

POST-MODERN GMRES

STEPHEN THOMAS*, ERIN CARSON†, MIRO ROZLOŽNÍK‡, ARIELLE CARR§, AND KASIA ŚWIRYDOWICZ¶

Abstract. The GMRES algorithm of Saad and Schultz (1986) for nonsymmetric linear systems relies on the Arnoldi expansion of the Krylov basis. The algorithm computes the QR factorization of the matrix $B = [\mathbf{r}_0, AV_k]$ at each iteration. Despite an $\mathcal{O}(\varepsilon)\kappa(B)$ loss of orthogonality, the modified Gram-Schmidt (MGS) formulation was shown to be backward stable in the seminal papers by Paige, et al. (2006) and Paige and Strakoš (2002). Classical Gram-Schmidt (CGS) exhibits an $\mathcal{O}(\varepsilon)\kappa^2(B)$ loss of orthogonality, whereas DCGS-2 (CGS with delayed reorthogonalization) reduces this to $\mathcal{O}(\varepsilon)$ in practice (without a formal proof). We present a post-modern (viz. not classical) GMRES algorithm based on Ruhe (1983) and the low-synch algorithms of Świrydowicz et al (2020) that achieves $\mathcal{O}(\varepsilon) \|Av_k\|_2/h_{k+1,k}$ loss of orthogonality. By projecting the vector Av_k , with Gauss-Seidel relaxation, onto the orthogonal complement of the space spanned by the computed Krylov vectors V_k where $V_k^T V_k = I + L_k + L_k^T$, we can further demonstrate that the loss of orthogonality is at most $\mathcal{O}(\varepsilon)\kappa(B)$. For a broad class of matrices, unlike MGS-GMRES, significant loss of orthogonality does not occur and the relative residual no longer stagnates for highly non-normal systems. The Krylov vectors remain linearly independent and the smallest singular value of V_k is not far from one. We also demonstrate that Henrici’s departure from normality of the lower triangular matrix $T_k \approx (V_k^T V_k)^{-1}$ in the modified Gram-Schmidt projector $P = I - V_k T_k V_k^T$ is an appropriate quantity for detecting the loss of orthogonality. Our new algorithm results in an almost symmetric correction matrix T_k .

1. Introduction. In the present study, linear systems of the form $A\mathbf{x} = \mathbf{b}$, with A an $n \times n$ real-valued matrix, are solved with a Krylov subspace method. Here, let $\mathbf{r}_0 = \mathbf{b} - A\mathbf{x}_0$ denote the initial residual with initial guess \mathbf{x}_0 . Inside GMRES, the Arnoldi- QR algorithm is applied to orthogonalize the basis for the Krylov subspace $\mathcal{K}_m(A, \mathbf{r}_0)$ spanned by the columns of the $n \times m$ matrix, V_m , where $m \ll n$, and after m iterations, produces the $(m + 1) \times m$ Hessenberg matrix, $H_{m+1,m}$, in the Arnoldi expansion such that

$$(1.1) \quad AV_m = V_{m+1}H_{m+1,m}.$$

At each iteration, the Arnoldi algorithm produces a QR factorization of $B = [\mathbf{r}_0, AV_k]$ and the columns of V_k form an orthonormal basis for the Krylov subspace $\mathcal{K}_k(A, \mathbf{r}_0)$ [1]. When the Krylov vectors are orthogonalized via the finite-precision MGS algorithm, their loss of orthogonality is related in a straightforward way to the convergence of GMRES. Orthogonality among the Krylov vectors is effectively maintained until the norm-wise relative backward error approaches the machine precision as discussed in Paige and Strakoš [2] and Paige et al. [1]. The growth of the condition number $\kappa(B)$ is related to the norm-wise relative backward error

$$(1.2) \quad \beta(\mathbf{x}_k) = \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{b}\|_2 + \|A\|_\infty \|\mathbf{x}_k\|_2}, \quad k = 1, 2, \dots$$

and it is observed in exact arithmetic that $\beta(\mathbf{x}_k)\kappa([\mathbf{r}_0, AV_k]) = \mathcal{O}(1)$. The orthogonality of the columns determines the numerical rank of the Krylov basis. However, in finite-precision arithmetic, the columns of V_k are no longer orthogonal, as measured by $\|I - V_k^T V_k\|_F$, and may deviate substantially from machine precision, $\mathcal{O}(\varepsilon)$. When linear independence is completely lost, the relative residual may stagnate at a certain level above $\mathcal{O}(\varepsilon)$ and this occurs when $\|S_k\|_2 = 1$, where $S_k = (I + L_k^T)^{-1} L_k^T$ and L_k is the $k \times k$ strictly lower triangular part of $V_k^T V_k$.

The development of low-synchronization Gram-Schmidt and generalized minimal residual algorithms by Świrydowicz et al. [3] and Bielich et al. [4] was largely driven by applications that need stable, yet scalable solvers. Both the modified and classical Gram-Schmidt algorithms with delayed re-orthogonalization (DCGS-2) are stable for a GMRES solver. Although the DCGS-2 results in an $\mathcal{O}(\varepsilon)$ loss of orthogonality, which suffices for GMRES to converge, stability has not been proven formally. Paige et al. [1] demonstrate that despite $\mathcal{O}(\varepsilon)\kappa(B)$ loss of orthogonality, MGS-GMRES is backward stable for the solution of

*National Renewable Energy Laboratory, Golden, CO

†Charles University, Prague, CZ

‡Czech Academy of Sciences, Institute of Mathematics, Prague, CZ

§Lehigh University, Bethlehem, PA

¶Pacific Northwest Laboratory, WA

linear systems. Here, the condition number of the matrix B is given by $\kappa(B) = \sigma_{\max}(B)/\sigma_{\min}(B)$, where $\sigma_{\max}(B)$ and $\sigma_{\min}(B)$ are the maximum and minimum singular values of the matrix B , respectively.

The purpose of the present work is to derive a post-modern (viz. not classical) formulation of the GMRES algorithm that uses an orthogonalization scheme based on the iterated solution of the normal equations in the Gram-Schmidt projector, as described by Ruhe [5], and the low-synch algorithms introduced by Świrydowicz et al. [3]. The basic idea developed here is to project the vector $A\mathbf{v}_k$ onto the orthogonal complement of the space spanned by the computed Krylov vectors represented by the (properly normalized) columns of $V_k \in \mathbb{C}^{n \times k}$, where $V_k^T V_k = I + L_k + L_k^T$ and $L_k \in \mathbb{C}^{k \times k}$ is strictly lower triangular. Ruhe [5, pg. 597] suggested applying LSQR, whereas Björck [6, pg. 312] recommended conjugate gradients to solve the normal equations (1.3). Instead, we apply two Gauss-Seidel iterations and note that Higham and Knight [7] proved norm-wise backward stability for such stationary iterations. We demonstrate that the loss of orthogonality may then approach $\mathcal{O}(\varepsilon)$ *without any need for reorthogonalization* in the Arnoldi- QR algorithm. Unlike the relative residual for MGS-GMRES, the stagnation shown in Paige and Strakoš [2] does not occur and yet the seminal backward stability result of Paige et al. [1] still applies.

An inverse compact WY modified Gram-Schmidt algorithm is presented in [3] and is based upon the application of the projector

$$P = I - V_k T_k V_k^T, \quad T_k \approx (V_k^T V_k)^{-1}$$

where V_k is again $n \times k$, I is the identity matrix of dimension n , and $T_k^{(1)} = (I + L_k)^{-1}$ is a $k \times k$ lower triangular matrix. To obtain a low-synch MGS algorithm, or one global reduction per GMRES iteration, the normalization is delayed to the next iteration. The matrix $T_k^{(1)}$ is obtained from the strictly lower triangular part of $V_k^T V_k$, denoted L_k . Note that because V_k has almost orthonormal columns, the norm of L_k is small, and $T_k^{(1)}$ is nearby I (here, the identity matrix of dimension k).

A Neumann series expansion for the inverse of the lower triangular matrix, $T_k^{(1)}$, results from the inverse compact WY form of the projector P , as discussed in Thomas et al. [8]. A post-modern GMRES (PM-GMRES) formulation based upon the matrix polynomial associated with two Gauss-Seidel iterations for the normal equations

$$(1.3) \quad V_{k-1}^T V_{k-1} \mathbf{r}_{1:k-1,k} = V_{k-1}^T A \mathbf{v}_k$$

can be derived, with the projector and correction matrix

$$(1.4) \quad P = I - V_k T_k^{(2)} V_k^T, \quad T_k^{(2)} = (I + L_k)^{-1} [I - L_k^T (I + L_k)^{-1}],$$

where the rows of L_{k-1} are constructed from the matrix-vector products $L_{k-1,:} = \mathbf{v}_{k-1}^T V_{k-1}$. The Neuman series for $(I + L_k)^{-1}$ is finite because the matrix L_k is nilpotent, as originally noted by Ruhe [5]. The loss of orthogonality will then be shown to follow $\mathcal{O}(\varepsilon) \|A\mathbf{v}_k\|_2/h_{k+1,k}$. For extremely ill-conditioned and non-normal matrices, the convergence history of the PM-GMRES algorithm has been found to be identical to the original MGS-GMRES algorithm introduced by Saad and Schultz [9], with the exception that the (implicit) relative residual continues to decrease monotonically without stagnating and thus matches the Householder (HH) GMRES of Walker [10].

