

Pele: An Exascale-Ready Suite of Combustion Codes

18th Numerical Combustion Conference, San Diego, CA

May 11, 2022

M. Day², L. Esclapez¹, M. Henry de Frahan¹, J. Rood¹,
H. Sitaraman¹, N. Wimer¹, R. Grout¹, J. Chen²


¹National Renewable Energy Laboratory, Golden, CO

²Sandia National Laboratories, Livermore, CA

What is the Exascale Computing Project (ECP)?

ECP is an accelerated research and development project funded by the US Department of Energy (DOE) to ensure all necessary pieces are in place to deliver the nation's first, capable, exascale ecosystem, including mission critical applications, an integrated software stack, and advanced computer system engineering and hardware components.

What is a “capable” exascale computing ecosystem?



Hardware

At least two diverse system architectures

Delivers 50x the performance of today’s 20 petaflop systems and 5x the performance of Summit, Oak Ridge National Laboratory’s supercomputer— i.e., allows at least a quintillion floating point operations per second

Functions with sufficient resiliency: an average fault rate of ≤ 1 per week



Software

Includes a software stack that meets the needs of a broad spectrum of applications and workloads

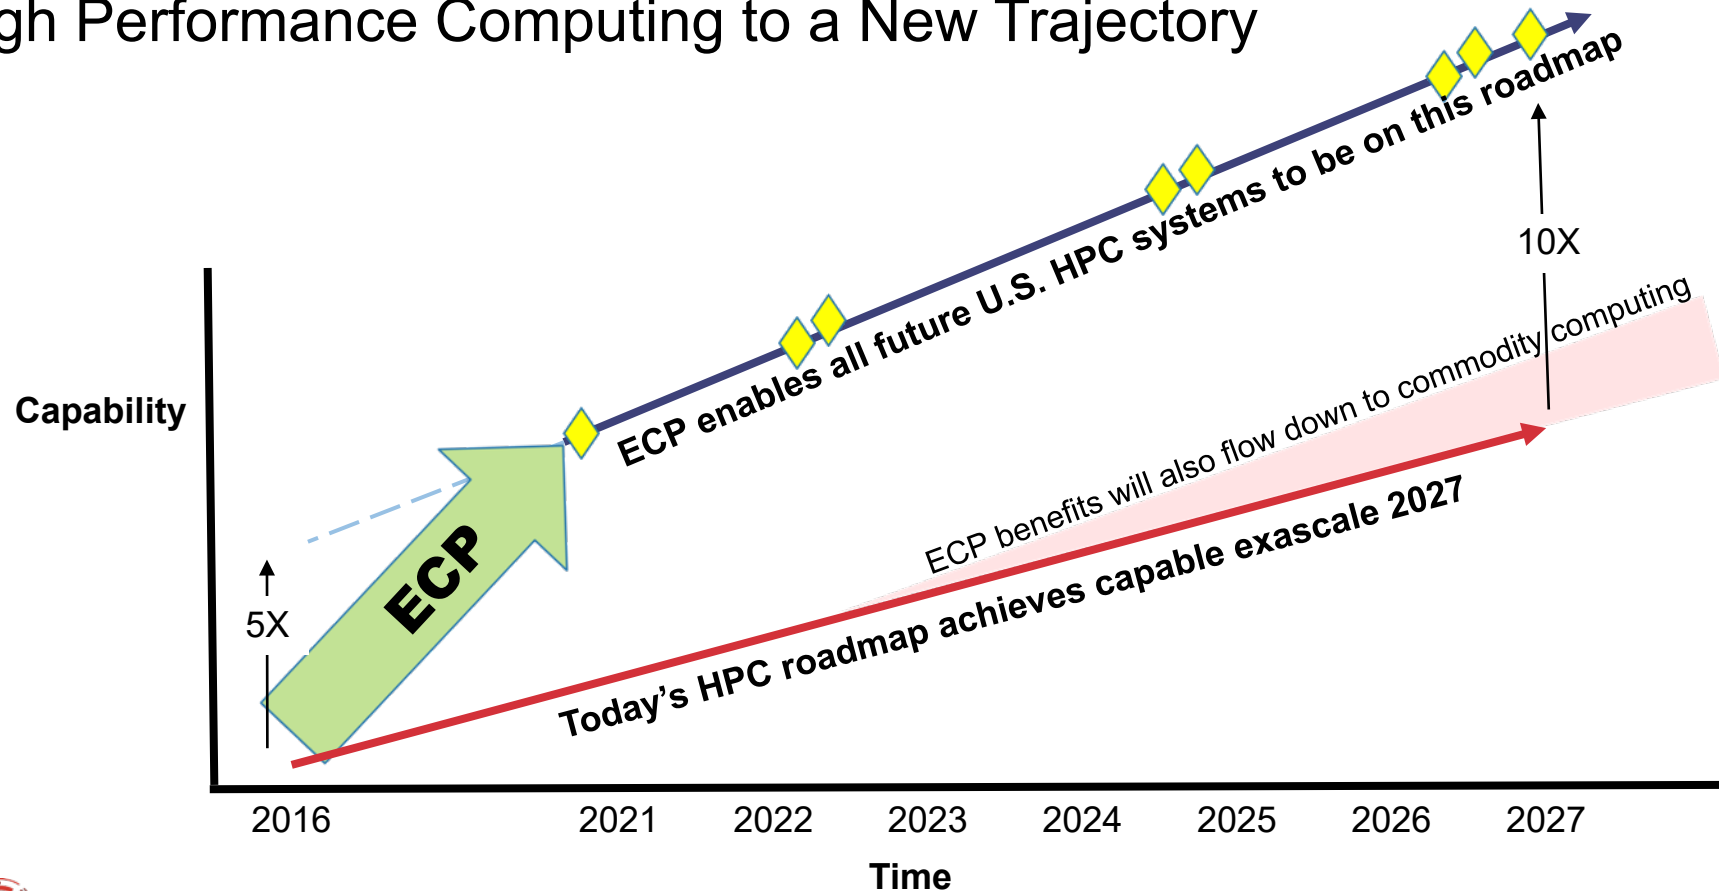


Applications

Supports a wide range of applications that deliver high-fidelity solutions in less time to problems of greater complexity

Exascale means real capability improvement in the science we can do, and how fast we can do it

Vision: Exascale Computing Project (ECP) Lifts all U.S. High Performance Computing to a New Trajectory





ECP applications target 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: Energy Security Applications

Exascale challenge problems and [impact](#)


• Turbine Wind Plant Efficiency

- Predictive simulation of the complex flow physics within a wind plant composed of 100s of multi-MW wind turbines to advance the fundamental understanding of the flow physics governing whole wind plant performance. Revolutionize the design and control of wind farms, as well as advance ability to predict the response of wind farms to a wide range of atmospheric conditions.
