
  - *Production runs on 32K cores*
  - *80% parallel efficiency on 124K cores (strong scaling)*
  - *Culmination of 3 year INCITE with results from lower pin counts*
- Historical Build
  - 7-pins: Strong correlation with LES and RANS providing path for saving compute resources in the future
  - 19-pins: Proof that time-saving boundary conditions can be used

- Innovations over 3 Years
  - Scalable spectral element multigrid solver for the pressure
  - 4th generation coarse-grid solver (algebraic multigrid)
  - Elimination of all arrays scaling with global element count
  - Communication algorithms to discover processor topology
  - Scalable grid partitioner
  - Parallel I/O rewrite (subcommunicator)
  - Parallel visualization

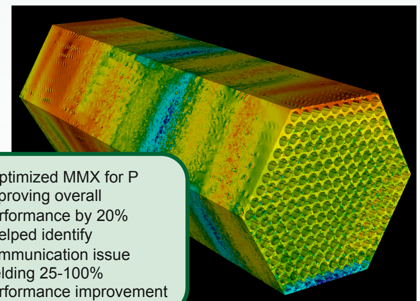

•Optimized MMX for P improving overall performance by 20%
•Helped identify communication issue yielding 25-100% performance improvement

CCSM Development on Blue Gene/P

Warren Washington NCAR

Science

- CCSM is a climate simulation code used by the DOE and NSF climate change experiments
- Ultra-high resolution atmosphere simulations on Intrepid
  - CAM-HOMME atmospheric component
  - 1/8th degree (12.km avg grid) coupled with land model at 1/4th degree and ocean/ice
  - Testing up to 56K cores, 0.5 simulated years per day with full I/O
- ½ degree finite-volume CAM with tropospheric chemistry and 399 tracers
- **New Coe Developments**
  - CAM performance doubled with threading and parallel decomposition of advection by species
  - CCSM4 nearly complete with the carbon cycle, long-term integrations and decadal forecasts
  - Parallel I/O using PIO and PNETCDF getting 700 MB/s
- 37 publications, presentations and posters

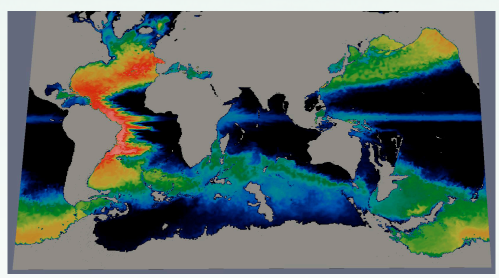

Figure. Value of an Impulse Boundary Propagator tracer 30 years after its introduction at the surface from the 1/10º POP simulation. The image plane descends from the surface in the north to the abyss in the south.

•Participated in the parallel I/O development
•Assisted with large-scale CAM-HOMME runs
•Ongoing work on OMP on BG/P

DD

Thousand Atom Nanostructures

Lin-Wang Wang
LBL

Science

- Design better materials for products including solar cells
- *Ab initio* electronic structure calculations
- Lin-Wang Wang, B. Lee, H. Shan, Z. Zhao, J. Meza, E. Strohmaier, D. Bailey, "Linear Scaling Divide-and-conquer Electronic Structure Calculations for Thousand Atom Nanostructures," SC08, to appear.

$Cd_{961}Se_{724}H_{948}$

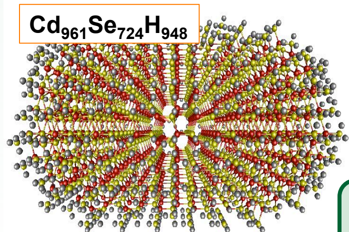

Dipole moment calculated on 2633 atom quantum rod

Methods and Challenges

- Novel divide & conquer approach to solve DFT but reducing $O(n^3)$ to $O(n)$
  - Many months to 30 hours
  - Direct DFT impractica
- Mapping critical
  - Linear scaling to 160K cores and a 10% improvement in per-core performance

ALCF consultation resulted in improving per-node performance over 200%

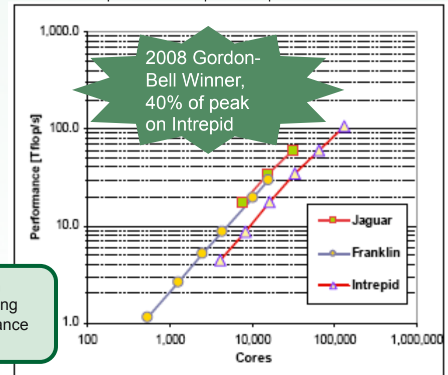

DD

Designing Better Materials for Nuclear Reactors

Priya Vashista
University of Southern California

- **47 million-atom ReaxFF MD simulations exploring fracture modes in materials**
  - 50M CPU-Hours
- **Key to design of next-generation nuclear reactor**
  - Revealed a missing link between sulfur-induced intergranular amorphization and embrittlement

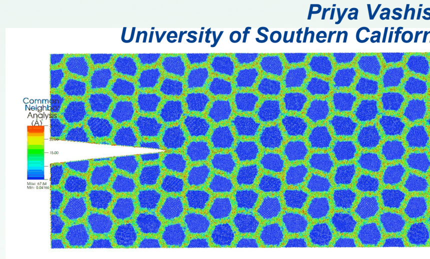

With Sulfur

Color: common-neighbor parameter that characterizes atomistic defects

".. the proposed upgrade of the ALCF resources to a 10-20 Petaflops Blue Gene Q system would enable our research group to make major breakthrough simulations on a number of applications of high relevance to DOE. We look forward to working with the ALCF on their new platform. "

(A) close-up fracture simulation in nanocrystalline Ni with amorphous sulfide GB layers

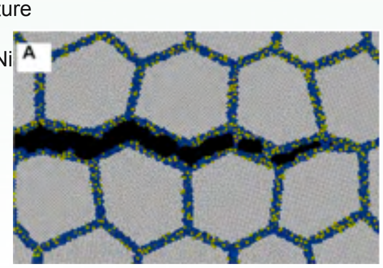

Identifying Potential Drug Targets

DD

Michael Wilde
ANL

Science

- Reduce dead ends in antibiotics and anticancer drugs with DOCK5 and DOCK6
  - 9 enzymatic proteins in core metabolism of bacteria and humans screened against 15,351 natural compounds and existing drugs
  - Study correlations and re-prioritize proteins for further study
- Able to complete 21.43 CPU-years of analysis in 2.01 wall-hours

