ECP Application Development Update

Andrew Siegel, Director, Application Development

ASCAC Virtual Meeting

Dec 12, 2018

ECP: Application Development Focus Area

- 33 Subprojects overall
- 25 Application Projects ← Domain science/engineering
  - 22 Office of Science
  - 3 NNSA
- 6 Co-Design Projects ← crosscutting infrastructure
- Proxy Applications, Application Assessment

EXASCALE COMPUTING PROJECT

What defines an application project?

1. Scientific or Engineering exascale *challenge problem*. ← impactful in community

2. Detailed completion criteria for successful execution of the challenge problem

3. Figure of Merit formula (FOM) quantifying performance of challenge problem

4. Where possible, a current *baseline* value of the FOM on an analogous problem

What defines a co-design project?

- Develop cross-cutting computational infrastructure based on HPC *motifs*
- Relevant to two or more ECP projects (and beyond)
- Require deeper engagement between application and software team than traditional third party libraries.

Exascale application projects

Challenge problem target and impact, lead institution, stakeholder

ExaWind: Turbine Wind Plant Efficiency

Harden wind plant design and layout against energy loss susceptibility; higher penetration of wind energy

Lead: NREL
DOE EERE

ExaAM: Additive Manufacturing of Qualifiable Metal Parts

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

Lead: ORNL
DOE NNSA / EERE

EQSIM: Earthquake Hazard Risk Assessment

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

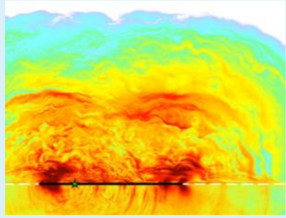

Lead: LBNL
DOE NNSA / NE, EERE

MFIX-Exa: Scale-up of Clean Fossil Fuel Combustion

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

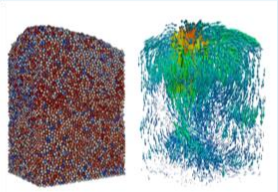

Lead: NETL
DOE EERE

GAMESS: Biofuel Catalyst Design

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

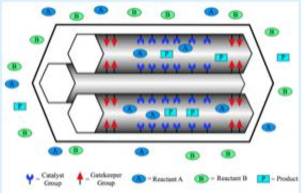

Lead: Ames
DOE BES

EXAALT: Materials for Extreme Environments

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

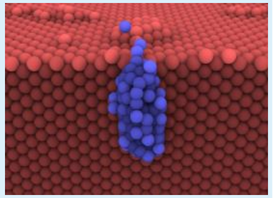

Lead: LANL
DOE BES, FES, NE

Exascale application projects, cont.

Challenge problem target and impact, lead institution, stakeholder

ExaSMR: Design and Commercialization of Small Modular Reactors

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

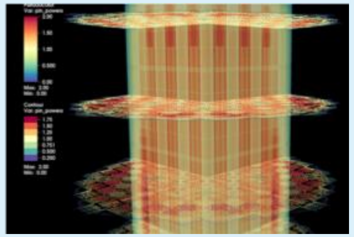

Lead: ORNL
DOE NE

Subsurface: Carbon Capture, Fossil Fuel Extraction, Waste Disposal

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

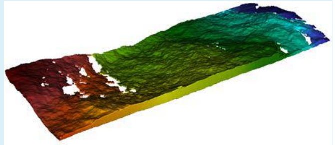

Lead: LBNL
DOE BES, EERE, FE, NE

QMCPACK: Find, Predict, Control Materials & Properties at Quantum Level

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

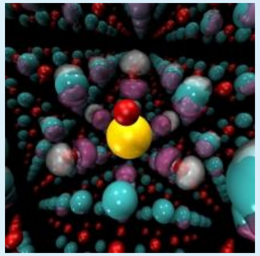

Lead: ORNL
DOE BES

ExaSGD: Reliable and Efficient Planning of the Power Grid

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

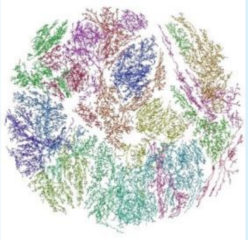

Lead: PNNL
DOE EDER, CESER, EERE

Urban: Urban Systems Science

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

Lead: ANL
DOE EERE, BER

Combustion-PELE: High-Efficiency, Low-Emission Combustion Engine Design

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

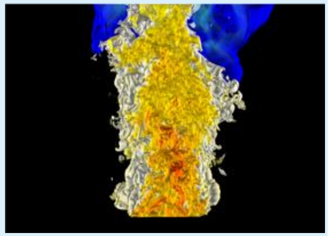

Lead: SNL
DOE BES, EERE

Exascale application projects, cont.

Challenge problem target and impact, lead institution, stakeholder

E3SM-MMF: Accurate Regional Impact Assessment in Earth Systems

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

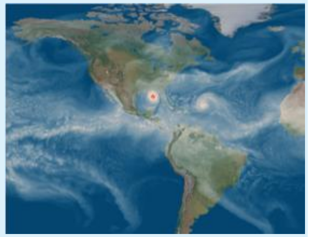

Lead: SNL
DOE BER

NWChemEx: Catalytic Conversion of Biomass-Derived Alcohols

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

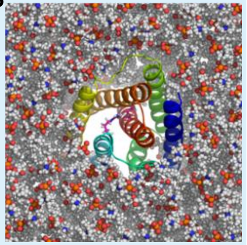

Lead: PNNL
DOE BER, BES

ExaBiome: Metagenomics for Analysis of Biogeochemical Cycles

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

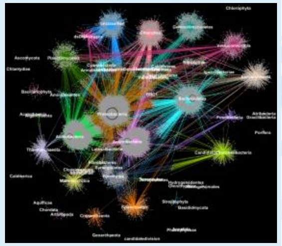

Lead: LBNL
DOE BER

ExaSky: Cosmological Probe of the Standard Model of Particle Physics

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

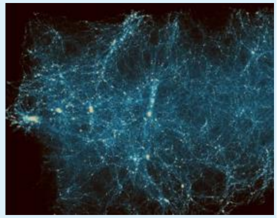

Lead: ANL
DOE HEP

LatticeQCD: Validate Fundamental Laws of Nature

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

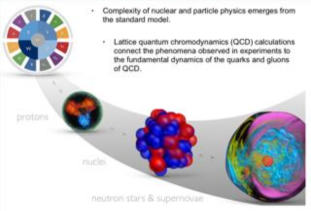

Lead: FNAL
DOE NP, HEP

WarpX: Plasma Wakefield Accelerator Design

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

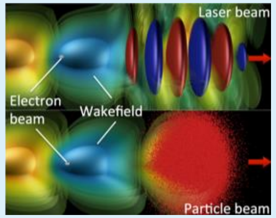

Lead: LBNL
DOE HEP

Exascale application projects, cont.