- [Harden wind plant design & layout against energy loss susceptibility; higher penetration of wind energy](#)




• Design and Commercialization of Small Modular Reactors (SMRs)

- Coupled neutronics and fluid dynamics will be used to create virtual experimental datasets for SMRs under varying operational scenarios. Validate fundamental design parameters including the turbulent mixing conditions necessary for natural circulation and steady-state critical heat flux margins between the moderator and fuel. Provide validation for low-order engineering simulations and reduce conservative operational margins resulting in higher updates and longer fuel cycles
- [Virtual test reactor for advanced designs via experimental-quality simulations of reactor behavior](#)



• Subsurface Carbon Capture, Fossil Fuel Extraction, Waste Disposal

- Safe and efficient use of the subsurface for geologic CO₂ sequestration, petroleum extraction, geothermal energy and nuclear waste isolation. Understand and predict reservoir-scale behavior as affected by the long-term integrity of the hundreds of thousands deep wells that penetrate the subsurface for resource utilization, needing integration of complex multi-physics processes occurring at multiple scales, from the micro- to the kilometer scale, in a high resolution reservoir simulator
- [Reliably guide safe long-term consequential decisions about storage, sequestration, exploration](#)




ECP: Energy Security Applications

Exascale challenge problems and [impact](#)


• Scaleup of Clean Fossil Fuel Combustion

- Develop an efficient high-fidelity multiphase flow modeling capability for the design of industrial-scale chemical reactors in order to reduce cost, reduce time to deployment, and provide risk mitigation at large scales. Aid in the development of carbon dioxide capture technology as well as unlock the ability to simulate a host of relevant problems in energy, chemical processing and pharmaceutical industries
- [Achieve commercial-scale demo of transformational energy technologies - curbing man-made CO2 emissions at fossil fuel power plants - by 2030](#)



