

## Recent Advances in Sparse Linear Solver Stacks for Exascale

NCAR Multi-core 9 Workshop

Stephen Thomas<sup>1</sup> Kasia Świrydowicz<sup>1</sup> Shreyas Ananthan<sup>1</sup> Michael Sprague<sup>1</sup>  
Julien Langou<sup>2</sup> Daniel Bielich<sup>2</sup> Ruidong Li<sup>3</sup> Rob Falgout<sup>3</sup> Ulrike M. Yang<sup>3</sup>  
Mark Hoemmen<sup>4</sup> Ichitaro Yamazaki<sup>4</sup> Eric Boman<sup>4</sup> James Elliot<sup>4</sup>

<sup>1</sup>High Performance Algorithms and Complex Fluids Group  
National Renewable Energy Laboratory, Golden CO

<sup>2</sup>Department of Mathematics University of Colorado, Denver, CO

<sup>3</sup>Center for Applied Scientific Computing  
Lawrence Livermore National Laboratories, Livermore, CA.

<sup>4</sup>Sandia National Laboratory,  
Livermore, CA and Albuquerque, NM

National Center for Atmospheric Research  
September 25, 2019

## Scalable Solvers

### Grand challenge ECP simulations

- High-fidelity wind turbine simulations, to capture wake vortex formation
- Nalu-wind Navier-Stokes incompressible flow solver.  $\mathcal{O}(100)$  Billion grid cells
- GMRES with symmetric Gauss-Seidel (SGS), and algebraic multigrid (AMG) preconditioners
- Integration rate of  $\mathcal{O}(10)$  sec time per time step.  $\mathcal{O}(1)$  hour per revolution

### Recent advances in scalable solvers and preconditioners

- Multi-MPI Multi-GPU implementation of **low-sync one-reduce** ICWY MGS-GMRES (**Hypre**)
- Multi-MPI Multi-GPU implementation of **low-sync one-reduce** s-step CA-GMRES (**Trilinos**)
- **2x faster** *Nalu-Wind* matrix assembly with optimal *hypre* CSR memory allocations
- **low-rank** approximate AINV *hypre* smoothers (NVIDIA V100 GPU cuda)
- **25x faster** than cuSparse lower triangular solve on GPU

## Gram-Schmidt and GMRES

Solve  $Ax = b$ , iterative Krylov methods in DOE libraries (*hypre*, Trilinos, PETSc)

- Start with  $x_0$  (initial guess),  $r_0 = b - Ax_0$ ,  $v_1 = r_0 / \|r_0\|_2$
- Search for update to  $x_0$  in *Krylov space*:  $\mathcal{K}_m = \text{span}\{v_1, Av_1, A^2v_1, \dots, A^{m-1}v_1\}$ ,
- Krylov vectors form columns of  $V_{:,j}$ , with  $v_j := V_{:,j}$ ,
- Arnoldi-GMRES solver based on the *QR* factorization of

$$\begin{bmatrix} r_0 & AV_m \end{bmatrix} = V_{m+1} \begin{bmatrix} \|r_0\|e_1 & H_{m+1,m} \end{bmatrix}$$

**Theorem.** One synch per column is sufficient for (MGS, CGS-2) Gram-Schmidt *QR* and Arnoldi-*QR*

- Inverse compact ICWY MGS-GMRES based on lower-triangular solve,  $\mathcal{O}(\varepsilon)\kappa(A)$  orthogonality
- Recursive classical Gram-Schmidt (CGS-2),  $\mathcal{O}(\varepsilon)$  orthogonality

**Corollary.** The backward stability and loss of orthogonality are equivalent to the original algorithms

## One-reduce Gram-Schmidt Algorithms

Björck (1967), Lemma 5.1 and Corollary, Ruhe (1983)

Modified Gram-Schmidt projector

$$P^M a_j = (I - Q_{j-1} T^M Q_{j-1}^T) a_j$$

inverse compact WY form  $T^M = (I + L_{j-1})^{-1}$

Classical Gram-Schmidt with reorthogonalization

$$P^{IC} a_j = (I - Q_{j-1} T^{IC} Q_{j-1}^T) a_j$$

where we are approximating  $P = I - Q_{j-1} (Q_{j-1}^T Q_{j-1})^{-1} Q_{j-1}^T$

$$T^{IC} = (Q_{j-1}^T Q_{j-1})^{-1} = I - L_{j-1} = I - \begin{bmatrix} 0 & 0 \\ -w^T & 0 \end{bmatrix}$$

