ALCF AURORA EARLY SCIENCE PROGRAM FOR DATA AND LEARNING

TIMOTHY J. WILLIAMS
Deputy Division Director, Computational Science Division
Manager, ALCF Early Science Program

Argonne Leadership Computing Facility

ALCF EARLY SCIENCE PROGRAM

ALCF APPLICATIONS READINESS: EARLY SCIENCE PROGRAM

- Mira ESP, Theta ESP gave us
  - Breakthrough science
  - Technical reports on code porting & tuning
  - Open community workshops (science & technology)
  - Synergy with Tools & Libraries project
  - Stable production platform (problems shaken out)
  - Persisting culture of apps readiness for next generation
  - Success stories for postdocs

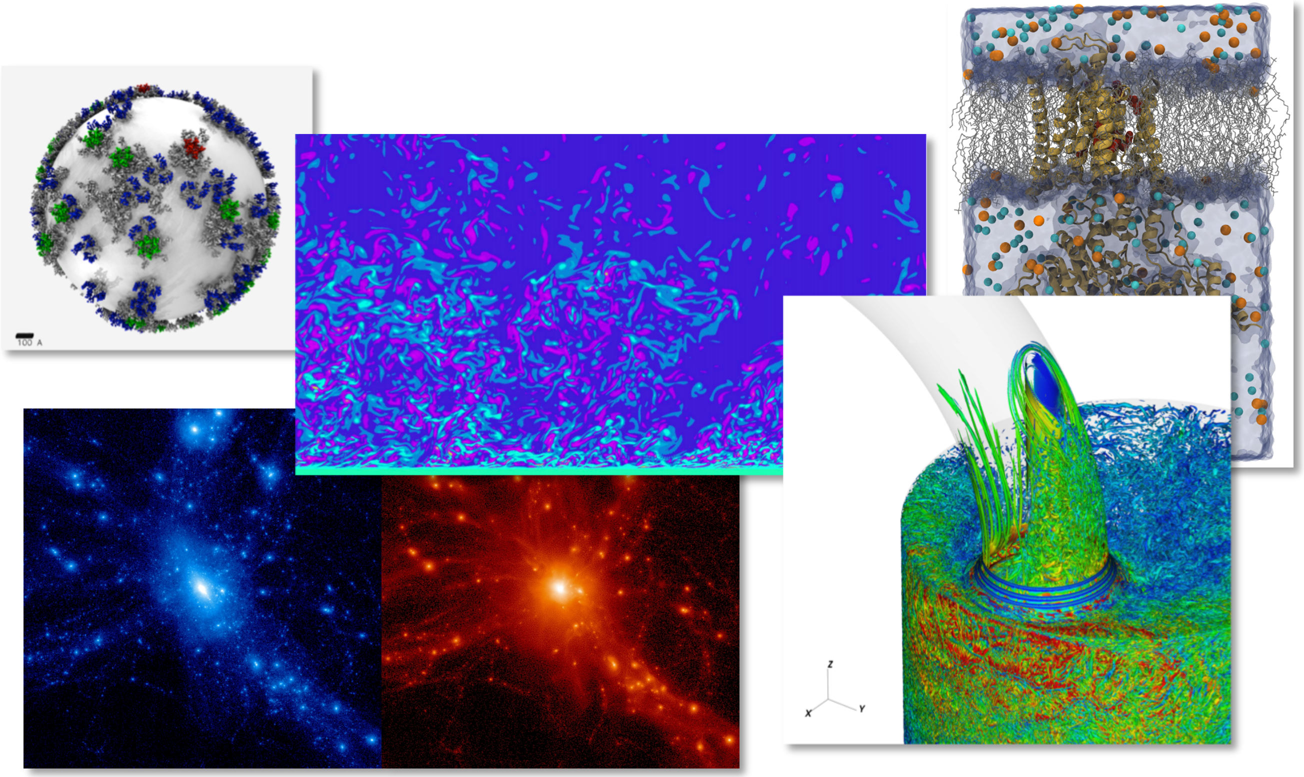

Argonne NATIONAL LABORATORY ANL/ALCF/ESP-13/17

ALCF-2 Early Science Program Technical Reports

Compendium of Individual Technical Reports for the 16 ESP Projects

Argonne Leadership Computing Facility

Argonne NATIONAL LABORATORY ANL/ALCF/ESP-17/1

Next-Generation Cosmology Simulations with HACC: Conquering the Baryon Problem

Technical Report for the ALCF Theta Early Science Program

Argonne Leadership Computing Facility

Argonne NATIONAL LABORATORY ANL/ALCF/ESP-17/

Free Energy Landscapes of Membrane Transport Proteins

Technical Report for the ALCF Theta Early Science Program

Argonne Leadership Computing Facility

Argonne NATIONAL LABORATORY ANL/ALCF/ESP-17/1

The Hadronic Contribution to the Anomalous Magnetic Moment of the Muon

Technical Report for the ALCF Theta Early Science Program

Argonne Leadership Computing Facility

ALCF EARLY SCIENCE PROGRAM (ESP)

Applications Readiness

- Prepare applications for next-gen system:
  - Architecture
  - Scale
- ~Two year lead time

Proposals

- Ambitious targeted science calculation
- Parallel performance
- Development needed
- Team

http://esp.alcf.anl.gov

Support

PEOPLE

- Funded ALCF postdoc
- Catalyst staff member support
- Vendor experts

TRAINING

- Training on HW and programming
- Community workshop to share lessons learned

COMPUTE RESOURCES

- Current ALCF systems
- Early next-gen hardware & simulators
- 3 months dedicated Early Science access
  - Pre-production (post-acceptance)
  - Large time allocation
  - Continued access for rest of year

ESP APPLICATION EFFORTS

Balance of optimization, scaling, development

Next-Generation Cosmology Simulations with HACC: Challenges from Baryons

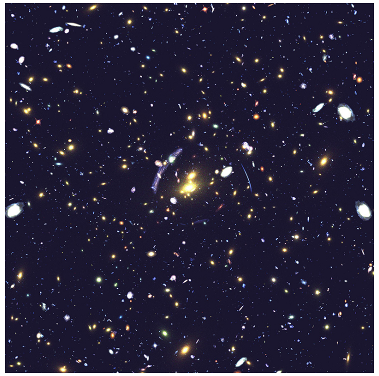

Code: HACC

PI: Katrin Heitmann (ANL)

N-body gravity + SPH hydro

Catalysts: H. Finkel, A. Pope

Postdoc: J.D. Emberson

- Tune kernel
- Develop CRK-SPY hydro
- Develop subgrid models

First-Principles Simulations of Functional Materials for Energy Conversion

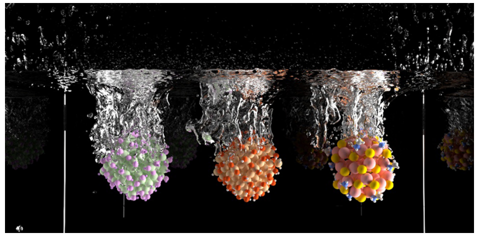

Codes: WEST & Qbox

PI: Giulia Galli (U. Chicago)

MBPT & ab initio MD

Catalyst: C. Knight

Postdoc: H. Zheng

- Use the right optimized libraries
- Address scaling
  - Optimize communication
  - Add 3rd parallelism layer to WEST