• Biofuel Catalyst Design

- Developing and optimizing "fragment" methods within the GAMESS package that allow the underlying methods to apply to large complex systems with tens of thousands of atoms and effectively take advantage of exascale systems utilizing different architectures - simulation of these complex systems will enable design of new catalysts - a grand challenge problem that will require the availability of exascale computers
- [Design more robust and selective catalysts and catalytic processes that are orders of magnitude more efficient at temperatures that are hundreds of degrees lower](#)



Pele Combustion Suite

PI: J. Chen, SNL

Pele is the Exascale Computing Project's (ECP's) application suite for high-fidelity detailed simulations of turbulent combustion in open and confined domains.

- Detailed physics, geometrical flexibility to evaluate design and operational characteristics of clean, efficient combustors for automotive, industrial, and aviation applications
- Targets simulation capabilities required to inform next-generation combustion technologies, for example:
 - Advanced internal combustion engines (e.g., RCCI)
 - Novel supercritical CO₂ power cycles
 - Rotating detonation engines
 - Supersonic cavity flame holders
 - Aviation combustors for sustainable drop-in JetA fuels


Pele combustion simulation and analysis suite:

- *PeleC* (compressible), *PeleLM*, *PeleLMEx* (low Mach)
- *PelePhysics* (thermodynamics, transport, chemistry models)
- *PeleAnalysis* (in-situ, post-processing/analysis)
- *PeleMP* [multi-physics] (soot, radiation, Lagrangian spray models)
- *PeleProduction* (collaboration hub)

Open-source code developed under the Exascale Computing Project:

<https://github.com/AMReX-Combustion>

Pele's KPP2 Combustion Challenge Problem: RCCI



PeleC, PeleLM, PeleLMEx

PelePhysics

Transport,
thermodynamics,
finite-rate chemistry

AMReX-Hydro

Hydrodynamics and
geometry

SUNDIALS

Implicit/explicit ODE
integrators

MAGMA

Batched linear solvers

AMReX

Block-structured AMR library

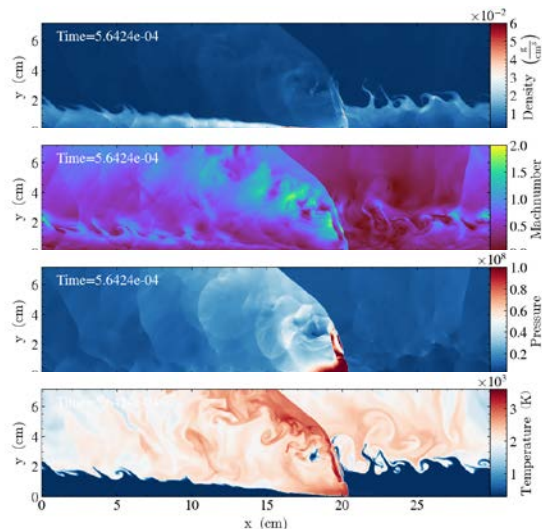
HYPRE

Distributed linear
solvers


PeleC

PeleC: Compressible flow solver

- Conservation of species mass densities, momentum, total energy
- Time-explicit Runge-Kutta (RK)-based advance
 - Time-explicit RK variants for diffusion and hyperbolics (PPM, PLM, WENO, MOL)
 - SUNDIALS-driven ODE integration for finite-rate chemical kinetics (CVODE, ARKODE)
- PelePhysics provides finite-rate chemistry models, equations of state and transport properties. Non-ideal thermo/chemistry modifications and tabulated lookup table models available
- PeleMP provides access to optional multiphase (spray) fuel models via AMReX particle capability
- <https://github.com/AMReX-Combustion/PeleC>



Rotating detonation engine
Sreejith NA et al., early results, 2022




Quad fuel jet into shaped piston bowl
Henry de Frahan et al.,
early results, 2022




Supersonic cavity flame holder
Sitaraman et al., Combustion and Flame, 2021

PeleLM(eX)



Gas turbine premixer, PeleLM
M. Vabre, B. Savard, et al.,
CICS Spring Technical Meeting, 2022

