Adaptive Mesh Refinement Simulations for Turbulent Reacting Flow

SIAM-CSE23, Amsterdam, Netherlands, March 1, 2023

M Day, M Henry de Frahan, J Rood, N Wimer, B Perry, H Sitaraman, S Yellapantula, L Esclapez, and S Appukutan

National Renewable Energy Laboratory, Golden, CO

with contributions from

The Sandia Pele Team (Chen, Soriano, Owen) - SNL
The AMReX team (Bell, Almgren, Zhang, Myers) – LBNL
The SUNDIALS Team (Woodward, Balos, Gardner) – LLNL
A Aspden and T Howarth, Newcastle University (UK)

Pele Combustion Suite

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 CO_2 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

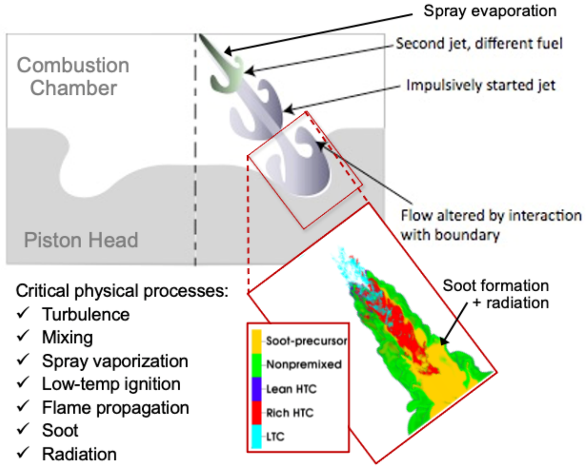

PeleC, PeleLM, PeleLMeX

PelePhysics Transport, thermodynamics, finite-rate chemistry	**AMReX-Hydro** Hydrodynamics and geometry	**SUNDIALS** Implicit/explicit ODE integrators **MAGMA** Batched linear solvers **HYPRE** Distributed linear solvers
AMReX Block-structured AMR library		

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

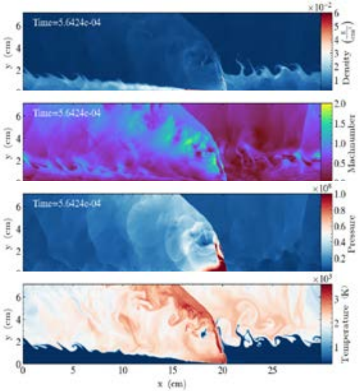

Rotating detonation engine
Sreejith NA et al., 2022

Compression ignition in a high-pressure combustion chamber, computed on Frontier as part of Exascale Computing Program Challenge Problem, 20223

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

PeleLM(eX)

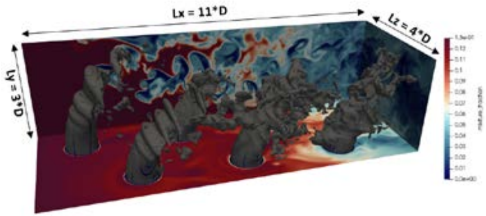

Gas turbine premixer, PeleLM
M. Vabre, B. Savard, et al.,
CICS Spring Technical Meeting, 2022

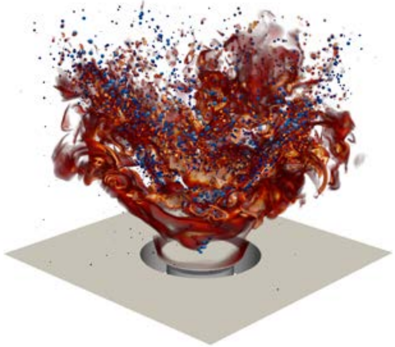

Aero-engine spray flame stabilization
with SAF fuel (C1-ATJ), B Soriano, et al

PeleLM/PeleLMeX: Low Mach flow solvers

- Conservation of species mass densities, momentum, enthalpy
- Iterative/implicit SDC-variants for tightly coupled ADR systems resulting from large dt 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*
- ***PeleLMeX's non-subcycled integrator supports AMR with closed-chamber pressurization due to fueling and heat release***

https://github.com/AMReX-Combustion/PeleLM
https://github.com/AMReX-Combustion/PeleLMeX

