CFD Applications in the DOE Exascale Computing Project

Douglas B. Kothe
Director, Exascale Computing Project
Oak Ridge National Laboratory

Future CFD Technologies Workshop
Sponsored by the AIAA CFD2030 Integration Committee and NASA's Transformative Tools and Technologies Project
Kissimmee, Florida
January 7, 2018

Components of the DOE Exascale Program

- **Exascale Computing Initiative (ECI)**
  - The ECI was initiated in FY 2016 to support research, development and computer system procurements to deliver an exascale (10^{18} ops/sec) computing capability by the early to mid-2020s.
  - It is a partnership between SC and NNSA, addressing science and national security missions.
  - In the FY 2018 President's Budget request, ECI includes the SC/ASCR and NNSA/ASC facility investments in site preparations and non-recurring engineering activities needed for delivery of early to mid-2020s exascale systems.
  - **Exascale Computing Project (ECP)**
    - Beginning in FY 2017, the ASCR ECI funding was transitioned to the DOE project (ECP), which is managed according to the principles of DOE Order 413.3B.
    - The ECP subprogram in ASCR (SC-ECP) includes only support for research and development activities in applications, and in partnership with NNSA, investments in software and hardware technology and co-design required for the design of capable exascale computers.
    - The NNSA/ASC Advanced Technology Development and Mitigation (ATDM) program supports the development of applications and, in collaboration with SC/ASCR, investments in software and hardware technology and co-design required for the design of exascale capable computers.

Relevant DOE Pre-Exascale and Exascale Systems for ECP

Pre-Exascale Systems

2013

Argonne
IBM BG/Q
Open

ORNL
Cray/NVidia K20
Open

LLNL
IBM BG/Q
Secure

2016

Argonne
Intel/Cray KNL
Open

LBNL
Cray/Intel Xeon/KNL
Open

LANL/SNL
Cray/Intel Xeon/KNL
Secure

2018

ORNL
IBM/NVidia
P9/Volta
Open

LLNL
IBM/NVidia
P9/Volta
Secure

2020

NERSC-9

LBNL
TBD
Open

Crossroads

LANL/SNL
TBD
Secure

Exascale Systems

2021-2023

Argonne
Intel/Cray TBD
Open

Frontier

ORNL
TBD
Open

El Capitan

LLNL
TBD
Secure

ORNL’s “Titan” Hybrid System: Cray XK7 with AMD Opteron and NVIDIA Tesla processors

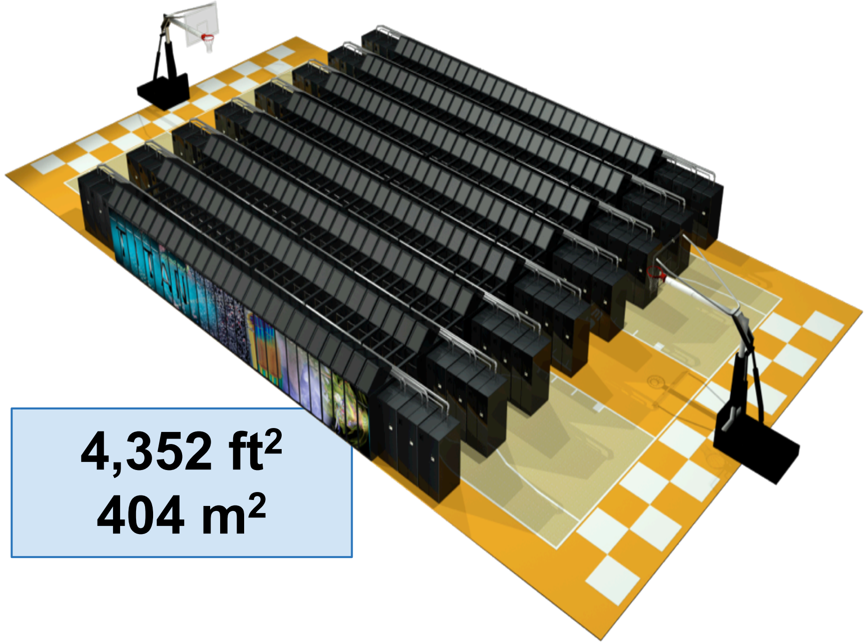

4,352 ft^2
404 m^2

SYSTEM SPECIFICATIONS:

- Peak performance of 27.1 PF (24.5 & 2.6)
- 18,688 Compute Nodes each with:
- 16-Core AMD Opteron CPU (32 GB)
- NVIDIA Tesla “K20x” GPU (6 GB)
- 512 Service and I/O nodes
- 200 Cabinets
- 710 TB total system memory
- Cray Gemini 3D Torus Interconnect

EXASCALE COMPUTING PROJECT

Summit vs Titan

Feature	Summit	Titan
Application Performance	5-10x Titan	Baseline
Number of Nodes	~4,600	18,688
Node performance	>40 TF	1.4 TF
Memory per Node	512 GB (HBM + DDR4)	38GB (GDDR5+DDR3)
NV Memory per Node	1600 GB	0
Node Interconnect	NVLink (5-12x PCIe 3)	PCIe 2
System Interconnect (node injection bandwidth)	Dual Rail EDR-IB (23 GB/s)	Gemini (6.4 GB/s)
Interconnect Topology	Non-blocking Fat Tree	3D Torus
Processors	2 IBM POWER9™ 6 NVIDIA Volta™	AMD Opteron™ NVIDIA Kepler™
File System	250 PB, 2.5 TB/s, GPFS™	32 PB, 1 TB/s, Lustre®
Peak power consumption	15 MW	9 MW

NVIDIA's Tesla V100

- 5,120 CUDA cores (64 on each of 80 SMs)
- 640 NEW Tensor cores (8 on each of 80 SMs)
- 20MB SM RF | 16MB Cache | 16GB HBM2 @ 900 GB/s
- 300 GB/s NVLink
- 7.5 FP64 TFLOPS | 15 FP32 TFLOPS | 120 Tensor TFLOPS

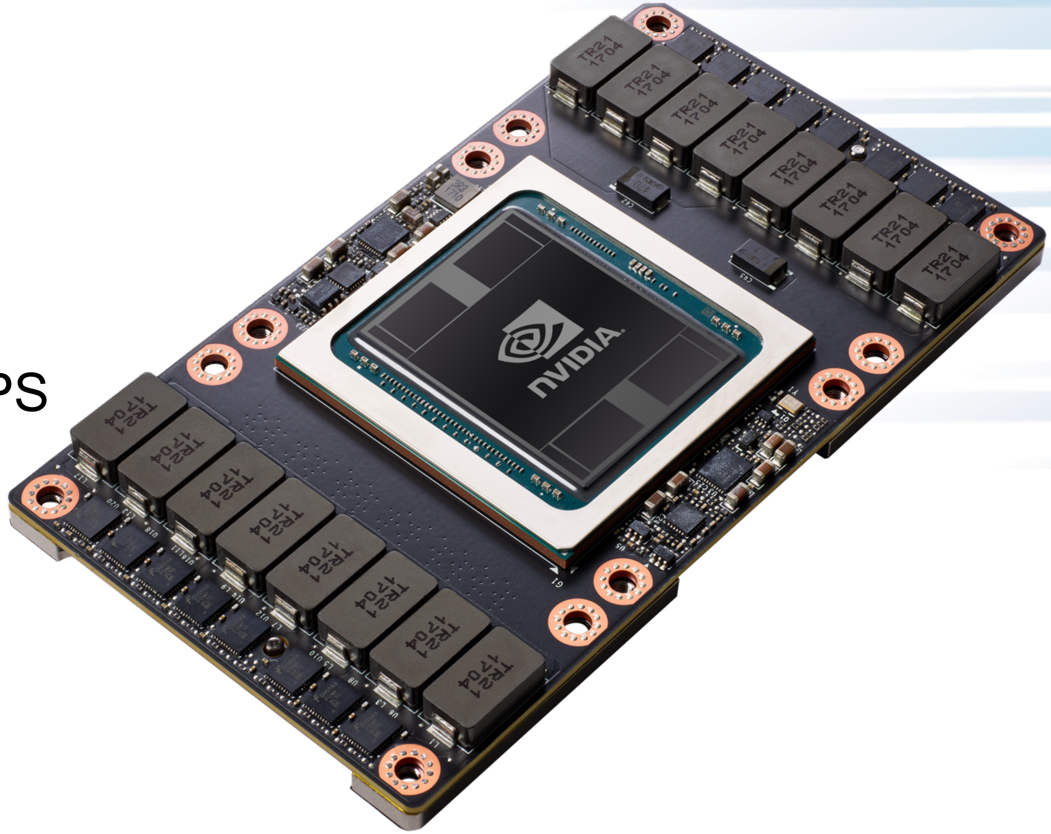

~57 times faster in 64-bit peak floating point performance than the CM-5 I worked on 25 years ago

ECP WBS

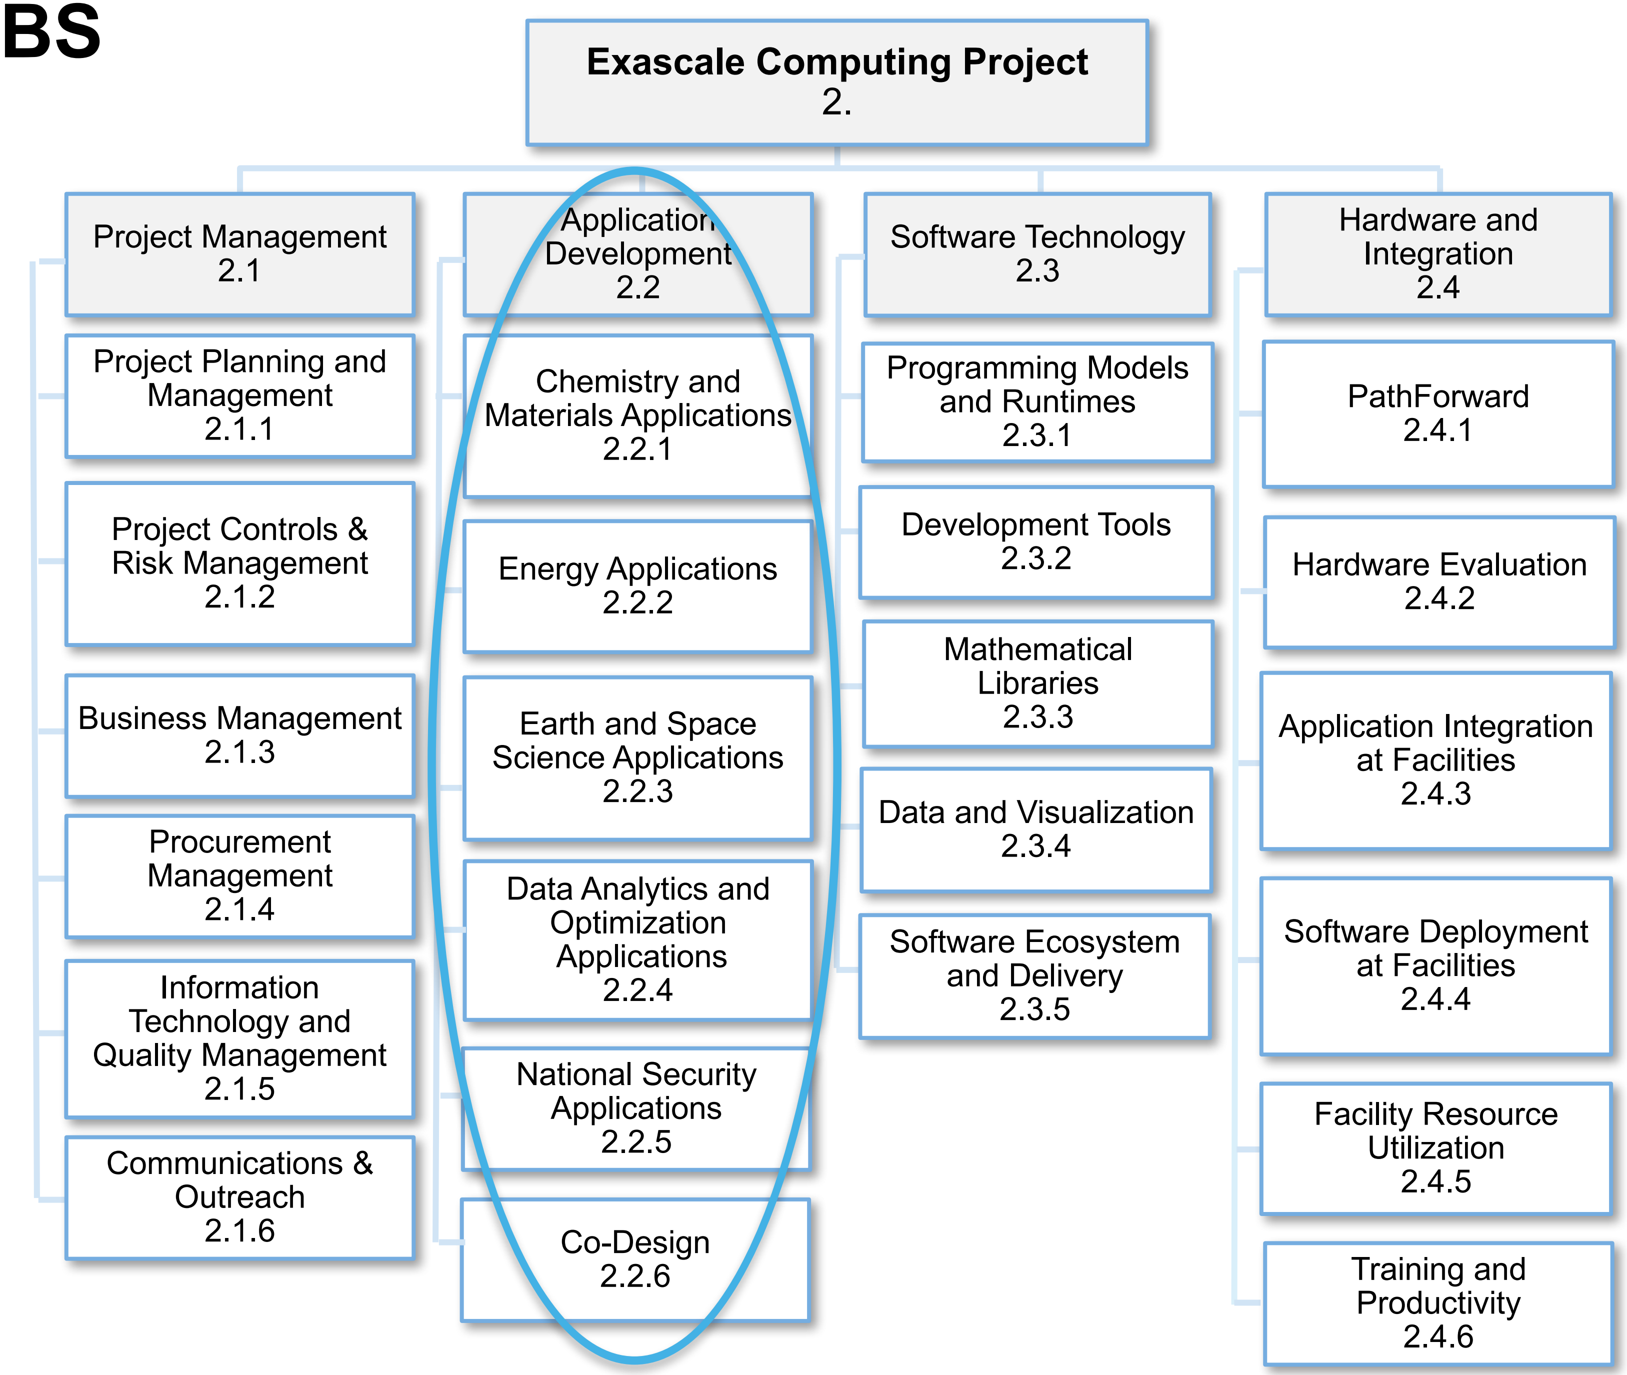

- Exascale Computing Project 2.
  - Project Management 2.1
    - Project Planning and Management 2.1.1
    - Project Controls & Risk Management 2.1.2
    - Business Management 2.1.3
    - Procurement Management 2.1.4
    - Information Technology and Quality Management 2.1.5
    - Communications & Outreach 2.1.6
  - Application Development 2.2
    - Chemistry and Materials Applications 2.2.1
    - Energy Applications 2.2.2
    - Earth and Space Science Applications 2.2.3
    - Data Analytics and Optimization Applications 2.2.4
    - National Security Applications 2.2.5
    - Co-Design 2.2.6
  - Software Technology 2.3
    - Programming Models and Runtimes 2.3.1
    - Development Tools 2.3.2
    - Mathematical Libraries 2.3.3
    - Data and Visualization 2.3.4
    - Software Ecosystem and Delivery 2.3.5
  - Hardware and Integration 2.4
    - PathForward 2.4.1
    - Hardware Evaluation 2.4.2
    - Application Integration at Facilities 2.4.3
    - Software Deployment at Facilities 2.4.4
    - Facility Resource Utilization 2.4.5
    - Training and Productivity 2.4.6

