The Exascale Computing Project

Approved for public release

ECP EXASCALE COMPUTING PROJECT

Lori Diachin (LLNL), ECP Deputy Director

BESAC Summer Meeting
July 11, 2019

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

The Exascale Computing Project (ECP) enables US revolutions in technology development; scientific discovery; healthcare; energy, economic, and national security

ECP mission

Develop exascale-ready applications and solutions that address currently intractable problems of strategic importance and national interest.

Create and deploy an expanded and vertically integrated software stack on DOE HPC exascale and pre-exascale systems, defining the enduring US exascale ecosystem.

Deliver **US HPC vendor technology advances and deploy ECP products** to DOE HPC pre-exascale and exascale systems.

ECP vision

Deliver **exascale simulation and data science innovations and solutions to national problems** that enhance US economic competitiveness, change our quality of life, and strengthen our national security.

The ECP is part of the broader DOE Exascale Computing Initiative (ECI)

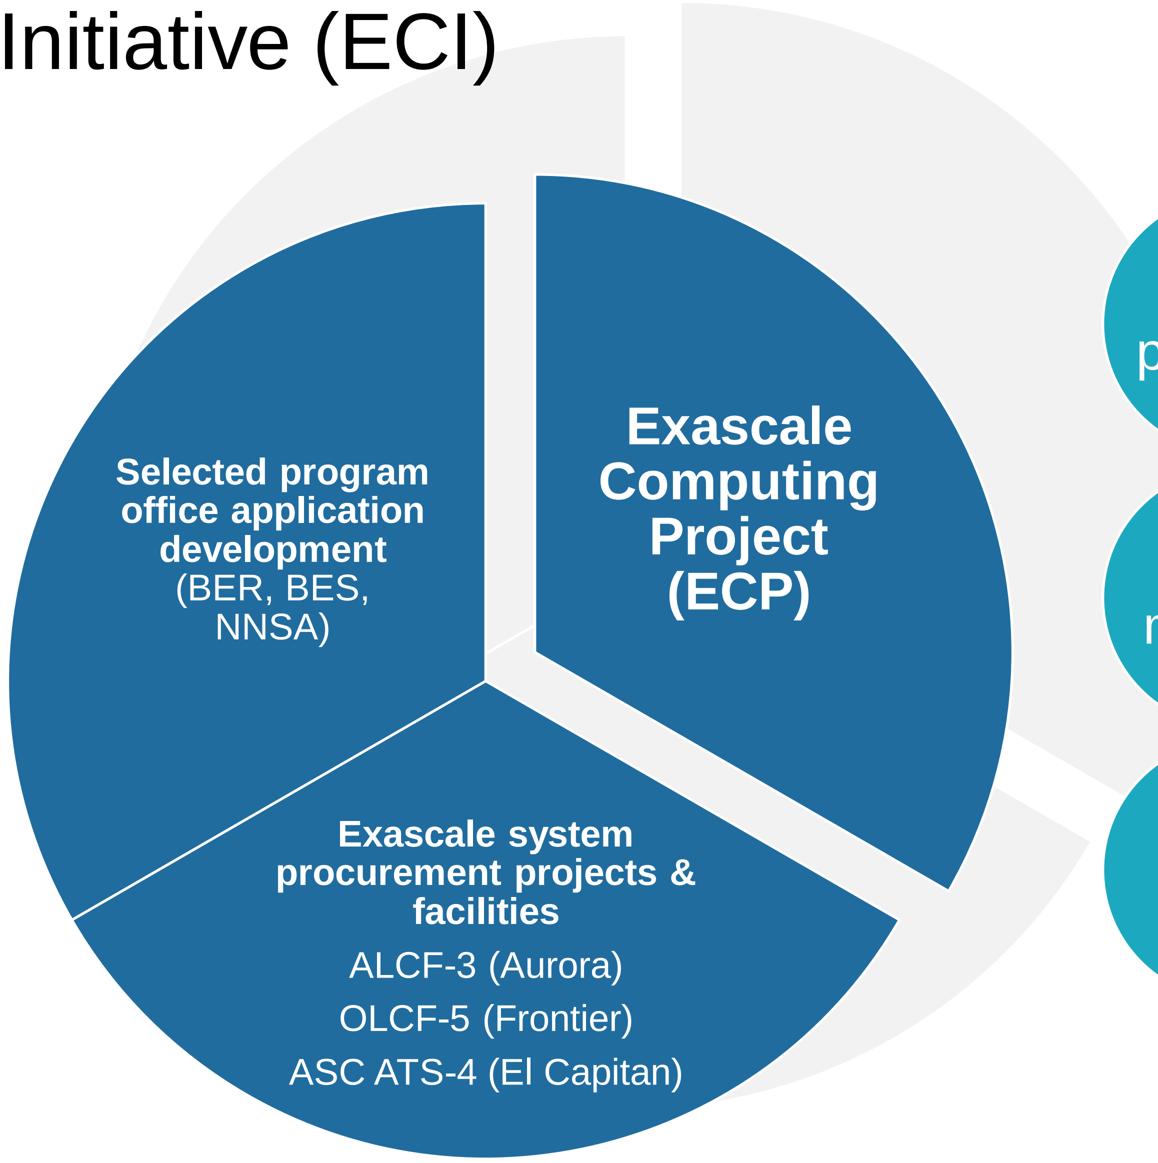

Three Major Components of the ECI

ECI partners: US DOE Office of Science (SC) and National Nuclear Security Administration (NNSA)

ECI mission: Accelerate R&D, acquisition, and deployment to deliver exascale computing capability to DOE national labs by the early- to mid-2020s

ECI focus: Delivery of an *enduring and capable exascale computing capability for use by a wide range of applications* of importance to DOE and the US

Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

Pre-Exascale Systems | **First U.S. Exascale Systems***

2012

ORNL
Cray/AMD/NVIDIA

ANL
IBM BG/Q

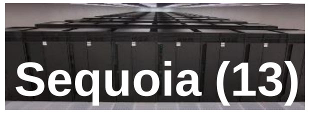

LLNL
IBM BG/Q

2016

ANL
Cray/Intel KNL

LBNL
Cray/Intel Xeon/KNL

LANL/SNL
Cray/Intel Xeon/KNL

2018

ORNL
IBM/NVIDIA

LLNL
IBM/NVIDIA

2020

2021-2023

FRONTIER

ORNL*
Cray/AMD

ANL*
Intel/Cray

LBNL
Cray/AMD/NVIDIA

CROSSROADS

LANL/SNL
TBD

LLNL*
TBD

ECP EXASCALE COMPUTING PROJECT

ECP by the Numbers

7
YEARS

$1.7B

A seven-year, $1.7 B R&D effort that launched in 2016

6
CORE DOE
LABS

Six core DOE National Laboratories: Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Sandia

- Staff from most of the 17 DOE national laboratories take part in the project

3
TECHNICAL
FOCUS
AREAS

Three technical focus areas: Hardware and Integration, Software Technology, Application Development all strongly partnered with specialists in project management

81
R&D TEAMS

1000
RESEARCHERS

More than 80 top-notch R&D teams

Hundreds of consequential milestones delivered on schedule and within budget since project inception

What do we mean by “exascale”?

