Challenges of Exascale Computing

Paul Messina

Argonne National Laboratory and
Director, Exascale Computing Project

ModSim 2017

August 9, 2017

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

The U.S. DOE Exascale Computing Project

- The ECP was established to accelerate delivery of capable exascale computing systems that integrate hardware and software capability to deliver approximately 50 times more performance than today's 20-petaflops machines on mission critical applications
  - DOE is a lead agency for this mission, along with DoD and NSF
    - Deployment agencies: NASA, FBI, NIH, DHS, and NOAA
- Timeline: at least one exascale system will be delivered in 2021 to a DOE Office of Science Leadership Computing Facility (Argonne LCF and/or Oak Ridge LCF)
  - ALCF and OLCF will have diverse architectures
  - A National Nuclear Security Administration (NNSA) facility will field an exascale system in 2022-2023; could be the ALCF or OLCF choice, or a third choice
- ECP's work encompasses
  - Applications
  - System software
  - Hardware technologies and architectures
  - Workforce development to meet scientific and national security mission needs

Key high-level technical challenges that must be tackled to achieve exascale

- *Massive Parallelism* – 100+ times greater than today's largest systems
- *Memory and Storage* – effective use of many levels of hierarchy
  - Memory and storage efficiencies consistent with increased computational rates and data movement requirements
- *Reliability* – system adaptation and recovery from faults in complex system components and designs
- *Energy Consumption* – Energy consumption reduced beyond current industry roadmaps
  - would be prohibitively expensive at this scale
  - hardware and software techniques for minimizing it

What Is a Capable Exascale Computing System?

- Delivers 50× the performance of today's 20 PF systems, supporting applications that deliver high-fidelity solutions in less time and address problems of greater complexity
  – NOTE: no LINPACK or peak FLOPS target
- Operates in a power envelope of 20–30 MW
- Is sufficiently resilient (perceived fault rate: ≤1/week)
- Includes a software stack that supports a broad spectrum of applications and workloads

This ecosystem will be developed using a co-design approach to deliver new software, applications, platforms, and computational science capabilities at heretofore unseen scale

Predictable performance challenges

- Node performance: always a major component of overall performance, now even more so
- I/O performance: almost always a bottleneck, now even more so
  - Hardware configuration
  - Software/methods
- Industrial HPC users often rely on commercial software that has not been scaled for use on large systems

Not yet known challenges: Where is that puck going?

- New computational patterns even in traditional simulations
  – UQ, more multi-physics, complex workflows
- New (to HPC) applications
  – Deep Learning, Machine Learning at huge scales
- New hardware architectures (well new as commercial products)
- Squeezing performance (déjà vu)
  – Half precision
  – Single precision
  – Mixed precision

Performance modeling and simulation are central aspects of the ECP

- Evaluate new architectural features
- Predict whether applications will meet Key Performance Parameters
  – 50X performance, performance portability
- Measure performance gains as defined in this project
- Determine reasons for unexpected performance
- Guide development of hardware, applications and supporting software

ECP is a Collaboration Among Six US DOE National Laboratories

- The ECP draws from the Nation's 6 premier computing national laboratories
- A Memorandum of Agreement for the ECP was signed by each Laboratory Director defining roles and responsibilities
- Funding comes from two sources: DOE Office of Science and NNSA Advanced Simulation and Computing (ASC) program