Free Energy Landscapes of Membrane Transport Proteins

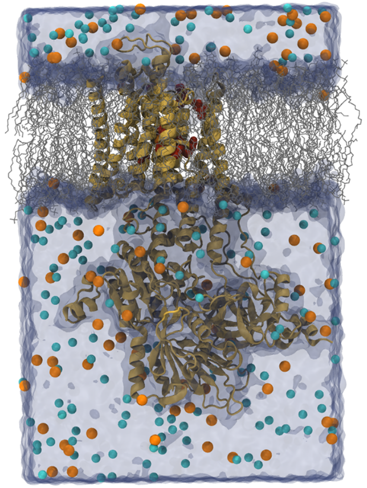

Code: NAMD

PI: Benoit Roux (U. Chicago, ANL)

MD with replica methods

Catalyst: W. Jiang

Postdoc: B. Radak

- Tune Charm++
- Develop constant-pH
- Develop statistical models

The Hadronic Contribution to the Anomalous Magnetic Moment of the Muon

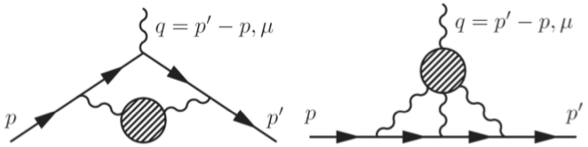

Codes: MILC & CPS

PI: Paul Mackenzie (FNAL)

Lattice QCD

Catalyst: James Osborn

- Developed/tuned KNL code
- Studied communication issues

FIRST-PRINCIPLES SIMULATIONS OF FUNCTIONAL MATERIALS FOR ENERGY CONVERSION

Case Study: Theta ESP application optimization

Codes: WEST, Qbox
PI: Giulia Galli (U. Chicago)
MBPT & ab initio MD
Catalyst: Chris Knight
Postdoc: Huihuo Zheng

Science Impact

- Properties of materials to be used for solar and thermal energy conversion will be optimized at an unprecedented level of accuracy—by combining ab initio molecular dynamics and post-density functional theory methods—thus providing truly predictive tools, ultimately for device performance, within a MGI material design framework.

Numerical Methods/Algorithms

- WEST implements Many Body Perturbation Theory at the GW level, starting from DFT inputs obtained either using GGA or hybrid functionals.
- Qbox is an ab initio molecular dynamics code. It is based on DFT and a plane wave basis.

Application Development

- Qbox: optimize FFT (MKL library), evaluate running entirely in MCDRAM
- WEST: vectorization, optimized linear algebra kernels such as DGEMM from MKL

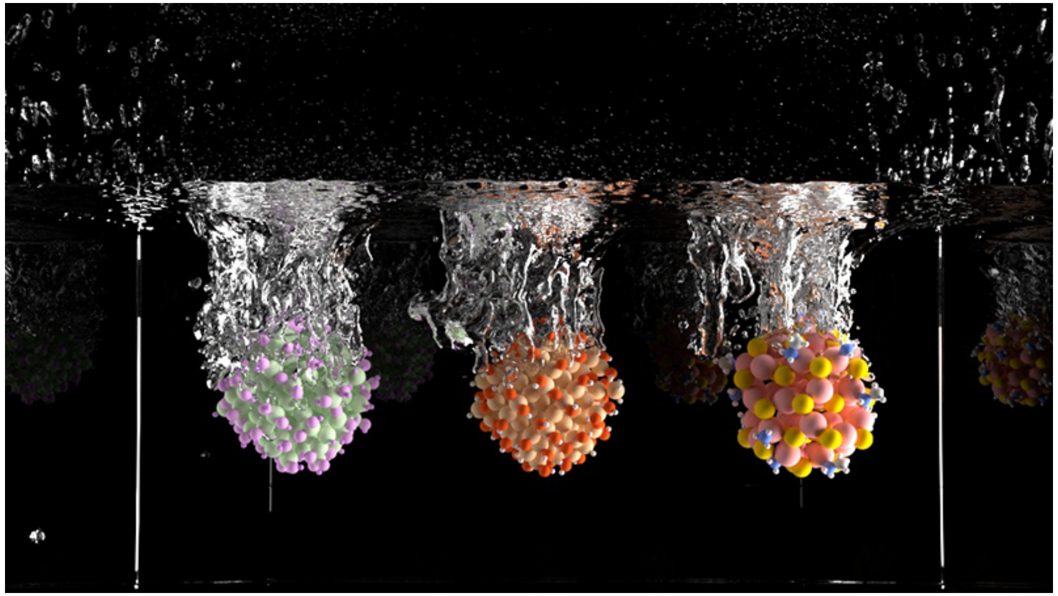

This project will focus on high-performance calculations of nanoparticles and aqueous systems for energy applications. Nicholas Brawand, Institute for Molecular Engineering, University of Chicago

QBOX OPTIMIZATION

SINGLE NODE

- Strong dependence on libraries
  - Linear algebra
  - FFT
- Replace ScaLAPACK eigenvalue solver with ELPA
  - **5-10X** speedup
- Intel MKL for FFT
- After these optimizations, focus on scaling/multimode issues

PROCESSOR REMAPPING

- Different optimal processor maps for linear algebra, FFTs
  - Swap between two layouts
  - Implement custom “gather & scatter” remap
    - 1000X faster than `pgemr2d`

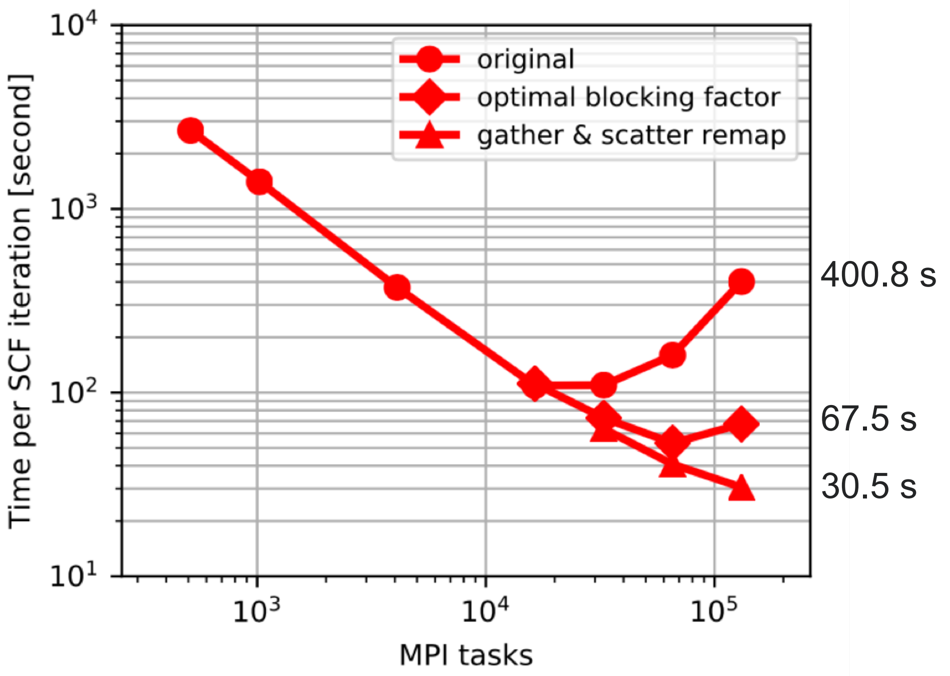

Zheng et al., IXPUG Fall 2017, APS March Meeting 2017

WEST OPTIMIZATION

NEW PARALLELISM LAYER

- Across bands N_b (in addition to across perturbations & plane waves $N_v \times N_z$)
- Good prognosis for 2-5K electron systems on pre-exascale/exascale

