

SHORT-TERM RECURRENCE KRYLOV SUBSPACE METHODS FOR NEARLY HERMITIAN MATRICES*

MARK EMBREE[†], JOSEF A. SIFUENTES[‡], KIRK M. SOODHALTER[§],
DANIEL B. SZYLD[§], AND FEI XUE[§]

Abstract. The progressive GMRES algorithm, introduced by Beckermann and Reichel in 2008, is a residual-minimizing short-recurrence Krylov subspace method for solving a linear system in which the coefficient matrix has a low-rank skew-Hermitian part. We analyze this algorithm, observing a critical instability that makes the method unsuitable for some problems. To work around this issue we introduce a different short-term recurrence method based on Krylov subspaces for such matrices, which can be used as either a solver or a preconditioner. Numerical experiments compare this method to alternative algorithms.

Key words. GMRES, MINRES, nearly Hermitian matrices, low-rank modifications

AMS subject classifications. 65F10, 65N12, 15B57, 45B05, 45A05

DOI. 10.1137/110851006

1. Introduction. A variety of applications warrant the solution of linear systems of equations where the coefficient matrix \mathbf{A} has a skew-Hermitian part $\frac{1}{2}(\mathbf{A} - \mathbf{A}^*)$ with low rank. Such systems arise, for example, from discretized integral equations derived from wave scattering applications and electrostatics [28], as well as path following methods [1]. In this paper, we consider efficient Krylov subspace methods for the solution of

$$(1.1) \quad \mathbf{Ax} = \mathbf{b},$$

where the nonsingular coefficient matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$ has the structure

$$(1.2) \quad \mathbf{A} = \mathbf{A}^* + \mathbf{FG}^* = \mathbf{H} + \frac{1}{2}\mathbf{FG}^*,$$

for full rank $\mathbf{F}, \mathbf{G} \in \mathbb{C}^{n \times s}$ with $s \ll n$; $\mathbf{H} := \frac{1}{2}(\mathbf{A} + \mathbf{A}^*)$ is the Hermitian part of \mathbf{A} .

Beckermann and Reichel [1] proposed a “progressive GMRES” algorithm based on a short recurrence for problems with the structure (1.2), which is mathematically equivalent to the much less efficient full GMRES method [23]. We have found that progressive GMRES, while theoretically elegant, can suffer from fundamental numerical instabilities that render it unsuitable for some otherwise benign matrices of the form (1.2). In this paper we carefully document these instabilities and propose an alternative algorithm that uses a Schur complement approach to solve systems with this

*Received by the editors October 11, 2011; accepted for publication April 17, 2012; published electronically June 5, 2012.

<http://www.siam.org/journals/simax/33-2/85100.html>

[†]Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005–1892 (embree@rice.edu). This author’s research was supported by the U. S. National Science Foundation under grant DMS-CAREER-0449973.

[‡]Courant Institute of Mathematical Sciences, New York University, New York, NY 10012-1185 (sifuentes@cims.nyu.edu). This author’s research was supported by the U. S. National Science Foundation under a Graduate Research Fellowship and grants DMS-CAREER-0449973, DMS06-02235 and Air Force NSSEFF Program award FA9550-10-1-018.

[§]Department of Mathematics, Temple University, Philadelphia, PA 19122-6094 (ksoodha@temple.edu, szyld@temple.edu, fxue@temple.edu). The research of these authors was supported in part by the U.S. Department of Energy under grant DE-FG02-05ER25672 and the U. S. National Science Foundation under grant DMS-1115520.

structure using short-term recurrence relations. While our method is not equivalent to full GMRES for the original system, it is based on an optimal method (MINRES) for the Hermitian part of \mathbf{A} , and can be applied as either a solver or a preconditioner.

We begin by introducing progressive GMRES (PGMRES) in the next section. In section 3 we present numerical experiments that demonstrate the instability of PGMRES for some well-conditioned linear systems, then comment on possible causes for this instability. We propose in section 4 an alternative algorithm for (1.1)–(1.2) based on existing Schur complement methods. In section 4.1, we use this method to precondition coefficient matrices that are small-norm perturbations of a matrix of the form (1.2). We present numerical experiments in section 5 that compare our new method to existing ones.

2. Preliminaries. The GMRES algorithm approximates the solution to (1.1) through an iterative process that, beginning with an initial iterate $\mathbf{x}_0 = \mathbf{0}$, constructs at iteration k the vector \mathbf{x}_k that solves the least squares problem

$$(2.1) \quad \min_{\tilde{\mathbf{x}} \in \mathcal{K}_k(\mathbf{A}, \mathbf{b})} \|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}\|$$

over the Krylov subspace $\mathcal{K}_k(\mathbf{A}, \mathbf{b}) = \text{span}\{\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}, \dots, \mathbf{A}^{k-1}\mathbf{b}\}$ [23]. To solve this problem, an orthonormal basis for the Krylov subspace is developed through the Arnoldi process. The first k steps of this procedure can be collected into the relationship

$$(2.2) \quad \mathbf{A}\mathbf{V}_k = \mathbf{V}_{k+1}\tilde{\mathbf{H}}_k$$

$$(2.3) \quad = \mathbf{V}_k\mathbf{H}_k + h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^T,$$