ECP Applications Target National Problems

National security

Stewardship of the nuclear stockpile

Material properties under extreme and energetic conditions

Nuclear forensics

Next-generation virtual flight testing for hypersonic re-entry vehicles

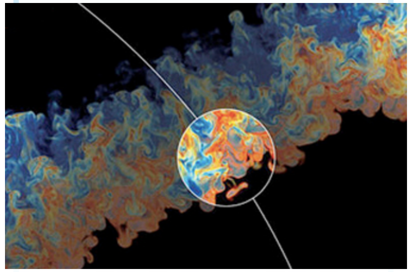

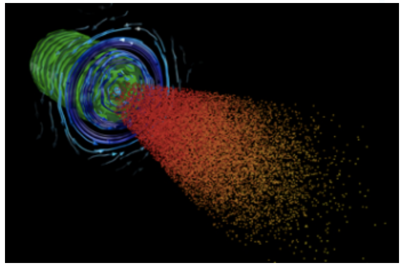

Energy security

Turbine wind plant efficiency

Design and commercialization of SMRs

Nuclear fission and fusion reactor materials design

Subsurface use for carbon capture, petroleum extraction, waste disposal

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

Carbon capture and sequestration scaleup

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Urban planning

Reliable and efficient planning of the power grid

Seismic hazard risk assessment

Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Plasma wakefield accelerator design

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

Find, predict, and control materials and properties

Predict and control stable ITER operational performance

Demystify origin of chemical elements

Earth system

Accurate regional impact assessments in Earth system models

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

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

Health care

Accelerate and translate cancer research

EXASCALE
COMPUTING
PROJECT

Addressing Simulation Gaps a Key Part of ECP Mission Need

Simulation Gap	Needed to Address the Gap	Impact of Gap Remaining
Simplified or incomplete physics	**• Higher fidelity models for all relevant physical phenomena** **• Compute memory and speed necessary to accommodate their numerical solutions and their non-linear coupling requirements**	• Cleaner, more efficient combustion engine designs delayed, or not discovered • Inefficient agriculture & energy production planning - insufficient regional water cycle assessments • Fewer (likely more expensive) options for CO_2 sequestration, petroleum extraction, geothermal energy, due to lack of understanding long-term reservoir-scale behavior • Astrophysics discoveries (origin of elements in the universe, gravity waves) remain elusive
Simulation detail insufficient	**• Larger simulation domain** **• Finer partitioning of the simulation domain** **• Computing speed, memory and I/O to accommodate larger simulation domains with finer partitioning** **• Workflows integrating an advanced technology and software toolset** **• Data streaming methods to steer simulation partition real time**	• Conservative earthquake retrofits more costly and over-designed • Inability to predict and control material properties at the quantum level precludes advances in high temperature superconductivity • Higher power grid operating margins and lost cost savings potential • Wind plant efficiencies lag theoretical energy extraction potential by 20-30% • Delays in scale-up of new chemical looping reactors for clean fossil fuel combustion • Key cosmology and nuclear physics discoveries – dark matter/energy, standard model of particle physics, inflation of the universe – remain elusive
Can only simulate subset of scenarios of interest	**• Robust and fast algorithms for the numerical solution of coupled multi-physics systems that expand limits of applicability** **• Workflow tools to analyze simulation ensembles**	• Continued high rejection rates of additively manufactured metal alloy parts with tight specifications, increasing waste and cost • Tools for retrofitting & improving urban districts remain empirical • Limited ability to influence ITER design decisions and ultimately operations
Uncertainty insufficiently quantified	**• High fidelity in situ data analytic techniques for reliable quantification of simulation uncertainties and sensitivities** **• Workflow tools to analyze simulation ensembles**	• Protracted deployment of small & advanced nuclear reactors • Delays in design of small, low cost, and ultra high intensity plasma wakefield accelerators • Engineering & materials design requires more expensive & time-consuming physical experiments
Unable to intersect to design cycle	**• Improved computer throughput to shorten simulation turn-around times** **• Large ensembles of calculations to enable optimization thru rigorous exploration of design space**	• Deploy materials for extreme environments with a cumbersome make-test cycle rather than by atomistic design • Efficient in silico design of new chemical catalysts not realized
Inadequate analysis and knowledge discovery in big data	**• Scalable AI (deep learning) networks of large size and complexity for efficient training on big datasets.** **• Efficient computational workflows seamlessly integrating simulation, data analytics, and big datasets**	• New cancer biology treatment options missed or delayed due to unrealized understanding of precision oncology • Analysis and reduction of data deluge from experimental science facilities requires weeks to months • Manufacture of new products and chemicals delayed or missed because microbiome DNA sequencing unable to keep up with available data • Inability to quantify uncertainties via a nexus of simulation and experimental facility data

Application Co-Design (CD)

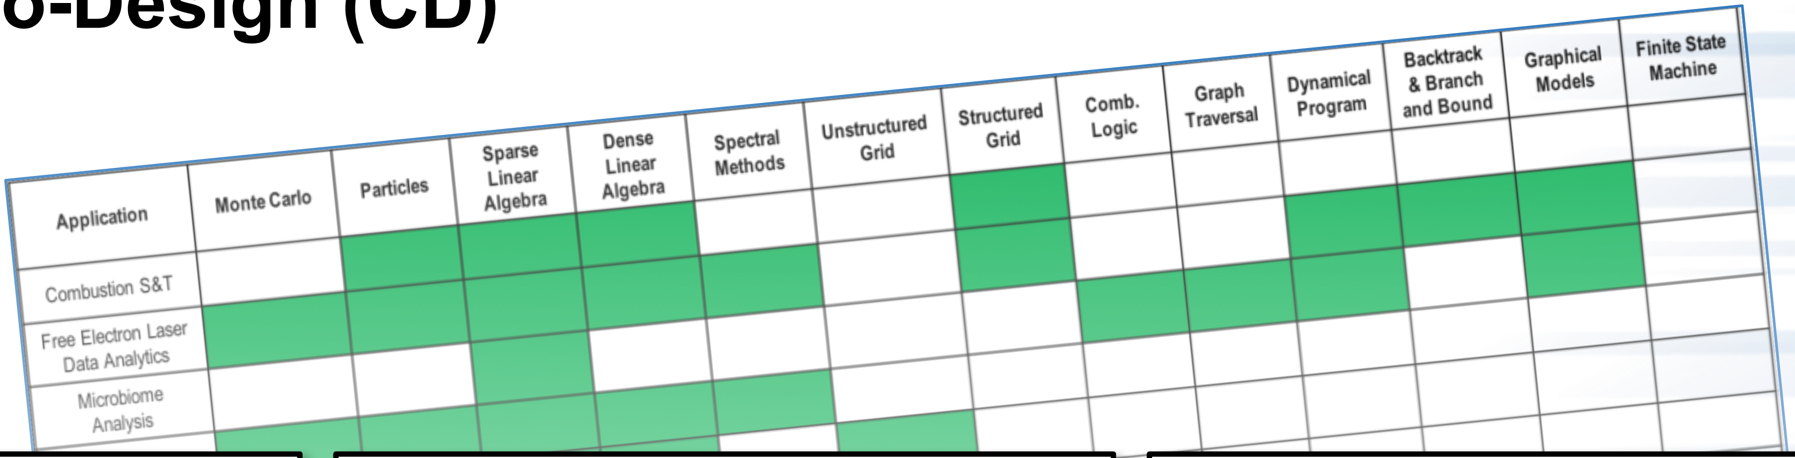

Application	Monte Carlo	Particles	Sparse Linear Algebra	Dense Linear Algebra	Spectral Methods	Unstructured Grid	Structured Grid	Comb. Logic	Graph Traversal	Dynamical Program	Backtrack & Branch and Bound	Graphical Models	Finite State Machine
Combustion S&T		🟩	🟩	🟩			🟩						
Free Electron Laser Data Analytics	🟩	🟩	🟩	🟩	🟩		🟩			🟩	🟩	🟩	
Microbiome Analysis			🟩					🟩	🟩	🟩		🟩	

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

Executed by several CD Centers focusing 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

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

Algorithmic methods that capture a common pattern of computation and communication

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

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

Survey of Application Motifs

Application	Monte Carlo	Particles	Sparse Linear Algebra	Dense Linear Algebra	Spectral Methods	Unstructured Grid	Structured Grid	Comb. Logic	Graph Traversal	Dynamical Program	Backtrack & Branch and Bound	Graphical Models	Finite State Machine
Cosmology		🟩	🟩		🟩		🟩		🟩				
Subsurface			🟩			🟩	🟩						
Materials (QMC)	🟩	🟩	🟩	🟩	🟩					🟩			
Additive Manufacturing		🟩	🟩	🟩	🟩	🟩				🟩			
Chemistry for Catalysts & Plants	🟩		🟩	🟩	🟩		🟩		🟩				
Climate Science		🟩	🟩			🟩	🟩						
Precision Medicine Machine Learning			🟩	🟩				🟩	🟩	🟩	🟩	🟩	
QCD for Standard Model Validation	🟩		🟩	🟩	🟩		🟩						
Accelerator Physics		🟩	🟩		🟩		🟩						
Nuclear Binding and Heavy Elements	🟩		🟩	🟩									
MD for Materials Discovery & Design		🟩	🟩	🟩	🟩		🟩						
Magnetically Confined Fusion		🟩	🟩		🟩	🟩	🟩						

Survey of Application Motifs

Application	Monte Carlo	Particles	Sparse Linear Algebra	Dense Linear Algebra	Spectral Methods	Unstructured Grid	Structured Grid	Comb. Logic	Graph Traversal	Dynamical Program	Backtrack & Branch and Bound	Graphical Models	Finite State Machine
Combustion S&T		🟩	🟩	🟩			🟩						
Free Electron Laser Data Analytics	🟩	🟩	🟩	🟩	🟩		🟩			🟩	🟩	🟩	
Microbiome Analysis			🟩					🟩	🟩	🟩		🟩	
Catalyst Design	🟩	🟩	🟩	🟩	🟩								
Wind Plant Flow Physics			🟩	🟩		🟩							
SMR Core Physics	🟩	🟩	🟩	🟩	🟩	🟩	🟩						
Next-Gen Engine Design		🟩	🟩		🟩	🟩							
Urban Systems			🟩		🟩	🟩	🟩						
Seismic Hazard Assessment			🟩		🟩	🟩	🟩			🟩			
Systems Biology	🟩		🟩	🟩				🟩	🟩	🟩	🟩	🟩	
Biological Neutron Science		🟩	🟩	🟩	🟩		🟩		🟩		🟩		
Power Grid Dynamics	🟩		🟩	🟩					🟩	🟩	🟩		

Survey of Application Motifs

Application	Monte Carlo	Particles	Sparse Linear Algebra	Dense Linear Algebra	Spectral Methods	Unstructured Grid	Structured Grid	Comb. Logic	Graph Traversal	Dynamical Program	Backtrack & Branch and Bound	Graphical Models	Finite State Machine
Stellar Explosions	🟩	🟩	🟩	🟩			🟩						
Excited State Material Properties				🟩	🟩								
Light Sources			🟩	🟩	🟩	🟩				🟩			🟩
Materials for Energy Conversion/Storage		🟩	🟩	🟩	🟩								
Hypersonic Vehicle Design	🟩	🟩	🟩	🟩		🟩							
Multiphase Energy Conversion Devices		🟩	🟩				🟩						

AMReX Co-Design Center Project Goals

- Develop and deploy software to support block-structured adaptive mesh refinement on exascale architectures
  - Core AMR functionality
  - Particles coupled to AMR meshes
  - Embedded boundary 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 AMReX when appropriate
- Interact with hardware technologies / vendors

John Bell (PI, LBNL)

Applications Using AMReX

- A wide range ECP applications are using AMReX to provide parallel AMR capability
  - Accelerators -- **WarpX**
  - Combustion – **PeleC, PeleLM**
  - Multiphase flow -- **MFIX-Exa**
  - Cosmology -- **Nyx**
  - Astrophyics – **Castro**
- Disparate applications have a number of common elements:

	Particles	ODE's	Linear Solvers	EB
Combustion	X	X	X	X
Multiphase	X		X	X
Cosmology	X	X	X	
Astrophysics	X	X	X	
Accelerators	X			

John Bell (PI, LBNL)

Release of New Infrastructure for Block-Structured Adaptive Mesh Refinement Developed by AMReX Co-Design Center

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
  - Linear solvers for AMR grids
  - Embedded boundary representation of complex geometry
  - Performance portability for different architectures
- Current milestone focused on:
  - Establishing support for core AMR functionality including documentation and tutorial examples