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

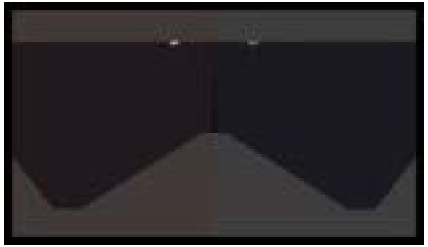

Quad n-dodecane jets into KPP PB
Wimer, Esclapez, et al., in prep, 2022

PelePhysics

An open-source combustion physics library

- https://github.com/AMReX-Combustion/PelePhysics
- EOS: ideal gas mixtures (CHEMKIN), Soave-Redlich-Kwong (SRK), EOS lookup tables / neural nets
- 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 sCO_2 jet with SRK EOS [2]

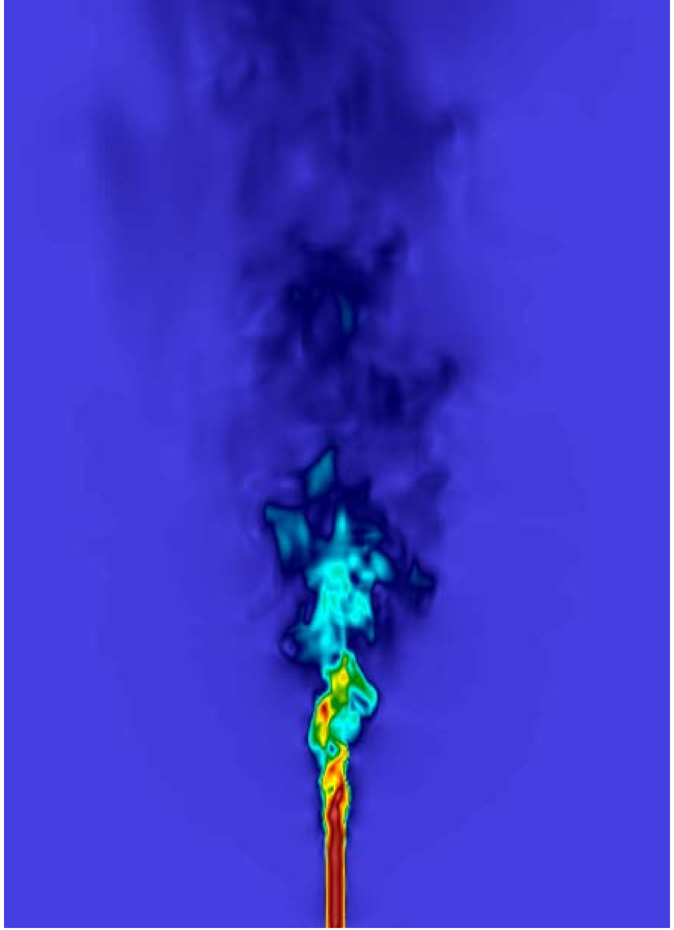

[2] Ream et al., SIAM CSE, 2021

Pele Methods and Tools

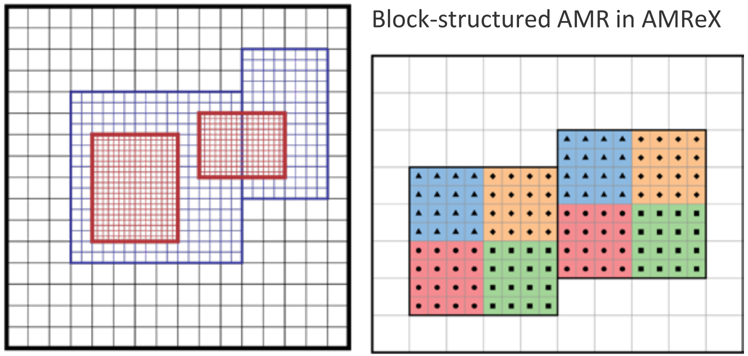

All Pele tools exploit block-structured AMR

- 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.)