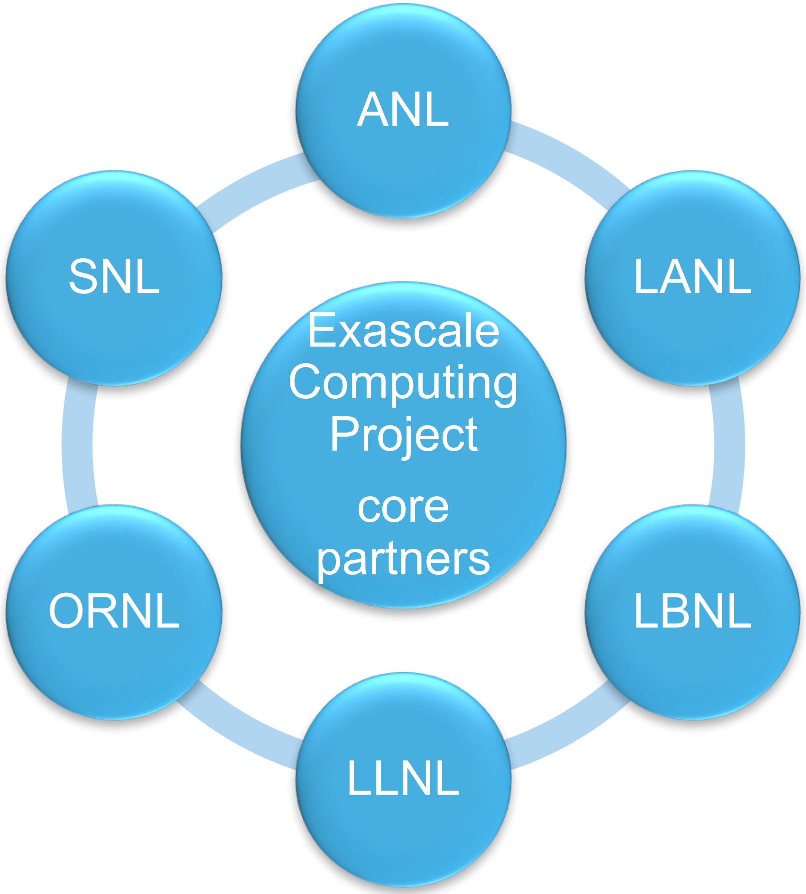

EXASCALE COMPUTING PROJECT

But the work is carried out at many institutions

EXASCALE COMPUTING PROJECT

DOE LABORATORIES & AGENCY PARTNERS
22

PRIVATE SECTOR PARTNERS
10

UNIVERSITY RESEARCH PARTNERS
39

INDUSTRY COUNCIL MEMBERS
18

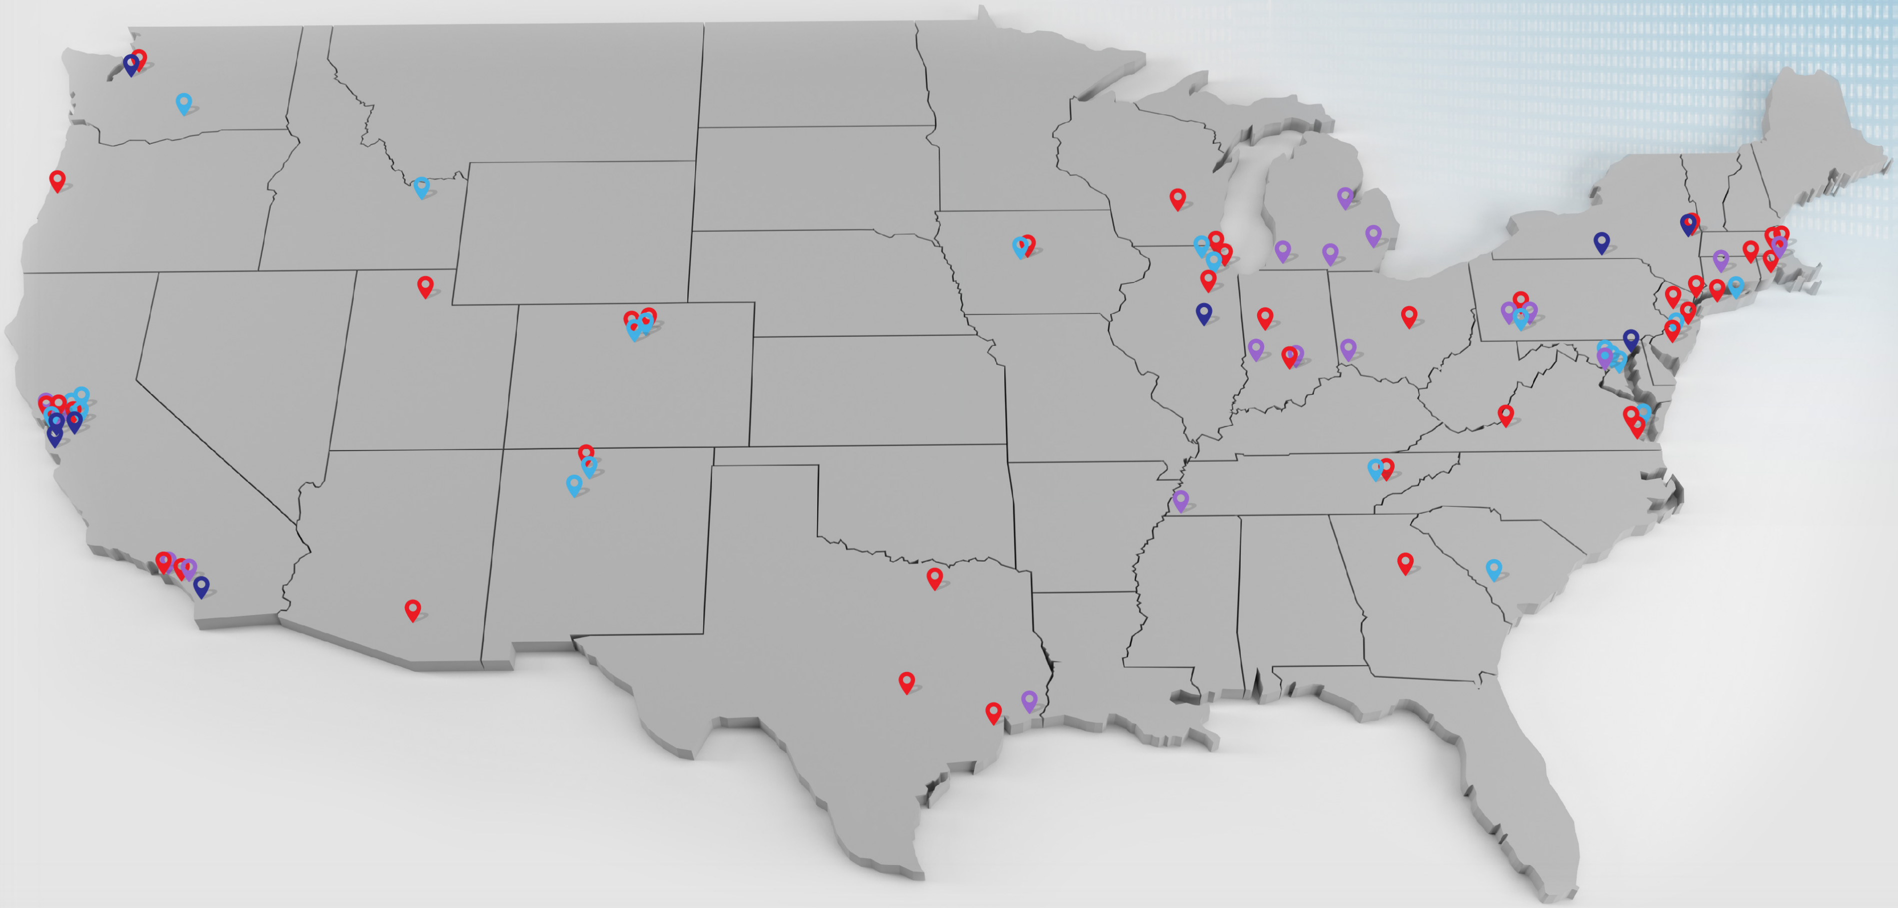

THE ECP ECOSYSTEM

- 800+ Researchers
- 25 Application Development Projects
- 66 Software Development Projects
- 5 Co-Design Centers

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

ECP uses co-design and integration to achieve exascale computing

Application Development	Software Technology	Hardware Technology
Science and mission applications	Scalable software stack	Hardware technology elements

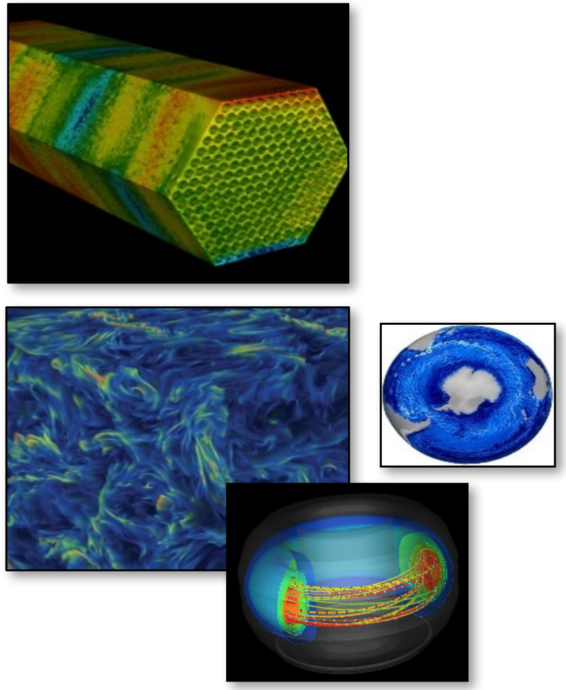

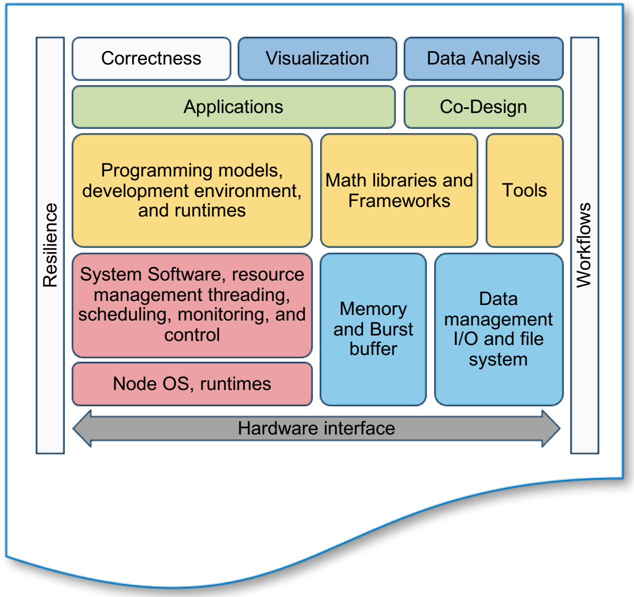

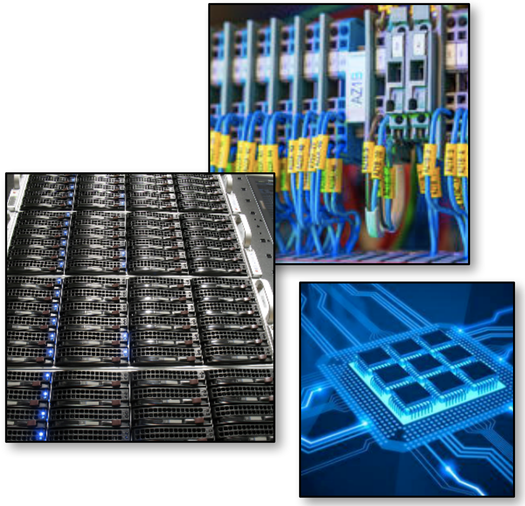

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Nuclear Energy (NE)

Accelerate design and commercialization of next-generation small modular reactors

