The Department of Energy Exascale Computing Project (ECP)

EXASCALE COMPUTING PROJECT

Douglas B. Kothe (ORNL), ECP Director

DOE Fusion Energy Sciences Advisory Committee (FESAC) Meeting
May 25, 2022
Video Conference

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

I'm a Big Fusion Fan

and not just in movies . . .

Credits: The “Cult of the Code”

- **FLIP-PHD**
  - **2D imploding inertial confinement fusion (ICF) targets (DOE)**
- NASA-VOF2D, NASA-VOF3D
  - 2D/3D micro-gravity free surface flows (NASA micro-gravity)
- RIPPLE
  - 2D free surface flows (NASA micro-gravity; Xerox inkjet, …)
- CFDLIB
  - 2D/3D multiphase flows in fluidized catalytic crackers (Exxon); diaper making
- PAGOSA
  - Armor/anti-armor program
- POP
  - Global ocean circulation
- TRUCHAS/TELLURIDE
  - Casting/welding processes; spray forming; coastal hydrodynamics; Corporate Lethality Program (MDA)
- VERA (Virtual Environment for Reactor Applications) [R&D 100 Winner!]
  - Nuclear reactors

Before Photo: Me prior to agreeing to lead the ECP

Imploding ICF Target Design

Kothe's dissertation work circa 1985-87

Fluid-Implicit-Particle (FLIP) particle-in-cell (PIC) method, now a growing base technology for some ECP apps & also aminated movies! (e.g., Frozen, Moana, etc.)

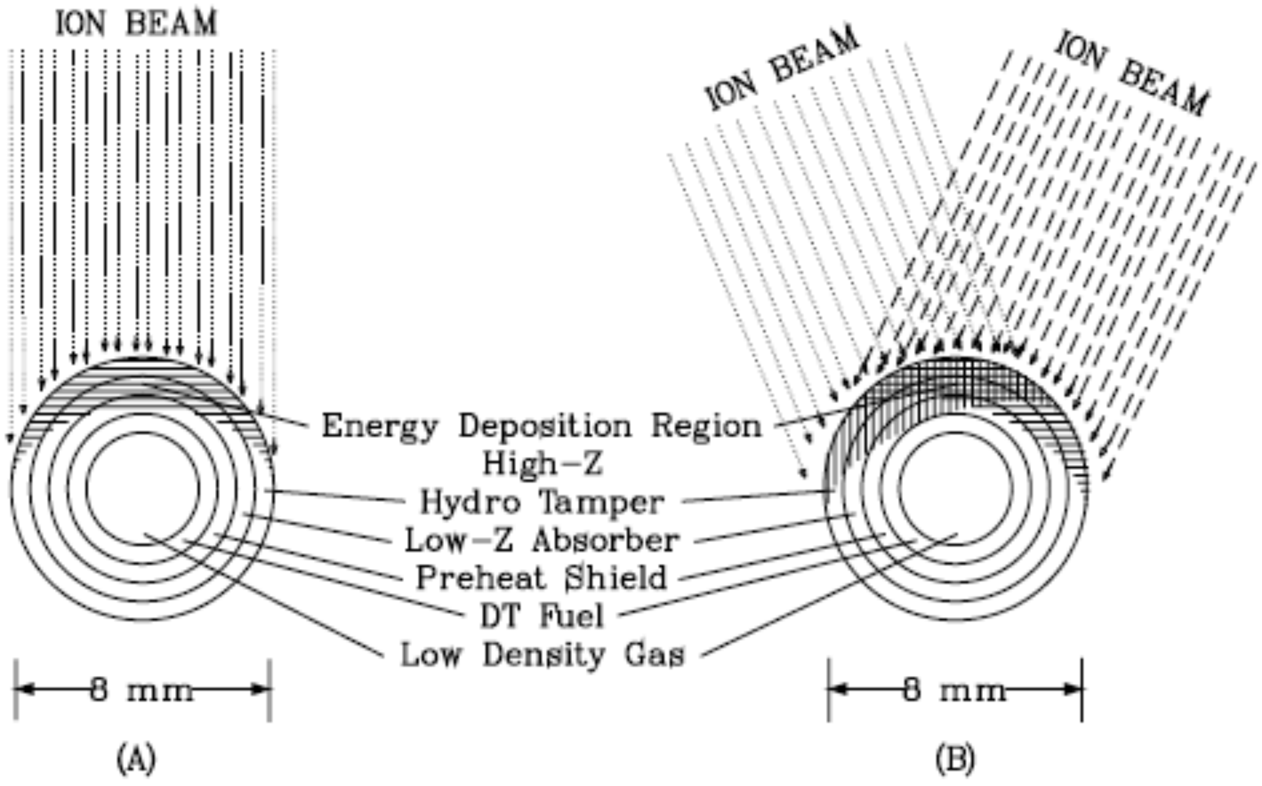

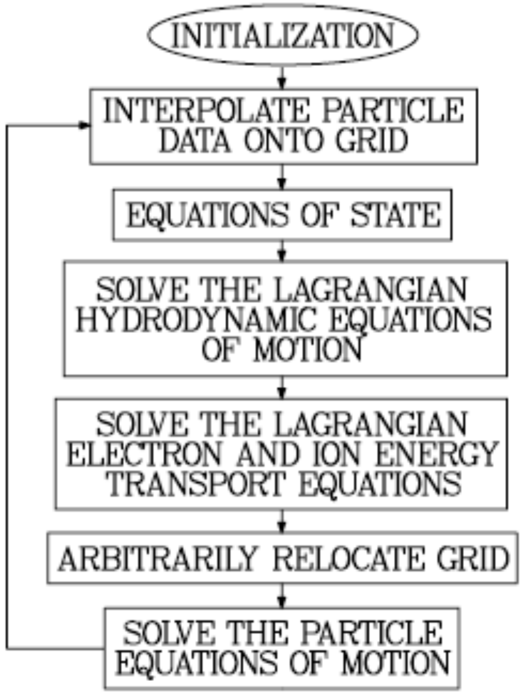

Unique characteristics of numerical method

- 2nd order time/space PIC method (FLIP) with inherent stability and fluid-like collisionality
- Adaptive grid (*not* AMR – it was the pre-AMR days)
- Discrete ray-tracing for ion beam penetration & energy deposition
- Natural ability to track interfaces via particle identity
- Innate sensitivity to unstable hydrodynamics (particle/grid Eulerian/Lagrangian duality)

Imploding ICF Target Design

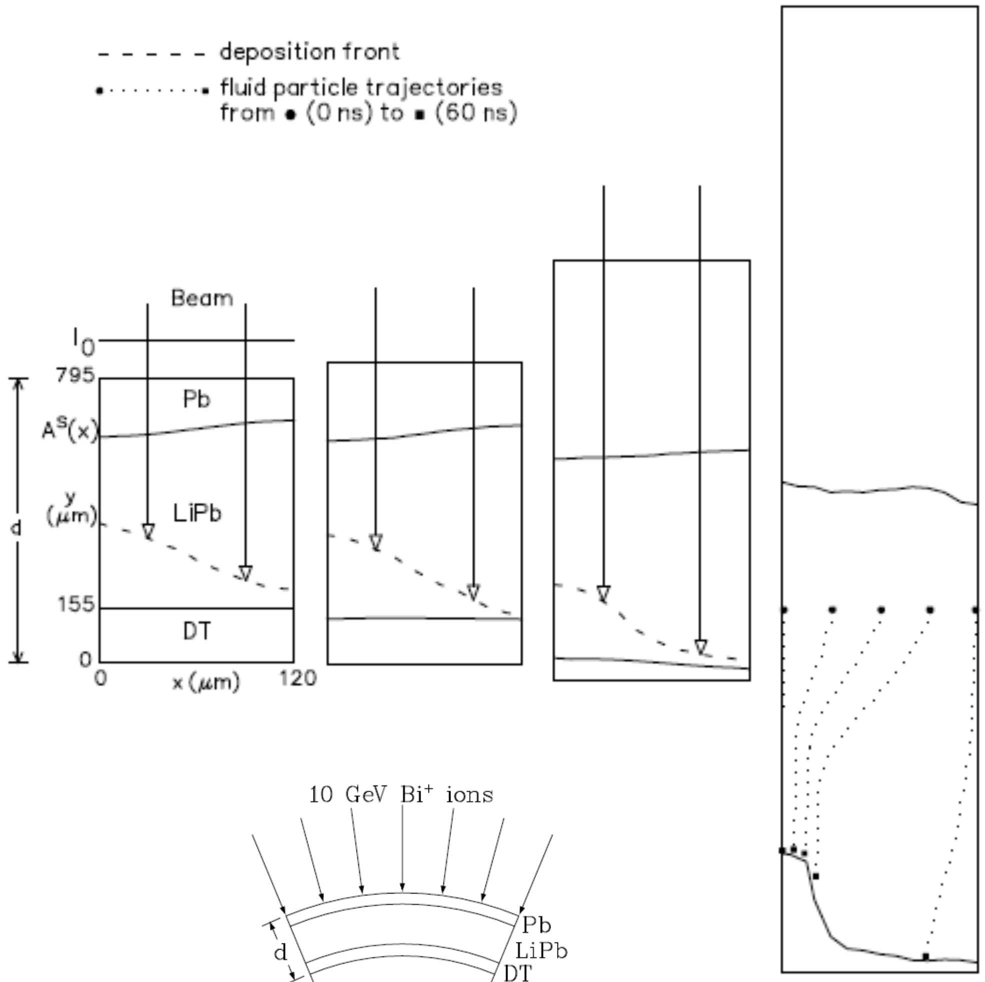

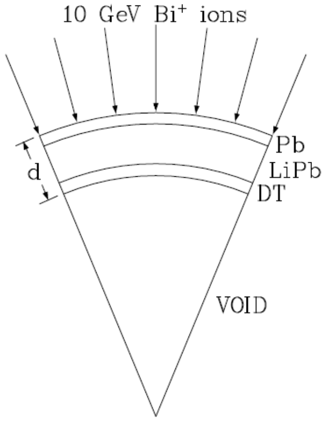

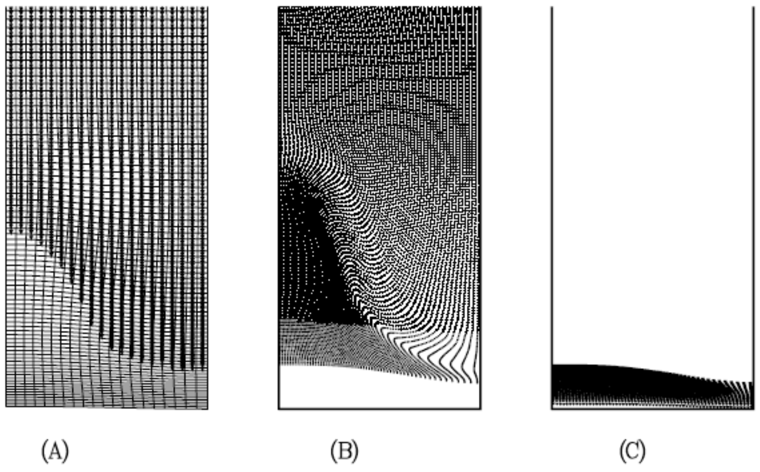

Characteristics of Implementation
- Fortran 77 (some inlined CAL)
- Linked lists for particles
- Kershaw's ICCG
- A memory/CPU hog
- CDC 7600, Cray XMP

DOE Exascale Computing Initiative (ECI)

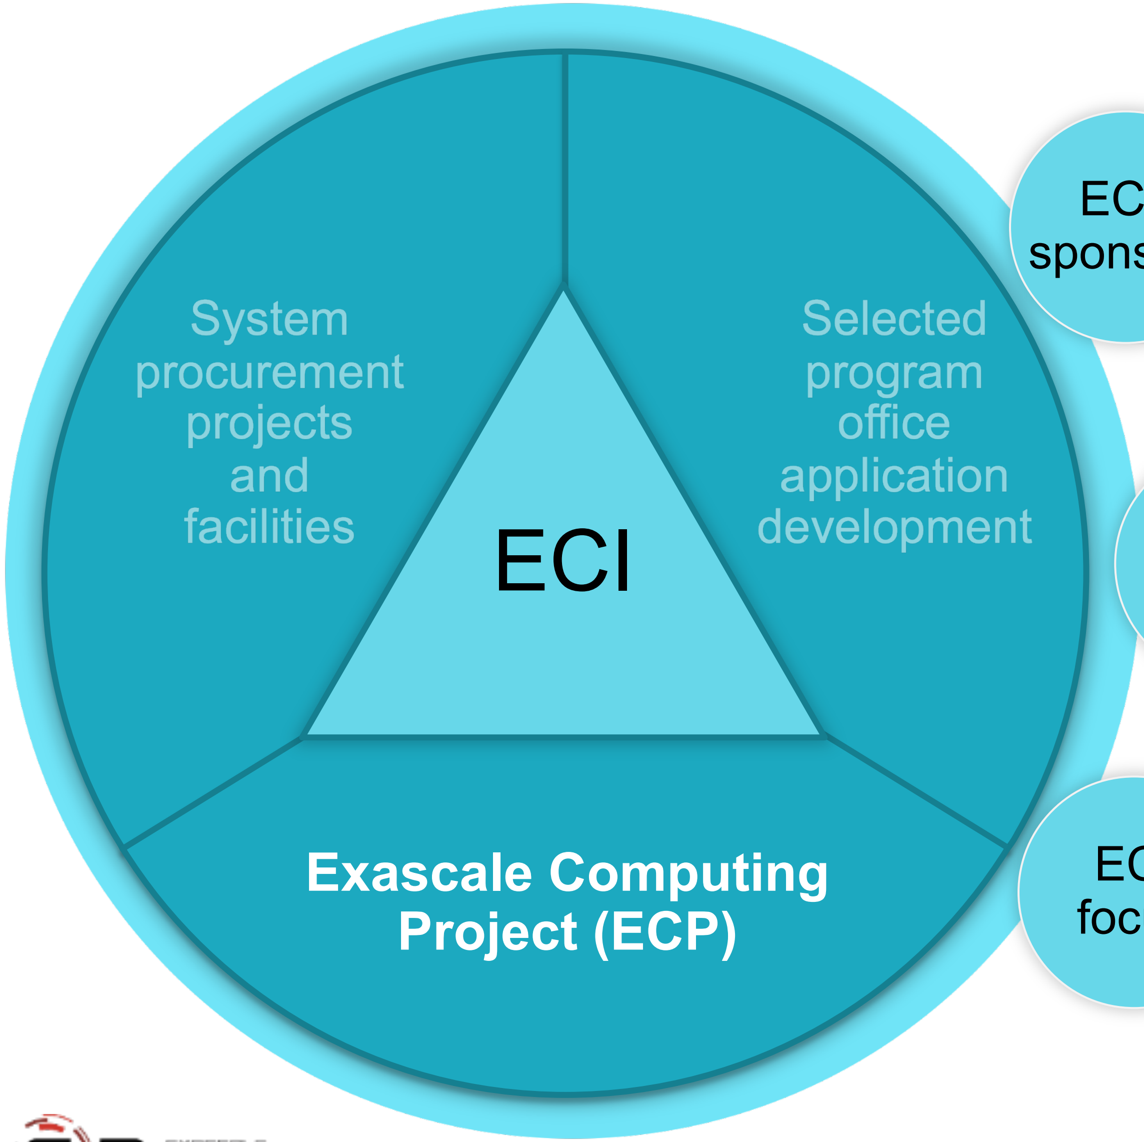

ECI sponsors: DOE Office of Science (SC); National Nuclear Security Administration (NNSA)

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

ECP is a critical component of the broader US ECI strategy

- Broader US Exascale Computing Initiative (ECI) elements are essential for success
  - Deploying exascale systems quickly enough to impact schedule-sensitive mission problems
  - Maintaining and advancing the "HPC ecosystem" after ECP
  - Developing U.S. industry and academia partnerships to ensure the benefits of advanced computing have broad and enduring impacts
- ECP Vision
  - *Accelerating innovation with exascale simulation and data science solutions that enhance US economic competitiveness, improve our quality of life, and strengthen our national security.*
- ECP Mission
  - *Deliver* 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
  - *Leverage* US HPC vendor R&D activities and products into DOE HPC exascale systems.

ECP enables future US revolutions in technology development, energy and national security, scientific discovery, economic security, and healthcare.