Contributions. In this paper, we present a new formulation of the MGS-GMRES algorithm of Saad and Schultz [9], and prove backward stability of the solutions, thus extending the results of Paige et al. [1]. The computed Krylov vectors maintain orthogonality to the level of machine precision by projection onto their orthogonal complement and this is accomplished with two Gauss-Seidel iterations in the low-synchronization Gram-Schmidt algorithms of Świrydowicz et al. [3]. The triangular matrix $T_k^{(1)}$ is an approximation of the matrix $(Q_k^T Q_k)^{-1}$ and is recognized as a Neuman series [8]. Two Gauss-Seidel iterations results in $T_k^{(2)}$ that is almost symmetric (see Appendix). This matrix was split and applied across two iterations to achieve $\mathcal{O}(\varepsilon)$ orthogonality for DCGS-2. Giraud et al. [11] demonstrated how a rank- k correction could be applied in an *a posteriori* step to improve orthogonality by computing the polar decomposition of Q_{k-1} , the matrix exhibited by Björck and Paige [12]. The algorithms described herein allow us to maintain the orthogonality of the properly normalized Q_{k-1} (from the

QR decomposition of A) at each iteration instead of as a post-processing step. Our paper is organized as follows. Low-synchronization Gram-Schmidt algorithms for the QR factorization are reviewed in Section 2 and two Gauss-Seidel iterations are applied to solve the normal equations (2.1). A rounding error analysis of the new Gram-Schmidt algorithm is presented in Section 3, leading to bounds on the representation error and the orthogonality of the columns of V_{k-1} . Section 4 extends these results to the Arnoldi- QR algorithm and post-modern GMRES. The relationship with Henrici’s departure from normality is explored in Section 5. Finally, numerical experiments on challenging problems studied over the past thirty-five years are presented in Section 6.

Notation. Lowercase bold letters denote a column vector and uppercase letters are matrices (e.g. \mathbf{v} and A , respectively). We use a_{ij} to represent the (i, j) scalar entry of a matrix A , and \mathbf{a}_k denotes the k -th column of A . A_k is a block partition up to the k -th column of a matrix. Subscripts indicate the approximate solution and corresponding residual (e.g. \mathbf{x}_k and \mathbf{r}_k) of an iterative method at step k . Throughout this article, the notation U_k (or L_k) and U_s (or L_s) will explicitly refer to *strictly upper/lower triangular matrices*.¹ Vector notation indicates a subset of the rows and/or columns of a matrix; e.g. $V_{1:k+1,1:k}$ denotes the first $m + 1$ rows and k columns of the matrix V and the notation $V_{:,1:k}$ represents the entire row of the first k columns of V . $H_{k+1,k}$ represents an $(k + 1) \times k$ matrix, and in particular H refers to a Hessenberg matrix. In cases where standard notation in the literature is respected that may otherwise conflict with the aforementioned notation, this will be explicitly indicated. Bars denote computed quantities such as \bar{R}_{k-1} , while \hat{Q}_{k-1} indicates a correctly (properly) normalized matrix as was introduced in Björck and Paige [12]. The matrix Q_k refers to the $A = QR$ factorization via Gram-Schmidt, whereas V_k refers to the orthogonal matrix produced in the Arnoldi- QR expansion.

2. Low-synchronization Gram-Schmidt Algorithms. Krylov subspace methods for solving linear systems are often required for extreme-scale physics simulations on parallel machines with many-core accelerators. Their strong-scaling is limited by the number and frequency of global reductions in the form of MPI_AllReduce operations and these communication patterns are expensive [13]. Low-synchronization algorithms are based on the ideas of Ruhe [5], and are designed such that they require only one reduction per iteration to normalize each vector and apply projections. The Gram-Schmidt projector applied to \mathbf{a}_k , the k -th column of A , in the $A = QR$ factorization can be written as

$$P\mathbf{a}_k = \mathbf{a}_k - Q_{k-1}\mathbf{r}_{1:k-1,k} = \mathbf{a}_k - Q_{k-1} (Q_{k-1}^T Q_{k-1})^{-1} Q_{k-1}^T \mathbf{a}_k,$$

where the vector $\mathbf{r}_{1:k-1,k}$ is the solution of the normal equations

$$(2.1) \quad Q_{k-1}^T Q_{k-1} \mathbf{r}_{1:k-1,k} = Q_{k-1}^T \mathbf{a}_k.$$

Ruhe [5] established that the iterated MGS algorithm employs a *multiplicative* Gauss-Seidel relaxation scheme with matrix splitting $Q_{k-1}^T Q_{k-1} = M_{k-1} - N_{k-1}$, where $M_{k-1} = I + L_{k-1}$ and $N_{k-1} = -L_{k-1}^T$ and $M_{k-1}^{-1} = T_{k-1}^{(1)}$. The iterated CGS is an *additive* Jacobi relaxation.

The inverse compact WY form for MGS was derived in Świrydowicz et al. [3], with strictly lower triangular matrix L_{k-1} . Specifically, these inverse compact WY algorithms batch the inner-products together and compute one row of L_{k-1} as

$$(2.2) \quad L_{k-1,:} = \mathbf{q}_{k-1}^T Q_{k-1}.$$

The resulting projector $P^{(1)}$ and correction matrix are given by

$$(2.3) \quad P^{(1)} = I - Q_{k-1} T_{k-1}^{(1)} Q_{k-1}^T, \quad T_{k-1}^{(1)} = (I + L_{k-1})^{-1}$$

and corresponds to one Gauss-Seidel iteration for the normal equations (2.1). The MGS algorithm with two Gauss-Seidel iterations is given as Algorithm 2.1 below. The compact and inverse compact WY forms of the correction matrix appearing in (2.3) are derived in the Appendix.

¹We note that the distinction between these two notations is crucial. For U_k , the size of the strictly upper triangular matrix changes with k , whereas the size of U_s remains fixed.

Algorithm 2.1 Inverse Compact WY MGS Algorithm with Normalization Lag

Input: Matrices Q_{k-1} , and R_{k-1} , $A_{k-1} = Q_{k-1}R_{k-1}$; column vector \mathbf{a}_k ; matrix L_{k-2}

Output: Q_k and R_k , such that $A_k = Q_k R_k$; L_{k-1} , \mathbf{w}_k

- 1: if $k = 1$ return
 - 2: $[L_{k-1,:}, \mathbf{r}_k] = [\mathbf{q}_{k-1}^T Q_{k-1}, Q_{k-1}^T \mathbf{a}_k]$ ▷ Global synchronization
 - 3: $\mathbf{r}_{k-1,k-1} = \|\mathbf{w}_{k-1}\|_2$
 - 4: $\mathbf{q}_{k-1} = \mathbf{w}_{k-1}/\mathbf{r}_{k-1,k-1}$ ▷ Lagged normalization
 - 5: $\mathbf{r}_{1:k-1,k}^{(0)} = \mathbf{r}_{1:k-1,k}/\mathbf{r}_{k-1,k-1}$ ▷ Scale for Arnoldi
 - 6: $L_{k-1,:} = L_{k-1,:}/\mathbf{r}_{k-1,k-1}$ ▷ Scale for Arnoldi
 - 7: $\mathbf{r}_{1:k-1,k}^{(1)} = (I + L_{k-1})^{-1} \mathbf{r}_{1:k-1,k}^{(0)}$ ▷ First Gauss-Seidel
 - 8: $\mathbf{r}_{1:k-1,k}^{(2)} = \mathbf{r}_{1:k-1,k}^{(1)} - (I + L_{k-1})^{-1} L_{:,k-1}^T \mathbf{r}_{1:k-1,k}^{(1)}$ ▷ Second Gauss-Seidel
 - 9: $\mathbf{w}_k = \mathbf{a}_k - Q_{k-1} \mathbf{r}_{1:k-1,k}^{(2)}$
-

The projection computed after one Gauss-Seidel iteration, where $\mathbf{r}_{1:k-1,k}^{(1)}$ is a solution of the normal equations, can be expressed as

$$(2.4) \quad \mathbf{w}_k = \mathbf{a}_k - Q_{k-1} \mathbf{r}_{1:k-1,k}^{(1)} = \mathbf{a}_k - Q_{k-1} (I + L_{k-1})^{-1} \mathbf{r}_{1:k-1,k}^{(0)}.$$

After two Gauss-Seidel iterations, the projection is given by

$$(2.5) \quad \mathbf{w}_k = \mathbf{a}_k - Q_{k-1} \mathbf{r}_{1:k-1,k}^{(2)} = \mathbf{a}_k - Q_{k-1} \mathbf{r}_{1:k-1,k}^{(1)} - Q_{k-1} (I + L_{k-1})^{-1} L_{k-1}^T \mathbf{r}_{1:k-1,k}^{(1)}.$$

The correction matrix for two Gauss-Seidel iterations is then found to be

$$T_{k-1}^{(2)} = M_{k-1}^{-1} [I + N_{k-1} M_{k-1}^{-1}],$$

which is close to a symmetric matrix (see Appendix for a proof) and is applied across Steps 7 and 8 in Algorithm 2.1. Algorithm 2.1 then consists of the following recurrence:

$$(2.6) \quad \begin{aligned} \mathbf{r}_{1:k-1,k}^{(0)} &= Q_{k-1}^T \mathbf{a}_k, \\ \mathbf{r}_{1:k-1,k}^{(1)} &= (I + L_{k-1})^{-1} \mathbf{r}_{1:k-1,k}^{(0)} = M_{k-1}^{-1} \mathbf{r}_{1:k-1,k}^{(0)}, \\ \mathbf{r}_{1:k-1,k}^{(2)} &= \mathbf{r}_{1:k-1,k}^{(1)} - (I + L_{k-1})^{-1} L_{k-1}^T \mathbf{r}_{1:k-1,k}^{(1)} = \mathbf{r}_{1:k-1,k}^{(1)} + M_{k-1}^{-1} N_{k-1} \mathbf{r}_{1:k-1,k}^{(1)}. \end{aligned}$$