Challenge problem target and impact, lead institution, stakeholder

WDMApp: High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas

Prepare for ITER experiments and increase ROI of validation data and understanding; prepare for beyond-ITER devices

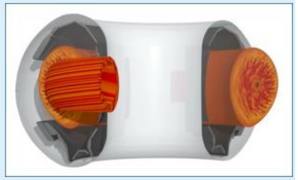

Lead: PPPL
DOE FES

ExaStar: Demystify Origin of Chemical Elements

What is the origin of the elements? Behavior of matter at extreme densities? Sources of gravity waves?

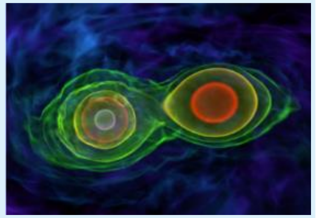

Lead: LBNL
DOE NP

ExaFEL: Light Source-Enabled Analysis of Protein and Molecular Structure and Design

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

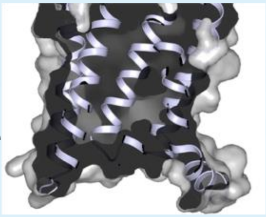

Lead: SLAC
DOE BES

CANDLE: Accelerate and Translate Cancer Research

Develop predictive pre-clinical models & accelerate diagnostic and targeted therapy thru predicting mechanisms of RAS/RAF driven cancers

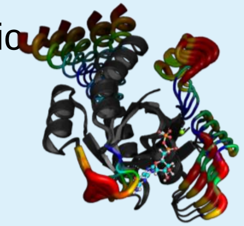

Lead: ANL
NIH

EXASCALE COMPUTING PROJECT

Exascale Multiphysics Applications for National Security Mission

ECP application projects in the NNSA ASC ATDM program element

The MARBL Multi-physics Code

Multi-physics simulations of high energy-density physics and focused experiments driven by high-explosive, magnetic or laser based energy sources

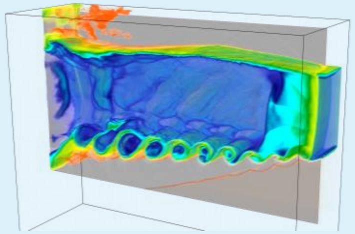

- Magneto-radiation-hydrodynamics at the exascale
- Next-generation pulsed power / ICF modeling
- High-order numerical methods

Lead: LLNL

Ristra: Next-Generation Multi-physics for National-Security Applications

– 3D multi-physics for national-security mission
– Mesoscale insight for extreme-condition materials
– Exascale high energy density physics simulations

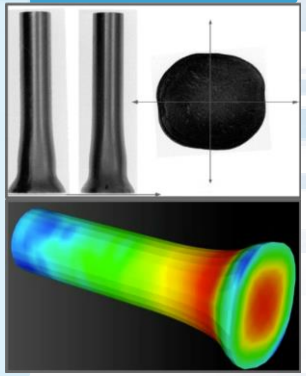

A Ristra hydrodynamics code with an advanced grain-structure-aware material model captures the asymmetric deformation in Taylor-Anvil experiments

Lead: LANL

EMPIRE for Electromagnetic Plasma Physics

Computing electronic effects induced by ionizing radiation interacting with materials under various re-entry flight conditions

Self-consistent plasma simulation including the radiation output of a hostile builder device, radiation transport, plasma generation and propagation down through the effects on ND system electronics

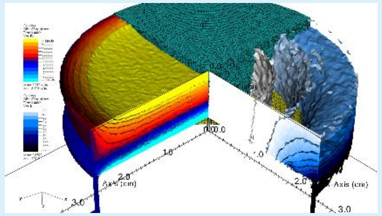

Lead: SNL

SPARC for Virtual Flight Testing

Virtual flight test of re-entry vehicles from exo-atmospheric bus separation to target for normal and hostile environments.

State-of-the-art hypersonic flight simulation capability on next-generation hardware, including thermo-chemical non-equilibrium gas ablation models, and hybrid RANS-LES turbulence models.

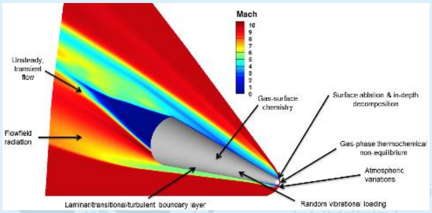

Lead: SNL

Co-Design

Development of 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	Co-Design center developing exascale toolkits for machine learning that facilitate the analysis of exascale data sets generated in ECP applications

Common R&D activities/challenges emerging from reviews

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

Common R&D activities/challenges emerging from reviews

1) **Porting to accelerator-based architectures** ExaSMR
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**

Porting to accelerator–based architectures: ExaSMR
Steve Hamilton (ORNL)