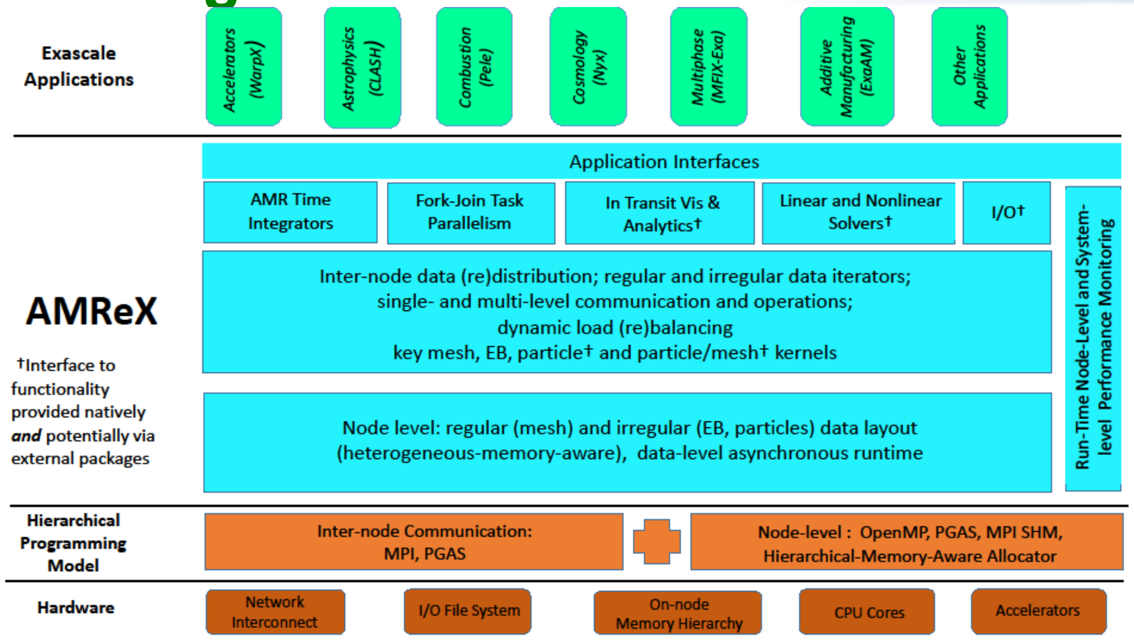

Exascale Applications

Accelerators (WarpX)
Astrophysics (CLASH)
Combustion (Pele)
Cosmology (Nyx)
Multiphase (MFIX-Exa)
Additive Manufacturing (ExaAM)
Other Applications

AMReX

†Interface to functionality provided natively *and* potentially via external packages

Application Interfaces

AMR Time Integrators
Fork-Join Task Parallelism
In Transit Vis & Analytics†
Linear and Nonlinear Solvers†
I/O†

Inter-node data (re)distribution; regular and irregular data iterators;
single- and multi-level communication and operations;
dynamic load (re)balancing
key mesh, EB, particle† and particle/mesh† kernels

Node level: regular (mesh) and irregular (EB, particles) data layout
(heterogeneous-memory-aware), data-level asynchronous runtime

Run-Time Node-Level and System-level Performance Monitoring

Hierarchical Programming Model

Inter-node Communication:
MPI, PGAS

Node-level : OpenMP, PGAS, MPI SHM,
Hierarchical-Memory-Aware Allocator

Hardware

Network Interconnect
I/O File System
On-node Memory Hierarchy
CPU Cores
Accelerators

Impact

- Established a next-generation framework for developing block-structured adaptive mesh refinement algorithms for current and emerging architectures
- Provides a common framework for multiple ECP applications that use AMR
- Provides a common focal point for software technology, hardware technology and vendors to leverage activities over multiple applications
- Broad constituency within Office of Science and NNSA

Project Accomplishment

- Designed and implemented basic mesh data structures and iterator interface to support tiled execution model
- Implemented particle data containers that provide alternative approaches to representing particle data to support different application requirements
- Developed initial documentation and tutorial examples
- AMReX code framework publicly released on github

Deliverables: AMReX software available at https://github.com/AMReX-Codes/amrex

Co-design Motifs

- PDE-based simulations on **unstructured grids**
- **high-order** and **spectral** finite elements

✓ *any order space on any order mesh* ✓ *curved meshes,*
✓ *unstructured AMR* ✓ *optimized low-order support*

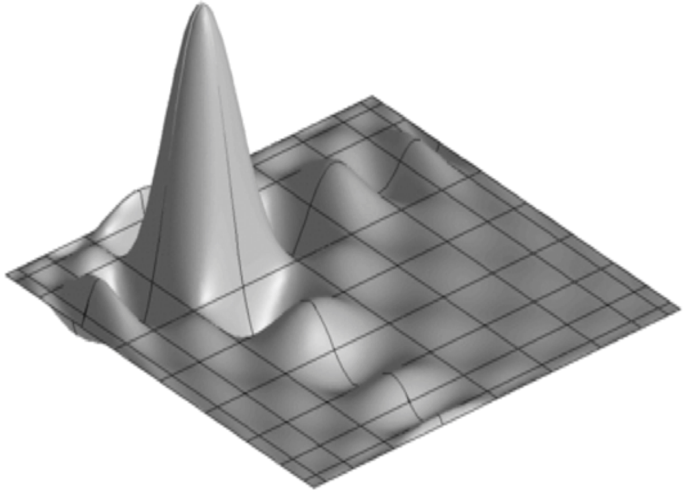

10th order basis function

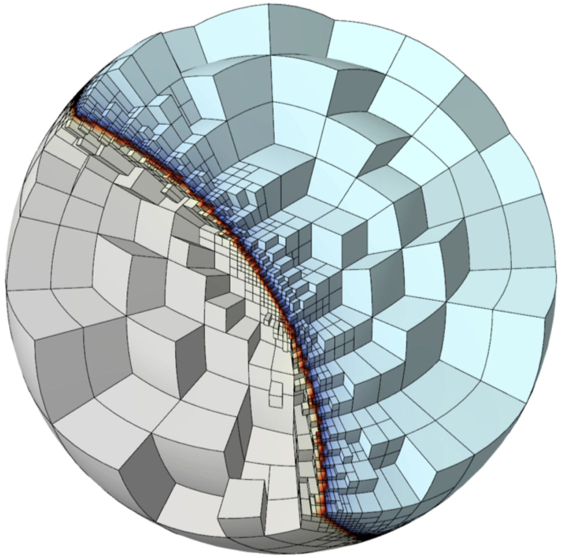

non-conforming AMR, 2nd order mesh

incompressible SEM / SC

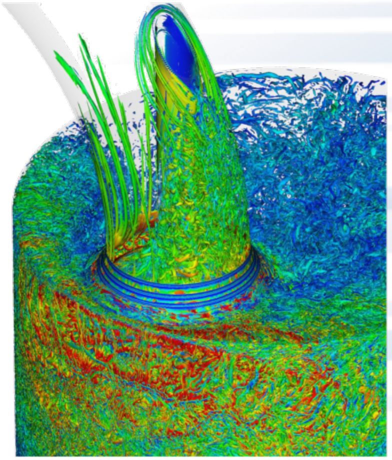

6th order DNS turbulence (Nek)

compressible FEM / NNSA

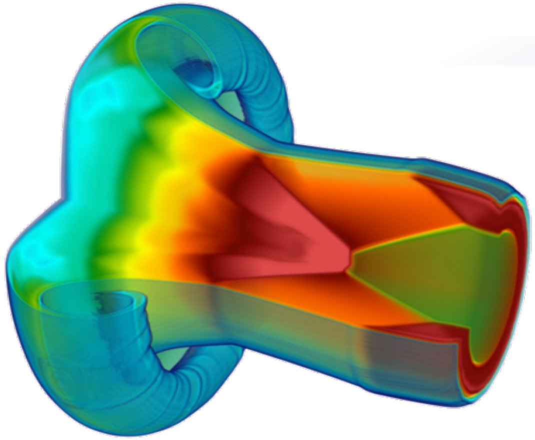

2nd order compressible shock hydro (MFEM)

Tzanio Kolev (PI, LLNL)

CEED Software Products

CEED's library model enables ECP apps to easily take advantage of the new discretization technologies

- state-of-the art CEED **discretization libraries**
  - ✓ *better exploit the hardware to deliver significant performance gain over conventional methods*
  - ✓ *based on MFEM/Nek, low & high-level APIs*

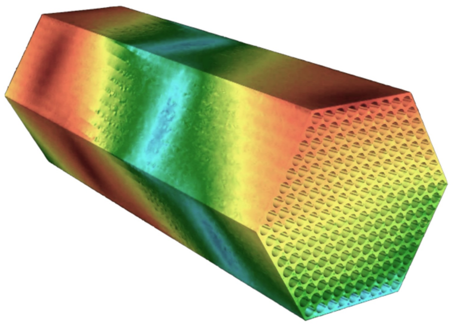

nek5000.mcs.anl.gov

High-performance spectral elements

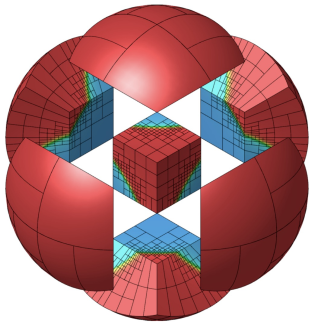

mfem.org

Scalable high-order finite elements

CEED's proxies and general purpose libs target ECP vendors, STs, broader community

- **Ceedlings** - CEED kernels, bake-off probs & miniapps
  - ✓ *main tools to engage vendors & external projects*
- CEED **broadly applicable libraries**

MAGMA

icl.cs.utk.edu/magma

LAPACK for GPUs, multi/many-core

libocca.org

Lightweight performance portability

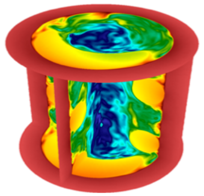

GSLIB

PUMi

Parallel Unstructured Mesh Infrastructure

Holmes

PETSc

Main deliverable: all CEED software freely available on GitHub at https://github.com/CEED
New releases: mfem-3.3, gslib, Laghos and NekCEM ceedling, …

Efficient Operator Evaluation

Mathematical framework for our computational motif

$$\boxed{A = P^T G^T B^T D B G P}$$

Finite element operators (matrices) can be decomposed into **parallel**, **mesh topology**, **basis**, and **geometry/physics** components:

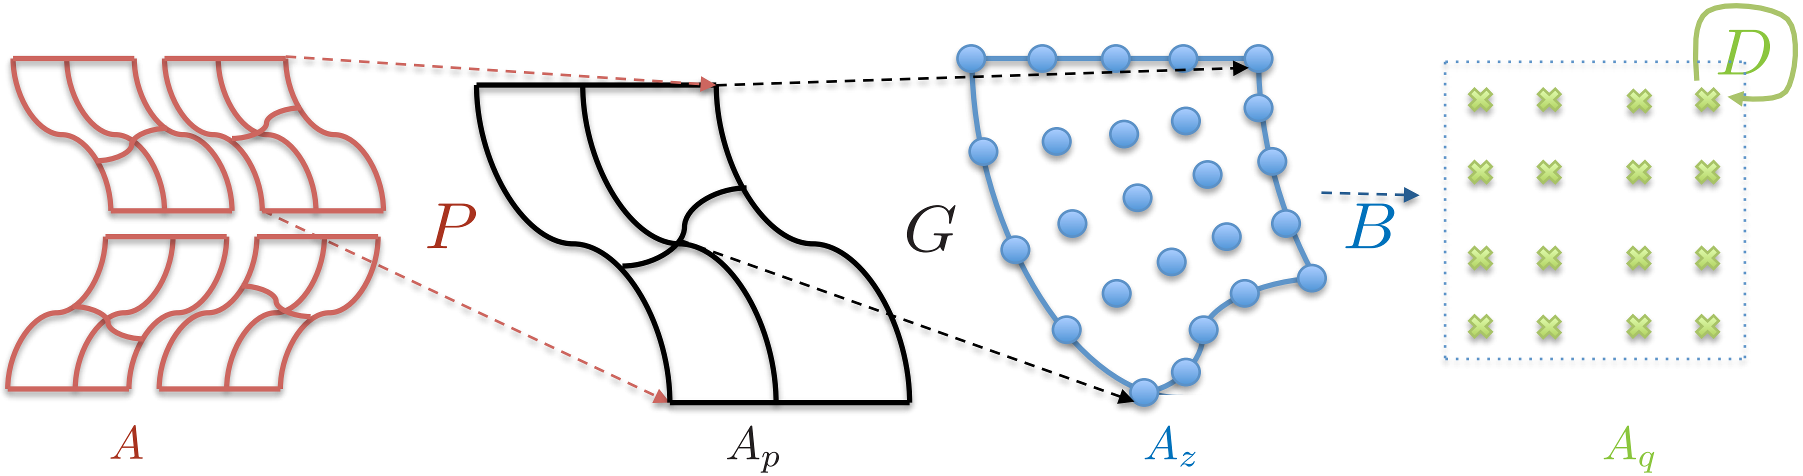

- ***partial assembly*** = store only D, evaluate B (tensor-product structure)
- better representation than A: *optimal memory*, *near-optimal FLOPs, hierarchical parallelism*
- *better mathematical algorithms trump code optimizations*

Batched Computing Technology

- Matrix-free basis evaluation needs efficient tensor contractions, e.g.,

$$C_{i1,i2,i3} = \sum_k A_{k,i1} B_{k,i2,i3}$$

- **CEED/MAGMA** designed batched methods to split the computation in many small high-intensity GEMMs, grouped together (batched) for efficient execution:

Batch_{ $C_{i3} = A^T B_{i3}$, for range of i3 }

- Developed techniques needed for autotuning, code inlining, code generation (reshapes, etc.), algorithmic variants for different architectures.
- Achieve 90+% of theoretically derived peaks.
- Significantly outperform vendor libraries.
- Released through MAGMA.