- We don’t just want advancements in theoretical peak performance. **10^{18} FLOP/s is not enough.**
- Massive amounts of theoretical computing does us no good if real applications can’t make use of it.
- The ECP is quantitatively measuring application capability improvement against today’s ~20 PF machines. **A 50x improvement is required for success.**

Sequoia IBM Blue Gene/Q

Theoretical peak:	20.1 PFLOP/s
LINPACK peak:	17.2 PFLOP/s
Application A:	11.9 PFLOP/s
Application B:	4.3 PFLOP/s
Application C:	0.6 PFLOP/s

$\left.\begin{array}{l}\text{Application A, B, C}\end{array}\right\}$ Specialized hardware works well for some applications, not so well for others

Exascale means real capability improvement in the science we can do and how fast we can do it.

Why high performance computing is hard, and getting harder

- Exascale architectures are specialized and represent a significant departure from traditional HPC application targets
- Applications will need to effectively use accelerators and find more and more concurrency to keep up
- Moving data becoming increasingly costly relative to computation
- I/O, vis, analysis becoming major bottlenecks
- New programming models are being developed to supplement traditional MPI+X approaches

Preparing applications for new architectures can be difficult and time-consuming, working together and learning from each other is crucial

The three technical areas in ECP have the necessary components to meet national goals

Performant mission and science applications @ scale

Aggressive RD&D Project	Mission apps & integrated S/W stack	Deployment to DOE HPC Facilities	Hardware tech advances

Application Development (AD)

Develop and enhance the predictive capability of applications critical to the DOE

24 applications including national security, energy, earth systems, economic security, materials, and data

Software Technology (ST)

Deliver expanded and vertically integrated software stack to achieve full potential of exascale computing

66 unique software products spanning programming models and run times, math libraries, data and visualization

Of particular interest to BES: sparse/dense solvers, FFTs, performance portability, data analysis

Hardware and Integration (HI)

Integrated delivery of ECP products on targeted systems at leading DOE HPC facilities

6 US HPC vendors focused on exascale node and system design; application integration and software deployment to facilities

Of particular interest to BES: early access to exascale hardware, performance engineers for applications, training, software deployment

EXASCALE COMPUTING PROJECT

ECP **Application Development (AD)**

Goal

Develop and enhance predictive capability of applications critical to DOE across science, energy, and national security mission space

- High impact science or engineering exascale *challenge problem*
- Detailed criteria for assessing successful completion of challenge problem
- A Figure of Merit (FOM) formula quantifying performance of challenge problem
- Demonstration and assessment of effective software integration

Chemistry and Materials

Earth and Space Science

Energy

Data Analytics and Optimization

National Security

Co-Design

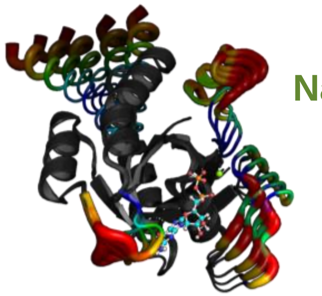

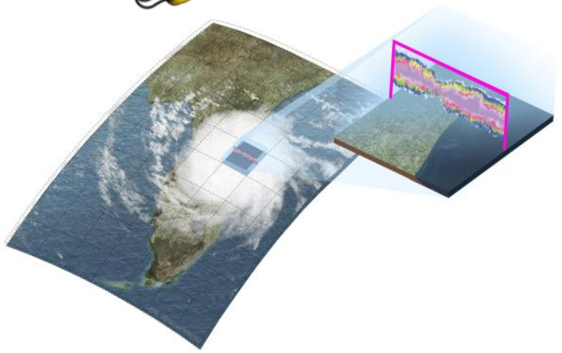

AD subprojects target national problems in DOE mission areas

National security

Next-generation, **stockpile stewardship** codes

Reentry-vehicle- environment simulation

Multi-physics science simulations **of high-energy density physics** conditions

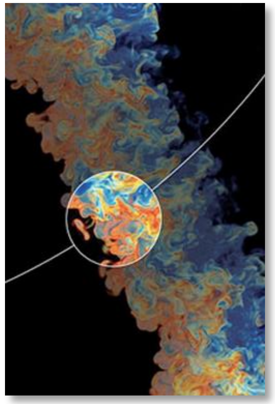

Energy security

Turbine **wind plant** efficiency

Design and commercialization of **SMRs**

Nuclear fission and fusion reactor **materials design**

Subsurface use for **carbon capture**, petroleum extraction, waste disposal

High-efficiency, low-emission **combustion engine** and gas turbine design

Scale up of **clean fossil fuel** combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Reliable and efficient planning of the **power grid**

Seismic hazard risk assessment

Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Plasma wakefield accelerator design

Light source-enabled **analysis of protein and molecular structure** and design

Find, predict, and control materials and properties

Predict and control **magnetically confined fusion plasmas**

Demystify **origin of chemical elements**

Earth system

Accurate regional impact assessments in **Earth system models**

Stress-resistant crop analysis and catalytic conversion of **biomass-derived alcohols**

Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

Health care

Accelerate and translate **cancer research** (partnership with NIH)

The 8 BES related applications represent ~30% percent of the total ASCR AD scope

EXASCALE COMPUTING PROJECT

Co-design centers address computational motifs common to multiple application projects