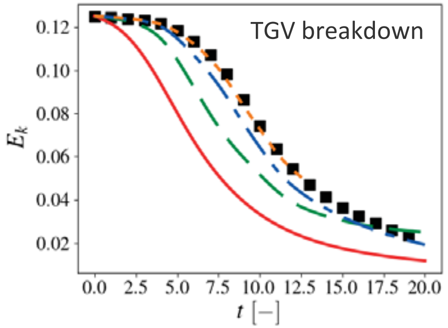

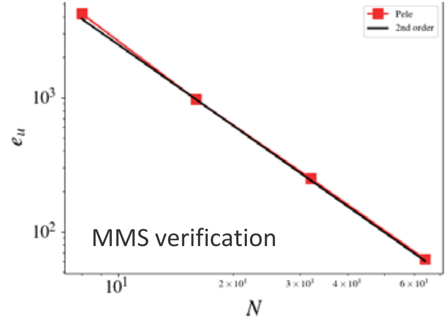

PeleC Scaling

Weak scaling, entire Summit

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

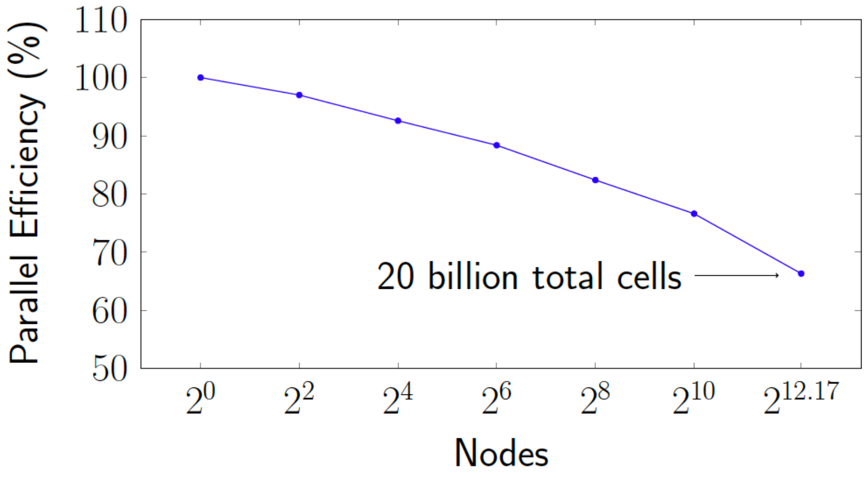

Parallel efficiency drop of 34% at 4608

Strong scaling

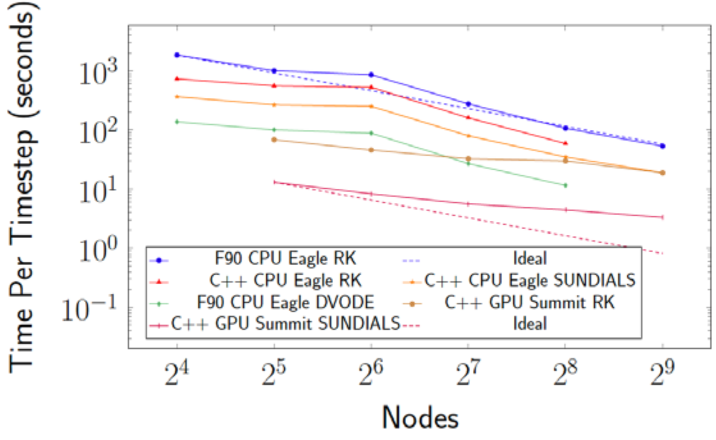

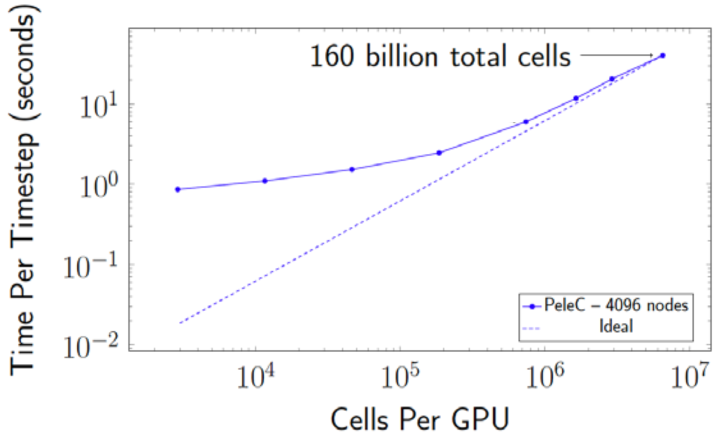

PeleLM(eX) Scaling

