Exascale Applications: Skin in the Game

EXASCALE COMPUTING PROJECT

Douglas B. Kothe (Oak Ridge National Laboratory)
Director, The US Department of Energy (DOE) Exascale Computing Project (ECP)

Numerical Algorithms for High-Performance Computational Science
The Royal Society
London, England
April 8, 2019

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

Sidebar

- *Skin in the game**: having "incurred risk by being involved in achieving a goal"
  - Where "*skin* is a synecdoche for the person involved and *game* is the metaphor for actions on the field of play under discussion"
- For exascale applications under development in the U.S Department of Energy (DOE) Exascale Computing Project (ECP), this is appropriate, with
  - The *skin* is exascale applications, and
  - The *game* is delivering comprehensive science-based computational applications that effectively exploit exascale HPC technologies to provide breakthrough modeling and simulation and data science solutions
- Exascale applications (and their companion co-designed computational motifs) are a foundational element of the ECP and are *the* vehicle for delivery of consequential solutions and insight from exascale systems
- Each ECP application is focused on targeted development to address a unique mission *challenge problem*, or one that
  - Possesses solution amenable to simulation insight
  - Represents a strategic problem important to a DOE mission program
  - Is currently intractable without the computational power of exascale.

*https://en.wikipedia.org/wiki/Skin_in_the_game_(phrase)

Exascale applications target US national problems in 6 strategic areas

National security

Stockpile stewardship

Next-generation electromagnetics simulation of hostile environment and virtual flight testing for hypersonic re-entry vehicles

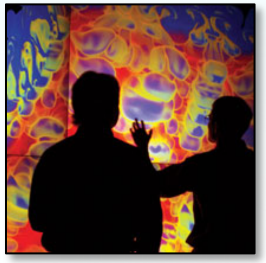

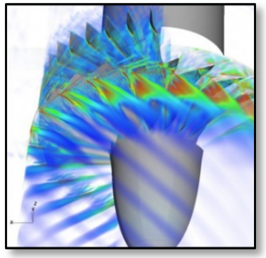

Energy security

Turbine wind plant efficiency

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

Materials design for extreme environments of nuclear fission and fusion reactors

Design and commercialization of Small Modular Reactors

Subsurface use for carbon capture, petroleum extraction, waste disposal

Scale-up of clean fossil fuel combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Reliable and efficient planning of the power grid

Seismic hazard risk assessment

Urban planning

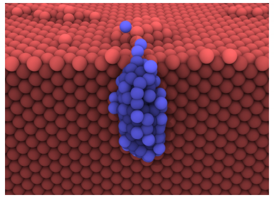

Scientific discovery

Find, predict, and control materials and properties

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Demystify origin of chemical elements

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

Whole-device model of magnetically confined fusion plasmas

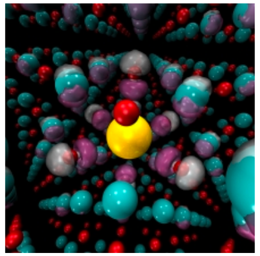

Earth system

Accurate regional impact assessments in Earth system models

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

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

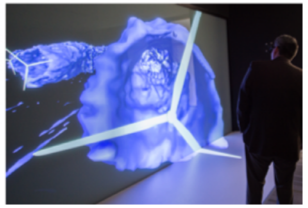

Health care

Accelerate and translate cancer research

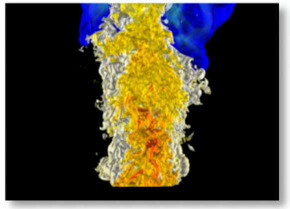

ECP EXASCALE COMPUTING PROJECT

Department of Energy (DOE) Roadmap to Exascale Systems

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

Pre-Exascale Systems [Aggregate Linpack (Rmax) = 323 PF!]

First U.S. Exascale Systems

2012	2016	2018	2020	2021-2023

ORNL
Cray/AMD/NVIDIA

ORNL
IBM/NVIDIA

FRONTIER

ORNL
TBD

ANL
IBM BG/Q

ANL
Cray/Intel KNL

ANL
Intel/Cray

LBNL
Cray/Intel Xeon/KNL

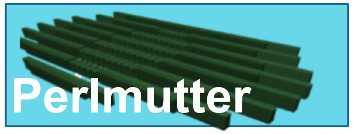

LBNL
Cray/AMD/NVIDIA

LLNL
IBM BG/Q

LANL/SNL
Cray/Intel Xeon/KNL

LLNL
IBM/NVIDIA

CROSSROADS

LANL/SNL
TBD

LLNL
TBD

ECP EXASCALE COMPUTING PROJECT

DOE Exascale Program: The Exascale Computing Initiative (ECI)

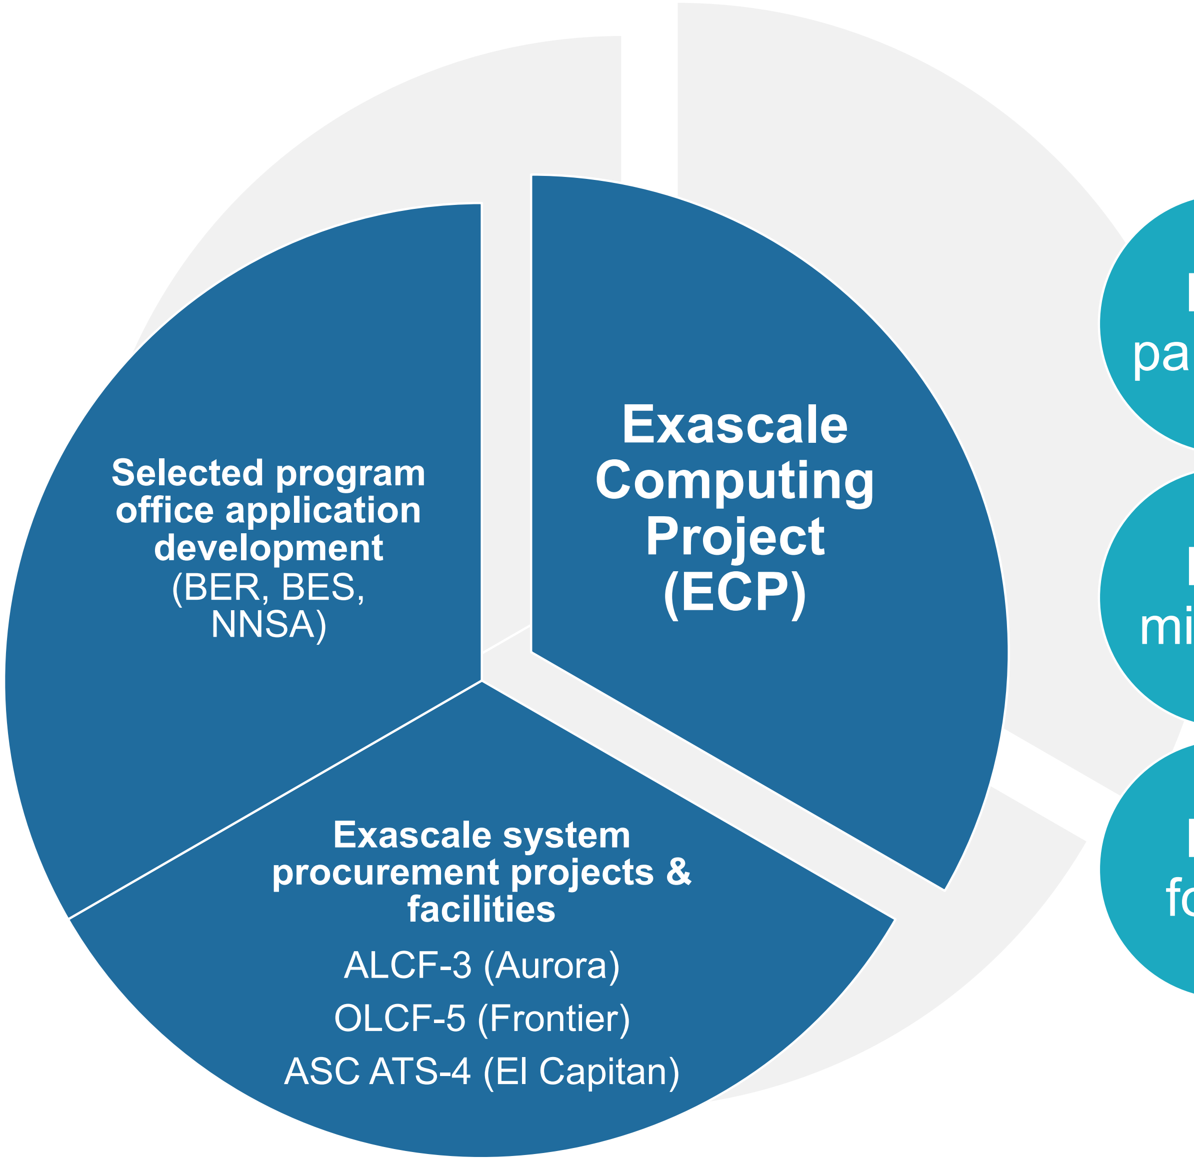

Three Major Components of the ECI

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

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

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

EXASCALE COMPUTING PROJECT

ECP will enable future US revolutions in technology development, scientific discovery, energy and economic security, and healthcare

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 vision

Accelerate innovation with exascale simulation and data science solutions that enhance US economic competitiveness, improve our quality of life, and strengthen our national security

ECP's Application Development Effort: Structured Around 6 Thrusts

Prepare key applications for exascale, execute challenge problems, measure performance

Chemistry and Materials	Seeks to describe underlying properties of matter needed to optimize and control the design of new materials and energy technologies
Energy	Advances modeling and simulation of existing and future technologies for the efficient and responsible production of energy to meet the growing needs of the U.S.
Earth and Space Science	Addresses fundamental scientific questions from the origin of the universe and chemical elements to planetary processes and interactions affecting life and longevity
Data Analytics and Optimization	Employs modern data analysis and machine learning techniques rather than strictly on approximate solutions to equations that state fundamental physical principles or reduced semi-empirical models
National Security	Stewardship of the US nuclear stockpile and assessment of future threats; related physics and engineering modeling and scientific inquiries consistent with that mission space
Co-Design	Focuses on crosscutting algorithmic methods that capture the most common patterns of computation and communication in ECP applications

ECP Application Development Leadership Team

Lead: Andrew Siegel (ANL) and Deputy Lead: Erik Draeger (LLNL)

Jack Deslippe, Chemistry and Materials Applications
Jack is the acting group lead of the Application Performance Group NERSC and leads the NERSC Exascale Science Applications Program (NESAP). Jack has a history of support and development of HPC apps in Materials Science and Ph.D. in Condensed Matter Physics from UC Berkeley.

Tom Evans, Energy Applications
Tom is a Distinguished R&D Staff member and Team Leader of the HPC Methods and Applications Team in Reactor and Nuclear Systems Division at ORNL. He over 20 years of experience in the development of single- and coupled-physics development on HPC hardware spanning Office of Science and NNSA applications.

Anshu Dubey, Earth and Space Science Applications
Anshu is a Computer Scientist in the Mathematics and Computer Science Division at ANL. She has a long history of working with all aspects of PDE based scientific software.

Bill Hart, Data Analytics and Optimization Applications
Bill is a senior research in the Center for Computing Research (CCR) at SNL. As a researcher and former manager in CCR, Bill has led and championed research in scalable optimization, data science and cybersecurity. Bill has a Ph.D. in Computer Science from UC San Diego.

Marianne Francois, National Security Applications
Marianne is the group leader of the Methods and Algorithms Group in the Computational Physics Division in the Weapons Physics Directorate at LANL. She has a history of developing methods for complex multi-material flow and Ph.D. in Aerospace Engineering from University of Florida.

Phil Colella, Co-Design
Phil is Senior Scientist at in the Computational Research Division at LBNL, and a Professor in Residence in the EECS Department at UC Berkeley. His research has been in the area of numerical methods and HPC software development for partial differential equations.

Application Co-Design (CD)

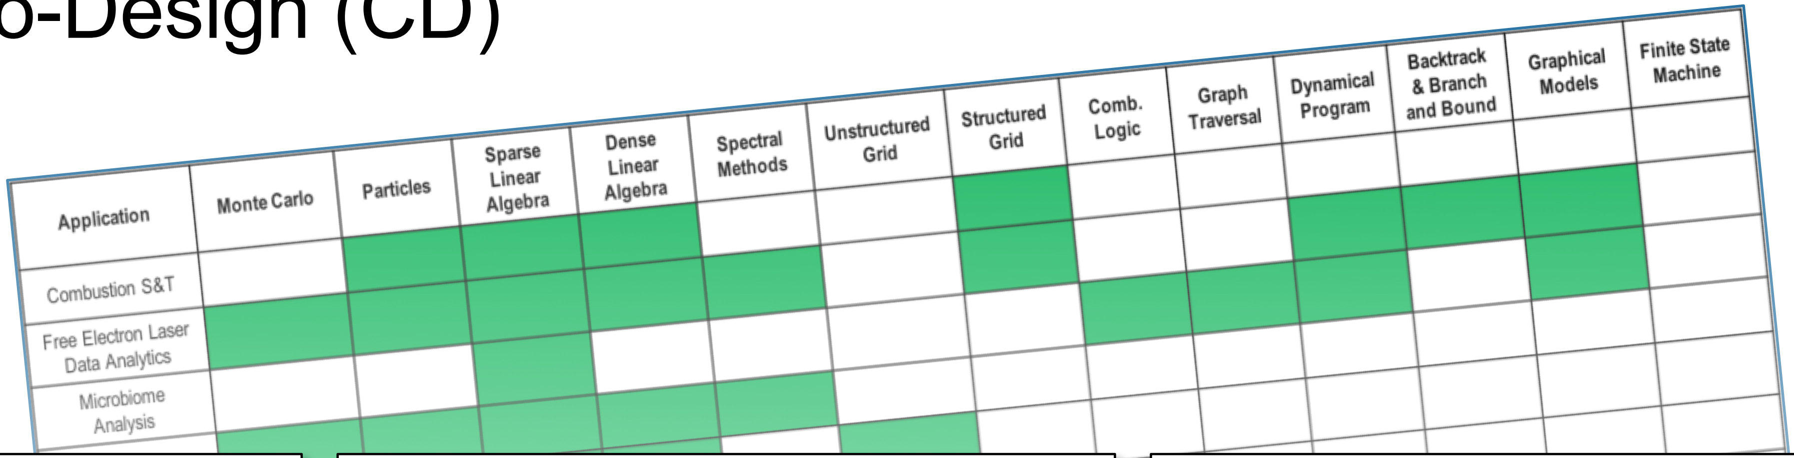

Essential to **ensure that applications effectively utilize** exascale systems

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

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

- Motif: algorithmic method that drives a common pattern of computation and communication
- CD Centers must address all high priority motifs invoked by ECP applications, including not only the 7 “classical” motifs but also the additional 6 motifs identified to be associated with data science applications

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

- Evaluate, deploy, and integrate exascale hardware-savvy software designs and technologies for key crosscutting algorithmic motifs into applications
- An appropriate nexus for reduced precision?

EXASCALE COMPUTING PROJECT

Application Co-Design

Develop efficient exascale libraries that address computational motifs common to multiple application projects

CODAR	Advance understanding of the constraints, mappings, and configuration choices that determine interactions of applications, data analysis and reduction, and exascale platforms
COPA	Create co-designed numerical recipes for particle-based methods that meet application team requirements within design space of STs and subject to constraints of exascale platforms
AMReX	Build framework to support development of block-structured adaptive mesh refinement algorithms for solving systems of partial differential equations on exascale architectures
CEED	Develop next-generation discretization software and algorithms that will enable a wide range of finite element applications to run efficiently on future hardware
ExaGraph	Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms
ExaLearn	Target learning methods to aid application and experimental facility workflows: deep neural networks (RNNs, CNNs, GANs), kernel & tensor methods, decision trees, ensemble methods, graph models, reinforcement learning
Proxy Apps	Improve the quality of proxy applications created by ECP and maximize the benefit received from their use. Maintain and distribute ECP Proxy App Suite.

