EXRSCRLE COMPUTING PROJECT

## Recent Advances in Sparse Linear Solver Stacks for Exascale

NCAR Multi-core 9 Workshop

Stephen Thomas ${ }^{1}$ Kasia Świrydowicz ${ }^{1}$ Shreyas Ananthan ${ }^{1}$ Michael Sprague ${ }^{1}$ Julien Langou ${ }^{2}$ Daniel Bielich ${ }^{2}$ Ruipeng Li ${ }^{3}$ Rob Falgout ${ }^{3}$ Ulrike M. Yang ${ }^{3}$ Mark Hoemmen ${ }^{4}$ Ichitaro Yamazaki ${ }^{4}$ Eric Boman ${ }^{4}$ James Elliot ${ }^{4}$
${ }^{1}$ High Performance Algorithms and Complex Fluids Group National Renewable Energy Laboratory, Golden CO
${ }^{2}$ Department of Mathematics University of Colorado, Denver, CO
${ }^{3}$ Center for Applied Scientific Computing
Lawrence Livermore National Laboratories, Livermore, CA.
${ }^{4}$ 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)
- 2× faster Nalu-Wind matrix assembly with optimal hypre CSR memory allocations
- low-rank approximate AINV hypre smoothers (NVIDIA V100 GPU cuda)
- $25 \times$ faster than cuSparse lower triangular solve on GPU


## Gram-Schmidt and GMRES

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

- Start with $x_{0}$ (initial guess), $r_{0}=b-A x_{0}, v_{1}=r_{0} /\left\|r_{0}\right\|_{2}$
- Search for update to $x_{0}$ in Krylov space: $\mathcal{K}_{m}=\operatorname{span}\left\{v_{1}, A v_{1}, A^{2} v_{1}, \ldots A^{m-1} v_{1}\right\}$,
- Krylov vectors form columns of $V_{\mathrm{i}, j}$, with $v_{i}:=V_{\mathrm{i}, i}$,
- Arnoldi-GMRES solver based on the $Q R$ factorization of

$$
\left[\begin{array}{ll}
r_{0}, & \left.A V_{m}\right]=V_{m+1}\left[\left\|r_{0}\right\| e_{1},\right. \\
H_{m+1, m}
\end{array}\right]
$$

Theorem. One synch per column is sufficient for (MGS, CGS-2) Gram-Schmidt $Q R$ 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

$$
P^{M} a_{j}=\left(I-Q_{j-1} T^{M} Q_{j-1}^{T}\right) a_{j}
$$

inverse compact $W Y$ form $T^{M}=\left(I+L_{j-1}\right)^{-1}$
Classical Gram-Schmidt with reorthogonalization

$$
P^{I C} a_{j}=\left(I-Q_{j-1} T^{I C} Q_{j-1}^{T}\right) a_{j}
$$

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

$$
T^{I C}=\left(Q_{j-1}^{T} Q_{j-1}\right)^{-1}=I-L_{j-1}=I-\left[\begin{array}{cc}
0 & 0 \\
-w^{T} & 0
\end{array}\right]
$$

## 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: \(\mathbf{f o r} j=1,2, \ldots n\) do
    \(q_{j}=a_{j}\)
    if \(j>1\) then
            \(T_{1: j-1, j-1}=Q_{:, 1: j-1}^{T} q_{j-1} \quad \triangleright\) Synchronization
            \(L_{1: j-1,1: j-1}=\operatorname{tril}\left(T_{1: j-1,1: j-1},-1\right)\)
            \(R_{1: j-1, j}=\left(I+L_{1: j-1,1: j-1}\right)^{-1} Q_{:, 1: j-1}^{T} a_{j} \quad \triangleright\) Lower triangular solve
            \(q_{j}=q_{j}-Q_{:, 1: j-1} R_{1: j-1, j}\)
        end if
        \(R_{j j}=\left\|q_{j}\right\|_{2}\)
        \(q_{j}=q_{j} / R_{j j}\)
    end for
```