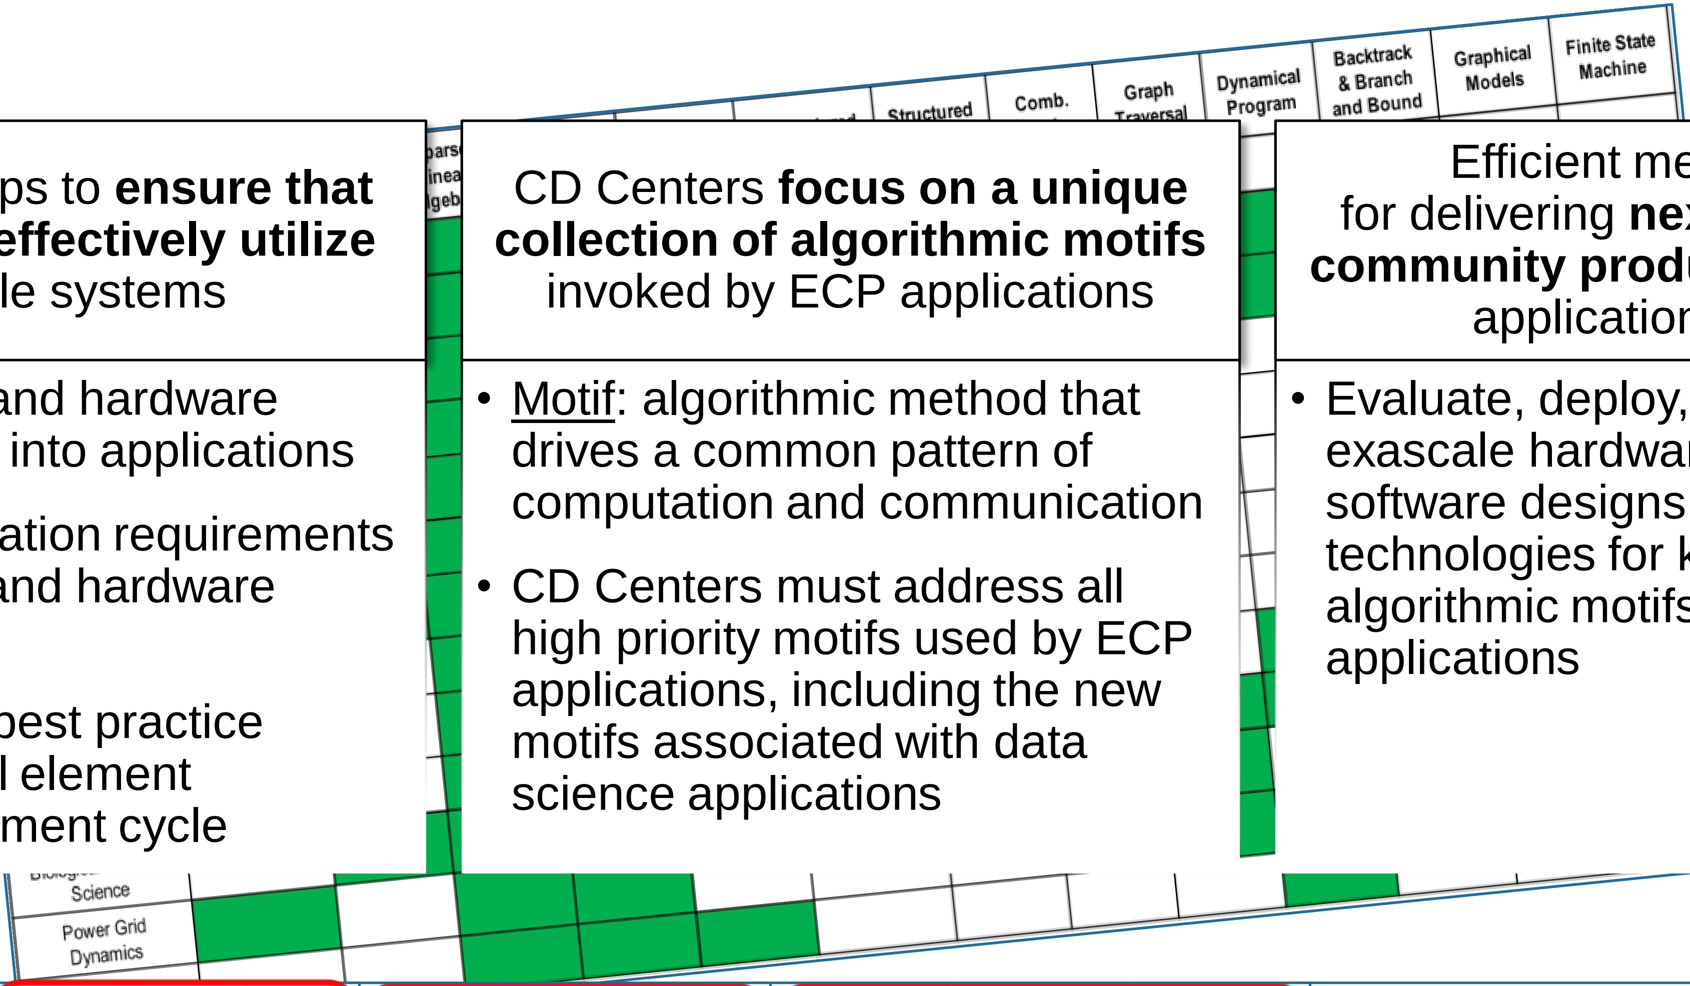

Co-design helps to **ensure that applications effectively utilize** exascale systems

- Pull software and hardware developments into applications
- Pushes application requirements into software and hardware RD&D
- Evolved from best practice to an essential element of the development cycle

CD Centers **focus on a unique collection of algorithmic motifs** invoked by ECP applications

- Motif: algorithmic method that drives a common pattern of computation and communication
- CD Centers must address all high priority motifs used by ECP applications, including the new motifs associated with data science applications

Efficient mechanism for delivering **next-generation community products** with broad application impact

- Evaluate, deploy, and integrate exascale hardware-aware software designs and technologies for key crosscutting algorithmic motifs into applications

CODAR	**COPA**	**AMReX**	**CEED**	**ExaGraph**	**ExaLearn**
Data and workflows	*Particles/mesh methods*	*Block structured AMR*	*Finite element discretization*	*Graph-based algorithms*	*Machine Learning*

Common challenges that are being addressed by the application teams

1) **Porting to accelerator-based architectures**
2) **Exposing additional parallelism**
3) **Coupling codes to create new multiphysics capability**
4) **Adopting new mathematical approaches**
5) **Algorithmic or model improvements**
6) **Leveraging optimized libraries**

Exascale hardware realities are forcing new thinking of algorithmic implementations and the move to new algorithms

Algorithmic Implementations

- Reduced communication/data movement
- Much greater locality awareness
- Much higher cost of global synchronization
- Value to mixed precision where possible

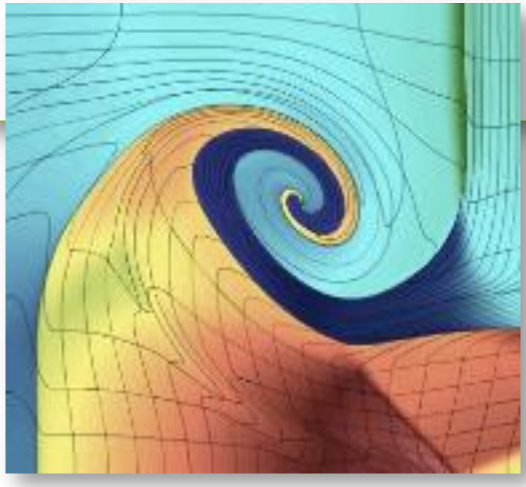

New Algorithms

- Adopting Monte Carlo vs. Deterministic approaches
- Exchanging on-the-fly recomputation vs. data table lookup (e.g. neutron cross sections)
- Moving to higher-order methods (e.g. CFD)
- Particle algorithms that favor collecting similar events together rather than parallelism though individual histories

Exascale apps can deliver transformative products and solutions for BES science needs