where $\mathbf{V}_k \in \mathbb{C}^{n \times k}$ and $\mathbf{V}_{k+1} := [\mathbf{V}_k \ \mathbf{v}_{k+1}] \in \mathbb{C}^{n \times (k+1)}$ have orthonormal columns and $\tilde{\mathbf{H}}_k \in \mathbb{C}^{(k+1) \times k}$ is upper Hessenberg with upper $k \times k$ block \mathbf{H}_k and $(k+1, k)$ entry $h_{k+1,k}$. If the Arnoldi process starts with $\mathbf{v}_1 = \mathbf{b}/\|\mathbf{b}\|$, then $\mathcal{K}_k(\mathbf{A}, \mathbf{b}) = \text{Ran}(\mathbf{V}_k)$, the range of \mathbf{V}_k . It is straightforward to then show that (2.1) is equivalent to the least squares problem

$$(2.4) \quad \min_{\mathbf{c} \in \mathbb{C}^k} \|\|\mathbf{b}\|\mathbf{e}_1 - \tilde{\mathbf{H}}_k\mathbf{c}\|,$$

and $\mathbf{x}_k = \mathbf{V}_k\mathbf{c}_k$ solves (2.1) if \mathbf{c}_k solves (2.4). Postmultiplying both sides of (2.3) by \mathbf{e}_k , we can write the Arnoldi vector update as

$$(2.5) \quad h_{k+1,k}\mathbf{v}_{k+1} = \mathbf{A}\mathbf{v}_k - \mathbf{V}_k\mathbf{H}_k\mathbf{e}_k,$$

where $h_{k+1,k} = \|\mathbf{A}\mathbf{v}_k - \mathbf{V}_k\mathbf{H}_k\mathbf{e}_k\| \geq 0$. Computing the updated Arnoldi vector \mathbf{v}_{k+1} through (2.5) is often carried out via modified Gram–Schmidt orthogonalization (see [22, Chap. 6] for details and alternatives).

Suppose the coefficient matrix is Hermitian, $\mathbf{A} = \mathbf{A}^*$. Then the Hessenberg matrix is also Hermitian, $\mathbf{H}_k = \mathbf{V}_k^*\mathbf{A}\mathbf{V}_k = \mathbf{V}_k^*\mathbf{A}^*\mathbf{V}_k = \mathbf{H}_k^*$, and thus tridiagonal. In such a case, the last column of \mathbf{H}_k has only two nonzero entries, implying that $\mathbf{V}_k\mathbf{H}_k\mathbf{e}_k$ is a linear combination of \mathbf{v}_k and \mathbf{v}_{k-1} . Explicitly using this fact in the Arnoldi relation (2.5) provides a three-term recurrence, known as the Lanczos iteration. When (2.4) is solved with the tridiagonal $\tilde{\mathbf{H}}_k$ constructed through the Lanczos process, GMRES reduces to the MINRES algorithm [20], and \mathbf{x}_k can be computed “progressively”: \mathbf{x}_k can be updated to \mathbf{x}_{k+1} without explicitly solving the least squares problem (2.4); see, e.g., [22, Chap. 6].

Is a similar simplification possible if \mathbf{A} is not Hermitian, but only “nearly Hermitian,” in the sense of (1.2)? Beckermann and Reichel [1] demonstrate that for matrices with such structure, the Arnoldi process can be updated through a short recurrence involving a rank- s projection. We provide here a simple derivation of their algorithm. The key idea is similar to the simplification that led to the Lanczos method: if \mathbf{A} is Hermitian, then so is \mathbf{H}_k ; if the skew-Hermitian part of \mathbf{A} is low-rank, then so is the skew-Hermitian part of \mathbf{H}_k .¹ To see this, note that for \mathbf{A} of the form (1.2),

$$\mathbf{H}_k = \mathbf{V}_k^*(\mathbf{A}^* + \mathbf{F}\mathbf{G}^*)\mathbf{V}_k = \mathbf{H}_k^* + (\mathbf{V}_k^*\mathbf{F})(\mathbf{V}_k^*\mathbf{G})^*.$$

Inserting this representation of \mathbf{H}_k into the Arnoldi update (2.5) yields

$$(2.6) \quad h_{k+1,k}\mathbf{v}_{k+1} = \mathbf{A}\mathbf{v}_k - \bar{h}_{k,k}\mathbf{v}_k - h_{k,k-1}\mathbf{v}_{k-1} - \mathbf{V}_k\mathbf{V}_k^*\mathbf{F}\mathbf{G}^*\mathbf{v}_k.$$

(The subdiagonal entry $h_{k,k-1}$, derived from the normalization of \mathbf{v}_k , is a nonnegative real number.) The orthogonal projection of \mathbf{F} onto the Krylov subspace, $\mathbf{V}_k\mathbf{V}_k^*\mathbf{F}$, can be computed progressively; that is, defining $\tilde{\mathbf{F}}_k := \mathbf{V}_k\mathbf{V}_k^*\mathbf{F}$, we see that

$$\tilde{\mathbf{F}}_k = \tilde{\mathbf{F}}_{k-1} + \mathbf{v}_k\mathbf{v}_k^*\mathbf{F},$$

so (2.6) is a three-term recurrence. From (2.6), $\mathbf{A}\mathbf{v}_k - \mathbf{V}_k\mathbf{V}_k^*\mathbf{F}\mathbf{G}^*\mathbf{v}_k \in \mathcal{K}_{k+1}(\mathbf{A}, \mathbf{b})$ is orthogonal to $\mathcal{K}_{k-2}(\mathbf{A}, \mathbf{b})$, i.e., the Arnoldi process can implicitly orthogonalize $\mathbf{A}\mathbf{v}_k$ against $\mathcal{K}_{k-2}(\mathbf{A}, \mathbf{b})$ by subtracting the (progressively computed) term $\mathbf{V}_k\mathbf{V}_k^*\mathbf{F}\mathbf{G}^*\mathbf{v}_k$; explicit orthogonalization against \mathbf{v}_k and \mathbf{v}_{k-1} is then required to form \mathbf{v}_{k+1} . Given the resulting short-term Arnoldi iteration, the iterate \mathbf{x}_k can be computed progressively without needing to access the complete set of Arnoldi vectors \mathbf{V}_k , and using only $s + 3$ vectors of storage. See [1] for a full description of this process. (An alternative short-term recurrence strategy for solving systems of this form was proposed earlier by Huhtanen [15, Thm. 2.10].)