- 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

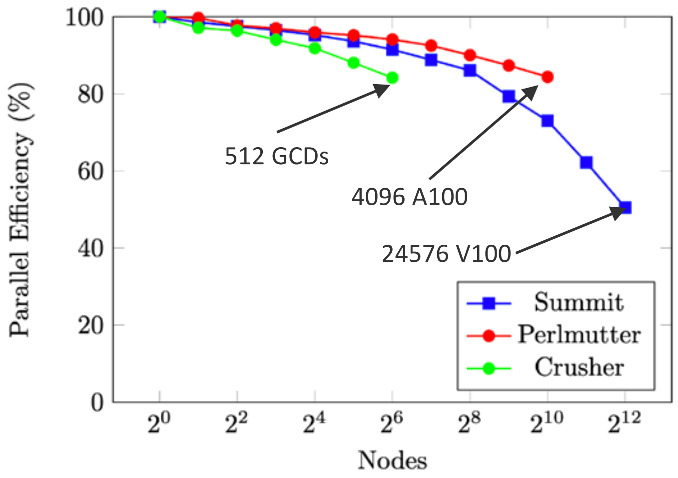

Weak scaling of PeleLM on a 53 species dodecane premixed flame

Wall-bounded turbulent flow models

- Classical periodic channel flow driven by a background pressure gradient
- AMR added in the wall viscous layer
- Algebraic LES closure models available in PeleLMeX: Smagorinsky, ***WALE***, Sigma

