Argonne NATIONAL LABORATORY

Leadership Computing Facilities

Katherine Riley
Argonne National Laboratory, ALCF

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

U.S. DEPARTMENT OF ENERGY

Outline

- What are the DOE Leadership Class Resources
- How to gain access and use the LCFs
- Collaborative model and why we need scientific communities to work with us

U.S. Department of Energy (DOE) Computing Facilities

Leadership Computing Facilities

Argonne Leadership Computing Facility
Allocations: 60% INCITE, 30% ALCC, 10% Discretionary
Resources: Mira 10PF BG/Q, Tukey (Data Analytics)
Users: 300+
Projects: 125 (37, 7 ALCC, 90 Discretionary)

Oak Ridge Leadership Computing Facility
Allocations: 60% INCITE, 30% ALCC, 10% Discretionary
Resources: Titan 10-20PF Cray , Lens (Data Analytics)
Users: 100s
Projects: 100+ (32 INCITE, 9 ALCC, ~100 Discretionary)

National Energy Research Scientific Computing Center
Allocations: 80% DOE SC, 10% ALCC, 10% Discretionary
Resources: Hopper 1.28PF, Euclid (Data Analytics)
Users: 1000s
Projects: 100s

INCITE resources: Mira at ALCF

- *Mira* - Blue Gene/Q System
  - 48K nodes / 768K cores
  - 786 TB of memory
  - Peak flop rate: 10 PF
- Storage
  - ~35 PB capacity, 240GB/s bandwidth (GPFS)
  - Disk storage upgrade planned in 2015
    - Double capacity and bandwidth
- New Visualization Systems
  - Initial system in 2012
  - Advanced visualization system in 2014
    - State-of-the-art server cluster with latest GPU accelerators
    - Provisioned with the best available parallel analysis and visualization software

INCITE resources: Titan at OLCF

- *Titan* – Upgrade of Jaguar Cray X5
  - 2.6PF Opteron performance
  - Accelerated node design using NVIDIA multi-core accelerators
  - Peak flop rate: 27 PF
  - 694 TB memory (32GB+6GB)
- Storage
  - ~1 TB/s
- Gemini interconnect
  - 3-D Torus
  - Globally addressable memory
  - Advanced synchronization features

Titan Specs	
Compute Nodes	18,688
Login & I/O Nodes	512
Memory per node	32 GB + 6 GB
NVIDIA “Kepler” (2012)	1.31 TFlops
Opteron	2.2 GHz
Opteron performance	141 GFlops
Total Opteron Flops	2.6 PFlops
Disk Bandwidth	~ 1 TB/s

Diversity of INCITE science

Simulating a flow of healthy (red) and diseased (blue) blood cells with a Dissipative Particle Dynamics method.
- George Karniadakis, Brown University

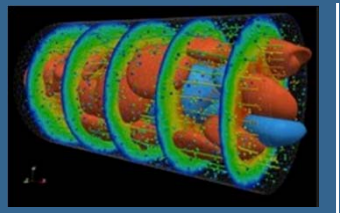

Provide new insights into the dynamics of turbulent combustion processes in internal-combustion engines.
-Jacqueline Chen and Joseph Oefelein, Sandia National Laboratories

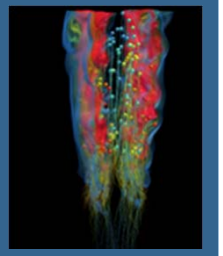

Demonstration of high-fidelity capture of airfoil boundary layer, an example of how this modeling capability can transform product development.
- Umesh Paliath, GE Global Research

Calculating an improved probabilistic seismic hazard forecast for California.
- Thomas Jordan, University of Southern California

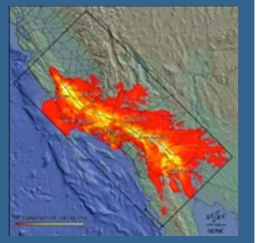

Modeling charge carriers in metals and semiconductors to understand the nature of these ubiquitous electronic devices.
- Richard Needs, University of Cambridge, UK

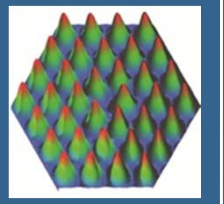

High-fidelity simulation of complex suspension flow for practical rheometry.
- William George, National Institute of Standards and Technology