Center for Efficient Exascale Discretizations (CEED)

Co-Design of unstructured mesh, FE-based PDE discretizations

CEED EXASCALE DISCRETIZATIONS

Goal

- Develop algorithms and software to enable more efficient HPC simulations in a wide range of PDE-based science applications.
- Focus on next-generation discretization methods: high-order finite elements on general unstructured grids.
- Target high performance on a variety of hardware: CPU, GPU, A21 in a flexible and user-friendly way.

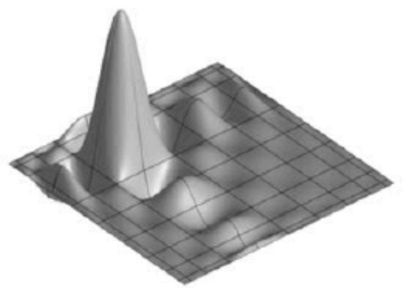

$$A = P^T G^T B^T D B G P$$

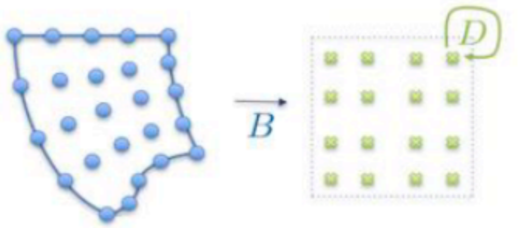

Approach

- Performance-enabling math foundation: high-order operator decomposition
- Fast kernels: CEED benchmarks, combine expertise, engage community
- Library integration: high-level API (MFEM, Nek5000), low-level API (libCEED)
- Application engagement: liaisons, CEED miniapps (Nekbone, Laghos)
- Collaborate with ECP/ST, broader community (SciDAC, xSDK, deal.ii, …)
- High-order software ecosystem: operator format, FMS, matrix-free solvers

```
CeedElemRestrictionCreate(ceed, nelem, 2, Nx, CEED_MEM_HOST,
                          CEED_USE_POINTER, indx, &Erestrictx);
CeedElemRestrictionCreate(ceed, nelem, P, Nu, CEED_MEM_HOST,
                          CEED_USE_POINTER, indu, &Erestrictu);

CeedBasisCreateTensorH1Lagrange(ceed, 1, 1, 2, Q, CEED_GAUSS, &bx);
CeedBasisCreateTensorH1Lagrange(ceed, 1, 1, P, Q, CEED_GAUSS, &bu);

CeedQFunctionCreateInterior(ceed, 1, 1, sizeof(CeedScalar),
                            CEED_EVAL_GRAD|CEED_EVAL_WEIGHT),
                            CEED_EVAL_NONE, setup, __FILE__ ":setup",
                            &qf_setup);
CeedQFunctionCreateInterior(ceed, 1, 1, sizeof(CeedScalar),
                            CEED_EVAL_INTERP, CEED_EVAL_INTERP,
                            mass, __FILE__ ":mass", &qf_mass);

CeedOperatorCreate(ceed, Erestrictx, bx, qf_setup, NULL, NULL, &op_setup);
CeedOperatorCreate(ceed, Erestrictu, bu, qf_mass, NULL, NULL, &op_mass);
```

CEED is targeting several ECP applications

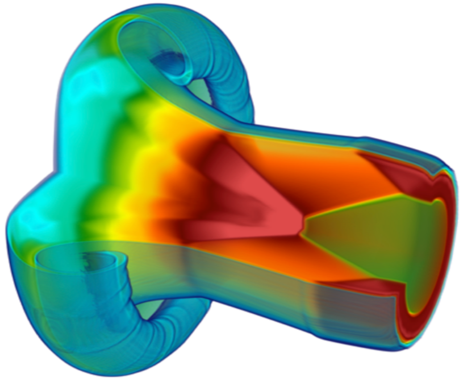
Compressible flow (MARBL)

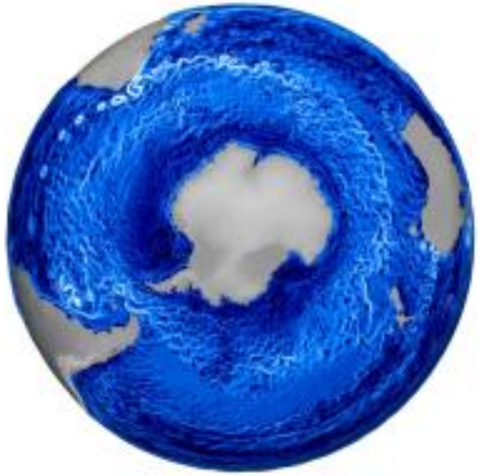
Climate (E3SM)

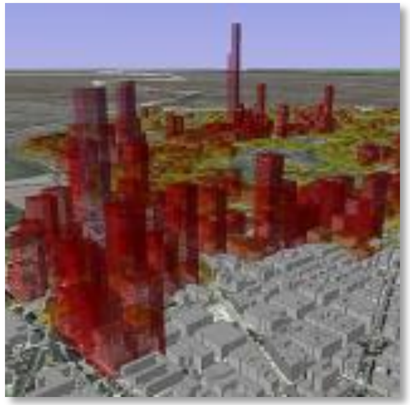
Urban systems (Urban)

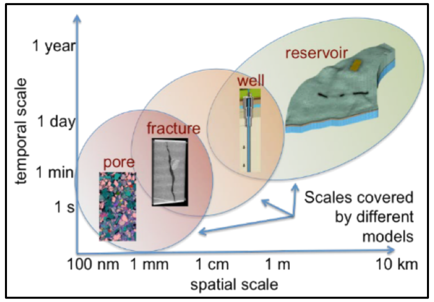

Subsurface (GEOS)

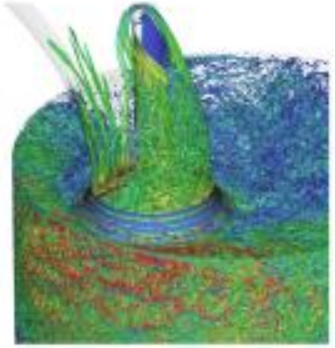
Combustion (Nek5000)

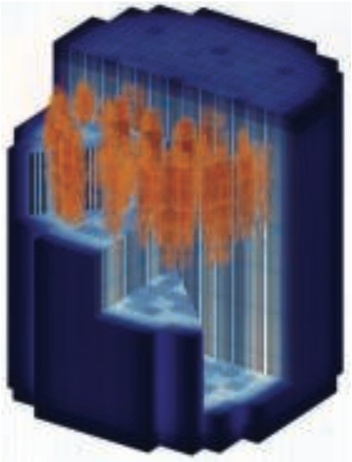
Modular Nuclear Reactors (ExaSMR)

Wind Energy (ExaWind)


Additive Manufacturing (ExaAM)

Magnetic Fusion (WDMApp)
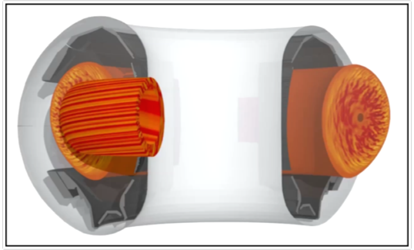

EXASCALE COMPUTING PROJECT

PI: Tzanio Kolev (LLNL)

ECP's Adaptive Mesh Refinement Co-Design Center: AMReX

- Develop and deploy software to support block-structured adaptive mesh refinement on exascale architectures
  - Core AMR functionality
  - Particles coupled to AMR meshes
  - Embedded boundary (EB) representation of complex geometry
  - Linear solvers
  - Supports two modalities of use
    - Library support for AMR
    - Framework for constructing AMR applications

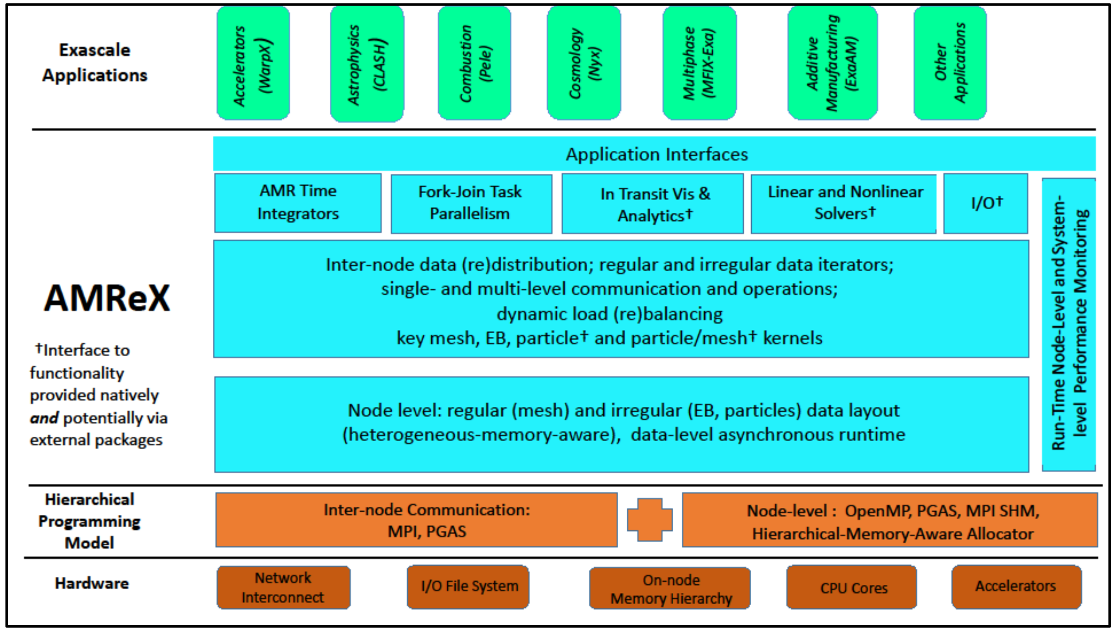

- Provide direct support to ECP applications that need AMR for their application
- Evaluate software technologies and integrate with AMReXwhen appropriate
- Interact with hardware technologies / vendors

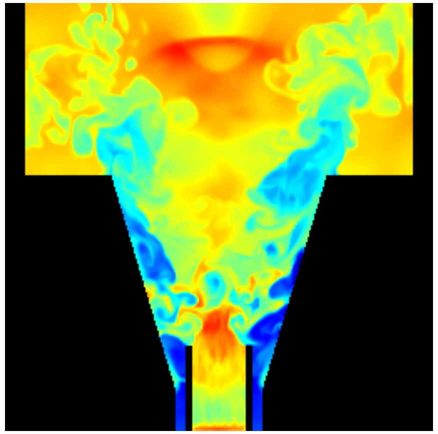

Application	Particles	ODEs	Linear Solvers	EB
Combustion	X	X	X	X
Multiphase	X		X	X
Cosmology	X	X	X	
Astrophysics	X	X	X	
Accelerators	X			

PI: John Bell (LBNL)

Why AMReX for Exascale?

- Data: Hierarchy of uniform mesh ”levels”, composed of unions of rectangular grids, 8-64 cells each dimension
- Accurate, well-characterized PDE discretizations, well-suited for multigrid-based linear system solvers
- Global metadata simplifies communication
- Hierarchical data & work map well to exascale hardware
  - Data distributed by grid, work distributed by **logical tile**
  - Implemented via local data iterators w/runtime controls, supports static and dynamic scheduling, avoids overheads (execution, programming efforts) of loop-based directives
  - Low-level kernels implemented in any language (i.e. Fortran: ideal “DSL” for multi-dimensional arrays)
  - Metadata-aware runtime (overlapping communication/computation) – Perilla
  - Flexible approach supports multiple on-node distribution strategies (OpenMP, CUDA, OpenACC, …)
    - Avoids need to reimplement complex kernels in limited languages supported by ”modern” on-node tools
    - However, does not preclude isolated use of alternative strategies, as appropriate

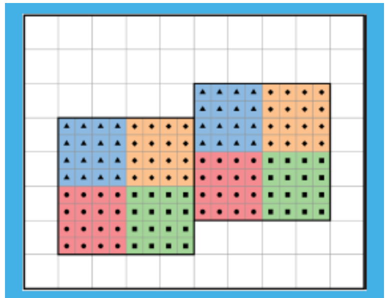

2-level tiled AMR hierarchy

ECP's Co-Design Center for Online Data Analysis and Reduction

CODAR

Traditional approach:
Compute…output…analyze [offline]

Write simulation output to secondary storage; read back for analysis

Decimate in time when simulation output rate exceeds output rate of computer

New approach:
Online data analysis and reduction

Co-optimize simulation, analysis, reduction for performance and information output

Substitute CPU cycles for I/O, via online data (de)compression and/or online data analysis

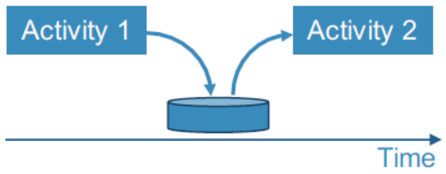

Provide the right information at the right time and place to accelerate discovery!

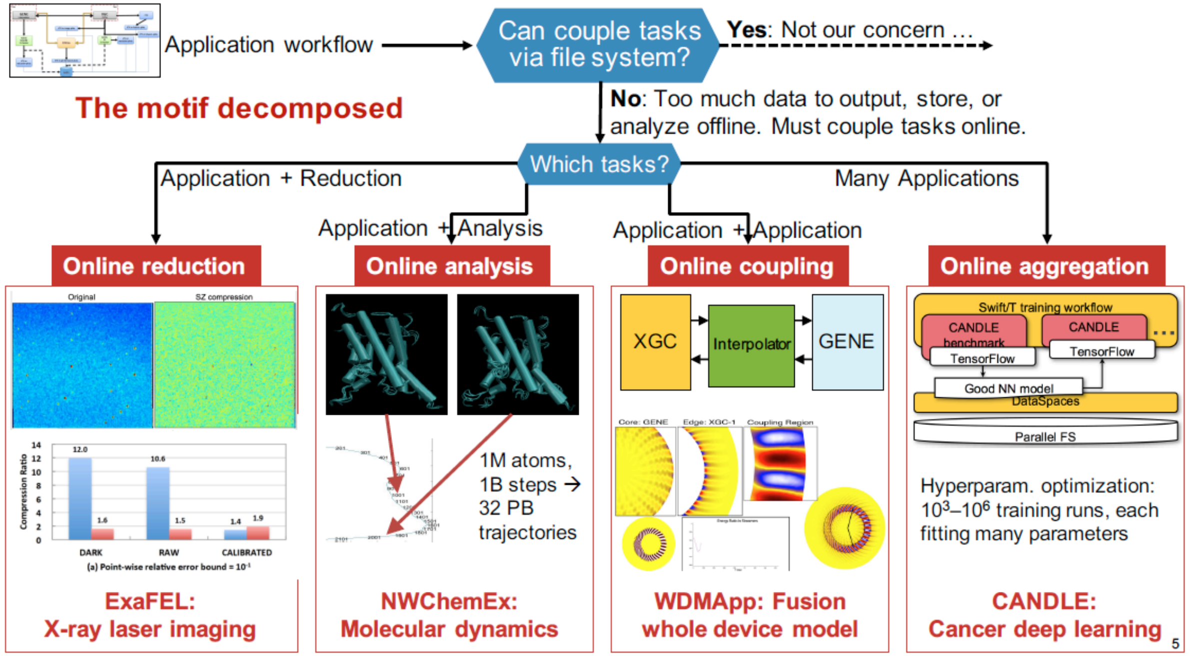

<u>Goal:</u> Replace the activities in HPC workflow that have been mediated through file I/O with in-situ methods / workflows. data reduction, analysis, code coupling, aggregation (e.g. parameter studies).