3. Backward Stability. We now present a rounding error analysis of the low-synch modified Gram-Schmidt algorithm with two Gauss-Seidel iterations. Björck [14] and Björck and Paige [6] analysed the traditional MGS algorithm. Our approach is more akin to the analysis of CGS presented by Giraud et al. [15]. We begin with the derivation of the representation error for the QR factorization using an induction argument. In particular, the triangular solve $M_k \mathbf{x} = \mathbf{b}$ implied by Step 7 of Algorithm 2.1 with $M_k = I + L_k$ is backward stable as shown by Higham [16], where

$$(M_k + E_k) \mathbf{x} = \mathbf{b}, \quad \|E_k\|_2 \leq \mathcal{O}(\varepsilon) \|M_k\|_2.$$

The associated error terms and bounds for the recurrence relations (2.6) are then given by

$$\begin{aligned} \bar{\mathbf{r}}_{1:k-1,k}^{(0)} &= \tilde{Q}_{k-1}^T \mathbf{a}_k + \mathbf{e}_k^{(0)}, & \|\mathbf{e}_k^{(0)}\|_2 &\leq \mathcal{O}(\varepsilon) \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2 \\ (M_{k-1} + E_{k-1}^{(1)}) \bar{\mathbf{r}}_{1:k-1,k}^{(1)} &= \bar{\mathbf{r}}_{1:k-1,k}^{(0)}, & \|E_{k-1}^{(1)}\|_2 &\leq \mathcal{O}(\varepsilon) \|M_{k-1}\|_2 \\ \bar{\mathbf{r}}_{1:k-1,k}^{(1/2)} &= N_{k-1}^T \bar{\mathbf{r}}_{1:k-1,k}^{(1)} + \mathbf{e}_k^{(1)}, & \|\mathbf{e}_k^{(1)}\|_2 &\leq \mathcal{O}(\varepsilon) \|N_{k-1}\|_2 \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2 \\ (M_{k-1} + E_{k-1}^{(2)}) \bar{\mathbf{r}}_{1:k-1,k}^{(3/2)} &= \bar{\mathbf{r}}_{1:k-1,k}^{(1/2)}, & \|E_{k-1}^{(2)}\|_2 &\leq \mathcal{O}(\varepsilon) \|M_{k-1}\|_2 \\ \bar{\mathbf{r}}_{1:k-1,k}^{(2)} &= \bar{\mathbf{r}}_{1:k-1,k}^{(1)} - \bar{\mathbf{r}}_{1:k-1,k}^{(3/2)} + \mathbf{e}_k^{(2)}, & \|\mathbf{e}_k^{(2)}\|_2 &\leq \mathcal{O}(\varepsilon) (\|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2 + \|\bar{\mathbf{r}}_{1:k-1,k}^{(3/2)}\|_2) \end{aligned}$$

It follows that the recurrence relations above can be re-written in computed form as

$$\bar{\mathbf{r}}_{1:k-1,k}^{(1)} = M_{k-1}^{-1} \tilde{Q}_{k-1}^T \mathbf{a}_k + M_{k-1}^{-1} \mathbf{f}_k^{(1)}, \quad \mathbf{f}_k^{(1)} = \mathbf{e}_k^{(0)} - E_{k-1}^{(1)} \bar{\mathbf{r}}_{1:k-1,k}^{(1)}.$$

After one Gauss-Seidel iteration, the following bound holds

$$\begin{aligned} \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2 &\leq (M_{k-1} + E_{k-1}^{(1)})^{-1} (\tilde{Q}_{k-1}^T \mathbf{a}_k + \mathbf{e}_k^{(0)}) \\ &\leq \frac{\|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2 (1 + \mathcal{O}(\varepsilon))}{\sigma_{k-1}(M_{k-1}) - \|E_{k-1}^{(1)}\|_2} \leq \frac{\|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2 (1 + \mathcal{O}(\varepsilon)) \|M_{k-1}^{-1}\|_2}{1 - \mathcal{O}(\varepsilon)\kappa(M_{k-1})}. \end{aligned}$$

The computed intermediate vector $\bar{\mathbf{r}}_{1:k-1,k}^{(3/2)}$ satisfies the recurrence

$$\bar{\mathbf{r}}_{1:k-1,k}^{(3/2)} = M_{k-1}^{-1} N_{k-1} \bar{\mathbf{r}}_{1:k-1,k}^{(1)} + M_{k-1}^{-1} \mathbf{f}_k^{(2)}, \quad \mathbf{f}_k^{(2)} = \mathbf{e}_k^{(1)} - E_{k-1}^{(2)} \bar{\mathbf{r}}_{1:k-1,k}^{(3/2)},$$

and an upper bound is given by

$$\begin{aligned} \|\bar{\mathbf{r}}_{1:k-1,k}^{(3/2)}\|_2 &\leq \frac{\|M_{k-1}^{-1} N_{k-1}\|_2 \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2 + \|\mathbf{e}_k^{(1)}\|_2}{1 - \|M_{k-1}^{-1}\|_2 \|E_{k-1}^{(2)}\|_2} \\ &\leq \frac{\|M_{k-1}^{-1} N_{k-1}\|_2 + \mathcal{O}(\varepsilon)\|N_{k-1}\|_2}{1 - \mathcal{O}(\varepsilon)\kappa(M_{k-1})} \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2. \end{aligned}$$

In order to bound the 2-norm of the iteration matrix, to have $\|M_{k-1}^{-1} N_{k-1}\|_2 < 1$, we assume that $\|N_{k-1}\|_2 < 1/2$, so that $\|M_{k-1}^{-1} N_{k-1}\|_2 \leq \|N_{k-1}\|_2 / (1 - \|N_{k-1}\|_2) < 1$ and the result follows. The individual terms are then bounded as follows

$$\begin{aligned} \|\mathbf{f}_k^{(1)}\|_2 &\leq \|\mathbf{e}_k^{(0)}\|_2 + \|E_{k-1}^{(1)}\|_2 \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2 \\ &\leq \mathcal{O}(\varepsilon) \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2 + \frac{\mathcal{O}(\varepsilon)\kappa(M_{k-1}) \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2}{1 - \mathcal{O}(\varepsilon)\kappa(M_{k-1})}, \end{aligned}$$

from which it follows that

$$\|M_{k-1}^{-1}\|_2 \|\mathbf{f}_k^{(1)}\|_2 \leq \frac{\mathcal{O}(\varepsilon)\kappa(M_{k-1}) \|M_{k-1}^{-1}\|_2}{1 - \mathcal{O}(\varepsilon)\kappa(M_{k-1})} \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2.$$

For the second term, a bound follows from

$$\begin{aligned} \|\mathbf{f}_k^{(2)}\|_2 &\leq \|\mathbf{e}_k^{(1)}\|_2 + \|E_{k-1}^{(2)}\|_2 \|\bar{\mathbf{r}}_{1:k-1,k}^{(3/2)}\|_2 \\ &\leq \mathcal{O}(\varepsilon) \|N_{k-1}\|_2 \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2 + \frac{\mathcal{O}(\varepsilon) \|M_{k-1}\|_2 (1 + \mathcal{O}(\varepsilon) \|N_{k-1}\|_2) \|\bar{\mathbf{r}}_{1:k-1,k}^{(1)}\|_2}{1 - \mathcal{O}(\varepsilon) \|M_{k-1}\|_2}, \end{aligned}$$

and therefore the last term in (3) is given by

$$\|M_{k-1}^{-1}\|_2 \|\mathbf{f}_k^{(2)}\|_2 \leq \frac{\mathcal{O}(\varepsilon)\kappa(M_{k-1})}{[1 - \mathcal{O}(\varepsilon)\kappa(M_{k-1})]^2} \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2.$$

Finally, the columns of the computed \bar{R} after two iterations of Gauss-Seidel relaxation satisfy

$$\begin{aligned} \bar{\mathbf{r}}_{1:k-1,k}^{(2)} &= M_{k-1}^{-1} \tilde{Q}_{k-1}^T \mathbf{a}_k + M_{k-1}^{-1} N_{k-1} M_{k-1}^{-1} \tilde{Q}_{k-1}^T \mathbf{a}_k + \mathbf{e}_k \\ \mathbf{e}_k &= M_{k-1}^{-1} \mathbf{f}_k^{(1)} + M_{k-1}^{-1} N_{k-1} M_{k-1}^{-1} \mathbf{f}_k^{(1)} + M_{k-1}^{-1} \mathbf{f}_k^{(2)}, \end{aligned}$$

and they are bounded according to

$$(3.1) \quad \|\bar{\mathbf{r}}_{1:k-1,k}^{(2)}\|_2 \leq \|M_{k-1}^{-1} (I + N_{k-1} M_{k-1}^{-1})\|_2 \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2 + \|\mathbf{e}_k\|_2.$$

The computed form of the projection in Step 9 of Algorithm 2.1 can be written as

$$(3.2) \quad \bar{\mathbf{w}}_k = \mathbf{a}_k - \tilde{Q}_{k-1} \bar{\mathbf{r}}_{1:k-1,k}^{(2)} + \mathbf{f}_k,$$

where

$$\|\mathbf{f}_k\|_2 \leq \mathcal{O}(\varepsilon) \left[\|\mathbf{a}_k\|_2 + \|\tilde{Q}_{k-1}\|_2 \|\bar{\mathbf{r}}_{1:k-1,k}^{(2)}\|_2 \right].$$