Beckermann and Reichel’s progressive GMRES (PGMRES) algorithm has two advantages over full GMRES: it only requires storage of three Arnoldi vectors, rather than the entire set, and it avoids the growing computational complexity of orthogonalizing the new Arnoldi vector against the previous Krylov subspace. Unfortunately, we have found that PGMRES suffers a critical drawback: it can introduce numerical instabilities that cause the residual to stagnate well before convergence—even for some examples where GMRES converges swiftly—as illustrated in the next section.

3. Illustrations and analysis of PGMRES instability. We begin by applying PGMRES to a problem from acoustic scattering, then move to a contrived class of well-conditioned normal matrices: in both cases our numerical experiments demonstrate instabilities that cause the residual norm to stagnate well before reaching convergence. For such examples, we compare the residuals produced by PGMRES to those generated by a standard implementation of GMRES based on the modified Gram–Schmidt Arnoldi process [23]. Since the iterate produced by PGMRES is computed progressively, for the sake of comparison we explicitly compute the residual norm $\|\mathbf{b} - \mathbf{A}\mathbf{x}_k\|$ at each step. To gain insight into the onset of instability, we also track the departure from orthonormality of the Arnoldi basis, and the rate at which these basis vectors drift toward linear dependence. We gauge the departure from orthonormality using a measure advocated by Paige et al. [18, 19]. Assume the

¹A similar observation can be exploited to efficiently reduce such \mathbf{A} to Hessenberg form for eigenvalue computations; see, e.g., [35].

columns of \mathbf{V}_k have norm 1, and let \mathbf{U}_k denote the strictly upper triangular part of $\mathbf{V}_k^* \mathbf{V}_k - \mathbf{I} = \mathbf{U}_k + \mathbf{U}_k^*$. Now define $\mathbf{S}_k := (\mathbf{I} + \mathbf{U}_k)^{-1} \mathbf{U}_k$. The departure from orthonormality can be measured by $\|\mathbf{S}_k\| \in [0, 1]$, which is zero when the columns of \mathbf{V}_k are orthonormal and one when those columns are linearly dependent [18, Thm. 2.1]. Moreover, [18, Cor. 5.2], [19, Lem. 5.1],

$$\text{cond}(\mathbf{V}_k) \leq \frac{1 + \|\mathbf{S}_k\|}{1 - \|\mathbf{S}_k\|},$$

where the condition number $\text{cond}(\cdot)$ is the ratio of the largest to smallest singular values. Each numerical experiment in this section shows $\|\mathbf{S}_k\|$ (though the columns of the computed \mathbf{V}_k are only normalized to machine precision).

3.1. The Lippmann–Schwinger equation. The Lippmann–Schwinger integral equation models acoustic scattering in one dimension at wave number κ :

$$(3.1) \quad (I + K)u(x) = u^i(x),$$

where K is the integral operator $K : L^2(0, 2\pi) \rightarrow L^2(0, 2\pi)$ given by

$$(3.2) \quad (Ku)(x) = \frac{i\kappa}{2} \int_0^{2\pi} e^{i\kappa|x-y|} m(y) u(y) dy,$$

$u^i(x)$ is an incident wave satisfying the one-dimensional Helmholtz equation $d^2u^i/dx^2 + \kappa^2u^i = 0$, and m is a function of the refractive index [3]. In the case of constant m , the skew-adjoint part of the operator (3.2) has a two-dimensional range [28]. Furthermore, a Nyström discretization of (3.2) (based on a quadrature rule on a uniform grid with weights equal to the interval length) produces a matrix whose skew-Hermitian part has rank two. (A higher order rule would not necessarily produce a discretization with a skew-Hermitian part of rank two; however, a simple preconditioning on the right by the diagonal matrix of quadrature weights gives a coefficient matrix whose skew-Hermitian part is rank two. Similarly, one could obtain the desired structure by using an inner product that incorporates the quadrature weights.)

Figure 3.1 shows the convergence of the GMRES and PGMRES residuals for the Lippmann–Schwinger problem (with a random \mathbf{b}) for four choices of the wave number κ , all with constant $m = -1$ and discretization dimension $n = 1000$. As the wave number increases, the PGMRES residual curve departs from the standard GMRES curve, stagnating at increasingly large residuals. One might argue that the instability is not significant at $\kappa = 3$; when $\kappa = 8$ the stagnation entirely compromises the utility of the algorithm.