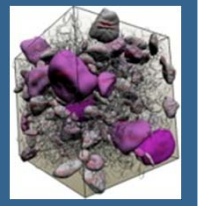

Other INCITE research topics

- Glimpse into dark matter
- Supernovae ignition
- Protein structure
- Creation of biofuels
- Replicating enzyme functions
- Global climate
- Regional earthquakes
- Carbon sequestration
- Turbulent flow
- Propulsor systems
- Membrane channels
- Protein folding
- Chemical catalyst design
- Combustion
- Algorithm development
- Nano-devices
- Batteries
- Solar cells
- Reactor design
- Nuclear structure

Allocation Programs at the LCFs

	60% INCITE		30% ALCC		10% Director's Discretionary	
Mission	High-risk, high-payoff science that requires LCF-scale resources*		High-risk, high-payoff science aligned with DOE mission		Strategic LCF goals	
Call	1x/year – (Closes June) ***2014 Call Open Now***		1x/year – (Closes February)		Rolling	
Duration	1-3 years, yearly renewal		1 year		3m,6m,1 year	
Typical Size	30 - 40 projects	50M - 500M core-hours/yr.	5 - 10 projects	10M – 300+M core-hours/yr.	100s of projects	.5M – 10M core-hours
Review Process	Scientific Peer-Review	Computational Readiness	Scientific Peer-Review	Computational Readiness	Strategic impact and feasibility	
Managed By	INCITE management committee (ALCF & OLCF)		DOE Office of Science		LCF management	
Readiness	**High**		**Medium to High**		**Low to High**	
Availability	Open to all scientific researchers and organizations ***Capability > 131,072 cores (16.7% of Mira)***					

Diverse ALCF User Community

- ALCF is targeted to a few very large science projects
- Minimal award in 2014 expected to be at least 50M core-hours

- Diverse base of users with diverse needs
- Scientific support helps **transition, plan and develop** for unique HPC resources
- Teams are from **diverse** scientific background
  - Scientists/engineers
  - Scalable algorithms
  - Performance engineering
  - Scientific code development