From equation (3.1), it therefore follows that because $\|I + M_{k-1}^{-1}N_{k-1}\|_2 < 2$,

$$\begin{aligned} \|\mathbf{f}_k\|_2 &\leq \mathcal{O}(\varepsilon) \left[\|\mathbf{a}_k\|_2 + \|M_{k-1}^{-1}\|_2 \|\tilde{Q}_{k-1}\|_2^2 \|\mathbf{a}_k\|_2 \right] + \mathcal{O}(\varepsilon^2) \\ (3.3) \quad &\leq \mathcal{O}(\varepsilon) \|\mathbf{a}_k\|_2 \left[1 + \frac{\|\tilde{Q}_{k-1}\|_2^2}{1 - \|N_{k-1}\|_2} \right] \leq \frac{\mathcal{O}(\varepsilon) \|\tilde{Q}_{k-1}\|_2^2 \|\mathbf{a}_k\|_2}{1 - \|N_{k-1}\|_2}. \end{aligned}$$

The lower bound for the norm of $\bar{\mathbf{w}}_k$ is now determined from the computed factorization

$$(3.4) \quad A_{k-1} = \tilde{Q}_{k-1} \bar{R}_{k-1} - F_{k-1},$$

$$(3.5) \quad \mathbf{a}_k = \tilde{Q}_{k-1} \bar{\mathbf{r}}_{1:k-1,k}^{(2)} + \bar{\mathbf{w}}_k - \mathbf{f}_k.$$

Combining these into an augmented matrix form, we obtain

$$A_k = \begin{bmatrix} A_{k-1} & \mathbf{a}_k \end{bmatrix} = \tilde{Q}_{k-1} \begin{bmatrix} \bar{R}_{k-1} & \bar{\mathbf{r}}_{1:k-1,k}^{(2)} \end{bmatrix} + \begin{bmatrix} -F_{k-1} & \bar{\mathbf{w}}_k - \mathbf{f}_k \end{bmatrix}$$

and the representation error in equations (3.4) and (3.5) has the lower bound

$$\|F_k\| = \left\| \begin{bmatrix} F_{k-1} & \mathbf{f}_k \end{bmatrix} \right\|_2 + \|\bar{\mathbf{w}}_k\|_2 \geq \left\| \begin{bmatrix} -F_{k-1} & \bar{\mathbf{w}}_k - \mathbf{f}_k \end{bmatrix} \right\|_2 \geq \sigma_k(A_k).$$

Thus, the 2-norm of the projected vector $\bar{\mathbf{w}}_k$ appearing in (3.2) is bounded from below by

$$\|\bar{\mathbf{w}}_k\|_2 \geq \sigma_k(A_k) - \|F_k\|_2,$$

where $\bar{\mathbf{r}}_{k-1,k-1} = \|\bar{\mathbf{w}}_k\|_2$ and $\bar{\mathbf{q}}_k = \bar{\mathbf{w}}_k / \bar{\mathbf{r}}_{k-1,k-1}$.

4. Loss of Orthogonality. Ruhe [5, pg. 597] notes that the deviation from orthogonality of the vector \mathbf{q}_k is represented by the residual of the normal equations (2.1). Therefore, the loss of orthogonality at the k -th iteration of the Gram-Schmidt algorithm with two Gauss-Seidel iterations is measured by the 2-norm of the vector

$$(4.1) \quad \|\tilde{Q}_{k-1}^T \bar{\mathbf{q}}_k\|_2 = \|\tilde{Q}_{k-1}^T \bar{\mathbf{w}}_k\|_2 / \|\bar{\mathbf{w}}_k\|_2.$$

Using (3.2), the projected vector is expanded as

$$\tilde{Q}_{k-1}^T \bar{\mathbf{w}}_k = \tilde{Q}_{k-1}^T \left[\mathbf{a}_k - \tilde{Q}_{k-1} \bar{\mathbf{r}}_{1:k-1,k}^{(2)} + \mathbf{f}_k \right].$$

We have already established in (3.3) that \mathbf{f}_k is bounded, however, at this point the residual of the normal equations and its 2-norm τ can be identified, and these are assumed to be solved to accuracy $\mathcal{O}(\varepsilon)$, i.e.,

$$(4.2) \quad \tau = \left\| \tilde{Q}_{k-1}^T \mathbf{a}_k - \tilde{Q}_{k-1}^T \tilde{Q}_{k-1} \bar{\mathbf{r}}_{1:k-1,k}^{(2)} \right\|_2 \leq \mathcal{O}(\varepsilon) \|\tilde{Q}_{k-1}\|_2 \|\mathbf{a}_k\|_2.$$

The largest singular value $\|\tilde{Q}_{k-1}\|_2^2 \leq 1 + 2\|L_{k-1}\|_2$ and $1 - \|N_{k-1}\|_2 \leq 1$. Therefore, provided that the inequality (4.2) holds, an appropriate bound on the loss of orthogonality at the k -th iteration of Algorithm 2.1 is given by

$$\|\tilde{Q}_{k-1}^T \bar{\mathbf{q}}_k\|_2 \lesssim \mathcal{O}(\varepsilon) \|\mathbf{a}_k\|_2 / \bar{\mathbf{r}}_{k,k}^{(2)}.$$

When applied to the linear system $A\mathbf{x} = \mathbf{b}$, assume $\mathbf{x}_0 = \mathbf{0}$, $\mathbf{r}_0 = \mathbf{b}$, $\|\mathbf{b}\|_2 = \rho$ and $\mathbf{v}_1 = \mathbf{b}/\rho$, the Arnoldi algorithm produces an orthonormal basis for the Krylov vectors spanned by the columns of the matrix V_m . After m iterations, in exact arithmetic, the algorithm produces the expansion

$$AV_m = V_{m+1} H_{m+1,m}, \quad V_{m+1}^T V_{m+1} = I.$$

Consider the properly normalized matrix \tilde{V}_k at iteration k , with Krylov vectors as columns. Recall that the strictly lower triangular matrix L_k is computed incrementally one row per iteration using (2.2) and is obtained from the relation

$$\tilde{V}_k^T \tilde{V}_k = I + L_k + L_k^T.$$

The essential result is based on the QR factorization of the matrix

$$B = [\mathbf{r}_0, AV_k] = V_{k+1} [\mathbf{e}_1 \rho, H_{k+1,k}].$$

Our backward error analysis is now applied to the Arnoldi- QR algorithm, where the vector AV_k is projected onto the orthogonal complement of the properly normalized Krylov vectors. Two Gauss-Seidel iterations reduce the error in equation (3.2)

$$(4.3) \quad \bar{\mathbf{w}}_{k+1} = A\tilde{\mathbf{v}}_k - \tilde{V}_k \bar{h}_{1:k,k} + \mathbf{f}_k.$$

For Arnoldi- QR , we multiply A times \mathbf{v}_k at iteration k . In effect, this is MGS (or our iterated Gauss-Seidel MGS) with \mathbf{a}_k in Algorithm 2.1 replaced by AV_k . By applying the Arnoldi- QR recurrence with AV_{k-1} in the Gram-Schmidt algorithm, we define the column vectors $\bar{\mathbf{h}}_{1:k-1,k} \equiv \bar{\mathbf{r}}_{1:k-1,k}^{(2)}$, and the representation error F_m for the computed Arnoldi expansion, after m iterations

$$A\tilde{V}_m - \tilde{V}_{m+1} \bar{H}_{m+1,m} = F_m$$

is a matrix that grows by one column in size at each iteration. It is important to note that the Arnoldi expansion represents the underlying recurrence relation based on the Krylov subspace, $\mathcal{K}_m(A, \mathbf{r}_0)$ and equation (4.3) is one column of the Arnoldi expansion. The bound $\|\mathbf{f}_k\|_2 \leq \mathcal{O}(\varepsilon) \|\bar{\mathbf{a}}_k\|_2$ in the error analysis is given above in (3.3) and thus a bound on the loss-of-orthogonality is given by

$$(4.4) \quad \|I - \tilde{V}_m^T \tilde{V}_m\|_2 \lesssim \mathcal{O}(\varepsilon) \|A\tilde{\mathbf{v}}_m\|_2 / \bar{h}_{m+1,m}.$$

In practice, this bound is near $\|N_m\|_2$ and is related to the departure from normality

$$\text{dep}(\tilde{V}_{m-1}^T \tilde{V}_{m-1}) = \|L_m\|_F.$$

5. Departure from Normality. A normal matrix $A \in \mathbb{C}^{n \times n}$ satisfies $A^*A = AA^*$. In the present study, Henrici's definition of the departure from normality

$$(5.1) \quad \text{dep}(A) = \sqrt{\|A\|_F^2 - \|\Lambda\|_F^2},$$

where $\Lambda \in \mathbb{C}^{n \times n}$ is the diagonal matrix containing the eigenvalues of A , [17] serves as a useful metric for the loss of orthogonality. While we find practical use for this metric for measuring the degree of (non)normality of a matrix, there are of course other useful metrics to describe (non)normality. We refer the reader to [17–19] and references therein. In particular, we have that the loss of orthogonality for MGS is signaled by the departure from normality as follows

$$\text{dep}(I + L_m)^2 = \|I + L_m\|_F^2 - \|I\|_F^2 = \|I\|_F^2 + \|L_m\|_F^2 - m = \|L_m\|_F^2.$$

We can then relate this to $S_m = (I + U_m)^{-1} U_m$ as we observe in practice that $\|S_m\|_F \approx \|L_m\|_F$ up to the first order in ε (noting that due to symmetry $\|L_m\|_F = \|U_m\|_F$).

