

# Performance-Portable Implicit Scale-Resolving Compressible Flow Using libCEED

SIAM CSE 2023

---

James Wright, Jed Brown, Kenneth Jansen, Leila Ghaffari

February 27, 2023

Ann and H.J. Smead Department of Aerospace Engineering Sciences



Smead Aerospace  
UNIVERSITY OF COLORADO **BOULDER**

Copyright © by James Wright, University of Colorado Boulder

# Outline

1. libCEED Overview
2. Compressible Fluid Equations in libCEED
3. Efficient Implicit Timestepping
4. Performance and Results of Flat Plate Boundary Layer Simulation



# libCEED Overview

---



# What is libCEED?

- C library for element-based discretizations
  - Bindings available for Fortran, Rust, Python, and Julia
- Designed for matrix-free operator evaluation
- Portable to different hardware via computational backends
  - Code that runs on CPU also runs on GPU without changes
  - Computational backend selectable at runtime, using runtime compilation
- Geared toward high-order finite element discretizations
- Performance demonstrated for solids in Brown *et al.* 2022<sup>1</sup>
  - Want to apply those methods and lessons-learned to fluids

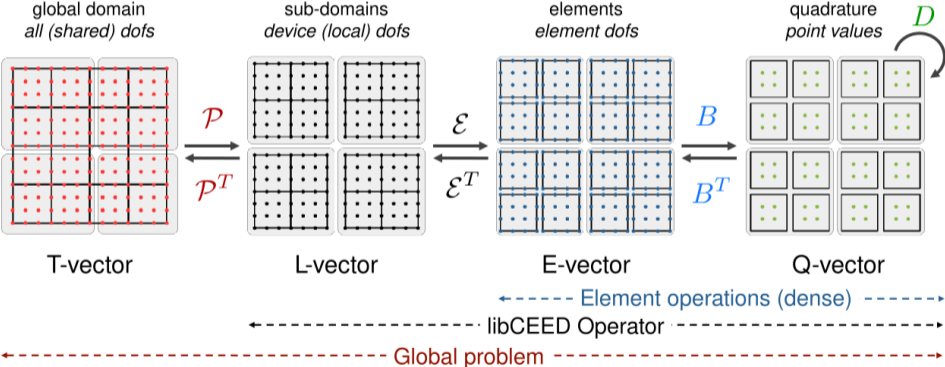
---

<sup>1</sup>*Performance Portable Solid Mechanics via Matrix-Free  $p$ -Multigrid*, Brown *et al.*, arXiv:2204.01722



# Finite Element Operator Decomposition

$$A = \mathcal{P}^T \mathcal{E}^T B^T D B \mathcal{E} \mathcal{P}$$



# Compressible Fluid Equations in libCEED

---



# Compressible Navier-Stokes

$$\mathbf{A}_0 \mathbf{Y}_{,t} + \mathbf{F}_{i,i}(\mathbf{Y}) - S(\mathbf{Y}) = 0$$

for

$$\mathbf{A}_0 \underbrace{\begin{bmatrix} p \\ u_i \\ T \end{bmatrix}}_{\mathbf{Y}} = \begin{bmatrix} \rho \\ \rho u_i \\ \rho e \end{bmatrix}, \quad \mathbf{F}_i(\mathbf{Y}) = \underbrace{\begin{pmatrix} \rho u_i \\ \rho u_i u_j + p \delta_{ij} \\ (\rho e + p) u_i \end{pmatrix}}_{\mathbf{F}_i^{\text{adv}}} + \underbrace{\begin{pmatrix} 0 \\ -\sigma_{ij} \\ -\rho u_j \sigma_{ij} - k T_{,i} \end{pmatrix}}_{\mathbf{F}_i^{\text{diff}}}, \quad S(\mathbf{Y}) = - \begin{pmatrix} 0 \\ \rho \mathbf{g} \\ 0 \end{pmatrix}$$



# Compressible Navier-Stokes for Continuous-Galerkin FEM

Find  $\mathbf{Y} \in \mathcal{S}^h$ ,  $\forall \mathbf{v} \in \mathcal{V}^h$

$$\int_{\Omega} \mathbf{v} \cdot [\mathbf{A}_0 \mathbf{Y}_{,t} - \mathbf{S}(\mathbf{Y})] \, d\Omega - \int_{\Omega} \mathbf{v}_{,i} \cdot \mathbf{F}_i(\mathbf{Y}) \, d\Omega + \int_{\partial\Omega} \mathbf{v} \cdot \mathbf{F}_i(\mathbf{Y}) \cdot \hat{\mathbf{n}}_i \, d\partial\Omega$$
$$+ \underbrace{\int_{\Omega} \mathcal{L}^{\text{adv}}(\mathbf{v}) \tau [\mathbf{A}_0 \mathbf{Y}_{,t} + \mathbf{F}_{i,i}(\mathbf{Y}) - \mathbf{S}(\mathbf{Y})] \, d\Omega}_{\text{SUPG}} = 0$$

Simplify into residual form:

$$\mathcal{G}(\mathbf{Y}_{,t}, \mathbf{Y}) = 0$$
$$\Rightarrow \mathcal{P}^T \mathcal{E}^T \mathbf{B}^T \mathbf{G} \mathbf{B} \mathcal{E} \mathcal{P} \begin{bmatrix} \mathbf{Y}_{,t} \\ \mathbf{Y} \end{bmatrix} = 0$$





# Efficient Implicit Timestepping

---



# Implicit Timestepping

Implicit timestepping requires solving:

$$\frac{d\mathcal{G}(\mathbf{Y},t, \mathbf{Y})}{d\mathbf{Y}} \Delta\mathbf{Y} = -\mathcal{G}(\mathbf{Y},t, \mathbf{Y})$$