<u>Cross-cutting tools:</u>

- Workflow setup, manager (Cheetah, Savanna); Data coupler (ADIOS-SST); Compression methods (MGARD, FTK, SZ), compression checker (Z-checker)
- Performance tools (TAU, Chimbuco, SOSFlow)

Example: WDMApp Workflow (fusion application)

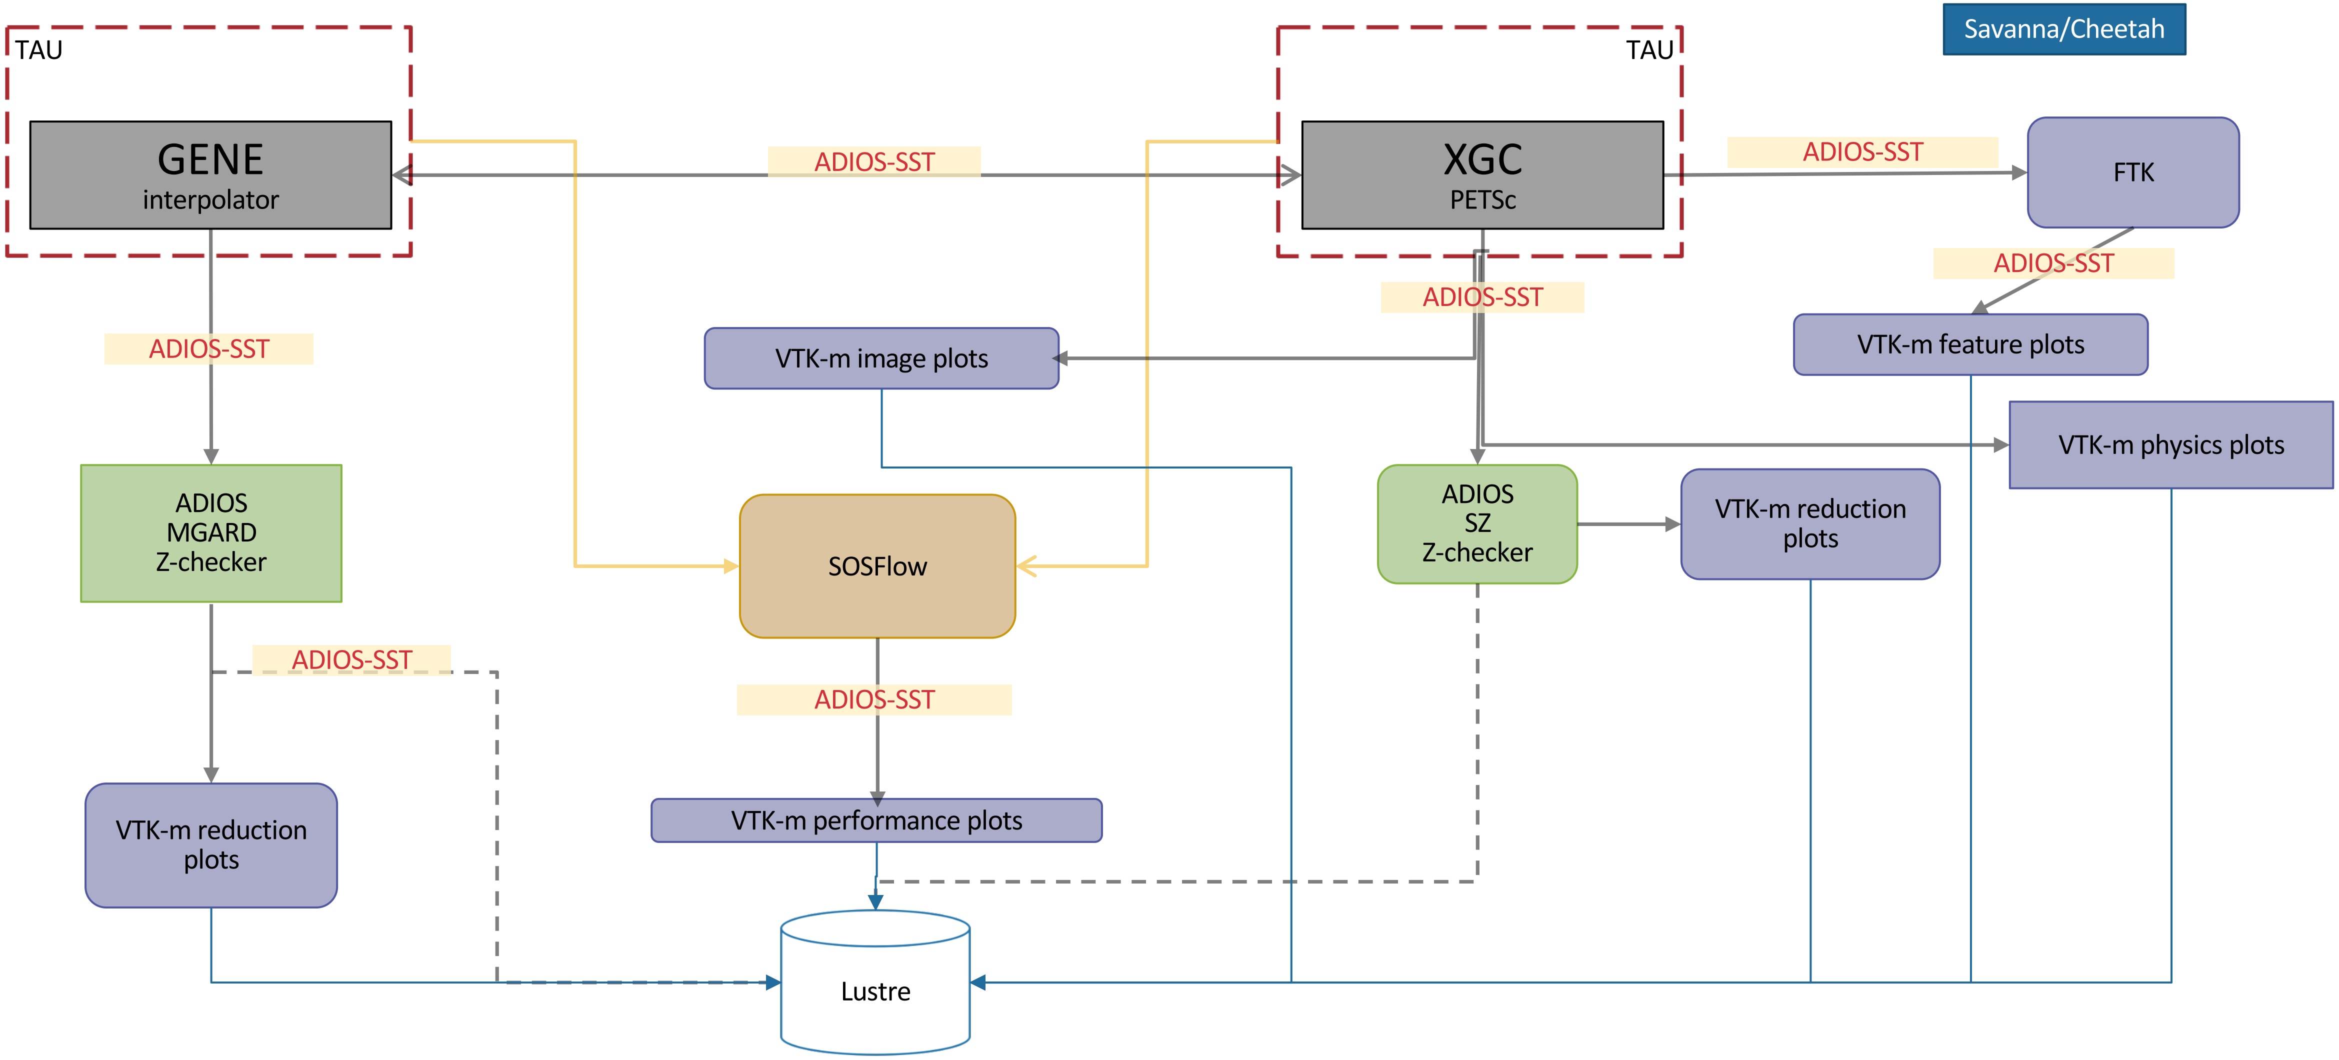

PI: Ian Foster (ANL)

ECP's Co-Design Center for Machine Learning: ExaLearn

Bringing together experts from 8 DOE Laboratories

- AI has the potential to accelerate scientific discovery or enable prediction in areas currently too complex for direct simulation (ML for HPC and HPC for ML)
- AI use cases of interest to ECP:
  - *Classification and regression*, including but not limited to image classification and analysis, e.g. scientific data output from DOE experimental facilities or from national security programs.
  - *Surrogate models* in high-fidelity and multiscale simulations, including uncertainty quantification and error estimation.
  - *Structure-to-function relationships*, including genome-to-phenome, the prediction of materials performance based on atomistic structures, or the prediction of performance margins based on manufacturing data.
  - *Control systems*, e.g., for wind plants, nuclear power plants, experimental steering and autonomous vehicles.
  - *Inverse problems* and optimization. This area would include, for example, inverse imaging and materials design.
- Areas in need of research
  - Data quality and statistics
  - Learning algorithms
  - Physics-Informed AI
  - Verification and Validation
  - Performance and scalability
  - Workflow and deployment

Expected Work Product: A Toolset That . . .

- Has a line-of-sight to exascale computing, e.g. through using exascale platforms directly, or providing essential components to an exascale workflow
- Does not replicate capabilities easily obtainable from existing, widely-available packages
- Builds in domain knowledge where possible "Physics"-based ML and AI
- Quantifies uncertainty in predictive capacity
- Is interpretable
- Is reproducible
- Tracks provenance

EXASCALE COMPUTING PROJECT

PI: Frank Alexander (BNL)