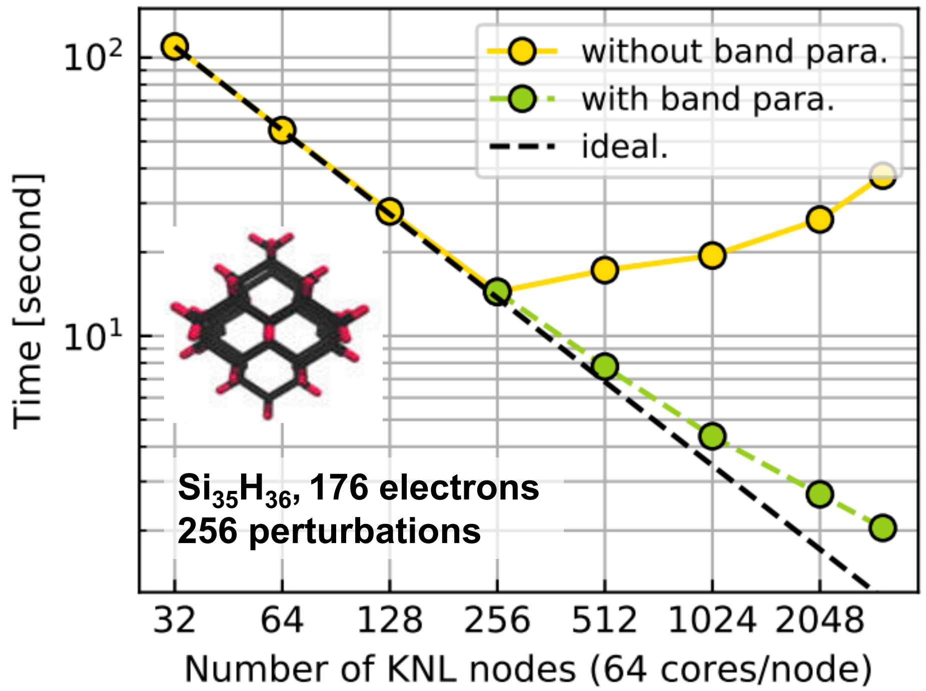

TASK-GROUPS FOR FFTS

- Distribute independent FFT operations on a smaller number of cores
  - Partition band group into task groups; each fits on single node
  - Shared memory MPI for transpose

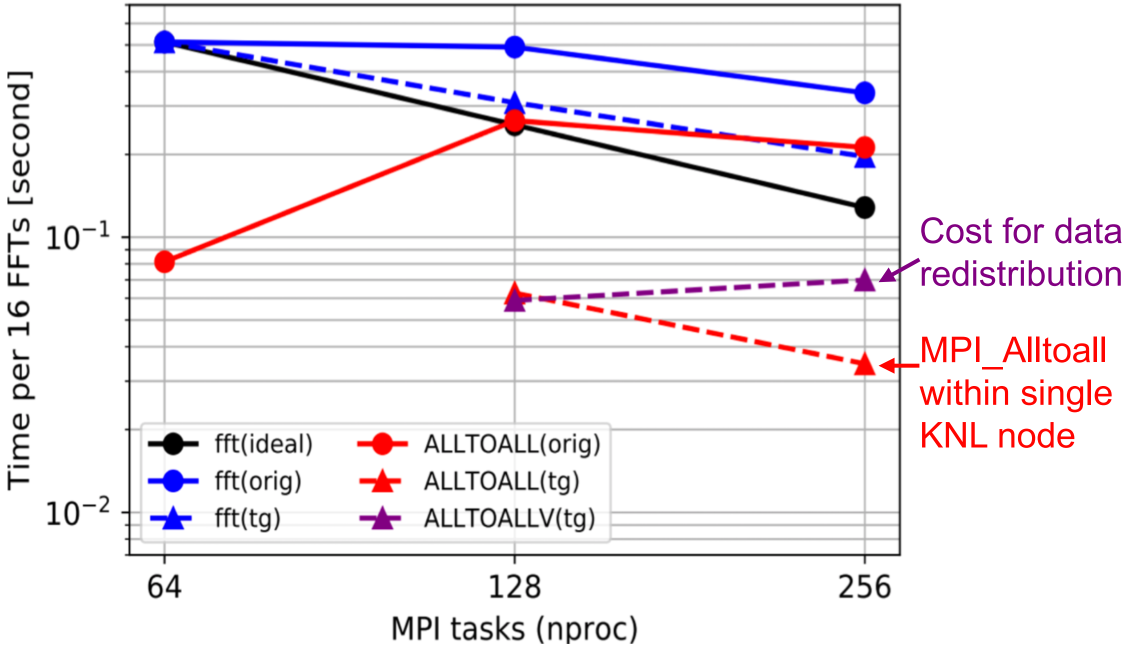

Zheng et al., IXPUG Fall 2017, APS March Meeting 2017

EXTREME SCALE UNSTRUCTURED ADAPTIVE CFD: AERODYNAMIC FLOW CONTROL

Case Study: Theta ESP science

Code: PHASTA • *PI: Kenneth Jansen (U. Colorado Boulder)* • CFD, unstructured mesh • *Catalyst: Hal Finkel*

