Berkeley Lab Computing Strategic Directions

CCS - LBNL Collaborative Workshop 12 May 2016

David L Brown, Director
Computational Research Division
Lawrence Berkeley National Laboratory

Much of Berkeley Lab Computing relocated to Wang Hall in late 2015

- **Four floors, 13,000 m²**
  - 300 offices,
  - 1,850m² HPC floor,
  - 12.5->40 MW
- **Promotes collaboration:**
  - LBNL, NERSC, CRD, Esnet, UCB
- **Exceptional energy efficiency**
  - Natural air and water cooling
  - Heat recovery
  - PUE < 1.1

Computing Sciences at Berkeley Lab in 2016

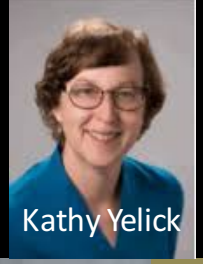
Kathy Yelick

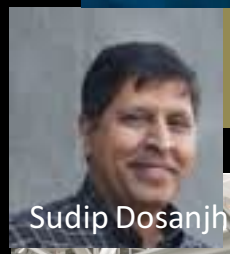
Sudip Dosanjh

NERSC


David Brown

Computational Research

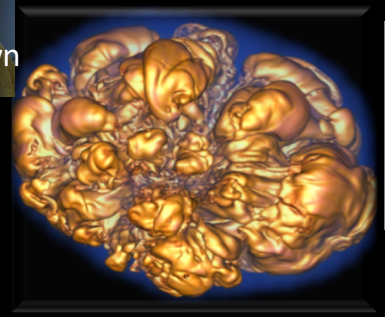

Computational Science

Applied Mathematics

$$\frac{dVar(t)}{dt} =$$
$$\int_{s_2(t)}^{s_1(t)+S} (-(g^{-1}F_s)_s - K^2Fg + gK^2F)ds - \int_{s_2(t)}^{s_1(t)+S} (-(g^{-1}F_s)_s - K$$
$$= -\int_{s_1(t)}^{s_2(t)} (g^{-1}F_s)_s \, ds + \int_{s_2(t)}^{s_1(t)+S} (g^{-1}F_s)_s \, ds$$
$$\left[g^{-1}F_s \mid_{s_2(t)} - g^{-1}F_s \mid_{s_1(t)}\right] + \left[g^{-1}F_s \mid_{s_1(t)+S} - g\right.$$
$$= -2\,(g^{-1}F_K K_s) \mid_{s_2(t)} + 2\,(g^{-1}F_K K_s) \mid_{s_1(t)}$$


Inder Monga (A)

ESnet

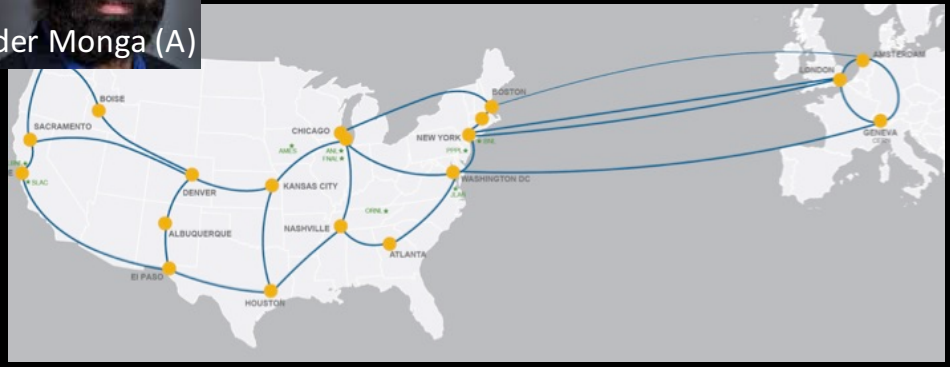

Computer Science

Data Science & Technology

Berkeley Lab Initiatives for Computing in the Next Decade

Enable transformative science through partnerships

Maximize science impact and prepare for post-Exascale

Transform experimental and observational science

Enable more complex science with mathematics

Effect broad impact through facilities

NERSC

ESnet
ENERGY SCIENCES NETWORK

Co-Invent and deploy Exascale systems and applications that accelerate science by 2025