Ipsen [18] characterizes the convergence of GMRES in terms of the (non)normality of the matrix A . A numerical measure of normality was presented by Ruhe [20], where normality implies that the eigenvalues and singular values of a matrix A are the same, namely $\sigma_i(A) = |\lambda_i(A)|$ for a certain ordering of the eigenvalues. Ruhe [20] establishes the connection between Henrici's metric and the distance between the eigenvalues and singular values

$$(5.2) \quad \text{dep}^2(A) = \sum_i \sigma_i^2(A) - \sum_i |\lambda_i(A)|^2.$$

The loss of orthogonality in MGS-GMRES is signaled by $\sigma_{\max}^2(I + L_{m-1})$ increasing above one.

6. Post-Modern GMRES. The MGS-GMRES orthogonalization algorithm is the QR factorization of the matrix B formed by adding a new column to V_k at each iteration. For the PM-GMRES Algorithm 6.1 below, note that bold-face (e.g. \mathbf{r}_k) denotes the residual vector at iteration k and subscripting (e.g. $\mathbf{r}_{1:k,k+1}$) denotes the corresponding column elements in the upper triangular matrix R .

The MGS-GMRES algorithm was proven to be backward stable for the solution of linear systems $A\mathbf{x} = \mathbf{b}$ in [1] and orthogonality is maintained to $\mathcal{O}(\varepsilon)\kappa(B)$, depending upon the condition number of the matrix $B = [\mathbf{r}_0, AV_k]$. The normalization for the Krylov vector \mathbf{v}_k at iteration k represents the delayed scaling of the vector \mathbf{v}_{k-1} in the matrix-vector product $\mathbf{v}_k = A\mathbf{w}_{k-1}$. Therefore, an additional Step 8 is required in the low-synch Algorithm 6.1, $\mathbf{r}_{1:k-1,k} = \mathbf{r}_{1:k-1,k}/\mathbf{r}_{k-1,k-1}$ and $\mathbf{v}_{k-1} = \mathbf{v}_k/\mathbf{r}_{k-1,k-1}$. The diagonal element $r_{k-1,k-1}$ of the R matrix, corresponds to $h_{k,k-1}$ in the Arnoldi- QR expansion (1.1) and is updated after the projection in Step 10 of the PM-GMRES Algorithm 6.1. Algorithm 6.1 applies the matrix $T_k^{(2)}$ in Step 9 as two Gauss-Seidel iterations in the form of a lower triangular solve and matrix-vector multiply where the matrix has dimension k .

Algorithm 6.1 Low-synchronization PM-GMRES

```

1:  $\mathbf{r}_0 = \mathbf{b} - A\mathbf{x}_0$ ,  $\mathbf{w}_0 = \mathbf{r}_0$ .
2: for  $k = 1, 2, \dots, m$  do
3:    $\mathbf{v}_k = A\mathbf{w}_{k-1}$  ▷ Matrix-vector product
4:    $[L_{k-1,:}, \mathbf{r}_k] = [ \mathbf{w}_{k-1}^T V_{k-1}, V_{k-1}^T \mathbf{v}_k ]$  ▷ Global AllReduce
5:    $\mathbf{r}_{k-1,k-1} = \|\mathbf{v}_k\|_2$ 
6:    $\mathbf{r}_{1:k,k+1}^{(0)} = \mathbf{r}_{1:k,k+1}/\mathbf{r}_{k-1,k-1}$  ▷ Scale for Arnoldi
7:    $\mathbf{v}_{k-1} = \mathbf{v}_k/\mathbf{r}_{k-1,k-1}$  ▷ Scale for Arnoldi
8:    $L_{k-1,:} = L_{k-1,:}/\mathbf{r}_{k-1,k-1}$ 
9:    $\mathbf{r}_{1:k-1,k}^{(2)} = T_{k-1}^{(2)} \mathbf{r}_{1:k-1,k}^{(0)}$  ▷ projection
10:   $\mathbf{w}_k = \mathbf{v}_k - V_{k-1} \mathbf{r}_{1:k-1,k}^{(2)}$ 
11:   $H_k = R_{k-1}$ 
12:  Apply Givens rotations to  $H_k$ 
13: end for
14:  $\mathbf{y}_m = \operatorname{argmin}(\|H_m \mathbf{y}_m - \|\mathbf{r}_0\|_2 \mathbf{e}_1\|)_2$ 
15:  $\mathbf{x} = \mathbf{x}_0 + V_m \mathbf{y}_m$ 

```

The low-synch DCGS-2 algorithm introduced by Świrydowicz [3] was recently employed to compute the Arnoldi- QR expansion in the Krylov-Schur eigenvalue and GMRES solvers by Bielich et al. [4]. Although a formal backward stability analysis is still in development, the algorithm exhibits desirable numerical characteristics including the computation of invariant subspaces of maximum size for the Krylov-Schur algorithm of Stewart [21]. However, the Arnoldi algorithm itself was modified with representation error correction terms in order to account for the delayed re-orthogonalization and it was not obvious how to adapt this approach to restarted GMRES. In contrast, the PM-GMRES algorithm presented herein does not require any modifications to the original Arnoldi- QR iteration.

Our first experiment illustrates that the bounds derived in the previous sections are representative and properly capture the behavior of the PM-GMRES algorithm in finite precision for a broad class of matrices. In particular, we examine the **fs1836** matrix studied by Paige and Strakoš [2]. Our Figure 1 should be compared with Figures 7.1 and 7.3 of their (2002) paper. In order to demonstrate empirically that the backward error is reduced by the Gauss-Seidel iteration matrix $M_{k-1}^{-1}N_{k-1}$, the quantity $\|S^{(2)}\|_2$ is computed, as defined by Paige et al. [1], which measures the loss of orthogonality for two Gauss-Seidel iterations. The spectral radius ρ_k of the matrix $M_k^{-1}N_k$ is highly correlated with and follows the metric $\|S_k^{(2)}\|_2$. The spectral radius ρ_k , $\|I - \bar{V}_k^T \bar{V}_k\|_F$ and $\|S_k^{(2)}\|_2$ are also plotted in Figure 1. Unlike the relative residual in Figure 7.1 of Paige and Strakoš [2] which stagnates near iteration forty-three at 1×10^{-7} before reaching $\mathcal{O}(\varepsilon)$, the (implicit) relative residual for PM-GMRES continues to decrease monotonically and the norm-wise relative backward error (1.2) reaches machine precision level, $\beta(\mathbf{x}_k) = 6.6 \times 10^{-17}$ at iteration fifty. Most notably, the 2-norm of A is large, where $\|A\|_2 = 1.2 \times 10^9$.




Fig. 1: fs1836 matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe’s metric (4.1), versus LHS of the loss of orthogonality relation (4.4). Spectral radius ρ_k of iteration matrix $M_k^{-1}N_k$ and Paige’s $\|S_k\|_2$.

7. Computational Complexity and Parallel Communication. The GMRES algorithm of Saad and Schultz [9] has a thirty-five year history and alternative formulations of the basic algorithm have been proposed over that time frame. A comprehensive review of these is presented by Zou [22]. In particular, pipelined s -step and block algorithms have been proposed which are better able to hide latency in parallel implementations and are described in Yamazaki et al. [23]. In the case of the DCGS2 algorithm, the symmetric correction matrix T_{k-1} was derived in Appendix 1 of [3] and is given by

$$T_{k-1} = I - L_{k-1} - L_{k-1}^T.$$

This correction matrix was employed in the s -step and pipelined GMRES. When the matrix T_{k-1} is split into $I - L_{k-1}$ and L_{k-1}^T and applied across two iterations of the DCGS2 algorithm, the resulting loss of orthogonality is $\mathcal{O}(\varepsilon)$ in this case. Indeed, it was conjectured in Bielich et al. [4] that two iterations of DCGS2 are needed to achieve $\mathcal{O}(\varepsilon)$ orthogonal vectors, however, our results demonstrate that iterated MGS is sufficient without re-orthogonalization.

The low-synch modified Gram-Schmidt and GMRES algorithms described in Świrydowicz et al. [3] improve parallel strong-scaling by employing one global reduction for each iteration, (see Lockhart et al. [13]). A review of compact WY Gram Schmidt algorithms and their computational costs is given in [4]. The triangular solve and matrix-vector multiply for the Gauss-Seidel iterations require $(k-1)^2$ flops at iteration k and thus lead to a slightly higher operation count compared to the original MGS algorithm which is $2m^2n$ for an $n \times m$ matrix. The two matrix-vector multiplies in Step 4 of Algorithm 6.1 have complexity $4nk$ per iteration k and the norm in Step 3 costs $2n$ flops, for a total of $2m^2n + 2mn$ flops. The number of global reductions is decreased from $k-1$ at iteration k to only one when combined with the lagged normalization of a Krylov vector. These costs can be compared with the DCGS2 algorithm requiring $4m^2n$ flops. Block generalizations of the DCGS2 and CGS2 algorithm are presented in Carson et al. [24, 25]. The authors generalize the Pythagorean trick to block form and derive BCGS-PIO and BCGS-PIP algorithms with the more favorable communication patterns described herein. An analysis of the backward stability of these block Gram-Schmidt algorithms is also presented.