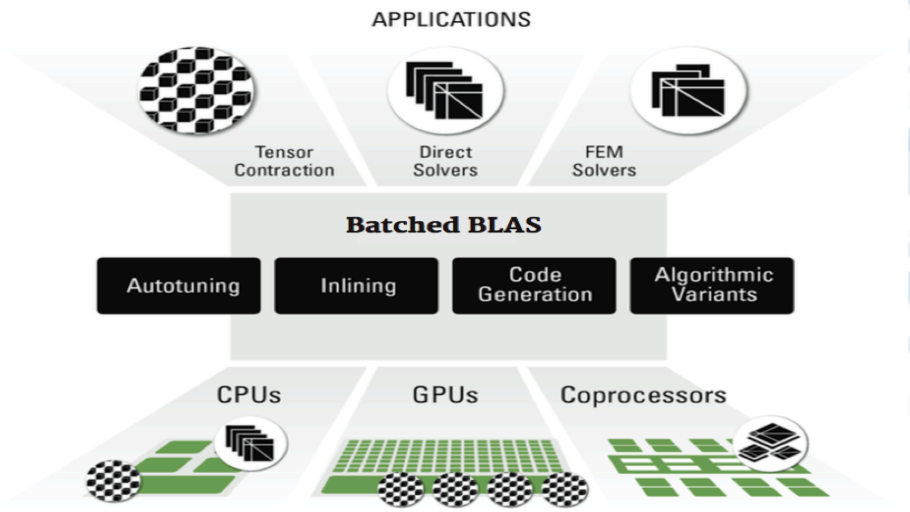

Batched DGEMMs on GPU
(P100, 100K)

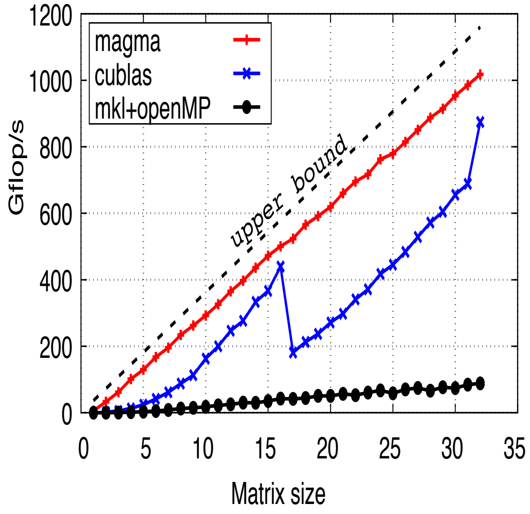

Batched DGEMMs on ARM
(Tegra X1 : 4-core Cortex A57)

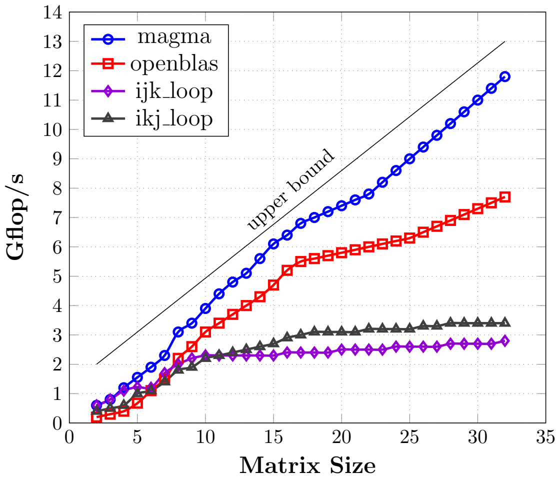

CEED-MS1

Engaged First Wave ECP/CEED Applications

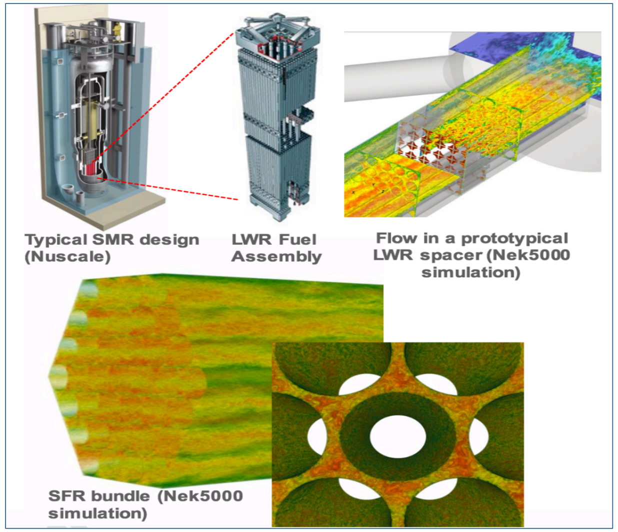

ExaSMR: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors.

- ☐ Goal: Coupled thermal-hydraulics/neutronics analysis
- ☐ Require scalable simulations up to one billion spectral elements and trillions of grid points; complex mesh structures, external and internal fluid mixing, advanced RANS/LES models, efficient time stepping, and improved parallel mesh I/O support.
- ☐ *Application liaison: Elia Merzari (ANL)*
- ☐ **Shared milestone: Nek5000 porting on GPUs, KNL performance**

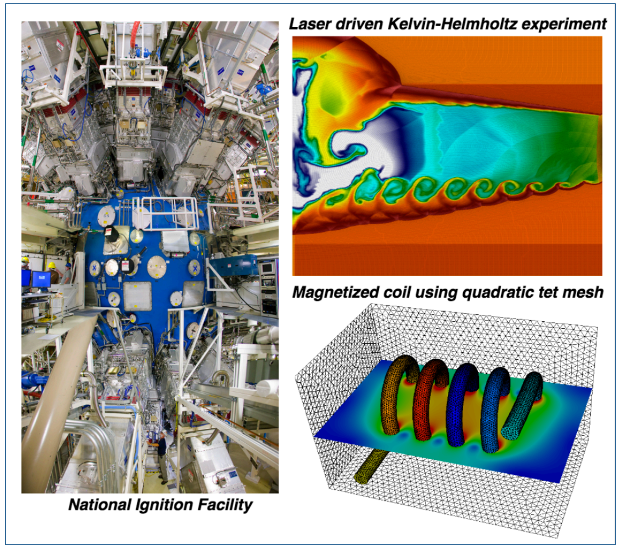

MARBL: Next-gen Multi-physics Simulation Code.

- ☐ Goal: Modular physics simulation capabilities with increased performance portability and flexibility.
- ☐ Require single-fluid multi-material hydrodynamics and radiation/magnetic diffusion simulation; inertial confined fusion, pulsed power experiments, and equation of state/material strength analysis
- ☐ *Application liaison: Vladimir Tomov (LLNL)*
- ☐ **Shared milestone: Developing Lagrangian hydro miniapp**

1st, 2nd and 3rd Wave of Applications

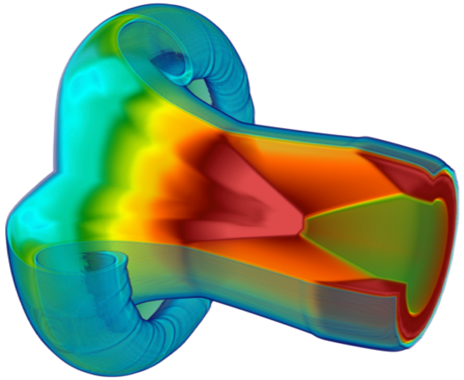

Compressible flow (MARBL)

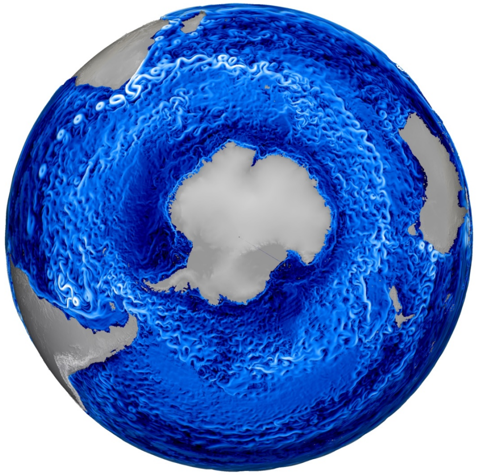

Climate (ACME)

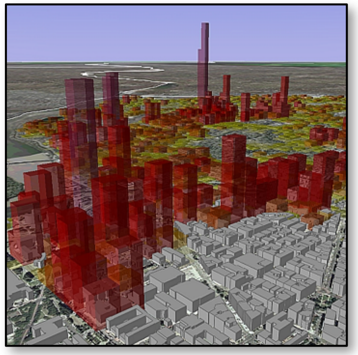

Urban systems (Urban)

Wind Energy (ExaWind)

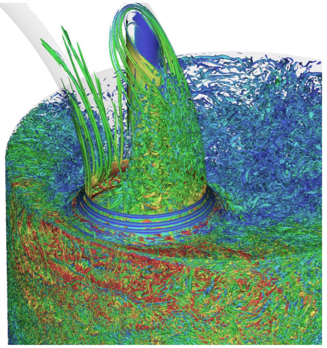

Combustion (Nek5000)

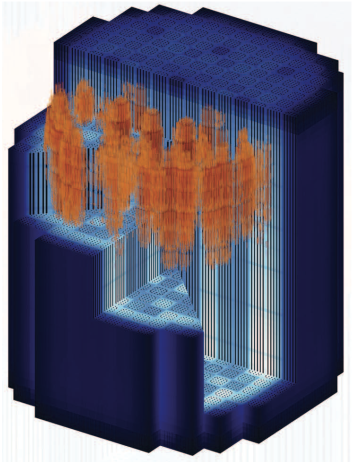

Modular Nuclear Reactors (SMRs)

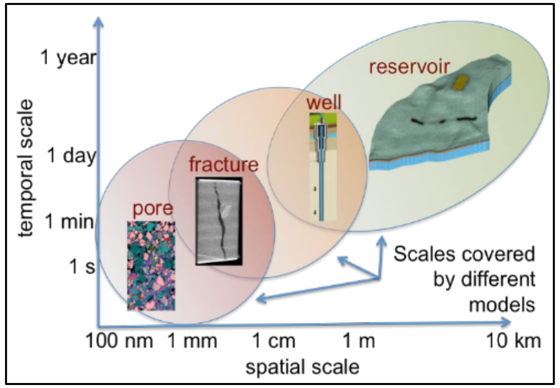

Subsurface (GEOS)

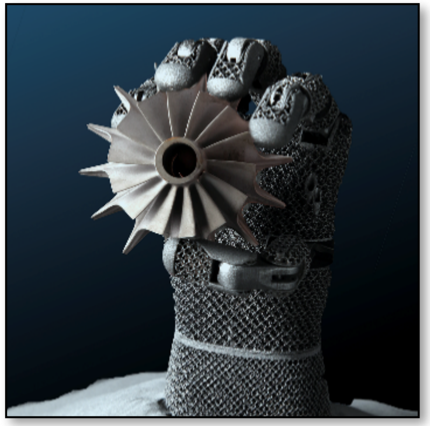

Additive Manufacturing (ExaAM)

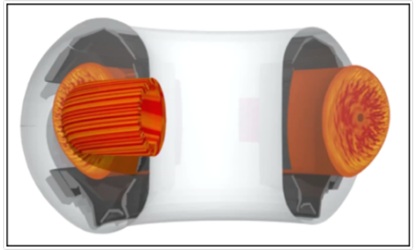

Magnetic Fusion (WDMApp)

Benefits

- Backed by well-developed theory.
- Naturally support unstructured and curvilinear grids.
  - Triangular, quadrilateral, tetrahedral and hexahedral; volume and surface meshes
  - Unstructured AMR
  - NURBS geometries and discretizations
- Increased accuracy for smooth problems
- Sub-element modeling for problems with shocks
- Bridge unstructured/structured grids
- Bridge sparse/dense linear algebra
- FLOPs/bytes increase with the order
- Demonstrated match for many applications: compressible, incompressible flow, ...

Benefits – Coarse Meshes

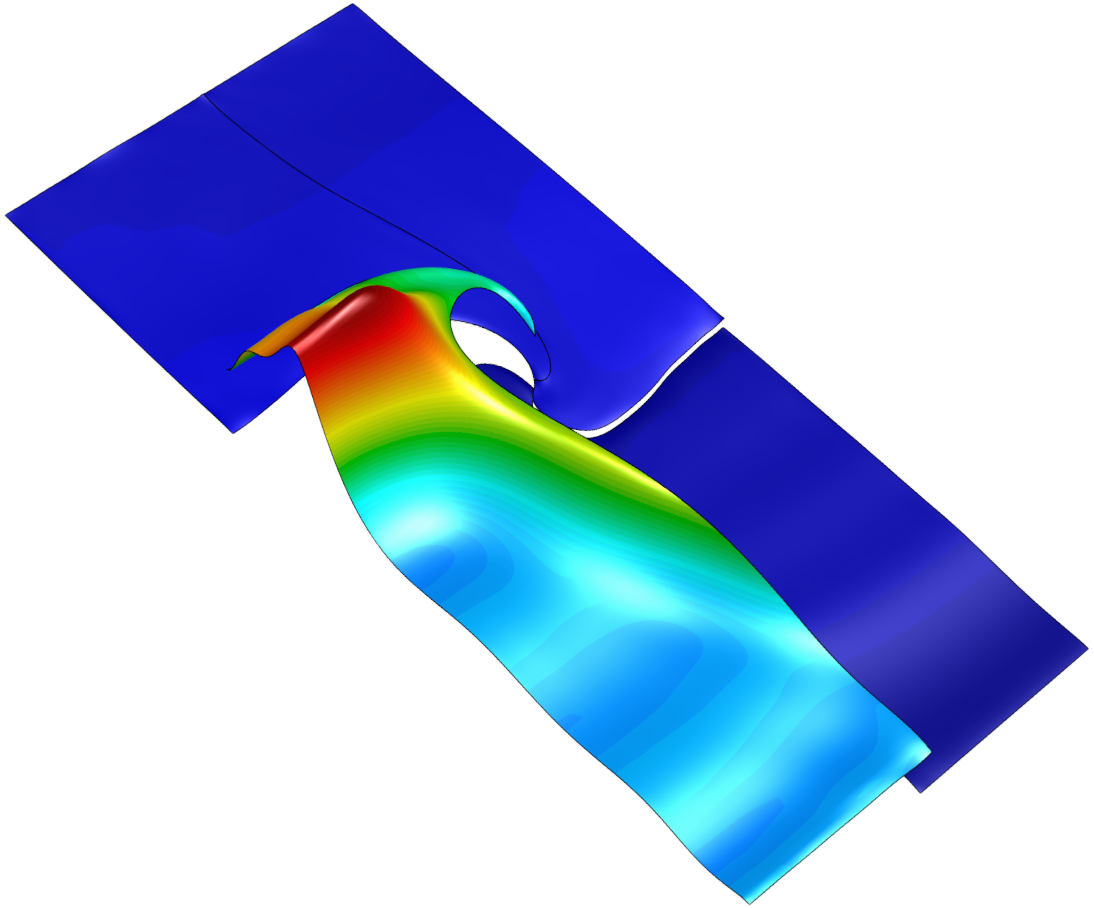

Shock triple-point interaction (4 elements)

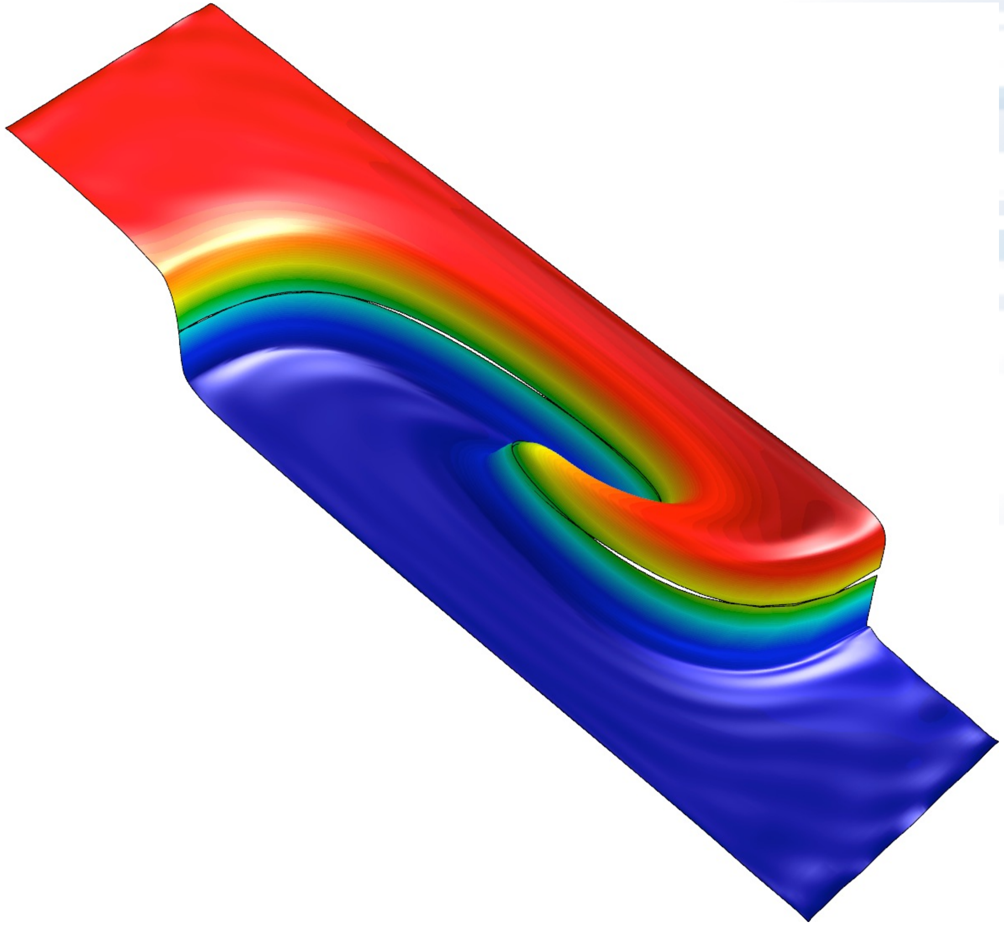

Smooth RT instability (2 elements)