Machine Learning in the Light Source Workflow

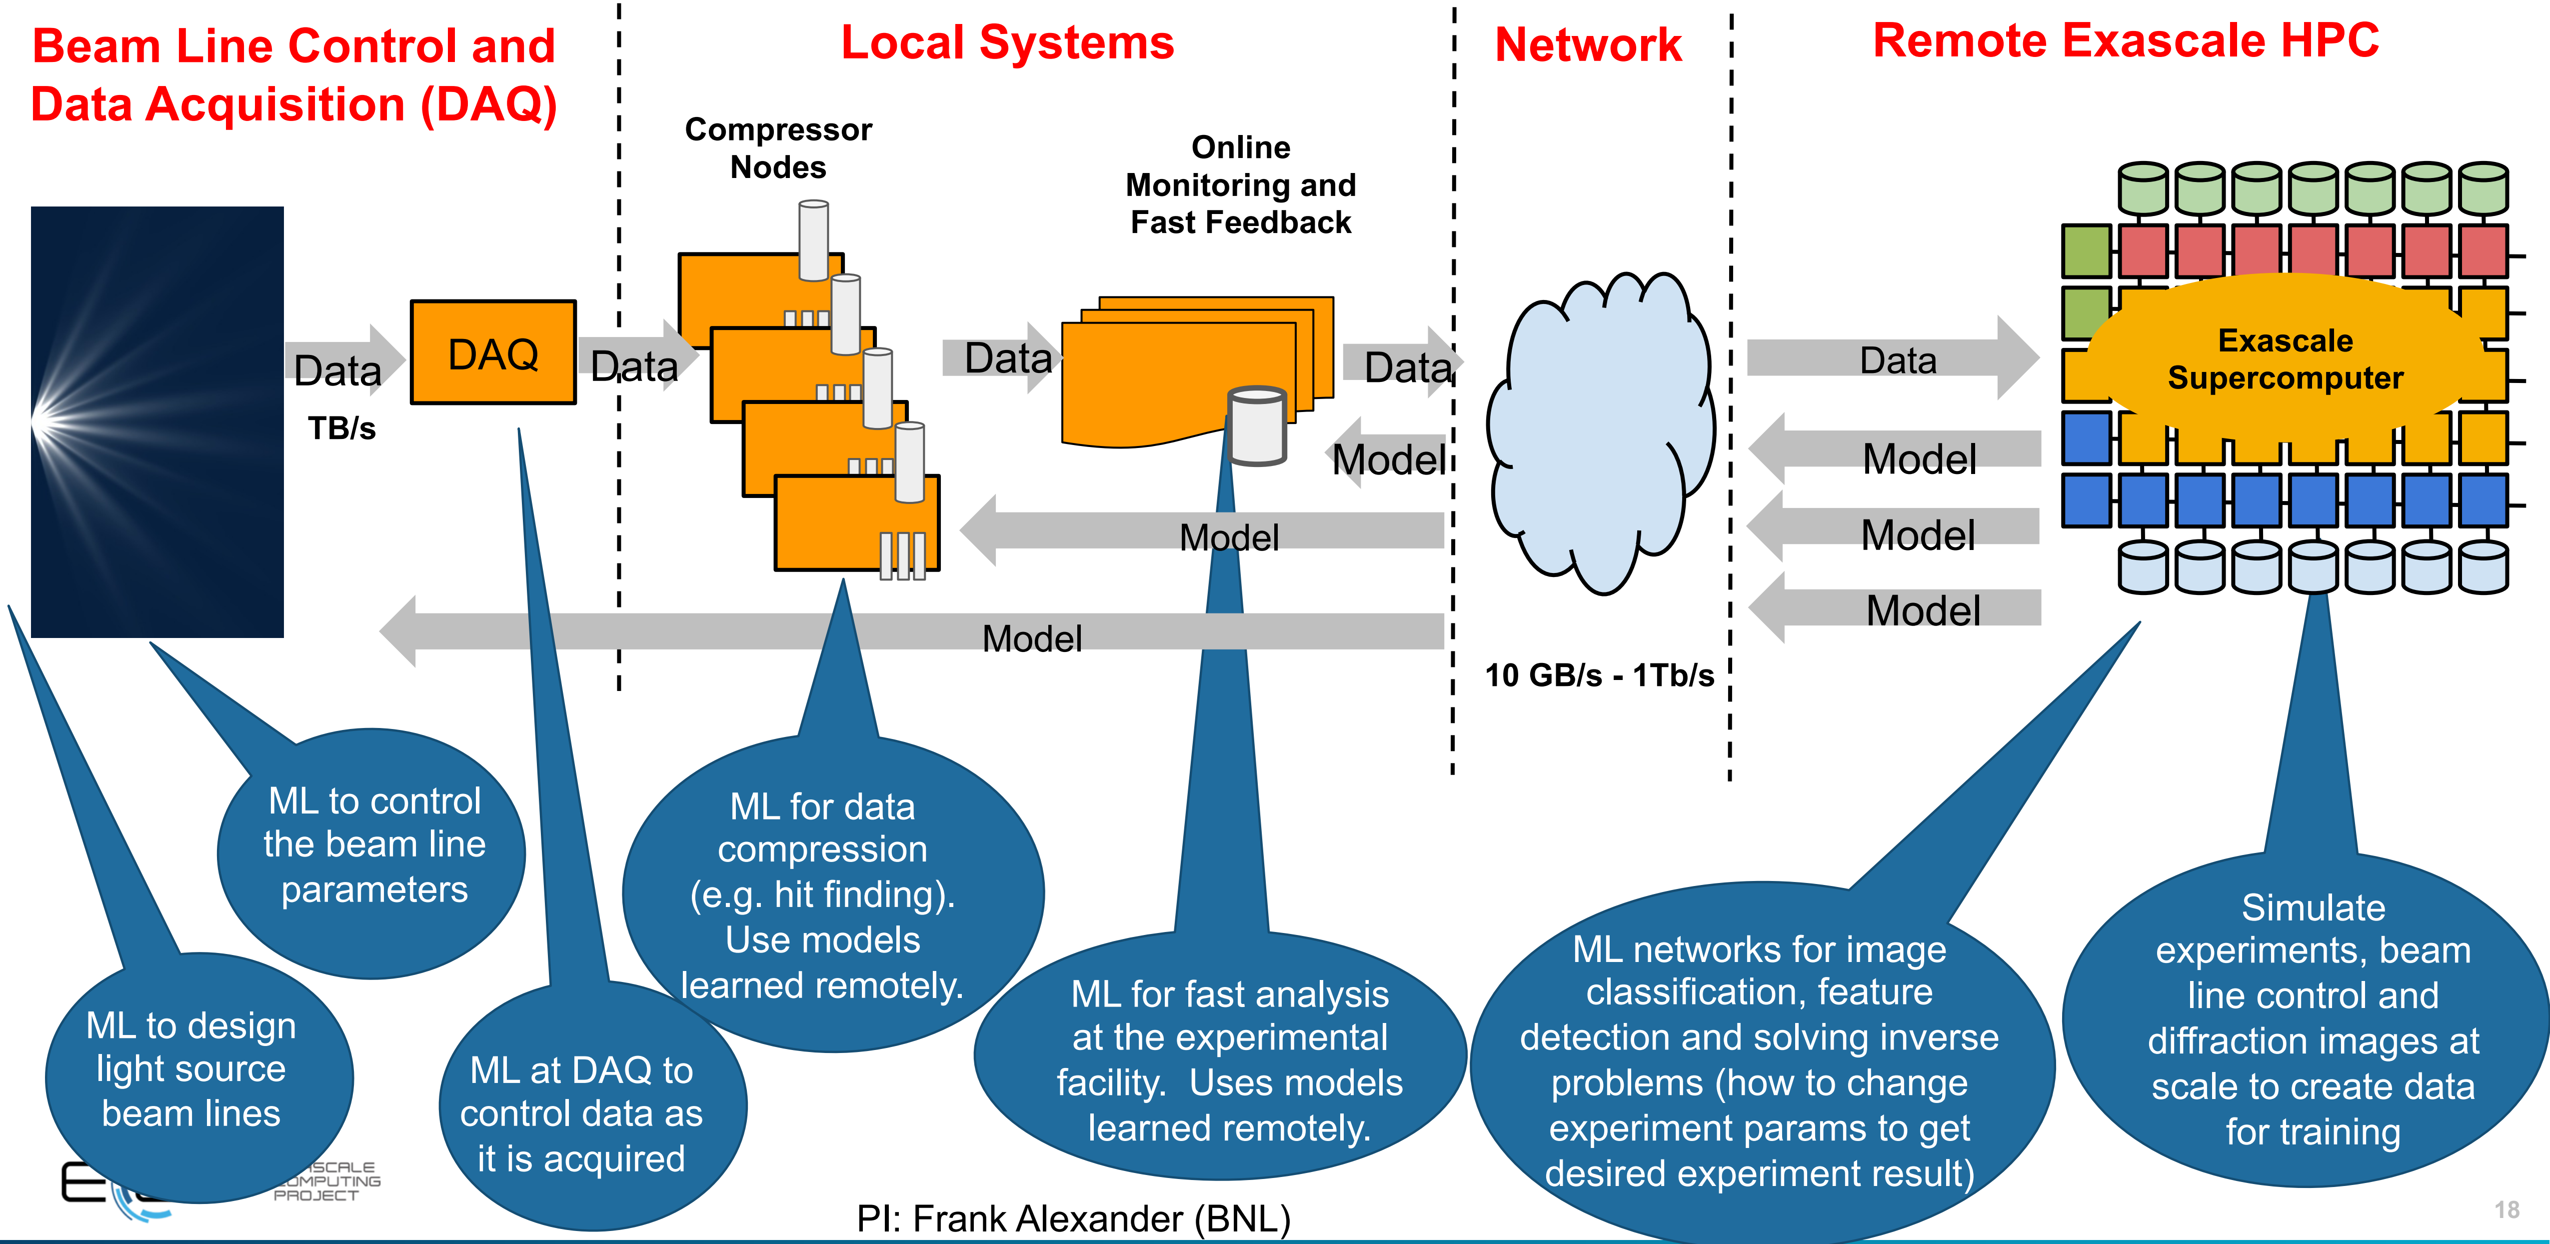

PI: Frank Alexander (BNL)

ECP's Co-Design Center for Particle Applications: CoPA

Goal: Develop algorithms and software for particle methods,

Cross-cutting capabilities:

- Specialized solvers for quantum molecular dynamics (Progress / BML).
- Performance-portable libraries for classical particle methods in MD, PDE (Cabana).
- FFT-based Poisson solvers for long-range forces.

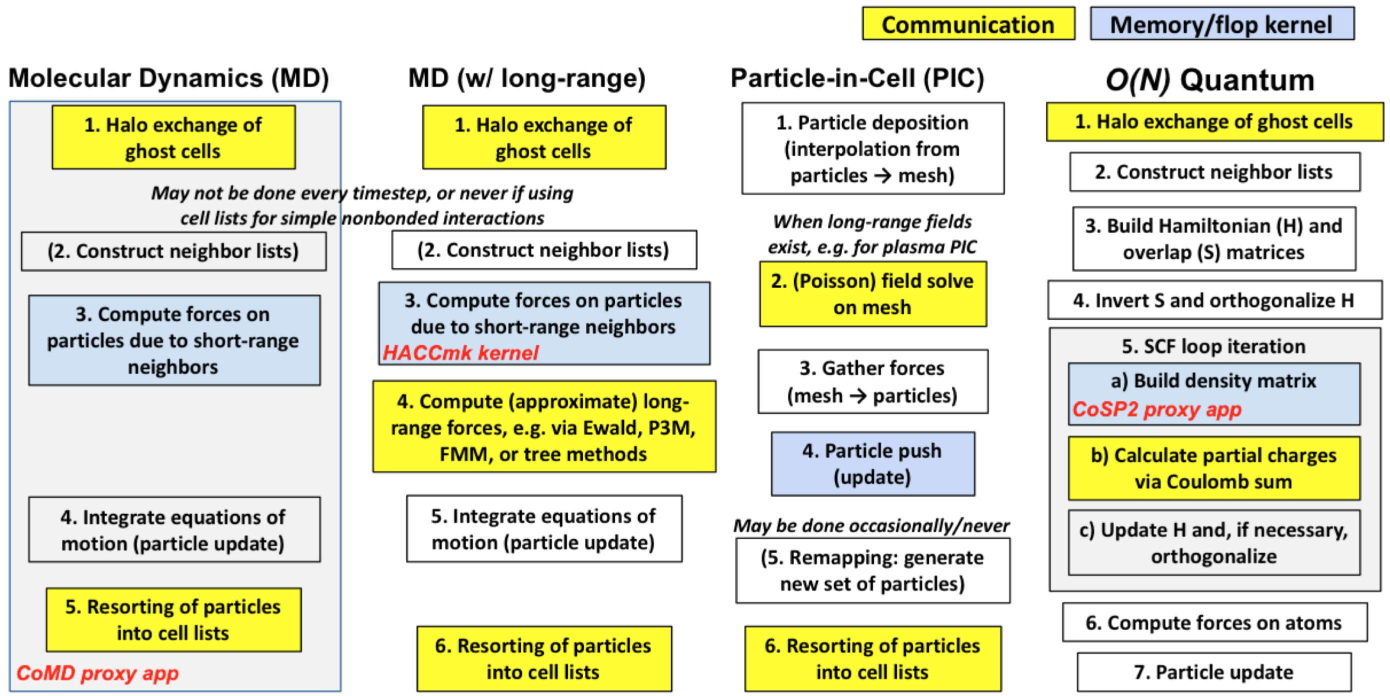

Technical approach:

- High-level C++ APIs, plus a Fortran interface (Cabana).
- Leverage existing / planned FFT software.
- Extensive use of miniapps / proxy apps as part of the development process.

EXASCALE COMPUTING PROJECT

PI: Tim Germann (LANL)

Exascale apps can deliver transformative products and solutions

ExaWind

Turbine Wind Plant Efficiency
(Mike Sprague, NREL)

- Harden wind plant design and layout against energy loss susceptibility
- Increase penetration of wind energy

Challenges: linear solver perf in strong scale limit; manipulation of large meshes; overset of structured & unstructured grids; communication-avoiding linear solvers

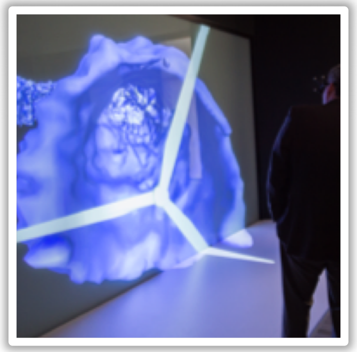

ExaAM

Additive Manufacturing (AM) of Qualifiable Metal Parts
(John Turner, ORNL)

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

Challenges: capturing unresolved physics; multi-grid linear solver performance; coupled physics

EQSIM

Earthquake Hazard Risk Assessment
(David McCallen, LBNL)

- Replace conservative and costly earthquake retrofits with safe purpose-fit retrofits and designs

Challenges: full waveform inversion algorithms

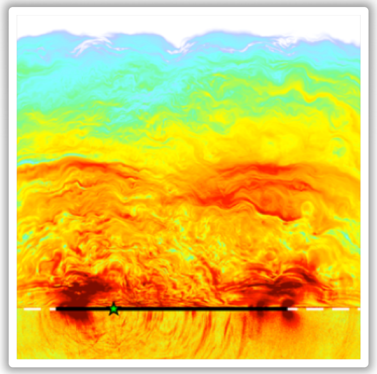

EQSIM: Understanding and predicting earthquake phenomenon

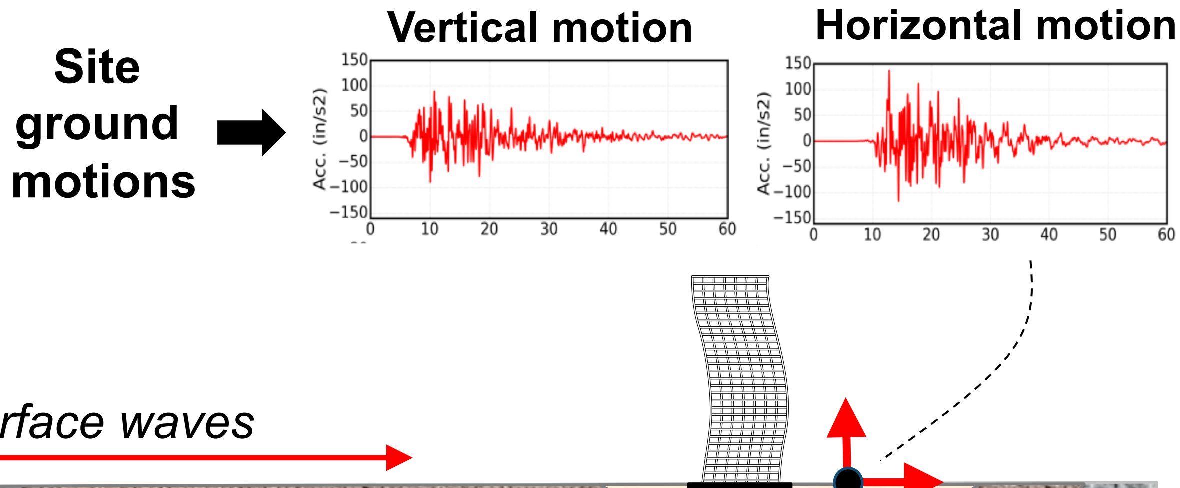

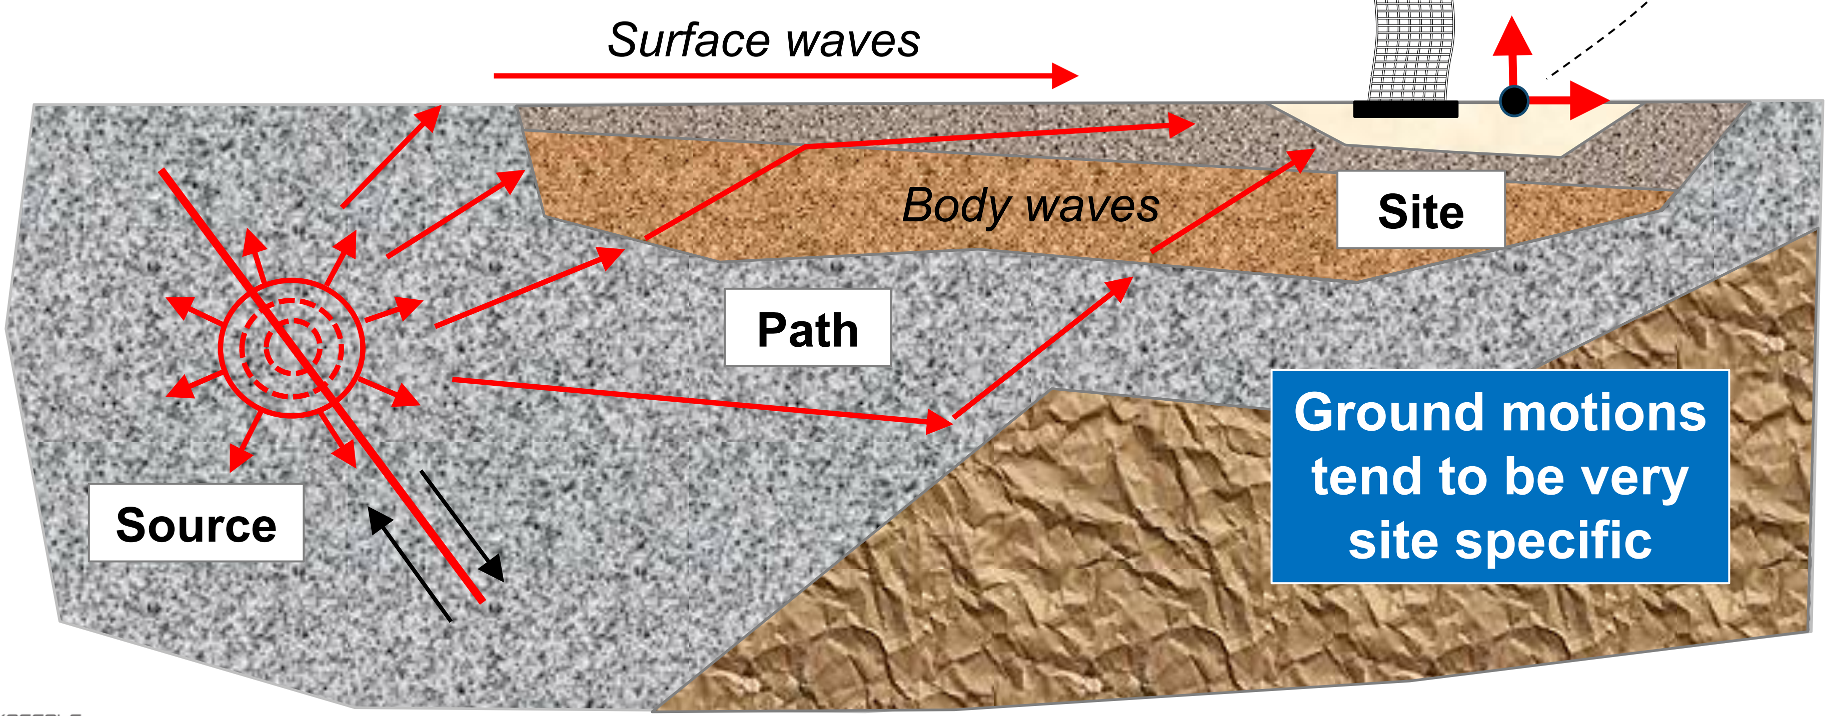

PI: David McCallen (LBNL)