Distribution of history lengths in SMR core

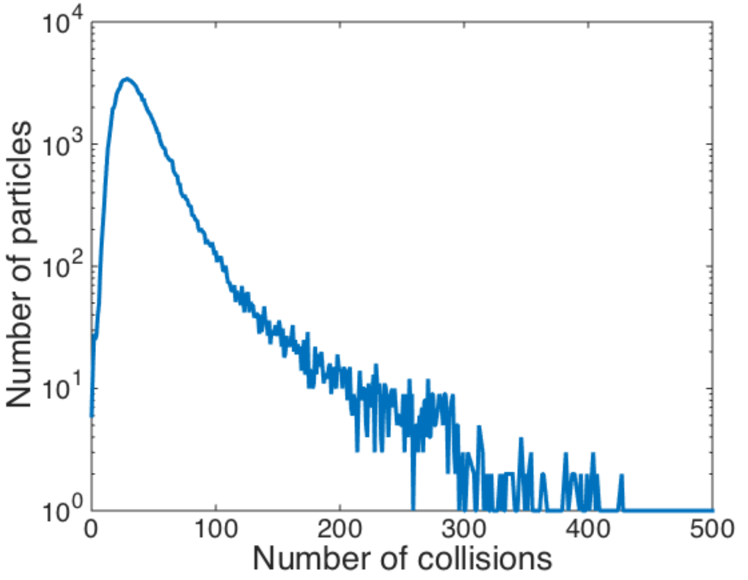

<u>Challenge</u>: Monte Carlo neutron particle transport is a stochastic method

- 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)

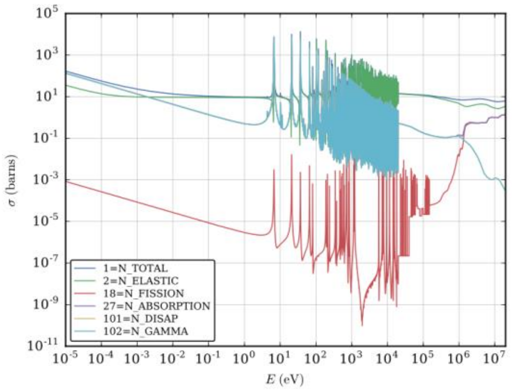

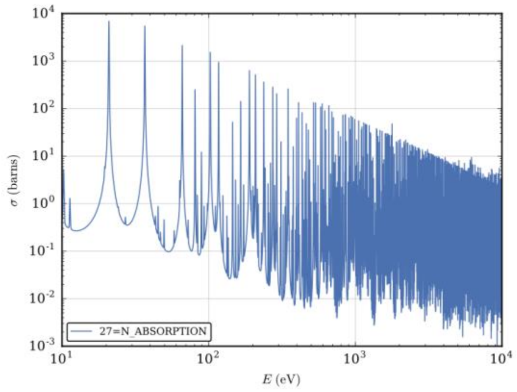

EXASCALE COMPUTING PROJECT

Algorithmic mapping to hardware – neutron particle transport

- 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
- Take advantage of other architectural improvements

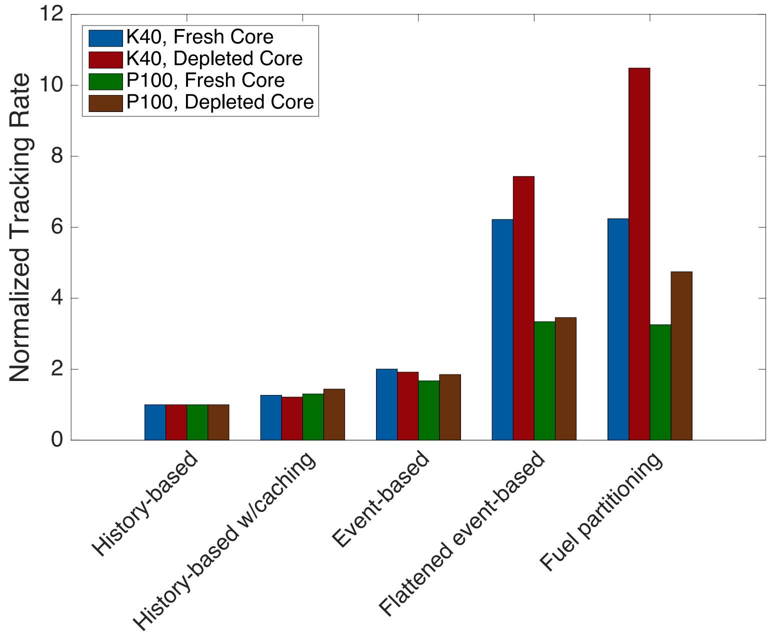

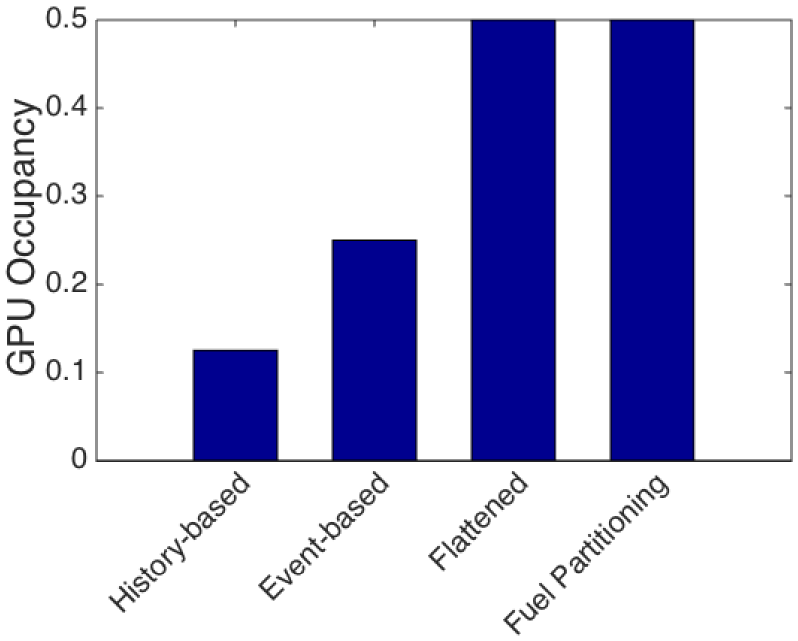

EXASCALE COMPUTING PROJECT

ExaSMR Figure of Merit (FOM)

FOM for coupled Monte Carlo neutronics + CFD modeling of SMRs

$$\mathrm{FOM}_{\mathrm{ExaSMR}} = \frac{1}{2}\mathrm{FOM}_{\mathrm{MC}} + \frac{1}{2}\mathrm{FOM}_{\mathrm{CFD}}$$