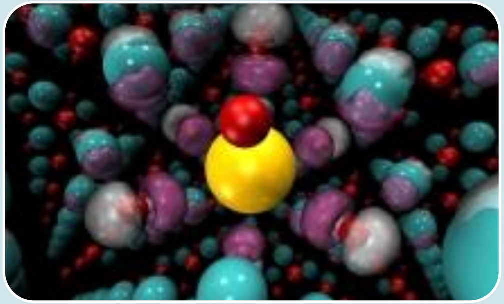

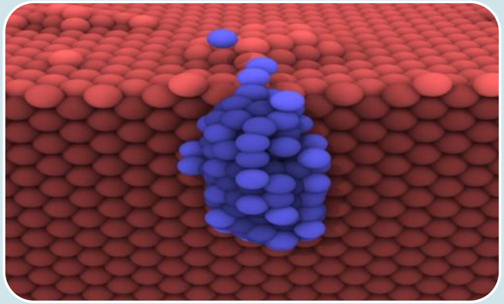

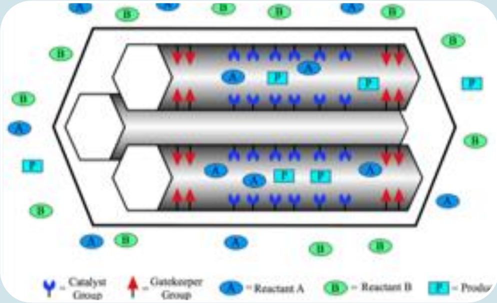

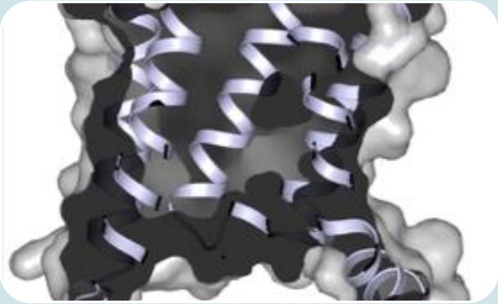

QMCPACK

Design and optimize next-generation materials from first principles with predictive accuracy

- Reference GPU implementation of AFQMC; excellent performance compared to equivalent CPU x10-70 on Summit compared to Theta
- QMCPACK v3.5.0 released
- Completed implementation of delayed updated algorithm

Lead: ORNL

EXAALT

Simultaneously address time, length, and accuracy requirements for predictive microstructural evolution of materials

- New sublattice ParSplice method benchmarked at scale on Theta
- Excellent scaling on 4096 nodes (262,144 cores)
- Test case used closely resembles grand challenge and FOM demonstration system

Lead: LANL

GAMESS

Design more robust and selective catalysts orders of magnitude more efficient at temperatures hundreds of degrees lower

- Completed initial demonstration of EFMO/QMC method; tested on a number of systems
- Developed and optimized RI-MP2 energy plus gradient code; run on Summit GPUs

Lead: Ames

ExaFEL

Process data without beam time loss; determine nanoparticle size & shape changes; engineer functional properties in biology and material science

- Increased number of images used for reconstruction by 2.7X for experiments
- 9:1 reduction via data compression
- Legion task-based programming model for scaling and portability

Lead: SLAC

Exascale apps can deliver transformative products and solutions for BES science needs

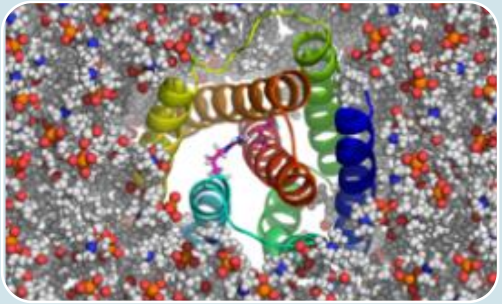

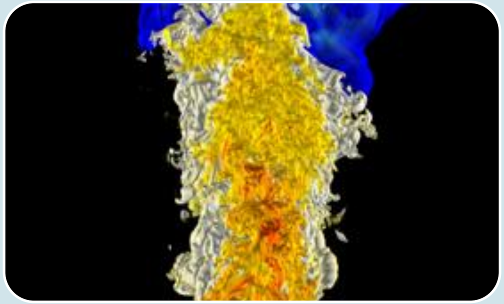

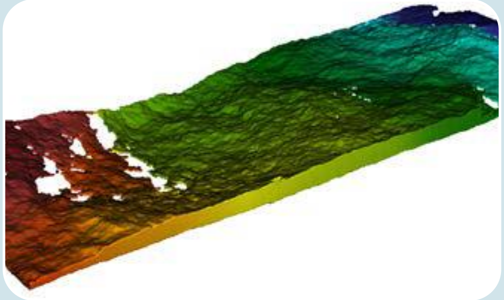

NWChemEx
Develop new optimal catalysts while changing the current design processes that remain costly, time consuming, and dominated by trial and-error

- Development and implementation of core computational kernels: dense and sparse tensors, canonical and local hartree-fock and DFT codes, ab initio MD module, classical MD module
- Initial port to GPU

Lead: PNNL

PELE
Reduction or elimination of current cut-and-try approaches for combustion system design

- Extended Pele with tiling approach for data-parallel shared memory and embedded boundaries
- Ported Pete compressible hydrodynamics to GPUs using OpenACC and CUDA
- Completed initial demonstration of multi-injection diesel ignition

Lead: SNL

Subsurface
Reliably guide safe long-term consequential decisions about storage, sequestration, and exploration

- Coupled structured finite volume Chombo code to unstructured finite element code GEOS using HDF5 API
- Performance portability: initial implementation of Proto layer, GEOS built on RAJA
- Kernels implemented on GPUS

Lead: LBNL

ExaAM
Accelerate the widespread adoption of AM by enabling routine fabrication of qualifiable metal parts

- Used co-design products from CEED and AMReX to create a new version of Truchas
- Implemented and demod initial coupled physics capabilities – melt pool thermal fluids + thermomechanics
- Ported many components to GPUs

Lead: ORNL

ExaFEL: Data Analytics at Exascale for Free Electron Lasers

PI: Amedeo Perazzo, SLAC
Institutions: SLAC, LBNL, LANL

The ExaFEL project will substantially reduce time to analyze molecular structure x-ray diffraction data

- Detector data rates at light sources advancing exponentially
  - LCLS by 3 orders of magnitude by 2025 with the LCLS-II-HE upgrade
- Fast analysis required to allow users to iterate on their experiments and extract the most value from beam line time
- Imaging of structural dynamics and heterogeneity in individual molecules using non-crystalline based diffractive imaging in x-ray free-electron lasers
  - Single particle imaging
  - Fluctuation x-ray scattering
- Reconstruction of 3D structure from 2D images via Multi-tiered Iterative Phasing (MTIP) Algorithm

***Must be able to keep up with data taking rates: fast feedback** (seconds / minute timescale) **is essential** to reduce the time to complete the experiment, improve data quality, and increase the success rate*

ExaFEL Key Performance Parameter: Number of events analysed per second

Pre ExaFEL capability (LCLS-I): 10 Hz

- LCLS-I operates at 120 Hz, hit rate is ~10% ⇨ 10 events/s for reconstruction

Target capability (LCLS-II & LCLS-II-HE): 5 kHz

- LCLS-II high rate detectors are expected to operate at 50 kHz by 2024-2026, hit rate ~10% ⇨ 5000 events/s for reconstruction (after DRP)
- DRP = Data Reduction Pipeline (for SFX and SPI: vetoes events which are not hits)