EXASCALE COMPUTING PROJECT

DOE HPC Roadmap to Exascale Systems

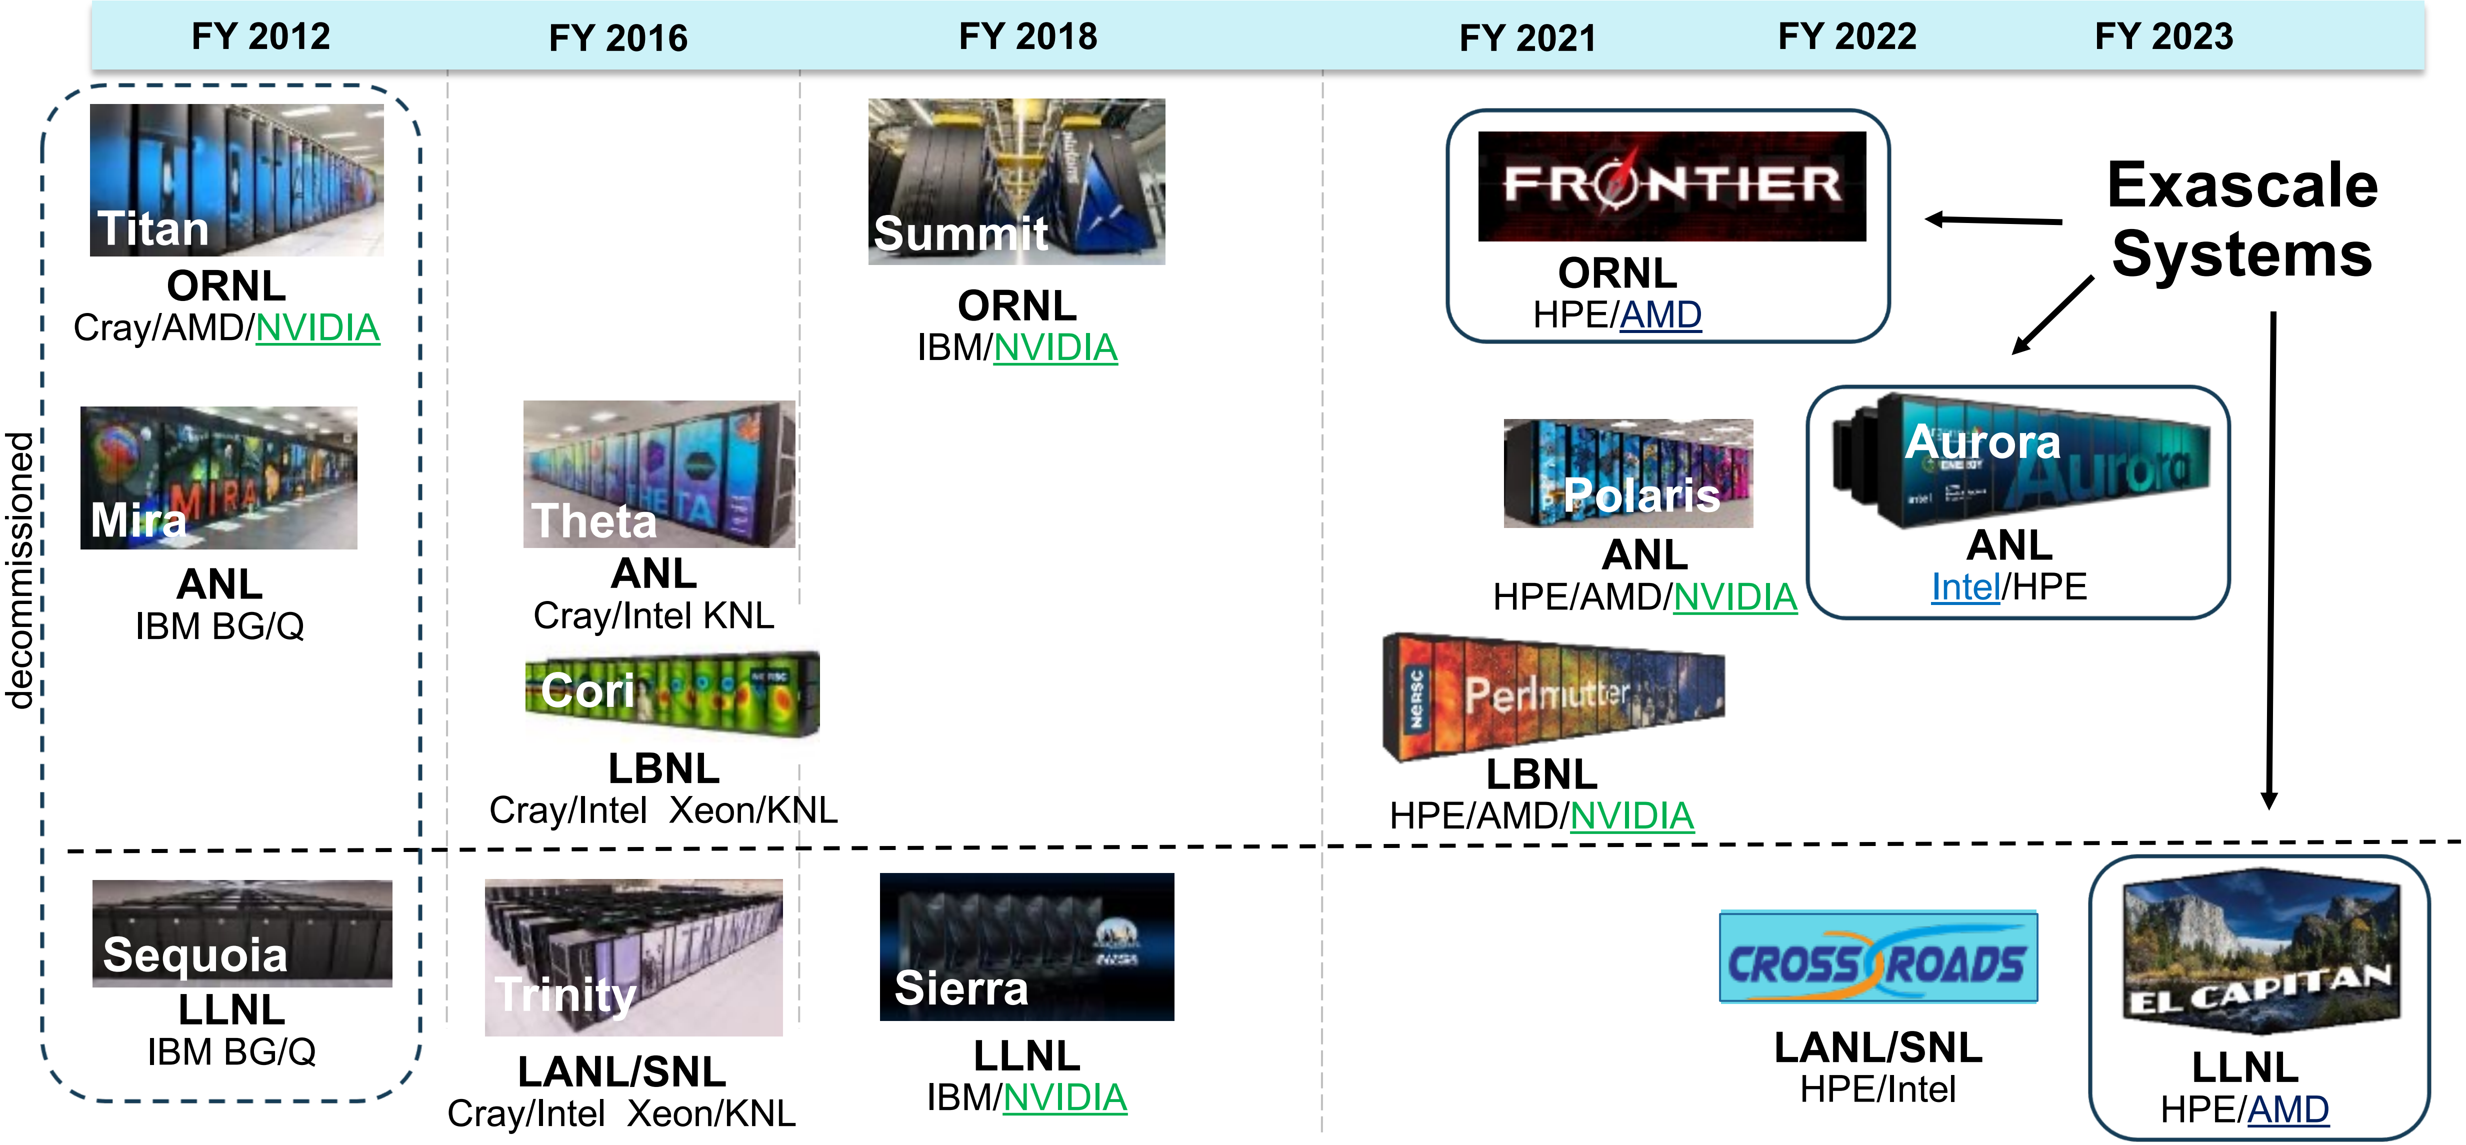

<u>Underlined</u>: GPUs

Each HPC system has served a vital role for ECP Teams

From benchmarking to development to now KPP demonstration

	Benchmark system for many ECP AD and ST teams	Multi-GPU system for scaling, algorithm & model dev, S/W design	Target system for KPP threshold demonstrations
System	**Titan (2012) Cray**	**Summit (2017) IBM**	**Frontier (2021) HPE**
Peak	27 PF	200 PF	> 1.5 EF
# nodes	18,688	4,608	9,408
Node	1 AMD Opteron CPU 1 NVIDIA Kepler GPU	2 IBM POWER9™ CPUs 6 NVIDIA Volta GPUs	1 AMD EPYC CPU 4 AMD Radeon Instinct GPUs
Memory		2.4 PB DDR4 + 0.4 HBM + 7.4 PB On-node storage	4.6 PB DDR4 + 4.6 PB HBM2e + 37 PB On-node storage, 66 TB/s Read 62 TB/s Write
On-node interconnect	PCI Gen2 No coherence across the node	NVIDIA NVLINK Coherent memory across the node	AMD Infinity Fabric Coherent memory across the node
System Interconnect	Cray Gemini network 6.4 GB/s	Mellanox Dual-port EDR IB 25 GB/s	Four-port Slingshot network 100 GB/s
Topology	3D Torus	Non-blocking Fat Tree	Dragonfly
Storage	32 PB, 1 TB/s, Lustre Filesystem	250 PB, 2.5 TB/s, IBM Spectrum Scale™ with GPFS™	695 PB HDD+11 PB Flash Performance Tier, 9.4 TB/s and 10 PB Metadata Flash
Power	9 MW	13 MW	29 MW

Frontier Node

All ECP teams have had access since Jan 2022

All GPUs and CPU are fully connected on node and have coherent shared memory

Custom AMD EPYC CPU (64 core)

- Supports Infinity Fabric
- Adds PCIe links for on node NVM (4 TB)
- 512 GB of DDR4 memory (1/2 TB per node)

Four AMD MI250X GPUs

- Announced by AMD November 8 2021
- 128 GB of HBM2e each (1/2 TB per node)
- 3.2 TB/s memory bandwidth

Each GPU is connected to a Slingshot NIC

- Eliminates GPU-CPU link bottleneck seen in Titan and Summit
- 1 GPU or CPU can use all NICS together

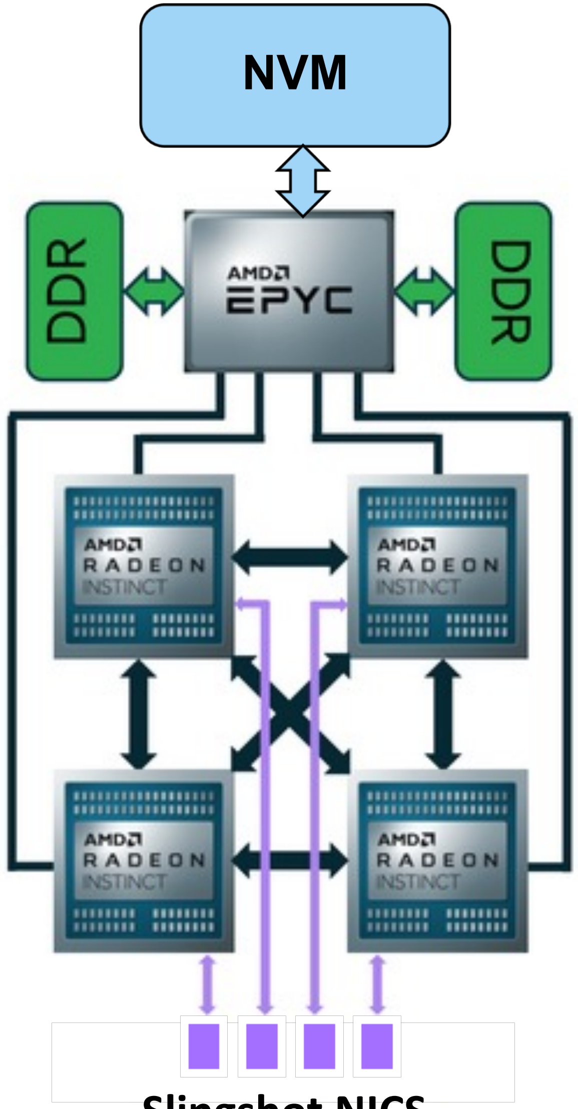

ECP by the Numbers

7
YEARS
$1.8B

A seven-year, $1.8B R&D effort that launched in 2016

6
CORE DOE
LABS

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

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

3
FOCUS
AREAS

Three technical focus areas: Hardware and Integration, Software Technology, Application Development supported by a Project Management Office

80+
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

ECP EXASCALE COMPUTING PROJECT

ECP's Technical Focus Areas

Providing the necessary components to meet national goals

Performant mission and science applications at scale

Aggressive RD&D project	Mission apps; integrated S/W stack	Deployment to DOE HPC Facilities	Hardware technology advances

Application Development (AD)

Develop and enhance the predictive capability of applications critical to DOE

24 applications
National security, energy, Earth systems, economic security, materials, data

6 Co-Design Centers
Machine learning, graph analytics, mesh refinement, PDE discretization, particles, online data analytics

Software Technology (ST)

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

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

Hardware and Integration (HI)

Integrated *continuous testing & 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

ECP Organization

Barb Helland
ASCR Program Manager

Thuc Hoang
ASC Program Manager

John Young
Federal Project Director

Board of Directors
Thomas Zacharia, Chair (Director, ORNL)
Thom Mason, Vice Chair (Director, LANL)

Laboratory Operations Task Force (LOTF)
Bruce Hendrickson (LLNL), Chair

Al Geist, ORNL
Chief Technology Officer

Exascale Computing Project
Doug Kothe, ORNL
Project Director
Lori Diachin, LLNL
Deputy Project Director

Industry and Agency Council
Dave Kepczynski, GE, and
Fran Hill, DoD, Co-Chairs
David Martin, ANL, and
Suzy Tichenor, ORNL,
Co-Executive Directors

DOE HPC Facilities
ALCF, OLCF, NERSC,
LANL, LLNL, SNL

Core Laboratories

Oak Ridge National Laboratory
Argonne National Laboratory
Berkeley Lab – Lawrence Berkeley National Laboratory
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Sandia National Laboratories

Julia White, ORNL
Technical Operations

Mike Bernhardt, ORNL
Communications & Outreach

Project Management
Tom Cook, ORNL
Director

Shaun Fomby
Project Controls & Risk

Jonathan Wilson
Information Technology

Project Office Support
Megan Fielden, Human Resources
Sam Howard, Export Control Analyst
Mike Hulsey, Business Management
Michael Johnson, Legal
Willy Besancenez, Procurement
Kim Milburn, Finance Officer
Susan Ochs, Partnerships
and Points of Contacts at the
Core Laboratories

Application Development
Andrew Siegel, ANL
Director
Erik Draeger, LLNL
Deputy Director

Software Technology
Mike Heroux, SNL
Director
Lois Curfman McInnes, ANL
Deputy Director

Hardware & Integration
Katie Antypas, LBNL
Director
Susan Coghlan, ANL
Deputy Director

ECP Industry **and Agency** Council Members

- Altair
- Ansys
- BOEING
- CASCADE TECHNOLOGIES
- Chevron
- Cummins
- EPRI
- ExxonMobil
- GE
- gm
- GOODYEAR
- KatRisk
- Lilly
- P&G
- Raytheon Technologies
- tae TECHNOLOGIES
- Westinghouse
- Whirlpool CORPORATION
- DEPARTMENT OF DEFENSE UNITED STATES OF AMERICA
- NOAA NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION U.S. DEPARTMENT OF COMMERCE
- NSF
- NASA
- NIH National Institutes of Health

ECP Broader Engagement Initiative

A multipronged initiative to expand the pipeline and workforce for DOE high-performance computing (HPC)

HPC Workforce Development and Retention Action Group

We are influencing culture in DOE labs and communities to promote the workforce pipeline for — and the retention of — a diverse DOE lab HPC workforce.

We are fostering a community, within

Intro to HPC

We are providing accessible introductory material to HPC — thereby addressing gaps in — and expanding the pipeline of — people with foundational HPC skills.

This becomes a pathway to build experience for (and interest in)

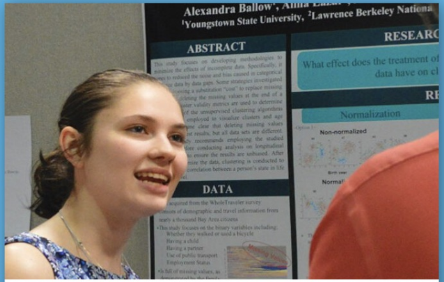

Sustainable Research Pathways for HPC (SRP-HPC)

We are establishing a multilab cohort of students from underrepresented groups (and faculty working with them), who are working side-by-side with ECP teams on world-class HPC projects:

https://www.exascaleproject.org/hpc-workforce

Reference: *A multipronged approach to building a diverse workforce and cultivating an inclusive professional environment for DOE high-performance computing*, response to DOE RFI on Software Stewardship, ECP Task Force on Broader Engagement, Dec 2021, https://doi.org/10.6084/m9.figshare.17192492

Partnership with **Sustainable Horizons Institute**
https://shinstitute.org/srp-hpc

SUSTAINABLE HORIZONS INSTITUTE

Strongly encouraged to apply: Students from (and faculty working with) underrepresented groups (Black or African American, Hispanic/Latinx, American Indian, Alaska Native, Native Hawaiian, and Pacific Islanders, women, persons with disabilities, first-generation scholars, and other underrepresented populations)

Why ECP? Unique multilab partnership across DOE computing sciences (apps / math / CS / facilities)

- Strength in spanning multiple institutions / strength in numbers / network beyond what any individual lab could do
- Proactive outreach and deployment of DOE HPC tools and technologies to communities beyond traditional targets

ECP Project Management (PM)

Measure progress and ensure execution within scope, schedule, and budget
Processes tailored according to DOE Order 413.3B

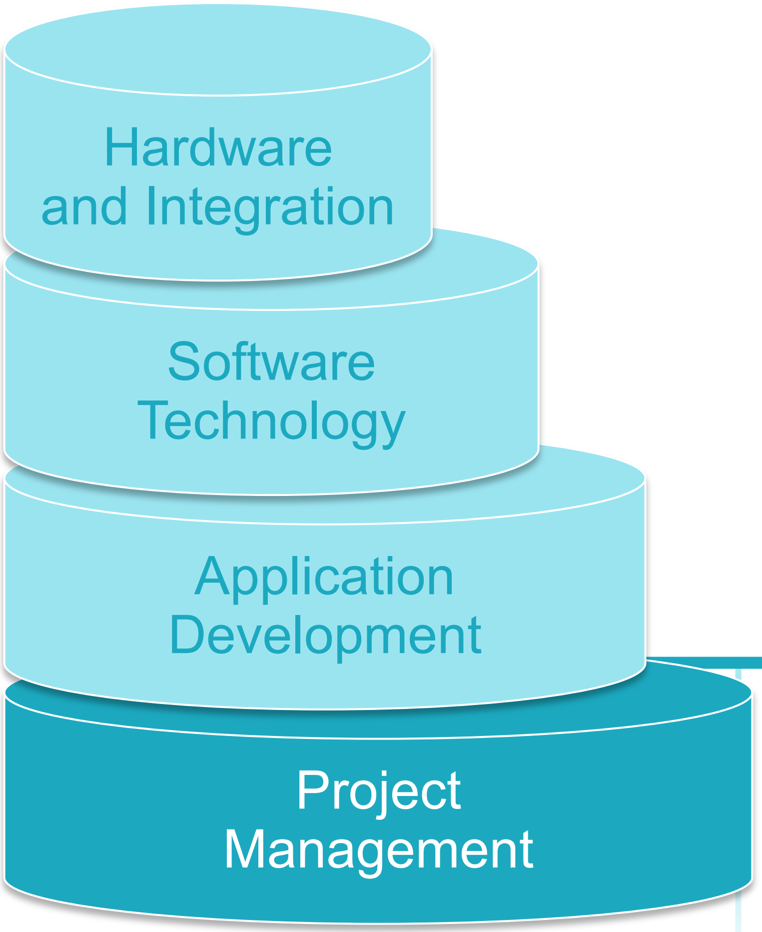

ECP overall

- 88 WBS L4 subprojects executing RD&D
- 2,561 L4 subproject (P6) milestones delivered in FY17–FY22Q1
  - 456 in FY20
  - 44 in FY22Q1 through November 2021
- 527 L4 subproject (P6) milestones planned in FY22–FY23
  - Plus 61 planning activities

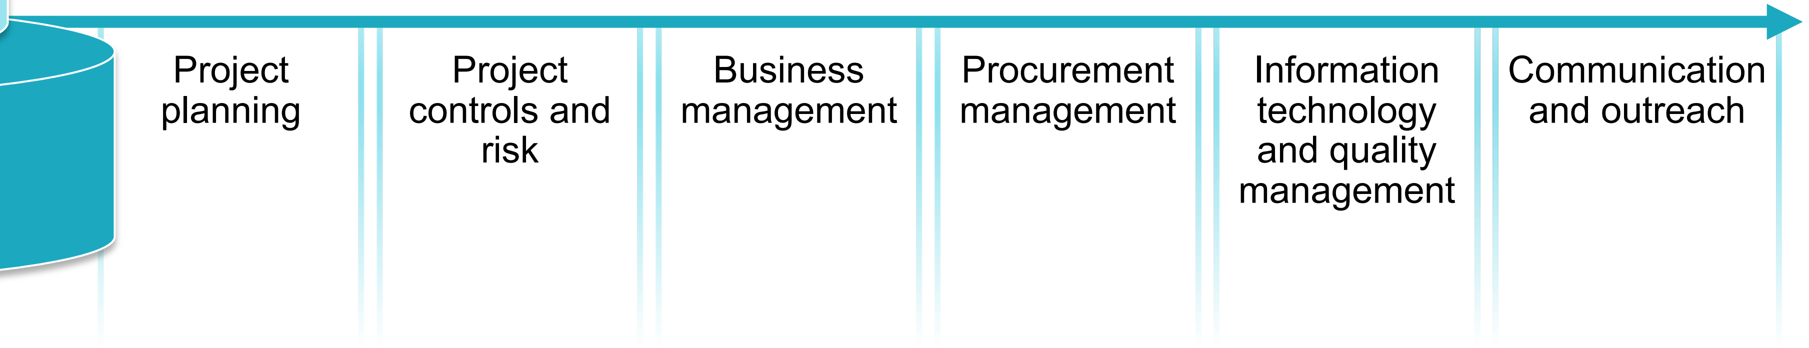

Merits of Effective Projectization of R&D

- Milestone-based planning and tracking imparts a sense of urgency & enforces accountability (if you have the “right” milestones)
- Drives setting and adhering to performance metrics that are regularly measured
- Can actually improve breadth and depth of science output
- Forces communication when it’s needed (e.g., when things are not going well)
- Rewards teaming to achieve project goals
- Helps to mentor and train next generation leaders
- Brings helpful process into potentially chaotic situations
- Forces decision points before it’s too late
- Requires active risk management when it’s often an oversight

ECP Application Development (AD) Focus Area

National security	Energy security	Economic security	Scientific discovery	Earth system	Health care
Next-generation, **stockpile stewardship** codes	Turbine **wind plant** efficiency	**Additive manufacturing** of qualifiable metal parts	**Cosmological probe** of the standard model of particle physics	Accurate regional impact assessments in **Earth system models**	Accelerate and translate **cancer research** (partnership with NIH)
Reentry-vehicle- environment simulation	Design and commercialization of **SMRs**	Reliable and efficient planning	Validate fundamental laws of nature	Stress-resistant crop analysis and catalytic	
Multi-ph simulat **ener physic**	combustion		Demystify **origin of chemical elements**		
	Biofuel catalyst design				

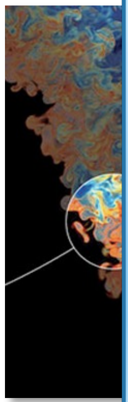

- Many complex apps on new first-of-kind HPC (exascale) systems.
- AD includes 24 projects ranging from materials science to the simulation of complex. engineered systems for energy generation to climate, astrophysics and cosmology
- AD teams: ~10 people with diverse, tightly integrated expertise
- AD depends heavily on hardware procurement/deployment and enabling software technologies managed/developed in other areas of ECP.