## Inverse Compact WY Modified Gram-Schmidt

$Q^T Q = I + L + L^T$ , compute  $L^T$  one column per iteration

---

**Algorithm 1** Triangular Solve Modified Gram-Schmidt left-looking (columns)

---

```
1: for  $j = 1, 2, \dots, n$  do
2:    $q_j = a_j$ 
3:   if  $j > 1$  then
4:      $T_{1:j-1, 1:j-1} = Q_{:, 1:j-1}^T q_{j-1}$ 
5:      $L_{1:j-1, 1:j-1} = \text{tril}(T_{1:j-1, 1:j-1}, -1)$ 
6:      $R_{1:j-1, j} = (I + L_{1:j-1, 1:j-1})^{-1} Q_{:, 1:j-1}^T a_j$ 
7:      $q_j = q_j - Q_{:, 1:j-1} R_{1:j-1, j}$ 
8:   end if
9:    $R_{jj} = \|q_j\|_2$ 
10:   $q_j = q_j / R_{jj}$ 
11: end for
```

▷ Synchronization

▷ Lower triangular solve

## Algorithm 2 One reduction MGS-GMRES with Lagged Normalization

```
1:  $r_0 = b - Ax_0, v_1 = r_0.$ 
2:  $v_2 = Av_1$ 
3:  $(V_2, R, L_2) = \text{mgs}(V_2, R, L_1)$ 
4: for  $i = 1, 2, \dots$  do
5:    $v_{i+2} = Av_{i+1}$ 
6:    $[L_{:,i+1}^T, r_{i+2}] = V_{i+1}^T[v_{i+1} \ v_{i+2}]$ 
7:    $r_{i+1,i+1} = \|v_{i+1}\|_2$ 
8:    $v_{i+1} = v_{i+1}/r_{i+1,i+1}$ 
9:    $r_{1:i+1,i+2} = r_{1:i+1,i+2}/r_{i+1,i+1}$ 
10:   $v_{i+2} = v_{i+2}/r_{i+1,i+1}$ 
11:   $r_{i+1,i+2} = r_{i+1,i+2}/r_{i+1,i+1}$ 
12:   $L_{:,i+1}^T = L_{:,i+1}^T/r_{i+1,i+1}$ 
13:   $r_{1:i+1,i+2} = (I + L_{i+1})^{-1} r_{1:i+1,i+2}$ 
14:   $v_{i+2} = v_{i+2} - V_{i+1} r_{1:i+1,i+2}$ 
15:   $H_i = R_{i+1}$ 
16:  Apply Givens rotations to  $H_i$ 
17: end for
18:  $y_m = \text{argmin} \| (H_m y_m - \|r_0\|_2 e_1) \|_2$ 
19:  $x = x_0 + V_m y_m$ 
```

- ▷ Matrix-vector product
- ▷ Global synchronization
  
- ▷ Lagged normalization
  - ▷ Scale for Arnoldi
  - ▷ Scale for Arnoldi
  
- ▷ Lower triangular solve

## Normwise Relative Backward Error

Stopping criterion for MGS-GMRES

- When does MGS-GMRES reach minimum error level ?

$$\frac{\|r_k\|_2}{\|b\|_2} = \frac{\|b - Ax_k\|_2}{\|b\|_2} < tol$$

flattens when  $\|S\| = 1$  and the columns of  $V_k$  become linearly dependent

- Paige, Rozložnik and Strakoš (2006), normwise relative backwards error

$$\frac{\|r_k\|_2}{\|b\|_2 + \|A\|_2 \|x\|_2} < \mathcal{O}(\varepsilon)$$

achieved when  $\|S\|_2 = 1$ .

- Paige notes: For a sufficiently nonsingular matrix

$$\sigma_{\min}(A) \gg n^2 \varepsilon \|A\|_F$$

can employ MGS-GMRES to solve  $Ax = b$  with NBRE stopping criterion.

