PELE - Transforming Combustion Science and Technology with Exascale Simulations

Jackie Chen – PI (SNL), Marc Day – CoPI (NREL), Ramanan Sankaran – CoPI (ORNL), Stephen Klippenstein – CoPI (ANL)

Marco Arienti, Marc Henry de Frahan, Swapnil Desai, Lucas Esclapez, Malik Hassanaly, Wenjun Ge, Bill Green, Ray Grout,, Tianfeng Lu, Landon Owen, Martin Rieth, Jon Rood, Michael Martin, Sreejith Nadakkal Appukuttan, Mohammad Rahimi, Bruno Soriano, Nicholas Wimer, Jordi Wolfson-Pou, Shashank Yellapantula

EXASCALE COMPUTING PROJECT

Sandia National Laboratories · NREL National Renewable Energy Laboratory · Berkeley Lab · Oak Ridge National Laboratory · Argonne National Laboratory · University of Connecticut · MIT

Pele project Overview

The Pele project provides a simulation capability for first-principles and near first-principles continuum modeling of practical combustion devices.

- Fundamental turbulence-chemistry interactions
- Exploration of processes leading to higher efficiency, reduced pollutant formation, and operational flexibility
- Sufficient model realism to account for design and operational variations

Pele Combustion Challenge Problem: RCCI

(Reactivity Controlled Compression Ignition)

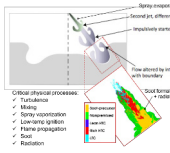

Key questions and sensitivities

- What is the *distribution of reactivity* in a realistic mixture prepared by multiple injections
- How are *reactivity gradients* affected by pulse characteristics:
  - Composition
  - Duration and timing
- How does the reactivity distribution and re-entrainment affect pyrolysis and soot generation near walls
- Are there *identifiable kinetics* that expand the operating map for RCCI combustion

KPP-2 Challenge Problem

Simulate interaction of impulsive injection of fuels, and impulsive prevaporized fuel-rich mixture of injection (baseline).

- Represents baseline enabling simulations to isolate effects of spray evaporation on mixture composition and temperature, effective use of alternative fuels, and strategies to control combustion phasing
- Scoped to consume 2-4 weeks on a significant fraction of Frontier's resources

- **Geometry:** Domain relevant to engine cylinder (see figure)
  ϕ 2.5 cm, flat cylinder head, shaped piston surface
- **Fuels:** *n*-dodecane and methane, ~35+ species
- **Strategy:** 4 symmetric jets, gas-phase injection (baseline), liquid sprays (stretch)
- **Resolution:** 1 μm cells (due to >50 atm environment)
- **Sim. Time:** 1 msec (based on jet transit, ignition delay)
- **Flow solver:** PeleLMeX (AMReX-based low Mach reacting flow)
  PeleC (AMReX-based compressible reacting flow)
- **AMR:** 4 levels of factor-of-2 refinement
  Level 0-4: volume = (100,20,10,5,1)%
  Cell count/level = (0.6,0.9,3.8,15,24) B

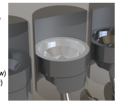

The Pele Combustion Suite

(https://github.com/AMReX-Combustion)

The RCCI Challenge Problem requires

- Time-dependent, adaptive mesh refinement (AMR) for large ranges of spatial and temporal scales in *turbulent reacting flow*
- Engine-relevant *geometry*
- Agile, performant computational kernels
  - Thermo, transport, chemistry (+ data)
  - Soot production/dynamics models
  - Radiation interactions/transport
  - Spray/fluid coupling models

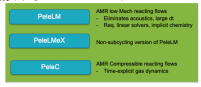

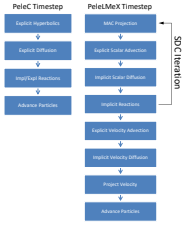

PeleC
- Time-explicit, subcycled AMR
- Second-order discretizations
- Explicitly coupled to chemistry and Lagrangian spray fuel particles
- Ideal and real gas mixture EOS

PeleLMeX
- Semi-implicit, SDC-based iterative non-subcycled time step
- Implicitly coupled to chemistry and Lagrangian spray fuel particles
- Ideal gas mixture EOS