To investigate the stagnation, we retain all Arnoldi vectors in PGMRES (in spite of only using three at a time to compute the new iterate), and compute the departure of the computed Arnoldi basis from orthogonality, $\|(\mathbf{I} + \mathbf{U}_k)^{-1} \mathbf{U}_k\|$, at each iteration. The right plots in Figure 3.1 compare this loss of orthogonality for PGMRES to that observed for standard GMRES (with a full-length Arnoldi recurrence based on the modified Gram–Schmidt process). Both PGMRES and GMRES produce bases that lose orthogonality, but GMRES does so at a slower pace that does not significantly destabilize the convergence, consistent with the analysis of Greenbaum, Rozložník, and Strakoš [10]. Observe that the PGMRES residual departs from that produced by GMRES when the departure from orthogonality approaches one and the Arnoldi vectors begin to approach linear dependence. (In section 4 we present an alternative




FIG. 3.1. Comparison of PGMRES (solid line) to standard (modified Gram-Schmidt) GMRES (dashed line) applied to the one-dimensional scattering problem for various values of the wave number κ . The left plot shows the relative residual norm for each method; the right plot shows the departure from orthonormality, along with the minimum singular value of the Arnoldi basis matrix V_k computed by the PGMRES recurrence (gray line). In exact arithmetic, the algorithms are identical; in finite precision, the performance of PGMRES degrades as κ increases.

method that applies MINRES to the Hermitian part of \mathbf{A} . While it is well known that the Lanczos recurrence for computing an orthonormal basis for the Krylov subspace can cause loss of orthogonality (see, e.g. [21]), this basis does not usually lose *linear independence*, and the MINRES algorithm derived from this basis does not typically exhibit the early stagnation seen in some of the PGMRES examples here.)

3.2. A class of simple examples. The instability of PGMRES is not confined to the application we have just illustrated. Our experiments suggest that striking examples exhibit the following features:

- (i) $\|\mathbf{F}\mathbf{G}^*\| \gg 0$, to stimulate the instability;
- (ii) GMRES should initially converge slowly, during which period the Arnoldi basis generated by PGMRES degrades;
- (iii) GMRES should then enter a phase of rapid convergence, which PGMRES cannot mimic due to its deficient basis.

We shall describe a class of examples that satisfies these three properties, while being normal, nearly diagonal, and even well conditioned (for appropriate parameter values). Consider block-diagonal matrices of the form

$$(3.3) \quad \mathbf{A} = \begin{bmatrix} \mathbf{\Lambda}_- & & \\ & \mathbf{\Lambda}_+ & \\ & & \mathbf{Z} \end{bmatrix},$$

where, for $p \ll n$ and positive constants $\alpha < \beta$ and γ ,

$$\mathbf{\Lambda}_- = \text{diag}(\lambda_1, \dots, \lambda_p), \quad \mathbf{\Lambda}_+ = \text{diag}(\lambda_{p+1}, \dots, \lambda_{n-2}), \quad \mathbf{Z} = \begin{bmatrix} 0 & \gamma \\ -\gamma & 0 \end{bmatrix},$$

with eigenvalues

- $\lambda_1, \dots, \lambda_p$ uniformly spaced in the negative real interval $[-\beta, -\alpha]$;
- $\lambda_{p+1}, \dots, \lambda_{n-2}$ uniformly spaced in the positive real interval $[\alpha, \beta]$;
- $\lambda_{n-1}, \lambda_n = \pm\gamma i$, from the skew-Hermitian matrix \mathbf{Z} .

Hence \mathbf{A} is a normal matrix with $\|\mathbf{H}\| = \beta$, $\|\frac{1}{2}\mathbf{F}\mathbf{G}^*\| = \gamma$, and condition number

$$\text{cond}(\mathbf{A}) := \|\mathbf{A}\|\|\mathbf{A}^{-1}\| = \frac{\max\{\beta, \gamma\}}{\min\{\alpha, \gamma\}}.$$

Figure 3.2 illustrates the spectrum of a representative \mathbf{A} . The qualitative description of GMRES convergence provided by Driscoll, Toh, and Trefethen [6] informs this construction. The purely imaginary eigenvalues $\pm\gamma i$ control $\|\mathbf{F}\mathbf{G}^*\|$: as γ gets large, this pair has little effect on GMRES (the pair delays convergence by roughly two iterations), yet, as indicated in Figure 3.3, the magnitudes of these entries induce the onset of instability. The p eigenvalues $\lambda_1, \dots, \lambda_p$ on the negative real axis associated with the block $\mathbf{\Lambda}_-$ add indefiniteness to the problem, and further delay convergence: for early iterations, GMRES will behave in a fashion similar to MINRES applied to a matrix whose spectrum falls in $[-\beta, -\alpha] \cup [\alpha, \beta]$ (see, e.g., [9, sect. 3.1]), whereby each iteration (asymptotically) reduces the residual norm by the factor

$$\rho_1 = \sqrt{\frac{\beta - \alpha}{\beta + \alpha}}.$$

After sufficiently many iterations to annihilate the outlying eigenvalues (roughly $2p+2$ steps), GMRES then converges at the much more rapid rate

$$\rho_2 = \frac{\sqrt{\beta} - \sqrt{\alpha}}{\sqrt{\beta} + \sqrt{\alpha}}$$

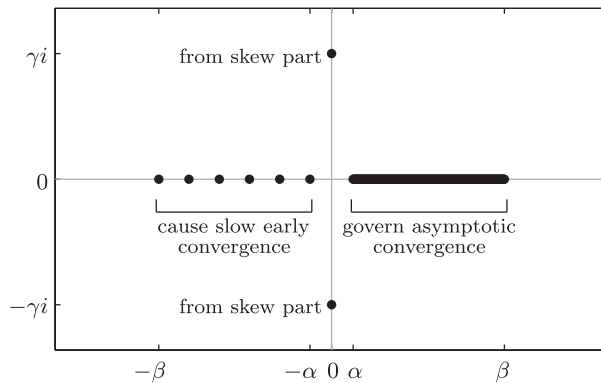


FIG. 3.2. Eigenvalues (●) of the matrix (3.3) in the complex plane, for $n = 200$ and $p = 6$.