8. Numerical Results. Numerically challenging test problems for GMRES have been proposed and analyzed over the past 35 years. These include both symmetric and non-symmetric matrices. Simoncini and Szyld [26] introduced a symmetric, diagonal matrix with real eigenvalues, causing MGS-

GMRES to stagnate. Highly non-normal matrices from Walker [10] were used to explore the convergence characteristics of HH-GMRES and then the non-normal fs1836 from Paige et al. [1] and west0132 from Paige and Strakoš [2] encounter stagnation. In addition to these, the `impcol_e` matrix from Greenbaum et al. [27], reaches the $\mathcal{O}(\varepsilon)$ relative residual level on the final iteration, unless it stagnates. Matrices with complex eigenvalues forming a disc inside the unit circle such as the Helmert matrix from Liesen and Tichý [28], are also evaluated. Results from a pressure continuity solver with AMG preconditioner and a circuit simulation with the ADD32 matrix from Rozložník, Strakoš and Tůma [29] are also presented.

8.1. Ill-Conditioned Diagonal Matrix. Simoncini and Szyld [26] consider several difficult ill-conditioned problems that can lead to stagnation of the GMRES residual before converging to the level of machine precision $\mathcal{O}(\varepsilon)$. In their example 5.5, they construct $A = \text{diag}([1e-4, 2, 3, \dots, 100])$, a diagonal matrix, and the right-hand side $\mathbf{b} = \text{randn}(100, 1)$ is normalized so that $\|\mathbf{b}\|_2 = 1$. The condition number of this matrix is $\kappa(A) = 1 \times 10^6$ and $\|A\|_2 = 100$. With the MGS-GMRES algorithm, the relative residual stagnates at the level 1×10^{-12} after 75 iterations, when $\|S_k\|_2 = 1$ indicating that the Krylov vectors are not linearly independent. In the case of the PM-GMRES algorithm, the convergence history is plotted in Figure 2, where it can be observed that the relative residual continues to decrease monotonically. Furthermore, the upper bound $\mathcal{O}(\varepsilon) \|A\tilde{\mathbf{v}}_k\|_2/h_{k+1,k}$ is plotted along with $\|I - \tilde{V}_k^T \tilde{V}_k\|_F$. The latter indicates that a significant loss of orthogonality does not occur.

8.2. Ill-Conditioned Symmetric and Non-Symmetric Matrices. Figures 1.1 and 1.2 from Greenbaum et al. [27] describe the results for STEAM1 (using the HH and MGS implementations, respectively). Similarly, Figures 1.3 and 1.4 from Greenbaum et al. [27] correspond to IMPCOLE. They emphasize that the convergence behavior illustrated in these plots is typical of the MGS-GMRES and HH-GMRES algorithms. The condition number of the system matrix is $\kappa(A) = 2.855 \times 10^7$ and $\|A\|_2 = 2.2 \times 10^7$ for STEAM1, whereas $\kappa(A) = 7.102 \times 10^6$ and $\|A\|_2 = 5.0 \times 10^3$ for IMPCOLE.

Greenbaum et al. [27] observe that although orthogonality of the Krylov vectors is not maintained near machine precision, as is the case for the Householder implementation, the relative residuals of the MGS-GMRES algorithm are almost identical to those of the HH-GMRES until the smallest singular value of the matrix V_k begins to depart from one. At that point the MGS-GMRES relative residual norm begins to stagnate close to its final precision level. This observation is demonstrated with the numerical examples for matrices STEAM1 ($N = 240$, symmetric positive definite matrix used in oil recovery simulations) and IMPCOLE ($N = 225$, nonsymmetric matrix from modelling of the hydrocarbon separation problem). In both experiments $\mathbf{x} = (1, \dots, 1)^T$, $\mathbf{b} = A\mathbf{x}$ and $\mathbf{x}_0 = \mathbf{0}$. The convergence histories for the PM-GMRES algorithm applied to these matrices are plotted in Figures 6 and 7. A loss of orthogonality is not observed until the last iteration. Otherwise the computed metric $\|I - \tilde{V}_k^T \tilde{V}_k\|_F$ and the bound associated with the solution of the normal equations (4.2) remain near $\mathcal{O}(\varepsilon)$.

8.3. Highly Non-Normal Matrices. Bidiagonal matrices with a δ off-diagonal were studied by Embree [30]. These are non-normal matrices where $0 < \delta \leq 1$ and are completely defective for all $\delta \neq 0$. A defective matrix is a square matrix that does not have a complete basis of eigenvectors, and is therefore not diagonalizable, and the pseudo-spectra [19] of these matrices are discs in the complex plane. Our PM-GMRES algorithm leads to convergence after 16 iterations without stagnation and orthogonality is maintained to machine precision as plotted in Figure 3. The matrix 2-norm is $\|A\|_2 = 1.1$ and condition $\kappa(A) = 1.2$. Walker [10] employed the highly non-normal matrix in equation (8.1) below to compare the Gram-Schmidt and Householder implementations of GMRES. The element α controls both the condition $\kappa(A)$ and departure from normality $\text{dep}(A)$ of the matrix of size $n \times n$. Here $\|A\|_2 = 2.0 \times 10^3$.

$$(8.1) \quad A = \begin{bmatrix} 1 & 0 & \cdots & 0 & \alpha \\ 0 & 2 & \cdots & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & n \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}$$

For large values of α , Walker found that the MGS-GMRES residual would stagnate and that the CGS algorithm led to instability. Furthermore, it was found that even CGS-2 with re-orthogonalization exhibited some instability near convergence. HH-GMRES maintains $\mathcal{O}(\varepsilon)$ orthogonality as measured by $\|I - \tilde{V}_k^T \tilde{V}_k\|_F$ and reduces the relative residual to machine precision.

In our experiments, the value $\alpha = 2000$ leads to a matrix with $\kappa(A) = 4 \times 10^5$. The departure from normality, based on Henrici’s metric, is large $dep(A) = 2000$. The convergence history for PM-GMRES is displayed in Figure 4. The loss of orthogonality remains near $O(\varepsilon)$ and our upper bound is close for this problem. Notably, the oscillations present in the relative residual computed by the CGS-2 variant of GMRES are not present in our PM-GMRES convergence history plots, where the relative residual decreases monotonically.

Paige and Strakoš [2] analysed convergence for the non-normal matrices, FS1836 and WEST0132. In all their experiments $\mathbf{b} = (1, \dots, 1)^T$. The matrix FS1836 has dimension $n = 183$, with $\|A\|_2 \approx 1.2 \times 10^9$, $\kappa(A) \approx 1.5 \times 10^{11}$. For the matrix WEST0132 with $n = 132$, $\|A\|_2 \approx 3.2 \times 10^5$, $\kappa(A) \approx 6.4 \times 10^{11}$. The MGS-GMRES algorithm is employed to produce the results reported by the authors. Their Figure 7.1 indicates that the relative residual for FS1836 stagnates at 1×10^{-7} at iteration 43 when orthogonality is lost. The relative residual for the WEST0132 matrix also stagnates at the 1×10^{-7} level after 130 iterations. These results contrast with our Figures 5 and 8. In both cases the relative residuals continue to decrease and $\|I - \tilde{V}_k^T \tilde{V}_k\|_F$ grows slowly or remains close to machine precision.

8.4. Complex Eigenvalues in a Disc. For their final experiment, Liesen and Tichy [28] employ the Helmert matrix generated by the Matlab command `gallery('orthog', 18, 4)`. Helmert matrices occur in a number of practical problems, for example in applied statistics. The matrix is orthogonal, and the eigenvalues cluster around -1 , as in the right panels of their Figure 4.4. The worst-case GMRES residual norm decreases quickly throughout the iterations and stagnates at the 12-th iteration, where the relative residual remains at 1×10^{-10} . From the PM-GMRES convergence history plotted in Figure 10, the loss of orthogonality remains near machine precision and the relative residual does not stagnate. The quantity $\mathcal{O}(\varepsilon) \|A\tilde{\mathbf{v}}_k\|_2 / \tilde{h}_{k+1,k}$ is an excellent predictor of the orthogonality with $\|A\|_2 = 1$.

8.5. Nalu-Wind Model. Nalu-Wind solves the incompressible Navier-Stokes equations, with a pressure projection. The governing equations are discretized in time with a BDF-2 integrator, where an outer Picard fixed-point iteration is employed to reduce the nonlinear system residual at each time step. Within each time step, the Nalu-Wind simulation time is often dominated by the time required to setup and solve the linearized governing equations. The pressure systems are solved using PM-GMRES with an AMG preconditioner, where a polynomial Gauss-Seidel smoother is now applied as described in Mulleney et al. [31]. Hence, Gauss-Seidel is a compute time intensive component, when employed as a smoother within an AMG V -cycle.

The McAlister experiment for wind-turbine blades is an unsteady RANS simulation of a fixed-wing, with a NACA0015 cross section, operating in uniform inflow. Resolving the high-Reynolds number boundary layer over the wing surface requires resolutions of $\mathcal{O}(10^{-5})$ normal to the surface resulting in grid cell aspect ratios of $\mathcal{O}(40,000)$. These high aspect ratios present a significant challenge. The simulations were performed for the wing at 12 degree angle of attack, 1 m chord length, denoted c , 3.3 aspect ratio, i.e., $s = 3.3c$, and square wing tip. The inflow velocity is $u_\infty = 46$ m/s, the density is $\rho_\infty = 1.225$ kg/m³, and dynamic viscosity is $\mu = 3.756 \times 10^{-5}$ kg/(m s), leading to a Reynolds number, $Re = 1.5 \times 10^6$. Due to the complexity of mesh generation, only one mesh with approximately 3 million grid points was generated.