# Normwise Relative Backward Error

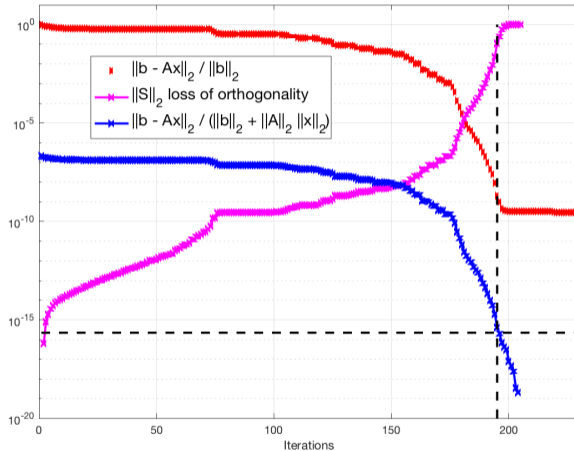



Figure: Greenbaum, Rozložnik and Strakoš (1997). `impcol_e` matrix




# Low-Synch GMRES Multi-MPI Multi-GPU

Świrydowicz et al (2019), Low-Synch Gram-Schmidt and GMRES (NLAA)



(a) Hydre Level-1 BLAS MGS-GMRES




(b) Low-Synch ICWY MGS-GMRES


Figure: Performance on Eagle Skylake + V100 GPU. NREL ABL Mesh  $n = 9M$

- Extended strong scaling roll-off by 4x for ExaWind ABL grid on Skylake + V100.

# Low-Synch GMRES Multi-MPI Multi-GPU



(a) Hyre Gram-Schmidt Times



(b) Low-Synch Gram-Schmidt Time

Figure: Gram-Schmidt kernels time on Eagle Skylake + V100 GPU. NREL ABL Mesh  $n = 9M$ .

Gram-Schmidt times are flat for low-synch algorithm(s)

## McAlister Blade




Figure: McAlister blade. Computational mesh.

# McAlister Blade




Figure: McAlister blade. Tip vortex

# GMRES-AMG reduced iterations with CGS-2 projection

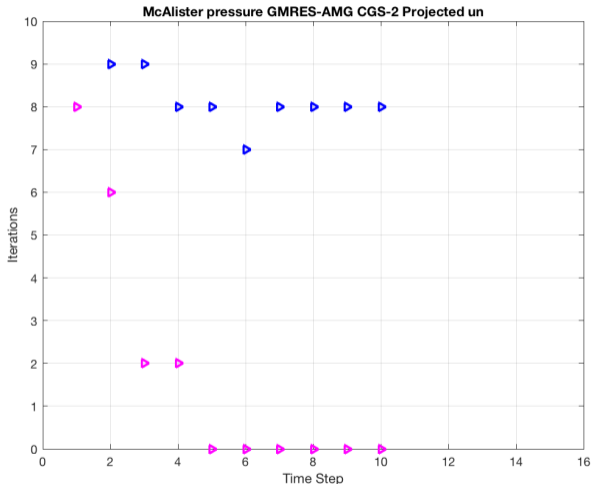


Figure: Projected  $x_0$ . CGS-2. Nalu-Wind pressure solve. McAlister blade.

## Iterative Classical Gram-Schmidt CGS-2

---

### Algorithm 3 Iterated Classical Gram-Schmidt (CGS-2) Algorithm

---

**Input:** Matrices  $Q_{j-1}$  and  $R_{j-1}$ ,  $A_{j-1} = Q_{j-1}R_{j-1}$ ; column vector  $q_j = a_j$

**Output:**  $Q_j$  and  $R_j$ , such that  $A_j = Q_jR_j$

1: **for**  $k = 1, 2$  **do**

2:  $s^{(k-1)} = Q_{j-1}^T a^{(k-1)}$

▷ Global synch

3:  $r_{1:j-1,j}^{(k)} = r_{1:j-1,j}^{(k)} + s^{(k-1)}$

4:  $q_j = q_j - Q_{j-1} s^{(k-1)}$

5: **end for**

6:  $q_j = a^{(k)} / r_{jj} = \|a^{(k)}\|_2$

▷ Global synch

7:  $r_j = [r_{1:j-1,j}, r_{jj}]$

---

## Recursive Classical Gram-Schmidt CGS-2

Orthogonal projector takes the form,  $A = QR$ ,

$$P^{IC} = I - [Q_{j-2}, q_{j-1} - Q_{j-2}Q_{j-2}^T q_{j-1}] T_{j-1} [Q_{j-2}, q_{j-1}]^T$$

Correction matrix  $(Q_{j-2}^T Q_{j-2})^{-1} = T_{j-1} = I - L_{j-1} - L_{j-1}^T$  takes the form

$$T_{j-1} = \begin{bmatrix} T_{j-2} & T_{:,1:j-2} \\ T_{1:j-2,:} & t_{j-1,j-1} \end{bmatrix} = \begin{bmatrix} I & 0 \\ -w^T \alpha^{-1} & 1 \end{bmatrix}$$