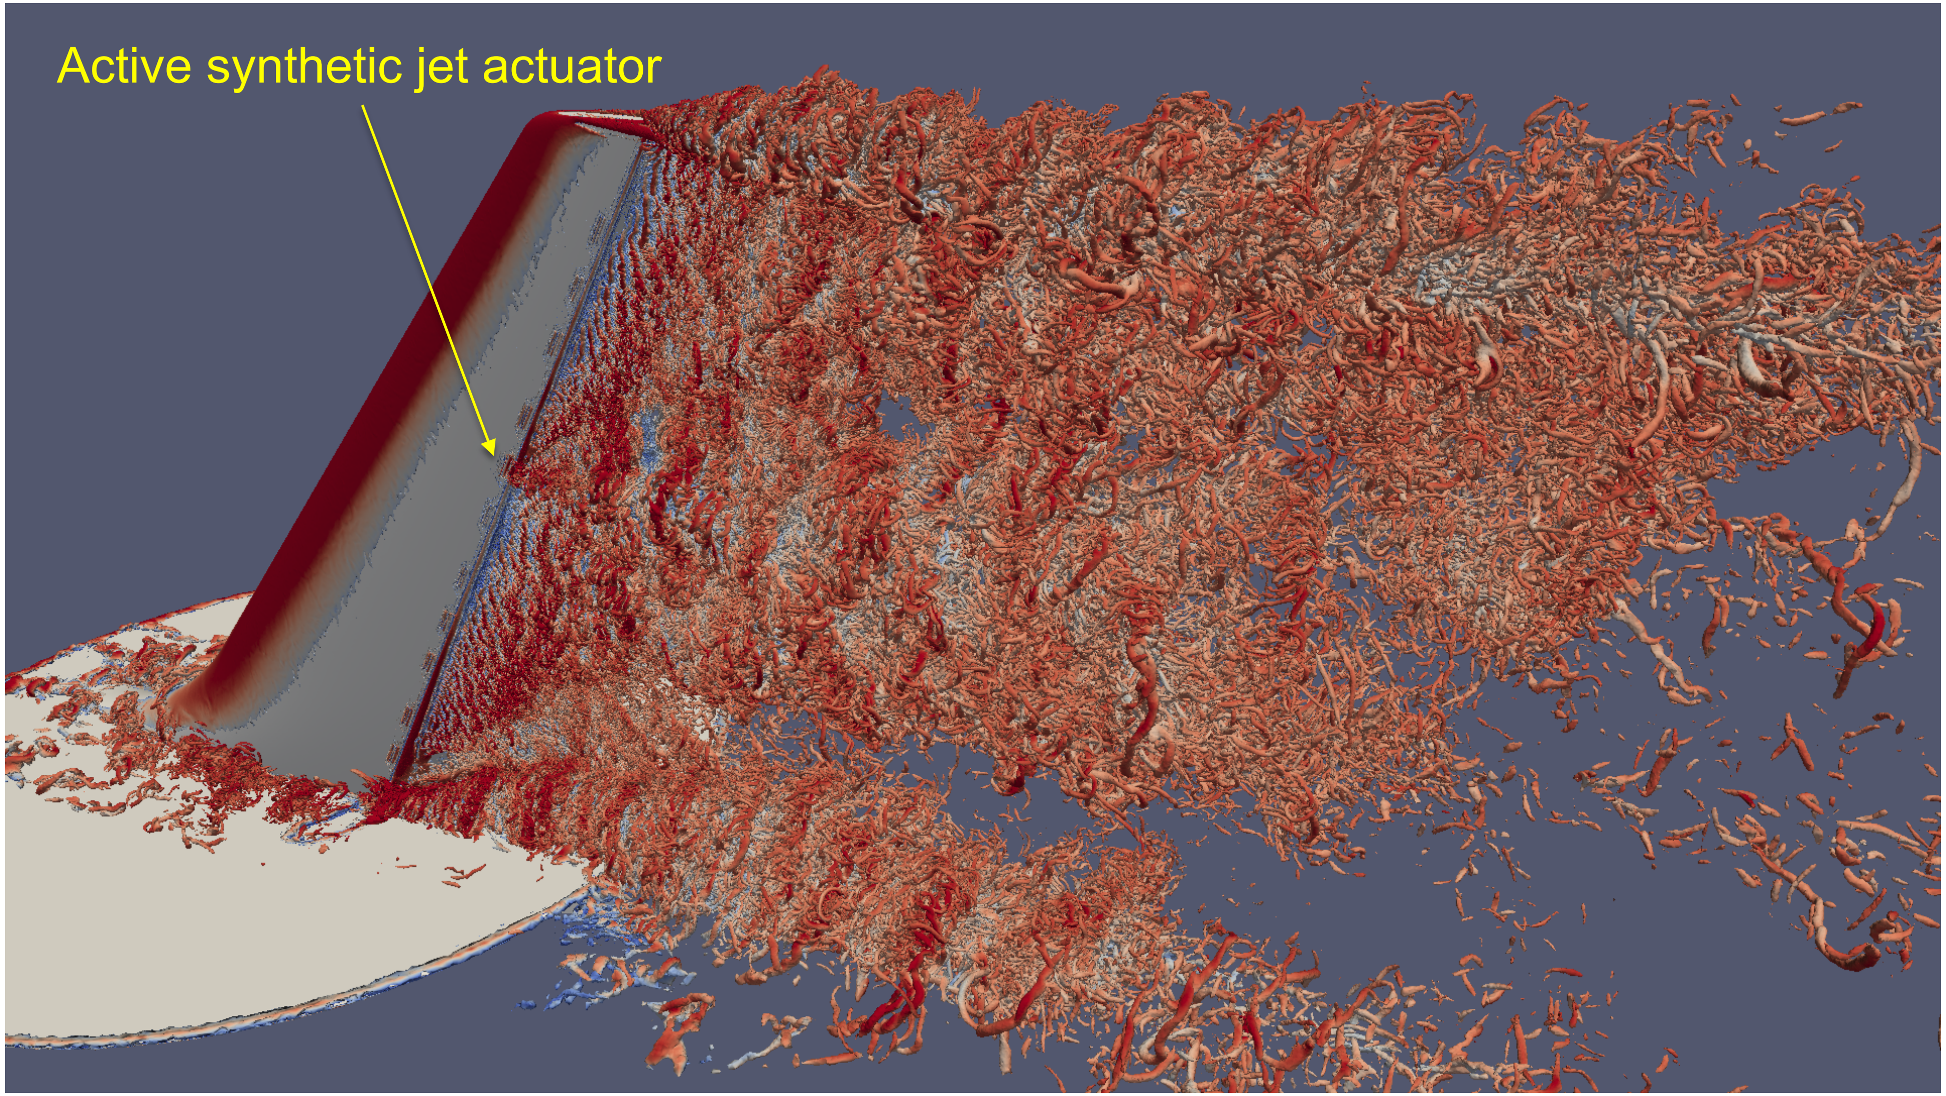

- 3D finite element
- unstructured adaptive mesh
- fully implicit
- 5 billion elements
- 2048 Theta nodes (128K KNL cores)

EXTREME SCALE UNSTRUCTURED ADAPTIVE CFD: AERODYNAMIC FLOW CONTROL

Code: PHASTA • *PI: Kenneth Jansen (U. Colorado Boulder)* • CFD, unstructured mesh • *Catalyst: Hal Finkel*