Spherical confined flame, PeleLM eX
L. Esclapez et al., in prep, 2022



PeleLM/PeleLM eX: Low Mach flow solvers

- Conservation of species mass densities, momentum, enthalpy
- Iterative/implicit SDC-variants for tightly coupled ADR systems resulting from large Δt enabled by the low Mach algorithm
 - Semi-implicit (Crank-Nicolson) diffusion, Godunov-based hyperbolics (BDS)
 - SUNDIALS-driven ODE integration for finite-rate chemical kinetics (CVODE)
- PelePhysics provides finite-rate chemistry models, equations of state and transport properties. Tabulated lookup table and neural-net-based models available
- PeleMP provides access to optional multiphase (spray) fuel models (Lagrangian, AMReX particles based), moment-based soot models, and radiation transport
- Critically, the low Mach model requires the solution of a linear elliptic system to compute the constrained spatially isobaric solution, and a set of linear systems for the implicit diffusion solve. *Due to geometry-induced ill-conditioning, the elliptic systems require HYPRE's BoomerAMG, with robustified, iterative smoothers*
- ***PeleLM eX's non-subcycled integrator supports AMR with closed-chamber pressurization due to fueling and heat release***



Lagrangian fuel sprays in PeleLM eX
(droplets colored by T)
Ariente et al., in prep, 2022



Quad n-dodecane jets into KPP PB
Wimer, Esclapez, et al., in prep, 2022


<https://github.com/AMReX-Combustion/PeleLM>
<https://github.com/AMReX-Combustion/PeleLM eX>

PelePhysics

An open-source combustion physics library


- <https://github.com/AMReX-Combustion/PelePhysics>
- EOS: ideal gas mixtures (CHEMKIN), Soave-Redlich-Kwong (SRK), and EOS lookup tables
- Models and parameters for thermodynamics
- Mixture-averaged and unity Le transport properties, including extensions for non-ideal gases
- Chemical reactions and finite-rate chemistry integration via SUNDIALS
- Python-based C++ generator to convert CHEMKIN combustion models into production rate and reaction Jacobian code for CPU/GPU evaluation, including optional quasi-steady-state assumptions (QSSA)

Turbulent sCO₂ jet with SRK EOS [2]



Manifold Models and ML in PeleLM


Exascale DNS Calculations
Using PeleC and PeleLM
>TB of data



Ex.: Spark Ignition (PeleC),
B. Soriano and J. Chen, SNL

Closing the Loop:
A Posteriori
Validation



High-Fidelity
Engineering
Calculations



Ex.: Oxy-Combustion Power Cycle Combustor

ML Model Training

Ex.: Co-Optimized Machine-Learned Manifolds [1]



A Priori Model Validation

Manifold-Based Models:

- Due to fast chemical time scales, the observed thermochemical states in a combustion system often lie on a low-dimensional manifold in state space:

$$[\dot{\omega}_j, Y_j, e, \rho, D, \dots] \rightarrow \phi_k = F(Y_j, T, p) \rightarrow \phi_k = \mathcal{F}(\xi_i) \leftarrow N_\xi \ll N_Y$$

- Manifold models: solve equations for ξ_i rather than Y_j and tabulate (or use a neural network to approximate) \mathcal{F} to compute all chemical quantities
- Physics-based examples: flamelet/progress variable, flamelet generated manifolds, etc.
- Data-driven examples: principal component analysis (PCA), co-optimized machine-learned manifolds


Implementation Targets:

- Leverage datasets from exascale simulations for data-driven modeling
- Enable high-fidelity LES of complex configurations
- Modularity and generality to enable physics-based and data-driven manifold models of various types, using both tabulation and neural networks
- Initial focus on PeleLM, but most routines implemented in shared PelePhysics library