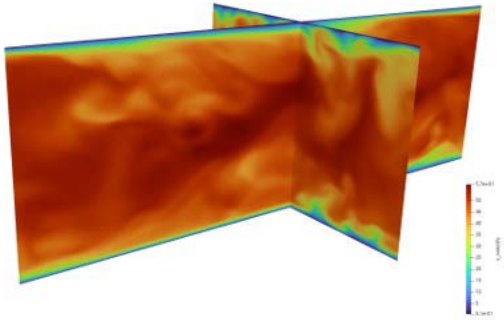

$Re_\tau = 180$ – Streamwise velocity

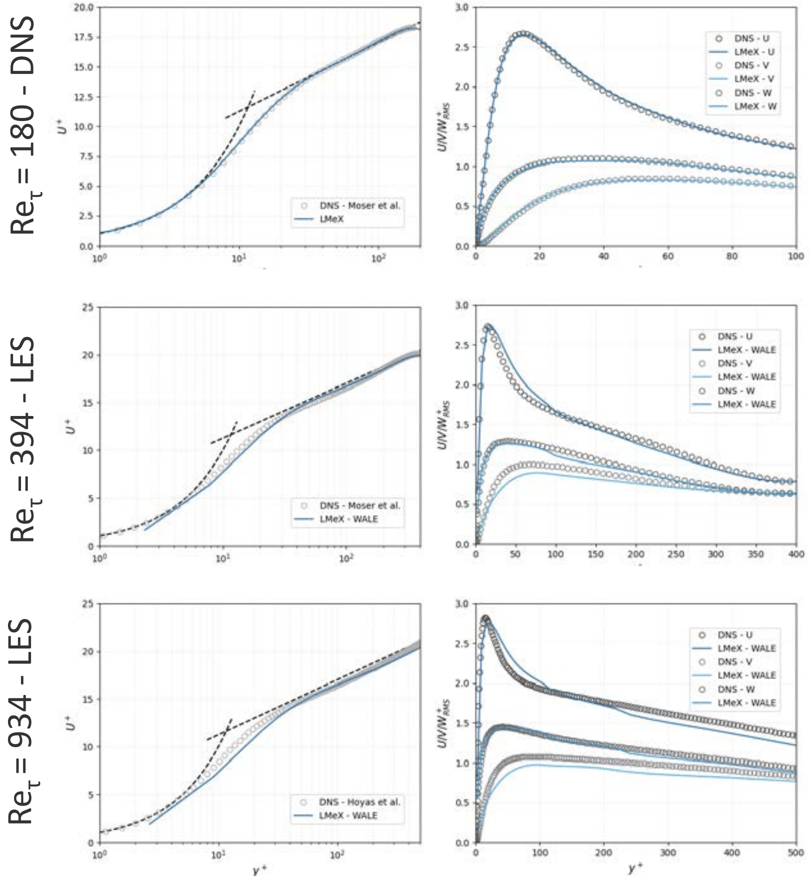

Pele's Adaptive Mesh Refinement

- With subcycled AMR, coarsest level data must "make sense" so that localized fine grid can capture/enhance computed detail
- Compressible and low Mach schemes are strictly mass and energy conservative, guaranteeing that feature locations are mesh independent
- Critically, this means that resolution requirements can depend on simulation goals, e.g.
  - Figure from PeleLM shows that coarse grid correctly predicts location/distortion of vortex-flame interaction
  - Additional resolution sharpens detail of finest scale flame intermediate
- Thus, we can judiciously refine to capture features of interest without resolving everything – the goal of any AMR run
- In turbulent flows, this leads to an implicit LES approach that is extremely useful for inert (regions of) flow

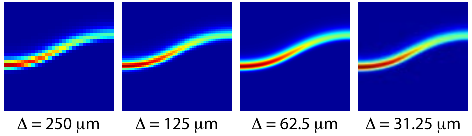

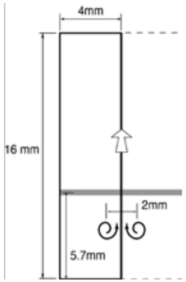

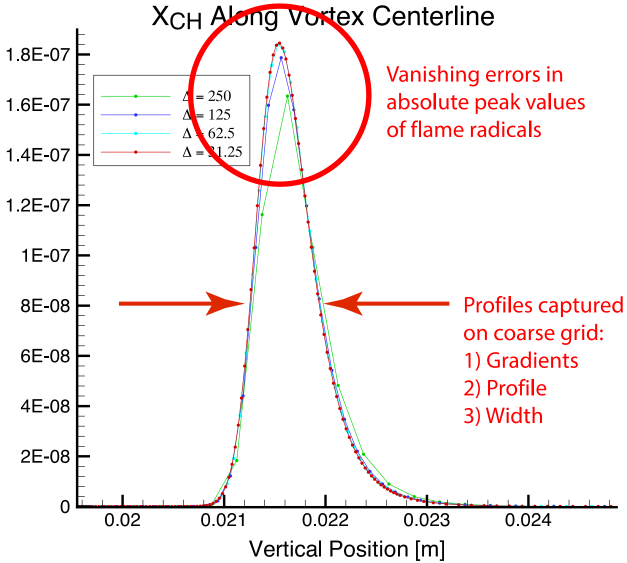

AMR Resolution Requirements

Numerical resolution requirements depend strongly on

1. Physical system investigated
2. Quantity of interest
3. Numerical scheme employed (not just order of accuracy)

For example:

- Lean hydrogen flames are thermodiffusively unstable, and spontaneously develop cellular burning features in 2D and 3D
- Properly capturing details of flame cells requires up to N=32 points across the mean flame thickness
  - ***AMR allows focusing this ONLY at the flame surface***
- In many cases, multi-physics process can adjust to added resolution over relatively short time scales, allowing fine grid additions **locally in space and time**

With a robust method and careful tuning, these observations can result in massive computational savings