EQSIM: The Exascale "Big Lift" – regional ground motion simulations at engineering frequencies

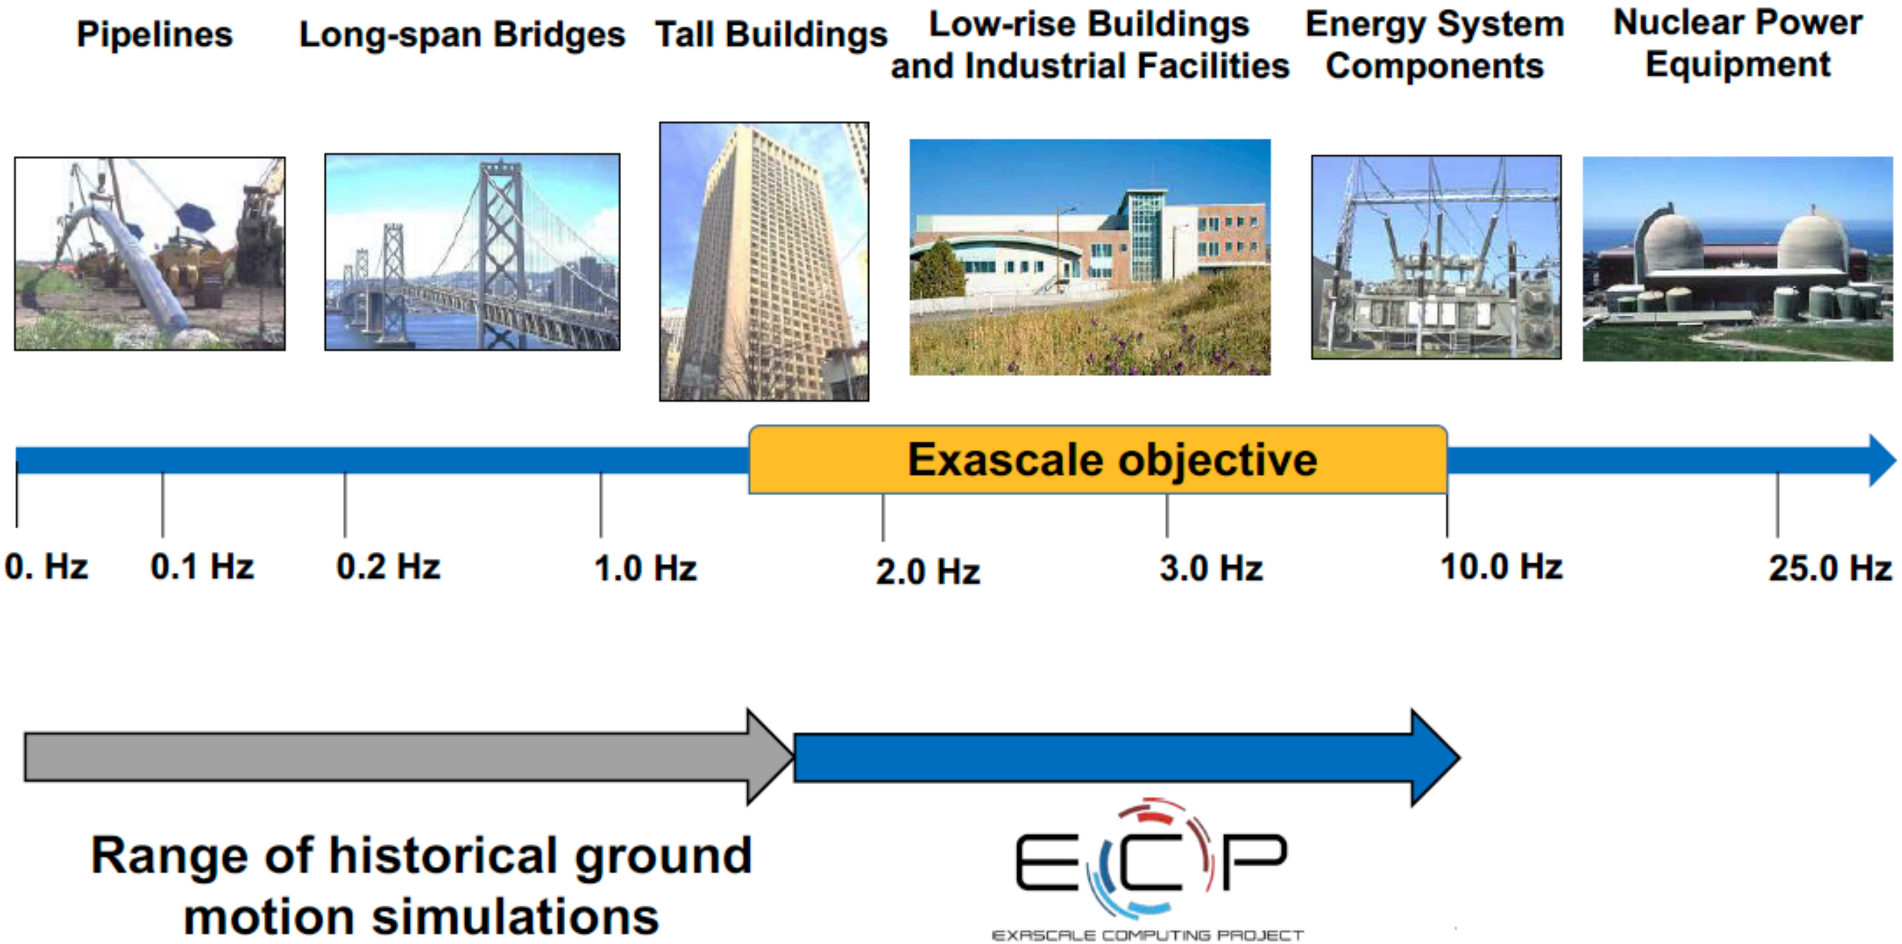

Doubling the frequency resolution = 16X computational effort!

PI: David McCallen (LBNL)

EQSIM: Advancing geophysics and infrastructure applications

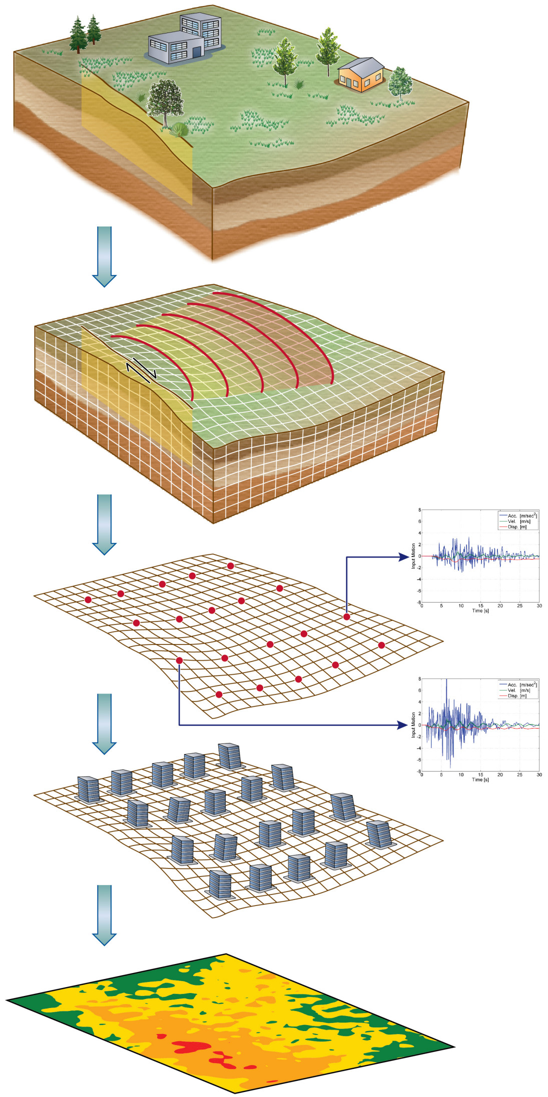

Earthquake Hazard

Earthquake Risk

***SW4* – 4^{th} order finite difference geophysics code for wave propagation**

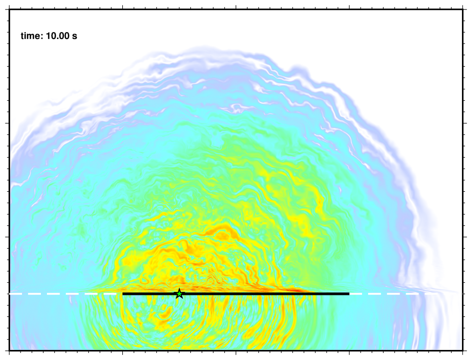

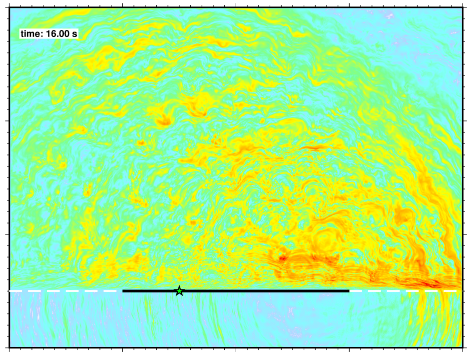

***NEVADA & MSESSI* – finite deformation, inelastic Finite Element codes for structures and soils**

Weak Coupling

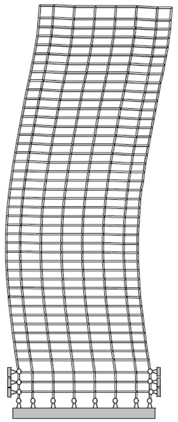

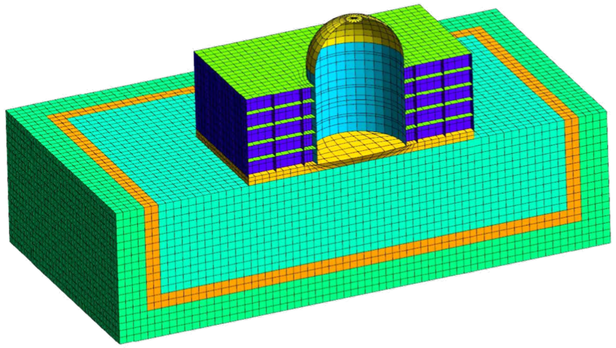

Strong Coupling

PI: David McCallen (LBNL)

EQSIM: Using RAJA to achieve performance portability

- RAJA is a C++ abstraction layer developed at LLNL.
- Same C++ source code for OpenMP and CUDA backends
  - Machine specific options in a policy file
- Coding complexity similar to OpenMP
- Currently running on the Sierra GPU machine at LLNL with low overhead
- August-2018: 1,024 nodes of Sierra, 4,096 GPUs, 6.9 Hz, giving an overall performance (Figure or Merit) improvement of 24.2

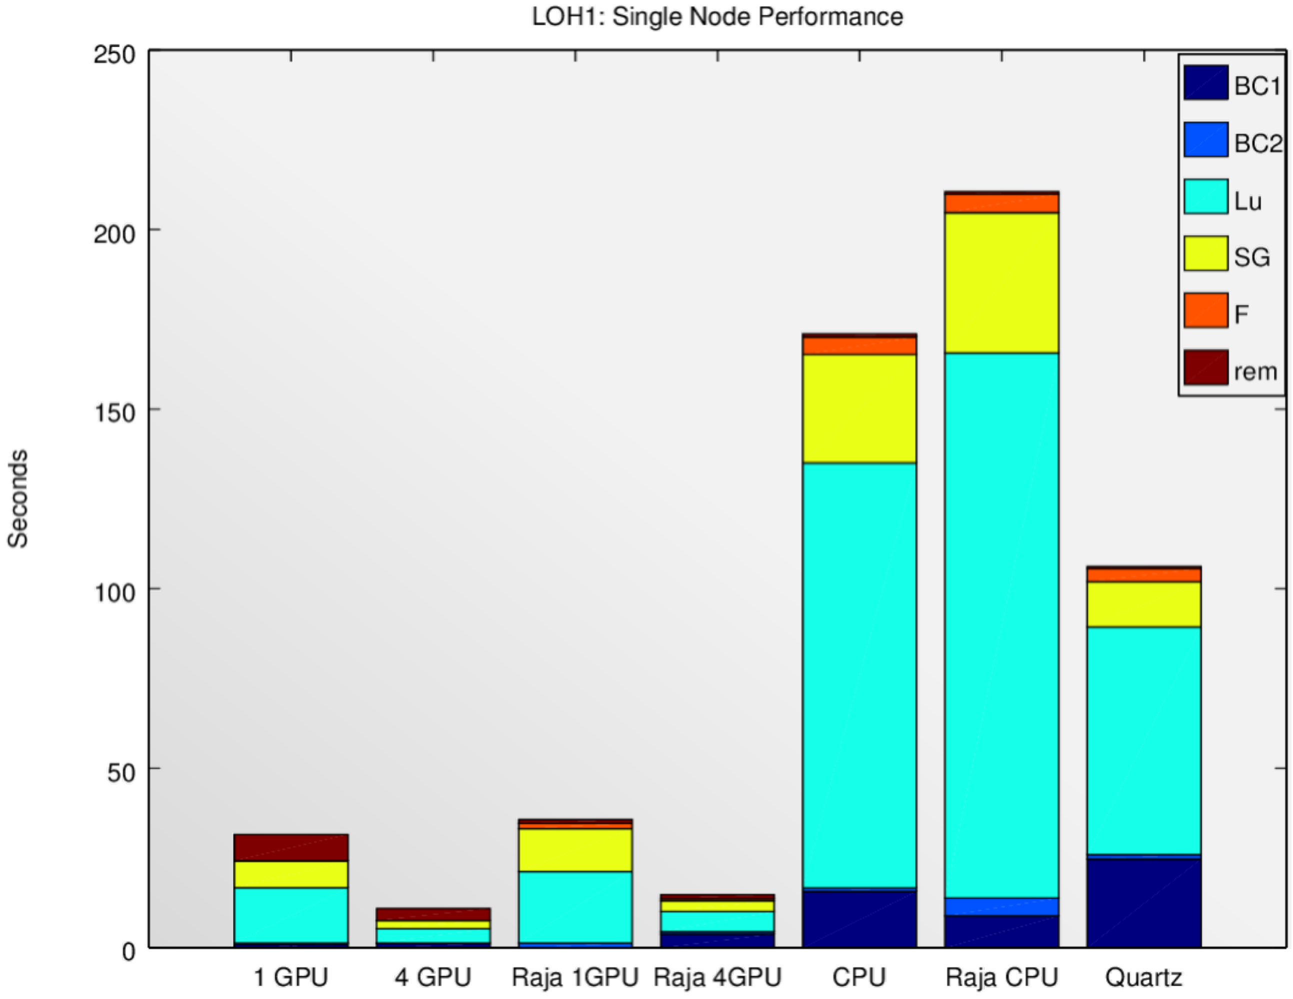

PI: David McCallen (LBNL)

Exascale apps can deliver transformative products and solutions

MFIX-Exa

Scale-up of Clean Fossil Fuel Combustion
(Madhava Syamlal, NETL)

- Commercial-scale demonstration of transformational energy technologies – curbing CO_2 emissions at fossil fuel power plants by 2030

Challenges: load balancing; strong scaling thru transients

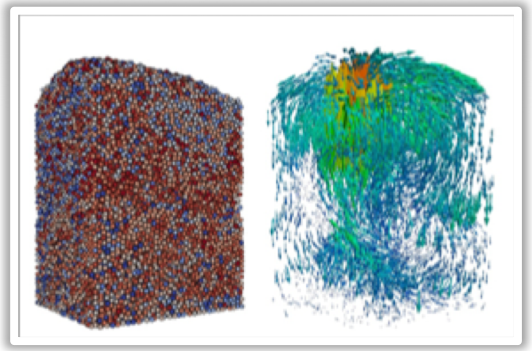

GAMESS

Biofuel Catalyst Design
(Mark Gordon, Ames)