Maximize the scientific impact of the Exascale program

Deploy Exascale Systems at NERSC

NERSC: broadest, most open, widely used computing center in US DOE

5000+ users, 1900+ publications, 700 codes, 5 associated Nobels

Cori is Exascale Roadmap system; balanced energy and productivity

Beyond 2025: Exploring Quantum and Neuromorphic Computing

Quantum Simulation for Materials, Chemistry and Physics

Multiple quantum device technologies at Berkeley

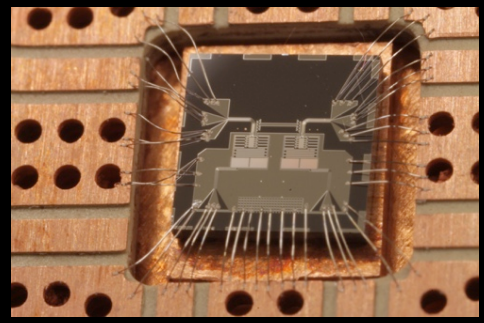

Siddiqi's Quantum Circuit

Experimentally implement chemical simulations on 3-qubit platform

Develop model for small user facility to explore device technology

THE HAMILTONIAN LANDSCAPE FOR QUANTUM SIMULATION

Neuromorphic computing provides low power options for data analysis

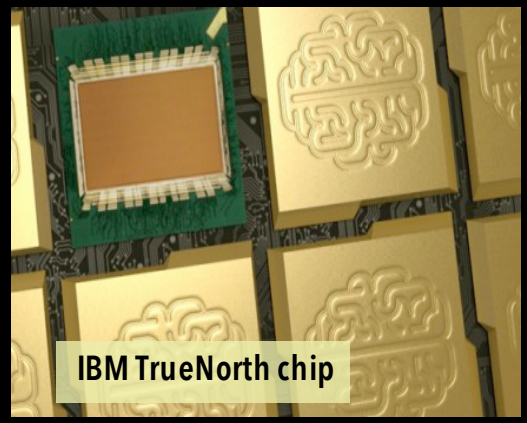

IBM TrueNorth chip

Low-power, real-time **data analysis** in materials, biology and cosmology

Studying image analysis and event detection with Kalman filters

Understand performance and power advantages, and limits on generality

Use Math and CS Leadership at Berkeley Lab to Maximize Exascale Science

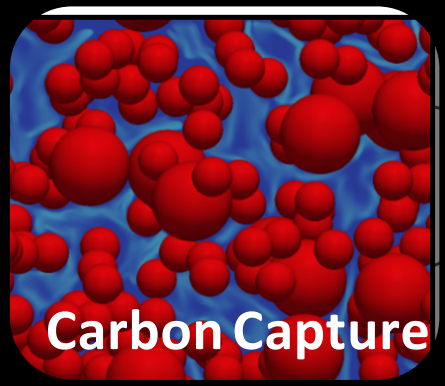

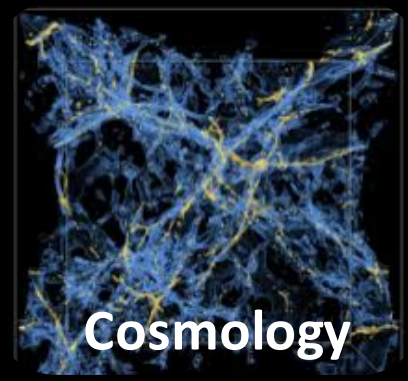

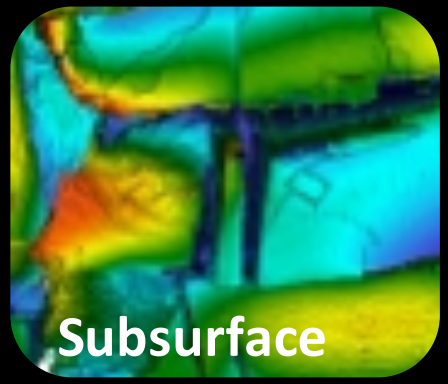

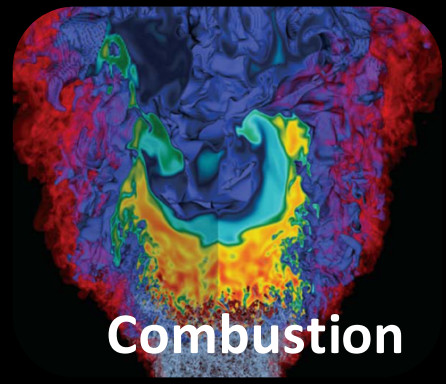

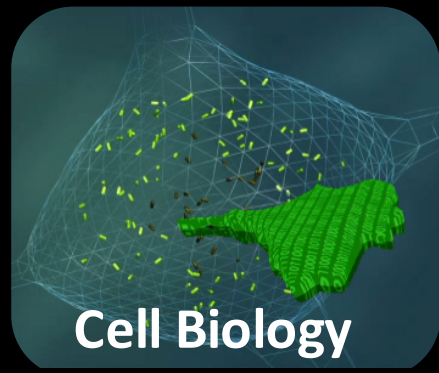

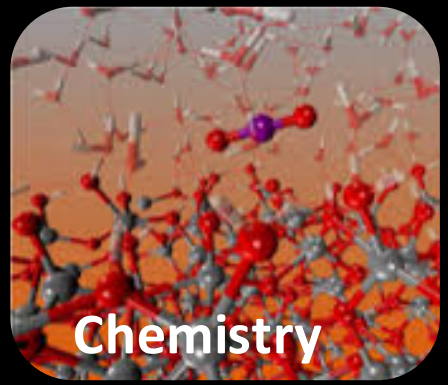

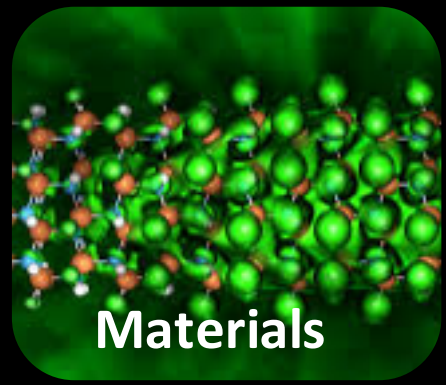

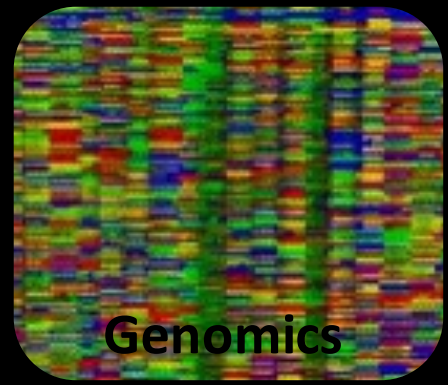

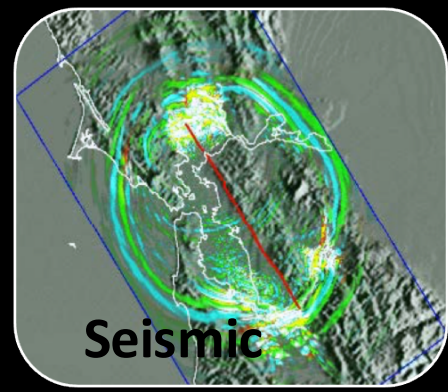

Many of our Science Applications use Berkeley Lab Adaptive Mesh Refinement (AMR) mathematics and software

