

SciDAC at the ASCR Facilities: Today and Tomorrow

Organizers:

Rob Ross, Argonne National Laboratory

Jeff Candy, General Atomics

Karen Devine, Sandia National Laboratories

Gautam Bisht, Pacific Northwest National Laboratories

Jim Kowalkowski, Fermi National Accelerator Laboratory

Format: Two Interactive Panels

- **Panel 1: SciDAC Use of DOE Facility Resources**

Purpose: Discuss how Partnerships are taking advantage of facility resources, hurdles, and unusual challenges

Moderator: Rob Ross

- **Panel 2: Accelerators**

Purpose: Discuss the implications of diversity of accelerators in future systems and methods to best take advantage of them

Moderator: Jeff Candy

We'll switch at around 11:00, feel free to stretch your legs.

SciDAC Use of DOE Facility Resources: Panelists

Corey Adams (ALCF) is an Assistant Computer Scientist with a background in Neutrino Physics at Argonne National Lab at the Leadership Computing Facility, where he works on applications of machine learning and deep learning at scale.

Richard Gerber (NERSC) is NERSC HPC Department Head and Senior Science Advisor. He manages three groups at NERSC: Application Performance, Advanced Technologies, and User Engagement. He also oversees many aspects of the NERSC 9 (Perlmutter) project.

Martin Head-Gordon (LBL, BES) is a Senior Faculty Scientist at LBNL, and the Kenneth S. Pitzer Distinguished Professor in the Chemistry Department at Berkeley; he works on computational quantum chemistry and is PI of the SciDAC Partnership on Advancing Catalysis Modeling.

Stephen Price (LANL, BER) is with the fluid dynamics and solid mechanics group at LANL. His current work focuses on ice sheet and climate modeling. He is the lead PI on the SciDAC4 ProSPect project, was a PI on the SciDAC3 PISCEES project, and is the lead for E3SM's Cryosphere science focus area.

Phil Roth (OLCF) is a computer scientist working on performance analysis and optimization in the Scientific Computing Group of the Oak Ridge Leadership Computing Facility, and is a member of the RAPIDS SciDAC Institute with strong ties to two SciDAC application partnership projects.

Carl Sovinec (UW Madison, FES) is a professor in the Engineering Physics Department at the University of Wisconsin-Madison and has led the national NIMROD (non-ideal MHD with rotation) code team for more than a decade.

SciDAC Use of DOE Facility Resources: Topics

- **Resources and Time**

- Partnerships: How are you using? do you need more time? for what sorts of runs? how do you typically get time?
- Facilities: Best strategies for getting time? How do users "misuse" or "miss" on allocations?

- **Machine/Deep Learning, AI**

- Partnerships: In what ways are learning, a part of your use of facility resources?
- Facilities: Are you seeing an increase in learning applications, what impact is this having on your systems today, designs for the future?

- **Scientific Workflows, Large and Small**

- Partnerships: In what ways are more complex "workflows" a part of your use of facility resources? What's unusual or challenging about how you use the facilities, work towards your science goals?
- Facilities: Are you seeing an increase in more complex workflows? what impact is this having on your systems today, designs for the future? What compromises do you see yourselves having to make to cater to this variety?

- **Institutes**

- All: What are the SciDAC Institutes missing that would help them have the highest impact?

- **What did we miss?**

- We'd like to give our participants a chance to help us refine this plan before the meeting proper, but we will also open up for discussion with the audience on topics beyond the ones listed here.

Argonne Leadership Computing Facility

The Argonne Leadership Computing Facility provides world-class computing resources to the scientific community.

- Users pursue scientific challenges
- Resources fully dedicated to open science
- In-house experts to help maximize results

SCIDAC-4 Partnership Projects are encouraged to apply for a DD allocation for ALCF resources to get started. INCITE and ALCC are avenues to pursue for production science.



Theta

4,392 nodes
281,088 cores
69 TiB
MCDRAM

824 TiB DDR4
549 TB SSD
Peak flop rate:
11.69 PF



October 1-3, 2018



TBA : ~May 2020

Training Programs

Martin Head-Gordon (LBNL)

Representing the BES partnership on
“Advancing catalysis modeling”



1st focus (very facility-relevant): Production atomistic simulations using a diverse software stack (NERSC).