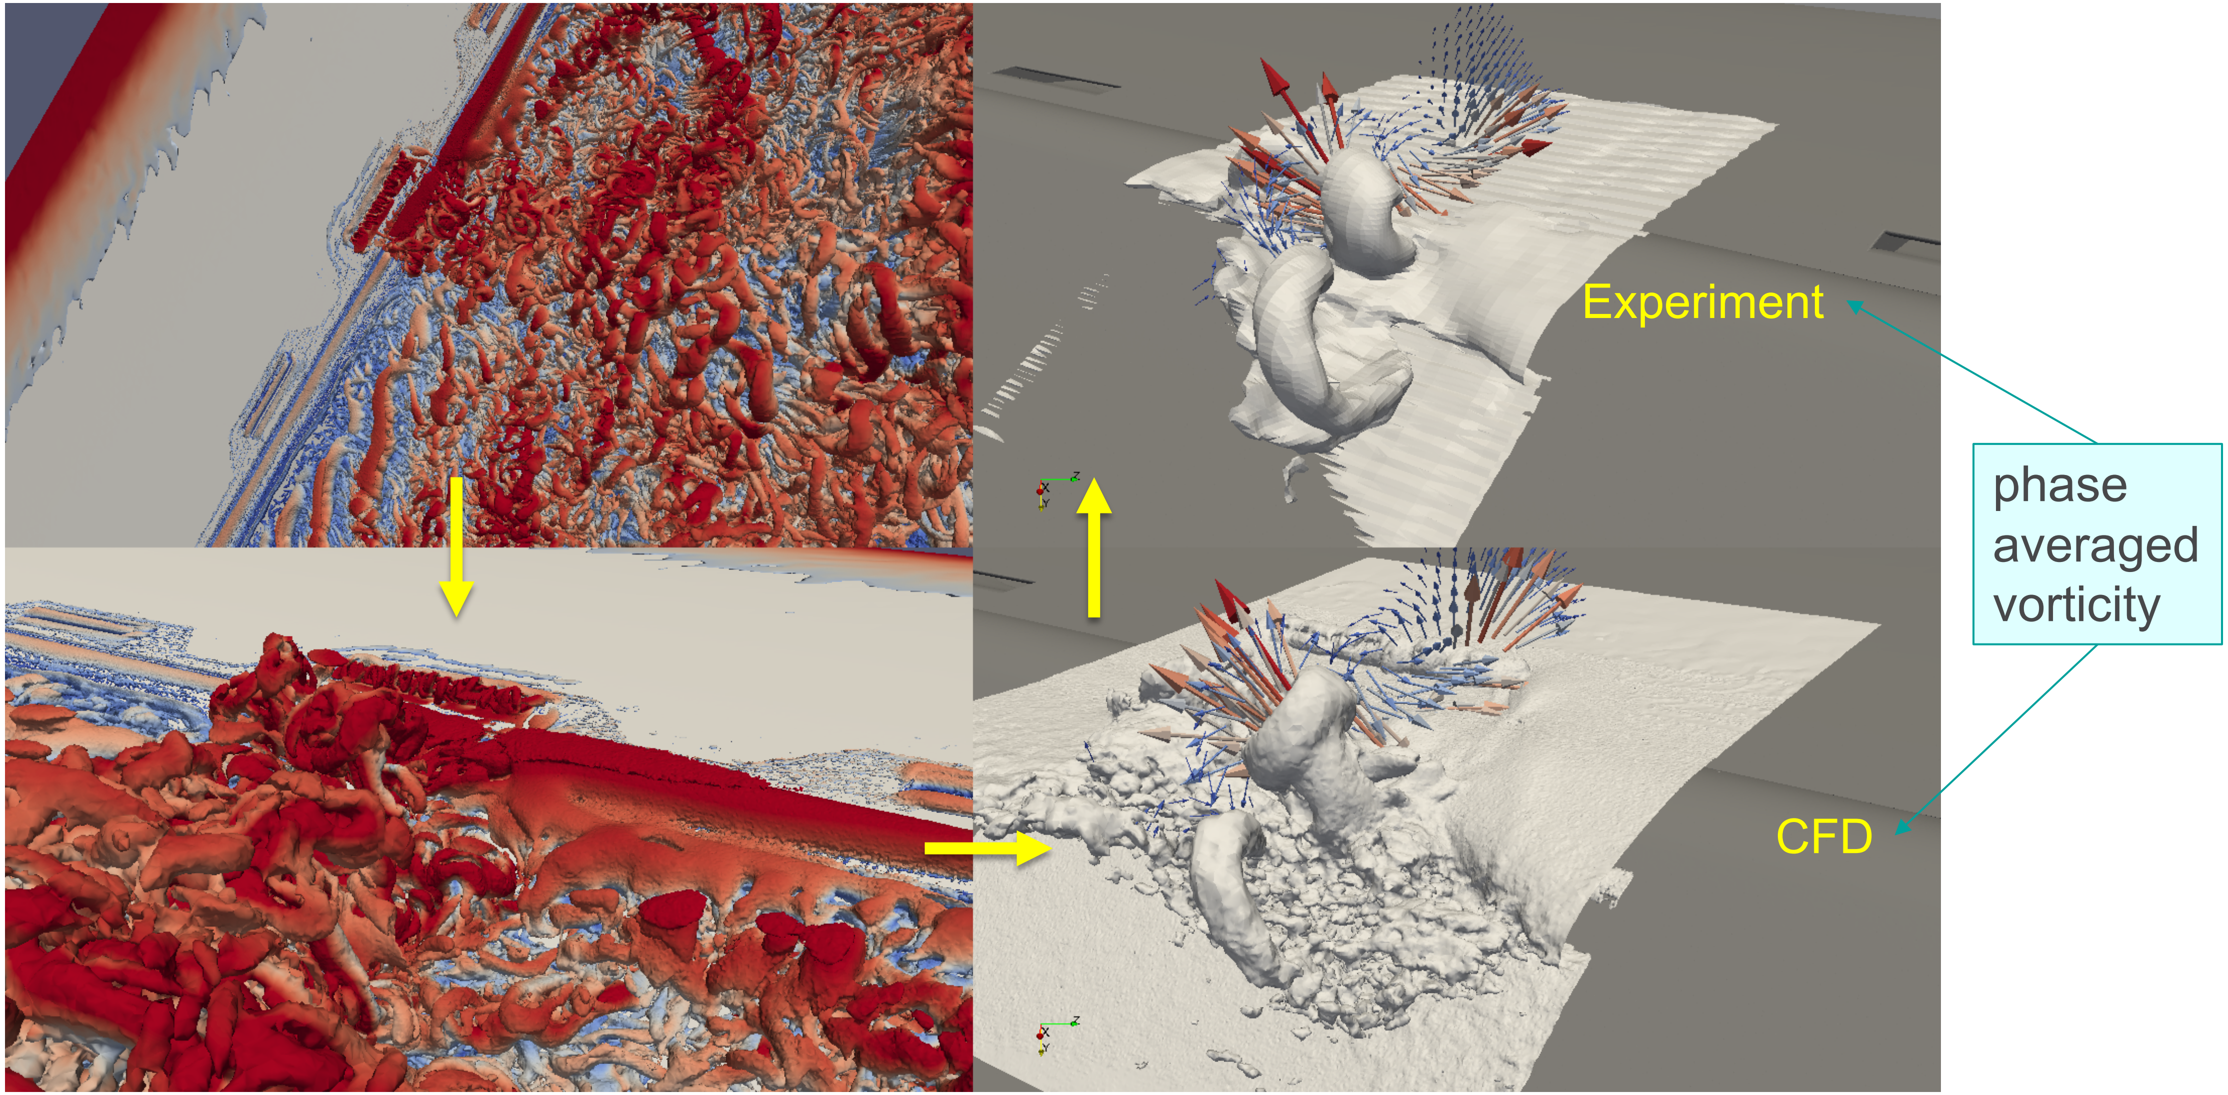

AURORA

ALCF 2021 EXASCALE SUPERCOMPUTER – AURORA

Intel Aurora supercomputer planned for 2018 shifted to 2021

Scaled up from **180 PF to over 1000 PF**

Support for three “pillars”