Both
- CEPTR Python-based code generator for thermodynamics, mixture-averaged transport, EOS
- PeleMP – multiphysics coupling for sprays, soot and radiation transport

Pele Applications

Pele code initially limited to simple configurations (periodic boxes, jets, …):

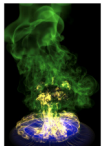

Turbulent puffing helium plume with PeleLM - 2018

Introducing EB allowed more complex & relevant cases:

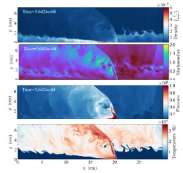

Rotating detonation engine with PeleC - 2022

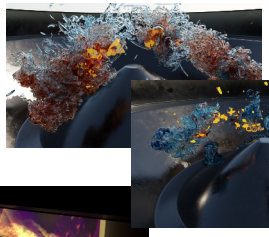

ECP KPP2 Challenge Problem with PeleLMeX (left) and PeleC (bottom) – 2022/2023

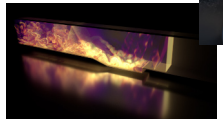

Cavity flame-holder with PeleC - 2020

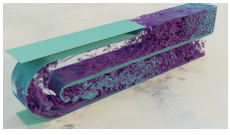

Turbulent U-turn gas turbine component- 2020

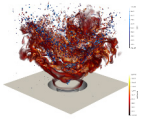

Simulation of SAF in a swirled liquid fueled burner in PeleLMeX - 2022

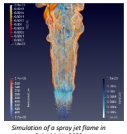

Simulation of a spray jet flame in PeleLMeX - 2022

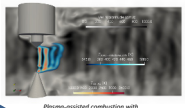

Plasma-assisted combustion with PeleC - 2022

Pele Suite of codes is the leading ExaScale-ready standard for the combustion community:

- Used by National Labs. to tackle upcoming reactive flow challenges
- Used by out academic partners in the US (UCI, UTAustin, USD) and abroad (Montreal (Canada), UNSW (Australia), SINTEF (Norway)
- Github repositories gather the growing community of Pele users

Pele Development Timeline Under ECP

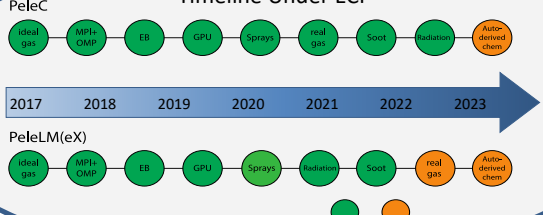

Pele Capabilities Added in 2022

CEPTR Code Generation in PelePhysics: Quasi-steady Model Reduction

Quasi-steady state reduced mechanism construction

Skeletal mechanism 53 species

Reduced mechanism 35 species

	SK53	Red35
Total	1022	533.8
Chemistry	926.8	459.1
Diffusion	53.76	41.14

Reduced QSSA performance on Crusher Flame sheet test, 6.5M cells on 8 GCD
SK53 - 53 species skeletal mechanism
Red35 - 35 species reduced mechanism

Efficient High-Dimensional Sampling

Iterative algorithm for instance selection. Given a large dataset, estimate the probability density function of the data, and devise a sampling probability to uniformly distribute downsampled data in phase space. The density estimation can use a subset of the full dataset. An iterative procedure can be used to compensate for the small dataset used.

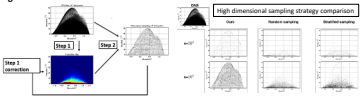

Reduced-Order Manifold Turbulent Combustion Models in PeleLMeX

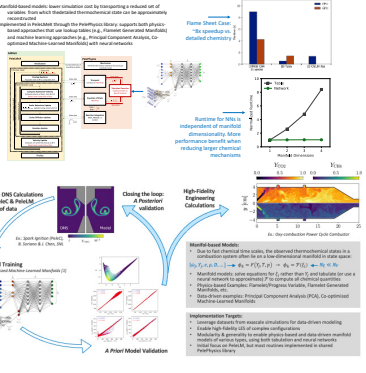

Predictive Automated Combustion Chemistry and Model Reduction

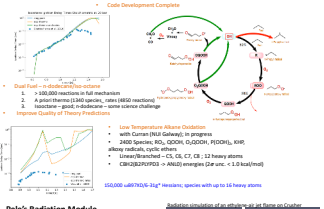

Pele's Radiation Module

- The radiation module is now coupled to PeleC and PeleLMeX
- The performances are benchmarked on Summit and Crusher
- Improvement of robustness to enable practical 1-D simulations