Benefits – Unstructured AMR

- We support both conforming (simpices) and non-conforming AMR (quad/hexes).
- Powerful and general:
  - any (high-order) finite element space on any (high-order) curved mesh
  - anisotropic refinement
  - serial and parallel

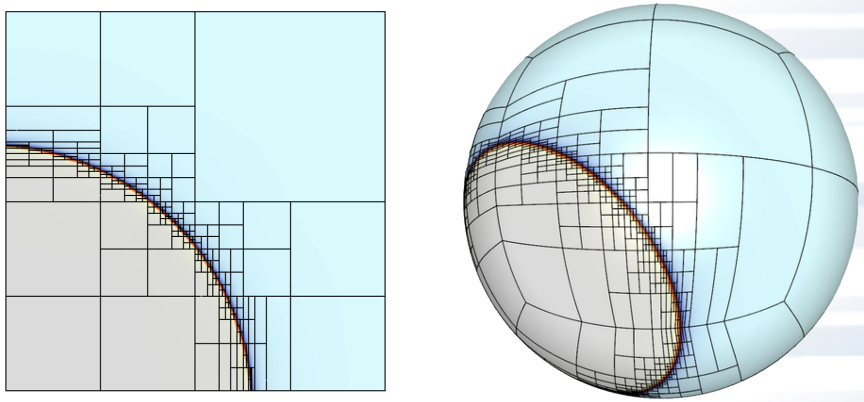

Adaptation to shock-like solution in 2D and 3D

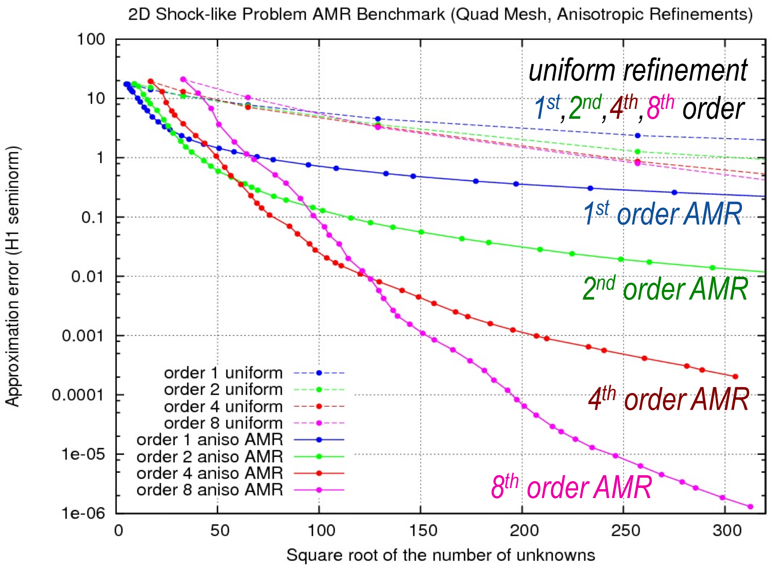

2D AMR error improves significantly with order

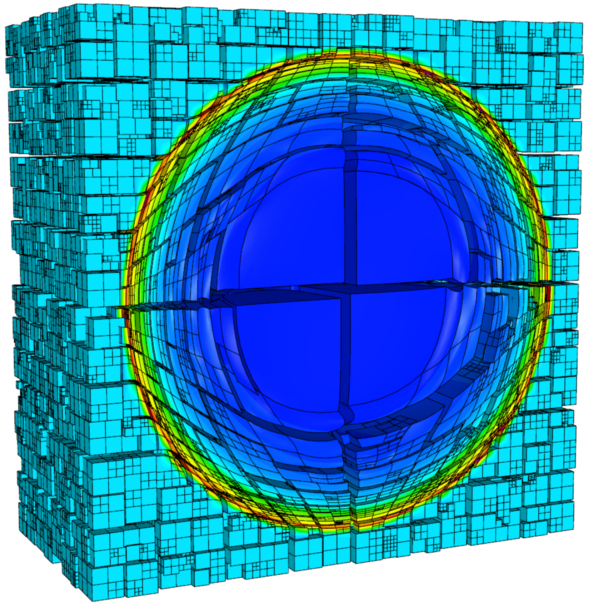

Shock propagation without mesh artifacts, 4096 cores

Finite volume/embedded boundary approach **to complex geometry/interfaces is a cut cell method**

Finite volume, embedded boundary (EB) method where irregular boundaries are represented intersecting domain with Cartesian grid:

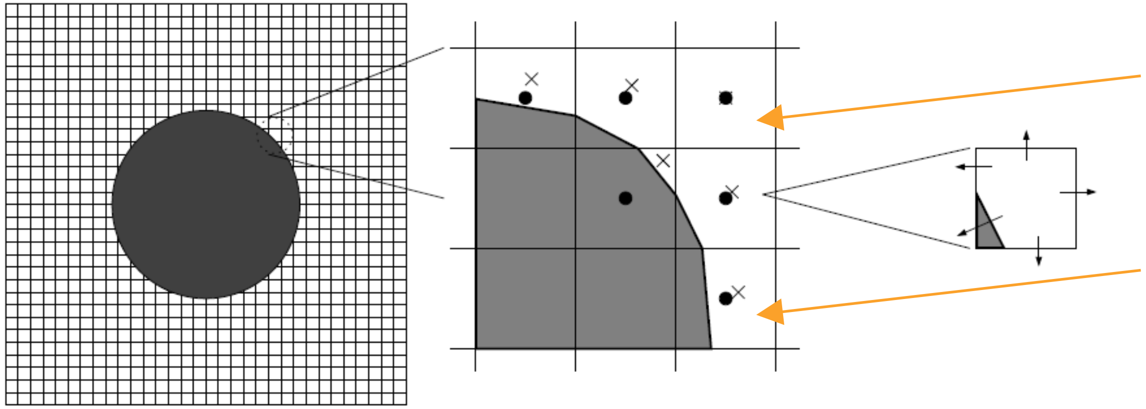

Majority of cells are regular

Special stencils used in small subset of irregular cut cells

Incompressible Navier-Stokes flow

$$\frac{\partial U}{\partial t} + U \cdot \nabla U = -\frac{1}{\rho}\nabla p + \nu \Delta U$$

$$\nabla \cdot U = 0$$

Fixed time divergence theorem

$$D(\vec{F})_v = \frac{1}{h\kappa_v}\left(\left(\sum_{d=1}^{D}\left(\alpha_{i+\frac{1}{2}\hat{e}^d}\widetilde{F^d}_{i+\frac{1}{2}\hat{e}^d} - \alpha_{i-\frac{1}{2}\hat{e}^d}\widetilde{F^d}_{i-\frac{1}{2}\hat{e}^d}\right) + \alpha_v^B F_v^B\right)\right.$$

Small cell problem avoided by volume-weighted scheme
(Trebotich and Graves, CAMCoS, 2015)

- EB is amenable to gridding of complex geometry and moving interfaces
- Fast stencil operations in cut cells

Moving embedded boundary method is based on level set method and conservative volume of fluid

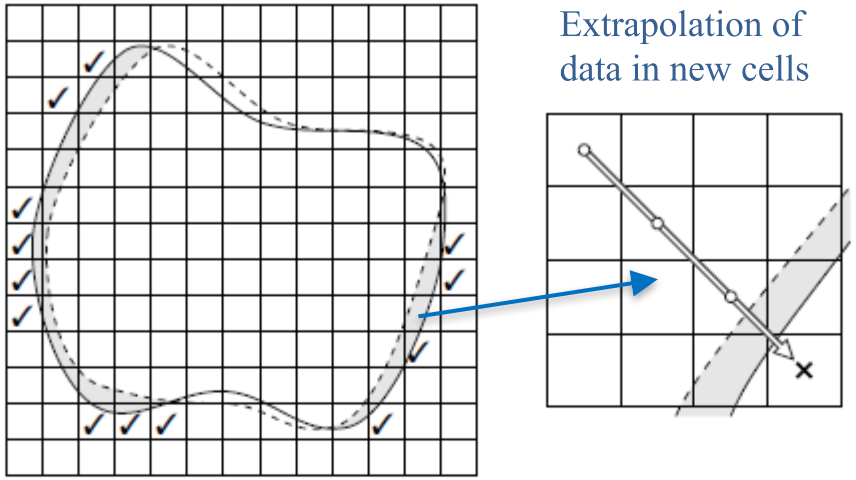

Extrapolation of data in new cells

- Consistent approach for interface tracking
  - fluid-solid interface (e.g., reactions)
  - fluid-fluid interface (e.g., multiphase)
- Space-time divergence theorem
  - inherently conservative (FV)
  - sharp interface (LS)

Space-time divergence theorem

$$
\begin{aligned}
0 &= \int_{t^n}^{t^{n+1}} dt \int_{\Omega_i(t)} dV \left(\frac{\partial}{\partial t}, \nabla\right) \cdot (u, F) \\
&\approx \kappa_i^{n+1} h^D u_i^{n+1} - \kappa_i^n h^D u_i^n \\
&\quad + \Delta t\, h^{D-1} \sum_d^D \left(\alpha_{i+1/2e_d} F_{d,i+1/2e_d} - \alpha_{i-1/2e_d} F_{d,i-1/2e_d}\right) + A_{i,\mathrm{EB}} n_{i,EB} \cdot (u, F)_{i,\mathrm{EB}},
\end{aligned}
$$

Carl Steefel (PI, LBNL), David Trebotich (LBNL)

A general capability for high performance CFD in complex geometries

Chombo-based CFD solver

Incompressible Navier-Stokes (INS) solver in complex geometries

- Embedded boundary (EB) method
  - Structured Cartesian grid
  - Sharp interface
  - Finite volume approach
  - Robust for wide range of flows
  - Dynamic local refinement (AMR)
  - AMG elliptic solvers (PETSc)
  - DNS from image data
  - HDF5 I/O
  - VisIt visualization/analytics

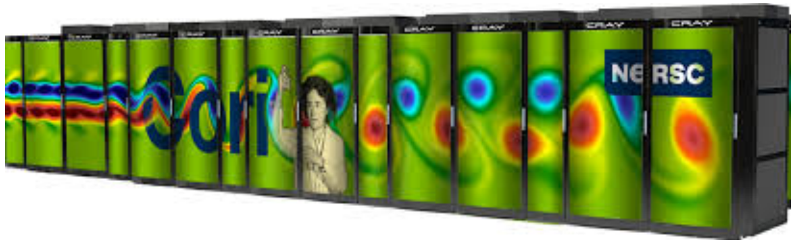


2D

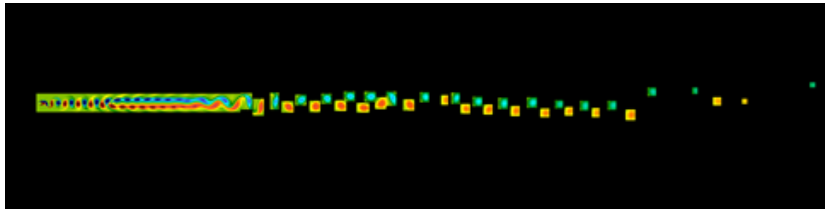
AMR

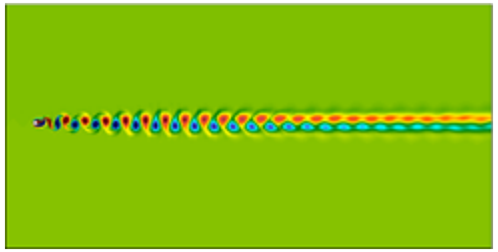
3D

Carl Steefel (PI, LBNL), David Trebotich (LBNL)

Resolved flow in fractured Marcellus shale

“Big Kahuna”

- 0.12 porosity, 0.00247242 cm^2 RSA
- 104 μm x 97 μm x 65 μm sample
- 48 nm image resolution (FIB/SEM)
- 24 nm grid resolution ⇒ O(TBs)/plot
- domain decomposes into >1M boxes (32^3 grid cells each)
- prune to 350,000 boxes/ KNL cores
- *100B degrees of freedom*