EXASCALE COMPUTING PROJECT

ECP Application Development (AD) Focus Area

National security

Next-generation, **stockpile stewardship** codes

Reentry-vehicle- environment simulation

Multi-ph simulat **energ physic**

Energy security

Turbine **wind plant** efficiency

Design and commercialization of **SMRs**

combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Reliable and efficient planning

Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Demystify **origin of chemical elements**

Earth system

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

Stress-resistant crop analysis and catalytic

Health care

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

The 24 AD application projects

- Include **78 separate codes**
- Represent over **10 million lines of code**
- In some cases support large user communities
- Covering broad range of mission critical science and engineering domains, **many of which directly apply to fusion science and fusion technology**
- Mostly started with MPI or MPI+OpenMP on CPUs

ECP's Key Performance Parameters (KPPs)

Quantified with Explicit Targets

KPP ID	Description of Scope	Threshold KPP	Objective KPP	Verification Action/Evidence
KPP-1	11 selected applications demonstrate performance improvement for mission-critical problems	6 of 11 applications demonstrate Figure of Merit improvement ≥50 on their base challenge problem	All 11 selected applications demonstrate their stretch challenge problem	Independent assessment of measured FOM results and base challenge problem demonstration evidence
KPP-2	14 selected applications broaden the reach of exascale science and mission capability	5 of 10 DOE Science and Applied Energy applications *and* 2 of 4 NNSA applications demonstrate their base challenge problem	All 14 selected applications demonstrate their stretch challenge problem	Independent assessment of base challenge problem demonstration evidence
KPP-3	76 software products selected to meet an aggregate capability integration score	Software products achieve an aggregate capability integration score of at least 34 out of a possible score of 68	Software products achieve the maximum aggregate capability integration score of 68	Independent assessment of each software product's capability integration score
KPP-4	Delivery of 267 vendor baselined milestones in the PathForward element	✓ Vendors meet 214 out of the total possible 267 PathForward milestones	✓ Vendors meet all 267 possible PathForward milestones	Independent review of the PathForward milestones to assure they meet the contract requirements; evidence is the final milestone deliverable

ECP Applications

Risks to achieving their base challenge problems are becoming more focused and measurable

Domain*	Base Challenge Problem	Risks and Challenges
Wind Energy	2x2 5 MW turbine array in 3x3x1 km^3 domain	Linear solvers; structured / unstructured overset meshes
Nuclear Energy	Small Modular Reactor with complete in-vessel coolant loop	Coupled CFD + Monte Carlo neutronics; MC on GPUs
Fossil Energy	Burn fossil fuels cleanly with CLRs	AMR + EB + DEM + multiphase incompressible CFD
Combustion	Reactivity controlled compression ignition	AMR + EB + CFD + LES/DNS + reactive chemistry
Accelerator Design	TeV-class 10^{2-3} times cheaper & smaller	AMR on Maxwell's equations + FFT linear solvers + PIC
Magnetic Fusion	Coupled gyrokinetics for ITER in H-mode	Coupled continuum delta-F + stochastic full-F gyrokinetics
Nuclear Physics: QCD	Use correct light quark masses for first principles light nuclei properties	Critical slowing down; strong scaling performance of MG-preconditioned Krylov solvers
Chemistry: GAMESS	Heterogeneous catalysis: MSN reactions	HF + MP2 + coupled cluster (CC) + fragmentation methods
Chemistry: NWChemEx	Catalytic conversion of biomass	CCSD(T) + energy gradients
Extreme Materials	Microstructure evolution in nuclear matls	AMD via replica dynamics; OTF quantum-based potentials
Additive Manufacturing	Born-qualified 3D printed metal alloys	Coupled micro + meso + continuum; linear solvers

EXASCALE COMPUTING PROJECT

***KPP-1 application, KPP-2 application**

ECP Applications

Risks to achieving their base challenge problems are becoming more focused and measurable

Domain*	Challenge Problem	Computational Hurdles
Quantum Materials	Predict & control matls @ quantum level	Parallel on-node perf of Markov-chain Monte Carlo; OpenMP
Astrophysics	Supernovae explosions, neutron star mergers	AMR + nucleosynthesis + GR + neutrino transport
Cosmology	Extract “dark sector” physics from upcoming cosmological surveys	AMR or particles (PIC & SPH); subgrid model accuracy; in-situ data analytics
Earthquakes	Regional hazard and risk assessment	Seismic wave propagation coupled to structural mechanics
Geoscience	Well-scale fracture propagation in wellbore cement due to attack of CO_2-saturated fluid	Coupled AMR flow + transport + reactions to Lagrangian mechanics and fracture
Earth System	Assess regional impacts of climate change on the water cycle @ 5 SYPD	Viability of Multiscale Modeling Framework (MMF) approach for cloud-resolving model; GPU port of radiation and ocean
Power Grid	Large-scale planning under uncertainty; underfrequency response	Parallel nonlinear optimization based on discrete algebraic equations; multi-period optimization
Cancer Research	Scalable machine learning for predictive preclinical models and targeted therapy	Increasing accelerator utilization for model search; exploiting reduced/mixed precision; resolving data management or communication bottlenecks
Metagenomics	Discover and characterize microbial communities through genomic and proteomic analysis	Graph algorithms, distributed hashing, matrix operations and other discrete algorithms
FEL Light Source	Protein and molecular structure determination using streaming light source data	Parallel structure determination for ray tracing and single-particle imaging

***KPP-1 application, KPP-2 application**

FY19-21 Summit Performance for ECP's KPP-1 Applications

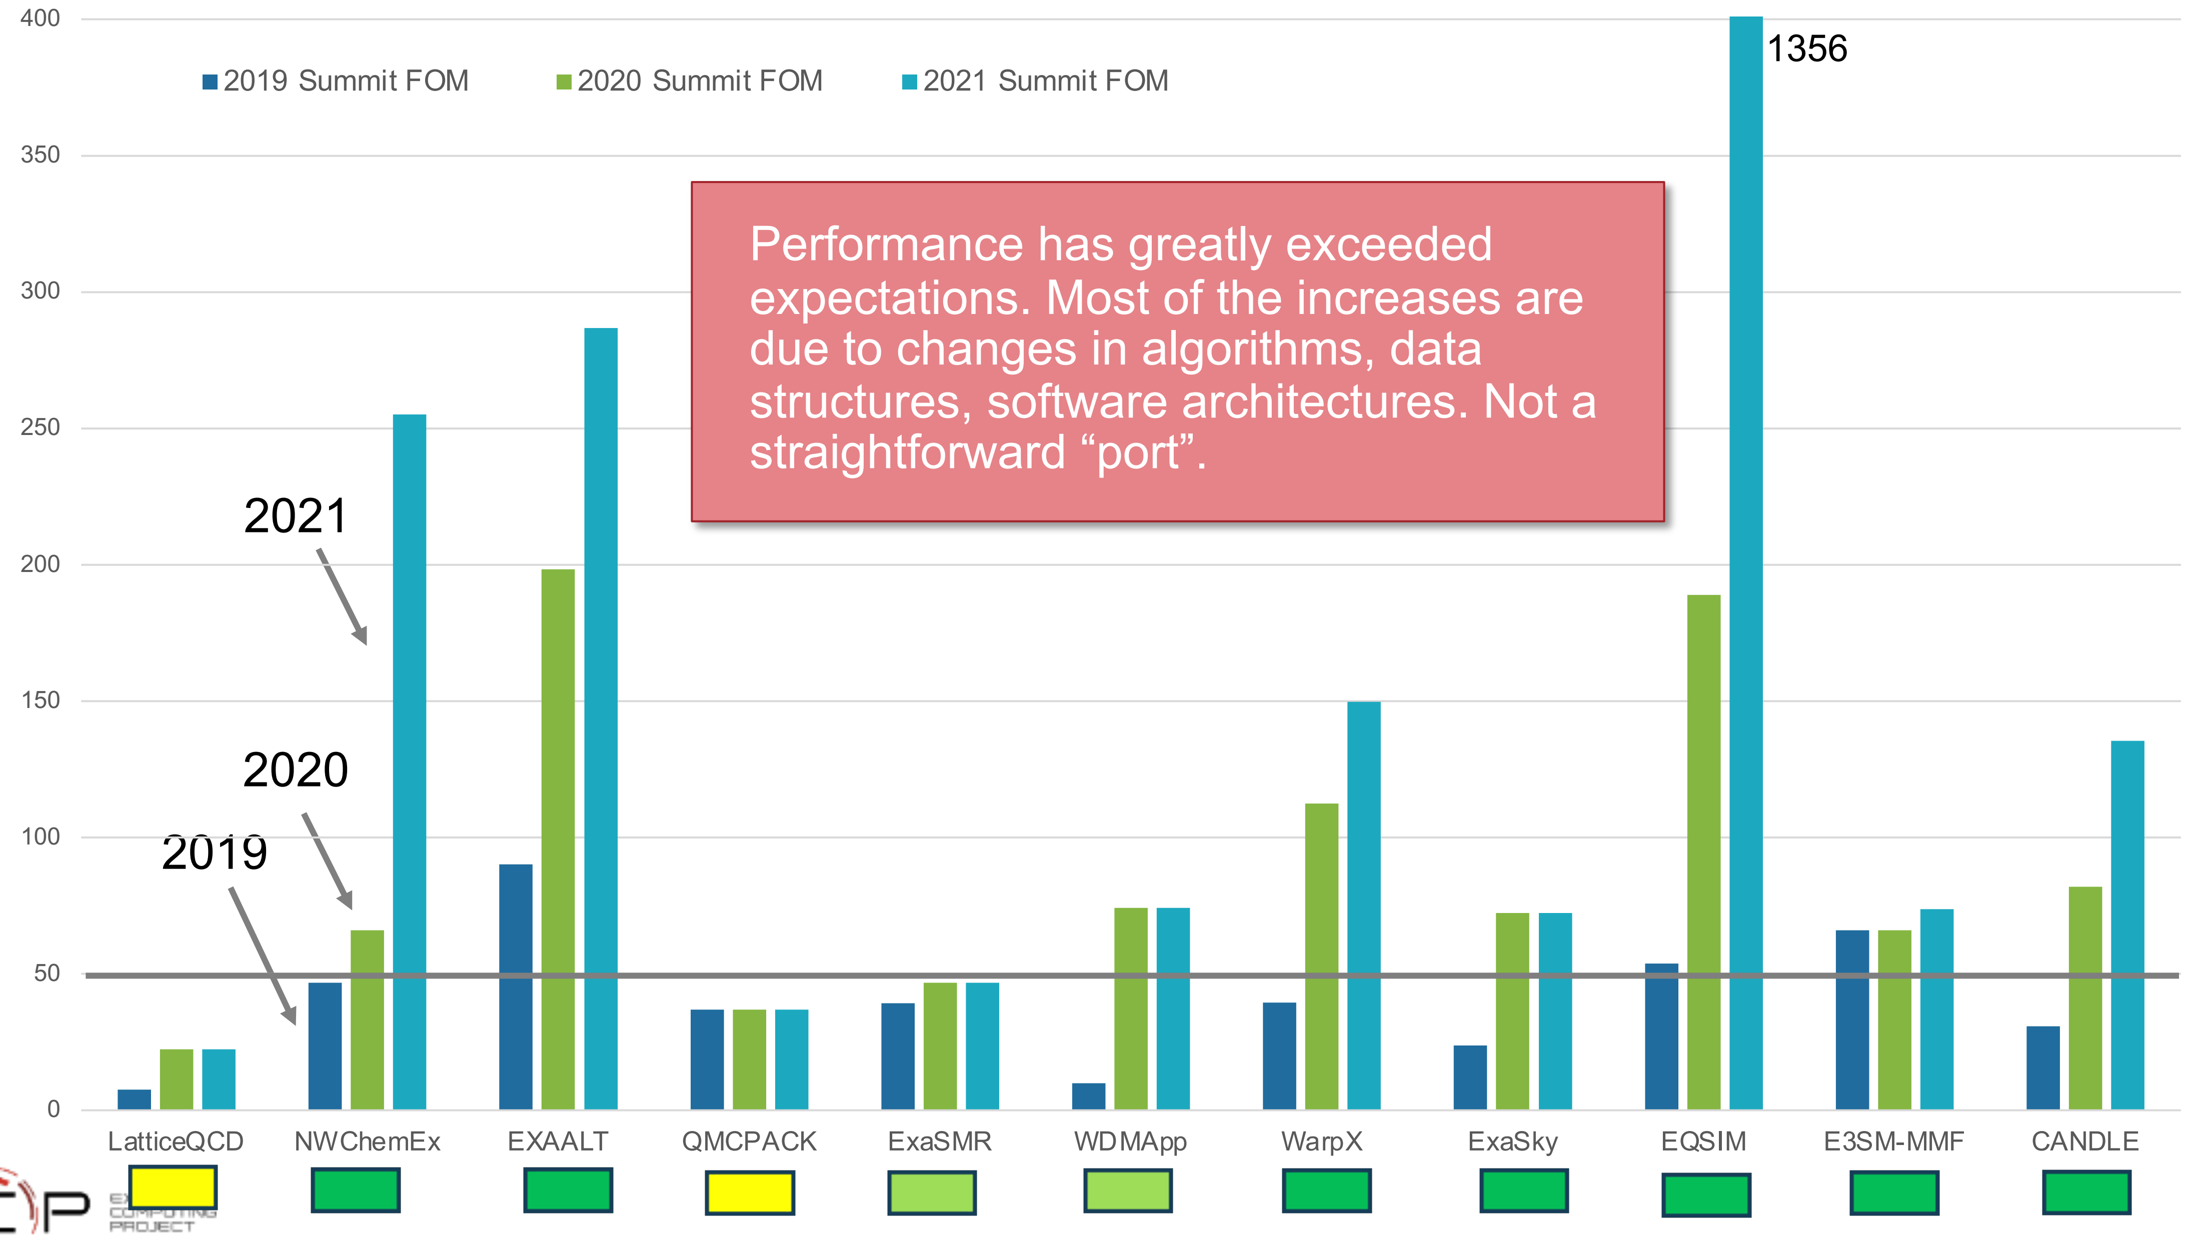

Performance has greatly exceeded expectations. Most of the increases are due to changes in algorithms, data structures, software architectures. Not a straightforward "port".

Application Motifs* (what's the app footprint?)

Algorithmic methods that capture a common pattern of computation and communication

1. **Dense Linear Algebra**
   - Dense matrices or vectors (e.g., BLAS Level 1/2/3)
2. **Sparse Linear Algebra**
   - Many zeros, usually stored in compressed matrices to access nonzero values (e.g., Krylov solvers)
3. **Spectral Methods**
   - Frequency domain, combining multiply-add with specific patterns of data permutation with all-to-all for some stages (e.g., 3D FFT)
4. **N-Body Methods (Particles)**
   - Interaction between many discrete points, with variations being particle-particle or hierarchical particle methods (e.g., PIC, SPH, PME)
5. **Structured Grids**
   - Regular grid with points on a grid conceptually updated together with high spatial locality (e.g., FDM-based PDE solvers)
6. **Unstructured Grids**
   - Irregular grid with data locations determined by app and connectivity to neighboring points provided (e.g., FEM-based PDE solvers)
7. **Monte Carlo**
   - Calculations depend upon statistical results of repeated random trials
8. **Combinational Logic**
   - Simple operations on large amounts of data, often exploiting bit-level parallelism (e.g., Cyclic Redundancy Codes or RSA encryption)
9. **Graph Traversal**
   - Traversing objects and examining their characteristics, e.g., for searches, often with indirect table lookups and little computation
10. **Graphical Models**
    - Graphs representing random variables as nodes and dependencies as edges (e.g., Bayesian networks, Hidden Markov Models)
11. **Finite State Machines**
    - Interconnected set of states (e.g., for parsing); often decomposed into multiple simultaneously active state machines that can act in parallel
12. **Dynamic Programming**
    - Computes solutions by solving simpler overlapping subproblems, e.g., for optimization solutions derived from optimal subproblem results
13. **Backtrack and Branch-and-Bound**
    - Solving search and global optimization problems for intractably large spaces where regions of the search space with no interesting solutions are ruled out. Use the divide and conquer principle: subdivide the search space into smaller subregions ("branching"), and bounds are found on solutions contained in each subregion under consideration

The Landscape of Parallel Computing Research: A View from Berkeley, Technical Report No. UCB/EECS-2006-183 (Dec 2006).

7 Computational Giants of Massive Data Analysis*

1. **Basic statistics**
   - Mean, variance (and other moments), median, sorting, clustering, # of distinct and frequently-occurring elements in a data set; O(N) calculations for N data points.
2. **Generalized *N*-body problem**
   - Distances, kernels, or other similarities between (all or many) pairs (or higher-order n-tuples) of points. Computational complexity $O(N^2)$ or $O(N^3)$
3. **Graph-theoretic computations**
   - Traversing a graph where the graph is the data or the statistical model takes the form of a graph. Common statistical computations include betweenness, centrality, commute distances; used to identify nodes or communities of interest
4. **Linear algebraic computations**
   - Linear systems, eigenvalue problems, inverses, many of which result from linear models, e.g., linear regression, PCA. Differentiator: statistical problems such as the optimization in learning, e.g., eigendecomposition of PCA to optimize a linear convex training error not necessarily requiring high accuracy. Differentiator: multivariate statistics arguably has its own matrix form, that of a kernel (or Gram) matrix

**Frontiers In Massive Data Analysis, National Research Council (The National Academies Press, 2013)*

7 Computational Giants of Massive Data Analysis*

5. **Optimization**
   - All subclasses of optimizations, from unconstrained to constrained, both convex and non-convex. Non-trivial optimizations in statistical methods, e.g., linear and quadratic programming, second-order cone programming in support vector machines, more recent classifiers, recent manifold learning methods, e.g. maximum variance unfolding

6. **Integration**
   - Integration of functions: a key class of computations within massive data analysis. Needed for Bayesian inference using any model and in non-Bayesian statistical settings, e.g. random effects models. Integrals that appear in statistics are often expectations with a special form.