expected for a matrix whose spectrum falls in the interval $[\alpha, \beta]$. (The improvement is substantial: for $\alpha = 1/8$ and $\beta = 1$, the early slow rate is $\rho_1 \approx 0.8819$, which is followed by the rate $\rho_2 \approx 0.4776$: thus it takes nearly six slow-phase iterations to reduce the residual as much as a single fast-phase iteration.) Since \mathbf{A} is normal, this discussion leads to a rigorous bound on GMRES convergence: for $k \geq 2p + 2$,

$$(3.4) \quad \frac{\|\mathbf{r}_k\|}{\|\mathbf{r}_0\|} \leq 2^{1+p/2} (1 + \beta/\alpha)^{p/2} (1 + \beta^2/\gamma^2) \rho_2^{k-p-2},$$

arrived at by bounding the GMRES residual polynomial with an inferior polynomial that has roots at the negative eigenvalues $\lambda_1, \dots, \lambda_p$ and the imaginary eigenvalues $\pm \gamma i$, and behaves like a Chebyshev polynomial on $[\alpha, \beta]$; see, e.g., [9, sect. 3.1], [34, pp. 70–71].

Figure 3.3 compares the performance of PGMRES and modified Gram–Schmidt GMRES for two instances of the matrix (3.3). The first instance is constructed to have a mild condition number; the second is more ill conditioned due to the large value of $\gamma = \|\mathbf{FG}^*\| = 10^6$. This extra magnitude brings forward the onset of instability, which is already significant at the fifth iteration. In both instances the PGMRES basis loses orthogonality, then linear dependence, just as for the Lippmann–Schwinger example shown in Figure 3.1.

3.3. Analysis of PGMRES orthogonalization. We next demonstrate how the local orthogonalization that gives PGMRES its distinct performance advantage over standard GMRES can cause the numerical instabilities exhibited in the previous computations. Begin with an *exact* decomposition resulting from $k - 1$ steps of the Arnoldi process: $\mathbf{AV}_{k-2} = \mathbf{V}_{k-1}\tilde{\mathbf{H}}_{k-1}$. Now suppose that the PGMRES Arnoldi process is used to compute subsequent Arnoldi vectors. We shall demonstrate how an error incurred at step k can be magnified by the next step of the PGMRES recurrence.

Let $\mathbf{V}_{k-1} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k-1}]$ denote the matrix whose columns are the first $k - 1$ exact Arnoldi vectors, an orthonormal basis for $\mathcal{K}_{k-1}(\mathbf{A}, \mathbf{b})$. Suppose the k th computed Arnoldi vector $\tilde{\mathbf{v}}_k$ is a unit vector with some error in direction, i.e., $\tilde{\mathbf{v}}_k = c\mathbf{v}_k + s\mathbf{d}_k$, where \mathbf{v}_k is the k th exact Arnoldi vector, $\|\mathbf{d}_k\| = 1$ and $\mathbf{d}_k \perp \mathbf{v}_k$, and $|c|^2 + |s|^2 = 1$ with $|s| \ll 1$. Now suppose $\tilde{\mathbf{v}}_k$ was computed by the PGMRES Arnoldi method (that is, via *local* orthogonalization of $(\mathbf{A} - \mathbf{V}_{k-1}\mathbf{V}_{k-1}^*\mathbf{FG}^*)\mathbf{v}_{k-1}$ against \mathbf{v}_{k-1} and \mathbf{v}_{k-2} ; see (2.6)). For a first-order analysis, it is reasonable to presume the explicit orthogonalization is exact, i.e., $\tilde{\mathbf{v}}_k \perp \{\mathbf{v}_{k-1}, \mathbf{v}_{k-2}\}$, and thus $\mathbf{d}_k \perp \{\mathbf{v}_{k-1}, \mathbf{v}_{k-2}\}$.




FIG. 3.3. PGMRES (solid line) and standard (modified Gram–Schmidt) GMRES (dashed line) as in Figure 3.1, applied to the simple example (3.3) with a vector $\mathbf{b} = [1, \dots, 1]^T$ and $n = 200$. The dotted line in each left plot shows the convergence bound (3.4). The gray line in each right plot shows the smallest singular value of the Arnoldi basis matrix \mathbf{V}_k computed by PGMRES. In the well-conditioned top example ($\text{cond}(\mathbf{A}) = 32$), the instability develops gradually; in the ill-conditioned bottom example ($\text{cond}(\mathbf{A}) = 1.25 \times 10^6$), the instability is apparent at the fifth iteration.

We shall use the notation \mathbf{P}_ℓ and $\tilde{\mathbf{P}}_\ell$ for the orthogonal projectors onto $\text{span}\{\mathbf{v}_\ell\}^\perp$ and $\text{span}\{\tilde{\mathbf{v}}_\ell\}^\perp$. To study the error in $\tilde{\mathbf{v}}_{k+1}$, let $\mathbf{u} := (\mathbf{A} - \tilde{\mathbf{V}}_k \tilde{\mathbf{V}}_k^* \mathbf{F} \mathbf{G}^*) \tilde{\mathbf{v}}_k$, and consider the unnormalized $(k+1)$ st Arnoldi vector $\tilde{\mathbf{w}}_{k+1} := \tilde{\mathbf{P}}_k \mathbf{P}_{k-1} \mathbf{u}$ computed by PGMRES. Then