In which

$$\mathrm{FOM}_{\mathrm{MC}} = \frac{\text{particles}}{\text{wall clock seconds}}$$

$$\mathrm{FOM}_{\mathrm{CFD}} = \frac{\text{degrees of freedom}}{\text{wall clock seconds}}$$

ExaSMR FOM baseline

- Both physics components use Titan for baseline calculations
  - Requires baseline calculations to utilize GPUs
  - Extrapolated projections using documented excellent scalability on Titan
- MC transport: 15.1 million particles per second
  - History-based GPU implementation (first history-based MC algorithm deployed on GPUs)
  - Extrapolated from 4096 node calculation
- CFD: 36.8 million degrees-of-freedom per second
  - OpenACC implementation
  - Extrapolated from 240 node calculation

ExaSMR FOM progress

- Preliminary FOMs calculated on Summit during early-access period
  – Same problem specification as baseline

- MC transport: 596 million particles per second
  – FOM increase from Titan: 39x
  – Uses optimized event-based GPU algorithm

- CFD: 150 million degrees-of-freedom per second
  – FOM increase from Titan: 4.1x
  – Optimized libParanumal implementation from CEED expected to provide additional 4x

- Overall projected FOM increase on Summit: 22x
  – With CEED CFD optimizations: 28x

Applications face common challenges in preparing for exascale

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

EXASCALE COMPUTING PROJECT

Coupling Structural Deformation and Wave Propagation Codes
EQSim (D. McCallen, LBL)