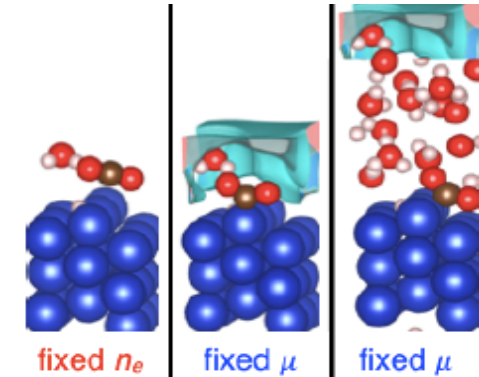
DFT codes (CP2K, VASP, Q-Chem)
Dynamics codes (LAMMPS)

Often augmented in workflows with extensions required for:

Advanced sampling
Nuclear quantum effects


2nd focus (not so facility-relevant): Pioneering new theory, algorithms and software for electronic structure, embedding, statistical mechanics and dynamics.

Typically uses mid-range local computers and clusters.



General comments:

- ✓ NERSC enables production calculations at a scale that is simply impossible with group-level mid-range computing. Critical for US simulation science competitiveness.
- ✓ Some challenges in efficiently using large allocations due to overall machine load and queue times.
- ✓ Ability to interact with RAPIDS experts (Ibrahim & Williams) is very helpful (we wish our budget could pay them more!).



ProSPect: Bounding Future Sea-Level Rise from Earth's Ice Sheets

Resources & Time

- Computing time generally obtained through annual ERCAP applications
- Because BER (climate) SciDAC projects are part of E3SM “ecosystem” some computing needs can piggy-back on E3SM computing resources
- ALCC and/or INCITE are more difficult to compete for if large amounts of computing time needed
- Long-term storage: could use long-term, semi-permanent storage for input and analysis datasets

Machine Learning & AI



- Potential for ProSPect to be heavy user *if* successful at improving emulator fidelity (to full phys sims) &/or reducing cost of emulators (design and use) in UQ workflow

Scientific Workflows

- non-simulation (=fwd model) workflows include optimization and UQ
- computing time for these is current small (not trivial) but will continue to increase as the project matures

Institutes


- Continued support for optimization and UQ needs associated with high-dimensional parameter spaces
- Support for particle methods (DEMSI – A. Turner)



Antarctic ice sheet 200 years after all floating ice shelves are removed. Shown are simulation results from *ProSPect* MALI (top) and BISICLES (bottom) ice sheet models.

View from the Oak Ridge Leadership Computing Facility

- User facility role taken seriously, with strong competition for access
- To improve likelihood of success:
 - Target appropriate program
 - Demonstrate readiness (e.g., existing ability to use GPUs) and plans to scale
 - Include justification (“show your work”)
- We want to “get to yes”
 - Follow online submission guidelines
 - Contact us for help, even before submitting
 - Director of Science Jack Wells, wellsjc@ornl.gov; SciComp Group Leader Judy Hill, hilljc@ornl.gov
- Be architecturally aware
 - Summit (POWER9 with NVIDIA GPUs) v. Frontier (x86_64 with AMD GPUs)
 - Portability with performance requires forethought, value judgement
- Conversation is ongoing about hosting SciDAC Institute software



Carl Sovinec, University of Wisconsin-Madison / FES / Center for Tokamak Transient Sim.

- Fusion requires a variety of computation; perspectives shared here pertain to 3D macroscopic stability.
- Computations deal with multiple spatial scales, multiple temporal scales (propagating), nonlinearity, evolving anisotropy.
 - Advanced numerical methods and algebraic solvers are critical.
 - Interaction with applied math groups pre-dates SciDAC, but SciDAC partnering led to major performance boosts.
- Theoretical models are still being developed; equations are not frozen.
- Concerns:
 - Codes and methods are complex; dealing with architectural changes are difficult.
 - NESAP training support will not be sufficiently inclusive; trickle-down approach for those not selected has long-term consequences.

SciDAC Use of DOE Facility Resources: Topics

- **Resources and Time**

- Partnerships: How are you using? do you need more time? for what sorts of runs? how do you typically get time?
- Facilities: Best strategies for getting time? How do users "misuse" or "miss" on allocations?

- **Machine/Deep Learning, AI**

- Partnerships: In what ways are learning, a part of your use of facility resources?
- Facilities: Are you seeing an increase in learning applications, what impact is this having on your systems today, designs for the future?

- **Scientific Workflows, Large and Small**

- Partnerships: In what ways are more complex "workflows" a part of your use of facility resources? What's unusual or challenging about how you use the facilities, work towards your science goals?
- Facilities: Are you seeing an increase in more complex workflows? what impact is this having on your systems today, designs for the future? What compromises do you see yourselves having to make to cater to this variety?

- **Institutes**

- All: What are the SciDAC Institutes missing that would help them have the highest impact?

- **What did we miss?**

- We'd like to give our participants a chance to help us refine this plan before the meeting proper, but we will also open up for discussion with the audience on topics beyond the ones listed here.