- ✓Climate Action Plan
- ✓SMR licensing support
- ✓GAIN

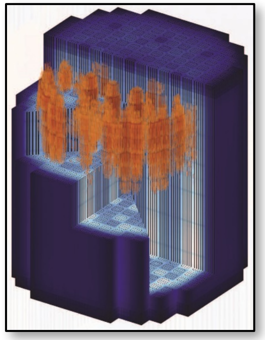

Climate (BER)

Accurate regional impact assessment of climate change

- ✓Climate Action Plan

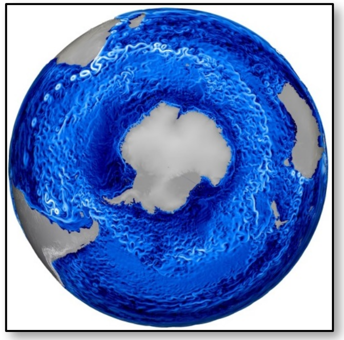

Chemical Science (BES, BER)

Biofuel catalysts design; stress-resistant crops

- ✓Climate Action Plan
- ✓MGI

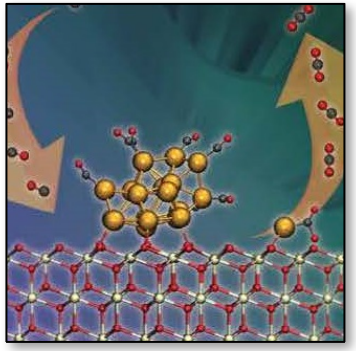

Wind Energy (EERE)

Increase efficiency and reduce cost of turbine wind plants sited in complex terrains

- ✓Climate Action Plan

Combustion (BES)

Design high-efficiency, low-emission combustion engines and gas turbines

- ✓2020 greenhouse gas and 2030 carbon emission goals

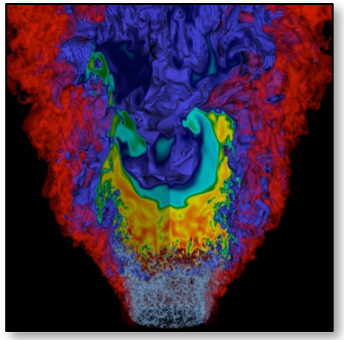

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Materials Science (BES)

Find, predict, and control materials and properties: property change due to hetero-interfaces and complex structures

✓MGI

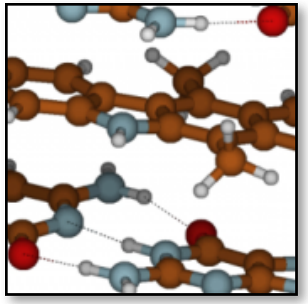

Nuclear Physics (NP)

QCD-based elucidation of fundamental laws of nature: SM validation and beyond SM discoveries

✓2015 Long Range Plan for Nuclear Science
✓RHIC, CEBAF, FRIB

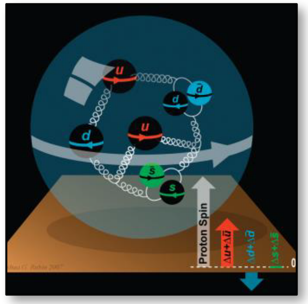

Nuclear Materials (BES, NE, FES)

Extend nuclear reactor fuel burnup and develop fusion reactor plasma-facing materials

✓Climate Action Plan
✓MGI
✓LWR Sustainability
✓ITER
✓Stockpile Stewardship Program

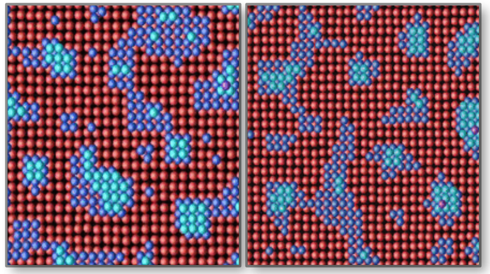

Accelerator Physics (HEP)

Practical economic design of 1 TeV electron-positron high-energy collider with plasma wakefield acceleration

✓>30k accelerators in industry, security, energy, environment, medicine

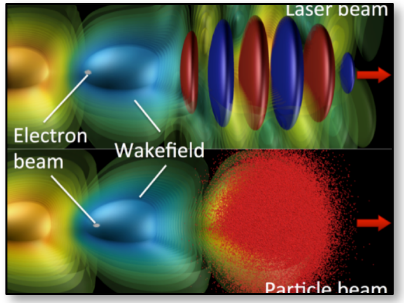

Materials Science (BES)

Protein structure and dynamics; 3D molecular structure design of engineering functional properties

✓MGI
✓LCLS-II 2025 Path Forward

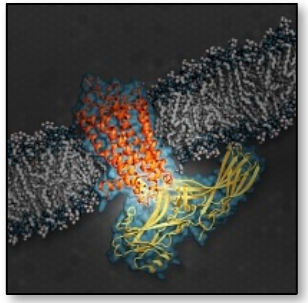

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Magnetic Fusion Energy (FES)

Predict and guide stable ITER operational performance with an integrated whole device model

- ✓ITER
- ✓Fusion experiments: NSTX, DIII-D, Alcator C-Mod

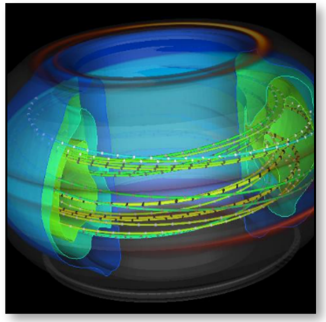

Advanced Manufacturing (EERE)

Additive manufacturing process design for qualifiable metal components

- ✓NNMIs
- ✓Clean Energy Manufacturing Initiative

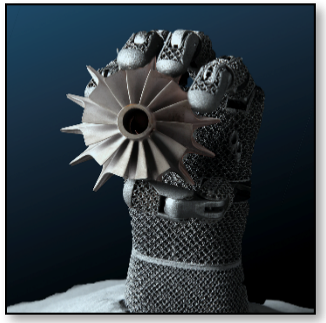

Cosmology (HEP)

Cosmological probe of standard model of particle physics: Inflation, dark matter, dark energy

- ✓Particle Physics Project Prioritization Panel (P5)

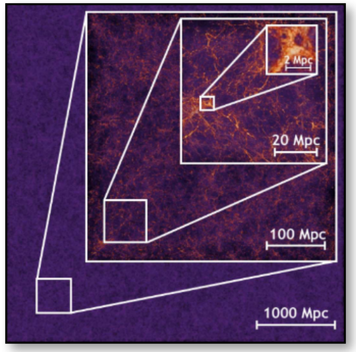

Geoscience (BES, BER, EERE, FE, NE)

Safe and efficient use of subsurface for carbon capture and storage, petroleum extraction, geothermal energy, nuclear waste

- ✓EERE Forge
- ✓FE NRAP
- ✓Energy-Water Nexus
- ✓SubTER Crosscut

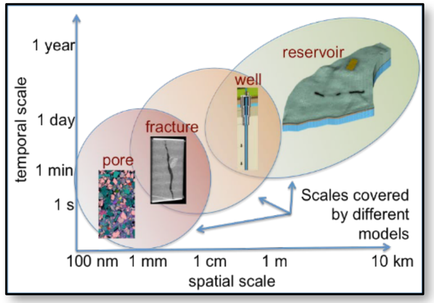

Precision Medicine for Cancer (NIH)

Accelerate and translate cancer research in RAS pathways, drug responses, treatment strategies

- ✓Precision Medicine in Oncology
- ✓Cancer Moonshot

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Seismic (EERE, NE, NNSA)	Carbon Capture and Storage (FE)	Chemical Science (BES)	Urban Systems Science (EERE)
Reliable earthquake hazard and risk assessment in relevant frequency ranges	Scaling carbon capture/storage laboratory designs of multiphase reactors to industrial size	Design catalysts for conversion of cellulosic-based chemicals into fuels, bioproducts	Retrofit and improve urban districts with new technologies, knowledge, and tools
✓DOE Critical Facilities Risk Assessment ✓Urban area risk assessment ✓Treaty verification	✓Climate Action Plan ✓SunShot ✓2020 greenhouse gas ✓2030 carbon emission goals	✓Climate Action Plan ✓SunShot Initiative ✓MGI	✓Energy-Water Nexus ✓Smart Cities Initiative