$w^T = L_{j-2,:}$ . Thus we have,  $r_{1:j-2,j} = z$ ,  $\alpha = r_{j-1,j-1}^{-1}$  and where

$$\begin{bmatrix} w & z \end{bmatrix} = \begin{bmatrix} Q_{j-2}^T q_{j-1} & Q_{j-2}^T a_j \end{bmatrix}, \quad \begin{bmatrix} r_{j-1,j-1} & r_{j-1,j} \end{bmatrix} = \begin{bmatrix} (q_{j-1}^T q_{j-1} - w^T w) & (q_{j-1}^T a_j - w^T z) \end{bmatrix}$$

Faster solver and more accurate eigenvalue computations based on Lanczos/Arnoldi algorithms  
PETSc-SLEPc. Corrects the Hernandez, Roman, Tomas (2005), (2007) DCGS-2 algorithm

*Rolling Stones Theorem: You can't always get what you want,  
But if you try sometime you find, you get what you need*

---

**Algorithm 4** Classical Gram-Schmidt (CGS-2) Algorithm with Normalization Lag and Reorthogonalization Lag

---

**Input:** Matrices  $Q_{j-1}$  and  $R_{j-1}$ ,  $A_{j-1} = Q_{j-1}R_{j-1}$ ; column vector  $q_j = a_j$

**Output:**  $Q_j$  and  $R_j$ , such that  $A_j = Q_jR_j$

1: if  $j = 1$  return

2:  $[r_{j-1,j-1}, r_{j-1,j}] = q_{j-1}^T [q_{j-1}, q_j]$

3: **if**  $j > 2$  **then**

4:  $[w, z] = Q_{j-2}^T [q_{j-1}, q_j]$ ,

5:  $[r_{j-1,j-1}, r_{j-1,j}] = [r_{j-1,j-1} - w^T w, r_{j-1,j} - w^T z]$

6:  $r_{1:j-2,j-1} = r_{1:j-2,j-1} + w$

7: **end if**

8:  $r_{j-1,j-1} = \{r_{j-1,j-1}\}^{1/2}$

9: **if**  $j > 2$   $q_{j-1} = q_{j-1} - Q_{j-2} w$

10:  $q_{j-1} = q_{j-1} / r_{j-1,j-1}$

11:  $r_{j-1,j} = r_{j-1,j} / r_{j-1,j-1}$

12:  $r_{1:j-2,j} = z$

13:  $q_j = q_j - Q_{j-1} r_j$

▷ Global synch

▷ same global synch

▷ Pythagorean

▷ Lagged  $R$  update

▷ Lagged norm

▷ Lagged Reorthogonalization

▷ Lagged Normalization

▷ Apply recursive projection

---

Pythagorean form (can) mitigate cancellation.  $a^2 - b^2 = (a - b)(a + b)$

Smoktunowicz, Barlow, Langou (2008). Ghysels et al (2013)