•MCS Computer Science teams, using discretionary allocation, to facilitate science
•Participated in submitting 2 INCITE proposals

Methods and Challenges

- Port of framework, Falkon, to manage run
- Falkon requires non-standard BG/P kernel (ZeptoOS)
- Huge demand on I/O system as each core is controlling multiple files
- 118,000 cores were used running nearly one million tasks

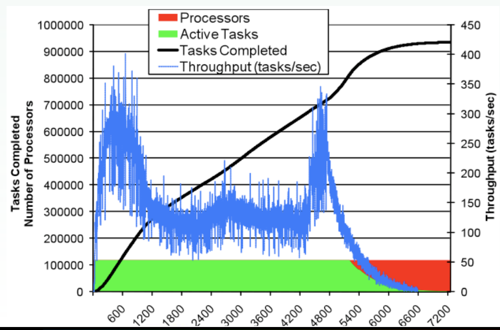

Gating Mechanism of Membrane Proteins

INCITE

Benoit Roux
ANL, University of Chicago

Science

- Understand how proteins work so we can alter them to change their function
- Validated the atomic models of Kv1.2 and first to calculate the gating charge in the two functional states

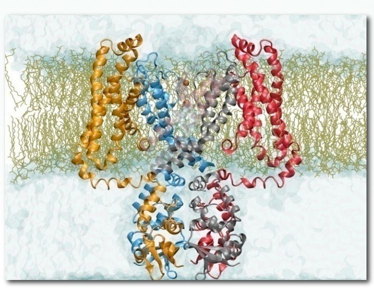

Methods and Challenges

- NAMD with periodicity and particle-mesh Ewald method
- Implemented new force fields

ALCF maintains NAMD and helped developed workflows for new force field work

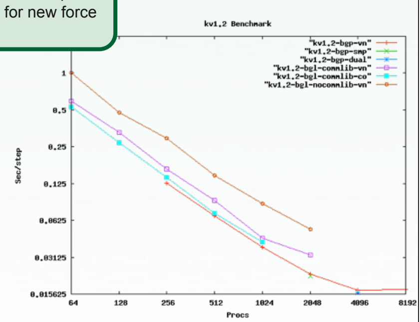

kv1.2 Benchmark (352K atoms)

15-20% gain over BG/L customized ("commlib") version.

INCITE

FLASH Project on Intrepid

Don Lamb, University of Chicago

Science

- Answered critical question on critical process in Type Ia supernovae
  - Shown it is *not* necessary to resolve the Gibson scale to adequately capture buoyancy-driven turbulent nuclear combustion
  - Moderate resolutions capture most of the burning rate
- Comparing methods of detonation for Type Ia

Methods and Challenges

- Operator split, multi-physics
- Block structured adaptive mesh
- Multi-pole and multi-grid gravity solves
- Load balancing with smaller memory footprint per core
- Single CPU performance optimizations

Collaborating with project to reconstruct their I/O

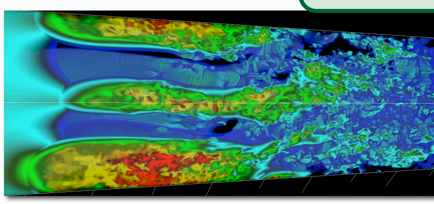

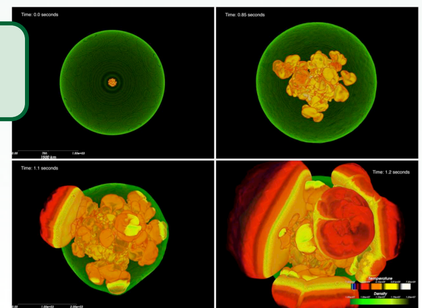

INCITE

Computational Nuclear Structure

David Dean
Oak Ridge National Laboratory
Steve Pieper
Argonne National Laboratory

- Green's Function Monte Carlo (GFMC)
  - *Ab initio* calculation of properties of light nuclei
  - Common benchmark for other methods
- Calculations of ^{12}C with complete Hamiltonian
  - INCITE time, 32K cores
  - Best converged *ab initio* calculations of ^{12}C ever
  - **Key calculation to launch into study of effects of different terms in the nuclear interaction and compute excited states**

- ADLB
  - SciDAC funded ANL-MCS developed multipurpose library for distributing load
  - Required for ^{12}C simulation
  - OpenMP and MPI programming model

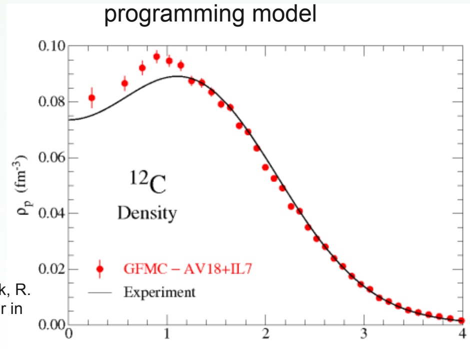

Outstanding paper award

P. Balaji, D. Buntinas, D. Goodell, W. Gropp, S. Kumar, E. Lusk, R. Thakur, and J.L. Traff, "MPI on a Million Processors", to appear in Proceedings of EuroPVM/MPI 2009, Springer-Verlag

INCITE

Lattice QCD

Bob Sugar and US-QCD

Science

- Addresses fundamental questions in high energy and nuclear physics, directly related to major experimental programs
- Determine parameters for Standard Model, including quark mass
- *MILC lattice generation for the lightest quark mass complete, moving onto HISQ quarks*