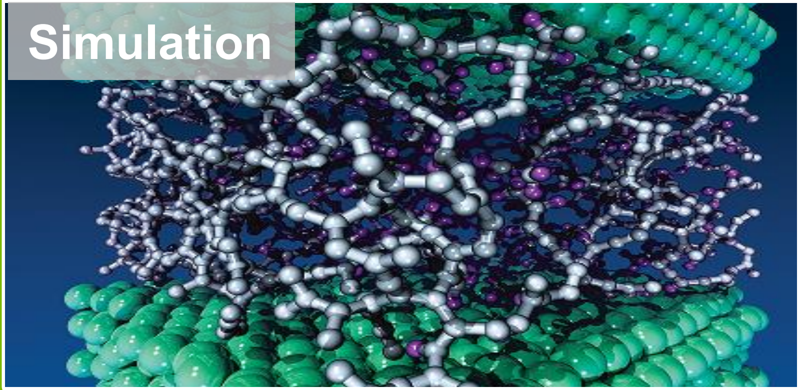

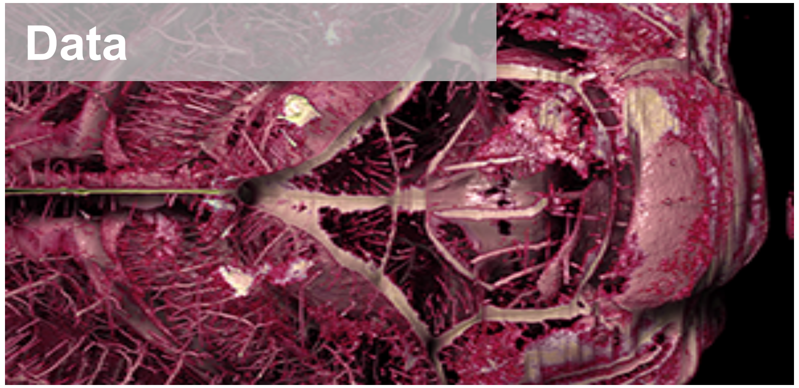

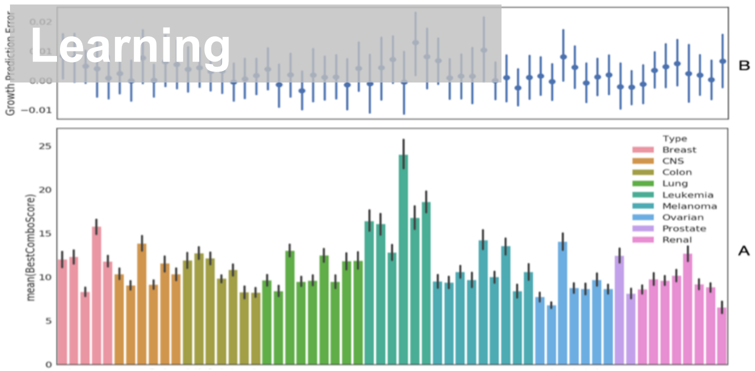

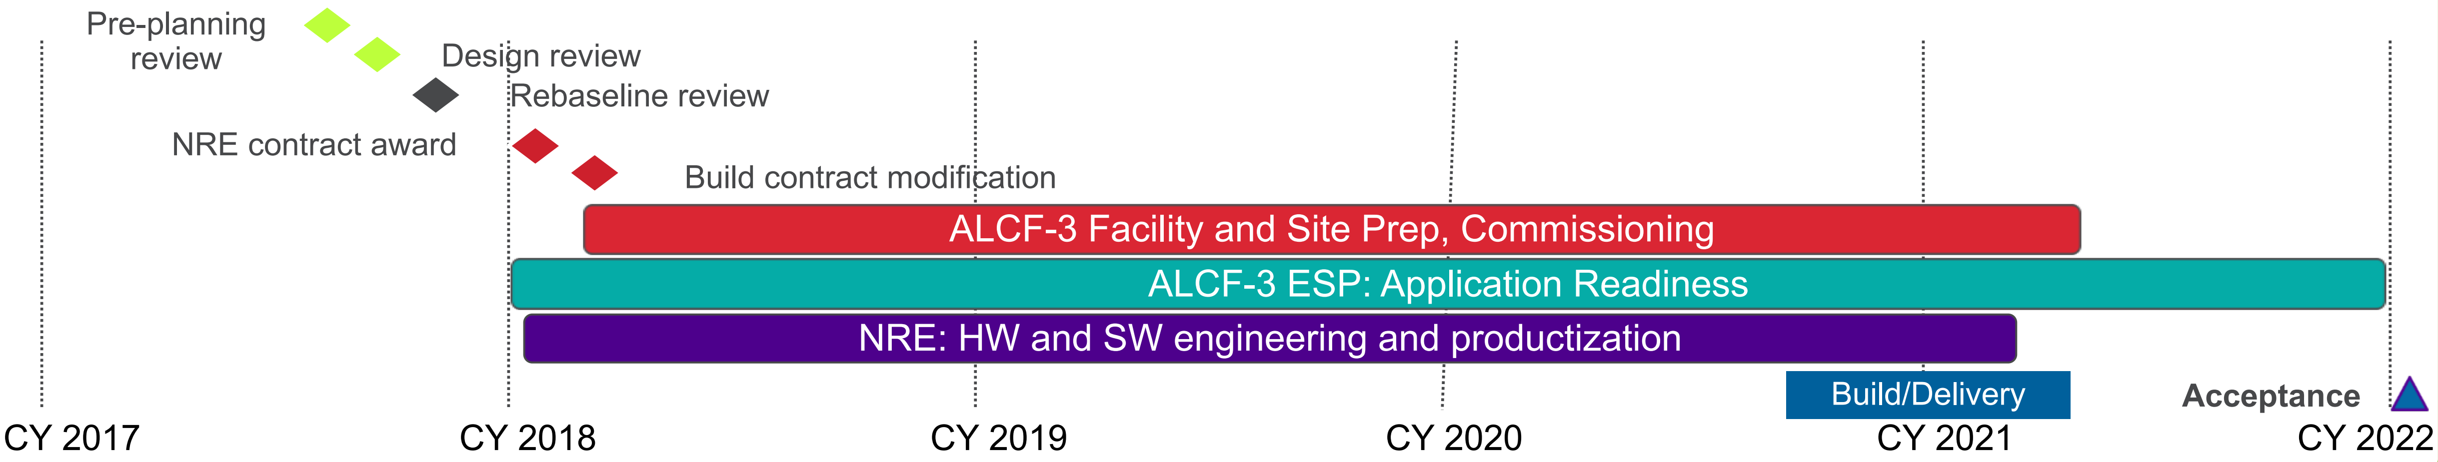

AURORA

Hardware and software optimized for Simulation, Data, and Learning

S
D
L

COMPUTE

- FLOPS
- Concurrency
- Memory performance
- ML/DL operations

I/O

- Speed
- Capacity
- Flexibility
  - Conventional I/O
  - Database
  - Analytics middleware

Programming Environment

- Optimizing compilers
- Latest OpenMP
- Key Big Data stack components
- Productivity languages
- ML/DL frameworks
- Optimized libraries
  - Math
  - Statistics
  - ML/NN

AURORA SYSTEM

- Nodes will have both high single thread core performance and the ability to get exceptional performance when there is concurrency of modest scale in the code
- Architecture optimized to support codes with sections of fine grain concurrency (~100 lines of code in a FOR loop e.g.) separated by serial sections
  - Degree of fine grain concurrency (e.g. number of loop iterations) that will be needed to fully exploit the performance opportunities is moderate. (~1000 for most applications)
  - Independence of these loops is ideal but not required for correctness
  - No limit on the number of such loops; overhead of starting/ending loops is very low
- Serial code (within an MPI rank) will execute very efficiently
- OpenMP 5 will likely contain the constructs necessary to guide the compiler to get optimal performance.
- The compute performance of the nodes will raise in a manner similar to the memory bandwidth
- The memory capacity will not grow as fast as the compute
  - The memory will all be high performance alleviating some concerns of explicitly managing multievel memory & data movement
  - The memory in a node will be coherent
- All compute will be first class citizens: equal access to all resources, memory and fabric etc.
- The fabric BW will be increasing similar to the compute performance for local communication patterns
  - Global communication BW will likely to not increase as fast as compute performance.

AURORA EARLY SCIENCE PROGRAM

AURORA ESP SIMULATION PROJECTS

Project	PI
Extending Moore's Law computing with Quantum Monte Carlo	Anouar Benali
Design and evaluation of high-efficiency boilers for energy production using a hierarchical V/UQ approach	Martin Berzins
High fidelity simulation of fusion reactor boundary plasmas	C.S. Chang
NWChemEx: Tackling Chemical, Materials & Biochemical Challenges in the Exascale Era	Thomas Dunning
Extreme-Scale Cosmological Hydrodynamics	Katrin Heitmann
Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Kenneth Jansen
Benchmark Simulations of Shock-Variable Density Turbulence and Shock-Boundary Layer Interactions with Applications to Engineering Modeling	Sanjiva Lele
Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics	Paul Mackenzie
Metascalable Layered Materials Genome	Aiichiro Nakano
Free Energy Landscapes of Membrane Transport Proteins	Benoit Roux