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

Challenges: weak scaling of overall problem; on-node performance of molecular fragments

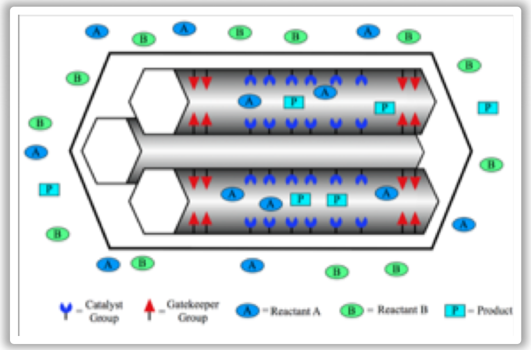

EXAALT

Materials for Extreme Environments
(Art Voter, LANL)

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

Challenges: SNAP kernel efficiency on accelerators; efficiency of DFTB application on accelerators

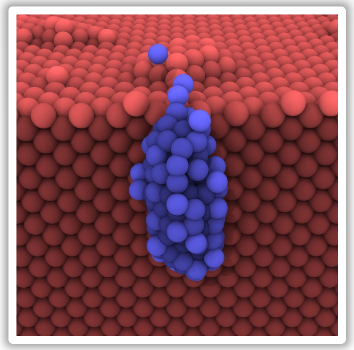

EXASCALE COMPUTING PROJECT

Exascale apps can deliver transformative products and solutions

ExaSMR

Design and Commercialization of Small Modular Reactors
(Steve Hamilton, ORNL)

- Virtual test reactor for advanced designs via experimental-quality simulations of reactor behavior

Challenges: existing GPU-based MC algorithms require rework for hardware less performant for latency-bound algorithms with thread divergence; performance portability with OCCA & OpenACC not achievable; insufficient node memory for adequate CFD + MC coupling

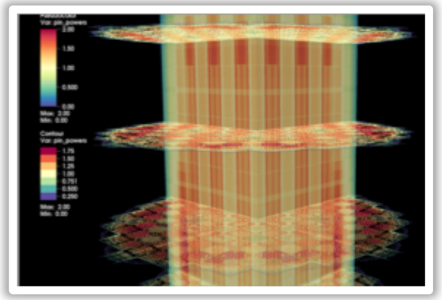

Subsurface

Carbon Capture, Fossil Fuel Extraction, Waste Disposal
(Carl Steefel, LBNL)

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

Challenges: performance of Lagrangian geomechanics; adequacy of Lagrangian crack mechanics) + Eulerian (reaction, advection, diffusion) models; parallel HDF5 for coupling

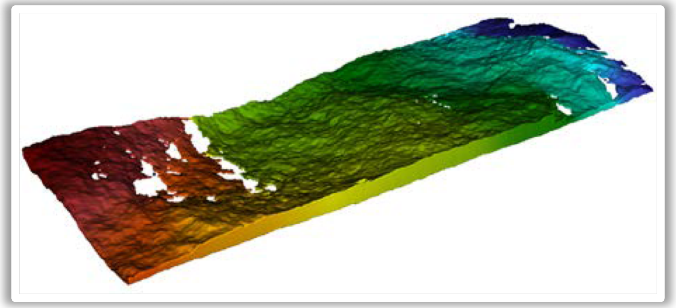

QMCPACK

Materials for Extreme Environments
(Paul Kent, ORNL)

- Find, predict and control materials and properties at the quantum level with unprecedented accuracy for the design novel materials that rely on metal to insulator transitions for high performance electronics, sensing, storage

Challenges: minimizing on-node memory usage; parallel on-node performance of Markov-chain Monte Carlo

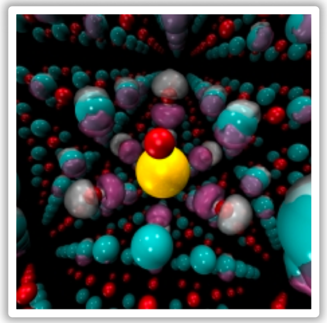

ECP EXASCALE COMPUTING PROJECT

Efficient Monte Carlo on accelerator–based architectures

<u>Challenge</u>: Monte Carlo neutron particle transport is a stochastic method

- Not amenable to single kernel optimization – no "high cost" kernel to optimize
- Independent random walks are not readily amenable to SIMT algorithms
- Sampling data (interaction cross sections) are:
  - randomly accessed
  - characterized by detailed structure
  - in standard applications consist of large point-wise representations (>1 – 5 GB per temperature point)

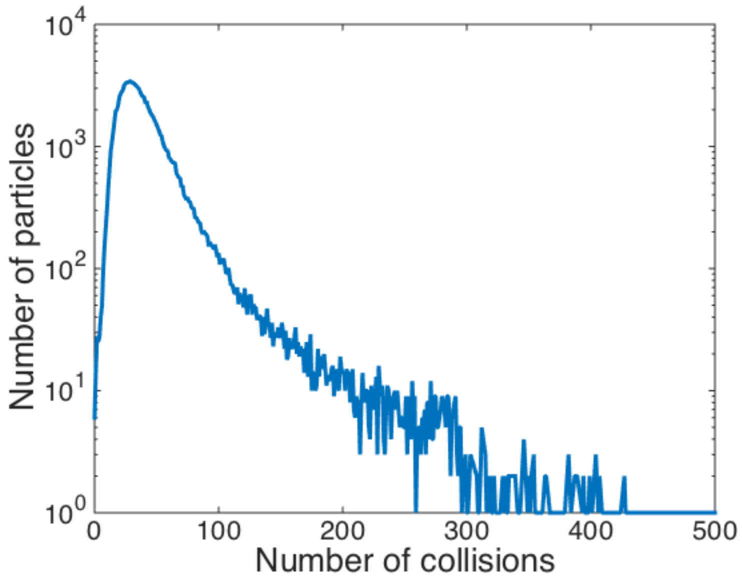

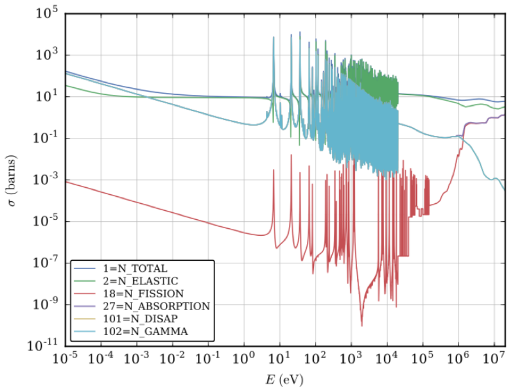

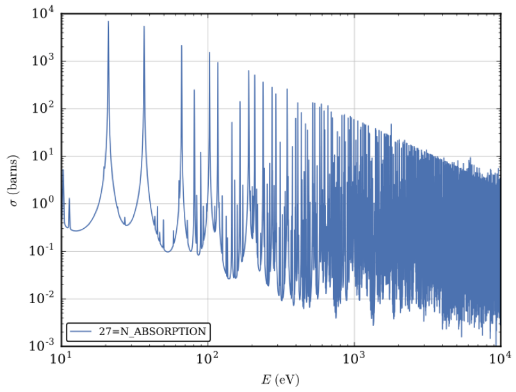

PI: Steve Hamilton (ORNL)

The Monte Carlo algorithm maps well to GPUs after changing from a history-based to an event-based algorithm

- Reduce thread divergence – change from history- to event-based algorithm
- Flatten algorithms to reduce kernel size; smaller kernels = higher occupancy
- Partition events based on fuel and non-fuel regions
- Debuted first comprehensive windowed multipole library for nuclear data (with temp correction)
- MC performance on Summit ~16x that achieved on Titan for the same algorithm
  - Significantly out-pacing gains in machine theoretical peak (7x on LINPACK)
- Overall MC performance has progressed from 15M to 600M particles/s!

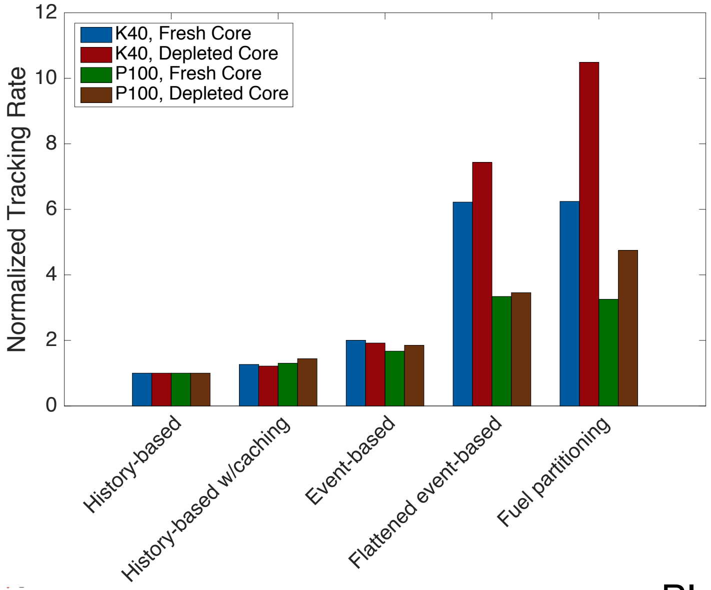

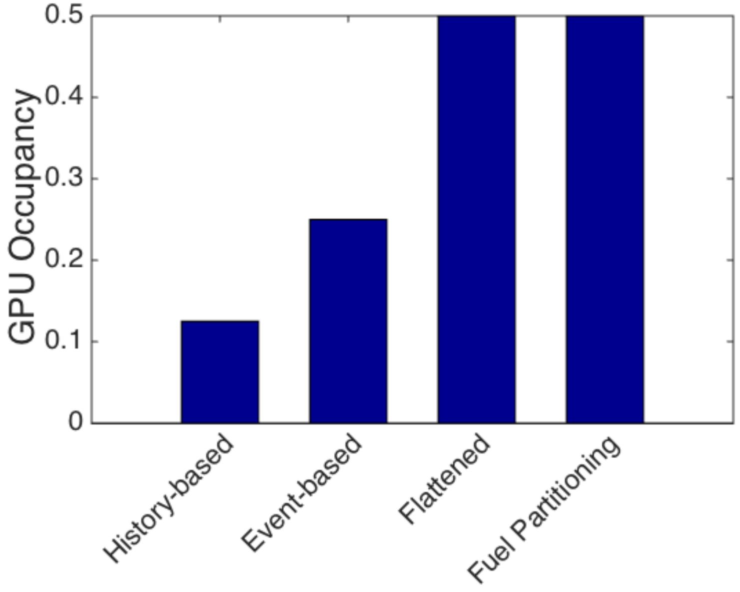

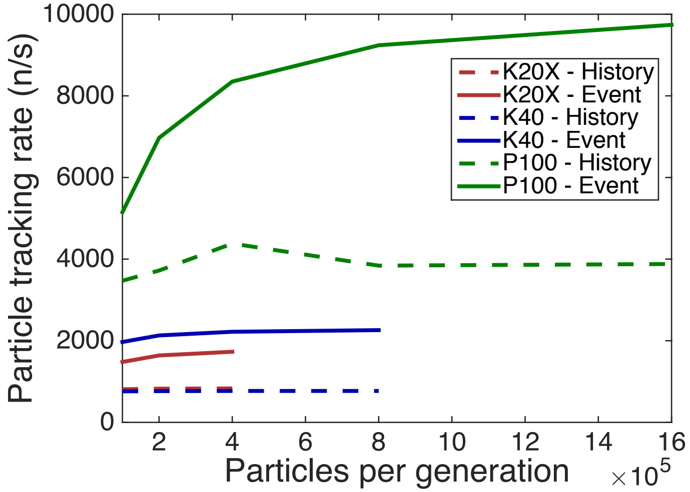

PI: Steve Hamilton (ORNL)

Exascale apps can deliver transformative products and solutions

ExaSGD

Reliable and Efficient Planning of the Power Grid
(Henry Huang, PNNL)

- Optimize power grid planning, operation, control and improve reliability and efficiency

Challenges: parallel performance of nonlinear optimization based on discrete algebraic equations and possible mixed-integer programming

Urban

Urban System Science
(Charlie Catlett, ANL)

- Evaluate energy codes and integration, retrofits, transportation, financing; integrate microgrids and renewables

Challenges: coupling of models (buildings, urban weather) with disparate time and length scales

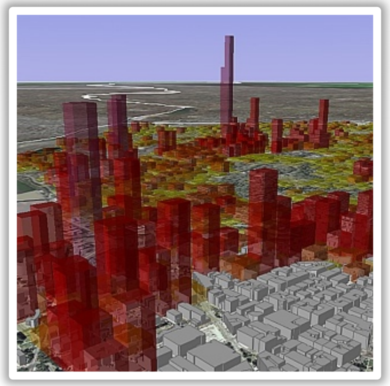

Combustion-PELE

High-Efficiency, Low-Emission Combustion Engine Design
(Jackie Chen, SNL)

- Reduce or eliminate current cut-and-try approaches for combustion system design

Challenges: performance of chemistry ODE integration on accelerated architectures; linear solver performance for low-Mach algorithm; explicit LES/DNS algorithm not stable

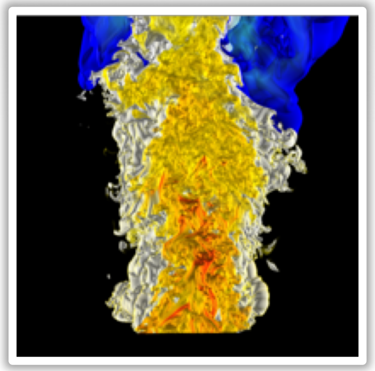

ECP EXASCALE COMPUTING PROJECT

Pele Code Design Overview

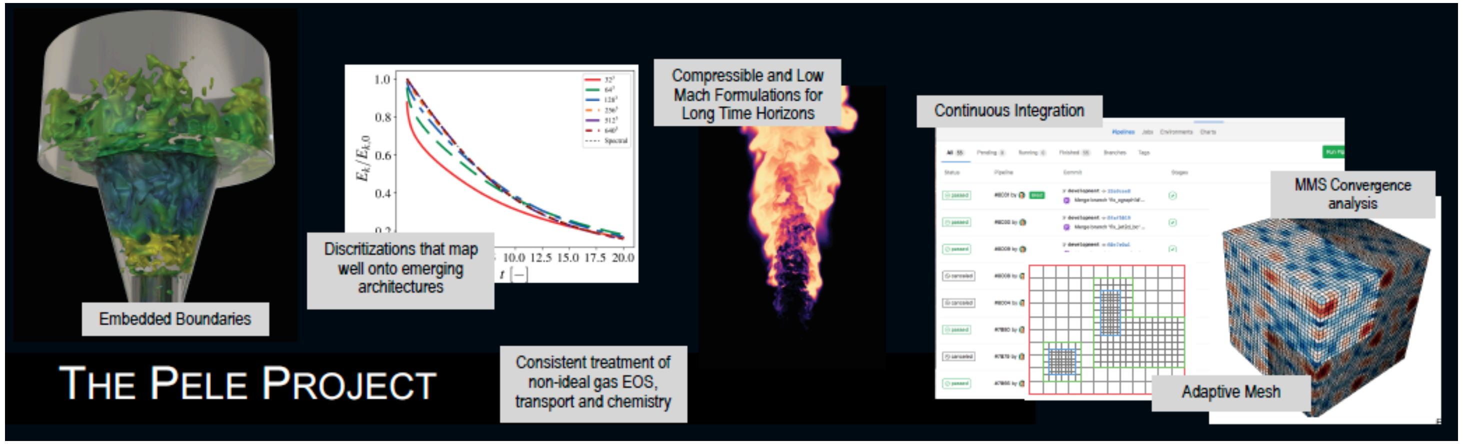

- Baseline algorithm design for multicomponent flowwith stiff reactions, AMR
  - **PeleC**: Comparable advection, diffusion time scales, motivates IMEX-type scheme based on Spectral Deferred Corrections (SDC) with time-implicit chemistry
    - Robust highly efficient time-explicit Godunov-type upwind advection, simple centered diffusion
    - BDF-style implicit chemistry ODE integration, with sources that approximate the other processes
  - **PeleLM**: acoustics filtered away analytically, but still want robust, time-explicit advection
    - Chemistry and diffusion are now time-implicit – iterative timestep simultaneously incorporates flow constraint (constant pressure), mutually coupled species/energy diffusion and chemistry. Entire system evolved stably on slower advection time scales across AMR grid hierarchy
- SDC-based iterative timestep – treats each process essentially independently, with accelerated iteration to couple everything together efficiently
- Robust baseline allows stable, well-behaved extensible time step
  - Switch 2nd order advection scheme with more accurate 4th order algorithm
  - Option for “destiffened” chemistry model that allows highly efficient time-explicit advance
  - Robust to other, potentially stiff, tightly coupled processes, such as sprays, radiation, soot, etc