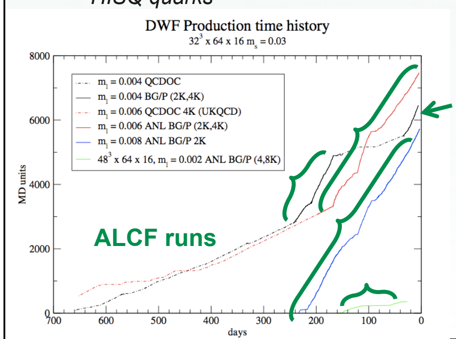

Methods and Challenges

- Rational Hybrid Monte Carlo
- For scalability and performance developed
  - QLA : 3x3 matrix linear algebra operations
  - QMP : low-level routines, partial MPI replacement
- Tuning complex algorithms after unexpected fall off in Monte Carlo acceptance rate

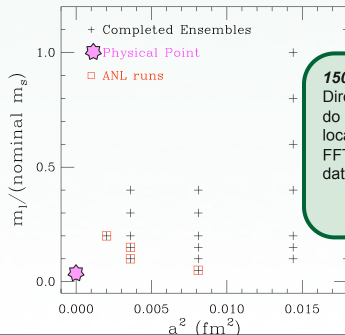

150x improvement
Directed MILC code to do a series of successive local 1d-FFTs using FFTW and shuffling the data as necessary.

INCITE

Faster Design of Better Jet Engines

Peter Bradly
Pratt&Whitney

Science

- Save cost and time by designing engines through simulation rather than building models
- Technologies from simulations now being applied to next generation high-efficiency low-emission engines
- A key enabler for the depth of understanding needed to needed to meet emissions goals

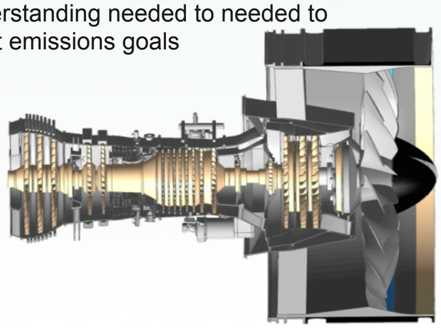

Challenges

- I/O algorithm redesign speeds up simulations by 3x

Insight into Parkinson's Disease

Igor Tsingelny
University of California, San Diego

Science

- Parkinson's Disease affects 5 million in US and Europe
- Increased aggregation of α-*synuclein* protein is thought to lead to harmful pore-like structures in human membranes
- Simulations show α-synuclein complexes, and β-synuclein prevents creation of propagating α-synuclein complexes
- Moving to membrane interactions which will simulate more than 1M atoms including the water box

Methods and Challenges

- Using NAMD and MAPAS on Blue Gene at ALCF and SDSC
- α-syn proteins are unstructured

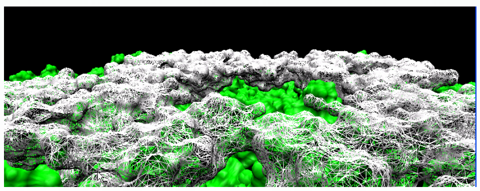

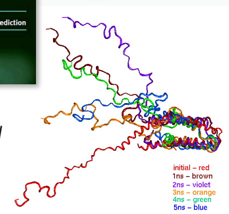

α-synuclein coformer at membrane (left), completed pentamer (above), evolution of α–synuclein over 5ns (right)

Computational Protein Structure Prediction and Protein Design

David Baker
University of Washington

- Computationally design protein-based inhibitors towards pathogens like H1N1
- Rapid turn around of huge campaigns on ALCF reinvented how the science is done and enables new research
- Rapidly determine an accurate, high-resolution structure of any protein sequence up to 150-200 residues
- Incorporating sparse experimental NMR data into Rosetta to allow larger proteins

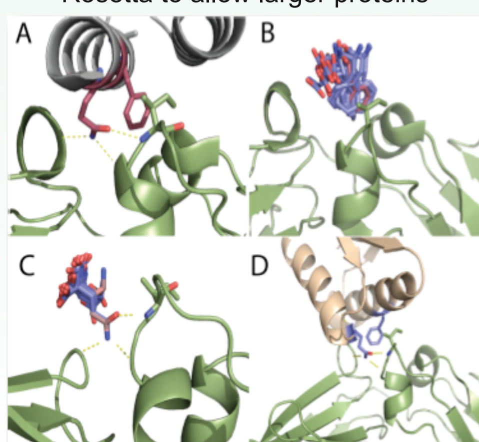

The interfaces of protein-protein complexes often exhibit a handful of key interactions, termed hot-spots. At right, the original protein (A) is replaced by an easy-to-manufacture custom scaffold (D)

INCITE

Molecular Simulation of Complex Chemical Systems

Christopher Mundy
LLNL

2008-2009 Progress

- Computational resolution of the conundrum of whether the hydroxide ion is more likely to found at a hydrophobic interface or in the solvation environment provided by bulk solution
- Computational insights to the nature of hydronium ion (H3O+) at the air water interface

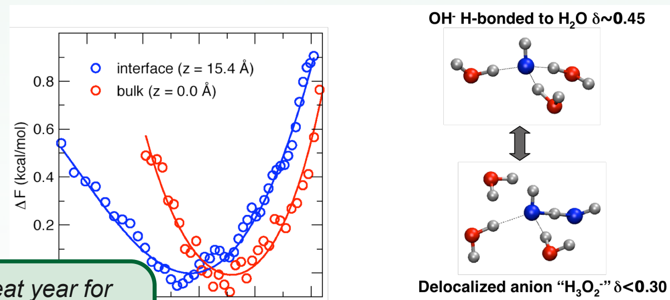

Figure 1: Free-energy surfaces and representative clusters as a function of the sampled reaction coordinate

"It has been another great year for science on INCITE resources. The staff at ORNL and ANL are first rate".

INCITE

Kinetics and Thermodynamics of Metal and Complex Hydride Nanoparticles

Christopher Wolverton
Northwestern University

- Explore the stability of a Li4BN3H10 nano-droplet
- Explore the reaction pathway for NH3 formation in bulk liquid Li4BN3H10.