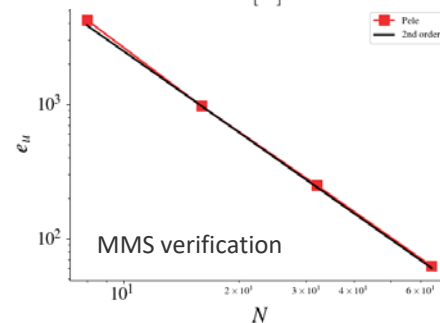


Pele Numerical Methods

All Pele tools exploit block-structured AMR

- Extensively leverage AMReX library (data structures, communication, parallelism, GPU acceleration, ...)
- Conservative cross-refinement finite-volume methods
 - PeleC: Time-explicit RK variants for diffusion and hyperbolics (PPM, PLM, WENO, MOL) with temporally split chemistry evolution
 - PeleLM(eX): Iterative/implicit SDC-variants for tightly coupled ADR systems resulting from large dt via the low Mach algorithm
- PeleC/PeleLM: Subcycling supports constant CFL time advance strategy across AMR hierarchy. PeleLMeX utilizes a non-subcycled time advance to support AMR with closed/pressurizing chambers
- Pele's CI supports formal design order verification through method of manufactured solutions (leveraging MASA library)
- AMR-aware in situ and post-processing tools (surface, slice and stream tubes/line extraction, high-dimensional sampling/statistics, CEMA and reaction path analysis, ROM/ML training, table/NN physics lookup, subsetting, demand-driven processing IO, etc.)





Block-structured AMR in AMReX



Embedded Boundaries

Complex geometries

- AMReX provides the functionality to define an embedded boundary (EB) representation of internal and external complex domain shapes in 2D and 3D
 - Geometry definition (simple shapes, CSG, surface meshes)
 - Data structures (cell, face, edge, node)
 - Numerical PDE operators (grad, div)
- The state is maintained in block-structured arrays, and away from the EB, the discretization stencils are standard, well-characterized uniform-grid operators
- Regular cells are “cut” or completely covered by the EB, and the corresponding stencils must be locally modified to account for partially or fully obstructed face areas and volumes (multi-cuts/multi-cells are not currently supported)
- Redistribution (discussed next) is used to avoid the instability arising from explicit updates of cut cells at the CFL-constrained dt based on the uniform cell size



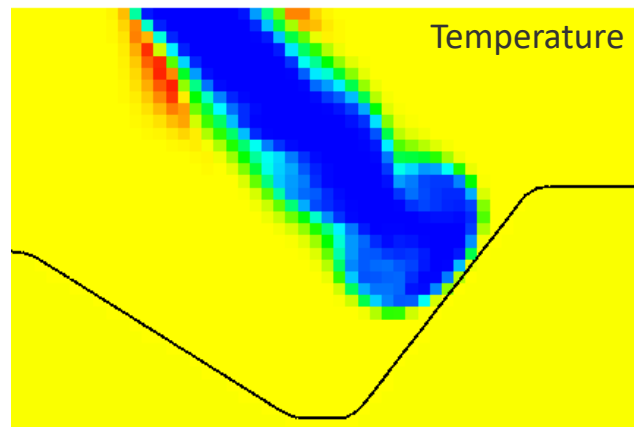
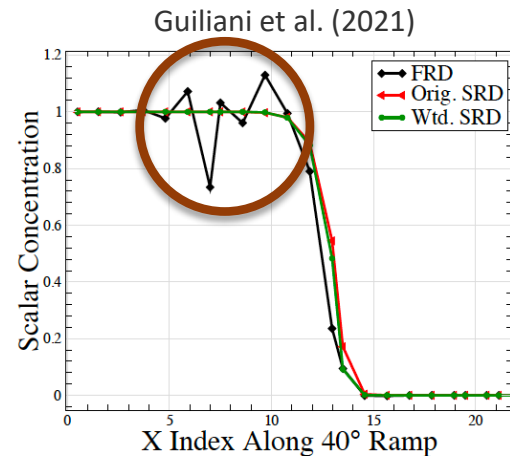
Pele Numerical Considerations

Additional considerations in the presence of complex geometries

- “Redistribution” is used to robustify the conservative update for cells cut by the EB
- However, there is some “art” to how best to carry this out—which cells to involve and how to weight the redistribution. AMReX provides several different redistribution options [3,4]
- For complex reacting flows, preserving mass fraction boundedness $[0, 1]$ during evolution is paramount. Out-of-bounds, non-physical state values lead to ill-defined:
 - Physical transport and thermodynamic properties
 - Chemical kinetics rate expressions
- Non-linearity introduced by flux interpolation and limiters in some of AMReX’s redistribution schemes may introduce or compound bounds violations
- The Pele team has worked extensively with AMReX to develop and test redistribution schemes that preserve the required constraints on the evolving state
- Also, for low Mach flows, PeleLM and PeleLMex make use of the newer **positivity-preserving BDS advection scheme** (NOTE: not available for compressible flows)