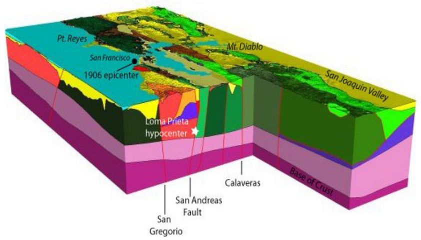

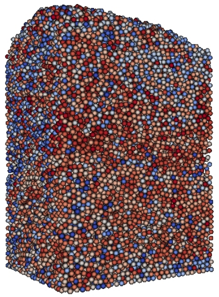

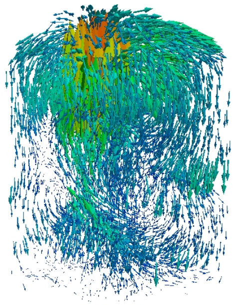

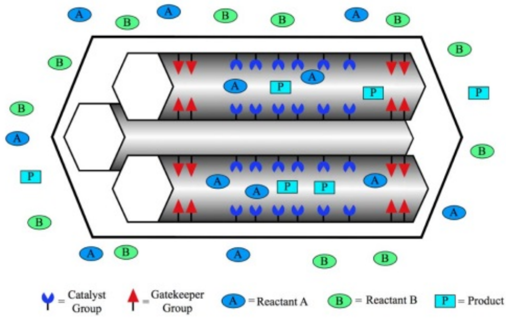

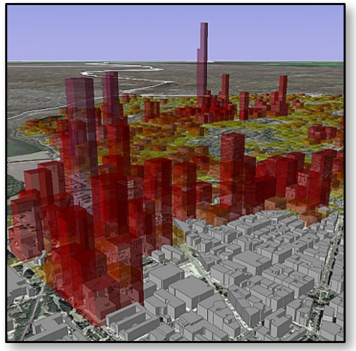

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Metegenomics (BER)	Astrophysics (NP)	Power Grid (EERE, OE)
Leveraging microbial diversity in metegenomic datasets for new products and life forms	Demystify origin of chemical elements (> Fe); confirm LIGO gravitational wave and DUNE neutrino signatures	Reliably and efficiently planning our nation's grid for societal drivers: rapidly increasing renewable energy penetration, more active consumers
✓Climate Action Plan ✓Human Microbiome Project ✓Marine Microbiome Initiative	✓2015 Long Range Plan for Nuclear Science ✓Origin of universe and nuclear matter in universe	✓Grid Modernization Initiative ✓Climate Action Plan

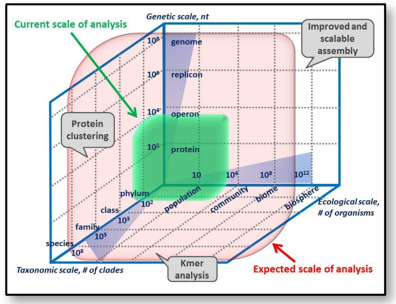

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE NNSA application development projects

Stockpile Stewardship

Gaps and Opportunities

- Complete understanding of thermonuclear boost
- Resolution of important length scales with appropriate fidelity

Simulation Challenge Problems

- 3D boost simulations with multiple coupled physical processes at unprecedented resolution
- Detailed highly resolved 3D nuclear safety simulations
- UQ performed in 3D at lower resolution with sub-grid models to capture unresolved physics

Prospective Outcomes and Impact

- Simulation of appropriately complex material at engineering scale through formal and rigorous validation of sub-grid models
- Improved interpretation and understanding of nuclear test data
- High-confidence predictions of thermonuclear boost less dependent upon 2D calibrations

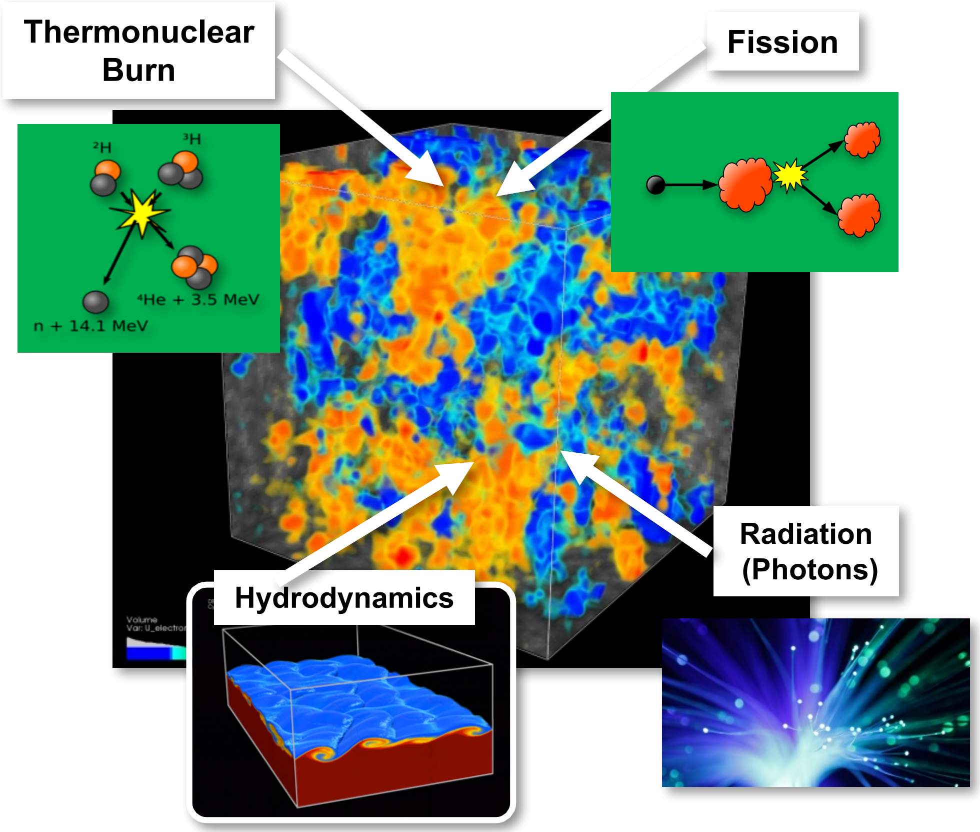

Each application has specified a challenge problem
Example: QMCPACK Challenge problem

“Calculate the cohesive energy of a 1024-atom supercell of NiO to an accuracy of 10 meV per NiO formula unit using a full LCF class machine in a reasonable and scientifically productive amount of wall clock time, ~1 day.”