EXASCALE COMPUTING PROJECT

Example of compute intensive algorithms for ExaFEL: M-TIP - a new algorithm for single particle imaging

M-TIP (**Multi-Tiered Iterative Phasing**) is an algorithmic framework that **simultaneously determines conformational states, orientations, intensity, and phase** from single particle diffraction images

The aim is to reconstruct a 3D structure of a single particle

- We can NOT measure: a) the orientations of the individual particles and b) phases of the diffraction patterns
- MTIP is an iterative algorithm that deduces these two sets of unknowns given some constraints

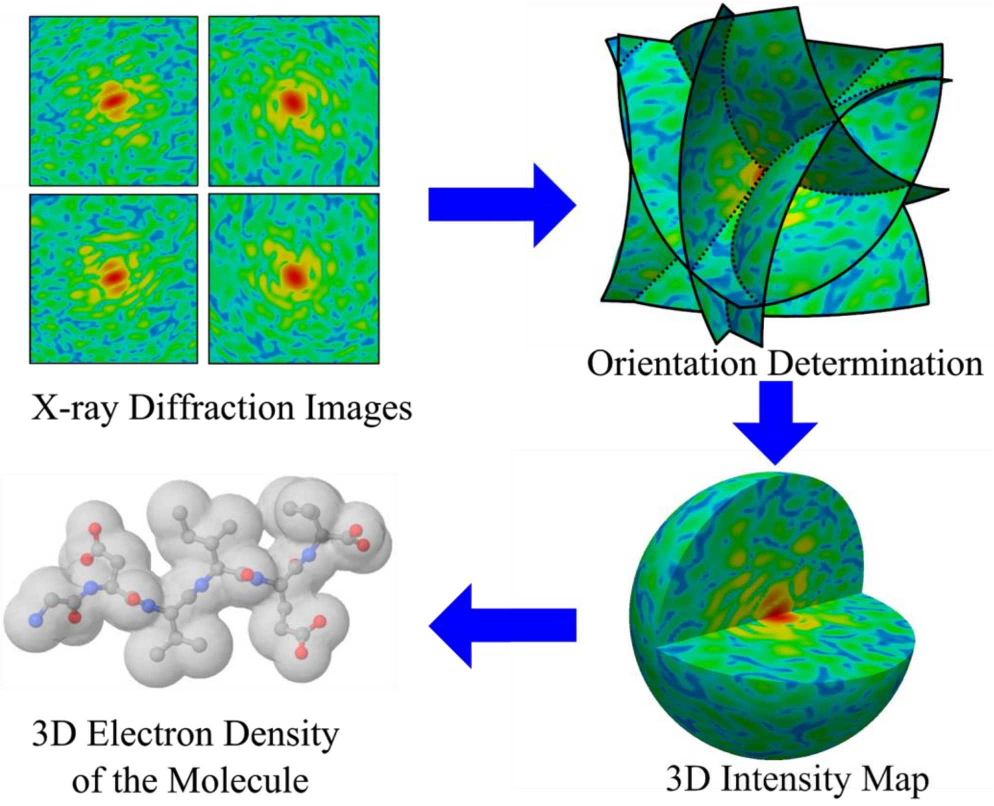

Progress:

- **Successful 3D reconstruction** of two viruses from LCLS SPI data
- **Designed new Cartesian** to non-uniform framework; replaces the current polar framework
  - **Improves scalability** - New approach can be fully parallelized over all images, whereas the old polar approach was done one image at a time
  - **Reduces complexity** from $O(DN^{4/3})$ to $O(D + N)$ (D = # data points, N = # grid points)

[1] Donatelli JJ, Sethian JA, and Zwart PH (2017) Reconstruction from limited single-particle diffraction data via simultaneous determination of state, orientation, intensity and phase. PNAS 114(28): 7222-7227.

Work done in collaboration with CAMERA

ExaFEL Challenge Problem

- **Goal:** Create an automated analysis pipeline for imaging of single particles via diffractive imaging
  - Provide rapid feedback to allow tuning sample concentrations during an experiment
  - Showcase scalability of the analysis pipeline by progressively increasing fraction of machine used while keeping number of diffraction images constant
  - Goal is for analysis time to be commensurate with data collection rates of 5 kHz
- **Challenge**:
  - Orientations and states of imaged particles are unknown
  - Images highly contaminated with noise
  - Number of useful images limited by achievable single-particle hit rates
  - Providing a workflow that allows for bursty data analysis
- **Criteria:**
  - Data Ingest: ingest diffraction images at 1 Tb/s
  - Memory: Store between .1 and 10 PB of event data. Each image is O(10 MB), each calculation requires 10^7 (minimum goal) to 10^9 (stretch goal) diffraction images
  - Workflow: ability to ingest data while calculation is ongoing, ability to exchange/average the electron density across nodes, ability offload the most compute intensive task to accelerators

Chemistry and Morphology of Combustion Aerosols (soot)

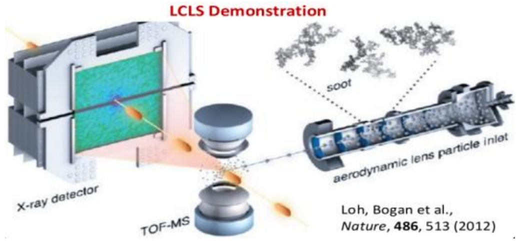

Understanding & Controlling Nano-materials Self-Assembly

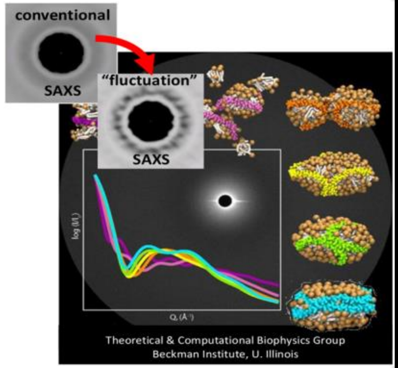

Coupled electronic & nuclear dynamics in

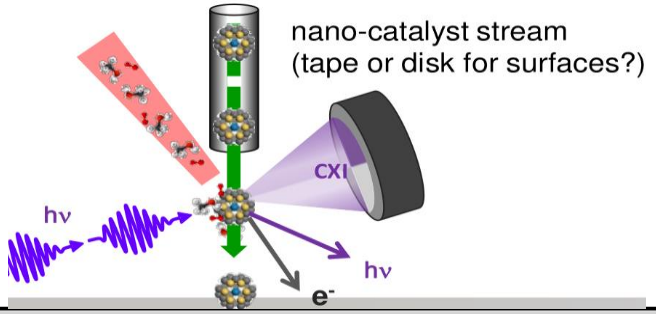

ExaFEL work combines algorithmic improvements, software architecture, workflow, and resource orchestration