7. **Alignment problems**
   - Matchings between two or more data objects or data sets, e.g., multiple sequence alignments in computational biology, matching of catalogs from different instruments in astronomy, matching of objects between images, correspondence between synonymous words in text analysis. Critical in data fusion – required before further data analyses is possible

**Frontiers In Massive Data Analysis, National Research Council (The National Academies Press, 2013)*

ECP Co-Design Centers for key computational motifs

Project	PI Name, Inst	Short Description/Objective
CODAR	Ian Foster, ANL	Understand the constraints, mappings, and configuration choices between applications, data analysis and reduction, and exascale platforms
AMReX	John Bell, LBNL	Build framework to support development of block-structured adaptive mesh refinement algorithms for solving systems of partial differential equations on exascale architectures
CEED	Tzanio Kolev, LLNL	Develop next-generation discretization software and algorithms that will enable finite element applications to run efficiently on future hardware
CoPA	Susan Mniszewski, LANL	Create co-designed numerical recipes and performance-portable libraries for particle-based methods
ExaGraph	Mahantesh Halappanavar, PNNL	Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms
ExaLearn	Frank Alexander, BNL	Deliver state-of-the-art machine learning and deep learning software at the intersection of applications, learning methods, and exascale platforms

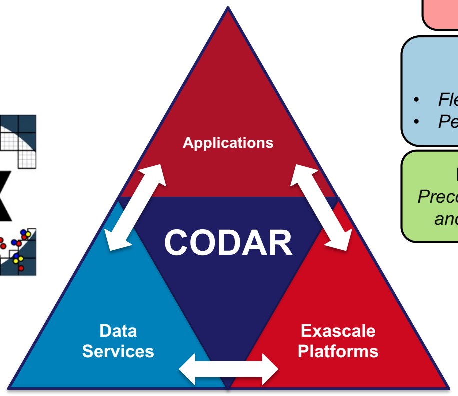

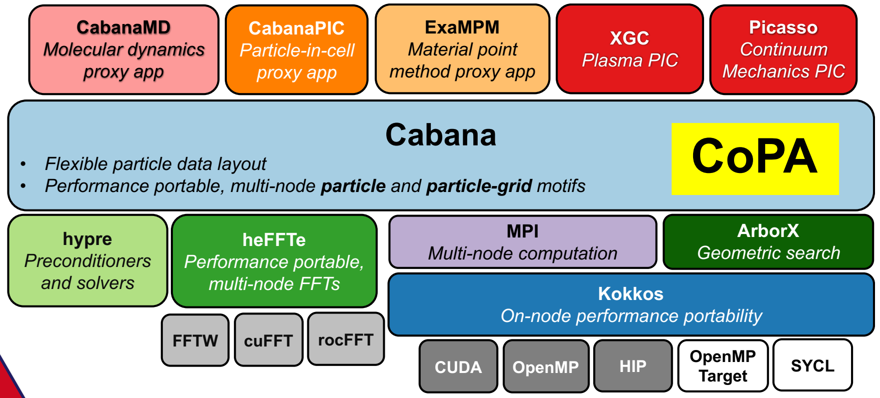

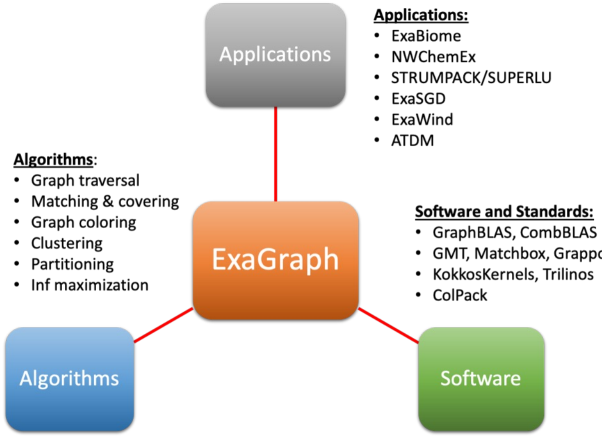

Exascale Modeling of Advanced Particle Accelerators

Toward compact and affordable particle accelerators. A laser beam or a charged particle beam propagating through ionized gas displaces electrons, creates a *wakefield* that supports electric fields orders of magnitude larger than with usual methods, accelerating a charged particle beam to high energy over a very short distance.

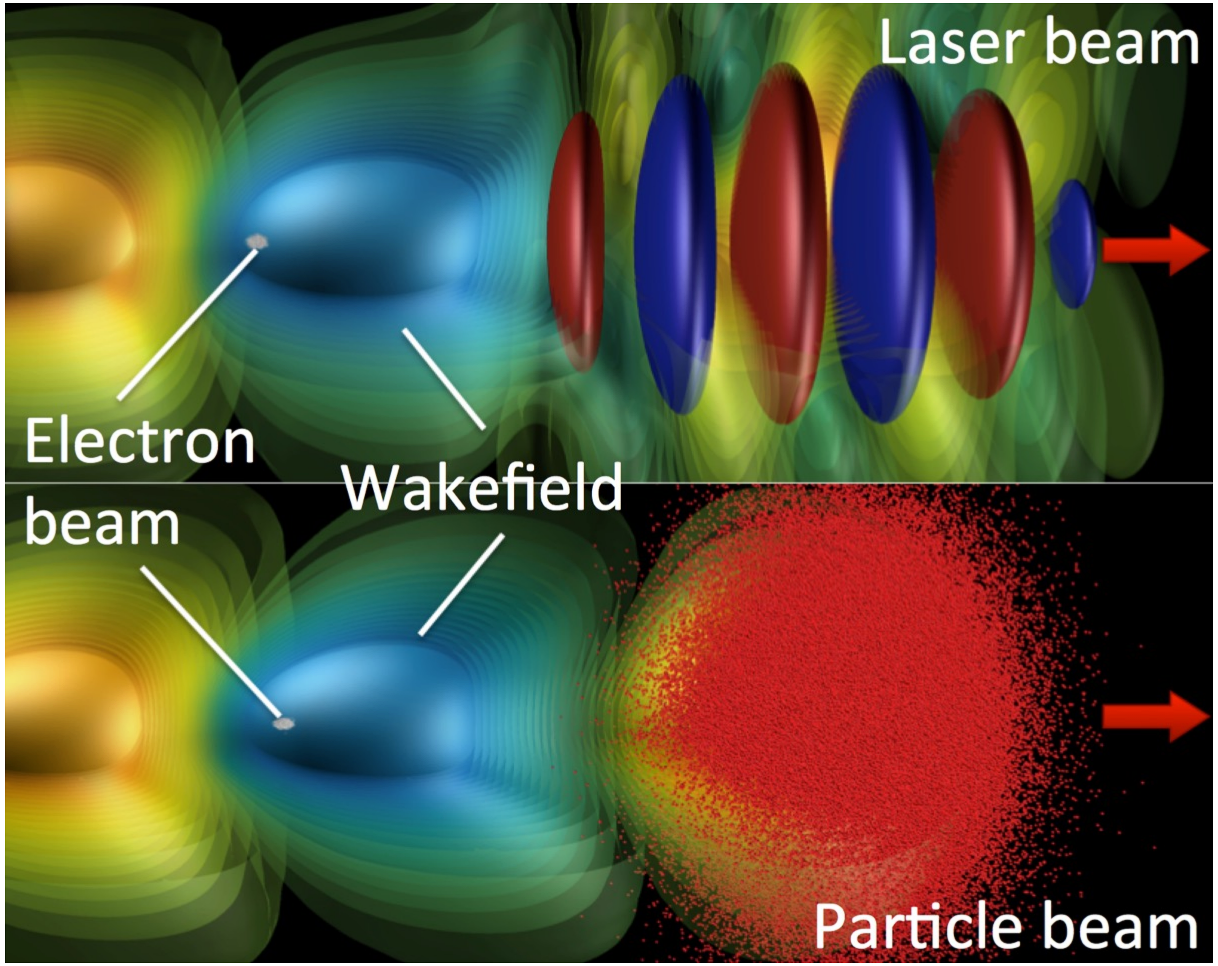

- Particle accelerators: a vital part of DOE infrastructure for discovery science and university- and private-sector applications - broad range of benefits to industry, security, energy, the environment, and medicine
- Improved accelerator designs are needed to drive down size and cost; plasma-based particle accelerators stand apart in their potential for these improvements
- Translating this promising technology into a mainstream scientific tool depends critically on exascale-class high-fidelity modeling of the complex processes that develop over a wide range of space and time scales
- Exascale-enabled acceleration design will realize the goal of compact and affordable high-energy physics colliders, with many spinoff plasma accelerator applications likely

PI: Jean-Luc Vay (LBNL)

ECP WarpX app is relevant to many FES topics
(developed for plasma acceleration of e⁻/e⁺ beams for HEP colliders)

New generation plasma code

WarpX

ECP EXASCALE COMPUTING PROJECT

https://github.com/ECP-WarpX/WarpX

Open source, developed by tightly integrated team of physicists + applied mathematicians + computer scientists

PI: J.-L. Vay; also co-lead of IFE BRN Workshop Theory & Simulations panel (June 2022)

international collaborators & industry

Runs on single user desktops/laptops up to largest CPU or GPU-based supercomputers

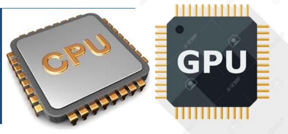

Applicable to the modeling of

- Laser-ion plasma acceleration.
- Laser-plasma interactions.
- Plasma mirrors.
- Collisionless shocks.
- Pulsars.
- Magnetic reconnections.
- Magnetic fusion sheaths.
- Intense particle beams.
- Accelerator designs.
- Particle sources.
- High-field physics (with QED)
- …

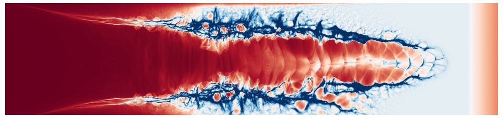

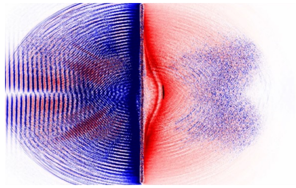

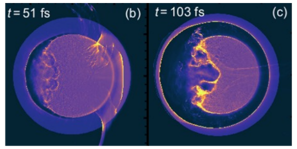

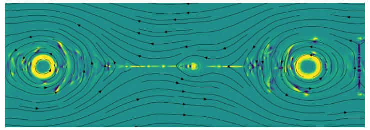

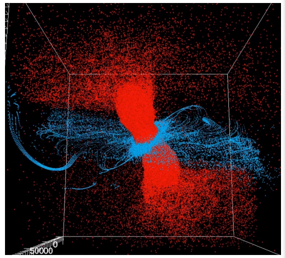

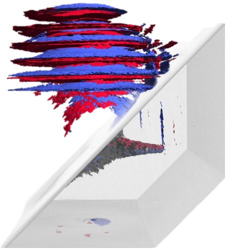

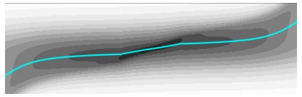

WarpX is also part of the inaugural pool of codes selected to support upcoming LaserNetUS experiments

LaserNetUS

PI: Jean-Luc Vay (LBNL)

Molecular Dynamics at the Exascale: Spanning the Accuracy, Length and Time Scales for Critical Problems in Materials Science (EXAALT)

Combining time-acceleration techniques, spatial decomposition strategies, and high accuracy quantum mechanical and empirical potentials

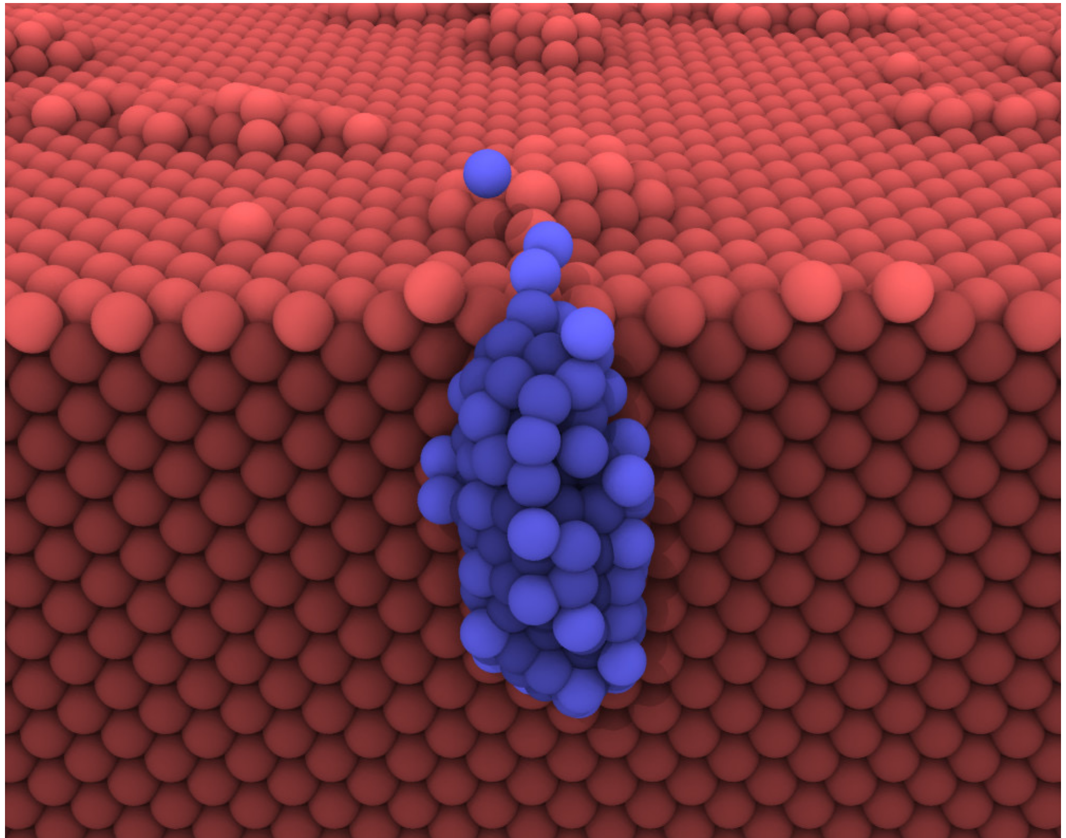

Slowly-growing He bubble in W at bursting

- Tackle materials challenges for energy, especially fission and fusion, by allowing the scientist to target, at the atomistic level, the desired region in accuracy, length, and time space
- Shown here is a simulation aimed at understanding tungsten as a fusion first-wall material, where plasma-implanted helium leads to He bubbles that grow and burst at the surface, ultimately leading to surface "fuzz" by a mechanism not yet understood
- At slower, more realistic growth rates (100 He/µsec), the bubble shows a different behavior, with less surface damage, than the fast-grown bubble simulated with direct molecular dynamics (MD)
- Atomistic simulation allows for complete microscopic understanding of the mechanisms underlying the behavior
- At the slower growth rate, crowdion interstitials emitted from the bubble have time to diffuse over the surface of the bubble, so that they are more likely to release from the surface-facing side of the bubble, giving surface-directed growth.

PI: Danny Perez (LANL)

ECP’s EXAALT Application: The Problem

- Molecular Dynamics (MD) is an ideal tool to investigate materials at the atomic scale
- MD is however computationally intensive
- Domain-decomposition weak-scales, but does not strong-scales
- Large-scale computers:
  - **Can** increase length-scales (trillions of atoms)
  - **Cannot** increase timescales (<microsecond)
- Damage/microstructure evolution is very difficult to study with MD

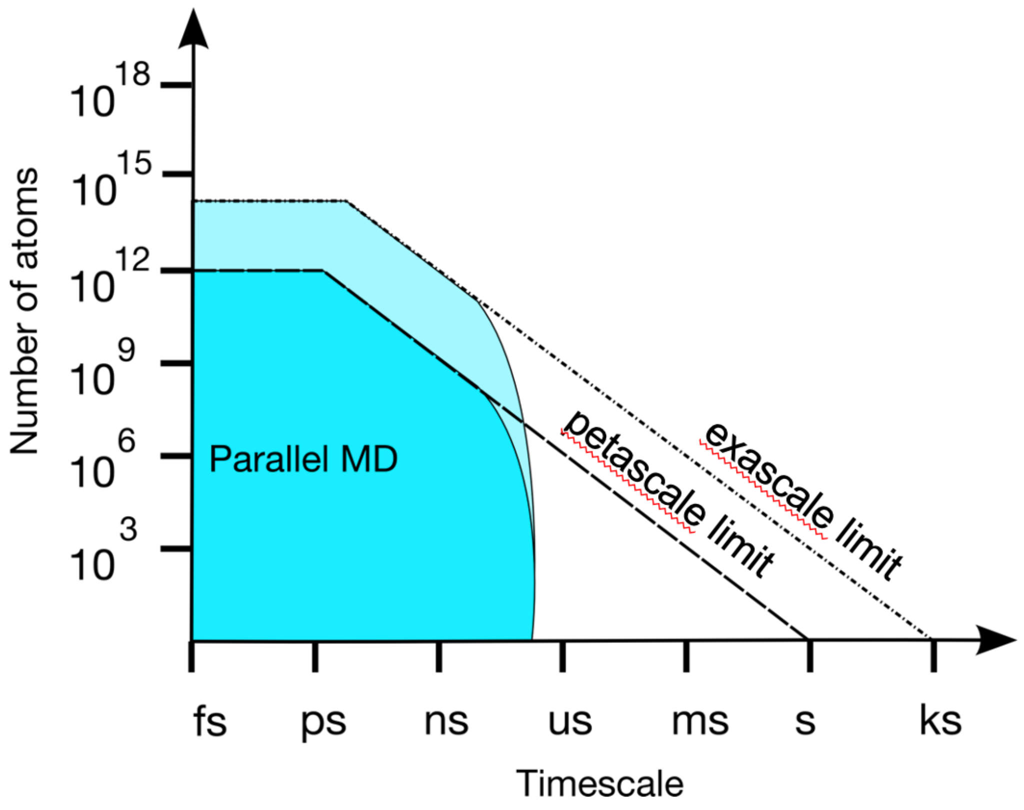

ECP's EXAALT Application: The Targets

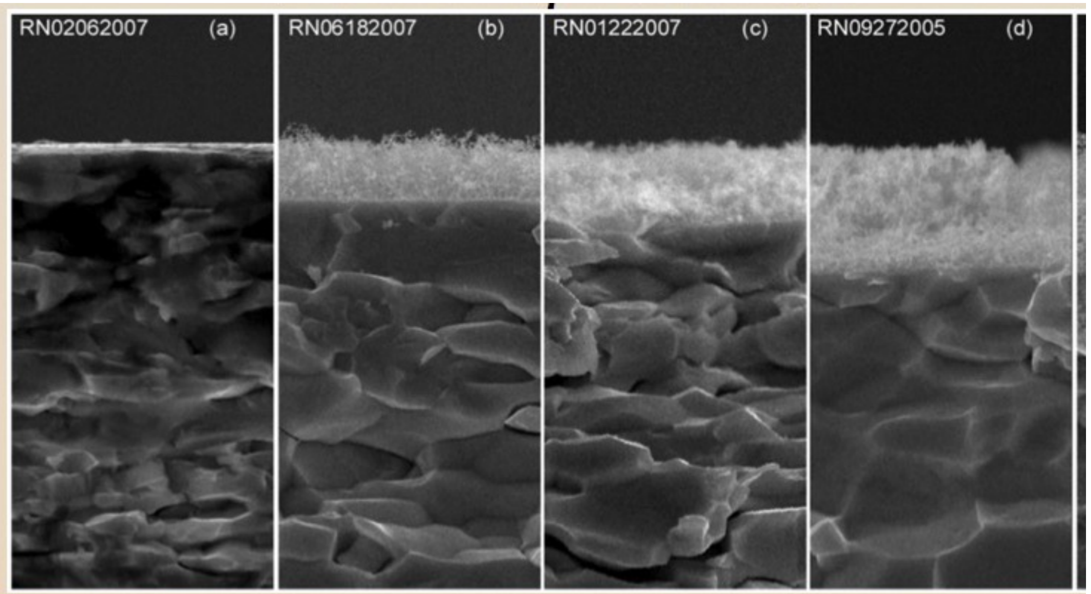

R. Doerner *et al.*

- Atomistic modeling of the microstructural evolution of material for energy applications
  - Surface damage evolution/fuzz formation
  - H/He interaction with defects
  - Effect of He/H on recrystallization