EXASCALE COMPUTING PROJECT

ECP Co-Design Centers

- **A Co-Design Center for Online Data Analysis and Reduction at the Exascale (CODAR)**
  – **Motifs**: Online data analysis and reduction
  – Address growing disparity between simulation speeds and I/O rates rendering it infeasible for HPC and data analytic applications to perform offline analysis. Target common data analysis and reduction methods (e.g., feature and outlier detection, compression) and methods specific to particular data types and domains (e.g., particles, FEM)
- **Block-Structured AMR Co-Design Center (AMReX)**
  – **Motifs**: Structured Mesh, Block-Structured AMR, Particles
  – New block-structured AMR framework (AMReX) for systems of nonlinear PDEs, providing basis for temporal and spatial discretization strategy for DOE applications. Unified infrastructure to effectively utilize exascale and reduce computational cost and memory footprint while preserving local descriptions of physical processes in complex multi-physics algorithms
- **Center for Efficient Exascale Discretizations (CEED)**
  – **Motifs**: Unstructured Mesh, Spectral Methods, Finite Element (FE) Methods
  – Develop FE discretization libraries to enable unstructured PDE-based applications to take full advantage of exascale resources without the need to “reinvent the wheel” of complicated FE machinery on coming exascale hardware
- **Co-Design Center for Particle Applications (CoPA)**
  – **Motif(s)**: Particles (involving particle-particle and particle-mesh interactions)
  – Focus on four sub-motifs: short-range particle-particle (e.g., MD and SPH), long-range particle-particle (e.g., electrostatic and gravitational), particle-in-cell (PIC), and additional sparse matrix and graph operations of linear-scaling quantum MD
- **Combinatorial Methods for Enabling Exascale Applications (ExaGraph)**
  – **Motif(s)**: Graph traversals; graph matching; graph coloring; graph clustering, including clique enumeration, parallel branch-and-bound, graph partitioning
  – Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms that play a critical enabling role in numerous scientific applications. The irregular memory access nature of these algorithms makes them difficult algorithmic kernels to implement on parallel systems

Additional Application Development activities

- Exascale Proxy Applications Suite
  - Assemble and curate a proxy app suite composed of proxies developed by other ECP projects that represent the most important features (especially performance) of exascale applications.
  - Friday morning Dave Richards will provide details on this part of the effort
- Application Assessment
- IDEAS-ECP – Advancing Software Productivity for Exascale Applications
  - Customize and curate methodologies for ECP app productivity & sustainability
  - Create an ECP Application Development Kit of customizable resources for improving scientific software development
- Training
  - Argonne Training Program on Extreme Scale Computing, many others

Application assessment

- Mission: independent assessment of applications development projects, relative to their stated project plan / milestones. Identify problems (if any) and propose methods to solve them.
- Two levels:
  - High-level review of all projects on an annual cycle.
  - Deep-dive assessment of two projects / quarter.
- Assessment areas:
  - Performance / scalability
  - Performance portability
  - Science capability
  - Integration
  - Software practices.
- Organization: two teams, each composed of experts in algorithms, software engineering, performance engineering.

ECP Training: 1.2.4.01
EPIC ADTR01-23: FY17 Q3 Training Activities

APPROVED

ECP WBS: Training and Productivity, 1.2.4
PI: Name, Ashley Barker
Members: Institutional Partners

Scope & Objectives

- The purpose of the ECP Training and Productivity (T&P) effort is to disseminate and transfer application development knowledge, lessons learned, and best practices across application teams and provide training on key ECP technologies.
- Goal is to facilitate at least two training events per quarter
- All ECP Training advertised on external ECP website: https://exascaleproject.org/
- Q3 included **Python in HPC and Intermediate Git (***Part of ECP IDEAS Best Practice Series***)** and **OpenMP** Tutorials

Python in HPC

https://exascaleproject.org/event/python-in-hpc-2/

Webinar Wrap-Up

- Occurred June 7, 2017
- 210 Registrations
- 132 Attendees
- Faciliated by
- The presentation materials and video are available on ECP IDEAS website and the ECP External Website.

OpenMP Tutorial

https://exascaleproject.org/event/openmp-tutorial/

Webinar Wrap-Up

- Occurred June 28, 2017
- 113 Registrations
- 60 Attendees
- Tutorial led by Oscar Hernandez and Tom Scogland of the ECP S&T Project, SOLLVE
- The presentation materials are available on the ECP External Website.

Intermediate Git

https://exascaleproject.org/event/intermediate-git/

Webinar Wrap-Up

- Occurred July 12, 2017
- 148 Registrations
- 77 Attendees
- Facilitated by Roscoe Bartlett of Sandia
- The presentation materials and video are available on ECP IDEAS website and the ECP External Website.

ECP Training: 1.2.4.01
EPIC ADTR01-24: FY17 Q4 Training Activities

IN PROGRESS

ECP WBS: Training and Productivity, 1.2.4
PI: Name, Ashley Barker
Members: Institutional Partners

Scope & Objectives

- The purpose of the ECP Training and Productivity (T&P) effort is to disseminate and transfer application development knowledge, lessons learned, and best practices across application teams and provide training on key ECP technologies
- Goal is to facilitate at least two training events per quarter
- All ECP Training advertised on external ECP website: https://exascaleproject.org/
- Q3 will include "**Using the Roofline Model and Intel Advisor**" (*Part of ECP IDEAS Best Practice Series*)

Using the Roofline Model and Intel Advisor

https://exascaleproject.org/event/using-the-roofline-model/

Using the Roofline Model and Intel Advisor

August 16 - August 16

The IDEAS Productivity project, in partnership with the DOE Computing Facilities of the ALCF, OLCF, and NERSC and the DOE Exascale Computing Project (ECP) is resuming the webinar series on Best Practices for HPC Software

Read More >

Webinar Updates

- Next in the ECP IDEAS Best Practice Series
- Scheduled for August 16, 2017
- Presented by Sam Williams of LBNL
- 65 registrations to-date

??

Your Event Could be Pictured HERE

EXASCALE COMPUTING PROJECT

ECP Requires Strong Integration to Achieve Capable Exascale

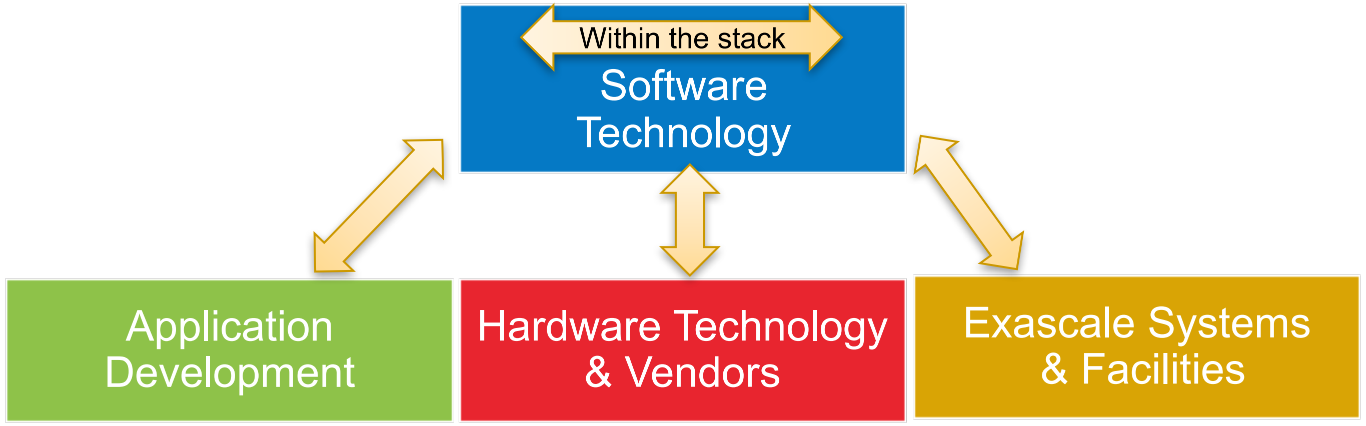

- To achieve a coherent software stack, we must integrate across all the focus areas
  - Understand and respond to the requirements from the apps but also help them understand challenges they may not yet be aware of
  - Understand and repond to the impact of hardware technologies and platform characteristics
  - Work with the facilities and vendors towards a successful stable deployment of our software technologies
  - Understand and respond to dependencies within the stack, avoiding duplication and scope creep
  - This is a comprehensive team effort — not a set of individual projects!