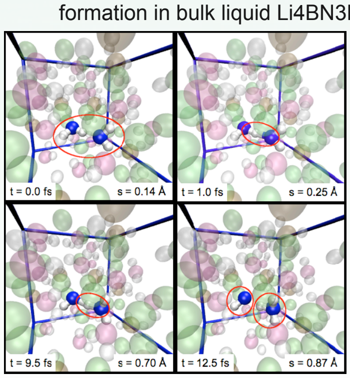

Figure 1 Snapshots of configurations relative to starting configuration for NH3 formation reaction. The reaction coordinate was defined as the distance between the "hopping" hydrogen traveled. The energy difference for the reaction is ~3.2 eV, a very large value. We attribute this to the energetic cost of the motion of the other atoms in the system.

- Explore GPAW as parallel bead Nudged Elastic Band and ensemble run capabilities of GPAW/ASE should benefit from the BG/P architecture.
- We have made changes to the ASE and GPAW source code that will allow us to run N-member MD ensembles across an arbitrary number of processors.

•"ALCF staff has been of great help in trying to find a workable solution to scalability issue of DFT codes. GPAW has allowed to carry large-scale runs "
•"The ALCF staff has been incredibly helpful and responsive to our requests. They have assisted in tuning codes, improving convergence of the electronic solver and compiling codes. "

INCITE

Modeling the Rheological Properties of Concrete

William George
NIST

Improving the Flow Properties of Cement Based Materials

- Develop a theoretical basis for understanding rheological properties of complex fluids like colloidal suspensions
- Identified a new physical mechanism '*supplementary stress*' needed when determining the yield stress of colloidal fluids
- Shown the application of a small strain can accelerate aging
- Quantified how stress and strain is transmitted throughput the suspension
- Progress in modeling a suspension with a non-Newtonian fluid matrix

Methods and Challenges

- Fundamental change in algorithms for avoid n^2 operation
- Implementation of parallel I/O

Suspension of spheres in a non-Newtonian Shear-Thinning Fluid Matrix
Experimental Data from SIKA (curves)
Simulation data (circles)

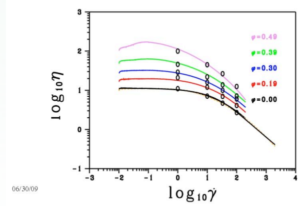

Baseline Area 4: Financial Summary

Metric: *The LCF will provide quarterly reports on the steady state costs to compare against plans as described in the OMB 300.*

Costs in FY2010 are expected to increase in line with a funding increase up to $42M. The majority of the increase (over $9M) will pay for ALCF-2 DME/Facilities, and the balance will go toward staffing increases. DME/Facility work is necessary preparation for the ALCF-2 BG/Q machine, and increases in staff will help meet demand for expanded science support on ALCF-1 as well as DME.

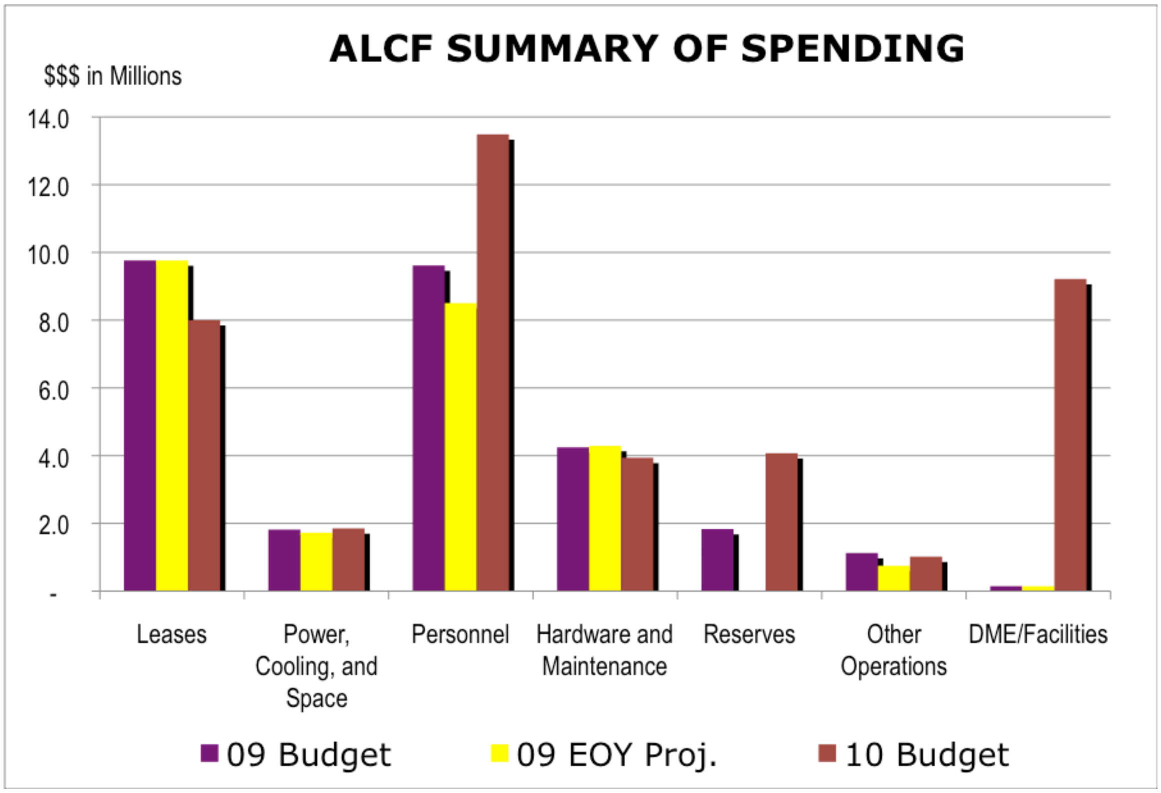

	2009 Budget	2009 Projections	2010 Budget
Leases	$9.8	$9.8	$8.0
Power, Cooling, and Space	$1.8	$1.7	$1.9
Personnel	$9.6	$8.5	$13.5
Hardware and Maintenance	$4.2	$4.3	$3.9
Reserves	$1.8	-	$2.6
Other Operations	$1.1	$0.8	$1.0

DME/Facilities		$0.1	$0.1	$9.2
Commitments		-	$3.3	$2.0
	TOTAL	28.5	$28.5	42.0

The largest budgeted expense categories in FY2010 are Personnel (32%), DME/Facilities (22%), and Leases (19%). This compares to FY2009 budgets: Personnel (34%), Leases (34%), and negligible DME/Facilities and FY2009 projected actual costs: Personnel (30%), Leases (34%), and negligible DME/Facilities. It is worth noting that because of the increased funds in FY2010, each percentage point represents ~$120K more in actual dollars.