- Target system: evolution of W first wall in fusion conditions
- Target regime: 10^5 atoms, >10 ms/wall-clock day
- Not accessible using standard MD techniques

ECP's EXAALT Application: Methods

- Long times accessed with Accelerated MD methods (Voter *et al.*)
  - Parallel Trajectory Splicing (Perez *et al.*)
  - TAMMBER (Swinburne *et al.*)
- Parallelizes in the *time domain* using replica-based techniques
- Dynamically accurate to arbitrary precision (Lelievre *et al.*)

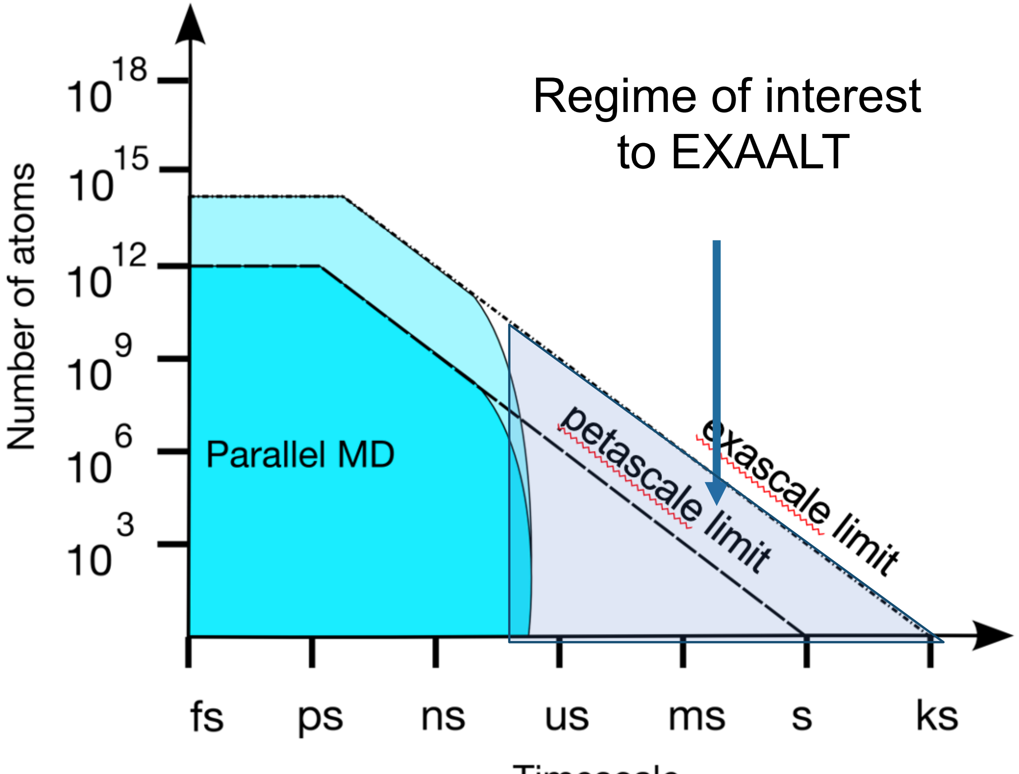

- Intermediate size/time regime through combined domain+replica decomposition (synchronous sub-lattice, Amar *et al.*)

ECP's EXAALT Application: Computational Capability

- AMD methods implemented through custom-made task and data management system
- Fully asynchronous execution: no blocking/all-to-all communications
- Can be used to implement a variety of complex workflows:
  - Kinetic model construction
  - Machine-learning potentials

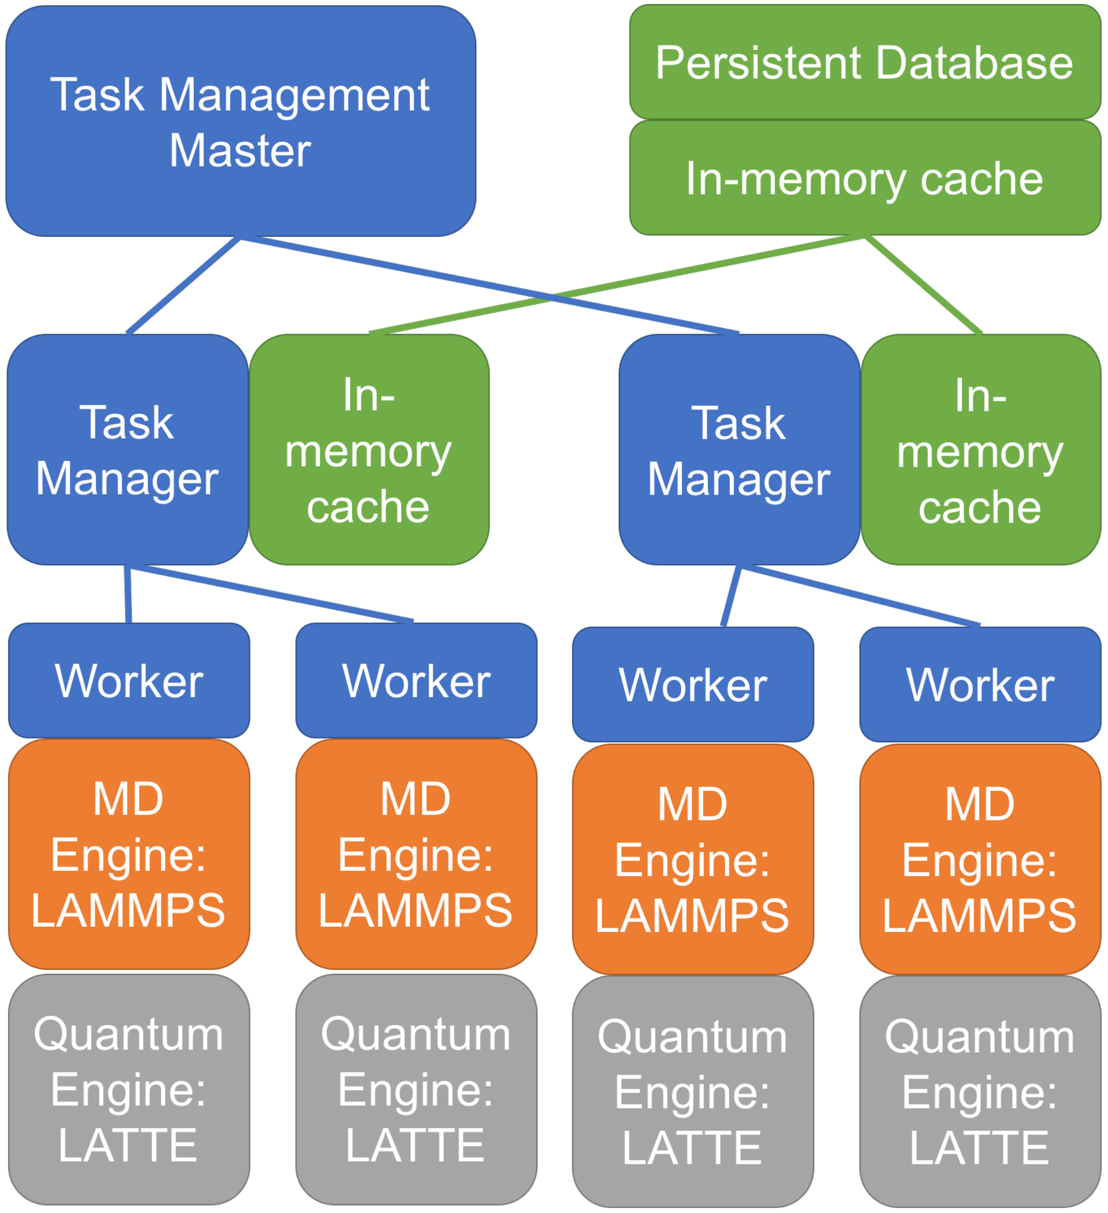

ECP's EXAALT Application: Performance optimization

- Next generation interatomic potentials rely on machine learning (SNAP, ACE)
- High computational complexity, deep nested loops, large implementation space
- Rewrote and optimized code using proxy app: **25x increase in GPU performance**
- Extremely high performance: Gordon-Bell 2021 Finalist
- Transferable to many other applications

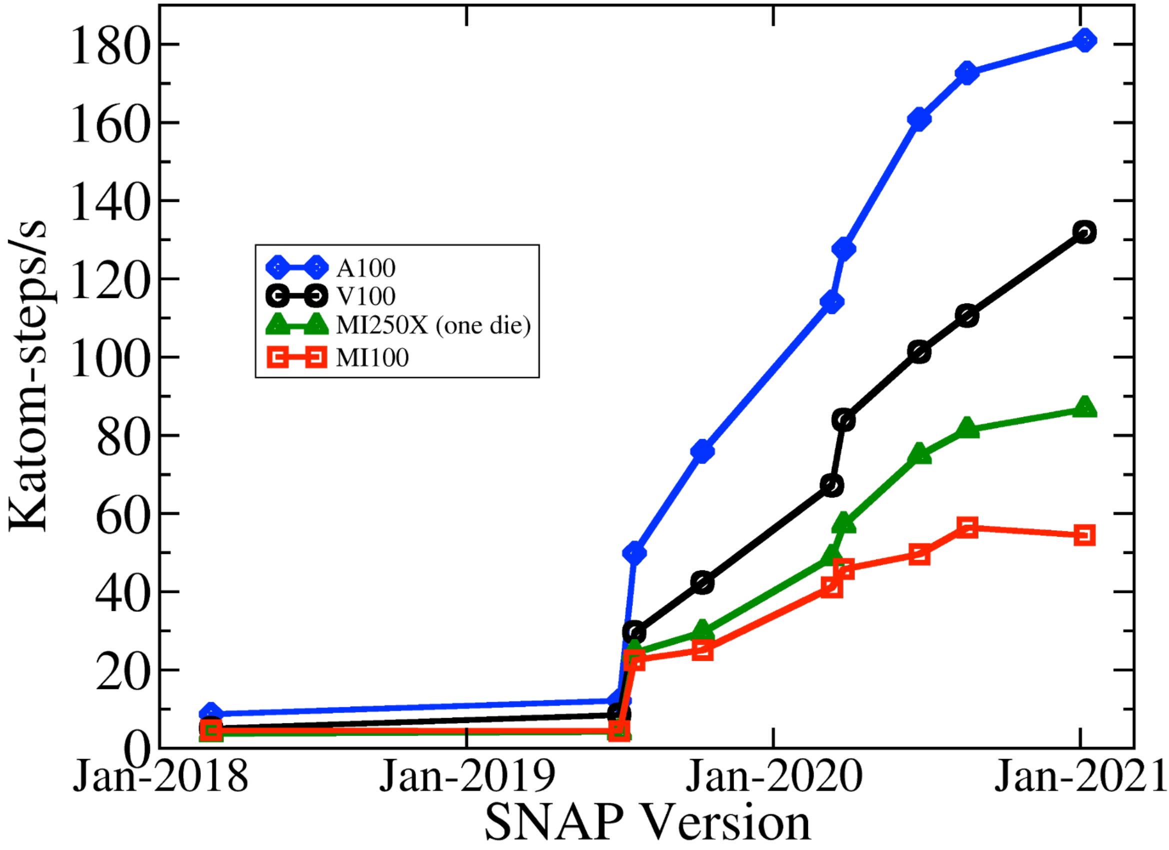

Rahul Gayatri (NERSC), Evan Weinberg (NVIDIA), Stan Moore, Aidan Thompson (SNL)
Nicholas Lubbers (LANL)

ECP's EXAALT Application: Machine Learning

- Obtaining good interatomic potentials is a critical, but time-consuming
- Machine-learning potentials can dramatically fail if not constrained by very diverse data
- Developed automated data generation procedure based on information entropy optimization
- Whole workflow being integrated with EXAALT
- In collaboration with FES FusMatML project (Aidan Thompson, SNL)

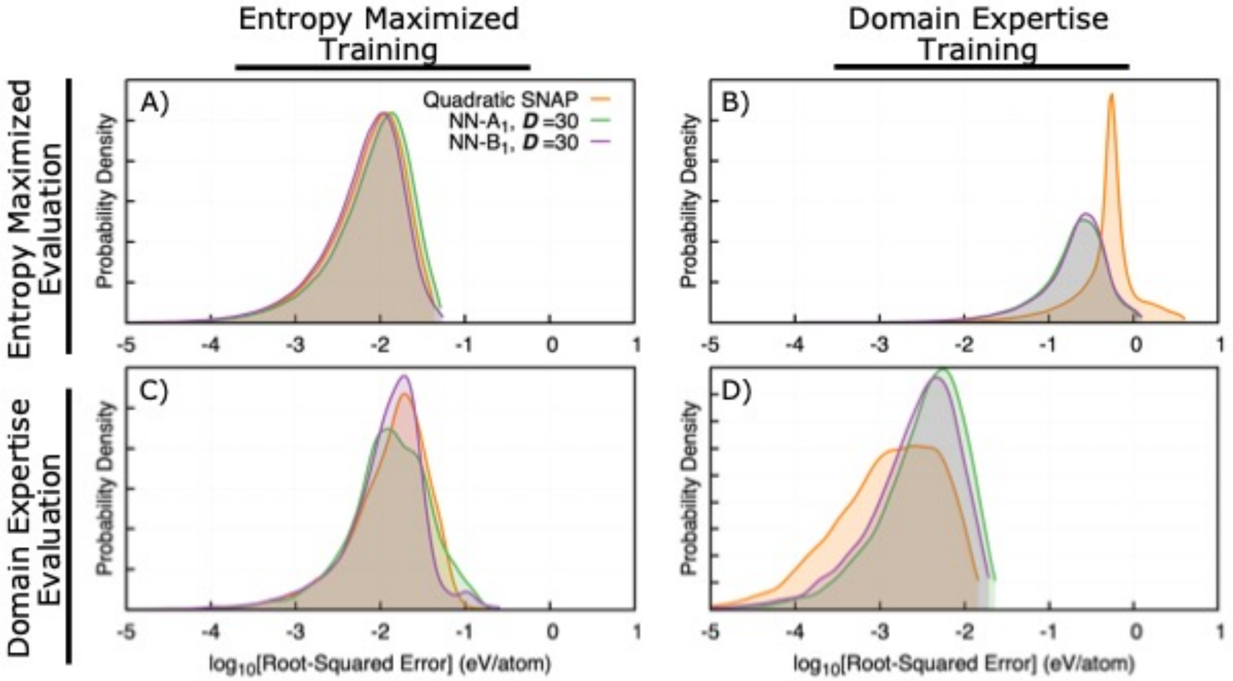

Entropy dataset:
Slightly higher local error
extreme transferability

Conventional dataset:
low local error, very
high extrapolation error

ECP's WDMApp Application: Goals

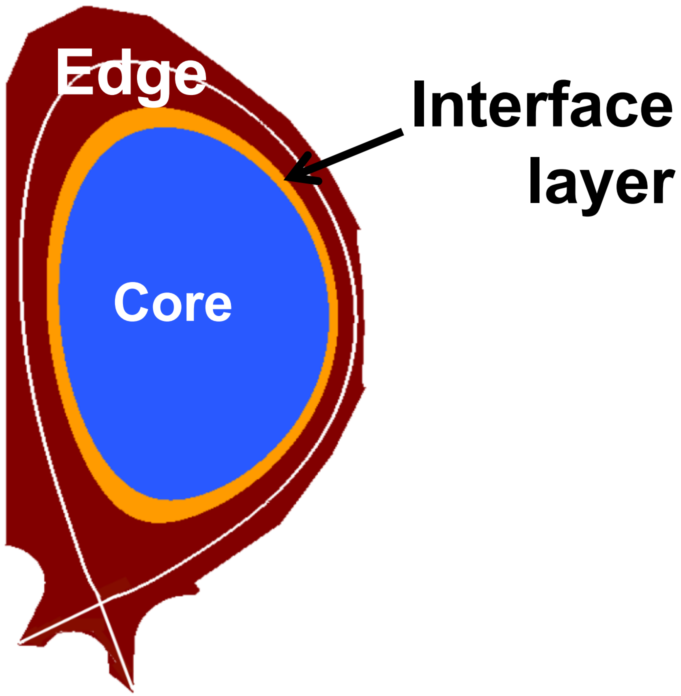

- Demonstration and assessment of WDM **gyrokinetic physics** on experimental transport time-scale in a challenge problem for pedestal formation
- **Figure of Merit (FOM) of >50** for coupled code on exascale platforms, accomplished through algorithmic advancement, performance engineering and hardware improvement
- Completion of **extensible integration framework** EFFIS 2.0 (End-to-End Framework for Fusion Integrated Simulations 2.0) and demonstration on exascale platform

ECP's WDMApp Application: Challenge Problem

High-fidelity simulation of a whole-device burning plasma (specifically, ITER with full plasma current) operating in "high-mode" (H-mode), and prediction of the plasma pressure "pedestal" shape (height and width)

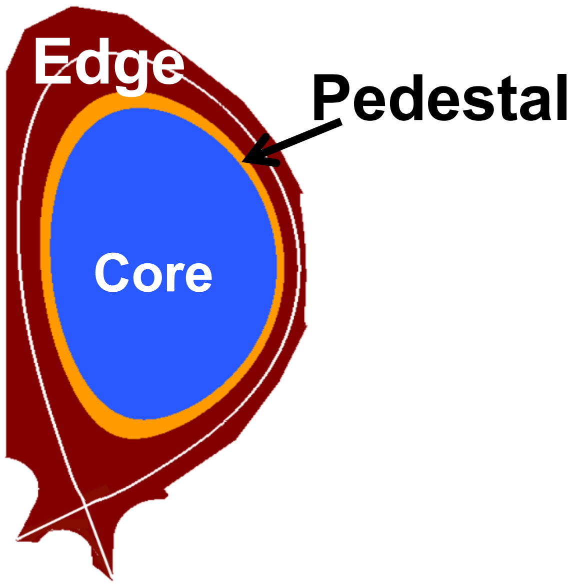

Pedestal determines the plasma pressure, hence fusion yield, in the burning core

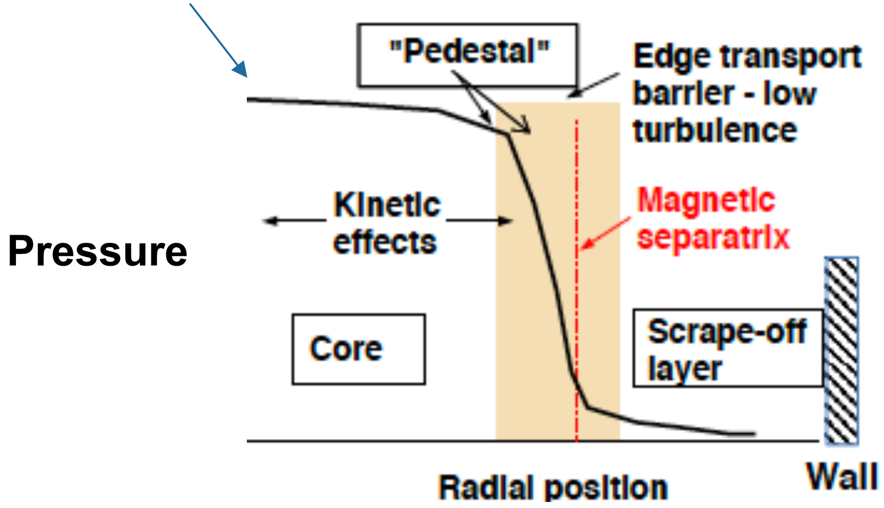

PI: Amitava Bhattacharjee (PPPL)