[3] Giuliani et al., submitted, 2021

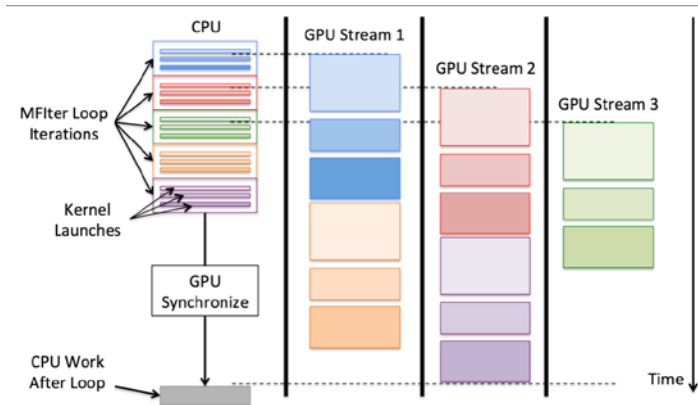
[4] github.com/AMReX-Codes/AMReX-Hydro



High-Performance Computing (HPC) With PeleC

HPC programming model from AMReX

- Pele codes are now pure C++
- Lambda-based C++ GPU performance portability framework
- Advantages: portability between CPUs and multiple GPU vendors, code readability, code reduction, and a single language solution




```
#pragma omp parallel if (amrex::Gpu::notInLaunchRegion())
for (amrex::MFIter mfi(mf, TilingIfNotGPU()); mfi.isValid(); ++mfi) {
    const amrex::Box& bx = mfi.tilebox();
    amrex::Array4<amrex::Real> const& fab = mf.array(mfi);
    amrex::ParallelFor(bx, ncomp,
        [=] AMREX_GPU_DEVICE (int i, int j, int k, int n) {
            fab(i,j,k,n) += 1.0;
        });
}
```

Example PeleC kernel using AMReX's C++ framework

Performance Testing at Scale

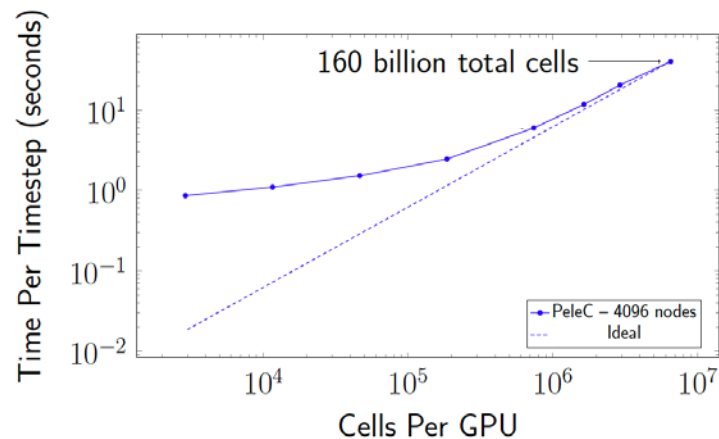
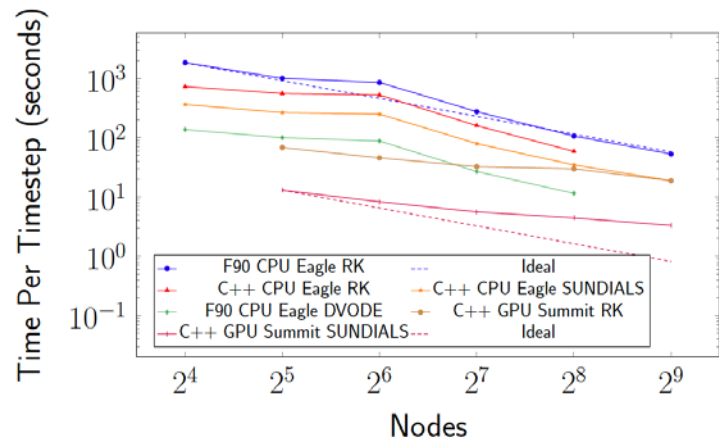
Weak scaling, entire Summit