Costs included in the other operations category cover expenditures for miscellaneous areas like education/training, user services expenses, travel, small value equipment (including PCs), telephones, etc.

"Commitments" represent contractually obligated costs in the following fiscal year for things like leases, maintenance contracts, etc.

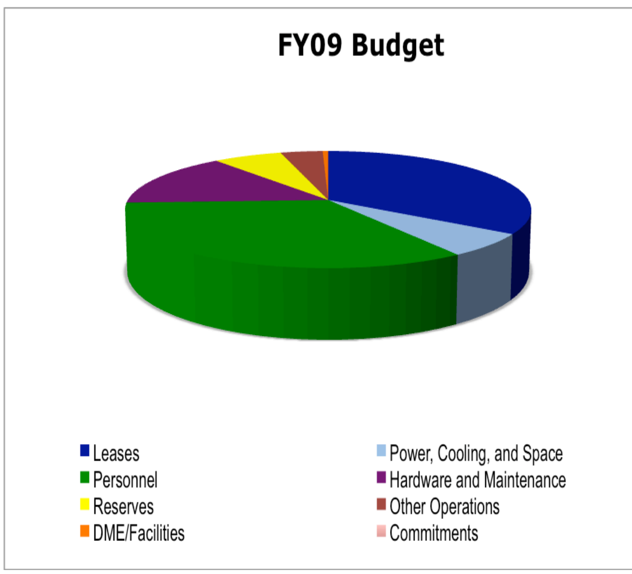

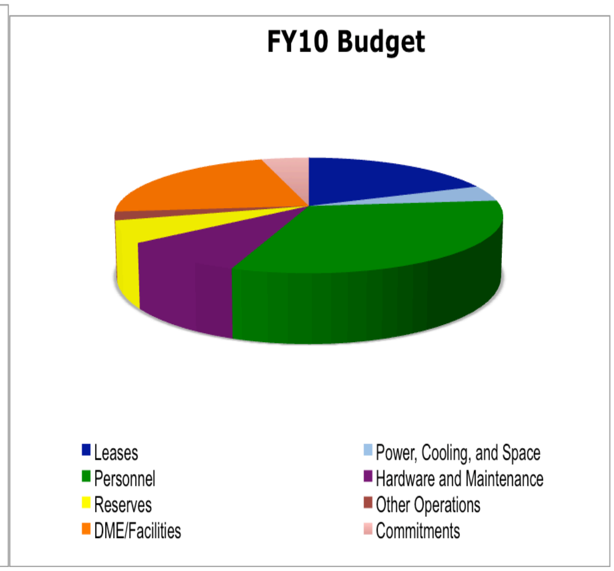

Personnel costs were below budget in FY2009 due to the Continuing Resolution and a slower ramp-up of staff than planned. Staffing levels increased toward year-end, in line with the plan. FY2010 has an aggressive hiring plan, which the funding supports. ALCF-1 personnel are budgeted to increase by 8 FTE in addition to 10 FTE for DME. ALCF maintains an ability to draw upon contractors and other parts of the Lab to provide support as required. For example, there are close ties to the Mathematics and Computer Science Division (MCS), with a corresponding interchange of technical knowledge and use of MCS staff as needed. There is also some reliance on other parts of the lab for operational support, including expertise in cyber security and networking (CIS), facilities (FMS), and technical services (TSD).

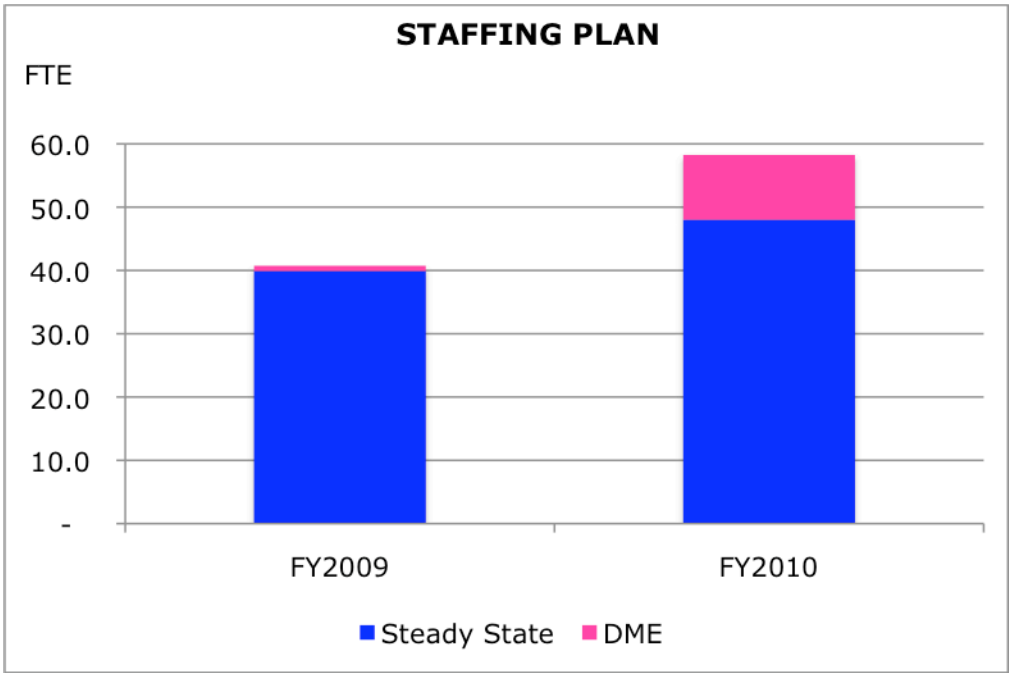

There is $2.6M of contingency/reserve in the FY2010 budget. Potential uses for these reserve funds have already been identified.