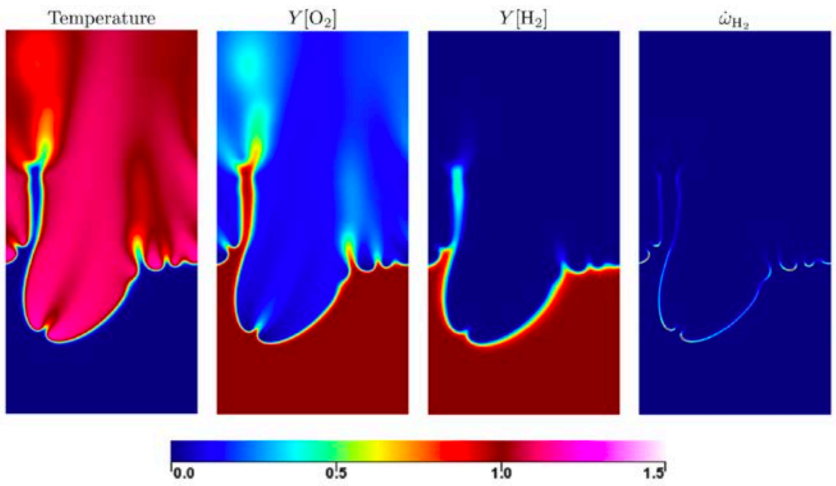

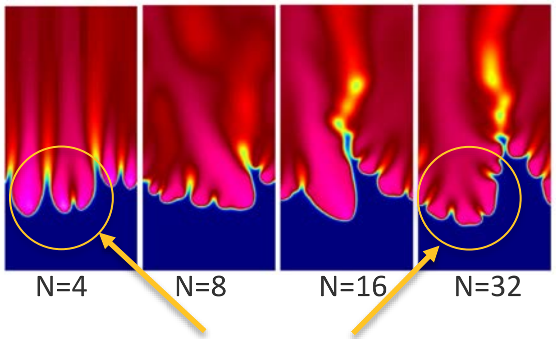

Flame features mesh dependent until N=32

"Expert" Refinement Strategies

Image of work-in-progress courtesy
A Aspden and T Howarth, Newcastle University

Example WIP: Turbulent H2 diffusion jet flame in a "micromixer" device using PeleLMeX on Polaris

H2 fuel
700K, 26 atm
100 m/s
(7mm diameter jet
into 14mm^2 boundary)

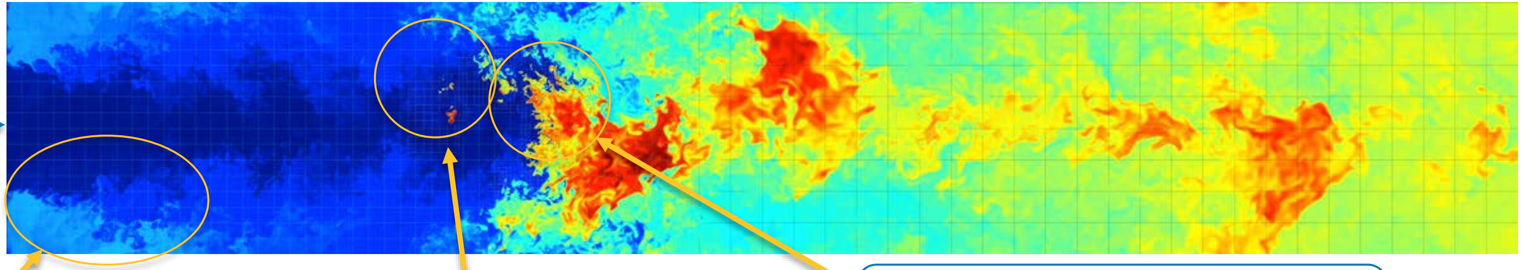

Capture Kolmogorov scales of turbulent mixing

Capture thermal mixing and autoignition

Capture narrow premixed flame surfaces at high pressure

Base grid is 256^2 x 1536, plus 3 levels of factor-of-two refinement at the flame surface, for effective resolution of 2048^2 x 12288 (1% domain refined at level 3)