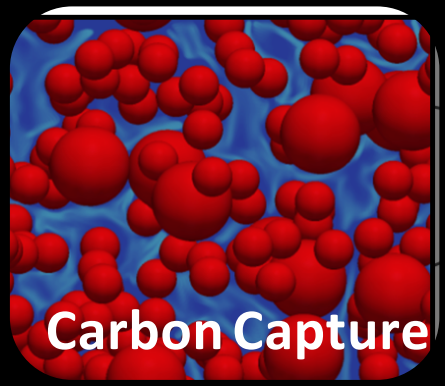

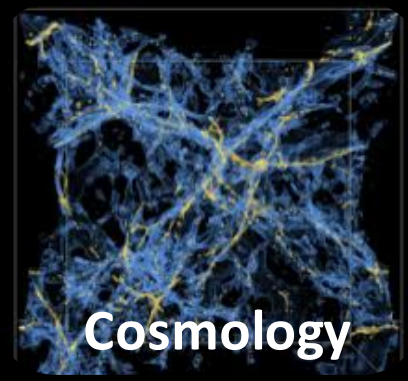

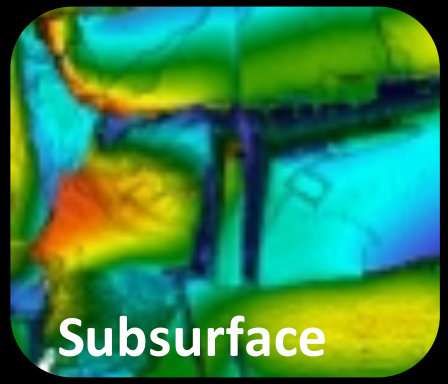

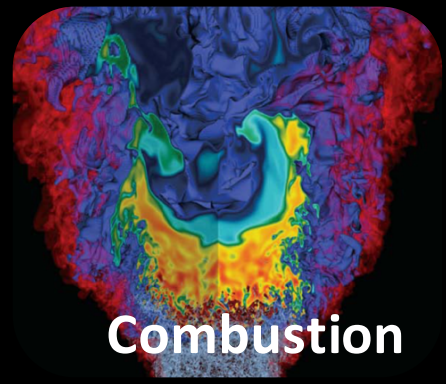

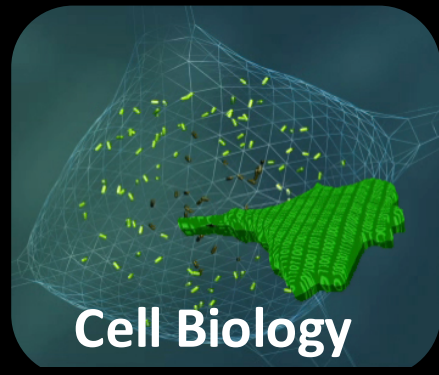

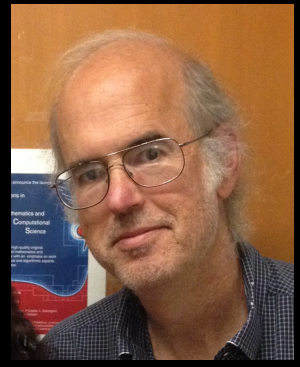

Transform experimental and observational science through advanced math, computation and networking

Center for Advanced Mathematics for Energy Research Applications (CAMERA)

CAMERA Applied Math

Today: Data analysis time-consuming.

Tomorrow: More data faster. More resolution.

Critical need: algorithms/analysis for *understanding*

New math: Transform experimental data into understanding

-Extract information from murky data
-Interpret, and optimize experiments
-Accelerate scientific discovery

computing *structure* from imaging → analyzing samples and proposed new materials → designing new materials

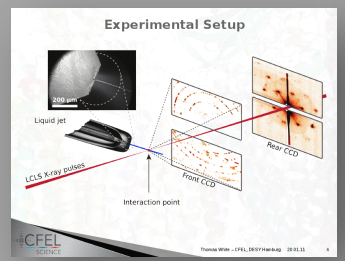

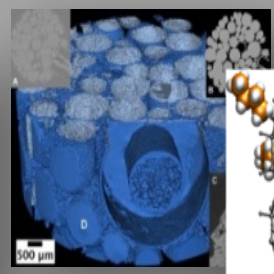

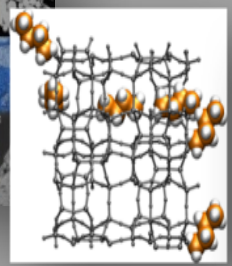

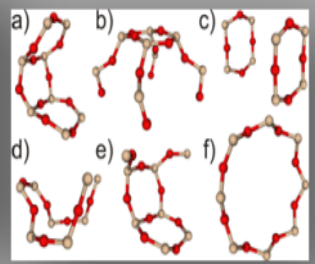

Partners

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB Lawrence Berkeley National Laboratory

ALS ADVANCED LIGHT SOURCE

Advanced Photon Source Argonne National Laboratory

SSRL

LCLS

NERSC Provides HPC and Data Resources for DOE Office of Science Research

U.S. DEPARTMENT OF ENERGY | Office of Science

Largest funder of physical science research in U.S.