ECP's WDMApp Application: Coupling the core and edge

The core evolves more slowly than the edge

Core Turbulence from GENE

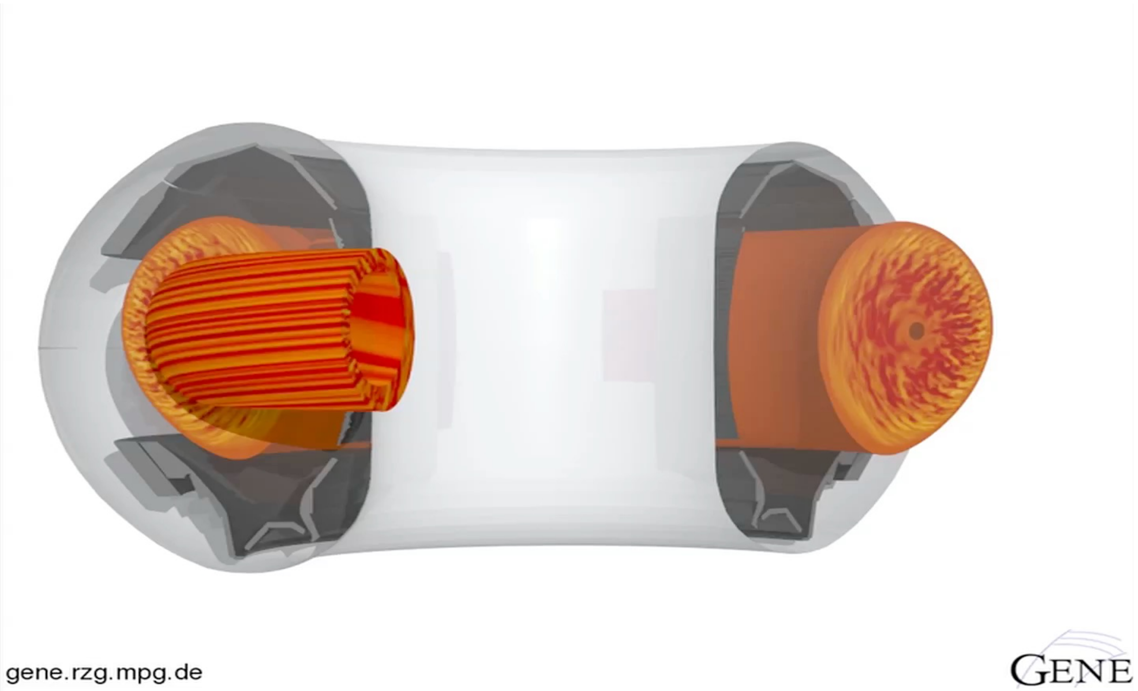

Edge Turbulence from XGC

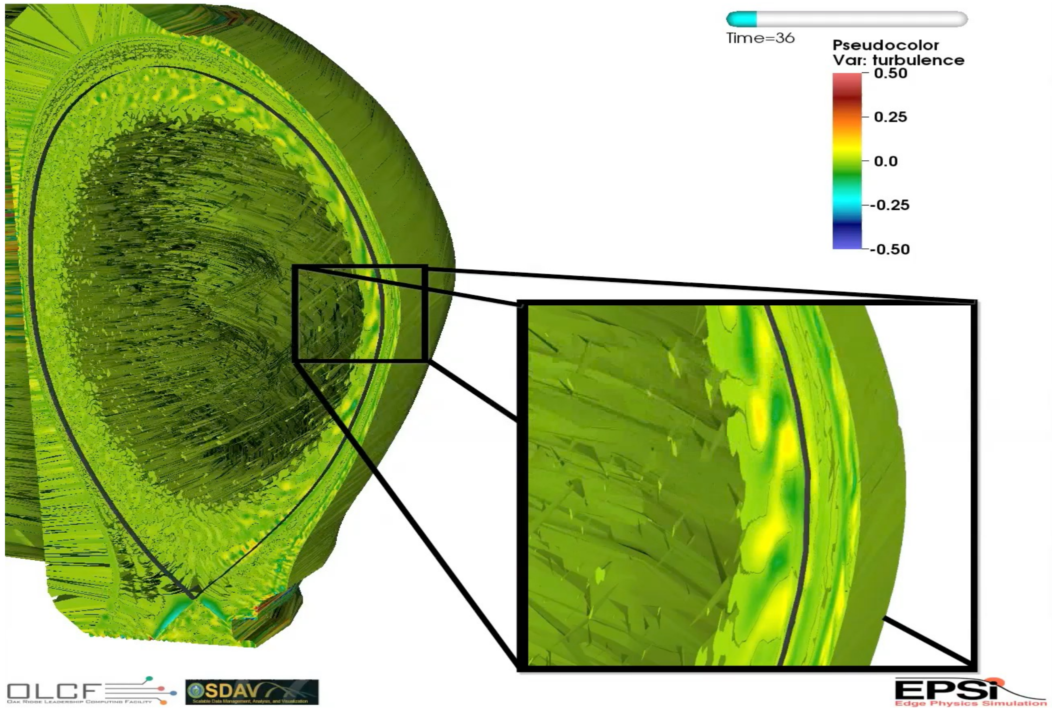

PI: Amitava Bhattacharjee (PPPL)

EFFIS 2.0 Workflow Framework

- EFFIS couples codes, analysis and visualization services, and automates post processing which enables scientists to
  - Helps measure the FOM for the WDMApp
  - Enables scientist to compose, execute and control the WDMApp, and allows for automated post and online processing
  - Visualize their data on a collaborative dashboard
- How
  - EFFIS is a workflow manager consisting of a library, utilities and runtime daemons
  - The goal is to enable end-to-end workflows, with easy integration of services
  - Application-to-application coupling, including through an external "coupler"
  - In situ and automated post-process for analysis and visualization
- EFFIS allows
  - **Low overhead** data movement via ADIOS (<<1% overhead): the figure on the right
  - Multiple codes/services to be executed on a single shared node or multiple nodes
  - **Platform portability** for composition, submission and execution
- ECP work that supports EFFIS workflows (ECP-ST-AD collaborations)
  - CODAR (Savanna, MGARD, FTK), VTK-M (VTK-M), ADIOS-ECP (ADIOS 2.0)

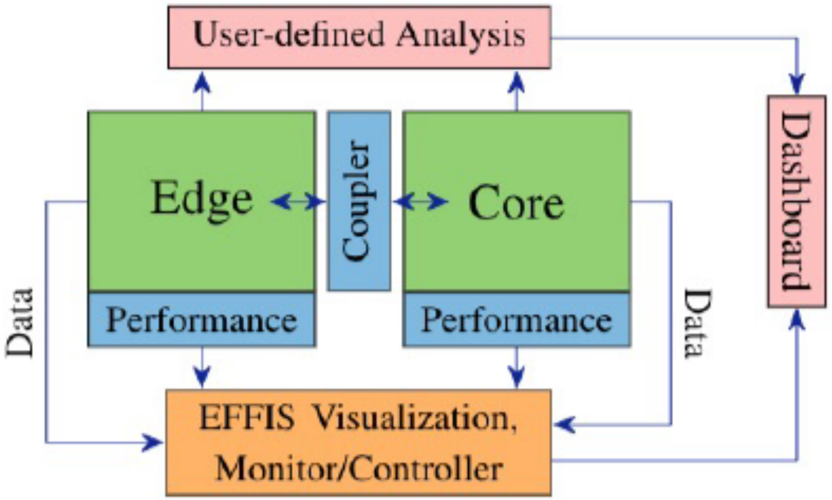

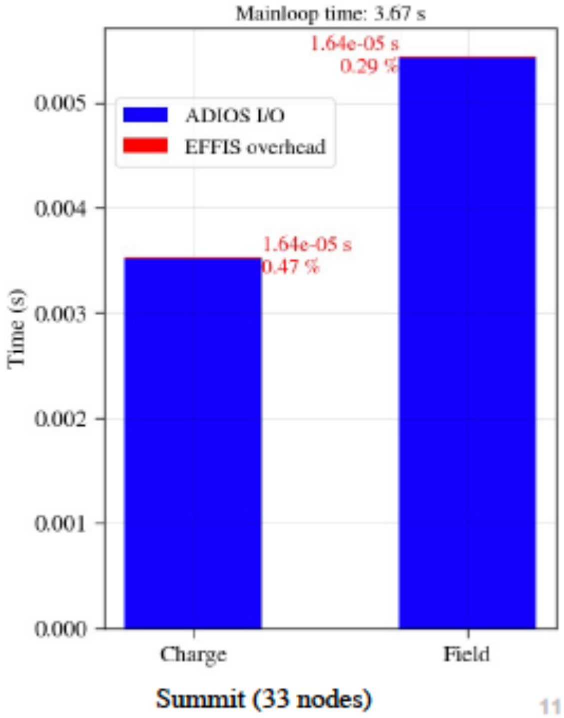

ECP’s ExaLearn Co-Design Center: Application Pillars

Surrogates

- ML-created models
- Faster and/or higher fidelity models
- Generative networks
- Using ML to replace complicated physics
- Cosmology

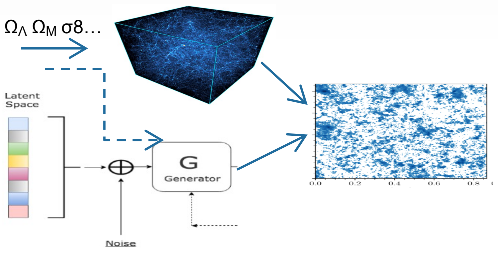

Control

- ML-controlled experiments
- Efficient exploration of complex space
- Reinforcement Learning
- Use RL agent to control light source experiments
- Temperature control for Block Co-Polymer (BCP) experiments

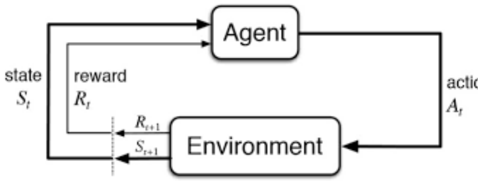

Image courtesy Sutton, Barto, Reinforcement Learning 2017

Design

- ML-created physical structures
- Optimized proposal for desired behavior of structure within complex design space
- Graph-Convnets
- Use Graph-CNN to propose new structures that respect chemistry
- Molecular Design

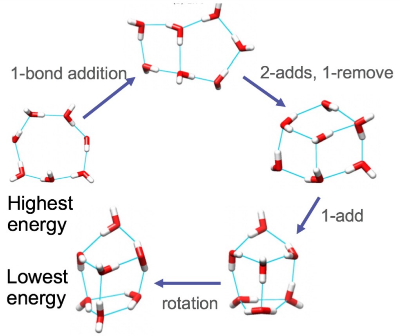

Inverse

- ML projection from observation to original form
- Back-out complex input structure from observed data
- Regression models
- Predicting crystal structure from light source imaging
- Material structure from neutron scattering

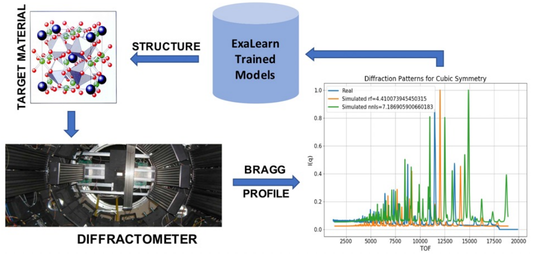

EXASCALE COMPUTING PROJECT

PI: Frank Alexander (BNL)

Applying ExaLearn’s Pillars to Tokamak Fusion: Some Examples

Surrogates

- ML-created models of conventional **HPC particle-in-cell codes** for plasma physics fusion dynamics (“GTC”)
- Synthetic “digital twins” enabling approximate real-time (RT) representation of large-scale HPC fusion simulations (“SGTC”)
- Synthetic plasma equilibrium state model (“SEFIT”)

Control

- ML-controlled fusion energy experiments with deployment in RT plasma control system (PCS) for DIII-D tokamak
- Reinforcement Learning (RL)
  - To alter temporal evolution in RT tokamak experiments
  - To enable desirable interventions, e.g., **disruption mitigation/ avoidance** in RT tokamak experiments

Design

- ML-generated improved future magnetic confinement fusion systems
- Optimization of proposed ideas for RT plasma behavior within “computer design space” for advanced tokamaks, including:
  - High-magnetic field
  - Superconducting/“steady-state”
  - Reduced-size “modular” systems
- **Exploration of design tools** for fusion applications, including:
  - Graph neural networks
  - Variational auto-encoders
  - Point-cloud-based convolutional neural nets (CNNs)

Inverse

- ML projection from experimental observations of **plasma states to possible earlier states** in evolution
- Complex input structure deduced/“backed out” from observed plasma state data
  - “Regression” models
- Predicting dynamical plasma structures from high-resolution diagnostic data, including:
  - Temperature profiles from Thompson-scattering and ECE emission measurements
  - Current profiles from Motional Stark Effect measurements
  - 2D micro-turbulence structures from Electron-Cyclotron-Emission-Imaging measurements

ECP Applications: potential outcomes and impact

Will be far-reaching for decades to come

- Predictive microstructural evolution of novel chemicals and materials for energy applications.
- Robust and selective design of catalysts an order of magnitude more efficient at temperatures hundreds of degrees lower.
- Accelerate the widespread adoption of additive manufacturing by enabling the routine fabrication of qualifiable metal alloy parts.
- Design next-generation quantum materials from first principles with predictive accuracy.
- Predict properties of light nuclei with less than 1% uncertainty from first principles.
- Harden wind plant design and layout against energy loss susceptibility, allowing higher penetration of wind energy.
- Demonstrate commercial-scale transformation energy technologies that curb fossil fuel plant CO_2 emission by 2030.
- Accelerate the design and commercialization of small and micronuclear reactors.
- Provide the foundational underpinnings for a ‘whole device’ modelling capability for magnetically confined fusion plasmas useful in the design and operation of ITER and future fusion reactors.

ECP Applications: potential outcomes and impact

Will be far-reaching for decades to come

- Address fundamental science questions such as the origin of elements in the universe, the behaviour of matter at extreme densities, the source of gravity waves; and demystify key unknowns in the dynamics of the universe (dark matter, dark energy and inflation).

- Reduce the current major uncertainties in earthquake hazard and risk assessments to ensure the safest and most cost-effective seismic designs.

- Reliably guide safe long-term consequential decisions about carbon storage and sequestration.

- Forecast, with confidence, water resource availability, food supply changes and severe weather probabilities in our complex earth system environment.

- Optimize power grid planning and secure operation with very high reliability within narrow operating voltage and frequency ranges.

- Develop treatment strategies and pre-clinical cancer drug response models and mechanisms for RAS/RAF-driven cancers.

- Discover, through metagenomics analysis, knowledge useful for environment remediation and the manufacture of novel chemicals and medicines.

- Dramatically cut the cost and size of advanced particle accelerators for various applications impacting our lives, from sterilizing food of toxic waste, implanting ions in semiconductors, developing new drugs or treating cancer.

ECP is delivering an open, hierarchical software ecosystem

More than a collection of individual products

Levels of Integration | **Product** | **Source and Delivery**

ECP ST Open Product Integration Architecture

- Build all SDKs
- Build complete stack
- Assure core policies
- Build, integrate, test

Source: ECP E4S team; Non-ECP Products (all dependencies)
Delivery: spack install e4s; containers; CI Testing, more

- Group similar products
- Make interoperable
- Assure policy compliant
- Include external products

Source: SDK teams; Non-ECP teams (policy compliant, spackified)
Delivery: Apps directly; spack install sdk; future: vendor/facility

- Standard workflow
- Existed before ECP

Source: ECP L4 teams; Non-ECP Developers; Standards Groups
Delivery: Apps directly; spack; vendor stack; facility stack

ECP ST Individual Products

ECP’s Software Development Kits (SDKs) Span All Technology Areas

ECP’s Extreme Scale Scientific Software Stack (E4S) embodies the latest Software Technology products developed in ECP and packaged in SDKs. The latest Feb 2022 release (https://e4s.io) includes 100 distinct products using the Spack package manager in a full-feature containerized release. E4S also supports AI/ML packages such as TensorFlor, PyTorch, and Horovod. E4S is available for download from Dockerhub, with bare metal and custom containers also supported using the E4S Spack build cache.