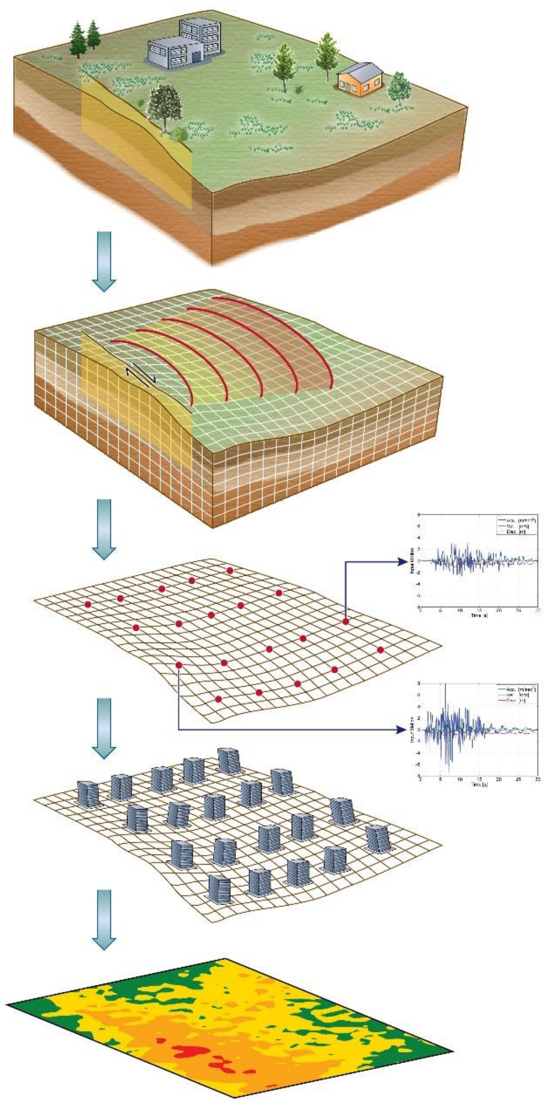

Earthquake Hazard

Earthquake Risk

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

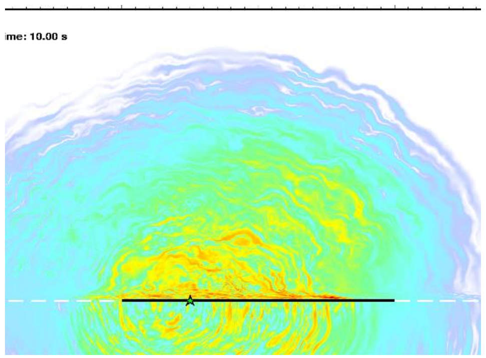

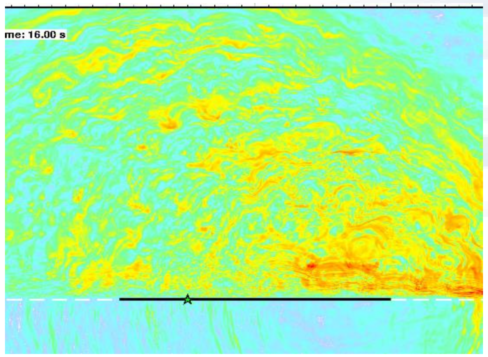

***NEVADA & MSESSI* – finite deformation, inelastic Finite Element codes for structures and soils**

Weak Coupling

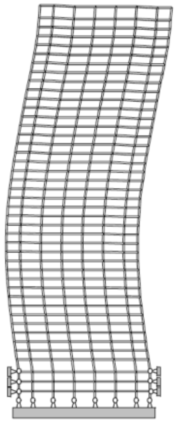

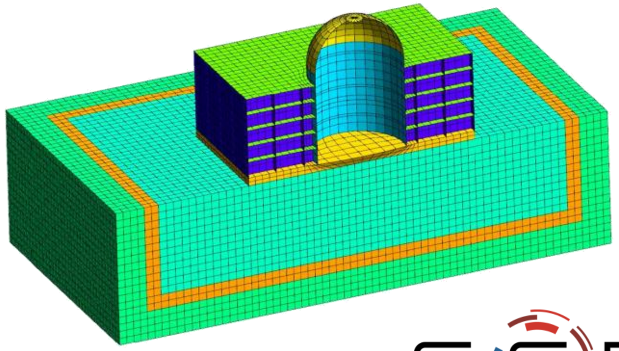

Strong Coupling

Developing thousands of building response simulations for near-field motions

Earthquake hazard

Surface motions from regional geophysics simulation

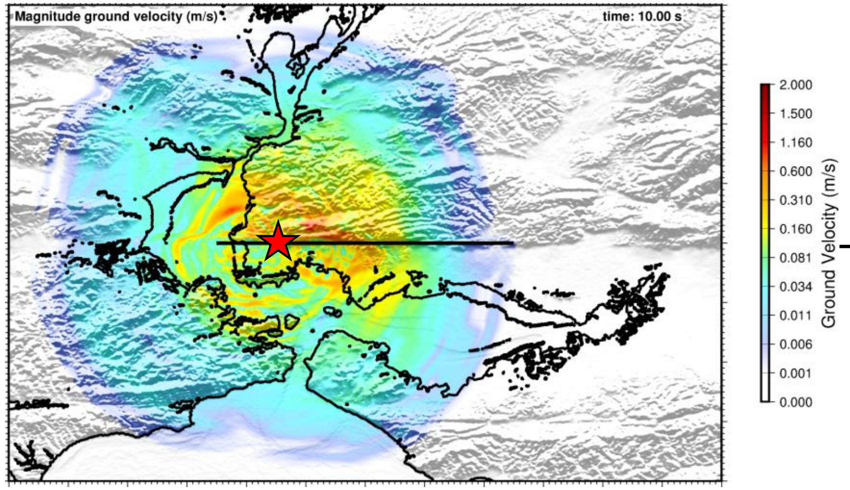

★ Rupture hypocenter

SW4

~ 2000 nonlinear building response history simulations

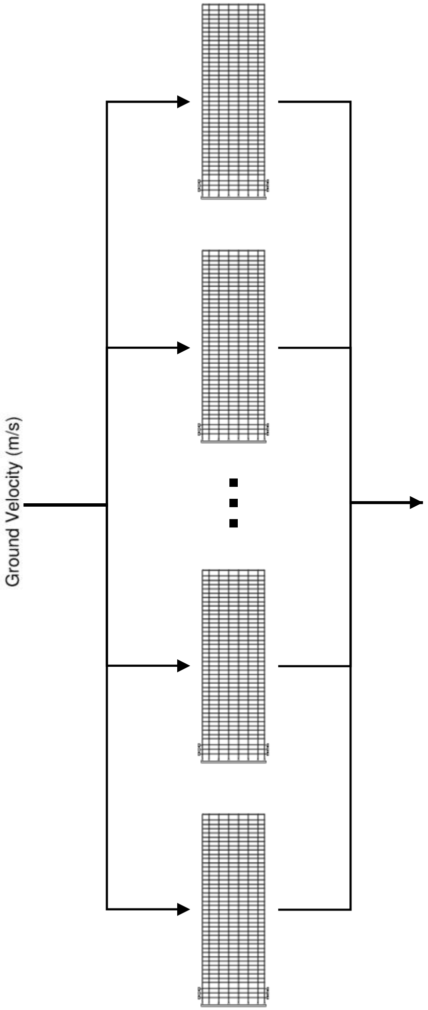

NEVADA

Earthquake risk

Distribution of building peak interstory drift

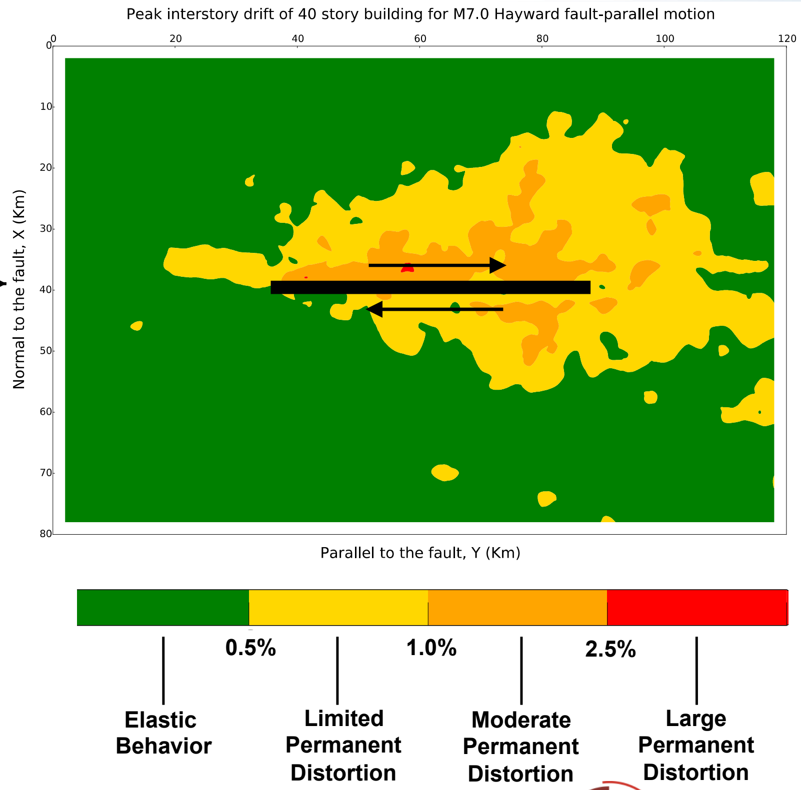

EXASCALE COMPUTING PROJECT

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, **FOM = 24.2**

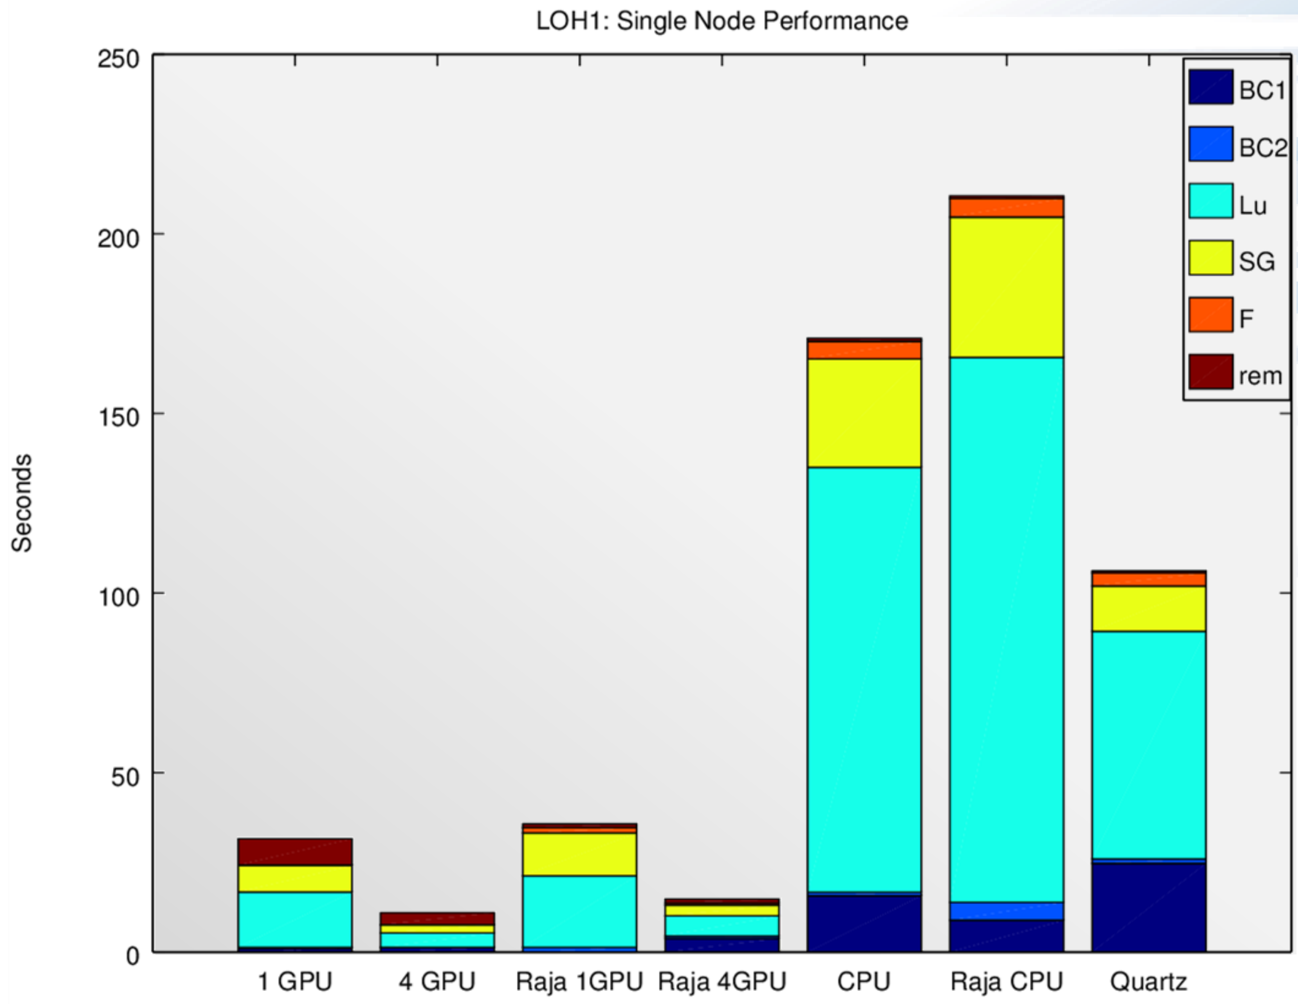