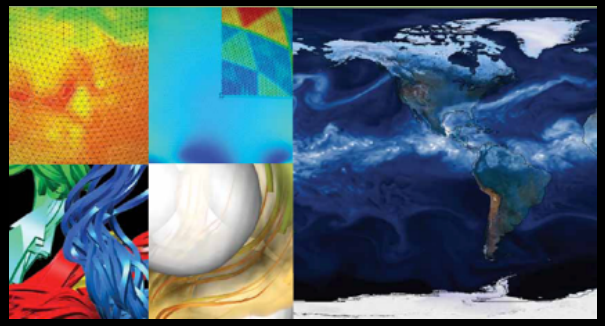

Biology, Environment

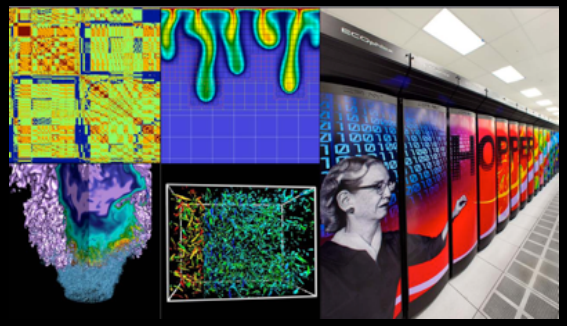

Computing

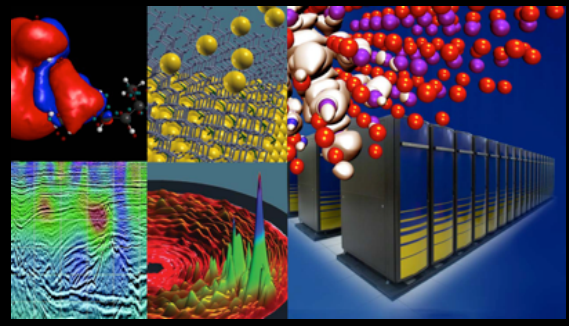

Materials, Chemistry, Geophysics

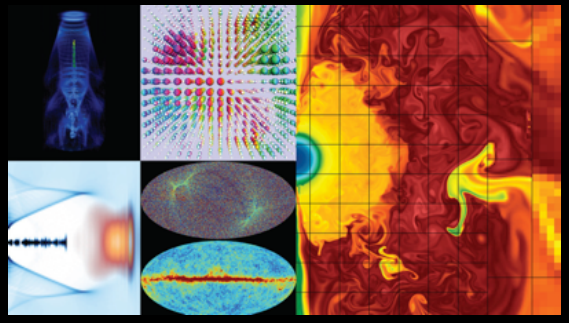

Particle Physics, Astrophysics

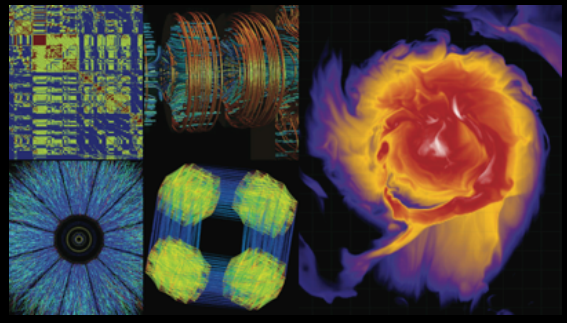

Nuclear Physics

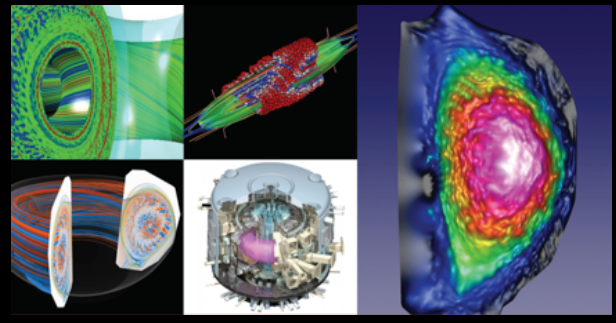

Fusion Energy, Plasma Physics

NERSC's newest machine Cori supports both the HPC Workload and Data-Intensive Science

- **Cray system with 9,300 Intel Knights Landing compute nodes – deliver in mid 2016**
  - Self-hosted, (not an accelerator) manycore processor > 64 cores per node
  - On-package high-bandwidth memory at >400GB/sec
- **Robust Application Readiness Plan**
  - Outreach and training for user community
  - Application deep dives with Intel and Cray
  - 8 post-docs integrated with key application teams
- **Data Intensive Science Support**
  - 10 Haswell processor cabinets (Phase 1) to support data intensive applications
  - NVRAM Burst Buffer with 1.5PB of disk and 1.5TB/sec
  - 28 PB of disk, >700 GB/sec I/O bandwidth in Lustre bandwith

NERSC is making significant investments on Cori to support data intensive science

- Haswell- 128GB memory nodes
- Support for real time and high-throughput queues
- Containers: Virtualization capabilities (Docker)
- SPARK software stack for data
- NVRAM Flash Burst Buffer as I/O accelerator
- High bandwidth external connectivity to experimental facilities from compute nodes through Software Defined Networking
- More (23) login nodes for managing advanced workflows

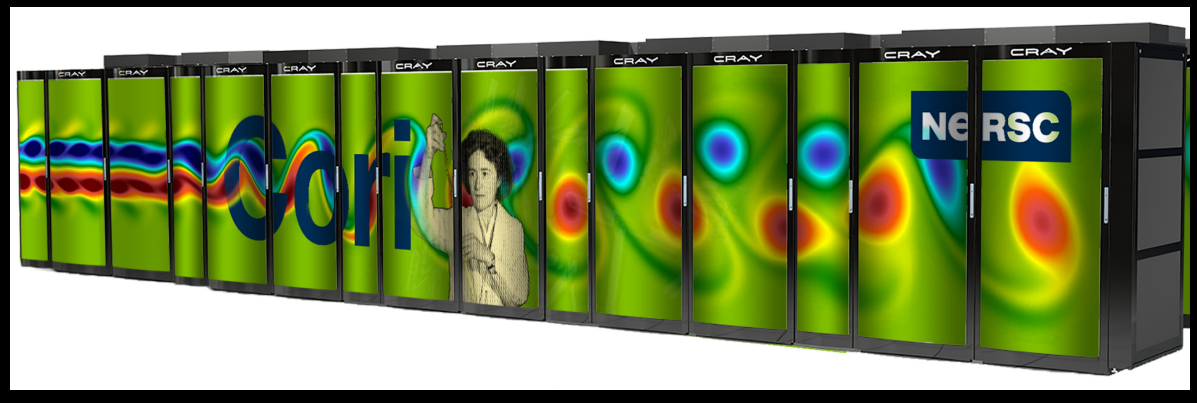

NERSC is exploring Burst Buffer Use Cases beyond checkpoint-restart

- **Accelerate I/O**
  - Checkpoint/restart or other high bandwidth reads/writes
  - Apps with high IOP/s e.g. non-sequential table lookup
  - Out-of-core applications
  - Fast reads for image analysis
- **Advanced Workflows**
  - Coupling applications, using the Burst Buffer as interim storage
  - Streaming data from experimental facilities
- **Analysis and Visualization**
  - In-situ/ in-transit
  - Interactive visualization

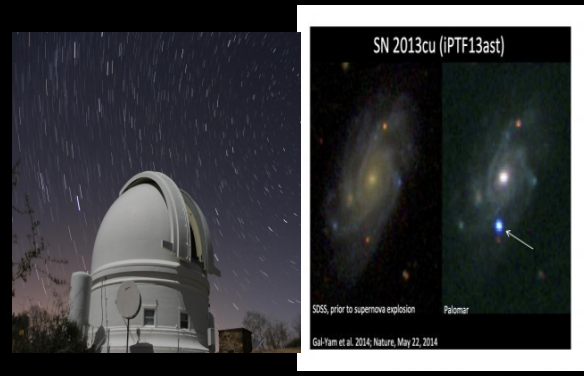

Palomar Transient Factory Pipeline:
Use Burst Buffer as cache for fast reads

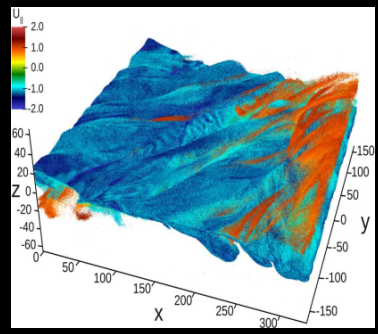

VPIC – in situ visualization of a trillion particles

Upgrading Cori's External Connectivity

LCLS

Enable 100Gb+ Instrument to Cori

- Streaming data to the supercomputer allows for analytics on data in motion
- Cori network upgrade provides SDN (software defined networking) interface to ESnet. 8 x 40Gb/s bandwidth.
- Integration of data transfer and compute enables workflow automation

Cori Network Upgrade Use Case:

- X-ray data sets stream from detector directly to Cori compute nodes, removing need to stage data for analysis.
- Software Defined Networking allows planning bandwidth around experiment run-time schedules
- 150TB bursts now, LCLS-II has 100x data rates