Software Technology Portfolio

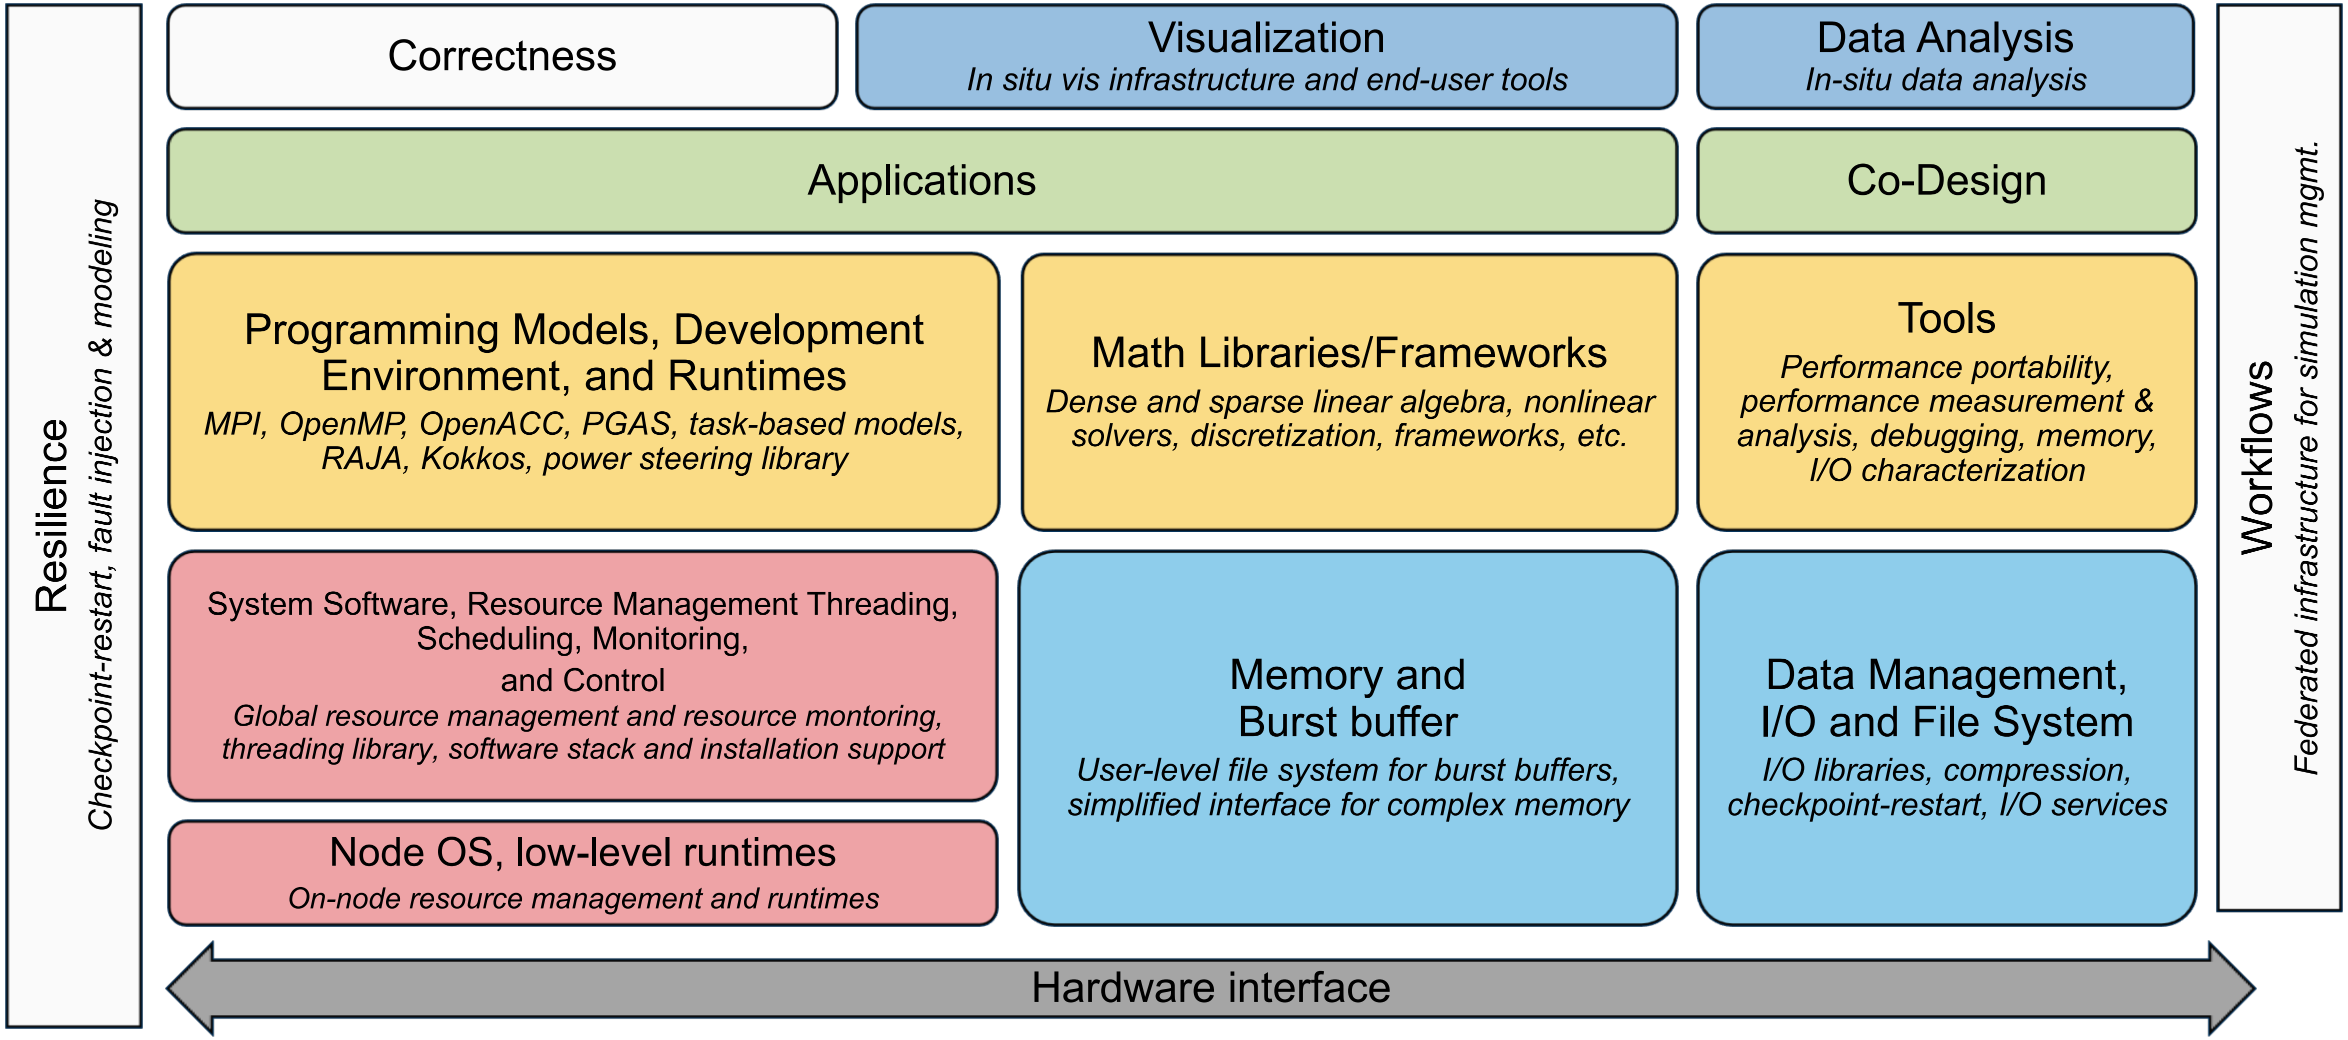

The ECP Software Technology Portfolio

Derived from

- Analysis of the software needs of exascale applications
- Inventory of software environments at major DOE HPC facilities (ALCF, OLCF, NERSC, LLNL, LANL, SNL)
  - For current systems and the next acquisition (CORAL, APEX)
- Expected software environment for an exascale system
- Requirements beyond the software environment provided by vendors of HPC systems

