Scientific Challenges at the Leadership Computing Facility over Next Decade

Vision of the LCF: CORAL procurement is the next step

LCF's next systems, OLCF-4 & ALCF-3 are being procured through CORAL

- CORAL: Collaboration of Oak Ridge, Argonne, Livermore.
- Created to jointly acquire leadership computing systems for DOE's National Nuclear Security Administration (NNSA) and Office of Science (SC).
- Formed on the basis of common acquisition timings.
- Offers a "win-win":
  - Reduces number of RFPs for vendors.
  - Allows pooling of R&D funds.
  - Supports sharing of technical expertise among labs.
  - Strengthens SC/NNSA alliance for exascale.

CORAL procurement: Execution process

- February 18, 2014: Vendor proposals submitted.
- Proposals were evaluated in a 2 step process.
  - 8 teams of technical experts assessed responsiveness to Draft Statement of Work and proposal instructions.
  - Buying team selected 2 proposals providing best value.
- Announcement of results will be made once contracts have been negotiated.
- Each system is expected to cost ~$125M and deliver 100–200 PF.

CORAL Procurement Model

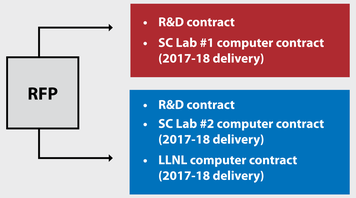

High Level System Targets

- Target speedup over current systems of 4x on Scalable benchmarks and 6x on Throughput benchmarks – highest priority.
- Peak Performance ≥ 100 PF.
- Aggregate memory of 4 PB and ≥ 1 GB per MPI task (2 GB preferred).
- Maximum power consumption of system and peripherals ≤ 20MW.
- Mean Time Between Application Failure that requires human intervention ≥ 6 days.
- Data centric capabilities.

CORAL benchmark categories represent DOE workloads and technical requirements

Scalable science benchmarks
- Expected to run at full scale on systems.

Throughput benchmarks
- Represent large ensemble runs.

Data centric benchmarks
- Represent emerging data intensive workloads.
- Integer operations, instruction throughput, indirect addressing.

Skeleton benchmarks
- Investigate network performance, threading overheads, I/O, memory, memory hierarchies, system software, and programming models.

Micro benchmarks
- Small code fragments that represent expensive compute portions of some of the scalable science and throughput applications.

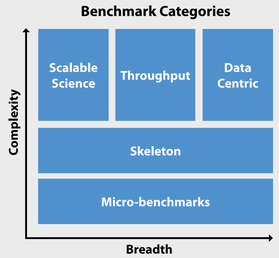

CORAL benchmarking suite uses mini-apps and a few larger applications

Categories	Scalable Science	Throughput	Data Centric	Skeleton
Marquee (TR-1)	LSMS QBOX NEKbone HACC	CAM-SE UMT2013 AMG2013 MCB	Graph500 Int sort Hashing	CLOMP IOR CORAL MPI Memory CORAL loops
Elective (TR-2)		QMCPACK NAMD LULESH SNAP miniFE	SPECint_ peak2006	Pynamic HACC I/O FTQ XSBench miniMADNESS
Elective Micro-Benchmarks (TR-3)	NEKbonemk HACCmk	UMTmk AMDmk MILCmk GFMCmk		

Application Readiness & Early Science Partnerships

ALCF and OLCF will issue calls for proposals for application development partnerships between community developers, LCF staff and the respective vendor's Center of Excellence located at LCF center.

- Performance analysis of community applications.
- Technical plan for code restructuring and optimization.
- Deployment on OLCF-4 & ALCF-3.
- Early Science Access path for applications.

More information on CORAL

https://asc.llnl.gov/CORAL
Provides all documents for the bidders

https://asc.llnl.gov/CORAL-benchmarks/
Provides all information on the benchmark codes

Selected key science drivers for LCF

Fusion Energy/ITER

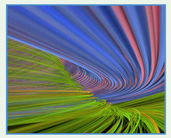

Key science challenges: Effectively model and control the flow of plasma and energy in a fusion reactor, scaling up to ITER-size. Develop predictive understanding of plasma properties, dynamics, and interactions with surrounding materials. Mitigate plasma disruptions.