The smoother is hybrid block-Jacobi with two sweeps of polynomial Gauss-Seidel relaxation applied locally on a subdomain and then Jacobi smoothing for globally shared degrees of freedom. The coarsening rate for the wing simulation is roughly $4\times$ with eight levels in the V -cycle for *hypr* [32]. Operator complexity C is approximately 1.6 indicating more efficient V -cycles with aggressive coarsening, however, an increased number of GMRES iterations are required compared to standard coarsening. The convergence history is plotted in Figure 9, where the loss of orthogonality is completely flat and close to machine precision.

8.6. Circuit Simulation. Rozložník et al. [29] study a typical linear system arising in circuit simulation (the matrix from a 32-bit adder design). In exact arithmetic the Arnoldi vectors are orthogonal. However, in finite precision computation the orthogonality is lost, which may potentially affect both the convergence rate and the ultimate attainable accuracy of the computed approximate solution. In their Figure 3, the authors have plotted the loss of orthogonality of the computed Krylov vectors for different implementations of the GMRES method (MGS, Householder and CGS). The equivalent results for the PM-GMRES algorithm are plotted in Figure 11, with a small matrix norm $\|A\|_2 = 0.05$ and $\kappa(A) = 213$.

9. Conclusions. The essential contribution of our work was to derive a post-modern (viz. not classical) formulation of the GMRES algorithm that employs an iterated solution of the normal equations appearing in the Gram-Schmidt projector, as described by Ruhe [5], and the low-synchronization algorithms introduced by Świrydowicz et al. [3]. The basic idea developed here was to project the vector $A\mathbf{v}_k$ onto the orthogonal complement of the space spanned by the computed Krylov vectors represented by the (properly normalized) columns of $\tilde{V}_k \in \mathbb{C}^{n \times k}$.

The insights gained from the seminal work of Ruhe [5] led to the conclusion that the iterated modified Gram-Schmidt algorithm is a *multiplicative* Gauss-Seidel iteration for the normal equations $Q_{k-1}^T Q_{k-1} \mathbf{r}_{1:k-1,k} = Q_{k-1}^T \mathbf{a}_k$. The classical Gram-Schmidt algorithm can be viewed as an *additive* Jacobi relaxation. The projector is then given by $P\mathbf{a}_k = \mathbf{a}_k - Q_{k-1} T_{k-1} Q_{k-1}^T \mathbf{a}_k$, where T_{k-1} is a correction matrix. In the case of DCGS-2, with delayed re-orthogonalization, Bielich et al. [4] split and apply the correction matrix across two Gram-Schmidt iterations and then apply Stephen's trick to maintain orthogonality. For MGS, the triangular matrix $T_{k-1} \approx (Q_{k-1}^T Q_{k-1})^{-1}$ appearing in Świrydowicz et al. [3] was identified as the inverse compact *WY* form of MGS with $T_{k-1} = (I + L_{k-1})^{-1}$, where the strictly lower triangular matrix L_{k-1} was computed from the loss of orthogonality relation

$$Q_{k-1}^T Q_{k-1} = I + L_{k-1} + L_{k-1}^T.$$

Q_{k-1} is the correctly (properly) normalized matrix as described in Björck and Paige [12]. The matrix T_{k-1} was also present, without having been explicitly defined, in the original rounding error analysis of Björck [14], in Lemma 5.1. In effect, the low-synch MGS algorithm presented in [3] represents one Gauss-Seidel iteration to construct the projector. When applied twice, the resulting correction matrix is close to a symmetric matrix

$$T_{k-1}^{(2)} = M_{k-1}^{-1} [I + N_{k-1} M_{k-1}^{-1}] = T_{k-1}^{(1)} - T_{k-1}^{(1)} L_{k-1}^T T_{k-1}^{(1)}.$$

When employed to compute the Arnoldi-*QR* expansion, two iterations of Gauss-Seidel relaxation results in an $\mathcal{O}(\varepsilon)$ residual error for the normal equations in the projection step and thereby reducing the loss of orthogonality of the Krylov vectors to $\mathcal{O}(\varepsilon) \|A\mathbf{v}_k\|_2 / h_{k+1,k}$ where $B = [\mathbf{r}_0, A\mathbf{v}_k]$. This is related to recent work on the iterative solution of triangular linear systems using Jacobi iterations that may diverge for highly non-normal matrices. Here, the departure from normality is an indicator of the loss of orthogonality. Both of these indicate a loss of numerical rank for the Krylov vectors with the smallest singular value decreasing from one. Our numerical experiments, on challenging problems proposed over the past thirty-five years, demonstrate that the relative residual does not stagnate. Furthermore, the loss of orthogonality for MGS-GMRES is at most $\mathcal{O}(\varepsilon)\kappa(B)$ and yet for PM-GMRES remains near machine precision.

We anticipate that PM-GMRES may facilitate the construction of backward stable iterative solvers based on block Gram-Schmidt algorithms including enlarged Krylov subspace methods and *s*-step methods. Mixed-precision formulations of GMRES may also benefit. Randomization and sketching may be applied to the normal equations in the projector, with an oblique inner-product $x^T B^T B x$ and sketching matrix B , leading to even greater computational efficiency. In addition, these developments could be relevant for Anderson acceleration of nonlinear fixed-point iterations, which is currently being applied to optimization algorithms in deep learning for artificial intelligence. The algorithms may also be useful for computing eigenvalues with the Krylov-Schur algorithm of Stewart [21], where $\mathcal{O}(\varepsilon)$ orthogonality is needed to obtain an invariant subspace of maximal dimension.

Acknowledgement. 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. The second author was additionally supported by Charles University PRIMUS project no. PRIMUS/19/SCI/11 and Charles University Research program no. UNCE/SCI/023. The third author was supported by the Academy of Sciences of the Czech Republic (RVO 67985840) and by the Grant Agency of the Czech Republic, Grant No. 20-01074S. Our work was inspired by the contributions of A. Björck, A. Ruhe, C. C. Paige, Z. Strakoš, and the legacy of Y. Saad and M. H. Schultz that is GMRES. We also wish to thank our friends and colleagues Julien Langou and Luc Giraud for their thoughtful insights over many years. A sincere thank-you to the organizers of the Copper Mountain conferences on Krylov and multigrid methods for their rich 38+ year tradition of innovative numerical linear algebra.

Appendix. The inverse of a block triangular matrix leads naturally to the correction matrices for the compact and inverse compact WY modified Gram-Schmidt projector. Björck wrote the Gram-Schmidt projector as $P_1 = I - Q L_1 Q^T$, in equation (17) of [6, page 313]. The inverse compact WY form of the MGS projector also appears in Björck and Paige [12] and Paige et al. [1, page 271], equation (3.15), as the (2, 2) block of the augmented Householder transformation \tilde{P}_k . By forming the matrix product $\tilde{P}_k \tilde{P}_k^T = I_k$, the authors obtain the strictly upper triangular part of $\tilde{Q}_k^T \tilde{Q}_k$ given by

$$\tilde{U}_k = (I - \tilde{S}_k)^{-1} \tilde{S}_k = \tilde{S}_k (I - \tilde{S}_k)^{-1}$$

For the augmented Householder matrix, \tilde{P}_k^T is applied, and within the (2,2) block the MGS projector is identified as, $\tilde{P}_k = I - \tilde{Q}_k (I - \tilde{S}_k)^T \tilde{Q}_k^T$

$$\begin{aligned} I - \tilde{S}_k^T &= I - \tilde{U}_k^T (I + \tilde{U}_k^T)^{-1} \\ &= I - \tilde{L}_k (I + \tilde{L}_k)^{-1} \\ &= I - \tilde{L}_k + \tilde{L}_k^2 - \tilde{L}_k^3 + \dots \\ &= (I + \tilde{L}_k)^{-1} \end{aligned}$$

Therefore, the Householder transformation for an augmented matrix naturally produces an inverse compact WY MGS projector, given by:

$$\tilde{P}_k = I - \tilde{Q}_k (I + \tilde{L}_k)^{-1} \tilde{Q}_k^T$$

with the lower triangular matrix $T_k = (I + \tilde{L}_k)^{-1}$. In the case of the block lower triangular matrix inverse, we have

$$(9.1) \quad \begin{bmatrix} A & 0 \\ X & B \end{bmatrix} \begin{bmatrix} A^{-1} & 0 \\ Y & B^{-1} \end{bmatrix} = \begin{bmatrix} I & 0 \\ XA^{-1} + BY & I \end{bmatrix}$$

where $Y = -B^{-1} X A^{-1}$. Therefore, for the inverse compact WY MGS, given the correction matrix $T_{k-1}^{(1)}$ generated by the recursion

$$(9.2) \quad T_{k-1}^{-1} = (I + L_{k-1}) = \begin{bmatrix} T_{k-2}^{-1} & 0 \\ \mathbf{q}_{k-1}^T Q_{k-2} & 1 \end{bmatrix}$$

it is possible to prove that

$$(9.3) \quad \begin{bmatrix} T_{k-2}^{-1} & 0 \\ \mathbf{q}_{k-1}^T Q_{k-2} & 1 \end{bmatrix} \begin{bmatrix} T_{k-2} & 0 \\ -\mathbf{q}_{k-1}^T Q_{k-2} T_{k-2} & 1 \end{bmatrix} = \begin{bmatrix} I_{k-2} & 0 \\ 0 & 1 \end{bmatrix}$$

Therefore, the lower triangular correction matrix $T_{k-1} = (I + L_{k-1})^{-1}$ from Świrydowicz et al. [3] is equivalent to the compact WY matrix L_1 from Björck [6, page 313].