Recent Accomplishments

- Introduction of Integration Optimization Triage and Analysis **algorithm increased number of images used for reconstruction by 2.7X** for nanocrystallography experiments
- **SZ data compression achieved 9:1 reduction** and speed of 180Mb/s, lossy compression shown to be useful for crystallography experiments
- Use **Legion task-based programming model** to improve scalability and portability for Psana framework for orchestration of data analysis at LCLS; showed performance comparable to MPI
- Architected a **solution to selectively request an uncongested data path** over ESnet between SLAC and NERSc
- Determined HDF5 unable to meet ExaFEL needs; **developed new custom data format XTC2**

Capability Development Plan for FY20-FY23

- Resource orchestration optimization: define and deploy capabilities for streaming data from beamline to HPC at 1 Tb/s, distributed images over supercomputer, **reduce job start up to <1 minute**
- **Generate simulated SPI data on Summit** – 1M images – used to test and benchmark M-TIP algorithms
- Identify and **port compute intensive steps in M-TIP to accelerators**; several candidates selected
- **Characterize levels of noise in data sets** and how it affects reconstruction
- Automate the workflow and **run M-TIP against actual experimental data for verification**
- Port and optimize for Aurora and Frontier architectures; **demonstrate challenge problem**

ExaFEL is on a good path to impact LCLS and LCLS-II

- Only ECP project tying BES facility resource needs to HPC to address unique needs of BES scientists
- Innovative and creative use of exascale computing resources: data science, workflow and resource orchestration aspects are all challenging
- Significant progress already and the next 4 years look very exciting - some of the development in team's ability to solve inverse problems at scale could be revolutionary for FEL science
- While the ExaFEL algorithmic development is focused on inverse problems, **all LCLS high throughput experiments will benefit from this development along with many other event driven experimental facilities**

EXAALT: Molecular Dynamics at the Exascale

PI: Danny Perez, LANL
Institutions: LANL, SNL, U Tennessee

EXAALT will improve molecular dynamics simulations for a wide range of problems

Goal: To develop a capability that extends the range of **accuracy**, **length** and **time** accessible via simulations; with a focus on what becomes possible at exascale levels of computing

Accuracy: from empirical to quantum
• **Low**: empirical models • **Medium/High**: data-driven/machine learned potentials (SNAP) • **High**: Density Functional Tight Binding (DFTB)

Length: from 10^2-10^14 atoms
• **Large size**: domain decomposition • **Intermediate size**: Force decomposition • **Small size**: Parallel solvers for DFTB

Time: from ps to s
• **Short**: MD with extended Lagrangian for DFTB • **Long**: Replica-based accelerated MD methods

EXAALT is creating an integrated computational framework that allows each of these dimensions of simulation space to be explored independently and together

EXASCALE COMPUTING PROJECT

EXAALT Approach – Replica-based Accelerated Molecular Dynamics

- Basic algorithm has existed for 20 years (Voter, BES support)
- Basic Algorithm:
  - Replicate system on M processors
  - Dephasing stage: Each Replica evolved forward with independent themostats for a time greater than the correlation time
  - Run M copies until a transition is detected on any processor; stop all processors
  - “Splice” trajectories from many parallel replicas into longer trajectory
- ECP is making this capability robust and available on exascale computing platforms; with a goal of obtaining wide use by the community

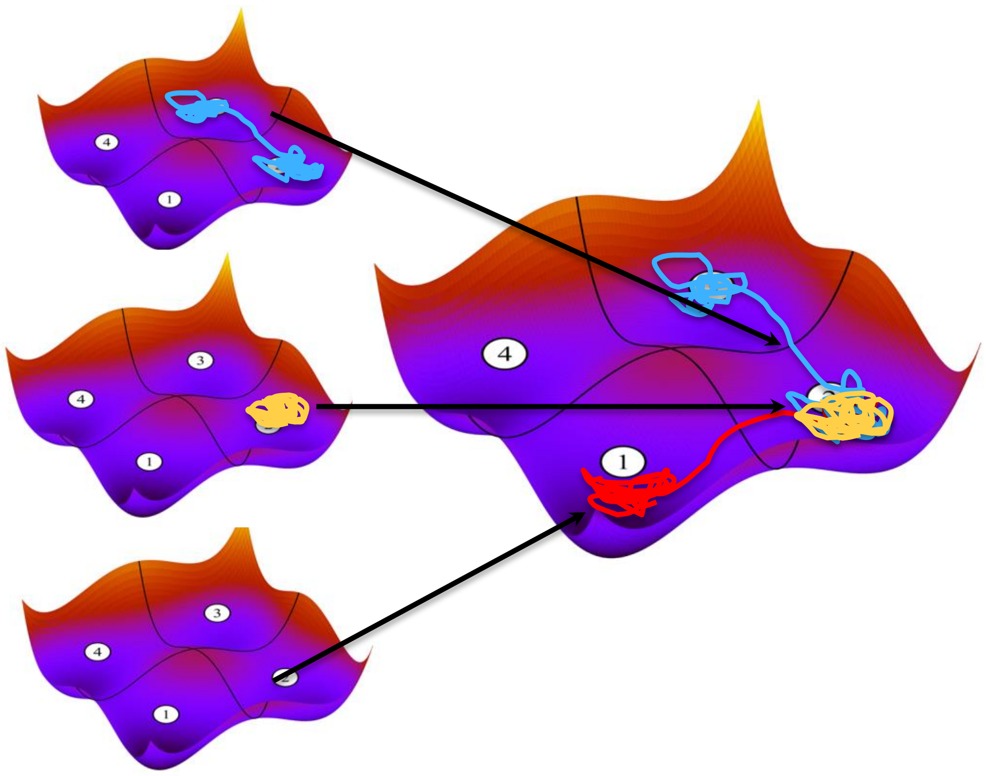

ParSplice
- Parallelizes in both space and time
- Accurately describes both the local energy environment and the transition statistics
- Trajectory can be made arbitrarily accurate