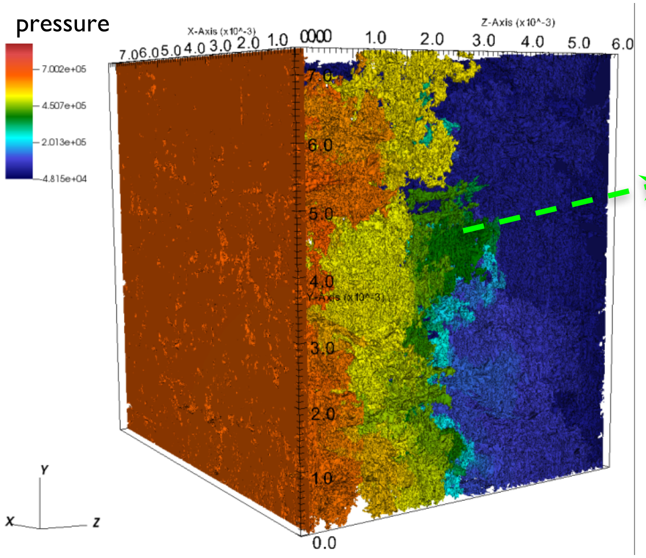

Computed permeability 7.8x10^{-17} m^2

74 μm x 74 μm x 65 μm cut for visualization

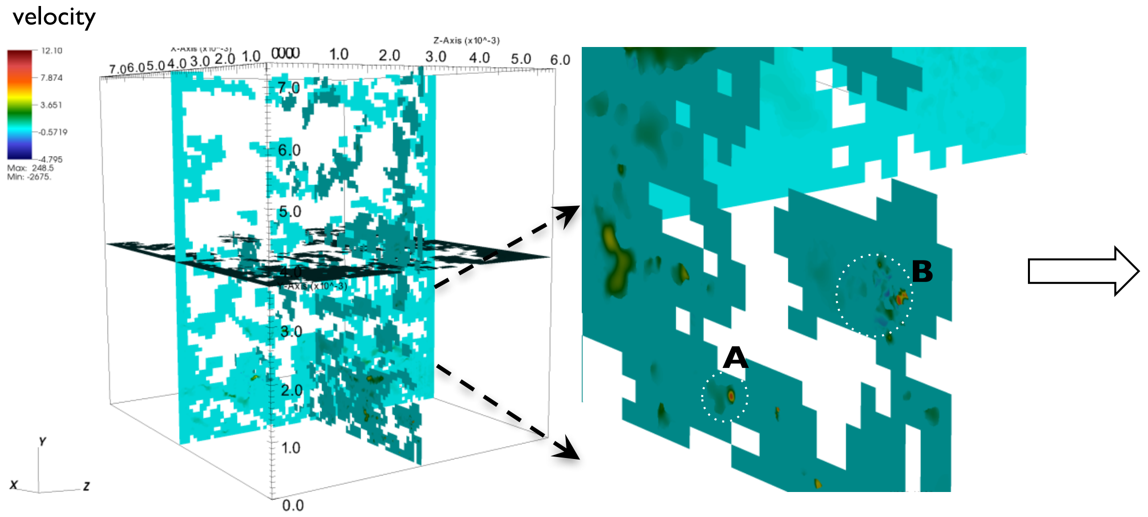

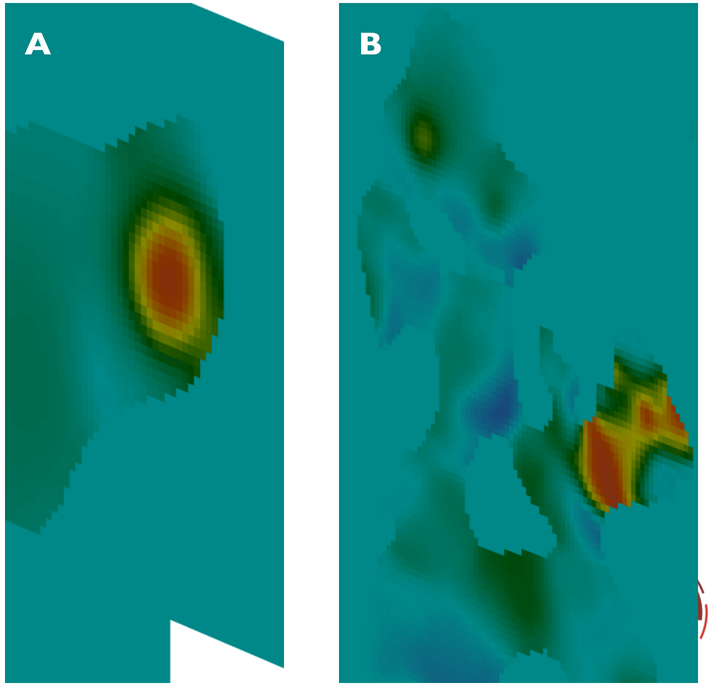

Carl Steefel (PI, LBNL), David Trebotich (LBNL)

Transforming Combustion Science and Technology with Exascale Simulations

Exascale Challenge Problem

- First-principles (DNS) and near-first principles (DNS/LES hybrids) AMR-based technologies to advance understanding of fundamental turbulence-chemistry interactions in device relevant conditions
- High-fidelity geometrically faithful simulation of the relevant in-cylinder processes in a low temperature reactivity controlled compression ignition (RCCI) internal combustion engine that is more thermodynamically favorable then existing engines, with potential for groundbreaking efficiencies yet limiting pollutant formation
- Also demonstrate technology with hybrid DNS/LES simulation of a sector from a gas turbine for power generation burning hydrogen enriched natural gas.
- High-fidelity models will account for turbulence, mixing, spray vaporization, low-temperature ignition, flame propagation, soot/radiation, non-ideal fluids

Applications, Motifs, PM, HPC systems

- **Physical Models**: turbulence, chemical kinetics, polydisperse spray, soot, thermal radiation, non-ideal fluids
- **Codes**: A new suite of compressible and low-Mach block-structured AMR codes, PeleC and PeleLM, with geometry capabilitiy will be developed based on petascale DNS codes: S3D, LMC, and the BoxLiB AMR framework.
- **Motifs:** Partial differential equations, AMR, method-of-lines, sparse iterative linear solvers, finite volume, spectral deferred correction impliciit-explicit schemes, ODE's, linear algebra
- **PM's:** C++, Fortran, Python, MPI, PGAS, OpenMP (gpu support in 2 yrs), Kokkos (particles), Legion (in S3D as proxy for physics in Pele)
- **HPC:** Cori, Edison, Titan, Summit

Risks and Challenges

- Risks with application performance portability –consider alternative PM's for tasking models for particles and for thermochemistry
- Chemistry DSL SINGE for performance portability, need NRE software PathForward for nVidia to develop DSL backend to emit fast kernels for GPU's
- Uncertainties in chemistry/thermo/transport in relevant turbulent environments – perform embedded UQ in Pele or S3D to assess uncertainties of rate constants
- Nonequilibrium chemistry effects on ignition delay – embedd molecular dynamics/DSMC to study competing thermalizing collisions with collision partners determined from local heterogeneity
- Semi-structured algebraic multi-grid may be a missing element
- Sensitivity of ignition to spray atomization in Pele

Milestones

FY17:(Completed) Released of baseline code suite (PeleC, PeleLM, PelePhysics) ADSE14-1; Implemented capability for combustion of non-ideal gas mixtures ADSE14-2; Developed Lagrangian spray capability in PeleC code and demonstrate it in a non-AMR setting ADSE14-3; Developed capability to simulate complex geometry using embedded boundary (cut cell) ADSE14-4; Developed capability to automatically predict the thermochemical properties for combustion of syngas, natural gas, and butane ADSE14-57; Performance assessment of task-based physics ADSE14-5.

FY18-20: (Planned) Lagrangian spray implementation with AMR in Pele; Implement soot l and thermal radiation models, Benchmark linear solvers on Summit, Perform end of year Pele performance benchmarks, simulate turbulent flame with complex geometry; Simulate gas turbine sector with real geometry

PI: Jacqueline Chen (SNL), Partner Sites (LBL, NREL, ORNL,ANL)

PeleC Embedded Boundary Capability

PI: Jackie Chen, SNL
ADSE14-3 Owner: Ray Grout (NREL)
Contributors: Jon Rood (NREL), John Bell, Marc Day, Dan Graves (LBNL).

Scope & Objectives

- The goal of this project is to provide a simulation capability for first-principles (DNS) and near-first principles (DNS/LES hybrids) simulations of turbulence-chemistry interactions in conditions relevant to practical combustion devices, including turbulence, mixing, spray vaporization, low-temperature ignition, and flame propagation.
- In this milestone, we developed the capability to simulate complex geometry using embedded boundary (cut cell)

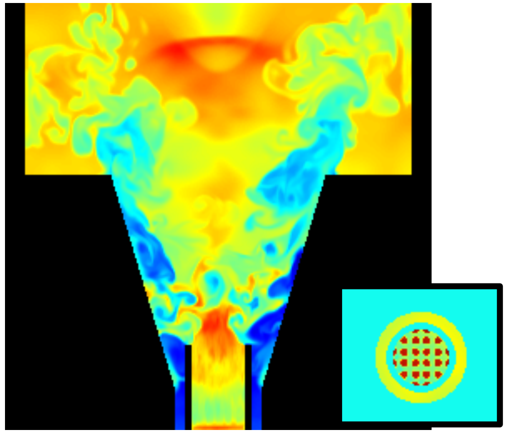

Z-momentum on cutting plane through center of combustor geometry (body cells not blanked, inlet velocity through central pipe in inset)

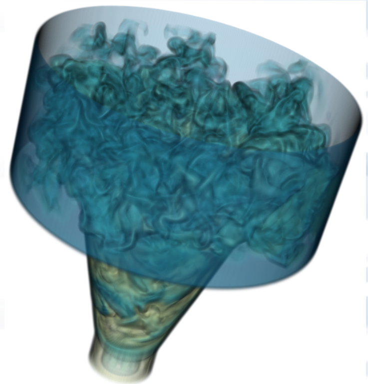

Volume rendering of density field matching image at left

Impact

- Accurate simulation of combustion at high pressure such as conditions in a diesel engine requires modeling non-ideal fluid behavior, particularly for large hydrocarbons
- Four year demonstration problem is a single sector of a gas turbine combustion; the geometry of the flame holder is needs to be captured to generate recirculation zones that anchor the flame.

Project Accomplishment and Next Steps

- Cartesian cut cell implementation in PeleC allows simulation of complex geometry using explicit diffusion treatment and method of lines approach to hyperbolic treatment
- Capability demonstration is ~30x faster than start of project baseline and 5x slower than proof of concept created by AMReX and tailored for gamma-law gas dynamics
- Calculation of diffusive and advective fluxes needs to be coordinated to improve computational throughput and reduce memory usage
- Performance engineering of initial code for more general cases (multivalued, vector potential) is next major step

Initial Release of the Pele Code Suite, Utilizing Block-Structured Adaptive Mesh Refinement Based on AMReX

PI: Jackie Chen, SNL
Members: SNL,LBNL,NREL,ORNL,ANL

Scope & Objectives

- The goal of this project is to provide a simulation capability for first-principles (DNS) and near-first principles (DNS/LES hybrids) simulations of turbulence-chemistry interactions in conditions relevant to practical combustion devices, including turbulence, mixing, spray vaporization, low-temperature ignition, and flame propagation. In addition, we will add capabilities to treat non-ideal gas effects, soot and radiation transport, sprays and non-trivial geometries.
- In this milestone, we released the code suite, Pele, consisting of the compressible code PeleC and the low Mach number code, PeleLM, and supporting physics implementations.

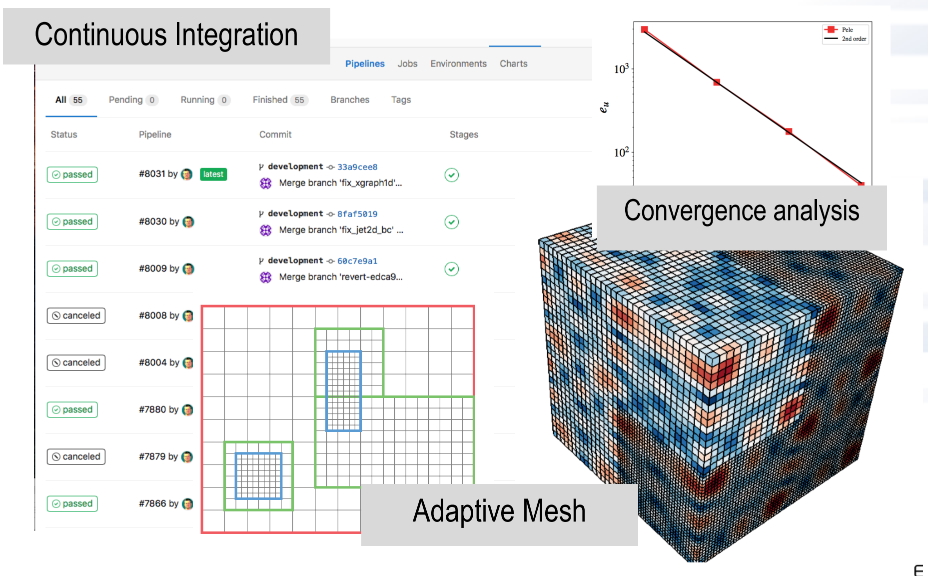

Status	Pipeline	Commit	Stages
passed	#8031 by latest	development 33a9cee8 Merge branch 'fix_xgraph1d'...	✓
passed	#8030 by	development 8faf5019 Merge branch 'fix_jet2d_bc' ...	✓
passed	#8009 by	development 60c7e9a1 Merge branch 'revert-edca9...	✓
canceled	#8008 by		
canceled	#8004 by		
passed	#7880 by		
canceled	#7879 by		
passed	#7866 by		

Impact

- Established unified interface code to incorporate multiphysics for evaporating liquid sprays, and embedded boundary geometry algorithms
- Baseline code for evaluating figures-of-merit as new multi-physics capabilities and algorithms are implemented
- Enables closer ties with ST and co-design centers on refactoring algorithmic component for performance portability

Project Accomplishment and Next Steps

- The release includes an agile development model, a portable build system, and a flexible testing framework with an automated continuous-integration strategy
- Deployed automated code generation for cross-platform evaluation of combustion kinetics, transport coefficients and thermodynamics models.
- Generated holistic performance benchmark data for low Mach and compressible baseline application scenarios
- Next: Add real-gas effects, demonstrate PeleC with sprays, integration with AMReX for non-trivial geometry capabilities

Deliverables: Pele code suite v1.0 at https://code.ornl.gov/Pele (2017/03/31).