$$\begin{aligned}
 (3.5) \quad \mathbf{u} &= \mathbf{A}(c\mathbf{v}_k + s\mathbf{d}_k) - \tilde{\mathbf{V}}_k \tilde{\mathbf{V}}_k^* \mathbf{F} \mathbf{G}^* (c\mathbf{v}_k + s\mathbf{d}_k) \\
 &= c(\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{v}_k + s(\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{d}_k \\
 &\quad - c(c\bar{s}\mathbf{v}_k \mathbf{d}_k^* + s\bar{c}\mathbf{d}_k \mathbf{v}_k^*) \mathbf{F} \mathbf{G}^* \mathbf{v}_k + \mathcal{O}(s^2),
 \end{aligned}$$

and one can show that

$$\begin{aligned}
 (3.6) \quad \tilde{\mathbf{w}}_{k+1} &= \tilde{\mathbf{P}}_k \mathbf{P}_{k-1} \mathbf{u} \\
 &= \mathbf{P}_k \mathbf{P}_{k-1} \mathbf{u} - (|s|^2(\mathbf{d}_k \mathbf{d}_k^* - \mathbf{v}_k \mathbf{v}_k^*) + c\bar{s}\mathbf{v}_k \mathbf{d}_k^* + s\bar{c}\mathbf{d}_k \mathbf{v}_k^*) \mathbf{P}_{k-1} \mathbf{u} \\
 &= c\mathbf{w}_{k+1} + s\mathbf{P}_k \mathbf{P}_{k-1} (\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{d}_k - \bar{s}c^2(\mathbf{d}_k^* \mathbf{w}_{k+1}) \mathbf{v}_k \\
 &\quad - s|c|^2(\mathbf{v}_k^* \mathbf{A} \mathbf{v}_k) \mathbf{d}_k + \mathcal{O}(s^2) \\
 &= c\mathbf{w}_{k+1} + \mathbf{f}_{k+1} + \mathcal{O}(s^2),
 \end{aligned}$$

where $\mathbf{w}_{k+1} := \mathbf{P}_k \mathbf{P}_{k-1} (\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{v}_k$ is the unnormalized $(k+1)$ st exact Arnoldi vector, and

$$(3.7) \quad \mathbf{f}_{k+1} := s \mathbf{P}_k \mathbf{P}_{k-1} (\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{d}_k - \bar{\sigma} c^2 (\mathbf{d}_k^* \mathbf{w}_{k+1}) \mathbf{v}_k - s |c|^2 (\mathbf{v}_k^* \mathbf{A} \mathbf{v}_k) \mathbf{d}_k,$$

so $\|\mathbf{f}_{k+1}\| = \mathcal{O}(s)$. Consider the orthogonal decomposition $c \mathbf{w}_{k+1} = \mathbf{w}_{k+1}^\perp + \mathbf{w}_{k+1}^\parallel$, where $\mathbf{w}_{k+1}^\perp \perp \tilde{\mathbf{w}}_{k+1}$ and $\mathbf{w}_{k+1}^\parallel$ is parallel to $\tilde{\mathbf{w}}_{k+1}$. It follows that

$$(3.8) \quad \|\mathbf{w}_{k+1}^\perp\| = \|c \mathbf{w}_{k+1} - \mathbf{w}_{k+1}^\parallel\| \leq \|c \mathbf{w}_{k+1} - \tilde{\mathbf{w}}_{k+1}\| = \|\mathbf{f}_{k+1}\| + \mathcal{O}(s^2),$$

and therefore $\sin \angle(\tilde{\mathbf{w}}_{k+1}, \mathbf{w}_{k+1}) = \|\mathbf{w}_{k+1}^\perp\| / \|c \mathbf{w}_{k+1}\| \leq \|\mathbf{f}_{k+1}\| / \|c \mathbf{w}_{k+1}\| + \mathcal{O}(s^2)$.

To simplify the analysis, we make one further assumption: $\tilde{\mathbf{v}}_k \in \mathcal{K}_k(\mathbf{A}, \mathbf{v}_1)$; in other words, the lack of explicit orthogonalization of $\mathbf{A} \mathbf{v}_{k-1}$ against $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k-3}\}$ in the PGMRES Arnoldi process only introduces errors in the space $\mathcal{K}_k(\mathbf{A}, \mathbf{v}_1)$. (We observe that this assumption is nearly attained in practical computations.) This implies that $\mathbf{d}_k = (\tilde{\mathbf{v}}_k - c \mathbf{v}_k) / s \in \mathcal{K}_k(\mathbf{A}, \mathbf{v}_1)$. Since $\mathbf{d}_k \perp \{\mathbf{v}_{k-2}, \mathbf{v}_{k-1}, \mathbf{v}_k\}$, it follows that $\mathbf{d}_k \in \mathcal{K}_{k-3}(\mathbf{A}, \mathbf{v}_1)$. In addition, note from (3.7) that since $\mathbf{f}_{k+1} \in \mathcal{K}_k(\mathbf{A}, \mathbf{v}_1)$, we have $\mathbf{w}_{k+1} \perp \mathbf{f}_{k+1}$. Hence $\tilde{\mathbf{w}}_{k+1} = c \mathbf{w}_{k+1} + \mathbf{f}_{k+1} + \mathcal{O}(s^2)$ from (3.6) is an orthogonal decomposition of $\tilde{\mathbf{w}}_{k+1}$, up to $\mathcal{O}(s^2)$ terms.