PMF, 20B cells, 750k cells per GPU, 2 levels of AMR



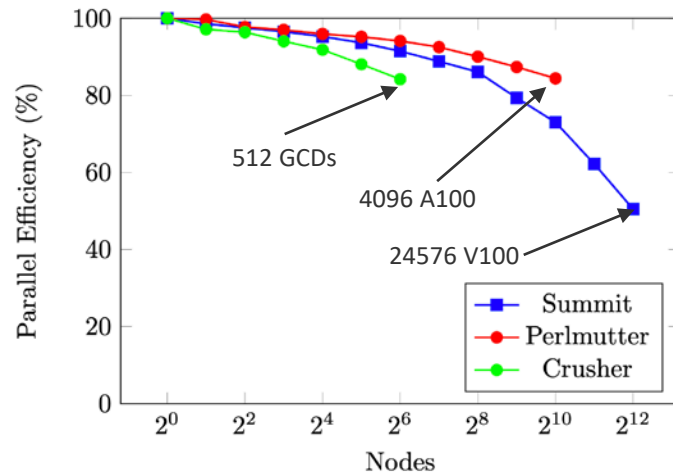
Parallel efficiency drop of only 34%

Strong scaling




PeleLM(eX) Update

- Finalized port of PeleLM and PeleLMeX capabilities to Crusher (turbulent jet injection, finite-rate chemistry, linear solvers)
- On a single Summit node, we see roughly 25x speedup comparing 68 CPU/MPI ranks vs. 6 MPI/GPUs
- “Good” weak scaling up to ~4096 Summit nodes
- “Moderate” strong scaling w/ linear solves rapidly plateauing while chemistry shows nearly optimal scaling




Weak scaling of PeleLM on a 53 species dodecane premixed flame

Node and Communication Performance Status



Similar performance stories for compressible and low Mach:

- With increasing node count (# MPI ranks), communication of grow cells begins to dominate
- We are now seeing consistent performance parity between Summit and Crusher



High-Performance Iterative Solvers for Exascale Combustion

Scientific Achievement


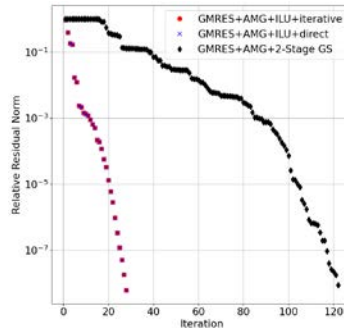
Highly ill-conditioned linear systems can be solved in a fast, scalar manner with AMG-ILU-preconditioned Krylov algorithms

Significance and Impact

3X performance gains achieved over state-of-the-art AMG-ILU preconditioned GMRES for systems whose eigenvalues span 15–16 orders of magnitude

Research Details

- Powerful smoothing algorithms such as ILU are required for highly ill-conditioned linear systems
- Correctly applied equilibration algorithms enable fast, Jacobi iterative (SpMV-based) solutions to the sparse triangular factors
- Jacobi iterative techniques for triangular solvers remove the need for expensive setup algorithms
- Result is highly scalable class of iterative algorithms for difficult linear systems, such as those arising in combustion simulations



Error reduction and scalability of GMRES+AMG+ILU:

(Left) Error reduction capacity of our new AMG-ILU-preconditioned GMRES solver for a highly ill-conditioned linear system, compared to previous high-performing schemes. This was run on 32 V100 GPUs and the total acceleration is 4x over 2-stage Gauss Seidel. The example system here arises from the projection update step in the low Mach combustion modeling code, PeleLMex [2]. PeleLMex utilizes a cut-cell strategy to represent complex internal combustion engine configurations. Tiny cells along the boundary lead to a matrix with eigenvalues that span 16 orders of magnitude.

(Right) Strong scaling performance of our AMG-ILU-preconditioned GMRES solver for the same ill-conditioned linear elliptic PeleLMex system, including setup and solve costs. Crusher vs. Summit.

[1] Mullowney et al., Proceedings of SuperComputing, 2021

[2] <https://github.com/AMReX-Combustion/PeleLMex>



Questions?

www.nrel.gov

NREL/PR-2C00-82880

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Science and National Nuclear Security Administration. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes. This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. A portion of the research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