Exascale apps can deliver transformative products and solutions

E3SM-MMF

Accurate Regional Impact Assessment in Earth Systems
(Mark Taylor, SNL)

- Forecast water resources and severe weather with increased confidence; address food supply changes

Challenges: MMF approach for cloud-resolving model has large biases; adequacy of Fortran MPI+OpenMP for some architectures; Support for OpenMP and OpenACC

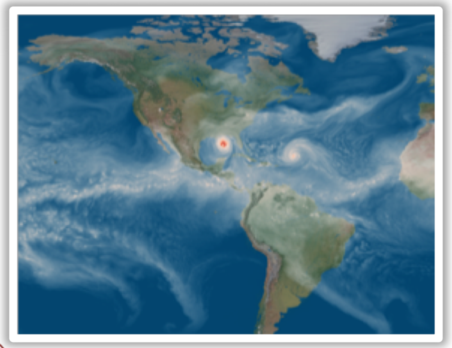

NWChemEx

Catalytic Conversion of Biomass-Derived Alcohols
(Thom Dunning, PNNL)

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

Challenges: computation of energy gradients for coupled-cluster implementation; on- and off-node performance

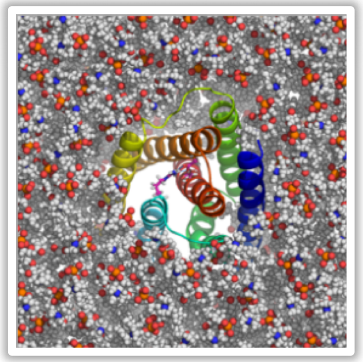

ExaBiome

Metagenomics for Analysis of Biogeochemical Cycles
(Kathy Yelick, LBNL)

- Discover knowledge useful for environmental remediation and the manufacture of novel chemicals and medicines

Challenges: Inability of message injection rates to keep up with core counts; efficient and performant implementation of UPC, UPC++, GASNet; GPU performance; I/O performance

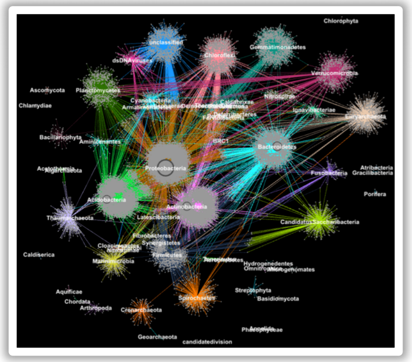

E3SM: Energy Exascale Earth System Model

- Global Earth System Model
- Atmosphere, Land, Ocean and Ice component models
- 8 DOE labs, 12 university subcontracts, 53 FTEs spread over 87 individuals
- Development driven by DOE Office of Scienc mission interests: Energy/water issues looking out 40 years
- Key computational goal: Ensure E3SM will run well on upcoming DOE pre-exascale and exascale computers
- E3SM is open source / open development
  - Website: www.e3sm.org
  - Github: https://github.com/E3SM-Project
  - DOE Science youtube channel: https://www.youtube.com/channel/UC_rhpi0lBeD1U-6nD2zvlBA

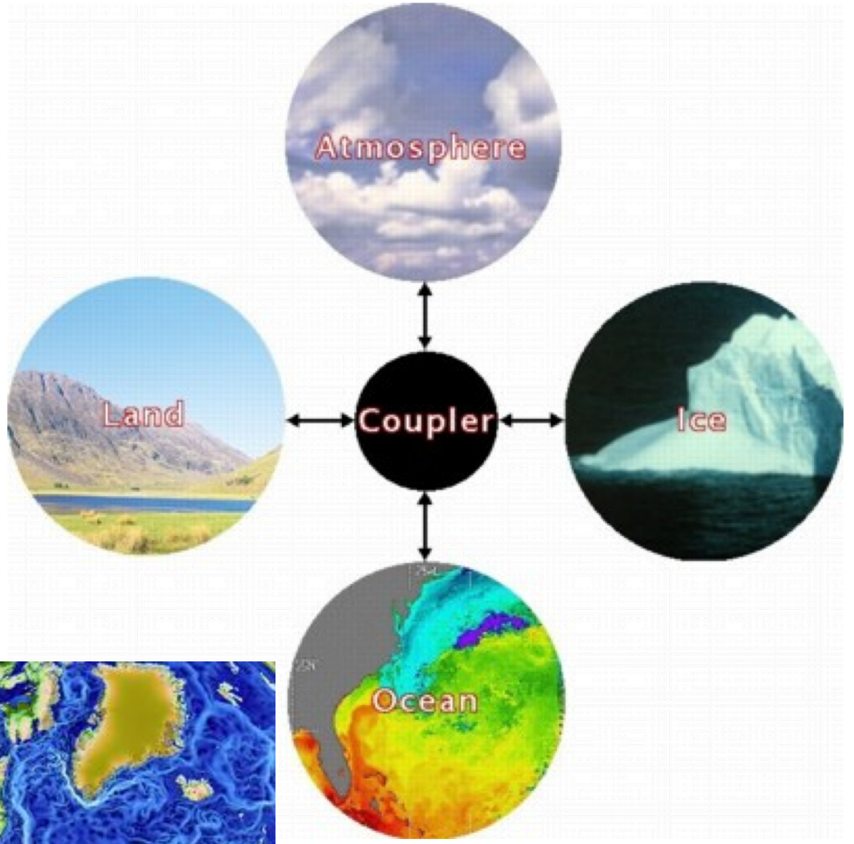

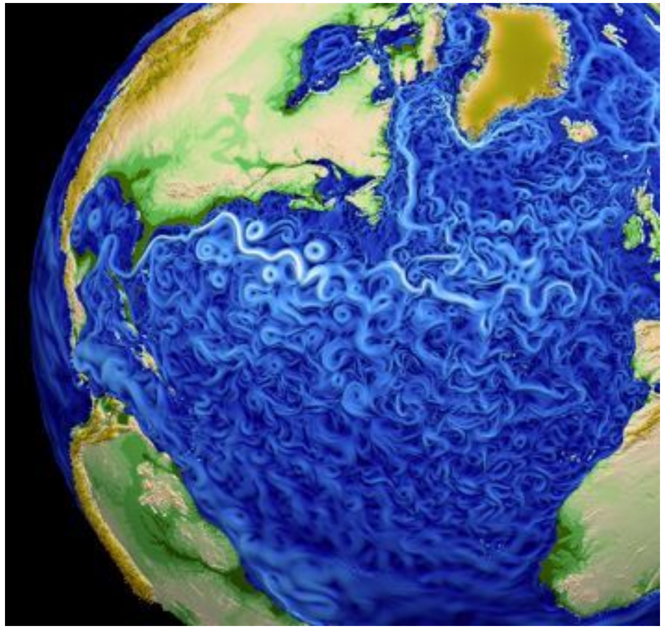

PI: David Bader (LLNL)

E3SM-Multiscale Modeling Framework (MMF)

Cloud Resolving Climate Model for E3SM

- Develop capability to assess regional impacts of climate change on the water cycle that directly affect the US economy such as agriculture and energy production.
- Cloud resolving climate model is needed to reduce major systematic errors in climate simulations due to structural uncertainty in numerical treatments of convection – such as convective storm systems
- Challenge: cloud resolving climate model using traditional approaches requires *zettascale* resources
- E3SM "conventional" approach:
  - Run the E3SM model with a global cloud resolving atmosphere and eddy resolving ocean.
    - 3 km atmosphere/land (7B grid points) and 15-5 km ocean/ice (1B grid points)
  - Achieve throughput rate of 5 SYPD to perform climate simulation campaigns including a 500 year control simulation
  - Detailed benchmarks on KNL and v100 GPUs show negligible speedups compared to conventional CPUs
    - Low arithmetic intensity of most of the code; throughput requirements lead to strong scaling and low work per node.
- E3SM-MMF: Use a multiscale approach ideal for new architectures to achieve cloud resolving convection on Exascale
  - Exascale will make "conventional" cloud resolving simulations routine for shorter simulations (process studies, weather prediction)
  - For cloud resolving climate simulations, we need fundamentally new approaches to take advantage of exascale resources

Convective storm system nearing the Chicago metropolitan area
http://www.spc.noaa.gov/misc/AbtDerechos/derechofacts.htm

PI: Mark Taylor (SNL)

The E3SM-MMF Exascale Challenge Problem

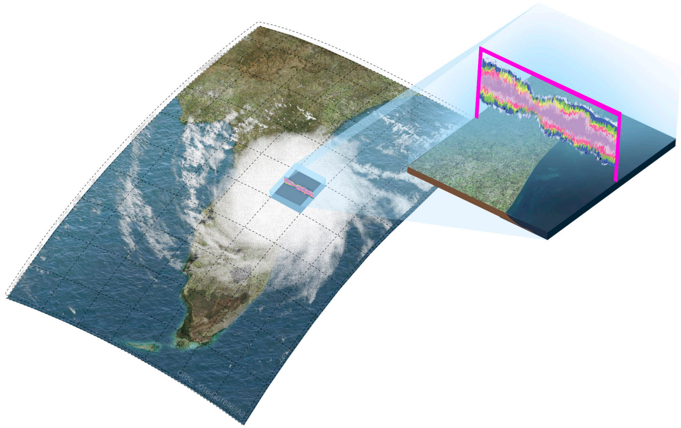

- E3SM-MMF approach addresses structural uncertainty in cloud processes by replacing traditional parameterizations with cloud resolving "superparameterization" within each grid cell of global climate model
- Super-parameterization dramatically increases arithmetic intensity, making the MMF approach one of the few ways to achieve exascale performance on upcoming architectures.
- Exascale + MMF approach will make it possible for the first time to perform climate simulation campaigns with some aspects of cloud resolving resolutions.
- Exascale challenge problem
  - Fully weather resolving atmosphere and cloud-resolving super-parameterization, an eddy resolving ocean and ice components, all while obtaining the necessary throughput to run 10-100 member ensembles of 100-year simulations.
  - Cloud-resolving: 1km grid spacing in both horizontal and vertical directions.
  - Weather resolving: 50-25km horizontal resolution, ~1km vertical (resolution of today's global operational forecast models)
  - Eddy resolving ocean/ice: minimum 18 km resolution in equatorial regions, decreasing to 6 km in polar regions to capture the reduction in eddy size with decreasing Rossby radius of deformation, with O(100) levels in the vertical.
  - Necessary throughput: 5 simulated-years-per-day (SYPD)
- Already achieved speedup of 20X relative to Titan (ORNL)
  - From 0.005 to 0.10 SYPD (goal of 5.0 on Exascale)

Exascale apps can deliver transformative products and solutions

ExaSky

Cosmological Probe of the Standard Model of Particle Physics
(Salman Habib, ANL)

- Unravel key unknowns in the dynamics of the Universe: dark energy, dark matter, and inflation

Challenges: subgrid model accuracy; OpenMP performance on GPUs; file system stability and availabilty

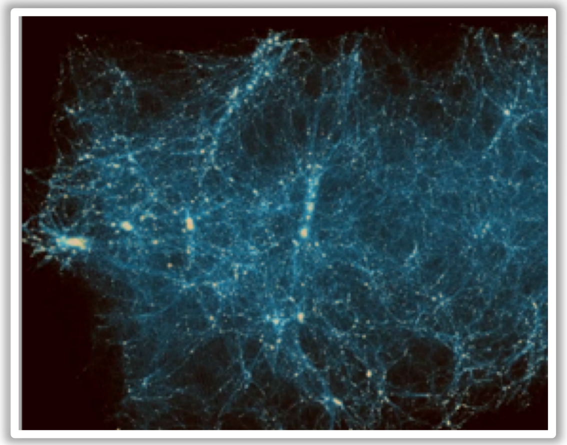

LatticeQCD

Validate Fundamental Laws of Nature
(Andreas Kronfeld, FNAL)

- Correct light quark masses; properties of light nuclei from first principles; <1% uncertainty in simple quantities

Challenges: performance of critical slowing down; reducing network traffic to reduce system interconnect contention; strong scaling performance to mitigate reliance on checkpointing

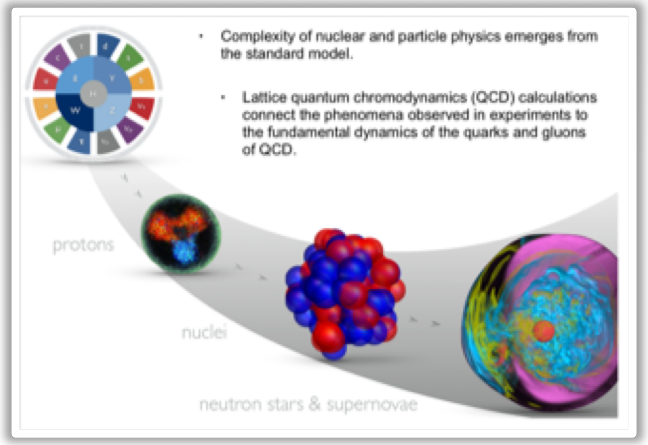

WarpX

Plasma Wakefield Accelerator Design
(Jean-Luc Vay, LBNL)

- Virtual design of 100-stage 1 TeV collider; dramatically cut accelerator size and design cost

Challenges: scaling of Maxwell FFT-based solver; maintaining efficiency of large timestep algorithm; load balancing

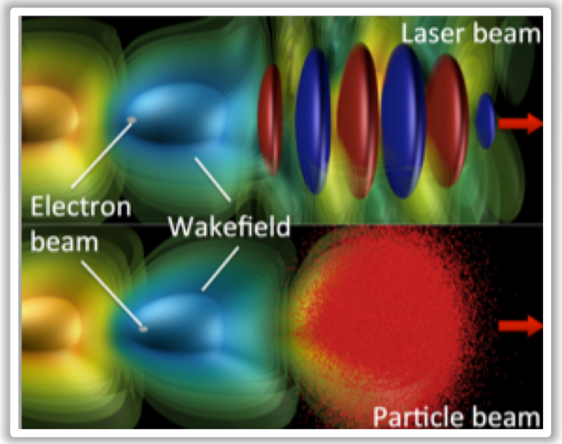

EXASCALE COMPUTING PROJECT

ExaSky: Computational Cosmology

Physics with cosmological surveys

- **'Dark Universe'**
  - Cosmic acceleration
  - Nature and properties of dark matter
  - Primordial fluctuations, inflation/B-modes
  - Neutrino properties
- **Cosmic Probes**
  - Gravitational lensing
  - Galaxy clustering
  - Clusters
  - CMB anisotropies
  - Supernovae
  - Cross-correlations