Math Libraries and Frameworks Projects (1)

Project	Description
Kokkos Kernels	Performance portable (CPU/GPU) sparse matrix and graph kernels
Trilinos Linear Solvers	Robust scalable preconditioned Krylov solvers for next-gen platforms
Trilinos PDE Components	Meshing, discretization, integration compatible with Trilinos solvers
Trilinos Embedded Analysis	Sensitivity analysis, optimization, uncertainty quantification
xSDK	Compatible, interoperable, turnkey installation & access to libraries
SUNDIALS	Preparing SUNDIALS for next-generation platforms
PETSc/TAO	Preparing PETSc/TAO for next-generation platforms

ECP EXASCALE COMPUTING PROJECT

Math Libraries and Frameworks Projects (2)

Project	Description
STRUMPACK/SuperLU	Next-generation sparse factorizations for multi-node CPU/GPU
ForTrilinos	Native, sustainable Fortran API to Trilinos next-gen capabilities
SLATE	Next-generation dense linear algebra for next-gen platforms
PEEKS	Latency-tolerant, production-quality iterative solvers and APIs
FleCSI	Framework for exploring next-gen component and execution models
MFEM	Advanced, high-order discretizations for next-gen platforms
ALExa	Multi-scale, multi-physics, sparse grid UQ

Tools

- **Goal:** A suite of tools and supporting unified infrastructure aimed at improving developer productivity

- **Current Portfolio:** Performance portability tools, performance measurement and analysis tools, debugging, memory tools, I/O characterization

Performance Portability Tools

- LLVM, PI: McCormick (LANL)
  - Augment and enhance the HPC toolchain with LLVM based tools and technology
- Autotuning, PI: Hall (Utah)
  - Performance portability across CPUs and GPUs through domain-specific optimization and autotuning using the CHiLL framework and SURF search space navigation
- ROSE, PI: Quinlan (LLNL)
  - Support the automated generation of code for current and future compute architectures
- PROTEAS, PI: Vetter (ORNL)
  - Pathfinding programming solutions based on directive-based methodologies directed at emerging architectural features such as heterogeneous and manycore processors, deep memory hierarchies, and nonvolatile memory systems (NVM)
  - Support for OpenACC

Performance Measurement and Analysis Tools

- PerfTools, PI: Schulz (LLNL)
  - Development of tools that help application and system developers understand the performance of their applications and HW/SW systems
- Caliper, PI: Schulz (LLNL)
  - Caliper is a simple annotation API, similar to timer libraries found in many codes, while under the hood Caliper is designed to combine the information from several additional sources
- EXA-PAPI, PI: Dongarra (UTK)
  - Consistent interface and methodology for the use of low-level performance counter hardware found across the system
- HPCToolkit, PI: Mellor-Crummey (Rice)
  - Extend HPCToolkit with improved capabilities for measurement and analysis of computation, data movement, communication, I/O, power, and energy consumption at extreme scale
- PerfAnal, PI: Hammond (SNL)
  - Tools for analysis of code performance on next generation architectures;
- PROTEAS/Tau, PI: Vetter (ORNL)
  - Extend Tau to support new architectural features and programming models at scale

EXASCALE COMPUTING PROJECT

ExaHDF5: Delivering Efficient Parallel I/O on Exascale Computing Systems

Suren Byna (LBNL, Lead PI), Quincey Koziol (LBNL), Scot Breitenfeld (The HDF Group),
Venkat Vishwanath, and Preeti Malakar (ANL)

- HDF5 is a mature technology being developed and released for 20 years
  - Among the Top-10 SW libraries used at NERSC, ALCF, and OLCF
- Features to be developed
  - Integration of Virtual Object Layer (VOL), data caching and prefetching using storage hierarchy, topology-aware I/O, async I/O, independent metadata updates, full single-writer – multiple-reader (SWMR), querying data and metadata, interoperability with netCDF and ADIOS file formats
- Several ECP applications either currently use or interested in using HDF5

AMReX – 5 apps

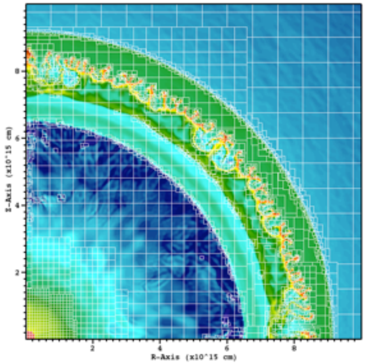

WarpX

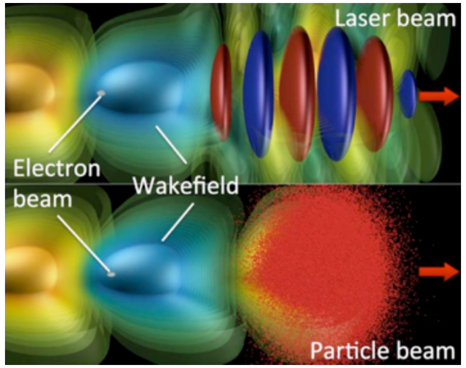

ExaFEL

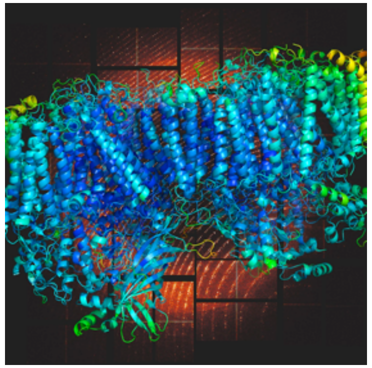

Subsurface simulation

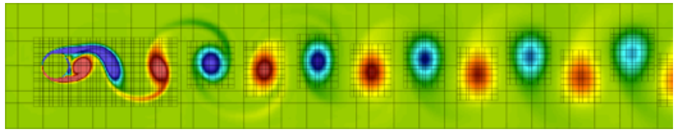

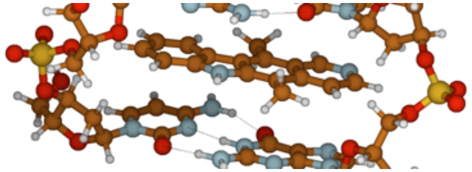

QMCPACK

CLASH

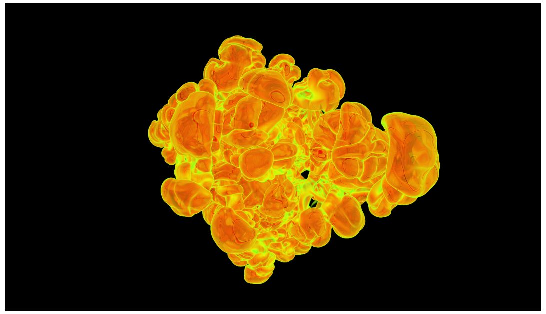

Hardware Technology R&D: PathForward Program

- The ECP PathForward program supports DOE-vendor collaborative R&D activities required to develop exascale systems with at least two diverse architectural features; quote from RFP:
  - PathForward seeks solutions that will improve application performance and developer productivity while maximizing energy efficiency and reliability of exascale systems.
- PathForward contracts were awarded recently to these companies:
  - Advanced Micro Devices (AMD)
  - Cray Inc. (CRAY)
  - Hewlett Packard Enterprise (HPE)
  - International Business Machines (IBM)
  - Intel Corp. (Intel)
  - NVIDIA Corp. (NVIDIA)