Envisioning the future of large scale experimental science

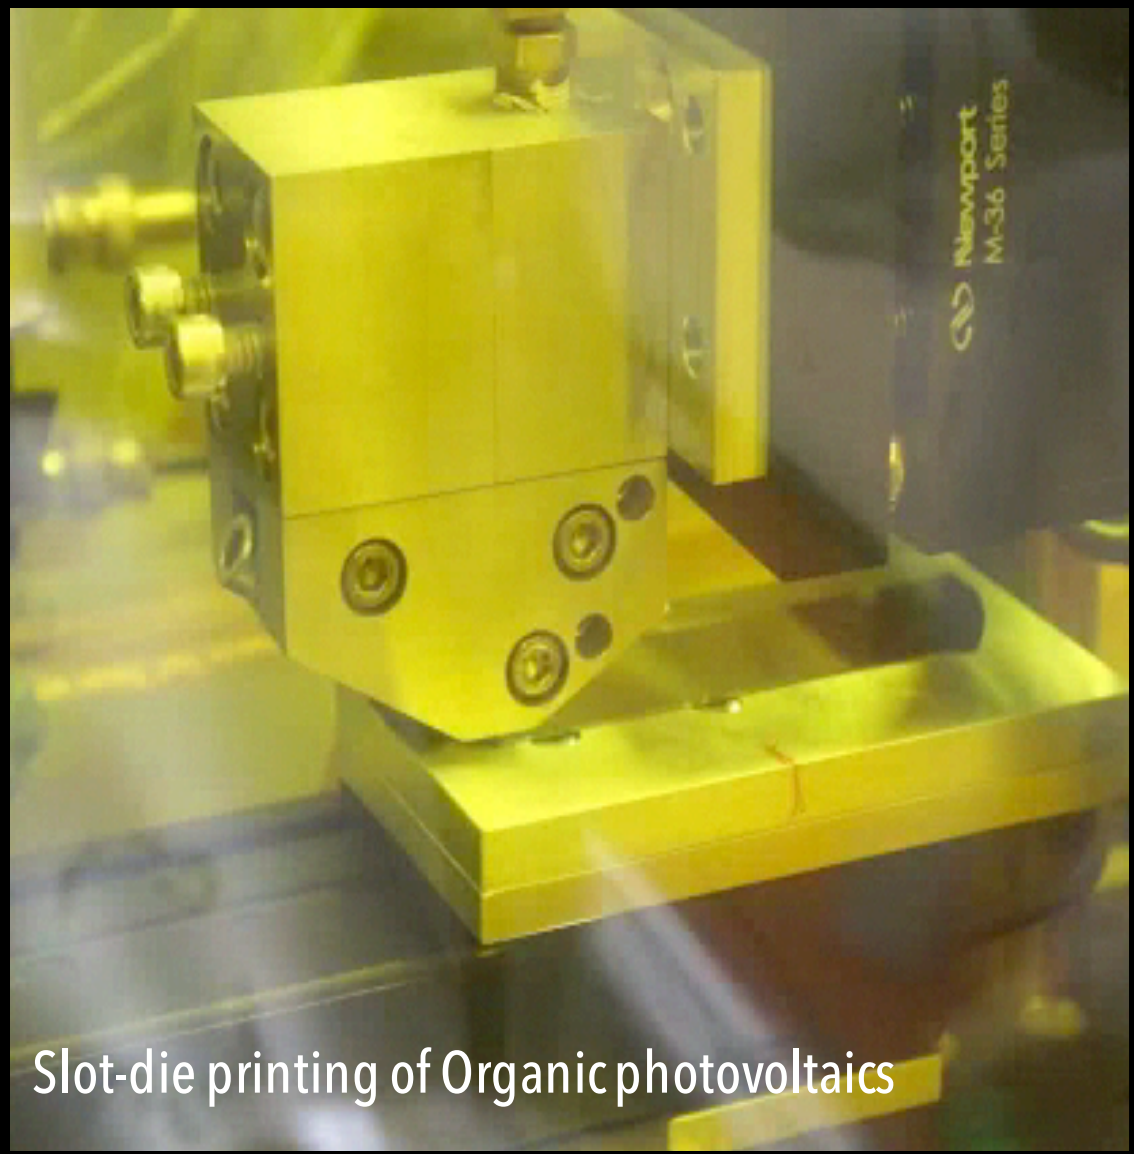

Slot-die printing of Organic photovoltaics

Vision: Scientists in 2025 will access experimental and computational capabilities through a "Superfacility" that seamlessly integrates experiments, networks and HPC