EXASCALE COMPUTING PROJECT

The ECP EXAALT project is creating an integrated code base with new

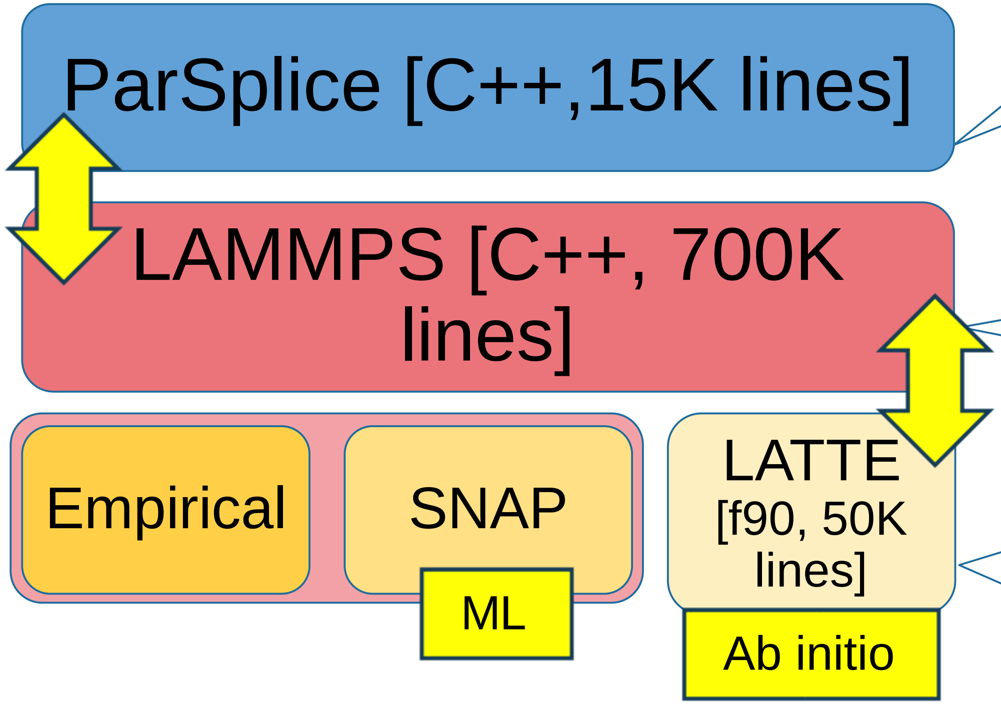

Task Management Layer

- **Implements replica-based accelerated MD algorithms**
- Data management through parallel key-value store

Molecular Dynamics Engine

- Integrates atomistic equations of motion
- Implements various parallelization strategies

Physical Models

- Describe materials properties at different levels of theory, from conventional empirical, data-driven, to semi-empirical quantum descriptions
- SNAP provides machine-learned forces and potentials that scale as O(N)
- LATTE DFTB ab initio approaches that scale as $O(N^3)$

Long range forces: Replica-based Accelerated Molecular Dynamics methods
(BES support : Voter, LANL)

EXAALT Nuclear Fusion Science Challenge Problem

- **Background:** Fusion as a commercially viable energy source requires advanced structural materials
  - Performance demands on plasma-facing components is beyond current day materials
  - Tungsten a leading contender, but experiments show the development of “fuzz” when exposed to low-energy plasmas containing helium
  - Impacts heat transfer and fuel retention, increases rates of erosion; embrittles divertor
  - No currently accepted models
- **Challenge problem**: Simulate microstructure evolution in fusion first wall
  - Develop a tool that can model the plasma-induced evolution of the surface microstructure of tungsten that leads to fuzz growth
  - Problem requires high accuracy, large system sizes and long time scales
- **Two cases**
  - Evolution from flat surface: empirical accuracy, 10^7 atoms, ~100 us
  - Evolution of helium in a tungsten tendril: machine learned SNAP, 10^5 atoms, 10 us

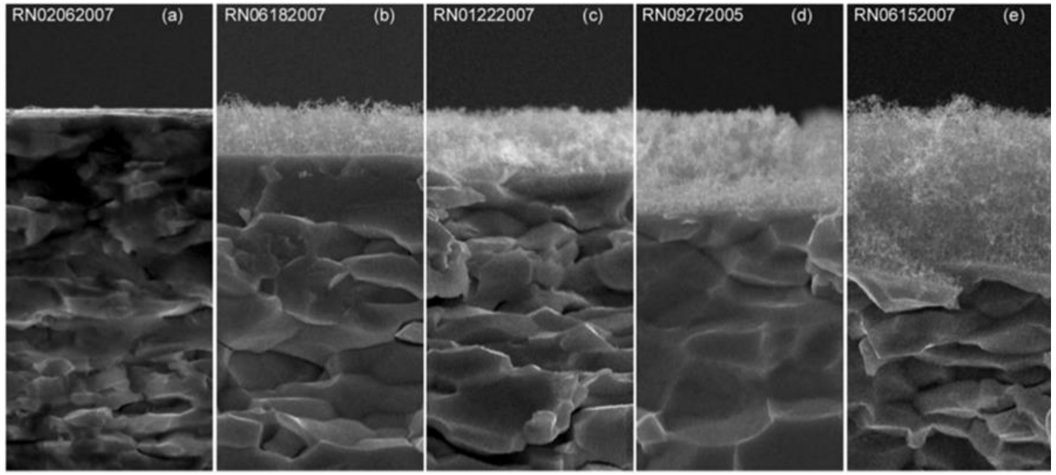

Nanoscale fuzz evolution on tungsten surface exposed to He plasma at T=1120 K.

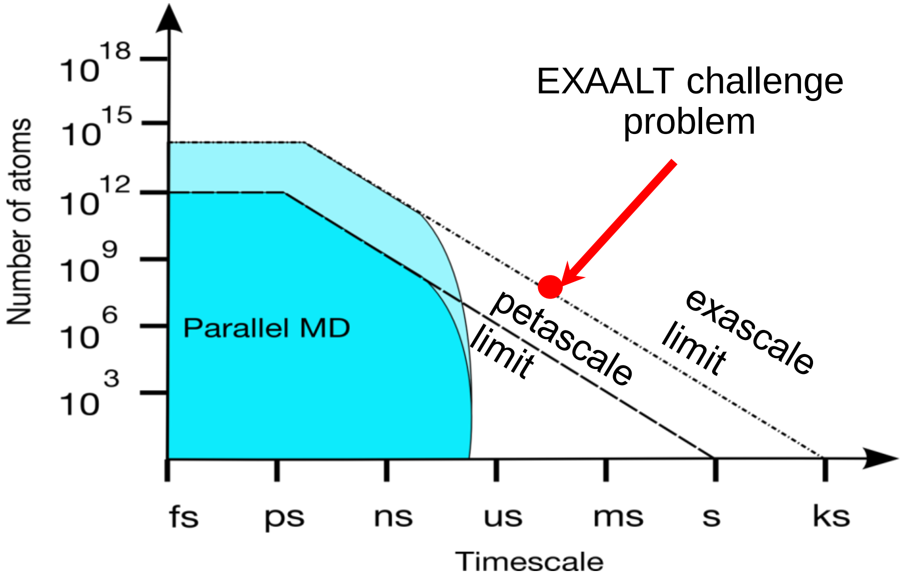

EXAALT work combines improvements of the physical models with optimization of the code base

Recent Accomplishments

- Integrated three main components of EXAALT
- **ParSplice** implements main replica-based accelerated MD methods; tested up to 300,000 worker processes **show scalability of techniques to extend time scales**
- **Domain decomposition** approaches with Sub-Lattice accelerated methods **extended to large numbers of atoms**
- Released and demonstrated baseline physical models for tungsten
- **Scaling demonstrated to full system** on Mira, Theta and Cori
- **Initial GPU port** to Summit architecture completed