Since $\mathbf{d}_k \in \mathcal{K}_{k-3}(\mathbf{A}, \mathbf{v}_1)$ implies $\mathbf{w}_{k+1} \perp \mathbf{d}_k$, we can use (3.7) to compute

$$(3.9) \quad \begin{aligned} \frac{\tan \angle(\mathbf{v}_{k+1}, \tilde{\mathbf{v}}_{k+1})}{\tan \angle(\mathbf{v}_k, \tilde{\mathbf{v}}_k)} &= \frac{\|\mathbf{f}_{k+1}\| / \|c \mathbf{w}_{k+1}\| + \mathcal{O}(s^2)}{|s|/|c|} \\ &= \frac{\|\mathbf{P}_k \mathbf{P}_{k-1} (\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{d}_k - |c|^2 (\mathbf{v}_k^* \mathbf{A} \mathbf{v}_k) \mathbf{d}_k\|}{\|\mathbf{P}_k \mathbf{P}_{k-1} (\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{v}_k\|} + \mathcal{O}(s). \end{aligned}$$

Recall from (1.2) that $\mathbf{H} = \mathbf{A} - \frac{1}{2} \mathbf{F} \mathbf{G}^*$. When the low-rank skew-Hermitian part of \mathbf{A} is large, $\|\frac{1}{2} \mathbf{F} \mathbf{G}^*\| \gg \|\mathbf{A} - \frac{1}{2} \mathbf{F} \mathbf{G}^*\|$, the formula (3.9) raises particular concern, as it suggests considerable growth in the directional error of the Arnoldi vectors generated by PGMRES.² In such scenarios, matrix-vector products $\mathbf{A} \mathbf{v}$ are dominated by $\mathbf{F} \mathbf{G}^* \mathbf{v}$ for generic \mathbf{v} , so one expects the Krylov space $\mathcal{K}_{s+1}(\mathbf{A}, \mathbf{b})$ to essentially contain the s -dimensional subspace $\text{Ran}(\mathbf{F})$. (The starting vector has no bias toward $\text{Ran}(\mathbf{F})$; this space emerges through the first s matrix-vector products with \mathbf{A} , i.e., $s + 1$ Krylov vectors.) For $k \geq s + 1$, we thus approximate $\mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^* \approx \mathbf{F} \mathbf{G}^*$, so using the form (1.2),

$$\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^* \approx \mathbf{A}^*.$$

Similarly, for $k > s + 1$ we expect $\mathbf{F}^* \mathbf{v}_k \approx \mathbf{0}$, and so, in the notation of (1.2),

$$\mathbf{A}^* \mathbf{v}_k = (\mathbf{H} + \frac{1}{2} \mathbf{G} \mathbf{F}^*) \mathbf{v}_k \approx \mathbf{H} \mathbf{v}_k,$$

where \mathbf{H} denotes the Hermitian part of \mathbf{A} . These observations suggest that the denominator in (3.9) can be approximated by $\|\mathbf{P}_k \mathbf{P}_{k-1} \mathbf{H} \mathbf{v}_k\| \leq \|\mathbf{H}\| \ll \|\mathbf{A}\|$. The same argument gives an approximation to the second term in the numerator of (3.9), via $|\mathbf{v}_k^* \mathbf{A} \mathbf{v}_k| \approx |\mathbf{v}_k^* \mathbf{H} \mathbf{v}_k| \ll \|\mathbf{A}\|$. Now the first term in that numerator behaves like

$$(3.10) \quad \mathbf{P}_k \mathbf{P}_{k-1} (\mathbf{A} - \mathbf{V}_k \mathbf{V}_k^* \mathbf{F} \mathbf{G}^*) \mathbf{d}_k \approx \mathbf{P}_k \mathbf{P}_{k-1} \mathbf{A}^* \mathbf{V}_{k-3} \mathbf{z}$$

²Note that such instability will not occur for MINRES applied to Hermitian \mathbf{A} .

for some $\mathbf{z} \in \mathbb{C}^{k-3}$, since $\mathbf{d}_k \in \mathcal{K}_{k-3}(\mathbf{A}, \mathbf{b})$. Presuming \mathbf{d}_k to arise from an arbitrary perturbation, we expect (3.10) to be on the order of $\|\mathbf{A}\|$, so (3.9) will have a large numerator and small denominator: for $k > s + 1$, we expect the k th iteration can magnify the angular error in the PGMRES Arnoldi basis vector on the order of $\|\mathbf{A}\|/\|\mathbf{H}\|$. Indeed, in numerical experiments like those shown in Figures 3.1 and 3.3, increasing $\|\mathbf{F}\mathbf{G}^*\|$ brings about earlier stagnation of PGMRES. Successful PGMRES computations seem to require that $\mathbf{F}\mathbf{G}^*$ be small in both rank *and* norm.

In summary, the step that causes the loss of orthogonality in the PGMRES Arnoldi procedure is the same step that makes PGMRES so computationally attractive, alleviating the need to preserve all Arnoldi basis vectors. In cases where the skew-Hermitian part of \mathbf{A} is small in both rank and norm, and GMRES converges steadily, our experience suggests that PGMRES can be viable; otherwise, numerical instabilities often induce stagnation before convergence to a reasonable tolerance. Though we cannot propose a repair for this instability, in the next section we suggest an alternative method for efficiently solving a nearly Hermitian linear system. This method will have similar storage characteristics, but will avoid the numerical problems endemic to PGMRES.