```
Algorithm 2 One reduction MGS-GMRES with Lagged Normalization
    \(r_{0}=b-A x_{0}, v_{1}=r_{0}\).
    \(v_{2}=A v_{1}\)
    \(\left(V_{2}, R, L_{2}\right)=\operatorname{mgs}\left(V_{2}, R, L_{1}\right)\)
    for \(i=1,2, \ldots\) do
        \(v_{i+2}=A v_{i+1}\)
        \(\left[L_{:, i+1}^{T}, r_{i+2}\right]=V_{i+1}^{T}\left[v_{i+1} v_{i+2}\right]\)
        \(r_{i+1, i+1}=\left\|v_{i+1}\right\|_{2}\)
        \(v_{i+1}=v_{i+1} / r_{i+1, i+1}\)
        \(r_{1: i+1, i+2}=r_{1: i+1, i+2} / r_{i+1, i+1}\)
        \(v_{i+2}=v_{i+2} / r_{i+1, i+1}\)
        \(r_{i+1, i+2}=r_{i+1, i+2} / r_{i+1, i+1}\)
        \(L_{:, i+1}^{T}=L_{:, i+1}^{T} / r_{i+1, i+1}\)
        \(r_{1: i+1, i+2}=\left(I+L_{i+1}\right)^{-1} r_{1: i+1, i+2}\)
14: \(\quad v_{i+2}=v_{i+2}-V_{i+1} r_{1: i+1, i+2}\)
15: \(\quad H_{i}=R_{i+1}\)
16: Apply Givens rotations to \(H_{i}\)
17: end for
\(y_{m}=\operatorname{argmin}\left\|\left(H_{m} y_{m}-\left\|r_{0}\right\|_{2} e_{1}\right)\right\|_{2}\)
\(x=x_{0}+V_{m} y_{m}\)
```

$\triangleright$ Matrix-vector product
$\triangleright$ Global synchronization
$\triangleright$ Lagged normalization
$\triangleright$ Scale for Arnoldi
$\triangleright$ Scale for Arnoldi
$\triangleright$ Lower triangular solve

## Normwise Relative Backward Error

Stopping criterion for MGS-GMRES

- When does MGS-GMRES reach minimum error level ?

$$
\frac{\left\|r_{k}\right\|_{2}}{\|b\|_{2}}=\frac{\left\|b-A x_{k}\right\|_{2}}{\|b\|_{2}}<t o l
$$

flattens when $\|S\|=1$ and the columns of $V_{k}$ become linearly dependent

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

$$
\frac{\left\|r_{k}\right\|_{2}}{\|b\|_{2}+\|A\|_{2}\|x\|_{2}}<\mathcal{O}(\varepsilon)
$$

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 $A x=b$ 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)


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

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


ExAECREME Eampurivs PRODETT

## Low-Synch GMRES Multi-MPI Multi-GPU



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

## 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 \(A_{j}=Q_{j} R_{j}\)
    for \(k=1,2\) do
        \(s^{(k-1)}=Q_{j-1}^{T} a^{(k-1)} \quad \triangleright\) Global synch
        \(r_{1: j-1, j}^{(k)}=r_{1: j-1, j}^{(k)}+s^{(k-1)}\)
        \(q_{j}=q_{j}-Q_{j-1} s^{(k-1)}\)
    end for
    \(q_{j}=a^{(k)} / r_{j j}=\left\|a^{(k)}\right\|_{2}\)
    \(r_{j}=\left[r_{1: j-1, j}, r_{j j}\right]\)
```


## Recursive Classical Gram-Schmidt CGS-2

Orthogonal projector takes the form, $A=Q R$,

$$
P^{\prime C}=I-\left[Q_{j-2}, q_{j-1}-Q_{j-2} Q_{j-2}^{T} q_{j-1}\right] T_{j-1}\left[Q_{j-2}, q_{j-1}\right]^{T}
$$

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

$$
T_{j-1}=\left[\begin{array}{cc}
T_{j-2} & T_{:, 1: j-2} \\
T_{1: j-2,:} & t_{j-1, j-1}
\end{array}\right]=\left[\begin{array}{cc}
l & 0 \\
-w^{T} \alpha^{-1} & 1
\end{array}\right]
$$

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

$$
\left[\begin{array}{ll}
w, & z
\end{array}\right]=\left[\begin{array}{ll}
Q_{j-2}^{T} q_{j-1}, & Q_{j-2}^{T} a_{j}
\end{array}\right], \quad\left[\begin{array}{ll}
r_{j-1, j-1}, & r_{j-1, j}
\end{array}\right]=\left[\begin{array}{ll}
\left(q_{j-1}^{T} q_{j-1}-w^{T} w\right), & \left(q_{j-1}^{T} a_{j}-w^{T} z\right)
\end{array}\right]
$$

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 \(A_{j}=Q_{j} R_{j}\)
    if \(j=1\) return
    \(\left[r_{j-1, j-1}, r_{j-1, j}\right]=q_{j-1}^{T}\left[q_{j-1}, q_{j}\right]\)
    if \(j>2\) then
        \([w, z]=Q_{j-2}^{T}\left[q_{j-1}, q_{j}\right], \quad \triangleright\) same global synch
        \(\left[r_{j-1, j-1}, r_{j-1, j}\right]=\left[r_{j-1, j-1}-w^{\top} w, r_{j-1, j}-w^{\top} z\right]\)
        \(r_{1: j-2, j-1}=r_{1: j-2, j-1}+w\)
    end if
    \(r_{j-1, j-1}=\left\{r_{j-1, j-1}\right\}^{1 / 2}\)
    if \(j>2 q_{j-1}=q_{j-1}-Q_{j-2} w\)
    \(q_{j-1}=q_{j-1} / r_{j-1, j-1}\)
    \(r_{j-1, j}=r_{j-1, j} / r_{j-1, j-1}\)
    \(r_{1: j-2, j}=z\)
    \(q_{j}=q_{j}-Q_{j-1} r_{j}\)
```

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. Kupal, 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

$$
x_{k+1}=x_{k}+D^{-1} W^{\top}\left(b-A x_{k}\right)
$$

Hybrid $M_{H}$ or $I_{1}$ sub-domain smoothers for Hypre-BoomerAMG

$$
G=I-M_{H}^{-1} A, \quad M_{H}=\operatorname{diag}\left(B_{k}\right)
$$

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.