Capability Development Plan for FY20-FY23

- Increased accuracy requires development of systematic approaches to generate necessary **training data**; massive parallelization of task management capability in **ParSplice will generate data with high information content and low redundancy**
- **Improve flexibility of SNAP potential** from quadratic terms in the descriptors of the atomic environments to general nonlinear combinations of descriptors **using neural networks**
- Develop specialized**, performance portable kernels** for exascale hardware based on Kokkos implementation developed in collaboration with CoPA
- Release **improved suite of solvers for LATTE**
- Optimize task-management framework at scale and **execute challenge problems**

Conclusions

- Robust capability for long time-scale molecular dynamics has potential to significantly extend reach of MD simulations
- Replica-based approach maps well to any scaled-up (replicated) platform
- Provides a new capability for real time interaction with quantum derived forces (LAMMPS – LATTE connection)
- Interesting incorporation of machine learning techniques to develop SNAP-based potentials

EXASCALE COMPUTING PROJECT

ECP is a DOE Order 413.3B project

"Provides DOE with program and project management direction for the acquisition of capital assets with the goal of delivering projects within the original performance baseline, cost and schedule, and fully capable of meeting mission performance, safeguards and security, and ES&H requirements." (https://www.directives.doe.gov)

- Mission Need Statement
  - Project requirements approved (Apr 14, 2016)
- Critical Decisions
  - ✓ CD-0: Mission Need (Jul 28, 2016)
  - ✓ CD-1: Alternative Selection & Cost Range (Jan 3, 2017)
  - CD-2: Performance Baseline (target Jan 2020)
  - CD-3: Start of Construction (target Jan 2020)
  - CD-4: Start of Operations or Project completion (target Sep 2023)
- Adherence to project management principles
  - Line management accountability, up-front planning, sound acquisition strategies, well-defined performance baselines, effective project management systems, integrated safety management, effective communication
- ECP has tailored many 413.3B requirements to meet the needs of an RD&D project

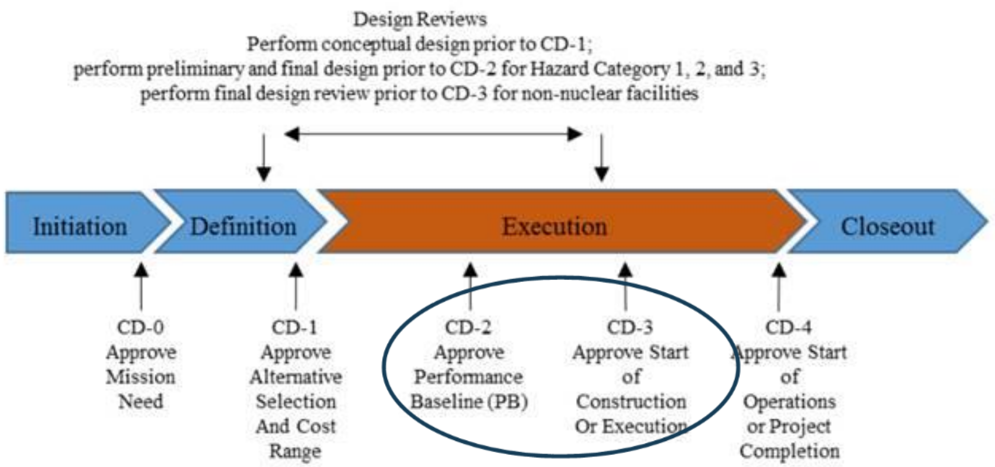

CD-2 & CD-3 are combined for ECP

Adhering to 413.3B formalisms has both opportunities and challenges for a large-scale RD&D project

Opportunities	Challenges

• Drives setting and adhering to performance metrics that are regularly measured
• Milestone-based management imparts a sense of urgency & enforces accountability
• Requires active risk management when it's often an oversight
• Forces communication when it's needed (e.g., when things are not going well)
• Forces decision points before it's too late
	• 413.3B not ideal for R&D as the order is currently laid out. Continuous improvement needed in overlaying project management formality with agile development
	  – Detailed milestone planning process often difficult for research and development project
	  – Challenging cultural shift for many ECP researchers; requires extensive education of ECP leadership
	• Critical dependence on DOE Facilities is a challenge; large projects in and of themselves with their own missions and priorities; required careful coordination and communication
	• Additional processes, while often adding value, take time and effort to implement and follow

Lessons learned by the ECP leadership team

A project as large as ECP has many benefits that could not be easily achieved with smaller projects

- Integration of activities and staff that depend (or should depend) upon and help one another
- Drives new efforts across the DOE complex, e.g., continuous integration or the extreme scale scientific software stack
- New or stronger staff R&D relationships and networks – career lasting – are established

Maintaining balance between formal project management and innovative R&D is a challenging endeavor

- Tailoring often drives innovative project management practices
- Preparing for and delivering on successful 413.3B reviews is a project in and of itself
- Relationship between formally-trained project managers and scientists should be peer, not customer-client

Communication is key to ECP success

- Use of Jira and confluence across ECP significantly improved cross-team and leadership to team communication
- Multi-institutional and geographically separated research teams and collaborations can function and execute well, even at this scale
- Dependence on the Facilities has required concentrated and ongoing interactions

What's next? Thoughts from the ECP perspective

ECP Software Artifacts (applications, libraries, tools)

- These code bases are "numerical laboratories" similar to a light source, research reactor, etc. so must be taken care of in a similar fashion; **maintenance and operations** to avoid bit-rot
- Most (and ideally all) ECP applications will be "ready to deliver science" (or more science). Need to support teams to **utilize the application codes for science** & engineering insights and breakthroughs
- ECP is catalyzing and nurturing a whole generation of "new and nontraditional" HPC applications (light source, additive manufacturing, grid, wind, urban, AI/ML, metagenomics). These **applications will need further aggressive developmen**t post ECP
- Software deployment pipeline (CI, packaging) will be first of it's kind for DOE software and should be continued but there is no natural home for these efforts

ECP Staff, Research, and Project Constructs

- Staff: A high risk for ECP is losing its staff in the latter two years. As ECP ramps down, staff **knowing what similar or related opportunities are next** helps to mitigate this risk (for example, other programs and projects utilizing or further extending ECP's products and solutions).
- Research: in applications and software stack should be nurtured; **re-examine the algorithms we've been using for decades**; possible to obtain more speed up than just hardware alone can provide
- Project Benefits: The **vertically integrated ECP structure** (apps, software, hardware working together) should be continued and leveraged going forward

We look forward to working with BES in this transition!

ECP EXASCALE COMPUTING PROJECT

Partnership with BES program management has helped ensure impact and relevance of BES application portfolio

- Ensuring impactful projects were selected for funding
- Ongoing interactions to ensure projects remain on point and impactful through…
  - PI briefings
  - Course corrections and improvement when needed
  - ECP leadership visits to DOE HQ
  - Attendance by BES program managers at the ECP annual meeting

- Special thanks to
  - Harriet Kung and her division leaders
  - Matthias Graf as POC for ECP

Questions?