2013 INCITE Allocations at ALCF (Mira & Intrepid)

736M Intrepid, 2.1B Mira

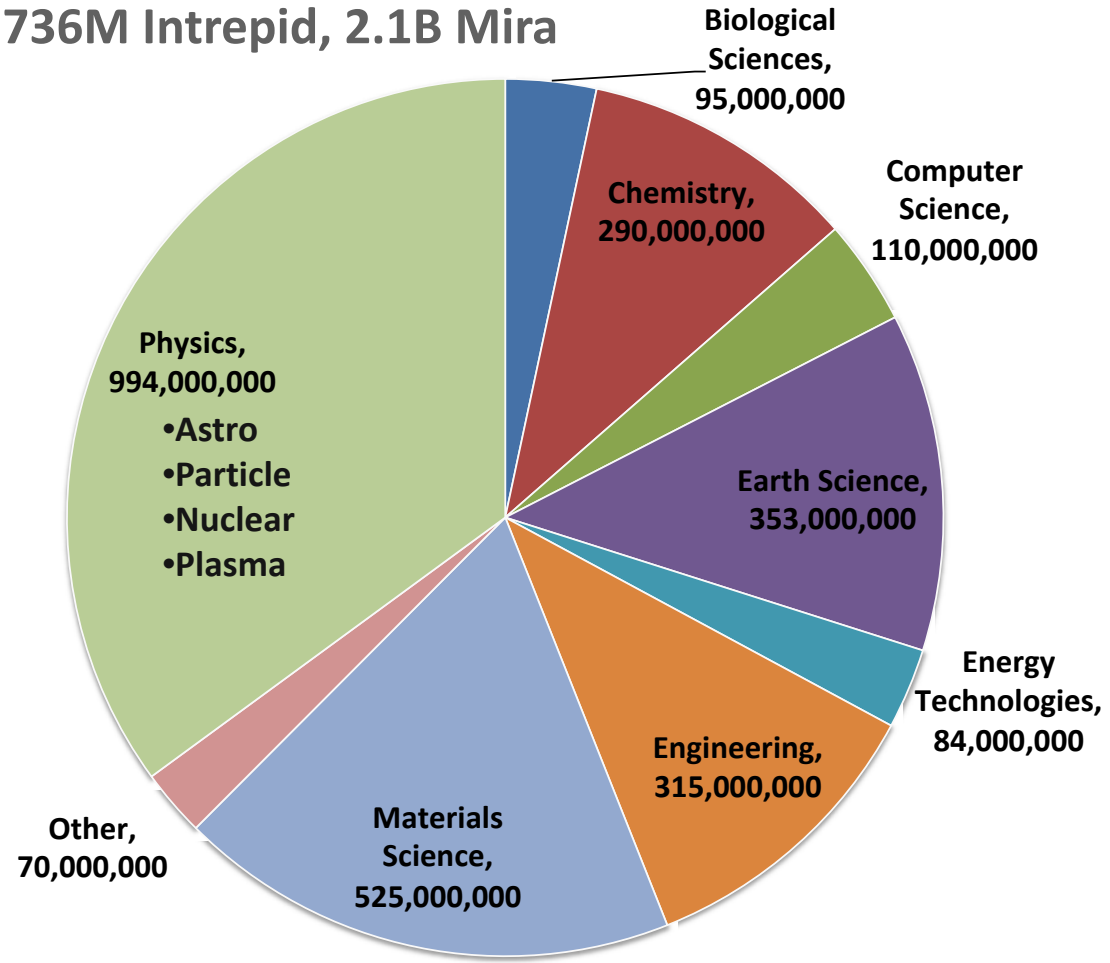

INCITE 60%, ALCC 60%, Discretionary 10%

Twofold review process

	New proposal assessment	Renewal assessment
Peer review: INCITE panels (1)	• Scientific and/or technical merit • Appropriateness of proposal method, milestones given • Team qualifications • Reasonableness of requested resources	• Change in scope • Met milestones • On track to meet future milestones • Scientific and/or technical merit
Computational readiness review: LCF centers (2)	• Technical readiness • Appropriateness for requested resources	• Met technical/ computational milestones • On track to meet future milestones
Award Decisions	• INCITE Awards Committee comprised of LCF directors, INCITE program manager, LCF directors of science, sr. management	

Key questions to ask yourself

- Is both the scale of the runs **and** the time demands of the problem of LCF scale?
  - Yes, I can't get the amount of time I need anywhere else.
  - Yes, I my simulations are too large to run on other systems.
- Do you need specific LCF hardware?
  - Yes, the memory and I/O available here are necessary for my work.
- Do you have the people ready to do this work?
- Do you have a workflow?
- Do you have a post-processing strategy?
- Do you use ensemble runs and need LCF resources?
  - My ensembles can run under the direction of a complex script or large MPI job, with I/O scaling on a parallel file system. -> yes
  - My ensemble expects to run millions of serial jobs on nodes with local disk available. -> probably no

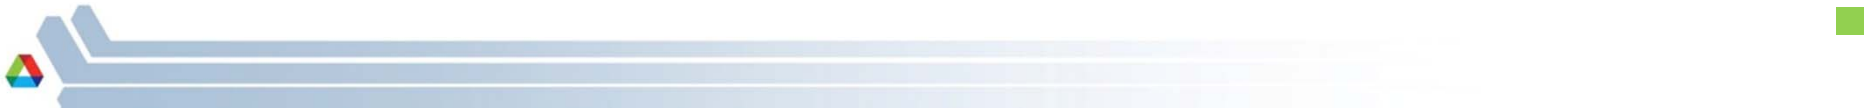

Scientific Support is Collaboration

The ALCF is staffed with a team of computational scientists, expert in their domain, scalable algorithms and performance engineering.

- Provide a "jump-start" in the use of ALCF resources
- Align the availability of ALCF resources with the needs of the project team
- Collaborate to maximize the value that ALCF can bring to our projects
- Connect the needs of the scientific community with future and current hardware

Two categories of collaboration and contribution to teams using the ALCF:

Tactical/Collaborative

- Short term, fast solutions
  - Compiling, Debugging, System Use
- Targeted problem resolution
  - Resolve a specific hard problem like restructuring I/O
- Long term collaborations
  - In depth work on a code that be over a long period of time
  - Constrained by staff

Strategic

- Training
  - Postdocs, students, community
- Understand HPC needs for different communities
- Plan for future needs
  - Help planning new facilities
  - Advise/Participate in long term code development paths

Transitioning Applications with New Hardware

What the LCFs offer

- Largest-scale computing resources
- Current and future hardware expertise
- Application level expertise and performance

What we need

- Engage future communities not just current ones
- Extract requirements and potential

Some Impacts From Approach

- Improved ease of use for current architectures
- Active efforts for improved workflows
- Planned significant efforts for data and it's relationship with exascale/extremescale