EXASCALE COMPUTING PROJECT

Applications face common challenges in preparing for exascale

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** — ExaWind
6) **Leveraging optimized libraries** — ExaWind

Leveraging optimized libraries for CFD
ExaWind (M. Sprague, NREL)

Challenge: Satisfying strong and weak-scaling limits for simulation of flow around wind farm using semi-implicit, projection-based CFD

- Large unstructured grids *O(100B)* points
- Large number of time steps *O(1M)*
- Pressure-Poisson and Momentum linear solves each timestep
  - Moving mesh around turbine blades requires matrix reinitialization every timestep

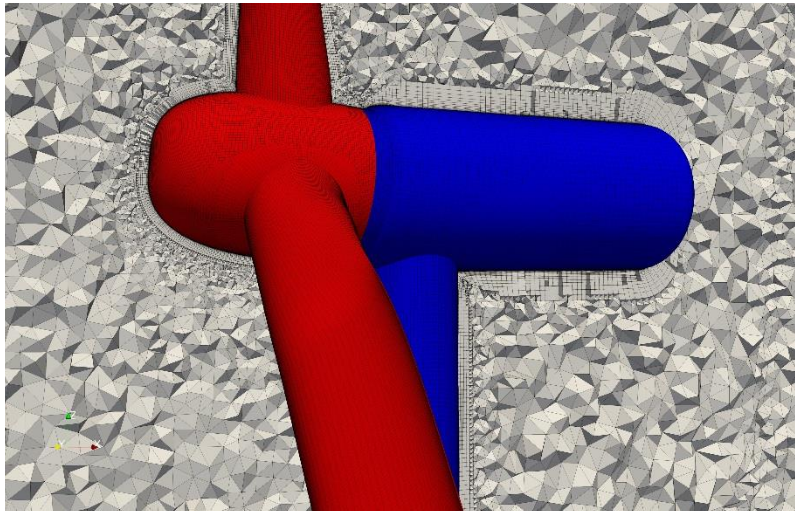

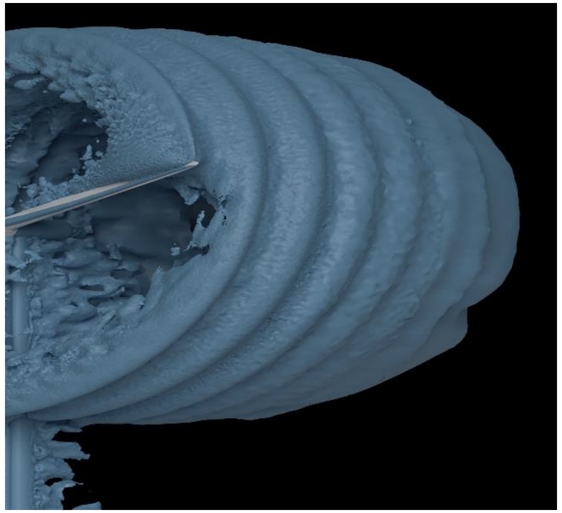

Primary application codes

Nalu-Wind

- Wind-specific version of the Nalu code developed at SNL
- https://github.com/exawind/nalu-wind
- Low-Mach-number computational fluid dynamics (CFD) code
- Unstructured-grid finite volume discretization
- C/C++
- Built on Trilinos