ALS ADVANCED LIGHT SOURCE

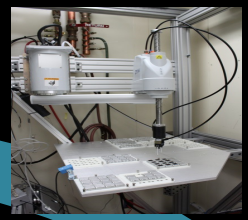

ESnet

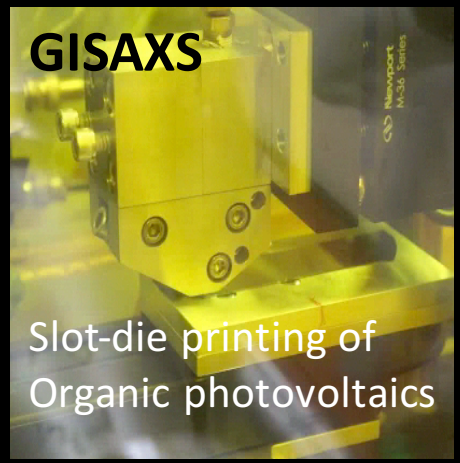

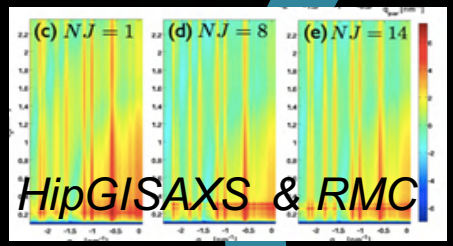

CAMERA Applied Math

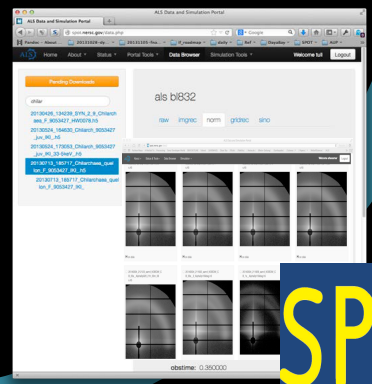

SPOT

Integrated Super-Facility Dashboard prototype

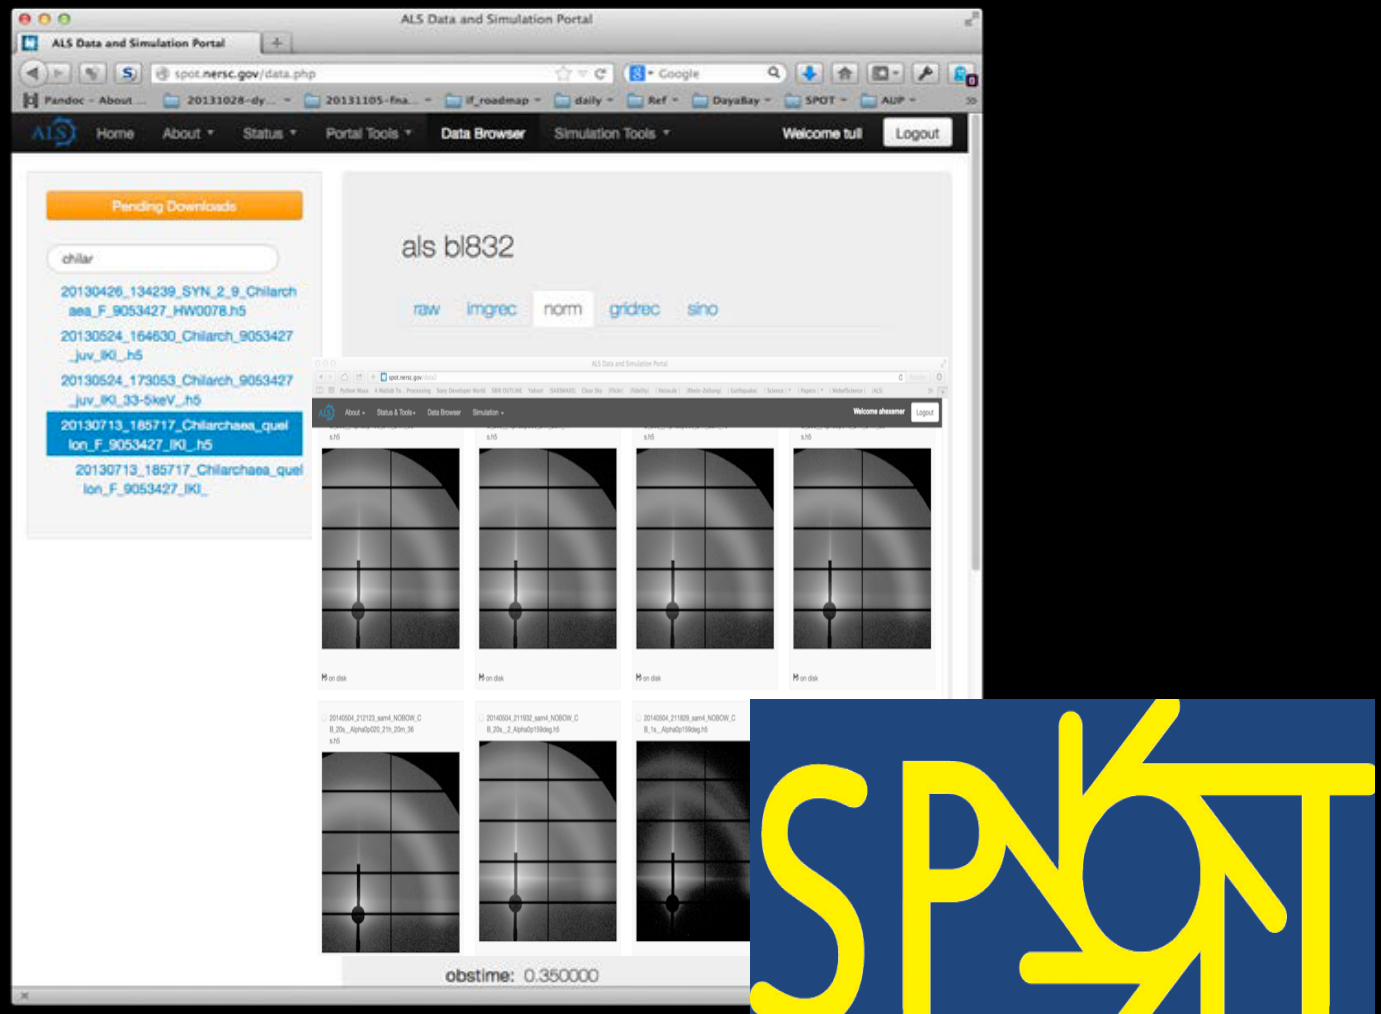

Science in embedded sensors: Internet of Things

Transportation Modeling

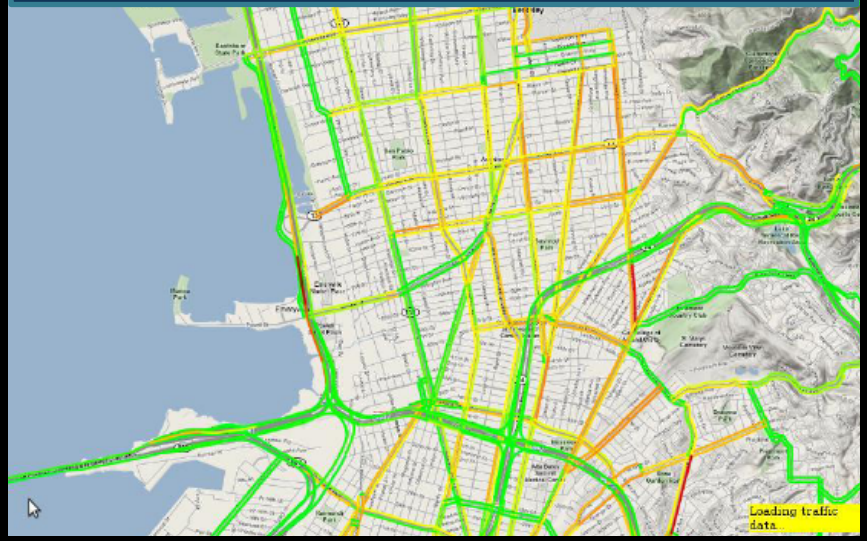

Power Grid Modeling

Scenario Prediction, Planning

Decision Science

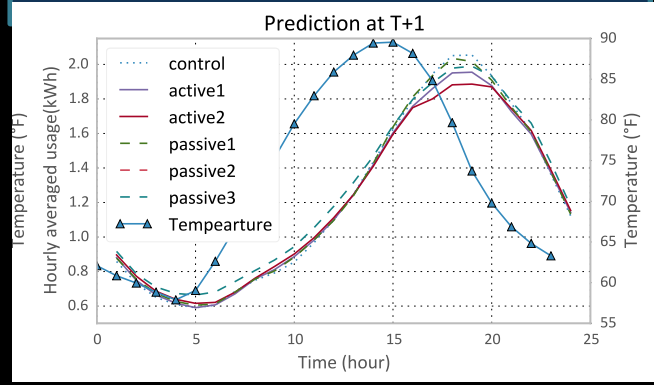

Handle increasingly complex science questions through better models and smarter algorithms