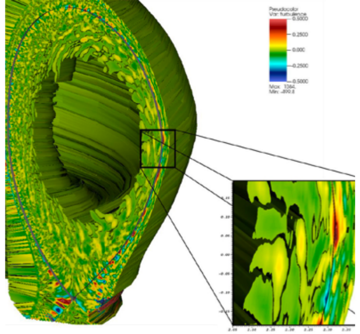

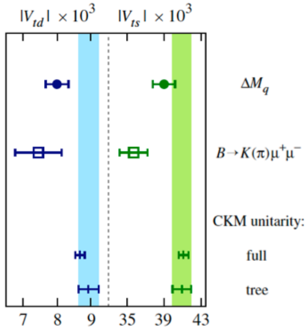

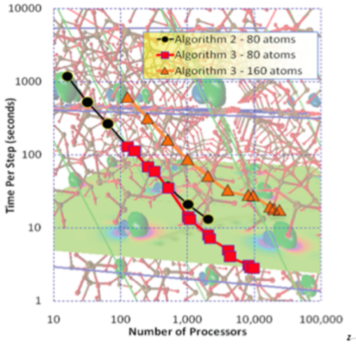

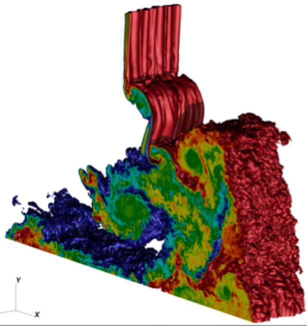

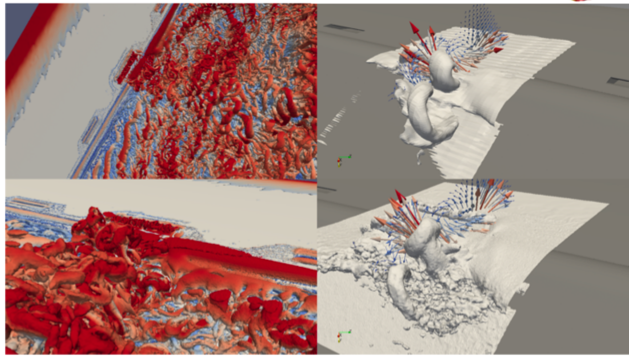

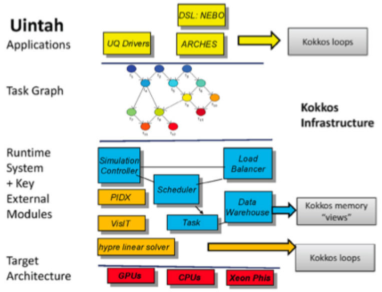

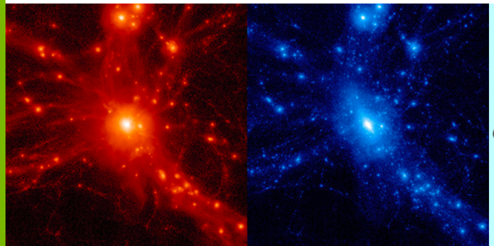

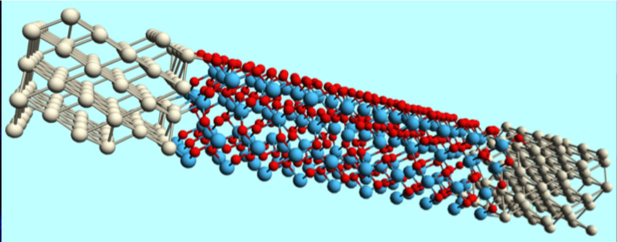

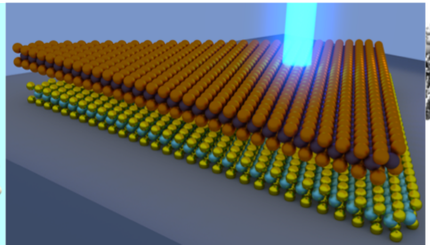

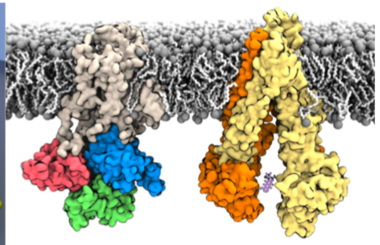

AURORA ESP TIMELINE (NOTIONAL)