## Backward (representation) error

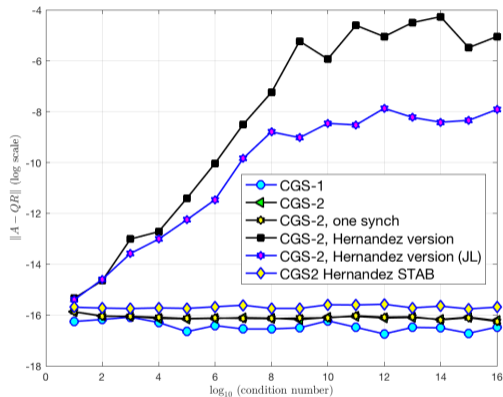


Figure: Representation error for classical Gram-Schmidt

# Orthogonality

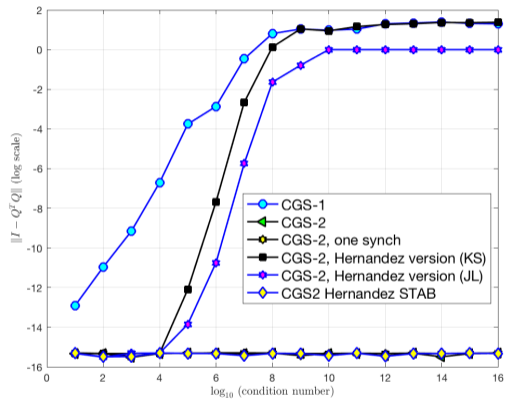


Figure: Orthogonality for classical Gram-Schmidt

# Trilinos s-step GMRES

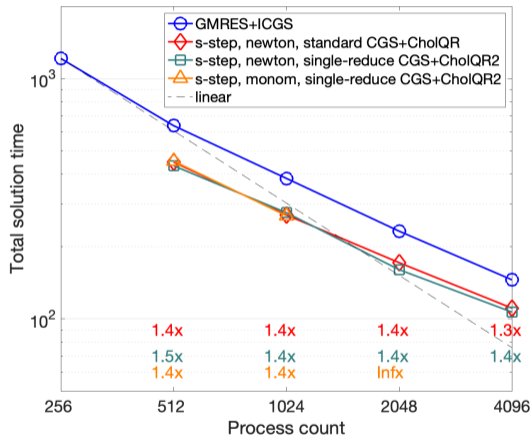


Figure: McAlister Solve Time.

# Trilinos *s*-step GMRES

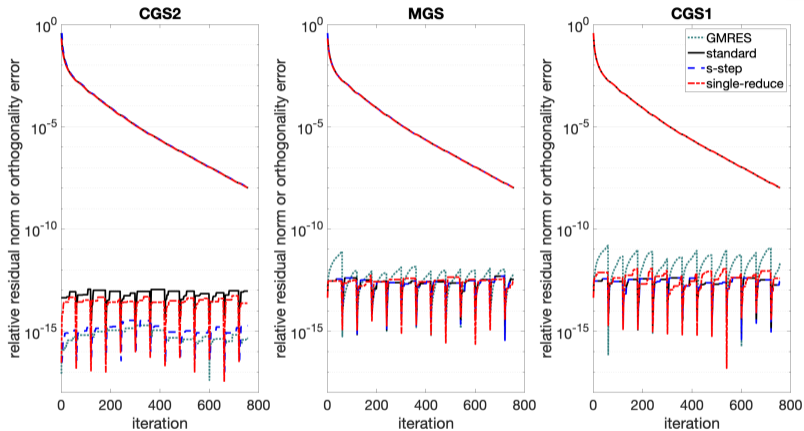


Figure: Convergence for *s*-step GMRES. Laplacian matrix

# Trilinos *s*-step GMRES




Figure: Convergence for *s*-step GMRES. Steam-1 matrix

# Trilinos *s*-step GMRES

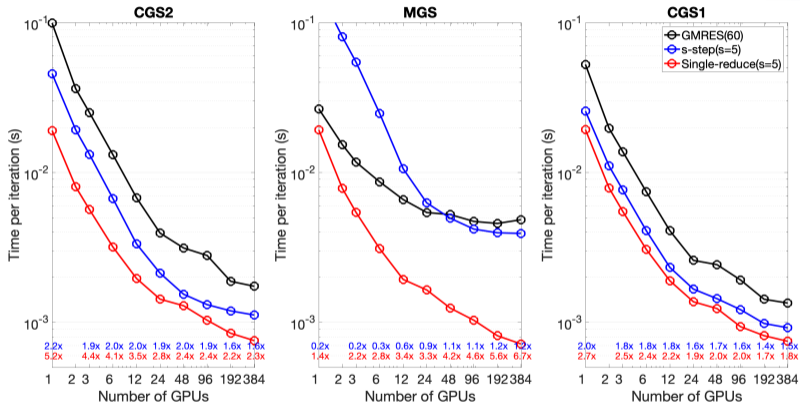


Figure: Time per iteration *s*-step GMRES.

# Trilinos *s*-step GMRES

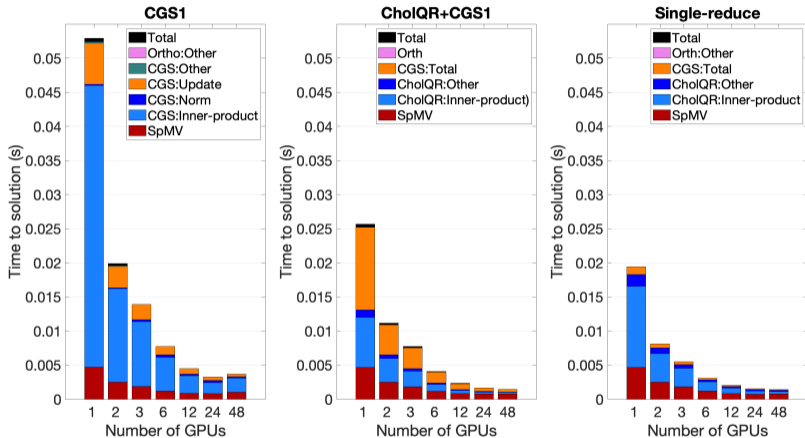


Figure: Time per iteration *s*-step GMRES.