Combustion-PELE KPP-2 Challenge

Dual pulse injection of combustion fuels w/varying reactivity into engine-relevant geometry

- Baseline enabling simulations to isolate effects of spray evaporation on mixture composition and temperature, use of alternative fuels, and combustion phasing control
- Scoped to consume 2-4 weeks on a significant fraction (~75%) of Frontier's resources

- **Geometry:** Domain relevant to engine cylinder (see figure)
  2.5 cm, flat cylinder head, shaped piston surface
- **Fuel:** *n*-dodecane/methane QSS model (35 species)
  Initial chamber gas: ϕ=0.4 CH_4 turbulent mixture, at 60 atm, 900K
  Jets: Re=14k, mixture *n*-dodecane(45%):chamber-gas(55%)
- **Strategy:** 4 symmetric jets, dual pulse, gas-phase injection
- **Resolution:** 0.85 μm cells (due to 60 atm environment)
- **Sim. Time:** 1 msec (based on jet transit, ignition delay)
- **Flow solver:** PeleC (AMReX-based compressible reacting flow)
- **AMR:** **6 levels** of factor-of-2 refinement
  Level 0-6: volume = (100,23,8.5,1.7,1.1,0.76,0.56)%
  Cell count/level = (0.03,0.06,0.8,0.3,1.5,8.4,49.5) B
  Total cell count ~ 60B (2.4T dofs)

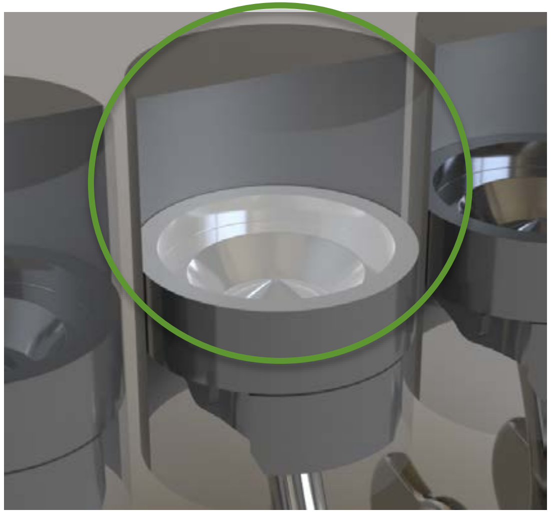

Combustion-PELE KPP-2 Challenge

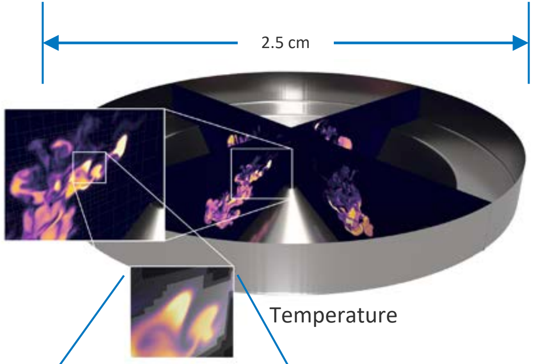

Weak scaling of PeleLMeX on Frontier

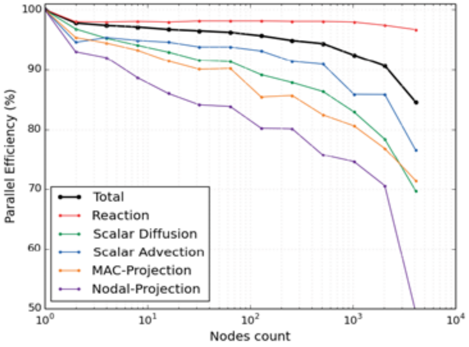

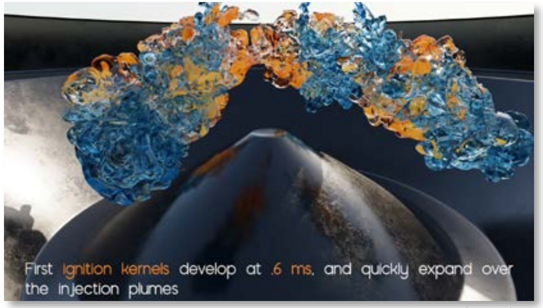

Precursor solution, showing jet-induced vorticity (blue) and ignition kernels (orange)