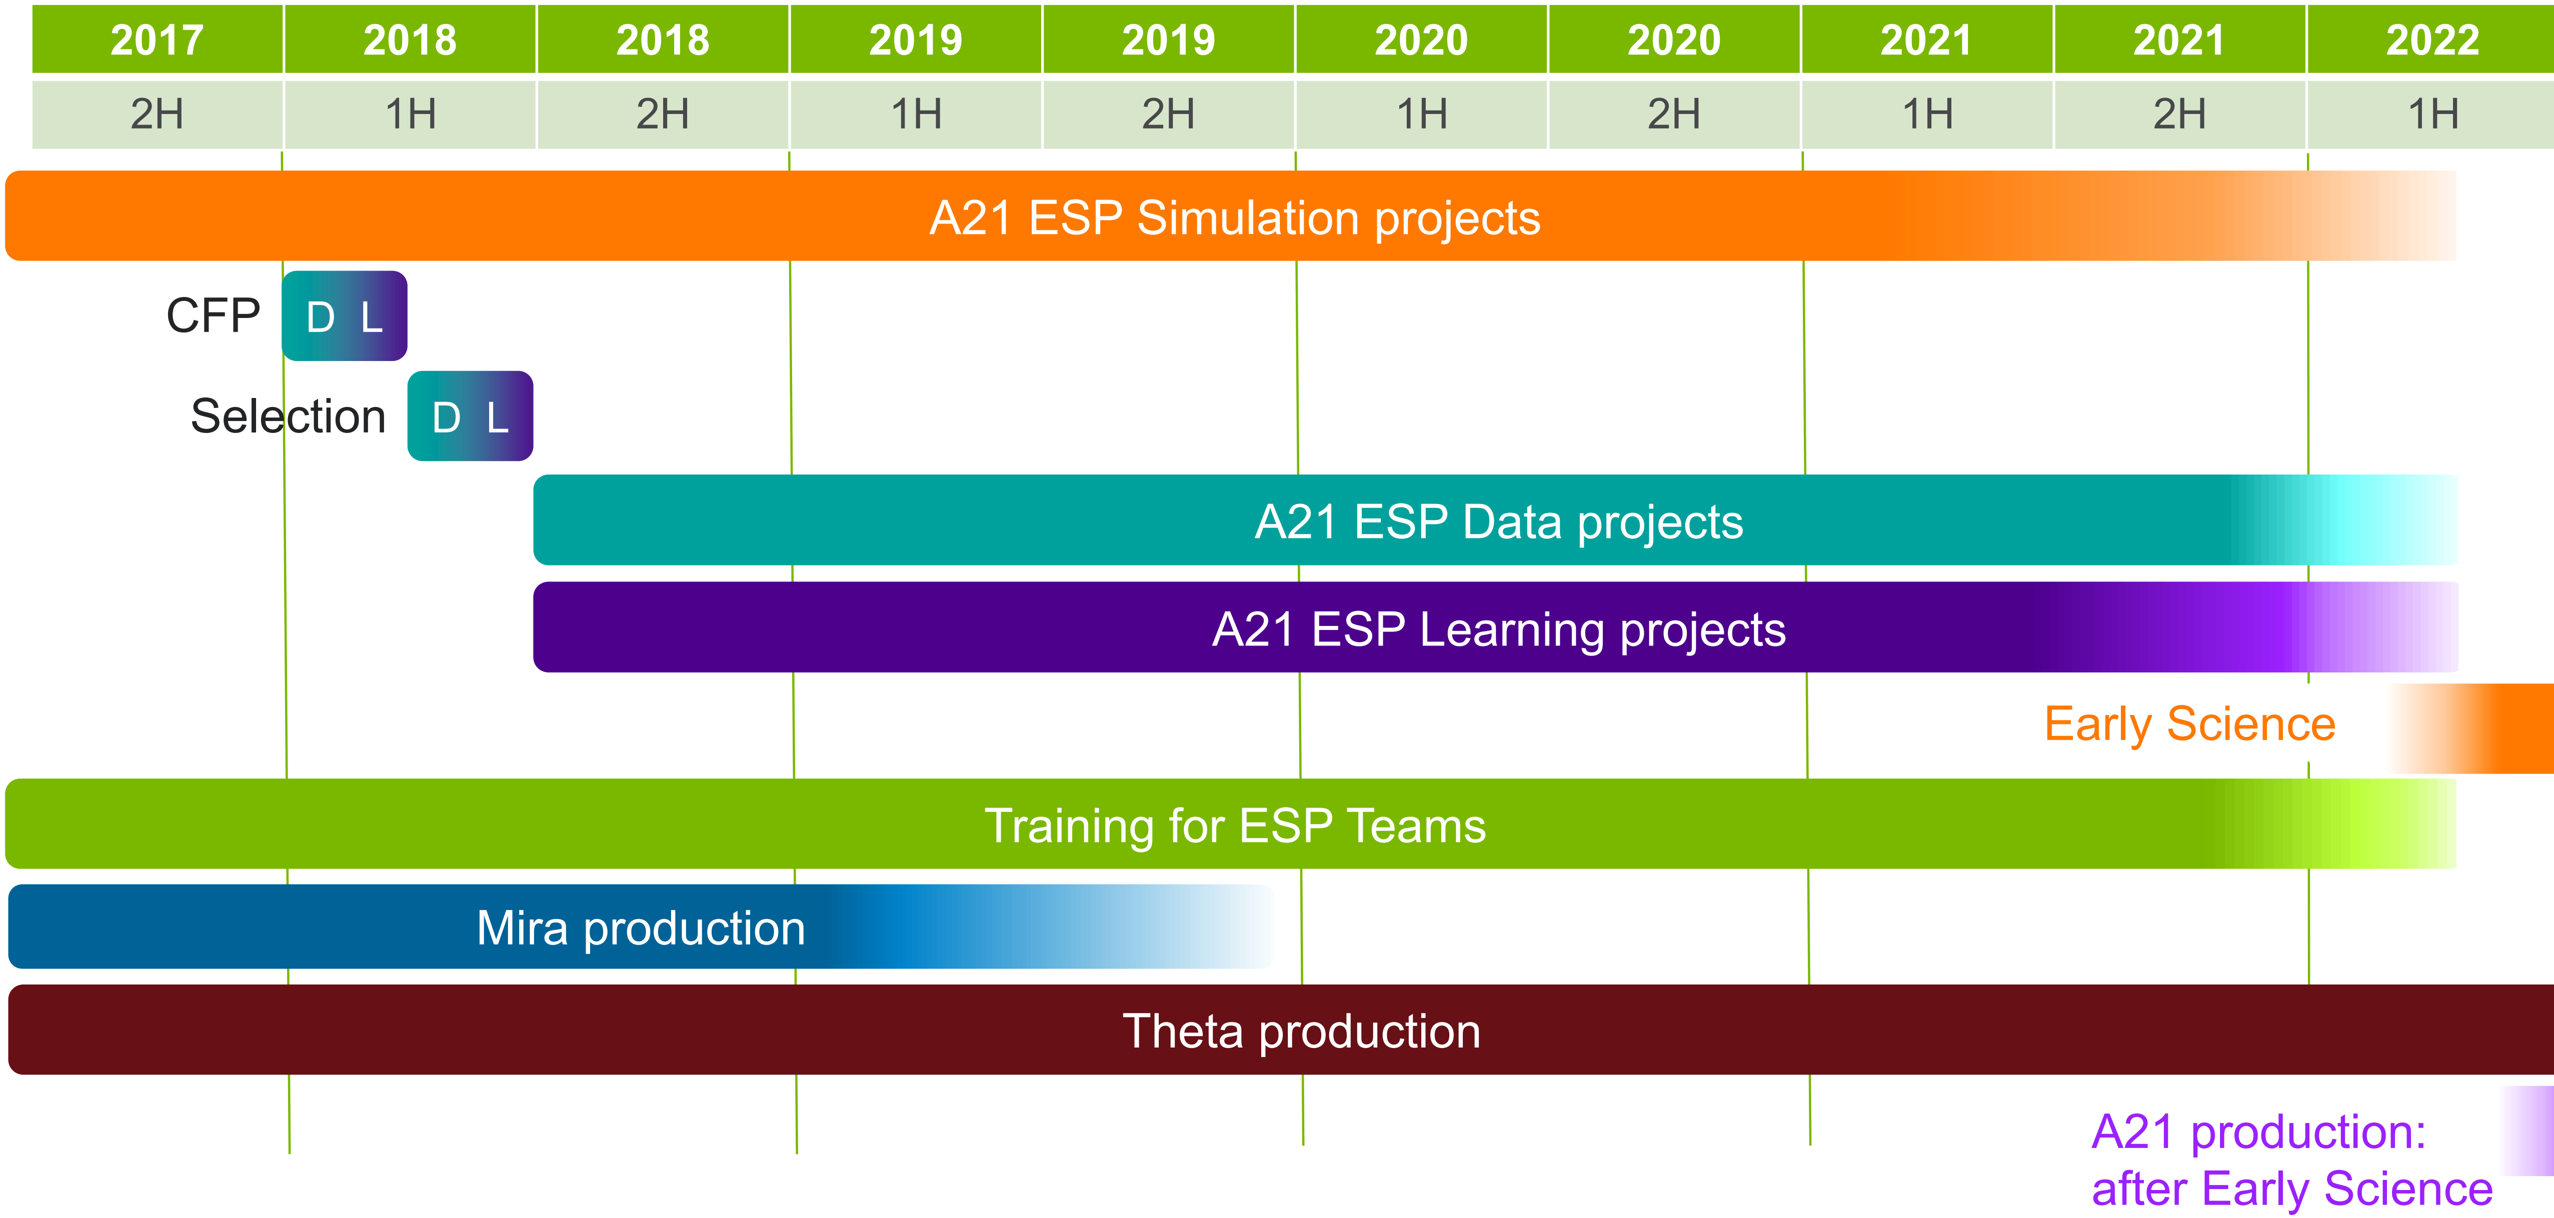

CALL FOR PROPOSALS: A21 ESP DATA, LEARNING PROJECTS

- **CFP January 2018**
  - Deadline 8 April 2018
- Selections June 2018
  - 5 Data projects
  - 5 Learning projects
- Two-year funded ALCF postdoc
- Cross-cutting proposals targeting the convergence of simulation, data and learning are very much encouraged.

DATA

- Experimental/observational data
  - Image analysis
  - Multidimensional structure discovery
- Complex and interactive workflows
- On-demand HPC
- Persistent data techniques
  - Object store
  - Databases
- Streaming/real-time data
- Uncertainty quantification
- Statistical methods
- Graph analytics

LEARNING

- Deep learning
- Machine learning steering simulations
  - Parameter scans
  - Materials design
  - Observational signatures
- Data-driven models and refinement for science using ML/DL
- Hyperparameter optimization
- Pattern recognition
- Reduced model derivation
- Bridging gaps in theory

https://www.alcf.anl.gov/alcf-aurora2021-esp-dl-cfp

THANK YOU

Upcoming Program Deadlines

Aurora Early Science Program for
Learning and Data
Call for Proposals in January 2018

ADSP Program
Call for Proposals in April 2018