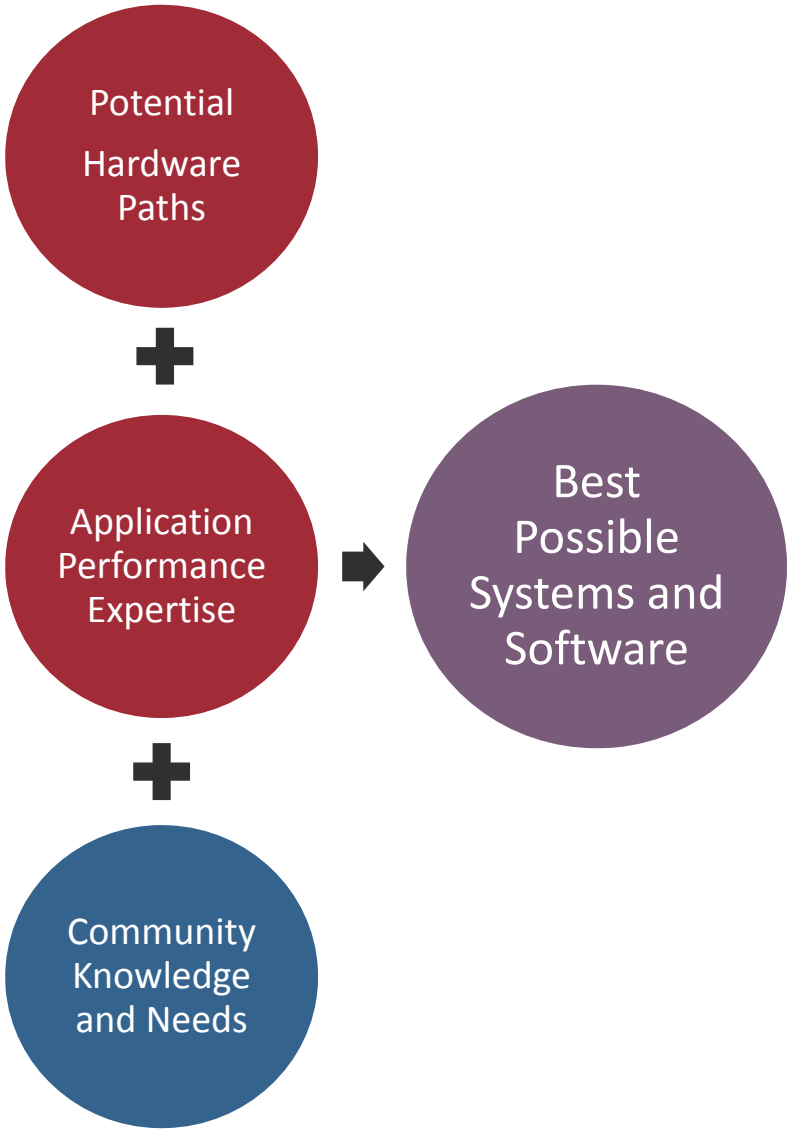

Strategies for the next 10 years

- The LCFs have 3 new systems planned, building to exascale
- Defining the platforms
  - Collaborating with vendors
  - Driving them based on what we think applications *can* do and *will* do
- Virtual Data Laboratory
  - Infrastructure for long distance, large scale, long term data sharing and storage
  - Integrated workflows into the computing resources
  - Included data management expertise and librarians in the facility
- Prepare the community
  - Preparation and experiments by the LCF staff and collaborations
  - Helping develop good techniques to move research forward

Prepare Scientific Applications for the Future

Two weeks, July 28 - August 9, 2013
The deadline for applying is May 22, 2013.

Training for the effective design and development of code for exascale

Renowned scientists, HPC experts, and leaders will lecture and guide the hands-on laboratory sessions.

- Programming methodologies effective for today's supercomputers and through exascale systems
- Multiple approaches on approaches for performance portability among current and future architectures
- Computer architectures, mathematical models, and numerical algorithms
- Community code techniques for HPC systems, and Big Data

Costs Covered, Chicago-land area

http://extremecomputingtraining.anl.gov

Explore and Impact the LCFs

- Identify possible projects for INCITE
- Engage us
  - Discretionary Allocations
  - Discuss needs!
- Embed the community in the ALCF
  - The ALCF has a continual pool of 9 post-docs, 3 per year
  - Long-term visitors are welcome for targeted research and collaboration
  - Staff hires - Given community demand, we can hire full or part-time staff members

Conclusion

- The LCFs want and need new communities
- If you think we you have a large scale problem, talk to us!
  - No DOE affiliation required

- Discretionary time:
- https://www.alcf.anl.gov/getting-started/apply-for-dd
  - Contact Katherine Riley or Paul Messina