The decision to do self-maintenance on parts of the facility, such as compute nodes and disk systems, results in savings of about $1M over contracting for full maintenance support with IBM. This was based on a best practice adopted from Lawrence Livermore National Laboratory (LLNL). The experience of time spent and costs of spare parts cache has certainly borne out this decision.

The competitive procurement for financing of equipment leases resulted in a significant savings over budget of about $9M total over the lifetime of the leases.

Electric rates increased in FY2009. Since Blue Gene/P uses less electricity in operations than expected (even for a low-power machine), electricity costs were still less than budgeted.

Baseline Area 5: Innovation

Metric: *The LCF will report annually on innovations and best practices developed, shared with the community, and imported from the community, as well as participate in the yearly ASCR-sponsored Facilities Best Practices Workshop.*

Best Practices Developed: Continued Partnership with Computer Science Developers

For faster application innovation, leading to breakthrough science on the BG/P, ALCF worked closely with computer scientists and vendors to provide necessary tools for application developers:

- ZeptoOS. Released version 2.0 on BG/P with new kernel features such as support for a new memory subsystem called big memory (flat, TLB miss free memory region). Also included was a port of BG/P SPI and DCMF low-level libraries and support for MPI applications in SMP mode. Kernel tools and analysis utilities (KTAU) were also integrated. The FALKON science applications on the BG/P have benefited from these changes.

- Cobalt. Several new production releases on the BG/P dramatically improved scheduling performance. Job preemption and scavenging have been implemented in Cobalt and will be deployed at ALCF soon. Utility-based scheduling has also been implemented and has greatly improved productivity at ALCF (see Innovation section). ALCF worked closely with the Cobalt team to characterize application performance on unwrapped mesh networks compared to torus networks. This study would lead to more efficient job scheduling on the BG/P in the future.

- TAU. ALCF worked with ParaTools to reduce overhead for automatic instrumentation of Python code. ALCF now has the capability of profiling Python and mixed Python-C codes on the same footing as conventional programming languages (C, Fortran).

- MPI. On ALCF's BG/P, there have been several new releases of MPICH2. Support for MPI-2.1, an efficient hybrid MPI + thread model (critical for ALCF-2), and memory and scalability optimizations for large-scale runs, have been included. Paul Fischer's NEK5000 application is one that greatly benefited from the memory optimizations. Also developed was an Asynchronous Dynamic Load Balancing Library (ADLB) that provides a scalable distributed work queue with no single master bottleneck. It enabled Steve Pieper's GFMC code to scale from 2,000 processes to more than 20,000 processors on the BG/P.

- MPI-IO. Working with IBM, MPI-IO was updated to MPICH-2.1.1. This fixed bugs that applications such as FLASH and VORPAL were encountering on the BG/P.

- Parallel Virtual File System (PVFS) . Developed Darshan – a lightweight tool for characterizing I/O behavior. This has been installed on ALCF's test and development BG/P system to profile I/O characteristics of our applications.

- Totalview. ALCF has a scalability agreement in place to improve performance of large- scale debugging jobs (32k MPI tasks). ALCF has been working jointly with LLNL on working group calls to improve debug turnaround time. Based on feedback from ALCF, Totalview developed Remote client.

Innovations:

- Bulk Power Module on the BG/Q. ALCF members and Argonne safety engineers worked with the IBM design team to have them incorporate safety and efficiency improvements in the design of the bulk power module for the ALCF-2. The power factor was improved by approximately 10%. This should result in substantial cost savings during operations. A breaker was incorporated into each BPM. This improves safety, as there is no potential for an arc flash when the BPM is changed. It is also an operational improvement, since it alleviates the need for a three-person team and high-voltage PPE during changes.

- Water Cooling. ALCF members worked with the IBM design team on the water cooling system for the ALCF-2. This water cooling system will utilize much warmer water than is typical (64°F – 77°F), allowing the opportunity for free cooling more during the year and overall reducing operational costs and environmental impact.

- ALCF continues to work on improving the Power Usage Effectiveness (PUE) in our facility. ALCF has run experiments, in cooperation with the local facilities team, to raise the incoming chilled water temperature as much as possible, while still maintaining temperature control in the room. ALCF hopes to begin experimenting with increasing the overall temperature in the room in the next year.

- Utility scheduling, in Cobalt, is a mechanism for prioritizing jobs based on competing scheduling goals. During each scheduling cycle, policies define a dynamic score for each job. These policies are site-specific python functions that return a pair of values. The first is a score for the job, essentially the value of running the job. The second value describes the extent to which the scheduler can reorder jobs, based on currently available resources. The scheduling process consists of scoring all idle jobs and ordering them by score. Cobalt assigns resources based on job scores, using a standard backfilling mechanism to drain resources. Also, the scheduler can use the fallback threshold to reorder jobs.

  These tools have been used to build scheduling policies that are tuned to the leadership-class workload and priorities of the ALCF. Several goals are explicitly supported by these policies:
    - Large jobs

- Reasonable response times, even for large jobs
- Good utilization of the machine
- Prioritizing of key projects

The production scheduling policy for Intrepid consists of the following formula:

(wait time/wall time)^3 * (requested nodes/total nodes) * project priority

This formula encodes several important ideas. The first is that shorter jobs should expect to wait less in the queue than longer jobs. The second is that large jobs should be priorities. And the final idea in this utility function is that projects have different relative priorities.

In conjunction with the Cobalt simulator, this function has been tuned specifically for the Intrepid workload. For example, it was determined empirically that cubing the time factor (as opposed to squaring it) produced better utilization and response times for our workload. Simulation of this kind has enabled us to ensure maximum scheduler performance as Cobalt and the ALCF's scheduling policy undergo modification.