- System too large for direct solve  $\rightarrow$  iterative solve
- Krylov subspace methods used most commonly

- Krylov solvers form solution basis from  $\text{span} \left\{ \left[ \frac{d\mathcal{G}(\mathbf{Y},t, \mathbf{Y})}{d\mathbf{Y}} \right]^n \Delta\mathbf{Y} \right\}_{n=0}$

## Bottom Line

Cost of  $\frac{d\mathcal{G}(\mathbf{Y},t, \mathbf{Y})}{d\mathbf{Y}} \Delta\mathbf{Y}$  dominates implicit timestepping cost



# Jacobian Matrix-Vector Multiply Options

How to compute  $\frac{d\mathcal{G}(\mathbf{Y},t, \mathbf{Y})}{d\mathbf{Y}} \Delta\mathbf{Y}$ ?

- Store  $\frac{d\mathcal{G}}{d\mathbf{Y}}$  directly (sparse matrix representation)
  - **Pros:** Opens up preconditioning options
  - **Cons:** Is large, expensive to store
- Finite difference matrix-free approximation:


$$\frac{d\mathcal{G}(\mathbf{Y},t, \mathbf{Y})}{d\mathbf{Y}} \Delta\mathbf{Y} \approx \frac{\mathcal{G}(\mathbf{Y},t, \mathbf{Y} + \epsilon\Delta\mathbf{Y}) - \mathcal{G}(\mathbf{Y},t, \mathbf{Y})}{\epsilon}$$

- **Pros:** Just need a residual evaluation, cheap (in programming and computation)
- **Cons:** Accuracy limited to  $\sqrt{\epsilon_{\text{machine}}}$ , preconditioning require partial assembly



# Exact Matrix-Free Jacobian via CeedOperator

$$\begin{aligned} \frac{d\mathcal{G}}{d\mathbf{Y}} \Delta\mathbf{Y} &= \frac{d}{d\mathbf{Y}} \left[ \overbrace{\mathcal{P}^T \mathcal{E}^T B^T G B \mathcal{E} \mathcal{P}}^{\mathcal{G}(\mathbf{Y},t,\mathbf{Y})} \right] \Delta\mathbf{Y} \\ &= \left[ \mathcal{P}^T \mathcal{E}^T B^T \frac{dG}{d\mathbf{Y}} B \mathcal{E} \mathcal{P} \right] \Delta\mathbf{Y} \end{aligned}$$



- **Pros:** Exact Jacobian matrix-vector product<sup>2</sup>
- **Cons:** Preconditioning requires partial assembly, requires coding Jacobian

<sup>2</sup>Affect of specific terms may be ignored from the Jacobian. This is done for  $d\boldsymbol{\tau}/d\mathbf{Y}$

# Performance and Results of Flat Plate Boundary Layer Simulation

---

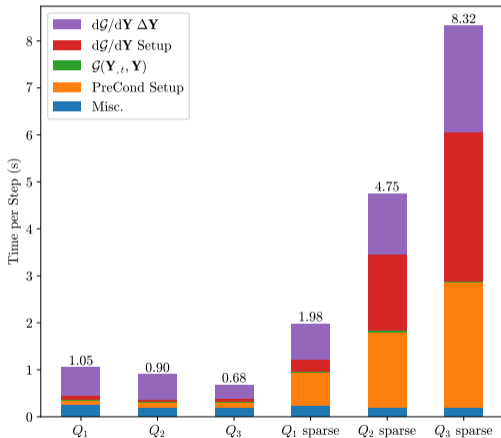


# Problem Description

- Flat plate boundary layer with zero pressure gradient
  - $Re_\theta \approx 970$  boundary layer at inflow,  $M \approx 0.1$
  - Synthetic turbulence generation (STG) used for inflow structures
  - Run at implicit large eddy simulation (ILES) resolution for linears (higher orders may be DNS level, tbd)
- Test 3 different order elements,  $Q_1, Q_2, Q_3$  tensor-product hexes
- Maintain *DOF resolution* (DoFs per physical length/ global DoF count)
- Performance results shown for two nodes of ALCF's Polaris (4x NVIDIA A100 per node)




# Exact Matrix-Free Jacobian vs Sparse



- Sparse  $d\mathcal{G}/d\mathbf{Y} \Delta\mathbf{Y}$  significantly slower than matrix-free
- Time to assemble  $d\mathcal{G}/d\mathbf{Y}$  quite large
- Associated costs rise with element order



# Fluids Performance Analysis




- Time of  $d\mathcal{G}/d\mathbf{Y} \Delta\mathbf{Y}$  decreases as order increases
- $d\mathcal{G}/d\mathbf{Y}$  setup time increases with order
  - Dominant cost is partial matrix assembly for preconditioning





# Results of Flat Plate Boundary Layer



- Spanwise statistics implemented to verify scale-resolving results
- Results *not* converged, but show realistic stress profiles

Zaki et al., 2013, *From Streaks to Spots and on to Turbulence: Exploring the Dynamics of Boundary Layer Transition*  
Schlatter et al., 2010, *Assessment of direct numerical simulation data of turbulent boundary layers*  
Wu et al., 2017, *Transitional-turbulent spots and turbulent-turbulent spots in boundary layers*



# Support and References

- Research supported by US Department of Energy through DE-SC0021411 FASTMath SciDAC Institute
- Argonne Leadership Computing Facility resources used for this research
- We thank PETSc and libCEED developers, especially Jeremy Thompson and Junchao Zhang among many others