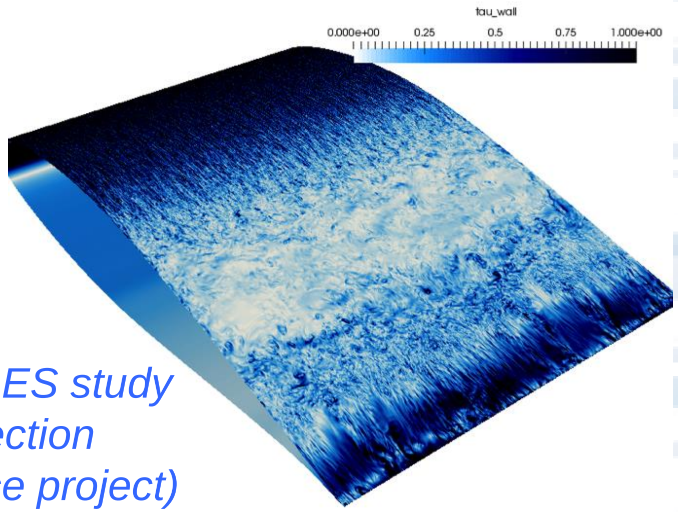

Nalu wall-resolved LES study of a turbine blade section (Trinity Open Science project)

OpenFAST

- https://github.com/openfast
- Whole-turbine simulation code
- Includes models for blades, control system, drivetrain, tower, etc.
- Fortran 90; dedicated Intel Parallel Computing Center (IPCC) for parallelization

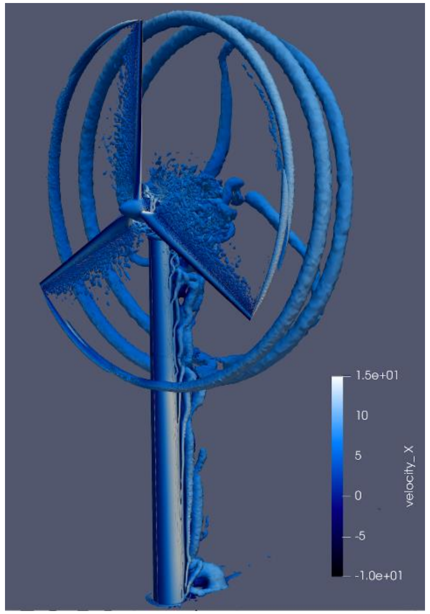

Using optimized solver libraries – Hypre and Trilinos

- Incorporate multiple solver stacks to leverage solver approaches
- Hypre
  - Classic (Ruge-Stüben) algebraic multigrid pre-conditioners
  - GMRES
- Trilinos
  - MueLu smoothed aggregation algebraic multigrid pre-conditioners
  - Belos Krylov solver library (GMRES)
  - Kokkos platform-portable layer
- Collaborate on architectural solver improvements
  - Hypre-GPU: Incorporate a version of GMRES optimized for GPUs using highly parallelizable versions of Gram-Schmidt for orthogonalization
  - Multi-threading in MueLu setup and solve phases

GPU Hypre development

- GMRES: memory intensive and hard to parallelize
- Modified Gram-Schmidt: the most expensive part
- Idea: Replace Hypre Level-1 BLAS MGS-GMRES with cheaper, yet stable alternatives
- Result: 35x speedup over HYPRE GPU enabled GMRES for V27 problem on a single GPU