Baseline Area 6: Risk Management

Metric: *The ALCF will use a documented risk management process.*

Risk management is the process, methods, and approach to measuring risk and developing strategies to manage it. Common risk management strategies include transferring the risk to another party, reducing the negative impact of the risk through specific actions or changes in plans, and accepting some or all of the implications of a particular risk.

The goal of the risk management strategy for the Argonne Leadership Computing Facility (ALCF) is aimed at maximizing the likelihood of success within the overall cost, technical, and schedule envelope.

Our strategy is based on the following approach:

- Establishing a context and framework for risk management, identifying the process scope, the stakeholders, the risk analysis basis, and an agenda for identification and analysis;
- Identifying potential risks and failure modes for the project and steady-state elements;
- Aggressively identifying sources of uncertainty in the project (scope, cost, schedule, functionality, etc.);
- Characterizing those uncertainties, including estimating the likelihood of occurrence, impact of occurrence, and prioritization based on overall impact;

- Executing risk management decisions—choosing which risks to actively manage via analysis of the options available (avoidance, transfer, mitigation, or retention);
- Developing a system that will track sources of risk and providing mechanisms for risk identification, risk management decisions, risk treatment tracking, and risk retirement.

The ALCF Risk Management Plan has been in place since November 2006. The ALCF has implemented the Risk Management Plan (RMP), Version 3.0, which provides the guidance for the ALCF's risk management strategy. The goal of the risk management strategy for the ALCF is aimed at maximizing the likelihood of success within the overall cost, technical, and schedule envelope.

The RMP outlines the risk management process and describes how risks are identified, assessed, tracked, updated, and retired. The Risk Management Program, PertMaster®, is the tool used to track risks, quantify and qualify risk impacts, and provide a risk register.

The ALCF continually evaluates operational risks through monthly risk status meetings, and shares major risks with the Program Office. These meetings enhance the effectiveness of risk management by evaluating impacts, contingencies, and mitigations as they develop. Some of the impacts and mitigations become supporting evidence in developing relevant lessons learned.

When analyzing steady state risks, we apply the criteria shown in the following figures:

Impact **Risk Area**	**Low**	**Moderate**	**High**
Cost	≤ $100K	$100k ≤ $500K	>$500K
Schedule	Impacts current FY OBM e300 specified INCITE production requirements by < 5%	Impacts current FY OBM e300 specified INCITE production requirements by 1 ≤ 3 months	Impacts current FY OBM e300 specified INCITE production requirements by >3 months
Technical	Negligible, if any, degradation to data security, storage, or customer support	Significant degradation to data security, storage, or customer support	Severe degradation to data security, storage, or customer support

	Impact		
Probability	**Low**	**Moderate**	**High**
High (p > 75%)	Low	Moderate	High
Moderate (25% < p < 75%)	Low	Moderate	High
Low (p < 25%)	Low	Low	Moderate

The following figure shows our current risk status:

	IMPACT		
PROBABILITY	**LOW**	**MODERATE**	**HIGH**
HIGH		**(3)**	
MODERATE	**(3)**	**(4)**	
LOW		**(27)**	**(2)**

To date, we have had a total of 63 steady state risks, the 39 active risks listed above, 23 that have been retired, and one recently added risk that is still in the proposed state for reporting purposes. Examples of retired risks and how they compared to our mitigations are as follows:

- 1,4,7 – Aggressive performance goals of (GPFS, PVFS, MPI-IO) not met
    - We reached good performance during the acceptance tests, but planned to tune the system and increase performance. We exercised our mitigation strategies and used system time and ALCF and MCS staff to tune the performance.
- 949, 1036 – Underestimated labor needs, Overworking staff
    - This risk was realized, we had underestimated labor needs. We exercised the mitigation of utilizing management reserves to hire additional staff, while remaining within budget.

Many risks did not manifest themselves. Examples include:

- 982 – Cobalt job scheduler does not scale
- 13 – Aggregate Remote Copy Interface (ARMCI) / Unified Parallel C (UPC) ports not available
- New BG/P service node security model insecure.

Our current moderate (we have none that are high) schedule and cost risks are:

- 142 - Funding is reduced in FY09-10 (low probability, high impact)
    - We mitigate this risk through regular communication with the program office, standard financial reports that monitor planned vs. actual costs, and in the event of a shortfall, cost cutting and use of management reserves.
- 990 - Electric Cost Could Increase Beyond the Budget (low probability, high impact)
    - We mitigate this risk by having negotiated annual “locked in” electric costs, standard financial reports that monitor planned vs. actual costs, and in the event of a shortfall, cost cutting and use of management reserves.
- 1018 - INCITE Users are not provided Adequate Support by ALCF (moderate probability, medium impact)
    - We monitor this via regular contact by the support team, monthly user calls, and metrics such as tickets addressed and the annual user survey. In the event we determine there is a support issue, we would try and optimize our support resources and/or expend management reserves to hire additional support resources.
- 1010 - We have more Intrepid INCITE projects than planned for 2010 (moderate probability, medium impact)
    - A significant mitigation to this risk is that we are now deeply involved in the selection of INCITE projects. Should we end up with more than planned, we would increase the use of our “tiered” support and reduce support to discretionary projects, reduce the number of discretionary projects, or use management reserves to hire more support staff.
- 978 - Component Failures Minor and Major (high probability, medium impact)
    - As noted elsewhere in this document, failures of Myricom optical transceivers have caused service interruptions. We have implemented advanced diagnostic procedures to rapidly diagnose and respond to these failures. We continue to work with Myricom to pro-actively replace optical transceivers before they fail. However, should the rate of failure not improve, we have sufficient management reserves to replace all of the transceivers associated with the manufacturing lot.
- 979 - Network Stability Problems (high probability, medium impact)
    - There are two primary components to this risk. Instability introduced into the network by the failures listed above, and firmware bugs in new releases. We mitigate component failures with advanced scripts. Mitigations for the second involve skipping software releases unless they have features or fixes we explicitly need and providing substantial maintenance time to allow for extensive testing and a rollback to a previous version of firmware should that be needed.