PMR Core (17)	Compilers and Support (7)	Tools and Technology (11)	xSDK (16)	Visualization Analysis and Reduction (9)	Data mgmt, I/O Services, Checkpoint restart (12)	Ecosystem/E4S at-large (12)
QUO	openarc	TAU	hypre	ParaView	SCR	mpiFileUtils
Papyrus	Kitsune	HPCToolkit	FleSCI	Catalyst	FAODEL	TriBITS
SICM	LLVM	Dyninst Binary Tools	MFEM	VTK-m	ROMIO	MarFS
Legion	CHiLL autotuning comp	Gotcha	Kokkoskernels	SZ	Mercury (Mochi suite)	GUFI
Kokkos (support)	LLVM openMP comp	Caliper	Trilinos	zfp	HDF5	Intel GEOPM
RAJA	OpenMP V & V	PAPI	SUNDIALS	VisIt	Parallel netCDF	BEE
CHAI	Flang/LLVM Fortran comp	Program Database Toolkit	PETSc/TAO	ASCENT	ADIOS	FSEFI
PaRSEC*		Search (random forests)	libEnsemble	Cinema	Darshan	Kitten Lightweight Kernel
DARMA		Siboka	STRUMPACK	ROVER	UnifyCR	COOLR
GASNet-EX		C2C	SuperLU		VeloC	NRM
Qthreads		Sonar	ForTrilinos		IOSS	ArgoContainers
BOLT			SLATE		HXHIM	Spack
UPC++			MAGMA			
MPICH			DTK			
Open MPI			Tasmanian			
Umpire			TuckerMPI			
AML						

Legend

- PMR
- Tools
- Math Libraries
- Data and Vis
- Ecosystems and delivery

Extreme-scale Scientific Software Stack (E4S)

- E4S: HPC Software Ecosystem – a curated software portfolio
- A **Spack-based** distribution of software tested for interoperability and portability to multiple architectures with support for GPUs from NVIDIA, AMD, and Intel in a single distribution
- Available from **source**, **containers**, **cloud, binary caches**
- Leverages and enhances SDK interoperability thrust
- Not a commercial product – an open resource for all
- Oct 2018: E4S 0.1 - 24 full, 24 partial release products
- Jan 2019: E4S 0.2 - 37 full, 10 partial release products
- Nov 2019: E4S 1.0 - 50 full, 5 partial release products
- Feb 2020: E4S 1.1 - 61 full release products
- Nov 2020: E4S 1.2 (aka, 20.10) - 67 full release products
- Feb 2021: E4S 21.02 - 67 full release, 4 partial release
- May 2021: E4S 21.05 - 76 full release products
- Aug 2021: E4S 21.08 - 88 full release products
- Nov 2021: E4S 21.11 - 91 full release products
- Feb 2022: E4S 22.02 – 100 full release products

Spack

https://spack.io

Spack lead: Todd Gamblin (LLNL)

E4S

https://e4s.io

Lead: Sameer Shende (U Oregon)

Also include other products .e.g.,
AI: PyTorch, TensorFlow (CUDA, ROCm)
Co-Design: AMReX, Cabana, MFEM

Growing functionality: May 2022 release (version 22.05) will have 100+ products

ECP's goal is a sustainable, reusable software ecosystem

- ECP is driving the creation of a portfolio approach for reusable scientific software:
  - Available to you from laptops to supercomputers
  - Portable across CPU and GPU architectures
  - Available as open source for you to use, contribute to, and collaborate with
- Creating a future software organization that is a first-class citizen in the leadership computing ecosystem
- Encourage the fusion community to consider
  - **Using E4S**: https://e4s-project.github.io/download.html
  - **Contributing to E4S**: https://e4s-project.github.io/join.html
  - **Contributing to one of the SDKs**, e.g.: https://xsdk.info

Dependency database helps to track internal & external dependencies

Types of consumers of ST products

- **ECP applications (AD)**
  - Chem/materials, energy, earth/space, data/opt, NNSA, co-design
- **Other ST products**
  - PMR, dev tools, math libs, data/viz, sw ecosystem, NNSA

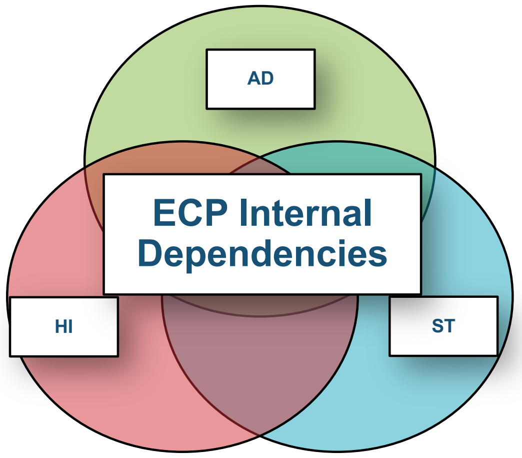

- **Facilities software stacks**
  - Aurora, Frontier, Perlmutter, El Capitan
- **Vendor software stacks**
  - NVIDIA, AMD, Cray, Intel
- **Community standards**
  - BLAS/LAPACK, C/C++, Fortran, LLVM, MPI, OpenACC, OpenMP, PowerAPI

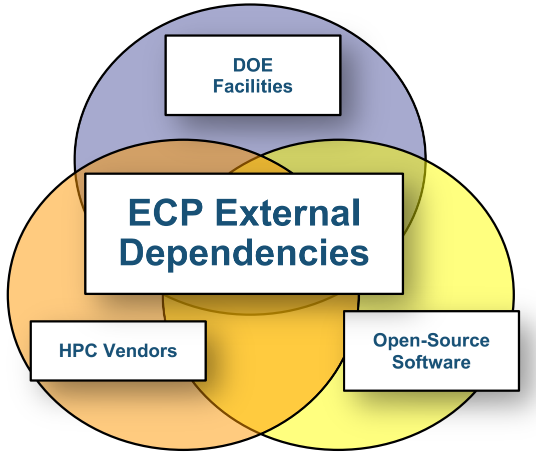

ECP has built a ST and AD code product dependency database

Proven indispensable for incentivizing, managing, and tracking critical dependencies

Product	URL	Description/Notes	Deployment Scope	Technical Area & WBS #	Point of Contact	KPP-3 Integration Categories
AML	https://xgitlab.cels.anl.gov/argo/aml	Hierarchical memory management library from Argo.	Experimental	PMR 2.3.1.19 (Argo)	@Pete Beckman	KPP-3
4. ALPINE (Ascent, ParaView, Catalyst, VisIt, LibSim, In Situ Algorithms)	https://github.com/Alpine-DAV/ascent https://www.paraview.org https://www.paraview.org/in-situ https://wci.llnl.gov/simulation /comput	Ascent: "Flyweight" *in situ* visualization and analysis runtime for multi-physics HPC	Moderate to Broad	Data & Viz 2.3.4.16	@James Ahrens @Terry Turton	AD
5. ArborX	https://g					
6. BEE	https://v https://g					

ST Product List

70 products with descriptions, **WBS**, **technical area**, points of contact, deployment scope, **integration categories**.

Use widely-recognized product names. Enables mapping between AD & Facilities dependencies and ST development efforts.

- MPI – MPICH, OpenMP
- C++/C/Fortran - LLVM
- Fortran – Flang
- hypre – hypre

AD Application Code Summary

Created by Erik Draeger, last modified by Tzanio Kolev on 2020-05-01

Below is a brief summary of the codes used by AD Application projects, the primary languages and what we believe to be their strategy for utilizing the GPUs. Please feel free to correct any errors or add additional information.

WBS	Application project	Code	Main language	GPU programming model	Notes	Point of Contact
2.2.1.01	LatticeQCD	Chroma	C++	Kokkos	Chroma is built on QDP++, which will be ported to Kokkos	Balint Joo
2.2.1.01	LatticeQCD	CPS	C++	GRID library	GRID currently uses CUDA, will port to AMD, Intel. May also try to port QUDA.	@Norman H. Christ
2.2.1.01						
2.2.1.02						

AD Code List

75 application codes used by project teams

Includes **WBS**, **application project**, languages used, GPU strategy, notes on integration, **point of contact**

ECP apps (AD) are primary consumers of ST products

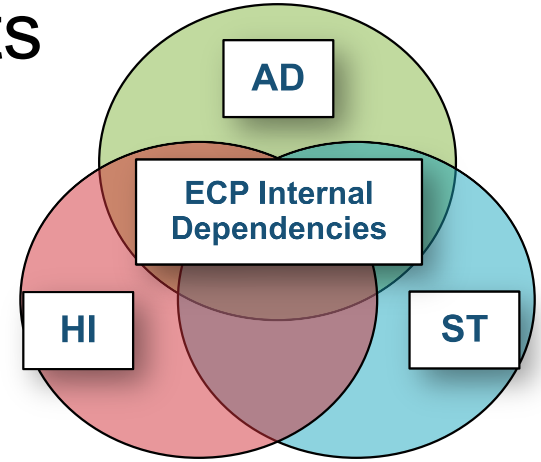

Dependency Database

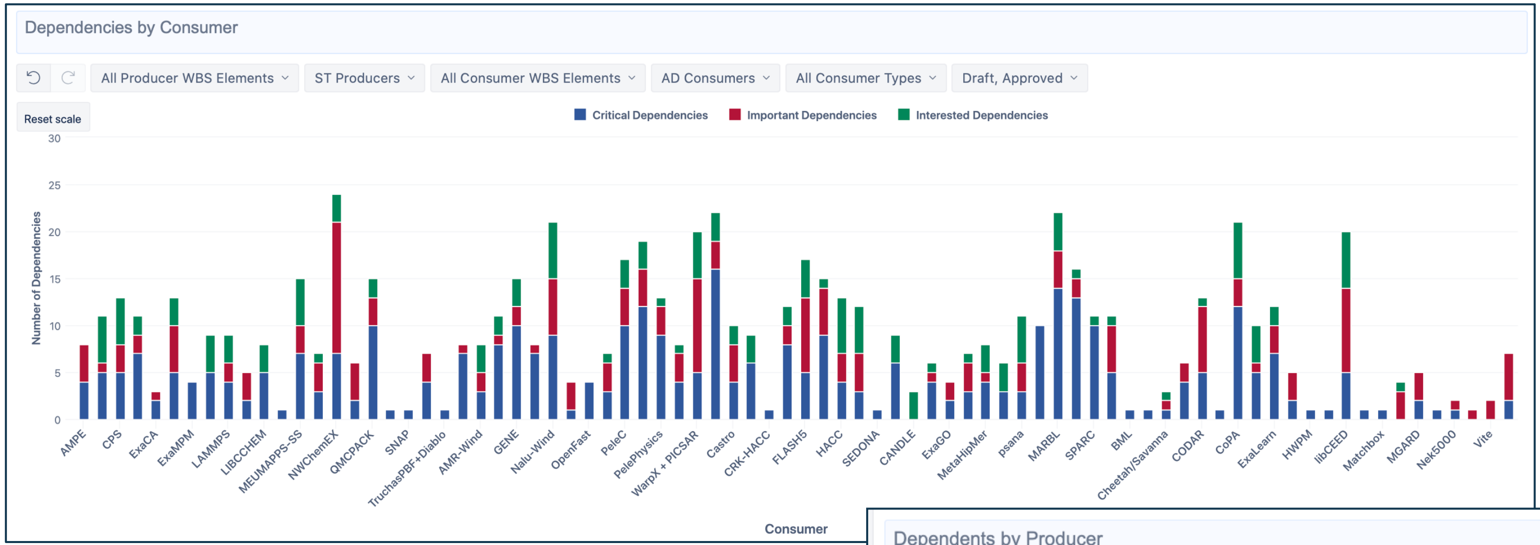

View by AD consumers

View by ST producers

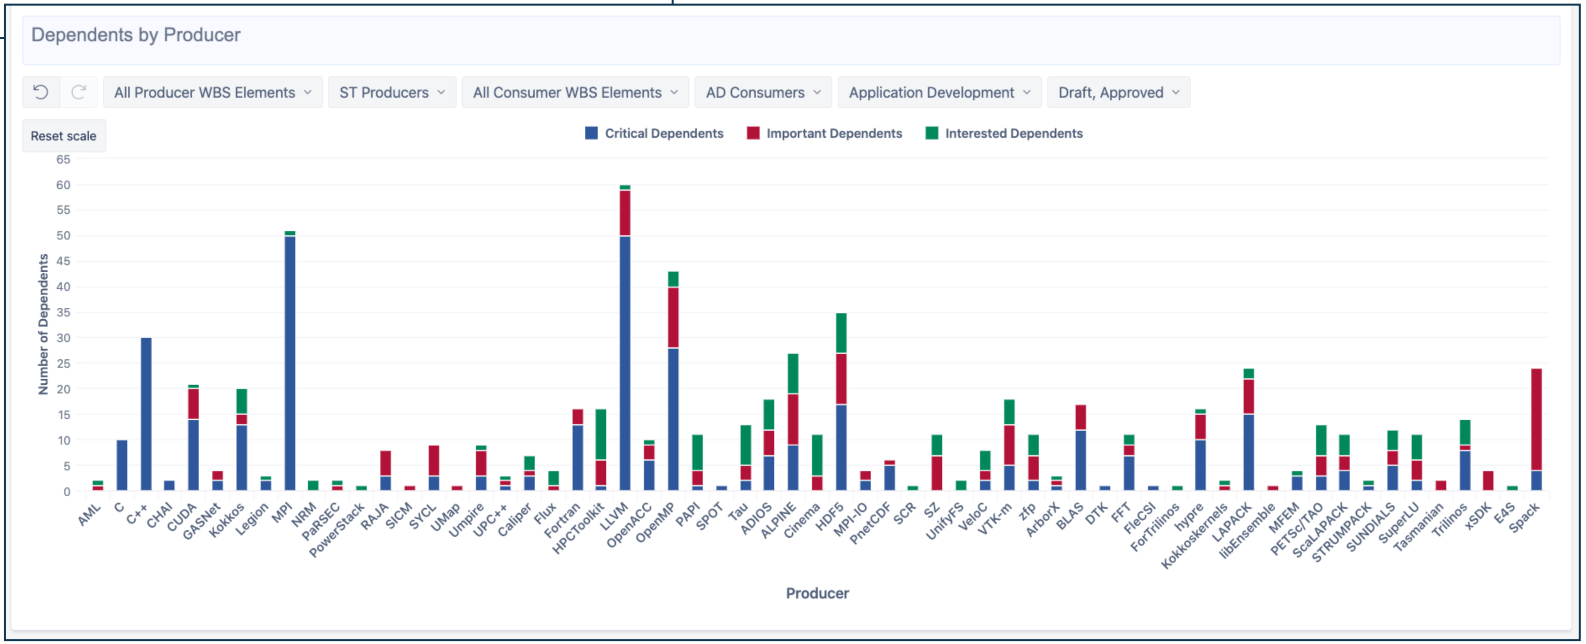

EXASCALE COMPUTING PROJECT

Closeup: WDMApp relies on several ST products

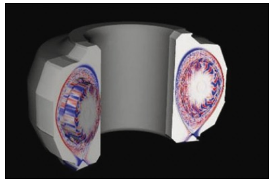

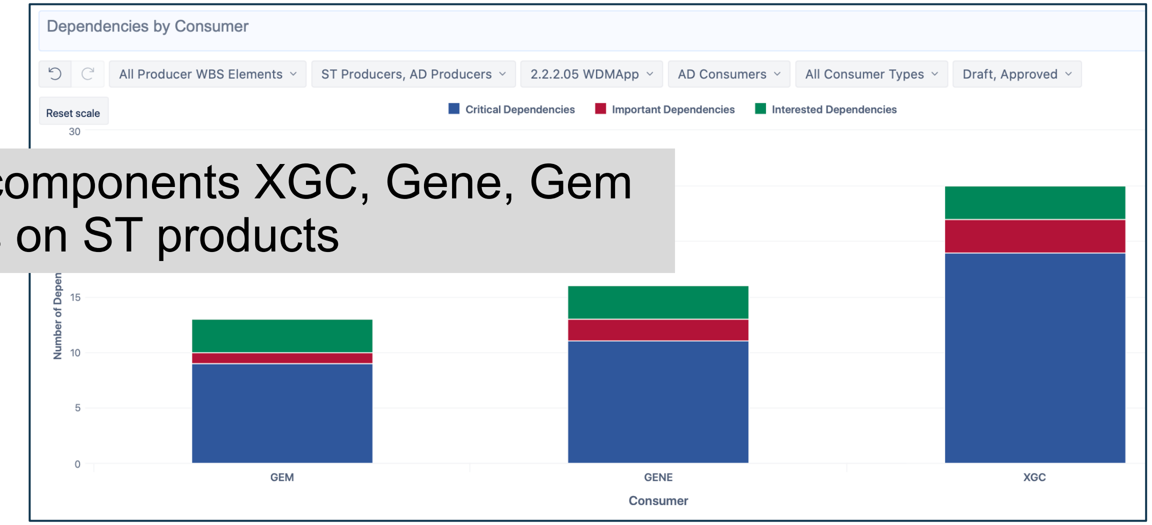

Multiphysics components XGC, Gene, Gem
- Each relies on ST products

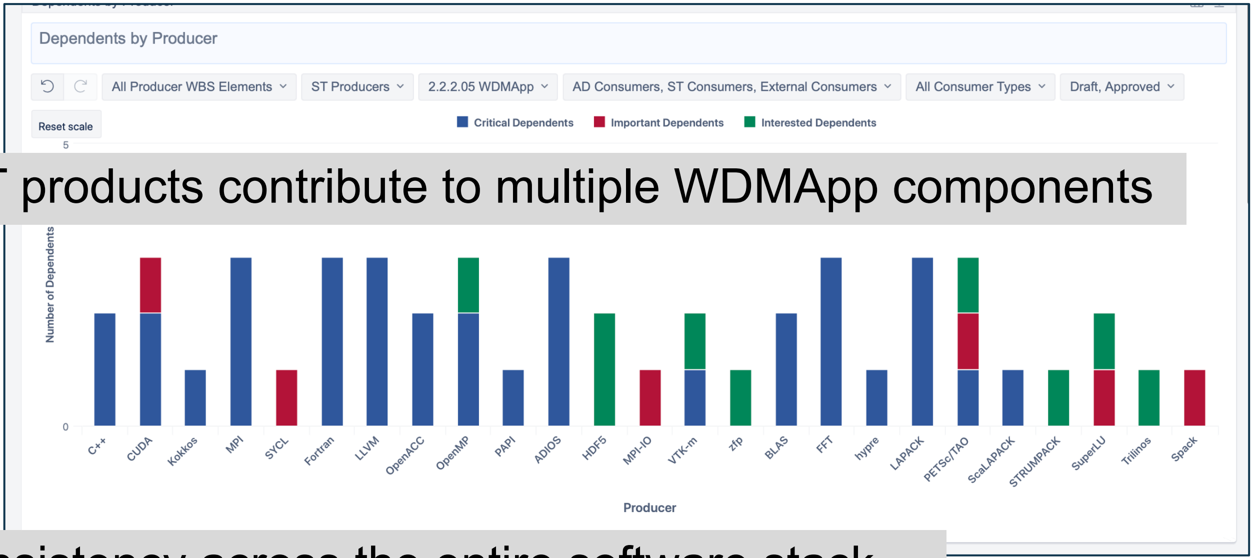

Some ST products contribute to multiple WDMApp components

Need consistency across the entire software stack

Details for XGC

			Critical Dependencies
Consumer	XGC	INT-155 LAPACK ↔ XGC	1
		INT-156 OpenMP ↔ XGC	1
		INT-154 hypre ↔ XGC	1
		INT-405 ADIOS ↔ XGC	1
		INT-609 BLAS ↔ XGC	1
		INT-635 PETSc/TAO ↔ XGC	1
		INT-1063 Kokkos ↔ XGC	1
		INT-1064 MPI ↔ XGC	1
		INT-1068 CUDA ↔ XGC	1
		INT-1069 FFT ↔ XGC	1
		INT-1070 VTK-m ↔ XGC	1
		INT-1071 PAPI ↔ XGC	1
		INT-1075 C++ ↔ XGC	1
		INT-1077 Fortran ↔ XGC	1
		INT-1177 OpenACC ↔ XGC	1
		INT-1212 LLVM ↔ XGC	1

			Important Dependencies
Consumer	XGC	INT-737 Spack ↔ XGC	1
		INT-1072 MPI-IO ↔ XGC	1
		INT-1074 SYCL ↔ XGC	1

			Interested Dependencies
Consumer	XGC	INT-404 STRUMPACK ↔ XGC	1
		INT-1073 Trilinos ↔ XGC	1
		INT-1076 SuperLU ↔ XGC	1

EXASCALE COMPUTING PROJECT

ECP: Key Takeaways

- The Exascale Computing Project (ECP) is not *just* about developing and demonstrating the ability of new and enhanced DOE mission critical applications to tackle currently unsolvable problems of National interest . . . but we also are building and deploying a new Extreme Scale Scientific Software Stack (E4S – e4s.io) that greatly lowers the barrier to adoption of new technologies and to porting on advanced hardware. We are building a scientific software ecosystem for decades to come that is present and supports scientific computing from laptops to desktops to clusters to leadership systems

- The fundamental tenant of ECP is not about building boutique applications and a software ecosystem that can only execute on the Nation's largest systems, but it is about *accelerated node computing*, namely designing, implementing, delivering, and deploying advanced agile software that effectively exploits heterogeneous node hardware on today and tomorrow's laptops and desktops

- We view *accelerators* as any compute hardware specifically designed to accelerate certain mathematical operations (typically with floating point numbers) that are typical outcomes of popular and commonly used algorithms. We often use the term GPUs synonymously with accelerators.

- Compute hardware, from laptop to the largest systems in the world (e.g., ORNL's Summit system), are made up of *accelerated nodes*. Accelerated-node computing is *here to stay*
  - Accelerators today: GPUs   Tomorrow: better GPUs or FPGAs or other ASICs?   Near future: quantum?

- ECP's first-mover applications & E4S software stack are available for testing (even on laptops) and have greatly demystified and lowered the barrier to productive utilization of heterogeneous accelerated-node hardware.