Biomass to Biofuels

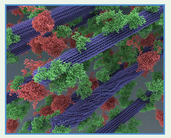

Key science challenges: Enhance the understanding and production of biofuels from biomass for transportation and other bio-products. The main challenge to overcome is the recalcitrance of biomass (cellulosic materials) to hydrolysis.

Nuclear Energy

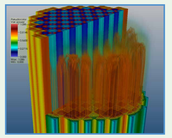

Key science challenges: For existing reactors provide safe, increased fuel utilization, power upgrades, and reactor lifetime extensions. Design new, safe, cost-effective reactors.

Energy Storage

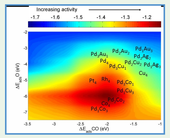

Key science challenges: Gain the fundamental understanding of reaction processes at the atomic and molecular level required for predictive design of new materials for energy storage and predictive engineering of safe, large-format rechargeable batteries.

Integration of compute and data at LCF

LCF integrates data analysis and simulation capabilities

- Simulation and data are critical to DOE.
- Both need more computing capability.
- Both have similar hardware technology requirements.
  - High bandwidth to memory.
  - Efficient processing.
  - Very fast I/O.
- Different machine balance may be required.

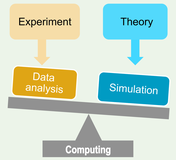

Big data
Analyzing and managing large complex data sets from experiments, observation, or simulation and sharing them with a community.

Simulation
Used to implement theory; helps with understanding and prediction.

Integrating Compute & Data Capabilities: Creating a Pathway for Scientific Discovery

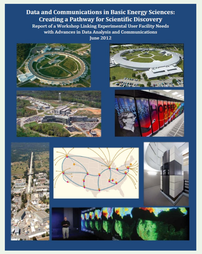

Accelerating discovery in neutron sciences and enhancing predictive capabilities

- Theory and analysis components should be integrated seamlessly within experimental workflow.
- Move analysis closer to experiment – future possibility of experiment steering.
- Match data management access and capabilities with advancements in detectors and sources.

ADIOS wins big: LCF project named one of top 100 tech products

- The LCF's Scott Klasky received an R&D 100 Award for his development of the Adaptable I/O System for Big Data, or ADIOS.
- Developed by ORNL, Georgia Institute of Technology, Rutgers University, and North Carolina State University, ADIOS is a portable, scalable, easy-to-use software framework conceived to solve "big data" problems.
- Klasky began developing ADIOS in 2008, and now the middleware assists numerous applications on Titan.
- ADIOS has demonstrated impressive I/O performance results on leadership class machines and clusters, sometimes showing an improvement of more than 1,000 times over well-known parallel file formats.
- The research was funded by DOE's LCF, the office of Advanced Scientific Computing Research, the Office of Fusion Energy Science, and the NSF.

PanDA Tool Provides Titan with High-Tech Scheduler

- Researchers with the ATLAS experiment in Europe have been integrating its scheduling and analysis tool, PanDa, with Titan and Mira.
- Workflow includes 1.8 million computing jobs each day distributed among 100 or so computing centers spread across the globe.
- PanDA's ability to efficiently match available computing time with high-priority tasks holds great promise for the LCF.
- Team developers redesigned parts of the PanDA system responsible for job submission on remote sites and gave PanDA new capability to collect information about unused worker nodes on LCF machines.
- Deployment of the tool could lead to a higher utilization of available hours on LCF resources.

Acknowledgements:
LCF Vendor Partners: IBM, Cray, AMD, NVIDIA, etc.
LCF Users: DOE ASCR Leadership Computing Program

Contact:
Jack Wells, wellsjc@ornl.gov
Richard Coffey, richardc@alcf.anl.gov

ORNL 2014-G00839/jbs