To prove that the correction matrix $T_{k-1}^{(2)}$ is close to a symmetric matrix, consider the matrix

$$T_{k-1}^{(2)} = (I + L_{k-1})^{-1} - (I + L_{k-1})^{-1} L_{k-1}^T (I + L_{k-1})^{-1} = M_{k-1}^{-1} [I + N_{k-1} M_{k-1}^{-1}]$$

and then $T_{k-1}^{(2)} = T_{k-1}^{(1)} - T_{k-1}^{(1)} L_{k-1}^T T_{k-1}^{(1)}$. From the block inverse (9.3),

$$T_{k-1}^{(1)} = \begin{bmatrix} T_{k-2}^{(1)} & 0 \\ -\mathbf{q}_{k-1}^T Q_{k-2} T_{k-2}^{(1)} & 1 \end{bmatrix}, \quad L_{k-1}^T = \begin{bmatrix} L_{k-2}^T & Q_{k-2}^T \mathbf{q}_{k-1} \\ 0 & 0 \end{bmatrix}$$

By subtracting these matrices, and dropping $\mathcal{O}(\varepsilon)$ terms, a symmetric matrix is obtained

$$(9.4) \quad T_{k-1}^{(2)} = \begin{bmatrix} T_{k-2}^{(1)} & -T_{k-2}^{(1)} Q_{k-2}^T \mathbf{q}_{k-1} \\ -\mathbf{q}_{k-1}^T Q_{k-2} T_{k-2}^{(1)} & 1 \end{bmatrix}.$$

It's singular values and eigenvalues remain close to one. However, as k increases $\sigma_{\max}(T_{k-1}^{(1)}) > 1$ and $\sigma_{\min}(T_{k-1}^{(1)}) < 1$, and thus $T_{k-1}^{(1)}$ departs from normality according to (5.2) when orthogonality is lost.




Fig. 2: Simoncini matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 3: Embree δ matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 4: Walker matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 5: fs1863 matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 6: steam1 matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 7: impcol.e matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 8: west0132 matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 9: Pressure matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 10: Helmert matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).




Fig. 11: Add32 matrix. Relative residual. Normal equations residual τ from (4.2) multiplied by the ratio in (4.4), and follows Ruhe's metric (4.1), versus LHS of the loss of orthogonality relation (4.4).

REFERENCES

- [1] C. C. Paige, M. Rozložník, Z. Strakoš, Modified Gram-Schmidt (MGS), least squares, and backward stability of MGS-GMRES, *SIAM Journal on Matrix Analysis and Applications* 28 (1) (2006) 264–284.
- [2] C. C. Paige, Z. Strakoš, Residual and backward error bounds in minimum residual Krylov subspace methods, *SIAM Journal on Scientific and Statistical Computing* 23 (6) (2002) 1899–1924.
- [3] K. Świrydowicz, J. Langou, S. Ananthan, U. Yang, S. Thomas, Low synchronization Gram-Schmidt and generalized minimal residual algorithms, *Numerical Linear Algebra with Applications* 28 (2020) 1–20.
- [4] D. Bielich, J. Langou, S. Thomas, K. Świrydowicz, I. Yamazaki, E. Boman, Low-synch Gram-Schmidt with delayed reorthogonalization for Krylov solvers, *Parallel Computing* (2021).
- [5] A. Ruhe, Numerical aspects of Gram-Schmidt orthogonalization of vectors, *Linear Algebra and its Applications* 52 (1983) 591–601.
- [6] Å. Björck, Numerics of Gram-Schmidt orthogonalization, *Linear Algebra and Its Applications* 197 (1994) 297–316.
- [7] N. J. Higham, P. A. Knight, Componentwise error analysis for stationary iterative methods, in: C. D. Meyer, R. J. Plemmons (Eds.), *Linear Algebra, Markov Chains, and Queueing Models*, IMA Volumes in Mathematics and its Applications, Springer-Verlag, 1993, pp. 29–46.
- [8] S. Thomas, A. Carr, P. Mullowney, R. Li, K. Świrydowicz, Neuman series in GMRES and algebraic multigrid, *SIAM Journal of Matrix Analysis and Applications* (2022).
- [9] Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal on scientific and statistical computing* 7 (3) (1986) 856–869.
- [10] H. F. Walker, Implementation of the GMRES method using Householder transformations, *SIAM Journal on Scientific and Statistical Computing* 9 (1) (1988) 152–163.
- [11] L. Giraud, S. Gratton, J. Langou, A rank- k update procedure for reorthogonalizing the orthogonal factor from modified Gram-Schmidt, *SIAM J. Matrix Analysis and Applications* 25 (4) (2004) 1163–1177.
- [12] Å. Björck, C. C. Paige, Loss and recapture of orthogonality in the modified Gram-Schmidt algorithm, *SIAM Journal on Matrix Analysis and Applications* 13 (1992) 176–190.
- [13] S. Lockhart, D. J. Gardner, C. S. Woodward, S. Thomas, L. N. Olson, Performance of low synchronization orthogonalization methods in Anderson accelerated fixed point solvers, in: *Proceedings of the 2022 SIAM Conference on Parallel Processing for Scientific Computing (PP)*, pp. 49–59.
- [14] A. Björck, Solving least squares problems by Gram-Schmidt orthogonalization, *BIT* 7 (1967) 1–21.
- [15] L. Giraud, J. Langou, M. Rozložník, J. v. d. Eshof, Rounding error analysis of the classical Gram-Schmidt orthogonalization process, *Numerische Mathematik* 101 (2005) 87–100.
- [16] N. J. Higham, The accuracy of solutions to triangular systems, *SIAM Journal on Numerical Analysis* 26 (5) (1989) 1252–1265.
- [17] P. Henrici, Bounds for iterates, inverses, spectral variation and fields of values of non-normal matrices, *Numerische Mathematik* 4 (1) (1962) 24–40.
- [18] I. C. Ipsen, A note on the field of values of non-normal matrices, Tech. rep., North Carolina State University. Center for Research in Scientific Computation (1998).
- [19] L. Trefethen, M. Embree, The behavior of nonnormal matrices and operators, *Spectra and Pseudospectra* (2005).
- [20] A. Ruhe, On the closeness of eigenvalues and singular values for almost normal matrices, *Linear Algebra and its Applications* 11 (1) (1975) 87–93.
- [21] G. W. Stewart, A Krylov-Schur algorithm for large eigenproblems, *SIAM Journal on Matrix Analysis and Applications* 23 (3) (2001) 601–614.
- [22] Q. Zou, GMRES algorithms over 35 years (2021). [arXiv:2110.04017](https://arxiv.org/abs/2110.04017).
- [23] I. Yamazaki, S. Thomas, M. Hoemmen, E. G. Boman, K. Świrydowicz, J. J. Elliott, Low-synchronization orthogonalization schemes for s -step and pipelined Krylov solvers in Trilinos, in: *Proceedings of the 2020 SIAM Conference on Parallel Processing for Scientific Computing*, SIAM, 2020, pp. 118–128. doi:10.1137/1.9781611976137.11.
- [24] E. Carson, K. Lund, M. Rozložník, S. Thomas, Block Gram-Schmidt algorithms and their stability properties, *Linear Algebra and its Applications* 638 (2022) 150–195.
- [25] E. Carson, K. Lund, M. Rozložník, The stability of block variants of classical Gram-Schmidt, *SIAM Journal on Matrix Analysis and Applications* 42 (3) (2021) 1365–1380.
- [26] V. Simoncini, D. B. Szyld, Theory of inexact krylov subspace methods and applications to scientific computing, *SIAM Journal on Scientific Computing* 25 (2) (2003) 454–477.
- [27] A. Greenbaum, M. Rozložník, Z. Strakoš, Numerical behaviour of the modified Gram-Schmidt GMRES implementation, *BIT* 37 (3) (1997) 706–719.
- [28] J. Liesen, P. Tichý, The worst-case GMRES for normal matrices, *BIT Numer Math* 44 (2004) 79–98.
- [29] M. Rozložník, Z. Strakoš, M. Tůma, On the role of orthogonality in the GMRES method, in: K. G. Jeffery, J. Král, M. Bartosek (Eds.), *SOFSEM '96: Theory and Practice of Informatics*, Proceedings, Vol. 1175 of Lecture Notes in Computer Science, Springer, 1996, pp. 409–416.
- [30] M. Embree, How descriptive are GMRES convergence bounds?, Tech. Rep. Tech. Rep. 99/08, Mathematical Institute, University of Oxford, UK (1999).
- [31] P. Mullowney, R. Li, S. Thomas, S. Ananthan, A. Sharma, A. Williams, J. Rood, M. A. Sprague, Preparing an incompressible-flow fluid dynamics code for exascale-class wind energy simulations, in: *Proceedings of the ACM/IEEE Supercomputing 2021 Conference*, ACM, 2021, pp. 1–11.
- [32] R. D. Falgout, J. E. Jones, U. M. Yang, The design and implementation of hypre, a library of parallel high performance preconditioners, in: A. M. Bruaset, A. Tveito (Eds.), *Numerical Solution of Partial Differential Equations on Parallel Computers*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 267–294.