EXASTAR

- Astrophysical CFD characteristics
  - Very high turbulent Re (10^{14} for Type Ia supernovae!)
  - Dynamics strongly coupled to small scale instabilities
    - Magneto-rotational instability (MRI; shear induced) in core-collapse supernovae (CCSN) - O(1 meter) resolution required
    - R-T driven turbulence
      - Flamelet burning in white dwarfs
      - Post-shock motion in CCSNe
  - Shocks ubiquitous
  - Complex geometry (e.g. walls, corners, etc.) not found
  - Huge disparity of scales – Red giant radii of several astronomical untis (AU) to meters or centimeters
    - AMR key to progress (AMREx)

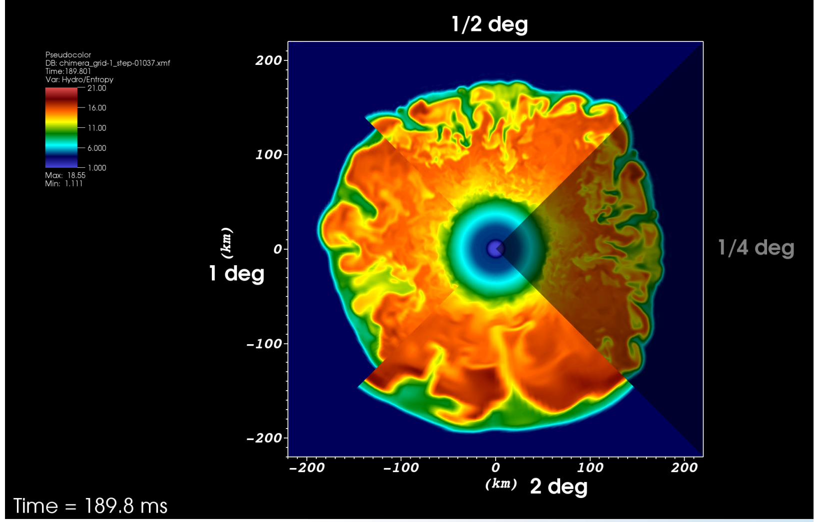

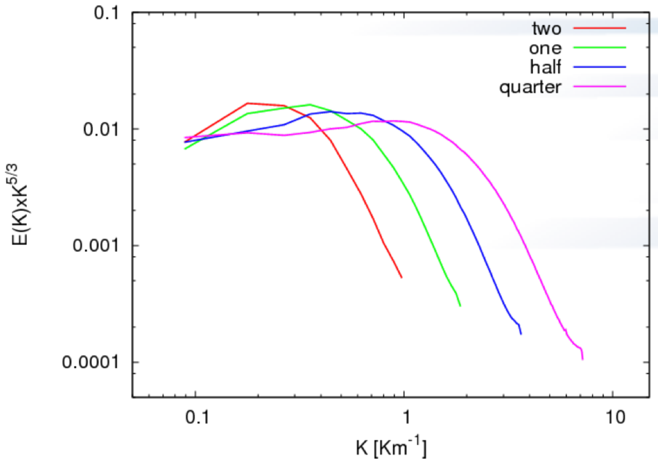

Dan Kasen (PI, LBNL), Bronson Messer (ORNL)

ExaSMR: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors

- Small modular reactors present significant computational challenges
  - Natural circulation (computing pressure drop across core and steam generator)
  - Fine geometric detail (mixing vanes)
  - High fidelity requires move from RANS to LES
  - Multiphase and transient modeling
  - Tightly coupled to neutronics (determines heat deposition into coolant)
- Nek5000 will provide CFD capability for ExaSMR
  - Spectral finite element
  - Demonstrated scaling to >1 million MPI ranks

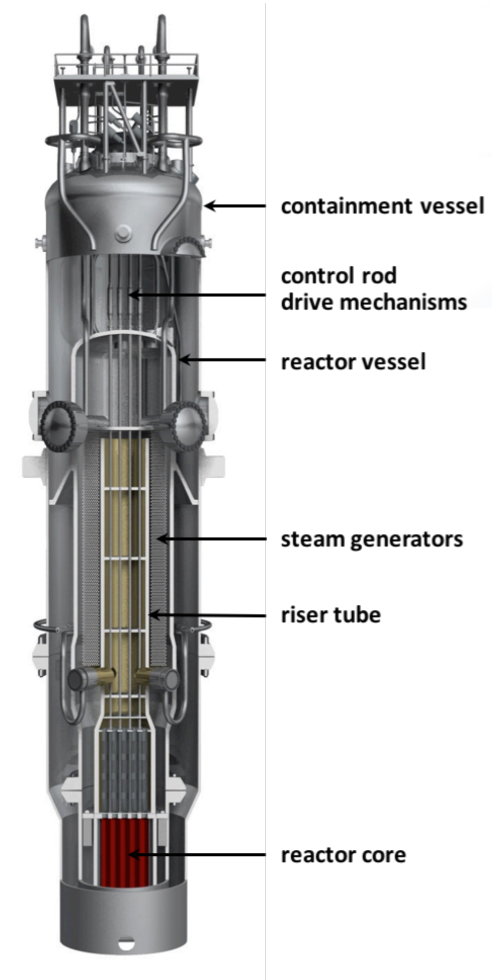

© NuScale Power, LLC. All Rights Reserved

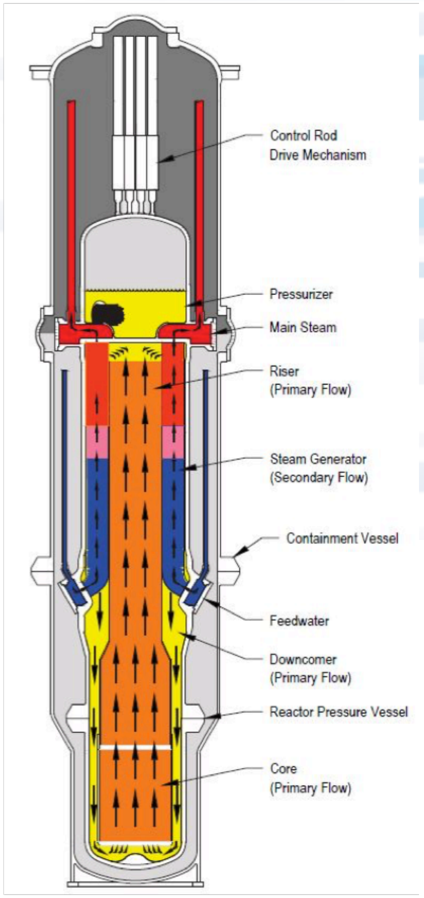

Steve Hamilton (PI, ORNL), Elia Merzari (ANL)

Nek5000 achievements through ECP

- First spectral element CFD code ported to GPU and KNL
  - OpenACC implementation offers portability across architectures
- New ensemble averaging delivers 10x reduction in time to solution
  - Multiple realizations with perturbed initial conditions
  - Allows use of additional computing resources beyond strong scaling limit
- Enhanced algebraic multigrid preconditioner improves weak scaling
  - Pressure Poisson solve is performance limiter in many problem regimes

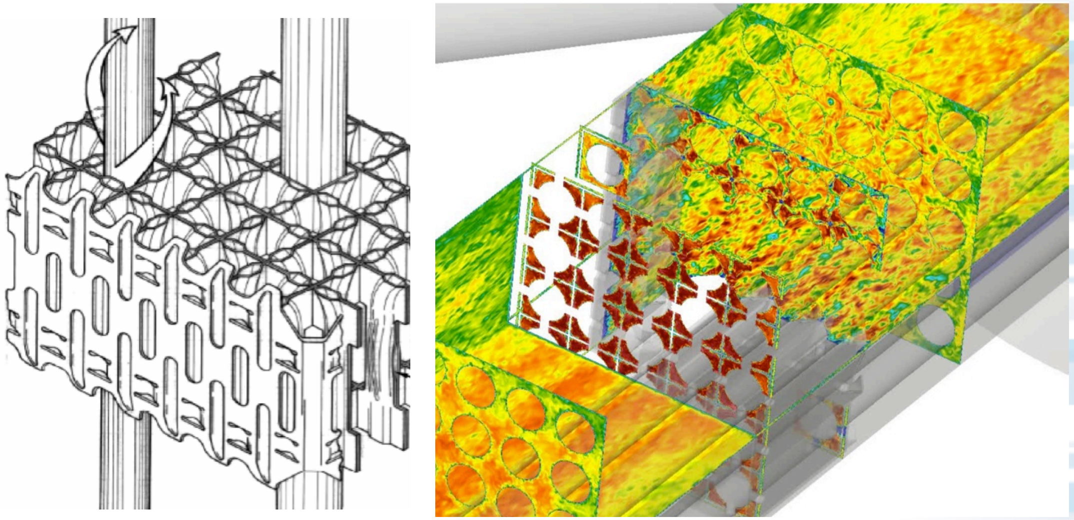

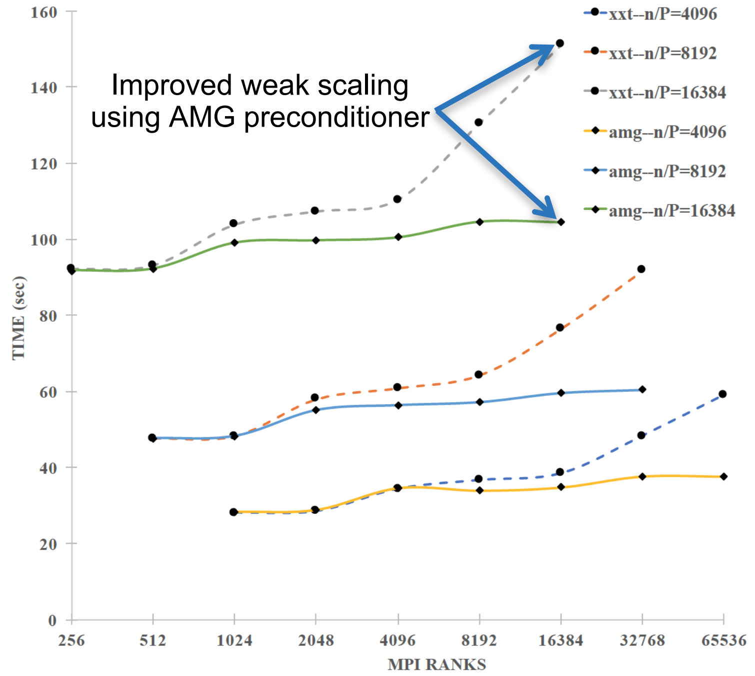

Improved weak scaling using AMG preconditioner

Steve Hamilton (PI, ORNL), Elia Merzari (ANL)

Exascale simulation for the design of industrial-scale chemical reactors

Goal: Develop an efficient high-fidelity multiphase flow modeling capability to aid in the design of industrial-scale chemical reactors

Simulation with high-fidelity, physics-based models is essential to scaling up from lab → pilot → commercial scale reactors

- Reduction in cost
- Reduction in time to deployment
- Risk mitigation at large scales

Proposed increase in fidelity will aid in the development of CO_2 capture technology (supported by DOE-FE) as well as unlock the ability to simulate a host of relevant problems in energy, chemical processing and pharmaceutical industries

Lab-scale testing of novel CO_2 capture materials at NETL (<1 kW)

Scale up?

Petra Nova, world's largest post-combustion CO_2 capture plant, began operation in January 2017 (240 MW)

Madhava Syamlal (PI, NETL)

MFIX-Exa challenge problem

Simulate 1 MWe chemical looping reactor with CFD-DEM

2017

Particle Count: 60×10^6
Time to Solution: 600 days

2023

Particle Count: 5×10^9
Time to Solution: 0.5 days

2026

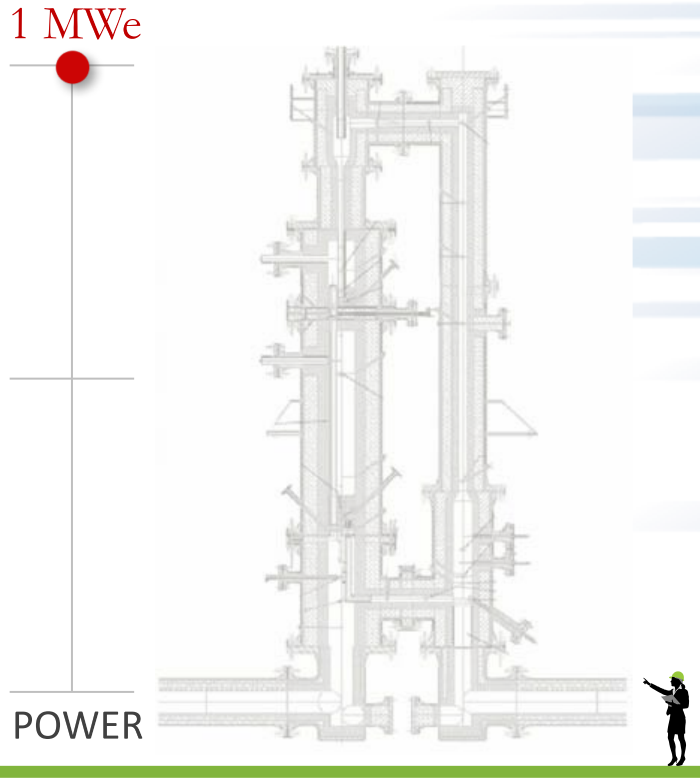

Particle Count: 100×10^9
Time to Solution: 2 days

Time-to-solution is estimated for 5 minutes of real time in all cases. 2017 estimate is for the Joule computer at NETL; 2023/2026 values are guestimates for exascale computers.

Madhava Syamlal (PI, NETL)

EXASCALE COMPUTING PROJECT

MFIX-Exa brings together three teams and two codes

NATIONAL ENERGY TECHNOLOGY LABORATORY

BERKELEY LAB

Boulder

- 60+ years of experience in multiphase modeling and MFIX (NETL and CU)
- 60+ years of experience in large-scale, multiscale multiphysics applications (LBNL)
- 90+ years of experience in high performance computing

MFiX Exa

MFiX + AMReX

- 30+ years of development
- 12 developers at NETL
- 4,000+ registered users
- 175+ downloads per month
- 200+ citations per year
- Applied for reactor design and troubleshooting in fossil, bio, nuclear, and solar energy; chemicals industry; and nuclear waste treatment

- Block-structured AMR software framework supported by ECP Co-Design Center
- Supports multiple DOE codes: accelerator modeling, astrophysics, combustion, cosmology, and subsurface
- Long development history

Madhava Syamlal (PI, NETL)

MFIX-Exa will be based on MFIX-DEM

MFiX DEM

Gas Phase – Navier-Stokes like equations

$$\frac{\partial}{\partial t}(\varepsilon_g \rho_g) + \frac{\partial}{\partial x_j}(\varepsilon_g \rho_g U_{gj}) = 0$$

$$\frac{\partial}{\partial t}(\varepsilon_g \rho_g U_{gi}) + \frac{\partial}{\partial x_j}(\varepsilon_g \rho_g U_{gj} U_{gi}) = -\varepsilon_g \frac{\partial P_g}{\partial x_i} + \frac{\partial \tau_{gij}}{\partial x_j} + f_{gi} + \varepsilon_g \rho_g g_i$$

Particles – Newton's law

$$\frac{dx_{pi}}{dt} = u_{pi}$$

$$m_p \frac{du_{pi}}{dt} = m_p g_i + f_{pi} + m_p A_{coll}$$

$$I_{ij} \frac{d\omega_{pj}}{dt} = T_{pi}$$

Advantages CFD-DEM over TFM

- No numerical diffusion in particle phase
- Does not require (uncertain) constitutive relations for the particle phase
- Able to describe slowly shearing or stagnant regions
- Able to track the evolution of particle-scale properties: size, density, chemical conversion …

Garg R, Galvin J, Li T, Pannala S. Documentation of open-source MFIX–DEM software for gas-solids flows, From URL https://mfix.netl.doe.gov/documentation/dem_doc_2012-1.pdf. (2012)

MFIX-DEM hydrodynamics migrated to AMReX framework

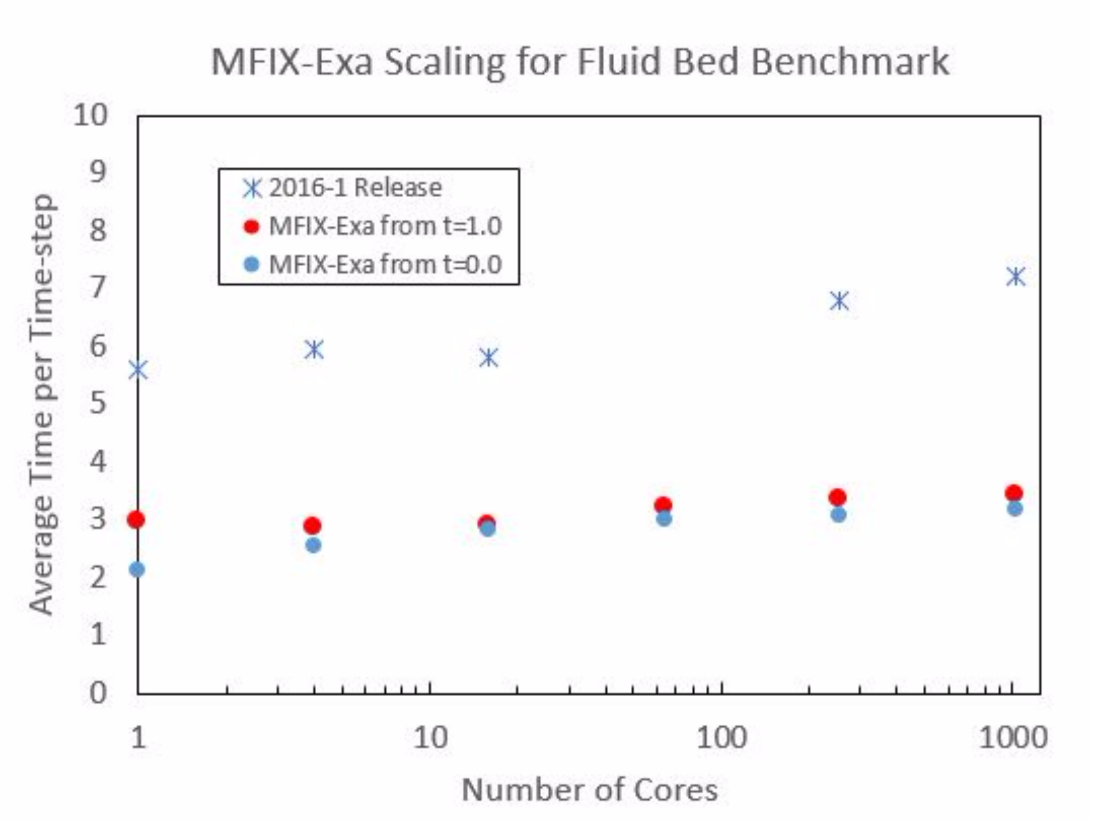

A preliminary scaling study on Cori at LBL shows a 2X performance improvement over MFIX-2016-1.