Mathematize new science problems

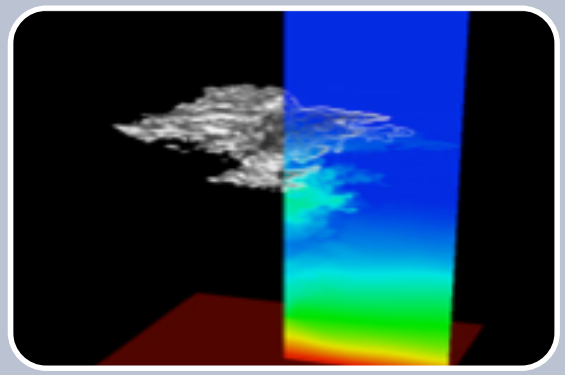

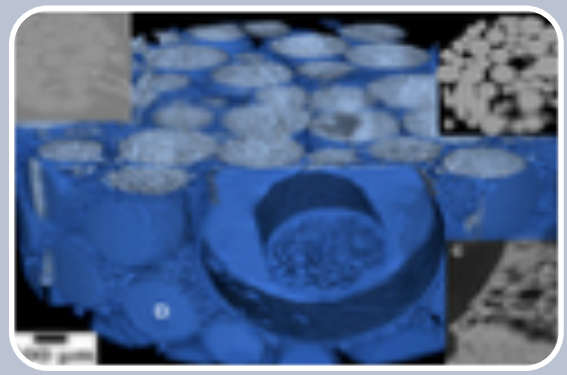

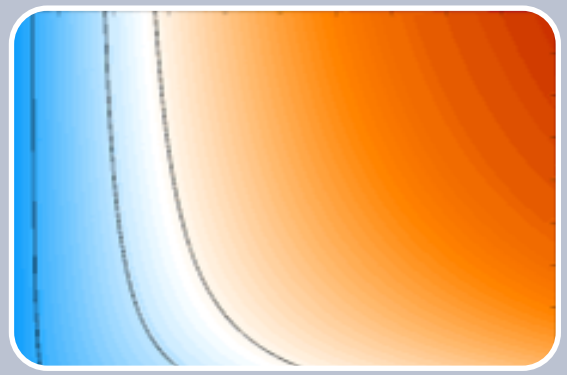

Mathematical models enabled by resolution and science	Advanced mathematics for experiment and observation	Mathematics at the boundary of simulation and observation

Complexity: Multi-physics models, need for multiscale methods, multi-modal and noisy data

Deliver Mathematics through new Partnership Models and Advanced Libraries

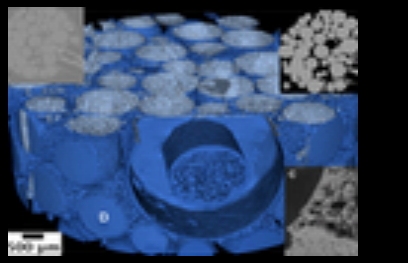

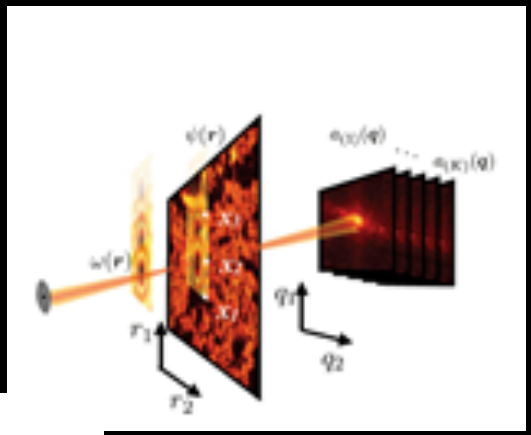

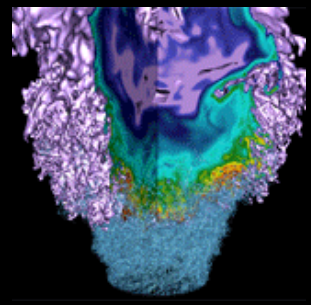

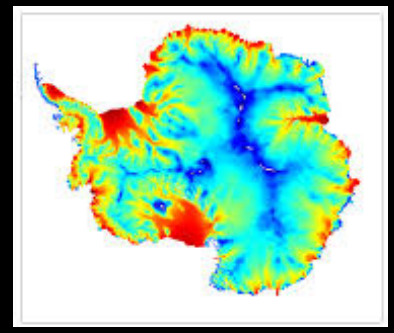

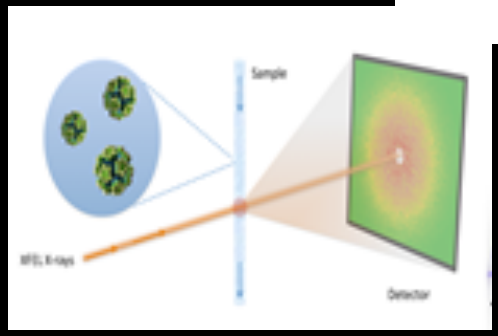

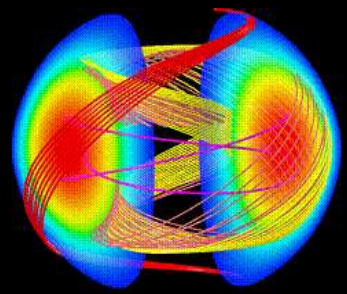

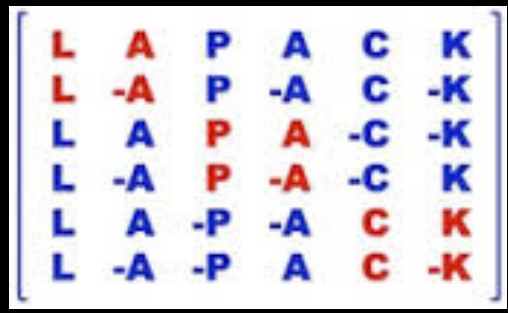

Expand CAMERA “facility” model of math collaboration **throughout experimental science**

Make advanced mathematics e accessible by the wide science through **software tools and lib**

In many areas, there are opportunities to combine simulation and observation for new discoveries