- 139 - Stability issues on the Data Direct Networks (DDN) 9900s (moderate probability, medium impact)
  - Mitigations include closely monitoring the issue and coordinating closely with DDN. We also have developed site-specific procedures to avoid known problems. For instance, write-back cache is enabled by default on replacement controllers, but this triggers a known firmware bug, so we have special procedures to work around that. We also encourage users to archive their data to tape regularly. We are also in discussions with DDN to provide an on-site engineer to provide better service.
- 137 - Lack of performance in the tape archive results in users not using tape archive as needed (high probability, medium impact)
  - The performance of the High Performance Storage System (HPSS) software has been lower than expected / desired. So far the bottleneck has been in the HSI and HTAR clients. Both are limited by a single host, and HTAR can not handle the large file sizes and path names that we sometimes have. Until recently, there were bugs in GridFTP causing it to be unstable when used with HPSS. Mitigations include scripts written to try and get performance through concurrency, working with the GridFTP to resolve issues (done) and increase performance (ongoing), working with Mike Gliescher to use the current tar specification to resolve file size and path length problems, and working with the HPSS team.
- 141 - Higher than anticipated Blue Gene spare part consumption due to aggressive diagnostics may require purchase of additional spare parts (moderate probability, medium impact)
  - We are very aggressive in the replacement of parts during diagnostics, replacing components before they actually fail. This aggressive preventative maintenance gives Intrepid a very good Mean Time To Interrupt (MTTI). However, this aggressive schedule does use more spare parts than waiting for a component to completely fail before being replaced. Mitigations in progress include varying the point at which we replace parts during diagnostics, as well as using marginal parts removed from our production resource (intrepid) in our test and development resource (surveyor). Finally, we track the consumption rate and project our usage. We have sufficient management reserves to buy additional parts if we continue to feel that the improved MTTI justifies the more aggressive parts replacement policy.

Baseline Area 7: Cyber Security

Metric: *The LCF will maintain an approved Cyber Security Program Plan and an approved Cyber Security Certification and Accreditation.*

The Advanced Leadership Computing Facility (ALCF) operates its cyber security program in conjunction with Argonne's Cyber Security Program Office. The ALCF leverages the central cyber security services provided by the Laboratory, and is today accredited at a FIPS 199 Moderate level under the Laboratory's Certification and Accreditation envelope. Accreditation is valid through January 23, 2011.

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

28 MAR 2008

Dr. Robert Rosner
Director, Argonne National Laboratory
President, UChicago Argonne, LLC
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Rosner:

SUBJECT: ACCREDITATION OF ARGONNE NATIONAL LABORATORY (ANL) CYBER ENCLAVES – AUTHORITY TO OPERATE (ATO)

References:
1. Expansion of Argonne National Laboratory (ANL) Cyber Enclaves – Authority to Operate (ATO) to include the Advanced Leadership Computing Facility (ALCF), C. Catlett to R. Lutha, February 25, 2008
2. Accreditation of Argonne National Laboratory (ANL) Cyber Enclaves – Authority to Operate (ATO), January 23, 2008

The submitted certification documentation (Ref. 1) has been reviewed for the General Computing Enclave and its major applications:

- Accelerator Control Systems (APS, IPNS)
- Business Systems
- Sensitive Information
- Advanced Leadership Computing Facility

This enclave and its associated major applications include all IT investments at ANL. The attached list designates the systems included in each enclave that have been reported under the requirements of the Federal Information Systems Management Act (FISMA).

Based on
- the testing performed by the July 2006 DOE Office of Science Site Assistance Visit (SAV),
- the FY2007 independent Security Test and Evaluation (ST&E) performed by Grant Thornton,
- the close-out of the DOE-IG's Certification and Accreditation finding, and
- the review of the Advanced Leadership Computing Facility performed by the DOE Chicago Integrated Support Center, Safeguard and Security Division.

A component of the Office of Science

Performance Results and Projections for the Next Year

Customer Results

The LCF will conduct a user satisfaction survey in December each year. The survey will be used as an annual tool to help gauge user satisfaction and identify areas for continued improvement. For the formal user satisfaction surveys going forward, the overall LCF score will be at a level of 3 or higher (out of 5.0).

In addition to our formal survey in December, we use other means of feedback (formal and informal) to better understand how LCF users view our services and capabilities. These include:

- Formal workshop surveys
- Formal requests for feedback as part of the quarterly INCITE updates
- Informal discussions over the course of the INCITE year to solicit feedback and comments on users' experiences with LCF, as well as discussions at workshops and outreach events on specific issues or concerns.

LCF problem resolution is measured on problem reports sent by the users to the LCF problem tracking system. For FY2010, 73% of LCF user problems shall be addressed within three working days, either by resolving the problem or informing the user how the problem will be resolved. The percentage addressed within three working days will reach 80% in FY2011.

LCF will track workshops, tutorials, monthly user teleconferences, and application support provided to users and will provide quarterly reports to the U.S. Department of Energy.

Business Results

LCF resource availability will be measured on major systems and reported quarterly. In FY10, the system will be assessed at 85% scheduled availability and 80% overall availability based on our standard method. Our standard method of determining availability: beginning the start of the quarter following general availability, new or upgraded systems will meet 75% scheduled availability and 70% overall availability; systems in operation more than one year since general availability or the last upgrade will meet 85% scheduled availability and 80% overall availability.

LCF capability usage of the major systems will be measured quarterly and reported annually. Beginning FY10, 250 million CPU-hours will be delivered to jobs 8,192 nodes or larger.

Strategic Results

The LCF will track the science output and accomplishments for each project, including milestone reports, presentations, publications, journal covers, and awards and will provide quarterly reports on the project results. The LCF will also track technology accomplishments, such as development of reusable code that results in a new tool for its discipline and new algorithm design ideas or programming methodologies.

Financial Performance

The LCF will provide monthly reports on the steady-state costs to compare against plans as described in the OMB 300.

The LCF will use a documented risk management process.

Innovation

The LCF will report annually on innovations and best practices developed, shared with the community, and imported from the community, as well as participate in the yearly ASCR-sponsored Facilities Best Practices workshop.

Security

The LCF will maintain an approved Cyber Security Program Plan and an approved Cyber Security Certification and Accreditation.