# Trilinos *s*-step GMRES

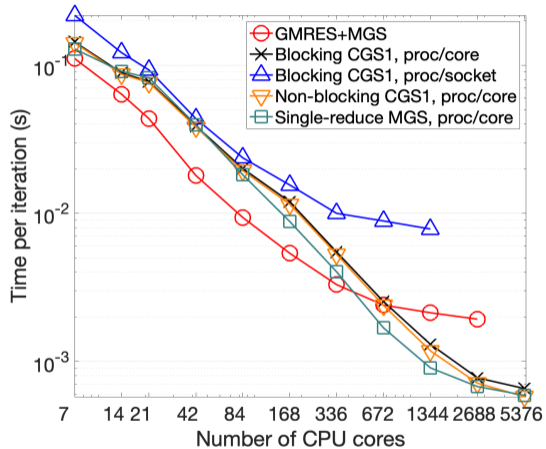


Figure: Strong scaling *s*-step GMRES versus one-sync GMRES.



## Benzi-Tuma AINV Biconjugation

Sparse approximate inverse for non-symmetric  $A$

$$W^T A Z = D, \quad A = W^{-T} D Z^{-1}, \quad A = L D U$$

Analogous to incomplete  $ILU$  factorization.

Benzi and Tuma (1998), Oblique-norm Gram-Schmidt algorithm

- L. Fox (1966), presented non-symmetric algorithm
- S. Thomas (1992) analysed the loss of orthogonality
- J. Kupał, M. Tuma, M. Rozložnik, A. Smoktunowicz (2014)
- R. Lowerey and J. Langou (2018), representation error

## AINV Lower Triangular Smoother

Smoother applied as Richardson iteration local to MPI-rank

$$x_{k+1} = x_k + D^{-1} W^T (b - Ax_k)$$

Hybrid  $M_H$  or  $l_1$  sub-domain smoothers for Hypre-BoomerAMG

$$G = I - M_H^{-1} A, \quad M_H = \text{diag}(B_k)$$

Speed advantage of matrix-vector multiplication on GPU

- Lower sweep count than Gauss-Seidel in V-cycle
- Faster GMRES-AMG convergence
- Triangular matvec is at least  $25\times$  faster the cuSparse triangular solver

# AINV Smoothers for GPU

McAlister 3M DOF. GMRES-AMG pressure solver convergence with AINV smoother

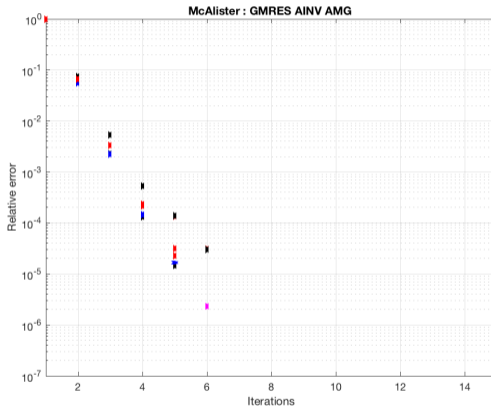


Figure: Pressure GMRES+AMG. AINV (blue 2 sweeps), Gauss-Seidel (black 2 sweeps)

## Challenges to be addressed

### Scientific

- AINV for Trilinos-Muelu smoothers in  $s$ -step CA-GMRES,
- AINV preconditioner to replace SGS in momentum

### Technical

- Testing performance on Summit (multi-node)
- Verify low AMG set-up time on GPU - Trilinos-MueLu and Hypre-BoomerAMG
- Integration with the Nalu-Wind model, Trilinos-Belos-Muelu and Boomer-AMG Hypre Stacks

## Thank you! Questions?

### Acknowledgements:

- This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy 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.
- 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. The NREL Eagle supercomputer was also used extensively for these studies
- NREL is a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, operated by the Alliance for Sustainable Energy, LLC under Contract No. DE-AC36-08GO28308.