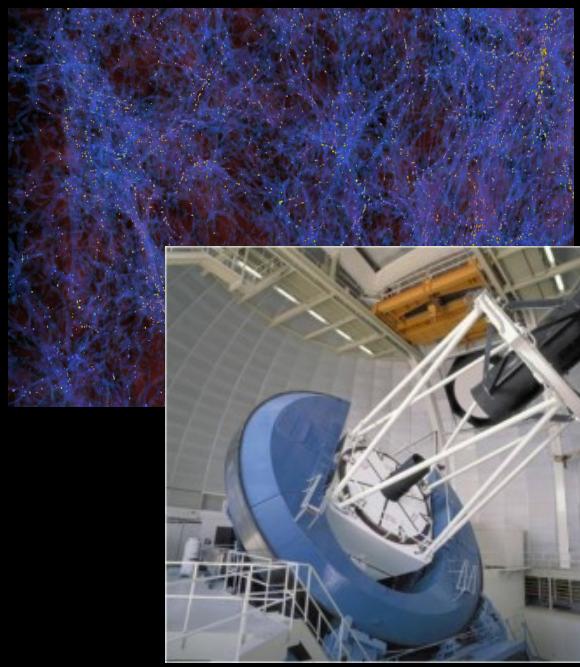

Cosmology

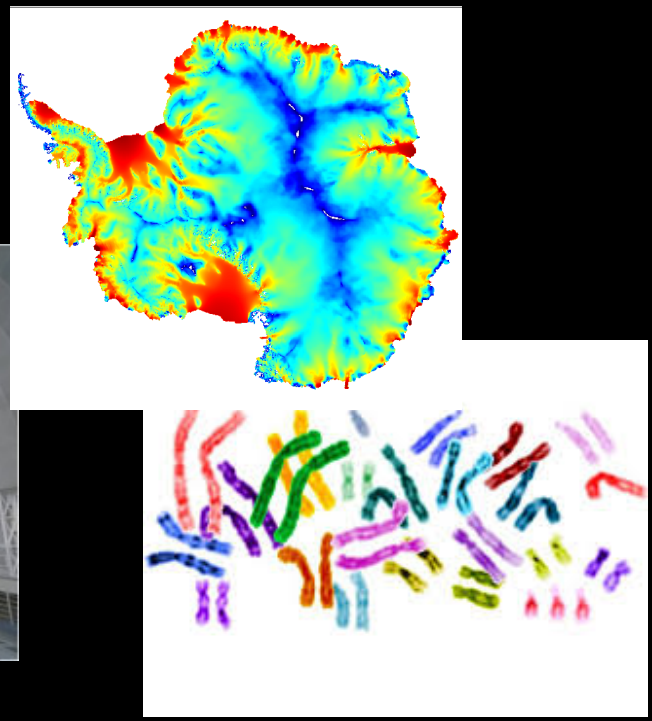

Environment

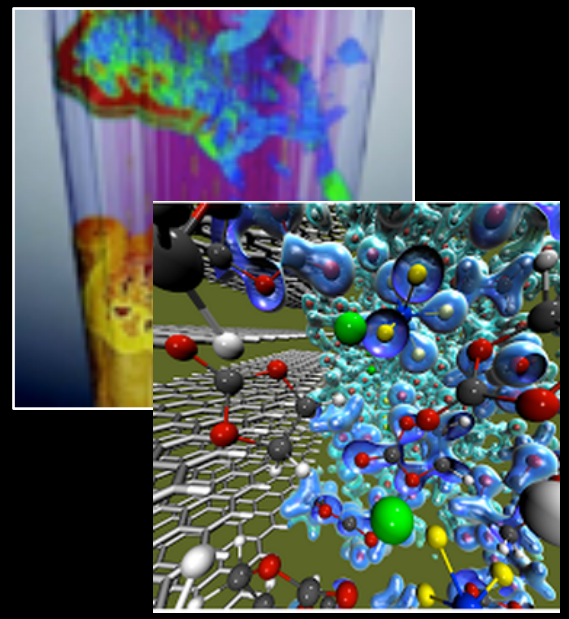

Materials

Our strategy leverages all of Berkeley Lab's strengths in Computing

ESnet: Unique and Essential

Applied Math: Strongest Program in DOE

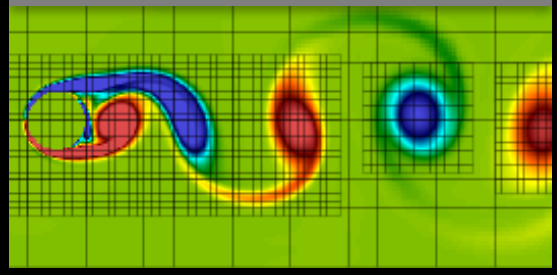

NERSC: More Science Impact than any HPC Facility

Integrated Science Partnerships

Computer and Data Science: Strength in Basic Research

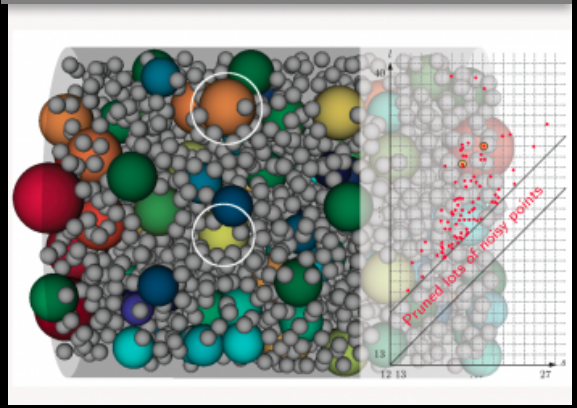

Software: Production Development and Support Expertise

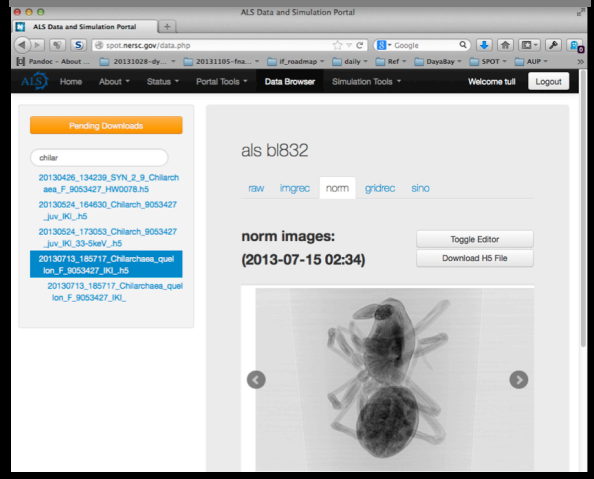

Berkeley Lab Computing Strategic Directions

CCS - LBNL Collaborative Workshop 12 May 2016

David L Brown, Director
Computational Research Division
Lawrence Berkeley National Laboratory