Retrospective

- The US Department of Energy (DOE) has been a leader in High Performance Computing and ”invented” it for the purposes of “design predictability” 80 years ago. **Lots of lessons learned and ROI evidence to share**. ☺
- **Development and application of advanced, predictive modeling and simulation (M&S) – both computational and data science – has long been a mainstay and critical crosscutting technology for the DOE** and its National Laboratories (17 of them!) in achieving its mission goals in science, technology, and national security. This has **never been more vibrant and foundational than today**.
- **Accelerated compute performance (FLOPS, memory, memory B/W, etc.) and enhanced physical models, numerical algorithms, and software architecture enabled by this performance directly correlate with more predictive M&S tools, technologies, outcomes, impact**. This does not come without difficulties, challenges, pain, and perseverance: from GF to TF to PF to EF. We celebrate these milestones - **each one comes with ”tipping points” that are disruptive for app and software stack development yet accompanied by (often unanticipated) high ROI**
- The **EF “exascale era” (>10^{18} floating operations / sec) is upon us, and many institutions and agencies have been preparing and investing** for this milestone for over a decade: DOE included!
- DOE’s Exascale Computing Initiative (ECI), of which the Exascale Computing Project (ECP) is a part, was initiated almost six years ago and is not only poised and ready to demonstrate the tremendous “science return” of this technology, but **we anxious to share and demonstrate this technology with the fusion community**

Now is the time to join us and jump in the deep end!

To *exascale* and beyond...

ECP

EXASCALE COMPUTING PROJECT

Courtesy Tim Germann, Los Alamos National Laboratory

Questions?

kothe@ornl.gov, https://www.exascaleproject.org/contact-us/

For more info

- Alexander F. et al. *Exascale Applications: Skin in the Game,* Phil. Trans. R. Soc. A 378: 20190056 (2020) (http://dx.doi.org/10.1098/rsta.2019.0056).
- Douglas Kothe, Stephen Lee, and Irene Qualters, *Exascale Computing in the United States,* Computing in Science and Engineering 21(1), 17-29 (2019).

Supplemental Material

ECP Project Office Organization

Board of Directors
Thomas Zacharia, Chair (Director, ORNL)
Thom Mason, Vice Chair (Director, LANL)

Al Geist, ORNL
Chief Technology Officer

Exascale Computing Project
Doug Kothe, ORNL, Project Director
Lori Diachin, LLNL, Deputy Project Director
Lora Wolfe, ORNL
Executive Assistant

Industry and Agency Council
Dave Kepczynski, GE, and
Fran Hill, DoD, Co-Chairs
David Martin, ANL, and
Suzy Tichenor, ORNL,
Co-Executive Directors

Julia White, ORNL
Technical Operations

Mike Bernhardt, ORNL
Scott Gibson, ORNL
Communications & Outreach

Andrew Siegel
Erik Draeger

Mike Heroux
Lois Curfman McInnes

Katie Antypas
Susan Coghlan

Project Office Support
Megan Fielden, Human Resources
Sam Howard, Export Control Analyst
Mike Hulsey, Business Management
Michael Johnson, Legal
Willy Besancenez, Procurement
Kim Milburn, Finance Officer
Debi Shepard, Doc Control
Susan Ochs, Partnerships and Points of Contacts at the Core Laboratories

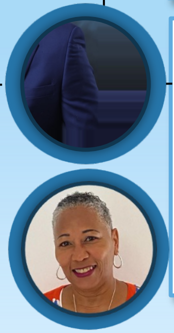

Project Management
Tom Cook,
ORNL Director
Corliss Thompson, ORNL
Sr. Administrative Assistant

Jonathan Wilson
Information Technology

Shaun Fomby
Project Controls and Risk

Andrew Arz
Joe Lawson
Nate Noll
Ryan Winter

Julia White,
Risk Manager
Jim Curley

Shaun Fomby, Lead
Monika Czajkowski

David Ellestad, Lead
Alicia Grundhoffer

Jordan Williams, Lead
Heather Aczel

EXASCALE COMPUTING PROJECT

HI leadership team : Accomplished technical leaders with Facility experience

Katie Antypas, HI Director (2.4)
15 years experiencing supporting HPC users and deploying HPC systems (LBNL)

Bronis de Supinski, PathForward (2.4.1)
5 years as the CTO for the Livermore Computing facility (LLNL)

Scott Pakin, HW Evaluation (2.4.2)
17 years in performance analysis and SW development at the ACES Facility (LANL)

Scott Parker, Application Integration at Facilities (2.4.3)
13+ years experience working on performance optimization for scientific applications (ALCF)

Susan Coghlan, HI Deputy Director (2.4)
30 years experience acquiring, deploying, managing extreme scale systems at DOE Facilities (Argonne)

Ryan Adamson, Software Deployment at Facilities (2.4.4)
12 years of systems and security administration, recently promoted to OLCF HPC Core Operations Group Lead (ORNL)

Haritha Siddabathuni Som, Facility Resource Utilization (2.4.5)
14 years in field and manager of the ALCF User Experience Team (ANL)

Ashley Barker, Training and Productivity (2.4.6)
8 years as a group leader of user assistance and outreach at the OLCF (ORNL)

ECP Software Technology Leadership Team

Mike Heroux, Software Technology Director
Mike has been involved in scientific software R&D for 30 years. His first 10 were at Cray in the LIBSCI and scalable apps groups. At Sandia he started the Trilinos and Mantevo projects, is author of the HPCG benchmark for TOP500, and leads productivity and sustainability efforts for DOE.

Lois Curfman McInnes, Software Technology Deputy Director
Lois is a senior computational scientist in the Mathematics and Computer Science Division of ANL. She has over 20 years of experience in HPC numerical software, including development of PETSc and leadership of multi-institutional work toward sustainable scientific software ecosystems.

Rajeev Thakur, Programming Models and Runtimes (2.3.1)
Rajeev is a senior computer scientist at ANL and most recently led the ECP Software Technology focus area. His research interests are in parallel programming models, runtime systems, communication libraries, and scalable parallel I/O. He has been involved in the development of open-source software for large-scale HPC systems for over 20 years.

Jeff Vetter, Development Tools (2.3.2)
Jeff is a computer scientist at ORNL, where he leads the Future Technologies Group. He has been involved in research and development of architectures and software for emerging technologies, such as heterogeneous computing and nonvolatile memory, for HPC for over 15 years.

Xaioye (Sherry) Li, Math Libraries (2.3.3)
Sherry is a senior scientist at Berkeley Lab. She has over 20 years of experience in high-performance numerical software, including development of SuperLU and related linear algebra algorithms and software.

Jim Ahrens, Data and Visualization (2.3.4)
Jim is a senior research scientist at the Los Alamos National Laboratory (LANL) and an expert in data science at scale. He started and actively contributes to many open-source data science packages including ParaView and Cinema.

Todd Munson, Software Ecosystem and Delivery (2.3.5)
Todd is a computational scientist in the Math and Computer Science Division of ANL. He has nearly 20 years of experience in high-performance numerical software, including development of PETSc/TAO and project management leadership in the ECP CODAR project.

Kathryn Mohror, NNSA ST (2.3.6)
Kathryn is Group Leader for the CASC Data Analysis Group at LLNL. Her work focuses on I/O for extreme scale systems, scalable performance analysis and tuning, fault tolerance, and parallel programming paradigms. She is a 2019 recipient of the DOE Early Career Award.

AD Leadership Team

Andrew Siegel, Director (2.2)

Erik Draeger, Deputy Director (2.2)

Jack Deslippe, Chemistry and Materials Applications (2.2.1)

Bill Hart, Data Analytics and Optimization Applications (2.2.4)

Tom Evans, Energy Applications (2.2.2)

Marianne Francois, National Security Applications (2.2.5)

Dan Martin, Earth and Space Science Applications (2.2.3)

Tim Germann, Co-Design (2.2.6)

ECP website is content-rich

www.exascaleproject.org

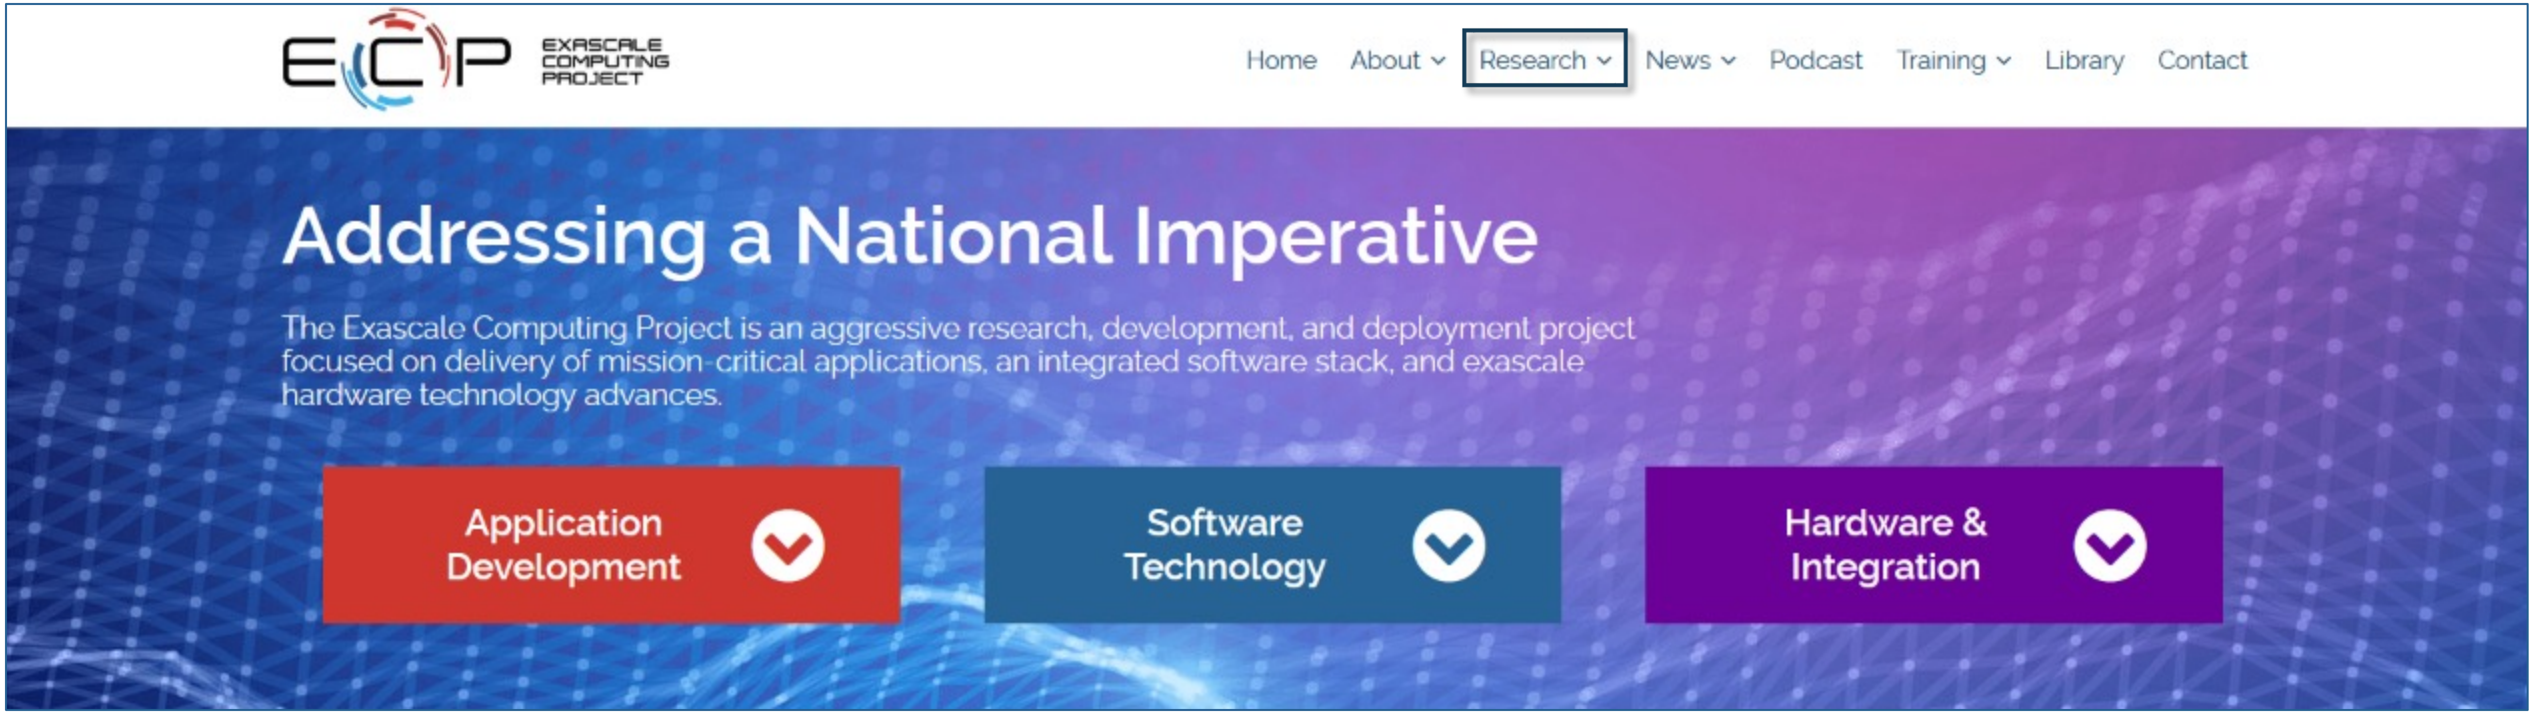

Descriptions of 24 Application Development projects, 6 Codesign Centers, Software Technology activities spanning programming models and run times, math libraries, data and visualization, and the integrated delivery of ECP products on targeted systems at leading DOE HPC facilities.

ECP website is content-rich

www.exascaleproject.org

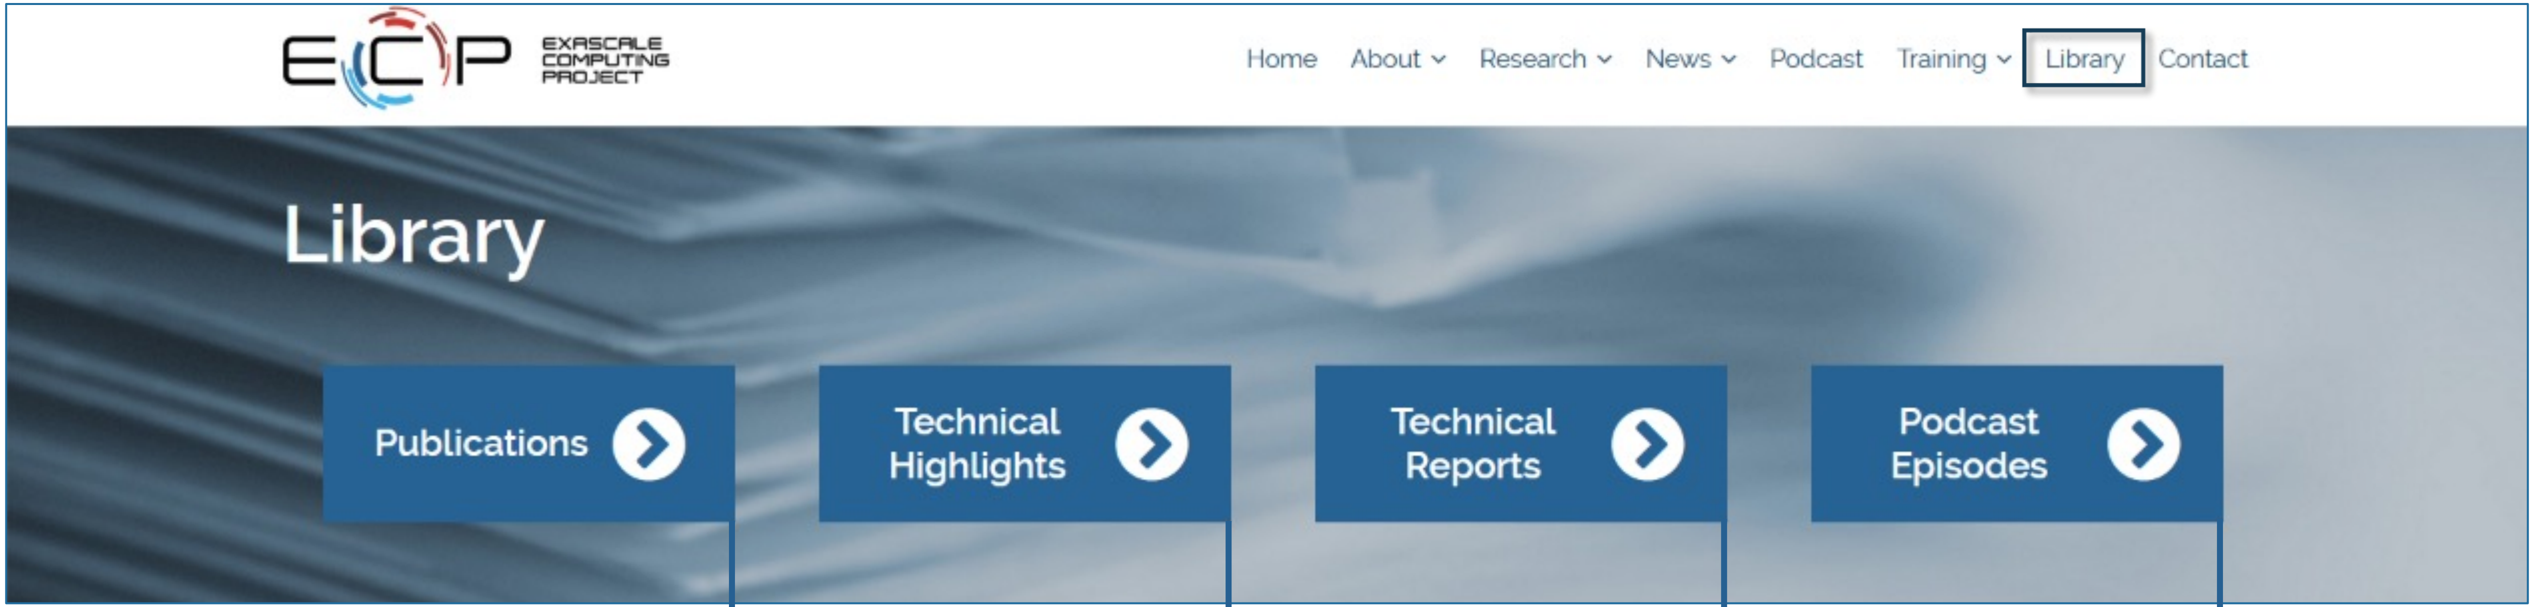

More than 370 journal publications and conference proceedings are listed on our Publications page.

Summaries highlighting publications are released every 1-2 weeks.

Technical highlight articles are produced by ECP writers on a monthly basis.

Technical reports include annual assessments by Application Development and Software Technology Teams.

The ECP communications team produces podcasts featuring ECP research activities.

ECP special journal issues: a component of proactive outreach

Appearing in the International Journal of High Performance Computing Applications

Codesign and Computational Motifs

- AMReX: Block-structured adaptive mesh refinement for multiphysics applications
- Co-design Center for Exascale Machine Learning Technologies
- Efficient exascale discretizations: High-order finite element methods
- Enabling particle applications for exascale computing platforms
- ExaGraph: Graph and Combinatorial Methods for Enabling Exascale Applications
- Online data analysis and reduction: An important Co-design motif for extreme-scale computers

Software Engineering

- MFIX-Exa: A path toward exascale CFD-DEM simulations
- Coupling of regional geophysics and local soil-structure models in the EQSIM fault-to-structure earthquake simulation framework
- ExaAM: Metal additive manufacturing simulation at the fidelity of the microstructure
- The Exascale Framework for High Fidelity coupled Simulations (EFFIS): Enabling whole device modeling in fusion science
- Unprecedented cloud resolution in a GPU-enabled full-physics atmospheric climate simulation on OLCF's summit supercomputer
- Exascale models of stellar explosions: Quintessential multi-physics simulation