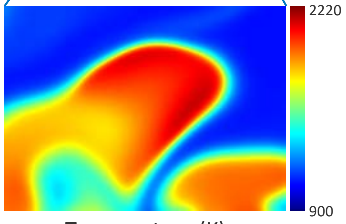

Temperature (K)

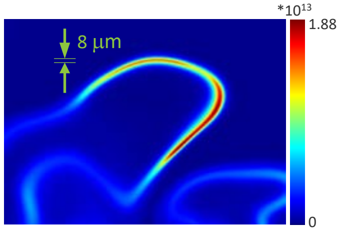

Heat Release (erg/cm^{3}s)

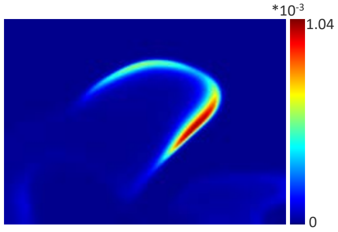

Y-OH

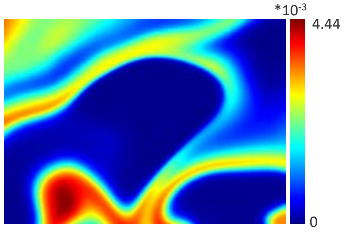

Y-CH2O

7-Level AMR PeleC simulation

- Initial domain: 60atm, 900K, ϕ=0.4 CH_4 turbulent mixture
- Effective resolution: $(32{,}768)^2$ x 8192 with 0.6% of the domain at dx_{Fine} = 0.85 μm
- Four Re=14,000 fuel jets (45% n-dodecane, 55% initial chamber gas)

Notes:

1. Weak scaling data on Frontier for PeleC not yet available
2. Chemistry component scales nearly perfectly (same for both codes)
3. Communication-heavy operations scale poorly at high node counts – we expect similar issues with PeleC:
   - Load imbalance
   - Network-dependent

Pele Applications: WIP

- Supporting certification efforts for (drop-in) sustainable aviation fuels for propulsion
- Enabling clean efficient dispatchable power generation with hydrogen and hydrogen blends in existing infrastructure
- Enabling chemical manufacturing with low energy heat sources from plasma catalysis (e.g., DOE Industrial Heat Shot Earthshot)
- Creating digital twins of hypersonic nonequilibrium reacting flows
- Reliably powering unmanned aerial systems with fuel diversity

Complex multi-physics insights in reactive flows and high-fidelity data for reduced-order models

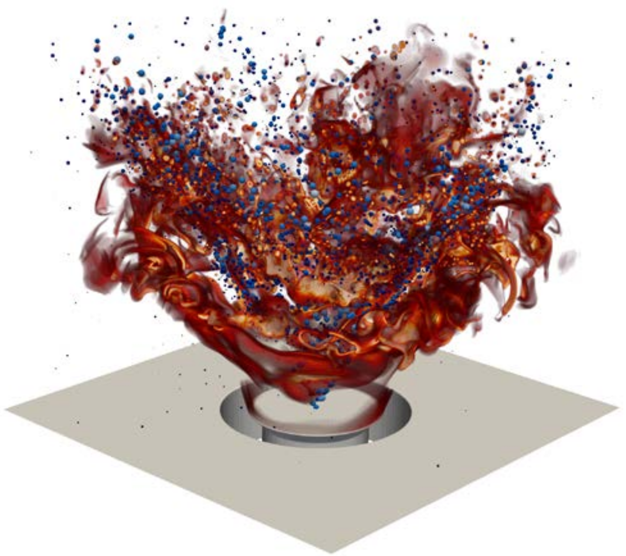

Aero-engine spray flame stabilization with SAF fuel (C1-ATJ) using PeleLMeX
B. Soriano, L Owen and J Chen - SNL

Questions?

www.nrel.gov

NREL/PR-2C00-85474

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.