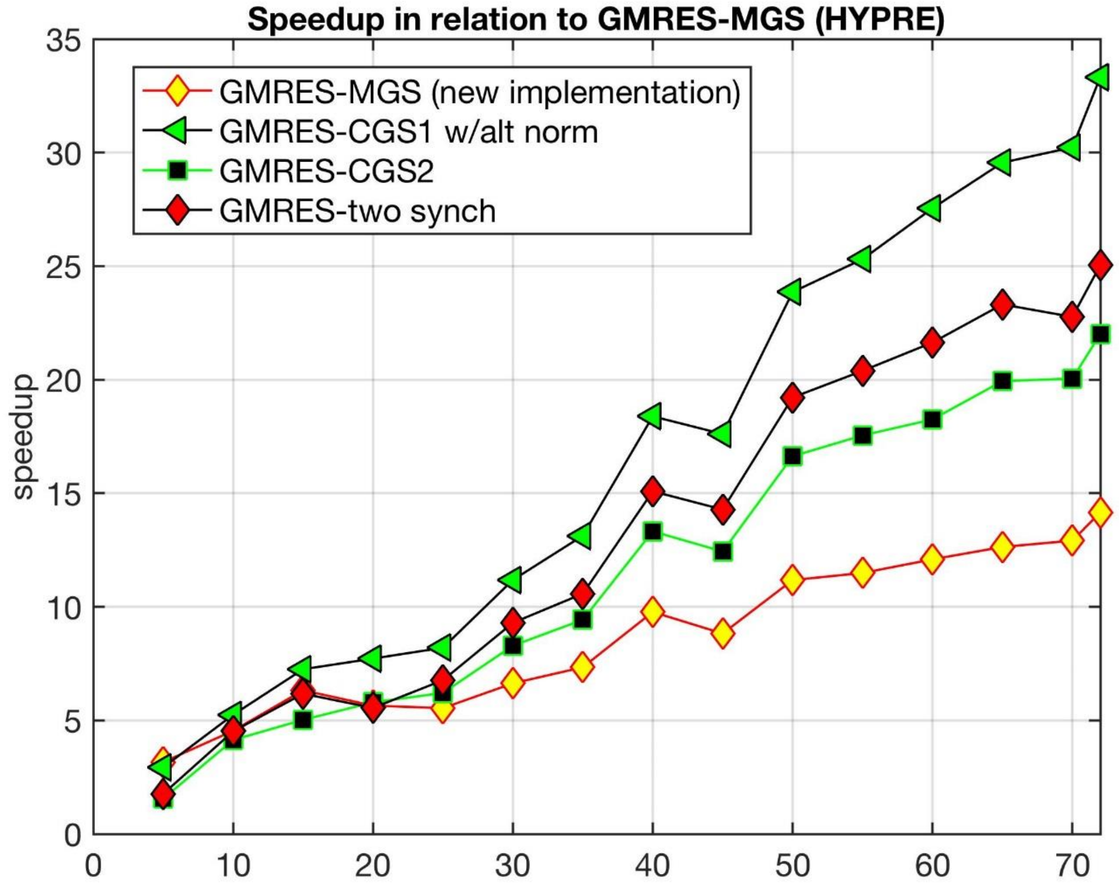

Number of vectors in Krylov subspace

Results for V27 problem, NVIDIA P100 16 GB

(Team: Swirydowicz, Thomas, Ananthan)

Co-Design
CEED (T. Kolev, LLNL)

CEED
EXASCALE DISCRETIZATIONS

Goal

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

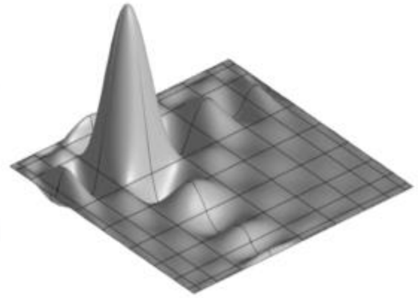

Approach

- 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

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

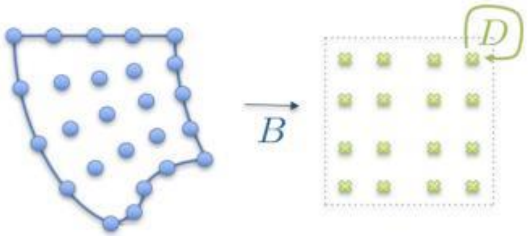

CEED is targeting several ECP applications

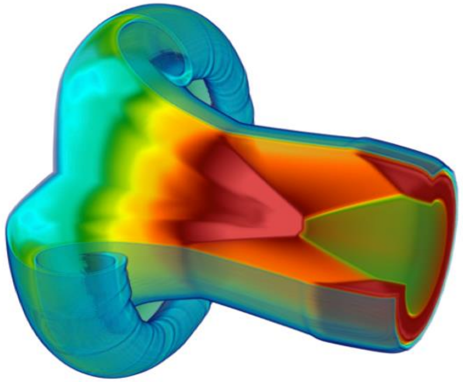

Compressible flow (MARBL)

Climate (E3SM)

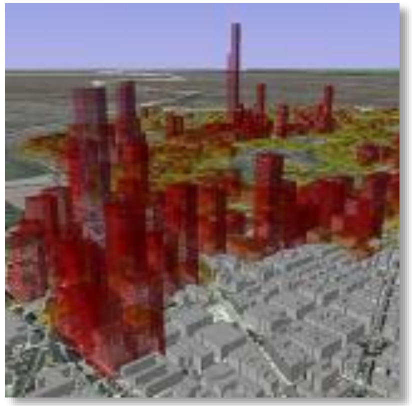

Urban systems (Urban)

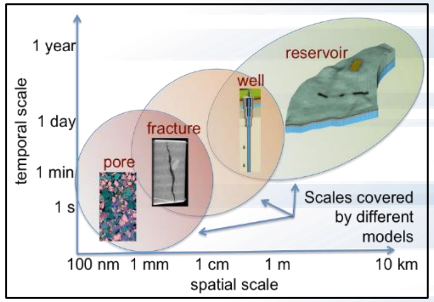

Subsurface (GEOS)

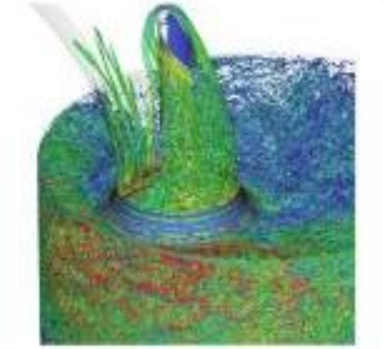

Combustion (Nek5000)

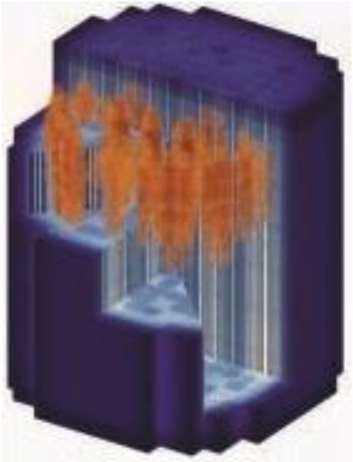

Modular Nuclear Reactors (ExaSMR)

Wind Energy (ExaWind)

Additive Manufacturing (ExaAM)

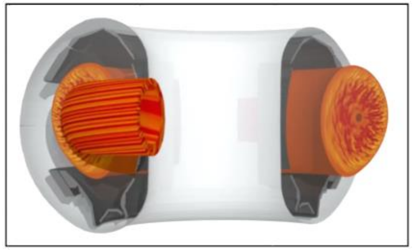

Magnetic Fusion (WDMApp)

EXASCALE COMPUTING PROJECT

Conclusions/Next Steps

- Very difficult project with highly complex and specialized node architectures enabling ambitious science and performance goals.

- Early adoption of intermediate (100PF) systems, test hardware, and hardware simulators critical to lowering risk by enabling progress tracking and early identification of issues.

- Considerable progress to date, need to continue to push early adoption of exascale-type hardware, ensure proper balance of domain expertise and performance engineering. Facilities engagement programs are critical to achieving this.