EXASCALE COMPUTING PROJECT

Hardware Technology R&D: Design Space Evaluation

- Apply laboratory architectural analysis capabilities and Abstract Machine Models to PathForward designs to support ECP co-design interactions
- Technology coverage areas and approaches
  - Memory technologies
  - Interconnect/system simulators
  - Analytical models
  - Node simulators
  - Abstract machine models and proxy architectures

Selected Challenges (1)

- How measure the 50X in performance or complexity?
- Programming models and languages
  - How many are viable to support?
  - Which ones to support in this project for existing applications?
  - What will new applications enabled by exascale need?
- Applications teams often feel they can't wait for math libraries to be available, have to roll their own – or they do not want to rely on software developed by others
  - The ECP is investing heavily on getting them ready on time AND satisfying application requirements

Measuring performance through Figures of Merit (FOMs)

- Doug Kothe will go into more detail on this topic in his presentation on Friday
- The short version follows

Measuring performance through Figures of Merit (FOMs)

- FOMs are defined so that they scale linearly (within measurement tolerances) with the delivered performance of the code.
  - For example, running a given benchmark 10x faster should result in a FOM that is ~10x larger.
  - Likewise, running a 10x larger problem in the same amount of time should also result in a ~10x increase in the measured FOM.
- The value of the FOM for each exascale science application will be described in the documentation for each application and calculated at the end of successful execution of each benchmark run of the application.
- The exascale FOMs will be compared to supplied baseline FOMs, which will be measured on today's largest computers in the US

EXASCALE COMPUTING PROJECT

Selected challenges (2)

- Many moving targets: applications, algorithms, system architectures
- Portability and performance portability
- Co-design and integration for applications, hardware, support software **at the scale of the ECP**, with the global computational science research community
- Guiding the vendors with our understanding of our applications and support software requirements while workload is changing
  - Difficult to design proxy applications that represent the key features of the full applications
  - Vendors can't deal with too many full or proxy applications

Challenges by area

- Application Development
  - Selecting challenge problems that require exascale and can realistically be developed in time
  - Portability with (partially) unknown target architectures
- Software Technology
  - Too many applications want deliverables from some of the ST projects
  - Will the applications use the ST results?
  - Coordinating with the vendors software stack
  - How meet needs of new types of applications or existing applications whose needs evolve due to exascale hardware features
- Hardware Technology
  - Will the vendor R&D projects result in better systems bid for facilities' RFP?

Challenges other than technical

- The "plumbing" aspects are a big challenge
- The exascale systems will be acquired by the usual DOE lab facilities
- Co-design with so many sub-projects in so many institutions
  - Different cultures, different terminologies
- Need tight ties with the facilities that will acquire and operate the exascale systems
- Creating an integrated software stack
- Dual funding sources
- Coordinating and collaborating with the vendors' software plans

Communication is the key approach to tackling those challenges

- Co-design is a contact sport
- In this project we create many venues and mechanisms for focused and deep discussions among the participants in all three focus areas
  - Meetings, phone calls
  - Confluence software that enables view into all project activities, discussions, milestones, reports, etc.
- Essential but it comes at a cost: communication takes time and needs to be done continually
- And there are signs of success

ECP Applications Adopt New Infrastructure for Block-Structured Adaptive Mesh Refinement Developed by AMReX Co-Design Center

ECP WBS 1.2.5.03: AMReX
PI: John Bell, LBNL
Members: LBNL, ANL, NREL

Scope & Objectives

- Develop infrastructure to enable block-structured adaptive mesh refinement on exascale architectures
  - Core mesh, particle & particle-mesh operations on adaptive mesh hierarchy
  - Support for multiple time-stepping approaches
  - Embedded boundary representation of complex geometry
  - Performance portability for different architectures
- Current activities focused on:
  - Establishing support for core AMR functionality
  - Engagement of applications

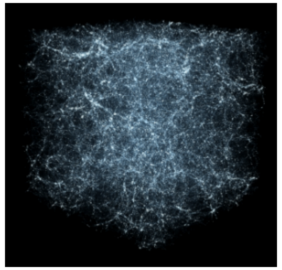
Cosmology

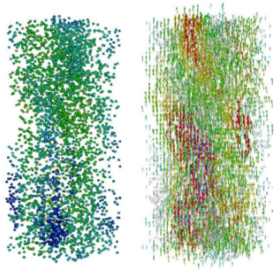
Multiphase flow in chemical reactors

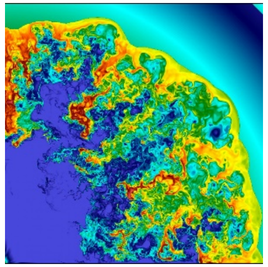
Astrophysics

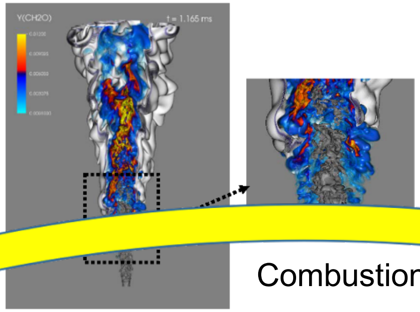
Combustion

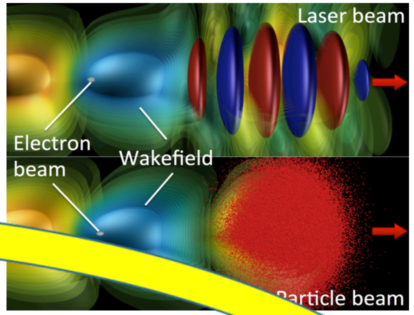

Accelerator design

Impact

- Established a next-generation framework for developin lock-structured adaptive mesh refinement algorithms for c ent and emerging architectures
- Provides a common framework for multiple ECP ap cations that use AMR
- Provides a common focal point for software technolo hardware technology and vendors to leverage activitie over multiple applications
- Broad constituency within Office of Science and NNSA

Project Accomplishment

- New AMReX code framework adopted by multiple ECP applications
  - Accelerator modeling – WarpX
  - Astrophysics -- ExaStar (CASTRO)
  - Combustion – PeleC and PeleLM
  - Cosmology – ExaSky (Nyx)
  - Multiphase flow – MFIX-Exa
- AMReX code framework publicly released

Deliverables: AMReX software available at h ithub.com/AMReX-Codes/amrex

EXASCALE COMPUTING PROJECT

Co-design in action: shared milestone examples

- OMPI-X ST project (Open MPI for exascale) will design, implement, and demonstrate new prototype APIs to facilitate improved coordination between MPI and OpenMP runtimes for use of threads, message delivery, and rank/endpoint placement. (Due 12/31/2018)
- Execution Plan:
  - Interact with ECP Application Development teams to identify requirements about runtime coordination.
  - Interact with appropriate ECP Software Technology teams to define coordination APIs for MPI and OpenMP (and other) runtimes.
  - Experiment and demonstrate the benefit running an ECP application or mini-application running on the ECP testbeds

Shared milestones examples

- SUNDIALS and the AMReX co-design center
  - Enable CVODE to be used from AMReX (Due 6/30/2017)
  - Description: AMReX will provide a build system and interface to CVODE within AMReX SUNDIALS will support AMReX in being able to link to CVODE from within the AMReX build system.
- Evaluation of VeloC (efficient checkpoint/restart) with the QCD application code
- Evaluation of VeloC (efficient checkpoint/restart) with the EXAALT application code (Molecular Dynamics at the Exascale application)

A request for additional help from you: User Training

- An additional contribution by the performance modeling and simulation community could be training
- I know that many of you engage in training activities
- But the ECP community could use more
- Your machine models might help with performance portability strategies

I think we can

- Create an exascale ecosystem that will support many applications and modes of usage
- Foster a culture of co-design, software engineering, and collaboration
- Yes, I am an optimist, but I truly think we can succeed

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering and early testbed platforms, in support of the nation's exascale computing imperative.

Thank you!

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science