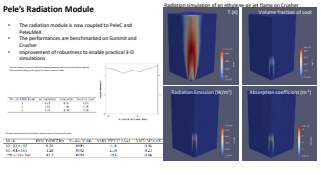

Pele Performance as of 2022

From Q3 2018 onwards, focus on GPUs, initially Nvidia (2019) then AMD/Intel (2021):

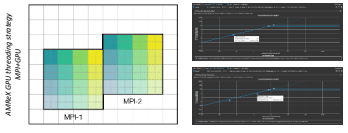

Initial MPI + X port of Pele codes

Upon creation of PeleC and PeleLM (from LMC) in 2017, initial effort concentrated on multicore CPUs (e.g. Intel Xeon Phi) using MPI+OMP:

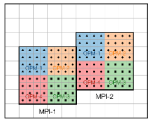

- Code base consists of C++ top-level functions with Fortran compute-intensive kernels
- Chemistry integrated point-wise with DVODE (PeleLM) or in-house RK scheme (PeleC)

- Extensive re-write of both codes, replacing Fortran kernels by C++ lambdas, removing managed memory usage
- Coupling the Pele codes with Sundials
- Both codes show 2x speed-on CPU introducing C++ lambda kernels
- CPU to GPU speed-up on a node basis ranging from 5x to 20x depending on the case/platform

Both codes show good weak scaling capabilities on Nvidia/AMD platforms:

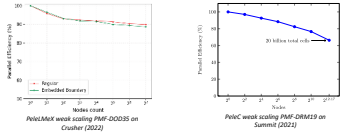

PeleLMeX weak scaling PMF-DOD35 on Crusher (2022)

PeleC weak scaling PMF-DRM19 on Summit (2021)

Pele codes show good scaling on Cori (NERSC) CPU-based system

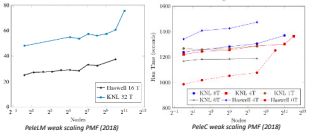

PeleLM weak scaling PMF (2018)

PeleC weak scaling PMF (2018)

Limited strong scaling ability for PeleLM(eX) due to communication-intensive linear solves

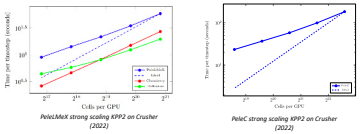

PeleLMeX strong scaling KPP2 on Crusher (2022)

PeleC strong scaling KPP2 on Crusher (2022)

PeleLMeX on CPU and GPU: Chemistry solver dominates execution

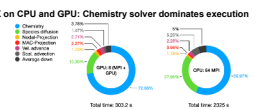

7.5x speed-up on a node basis from CPU to GPU

Performance bottleneck identified and currently under investigation

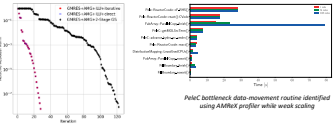

PeleC bottleneck data-movement routine identified using AMReX profiler while weak scaling

Contact: Jackie Chen jhchen@sandia.gov & Marc Day Marcus.Day@nrel.gov

GitHub: https://github.com/AMReX-Combustion

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy (DOE) organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation's exascale computing imperative. The research was performed using computational resources sponsored by DOE Office of Energy Efficiency and Renewable Energy and located at National Renewable Energy Laboratory and resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the DOE Office of Science under Contract No. DE-AC05-00OR22725.

Exascale Computing Project Annual Meeting
Royal Sonesta Galleria, Houston, Texas
January 17-20, 2023

NREL/PO-2C00-85060