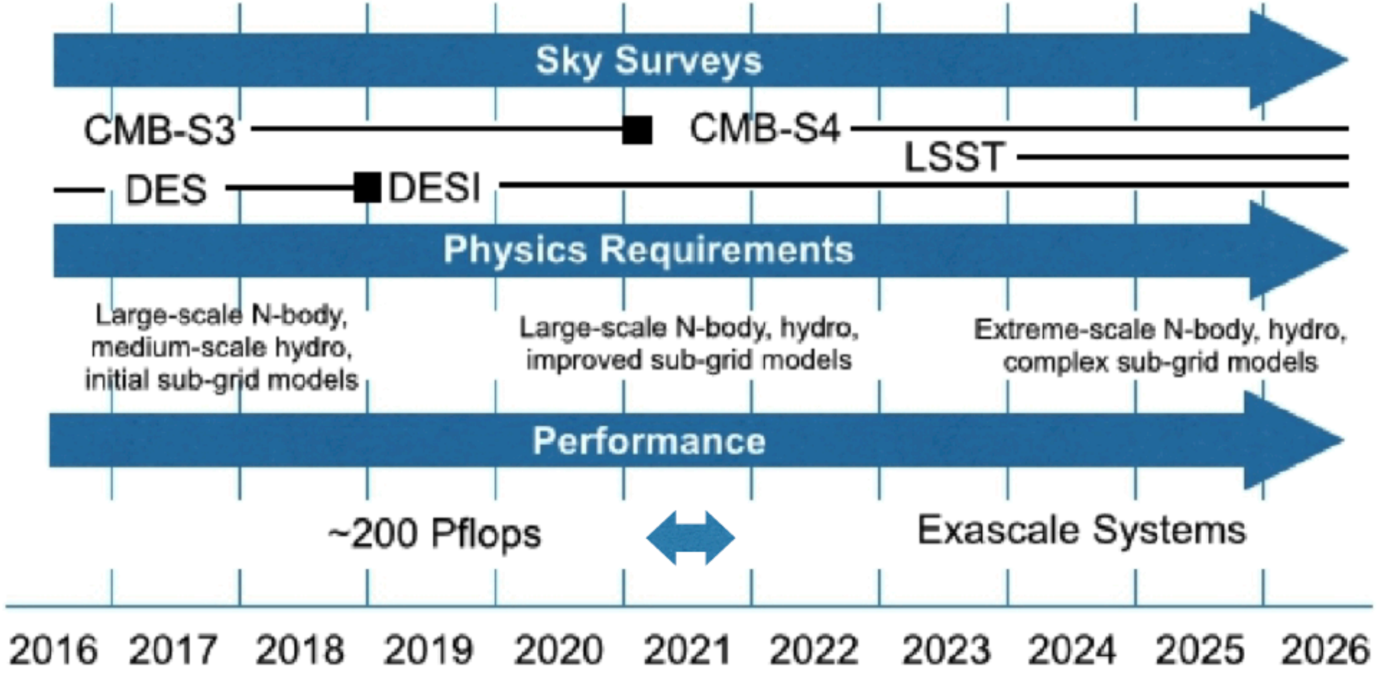

PI: Salman Habib (ANL)

ExaSky Codes: HACC and Nyx

- **HACC (FOM/Challenge Problem Code, 2008)**
  - Extreme-scale N-body/hydro (Lagrangian)
  - Hybrid SPM (SWFFT)-Tree/FM/P^3M gravity
  - CRK-SPH, next-gen SPH method
  - New subgrid models (cooling/heating, star formation, AGN/SN feedback)
  - MPI + X (OpenMP, CUDA, OpenCL, assembly, —)

Habib et al., New Astron. 42, 49 (2016); Comm. ACM 60, 97 (2017) [Research Highlight]

- **Nyx (2013)**
  - N-body/hydro (Eulerian, AMR)
  - PM (multigrid), AMR gravity
  - Unsplit PLM/PPM
  - New subgrid model suite
  - MPI + OpenMP (GPU extension underway)

Almgren et al., Astrophys. J. 765, 39 (2013)

High-z universe with Nyx: redshifts of reionization shown in a single slice of a simulation run on Cori

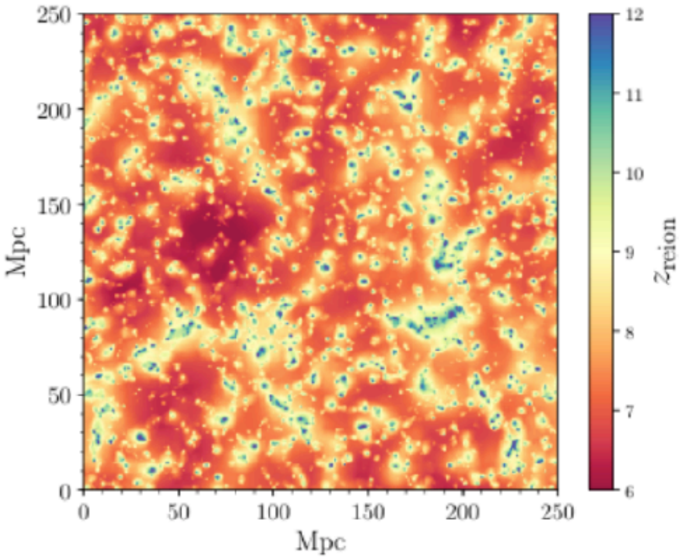

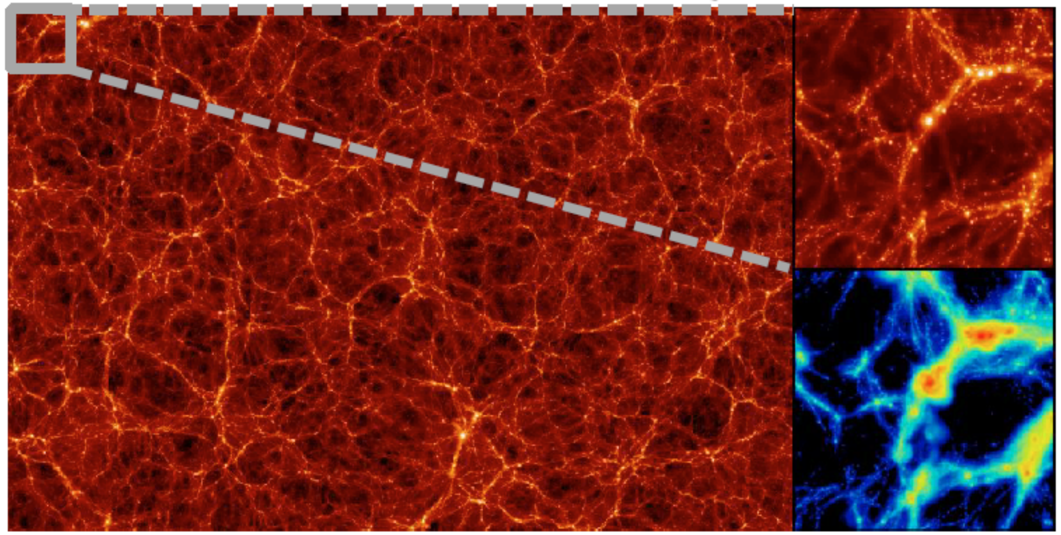

Production hydro run with the CRK-SPH version of HACC on Theta

PI: Salman Habib (ANL)

ExaSky: First HACC Runs on Summit

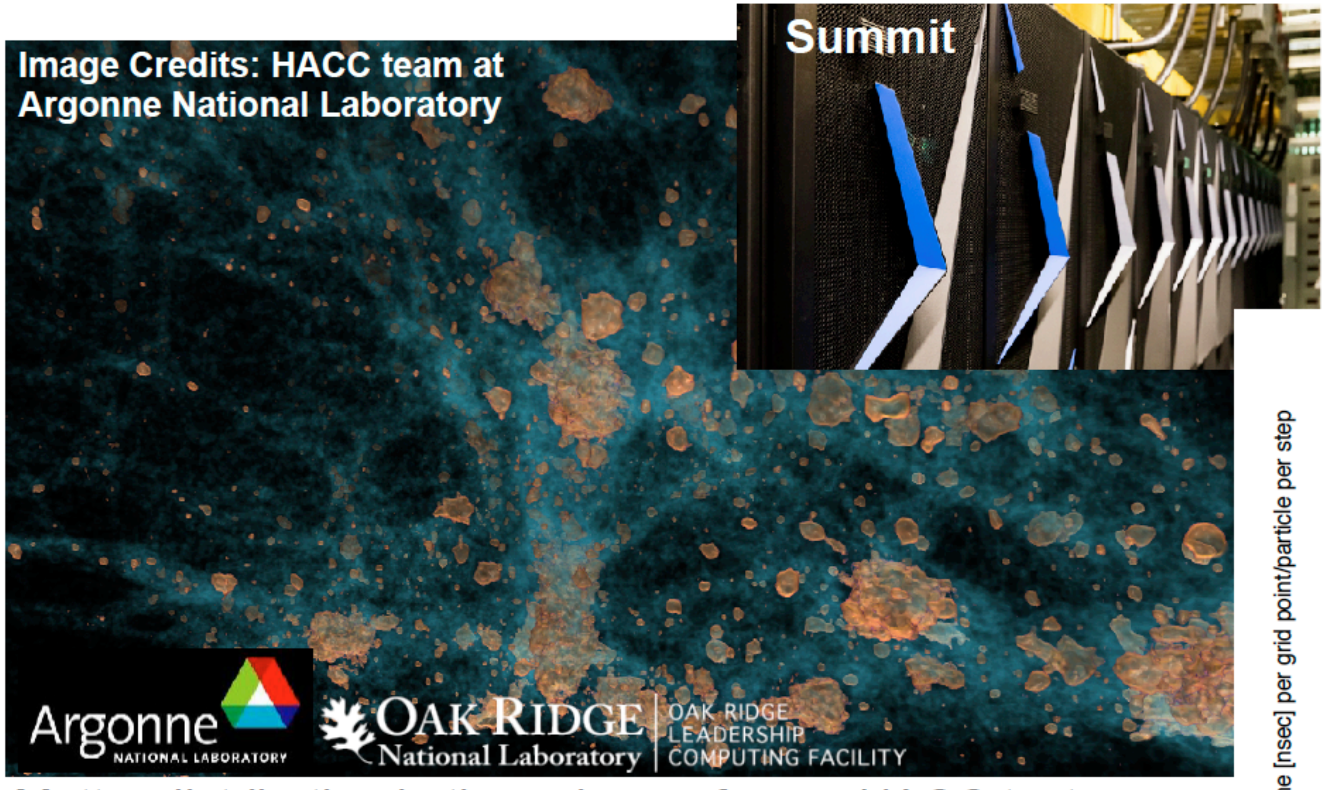

Matter distribution in the universe from a HACC test run on Summit at ORNL; this work is part of an OLCF Center for Accelerated Application Readiness (CAAR) supported project

Demonstration of GPU acceleration on Summit with ExaSky's extreme-scale HACC simulation framework

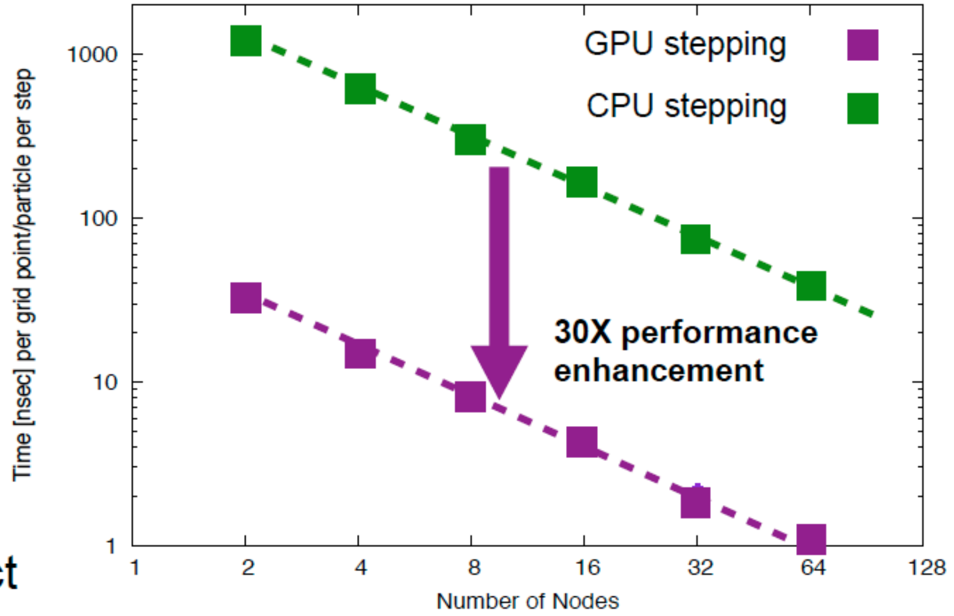

PI: Salman Habib (ANL)

Exascale apps can deliver transformative products and solutions

WDMApp

High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas
(Amitava Bhattacharjee, PPPL)

- Prepare for ITER exps and increase ROI of validation data and understanding
- Prepare for beyond-ITER devices

Challenges: robust, accurate, and efficient code-coupling algorithm; reduction in memory and I/O usage

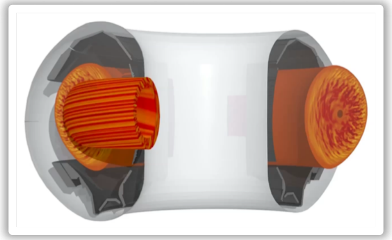

ExaStar

Demystify Origin of Chemical Elements
(Dan Kasen, LBNL)

- What is the origin of the elements?
- How does matter behave at extreme densities?
- What are the sources of gravity waves?

Challenges: delivering performance on accelerators; delivering fidelity for general relativity implementation

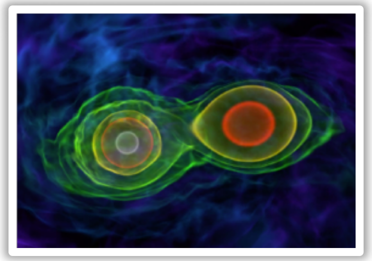

ExaFEL

Light Source-Enabled Analysis of Protein and Molecular Structure and Design
(Amadeo Perazzo, SLAC)

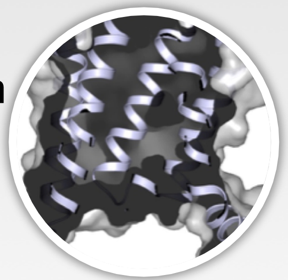

- Process data without beam time loss
- Determine nanoparticle size and shape changes
- Engineer functional properties in biology and materials science

Challenges: improving the strong scaling (one event processed over many cores) of compute-intensive algorithms (ray tracing, M-TIP) on accelerators

CANDLE

Accelerate and Translate Cancer Research
(Rick Stevens, ANL)

- Develop predictive preclinical models and accelerate diagnostic and targeted therapy through predicting mechanisms of RAS/RAF driven cancers

Challenges: increasing accelerator utilization for model search; effectively exploiting HP16; preparing for any data management or communication bottlenecks

Exascale Application Development Challenges Overall

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

ECP EXASCALE COMPUTING PROJECT

Overarching theme: must intensify efforts to effectively exploit reduced precision representations of hardware-accelerated operations

Exascale Applications: Take the Red Pill!

This is your last chance. After this, there is no turning back. You take the blue pill - the story ends, you wake up in your bed and believe whatever you want to believe. You take the red pill - you stay in Wonderland and I show you how deep the rabbit-hole goes.

EXASCALE COMPUTING PROJECT

Discussion