The hourglass simulation demonstrates MFIX-Exa's newly developed (FY18-Q1) capability for embedded boundaries (EB).

MFIX-Exa with hybrid parallelism and dynamic load balancing released to ECP.

Madhava Syamlal (PI, NETL)

ExaSky: Cosmological Hydrodynamic Simulations with CRK-HACC

- Next generation ExaSky simulations supplement the highly-tuned, gravity-only version of HACC (Hardware Accelerated Cosmology Code) with baryonic capabilities (CRK-HACC).
  - Multi-species approach with both dark matter (DM) and baryons interacting gravitationally.
  - Baryonic gasdynamics modeled using the Conservative Reproducing Kernel Smoothed Particle Hydrodynamics (CRK-SPH) algorithm.

- CRK-SPH improves upon many aspects of traditional SPH (see Frontiere et al. 2017, J. Comp. Phys.)
  - Uses higher order reproducing kernels which can exactly interpolate constant and linear order fields (removes the “E0 error” of traditional SPH).
  - Uses a conservative reformulation of the dynamic equations that maintain machine precision energy and momentum conservation (a previous problem encountered when using higher order kernels).
  - Uses a new artificial viscosity form that capitalizes on the increased accuracy calculation of the velocity gradients (improves the excessive diffusion normally seen in SPH).

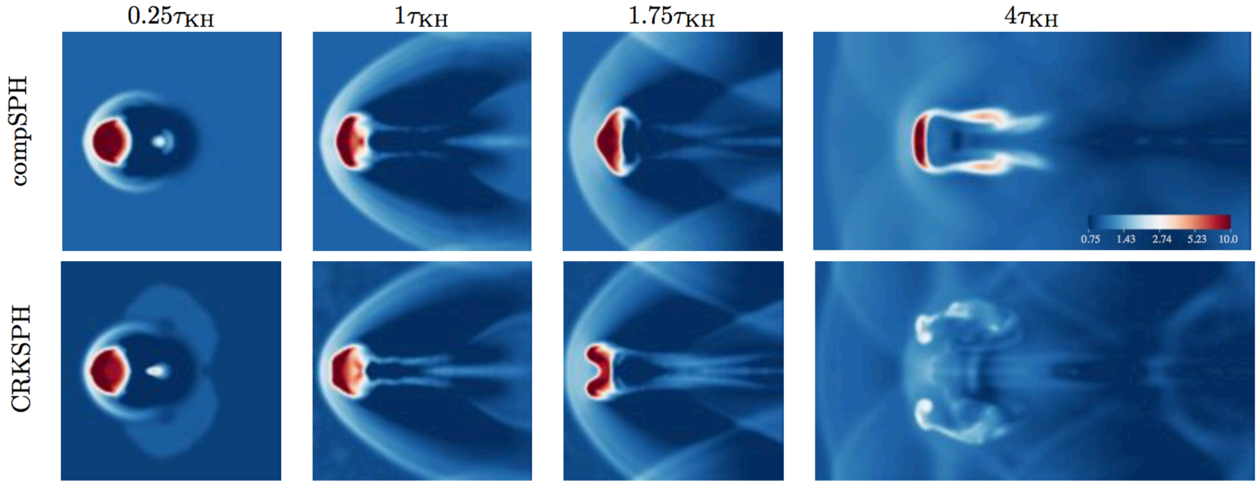

Blob test: Traditional SPH (top) artificially suppresses the disruption and mixing of the blob into the background medium. CRK-SPH (bottom) leads to a full evaporation of the blob.

ExaSky: Cosmological Hydrodynamic Simulations with CRK-HACC

BorgCube: First large-scale CRK-HACC simulation containing 2x2304^3 DM plus baryonic particles in a cubic volume of side length 800 Mpc/h. Run on 3072 nodes of Theta (KNL system) at ALCF.

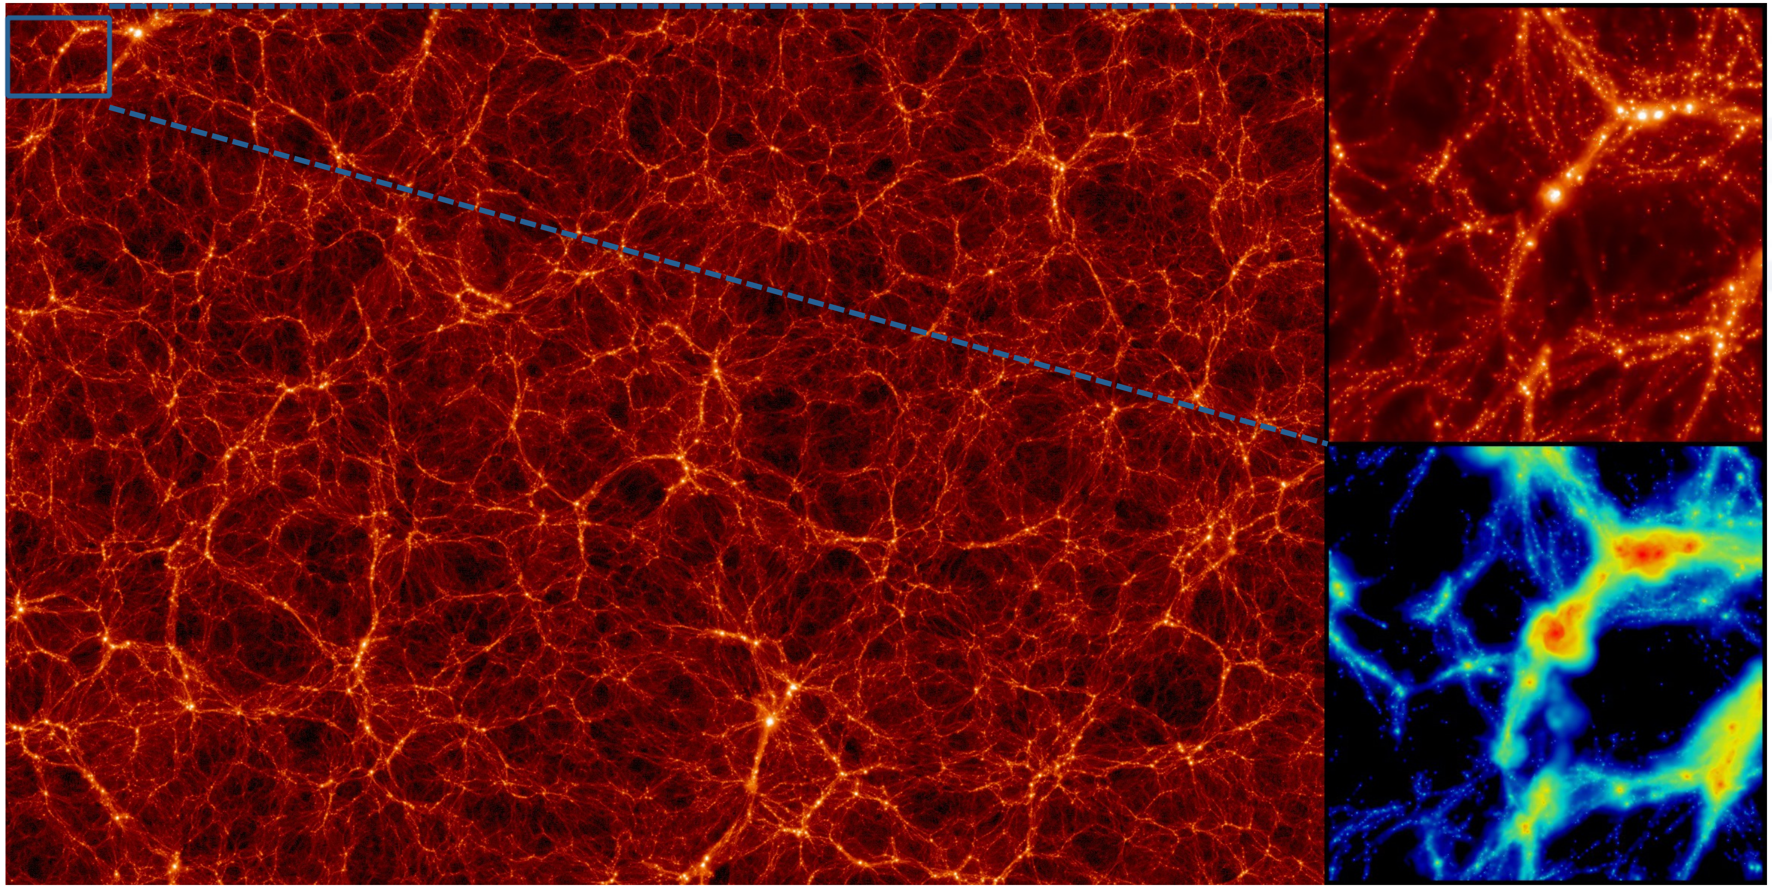

Baryonic density (left and top-right panels) and temperature (bottom-right panel) at z = 0. Left panel is 600x400 Mpc/h while the right panels zoom in on a 50x50 Mpc/h region. The depth of each slice is 10 Mpc/h.

ExaWind: Challenge Problem, Motivation, Goals

10-Year challenge problem: Predictive simulation of the complex flow physics within a wind plant composed of O(100) multi-MW wind turbines sited within a 10 km x 10 km area with complex terrain, involving simulations with O(100) billion grid points.

Goals/motivation for predictive simulations:

- Advance our fundamental understanding of the flow physics governing whole wind plant performance
- Create modeling and simulation capability that will revolutionize the design and control of wind farms
- Advance our ability to predict the response of wind farms to a wide range of atmospheric conditions

Where are we today with wind farm simulation? State of the art:

- Turbine blades are represented as “line forces” in the fluid domain; based on 2D look-up tables
- Turbulence models are incapable of capturing the complex flow separation and wake generation

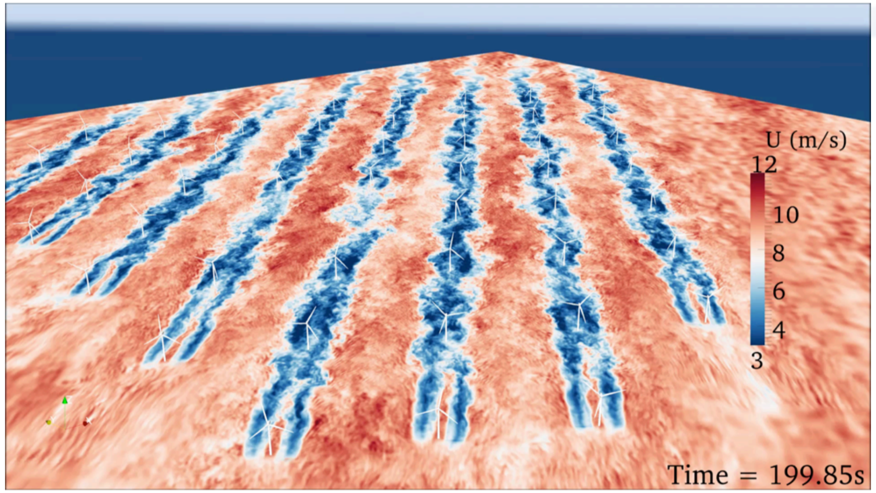

SOWFA actuator-line simulation of the Lillgrund wind farm (Churchfield)

Application Metrics: Let’s Boil it Down . . .

1. **Deliver improved and impactful science & engineering (*performance*)**
   - New or improved (ideally step change in) predictability on a problem of national importance (a “challenge problem”)
2. **As performance portable as possible and reasonable (*portability*)**
   - No “boutique” one-off applications able to only execute on one (and likely ephemeral) system
3. **Able to make effective use of a capable system (*readiness*)**
   - *Effective* is app specific (weak, strong, ensembles, single-node performance)
4. **Able to integrate latest relevant software technologies (*modern*)**
   - Needed to demonstrate agility, flexibility, modern architecture; overall app portfolio must apply pressure to all key attributes of the system design characteristics
5. **High priority (*strategic*)**
   - Some key stakeholder somewhere really cares about using application to make consequential decisions

How can Applications be “Impactful”?

- **Target development toward a few specific *Challenge Problems (CP)***
  - If you try to be good at everything, you’ll be good at nothing
  - Allows development progress to be measured more easily and objectively
    - E.g.: ability to deliver on yearly “progression problems” capability (marching toward the CP)
- **Yet try to design architecture to be more general**
  - To tackle that next *way cool* problem you’d like to tackle, but . . .
- **A Challenge Problem is one that . . .**
  - Has a solution amenable to M&S and/or data analytic computing
  - Is currently intractable on today’s hardware (requires capable exascale)
  - Matters immensely to a key stakeholders (addresses a top 3 strategic priority)

Delivering on the Challenge Problem (CP)

- **Define a specific and quantitative (if possible) CP to target**
  - Flow requirements down from business (why?) to functional (what?) to design (how?) to quality (how?) requirements
  - Spec out as much of the CP as possible while still allowing R&D flexibility
- **Measure the application's *rate of executing useful work (W/t) toward* addressing the CP**
  - W/t can go up by increasing the quality of the science (W↑) or accelerate the return on the science (t↓)
  - For most applications, increasing W and decreasing t both apply
  - $(W/t)_{\text{exascale}} = N(W/t)_{\text{today}}$
  - What should N be as the target? Probably *not* peak flop ratio, but . . . [sigh]
  - How to best measure (W/t)? What does "today" mean (given the half life)
- **CPs are *very* application specific (and must be)**

Questions?