4. An alternative approach. We seek an efficient alternative method for solving $\mathbf{A}\mathbf{x} = \mathbf{b}$ when \mathbf{A} has the “nearly Hermitian” form (1.2) that avoids the unstable performance of PGMRES. Here we present such a method that requires only three-term recurrences. This algorithm follows from the simple observation that we can express nearly Hermitian matrices in the form

$$(4.1) \quad \mathbf{A} = \mathbf{H} + \mathbf{F}\mathbf{C}\mathbf{F}^*,$$

where $\mathbf{H} := (\mathbf{A} + \mathbf{A}^*)/2$, the Hermitian part of \mathbf{A} , is assumed to be invertible, and $\mathbf{F}\mathbf{C}\mathbf{F}^* = (\mathbf{A} - \mathbf{A}^*)/2$ is a decomposition of the low-rank skew-Hermitian part, with $\mathbf{C} \in \mathbb{C}^{s \times s}$ skew-Hermitian and $\mathbf{F} \in \mathbb{C}^{n \times s}$. For example, one can think of $\mathbf{F}\mathbf{C}\mathbf{F}^*$ as a reduced unitary diagonalization of the skew-Hermitian part, or one can compute \mathbf{C} from the representation (1.2) by solving $\mathbf{F}\mathbf{C}^* = \frac{1}{2}\mathbf{G}$. Other decompositions may follow more naturally from the underlying mathematical model.

When the Hermitian part \mathbf{H} is invertible, one can use it as a preconditioner, e.g.,

$$\mathbf{H}^{-1}\mathbf{A} = \mathbf{I} + \mathbf{H}^{-1}\mathbf{F}\mathbf{C}\mathbf{F}^*.$$

A rank- s perturbation of the identity, this matrix will require no more than $s + 1$ iterations of GMRES to reach convergence (provided \mathbf{H} is nonsingular). The philosophy behind this preconditioner is the same as in the CGW method [4, 37], except that in the latter, the Hermitian part needs to be positive definite; see also [15, Thm. 2.10]. To form the preconditioned system, one can first compute $\mathbf{H}^{-1}\mathbf{F}$ and $\mathbf{H}^{-1}\mathbf{b}$ by solving $s + 1$ Hermitian linear systems *for the same coefficient matrix* \mathbf{H} using the short-recurrence MINRES [20] or conjugate gradient [13] algorithms (possibly as a block method, or in parallel). Inexact computation of the preconditioner can affect GMRES convergence, as discussed in [27], [30, sect. 6], spoiling exact convergence. Here we describe an alternative strategy, based on the Schur complement, that requires the solution of $s + 1$ systems with MINRES, followed by the solution of an $s \times s$ linear system; no GMRES iterations are necessary.

Note that (4.1) is the Schur complement of $-\mathbf{C}^{-1}$ in the matrix

$$\Phi = \begin{bmatrix} \mathbf{H} & \mathbf{F} \\ \mathbf{F}^* & -\mathbf{C}^{-1} \end{bmatrix},$$

so solving $\mathbf{Ax} = \mathbf{b}$ is equivalent to solving

$$(4.2) \quad \begin{bmatrix} \mathbf{H} & \mathbf{F} \\ \mathbf{F}^* & -\mathbf{C}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix}.$$

Schur complement methods for solving systems of the form (4.2) are well known; see, e.g., [2, 5]. One such method eliminates \mathbf{x} by inserting $\mathbf{x} = \mathbf{H}^{-1}(\mathbf{b} - \mathbf{F}\mathbf{y})$ into $\mathbf{F}^*\mathbf{x} - \mathbf{C}^{-1}\mathbf{y} = \mathbf{0}$, and solving for \mathbf{y} through the s -dimensional system

$$(4.3) \quad (\mathbf{F}^*\mathbf{H}^{-1}\mathbf{F} + \mathbf{C}^{-1})\mathbf{y} = \mathbf{F}^*\mathbf{H}^{-1}\mathbf{b}.$$

When the formulation (1.2) is more natural, this last equation takes the form

$$(4.4) \quad (\mathbf{G}^*\mathbf{H}^{-1}\mathbf{F} + 2\mathbf{I})\mathbf{y} = \mathbf{G}^*\mathbf{H}^{-1}\mathbf{b}.$$

This approach is equivalent to applying the Sherman–Morrison–Woodbury formula to (4.1); see, e.g., [12, 26, 38].

The solution of (1.1) via the method just described requires the solution of

$$(4.5) \quad \mathbf{HW} = \mathbf{F} \quad \text{and} \quad \mathbf{Hu} = \mathbf{b},$$

for $\mathbf{W} \in \mathbb{C}^{n \times s}$ and $\mathbf{u} \in \mathbb{C}^n$, as well as the $s \times s$ system (4.3) or (4.4) for \mathbf{y} . From these ingredients, one can construct $\mathbf{x} = \mathbf{u} - \mathbf{W}\mathbf{y}$. This approach is described in a more general setting in, e.g., [12, 39].

Since \mathbf{H} is Hermitian, one can approximate the solutions to (4.5) via a Krylov subspace method driven by the three-term Lanczos recurrence. One can use MINRES or, for positive definite \mathbf{H} , the conjugate gradient (CG) method. Thus, constructing an approximate solution to (1.1) requires $s + 1$ Hermitian solves. Methods for solving a Hermitian system with several right-hand sides (see, e.g., [11, 17, 24, 29]) could potentially expedite this calculation.

The decomposition of the skew-Hermitian part into \mathbf{FG}^* or \mathbf{FCF}^* is not unique, so one could, in principle, select the factorization in a way that optimizes the conditioning of the system in (4.3) or (4.4). For example, a result of Yip [39, Thm. 1] implies that some choice of \mathbf{F} and \mathbf{G} ensures that the condition number $\text{cond}(\mathbf{G}^*\mathbf{H}^{-1}\mathbf{F} + 2\mathbf{I})$ is bounded above by $\text{cond}(\mathbf{A})\text{cond}(\mathbf{H})$.

We consider next the matter of choosing suitable stopping criteria for the Hermitian solves. For simplicity, assume henceforth in this section that \mathbf{F} is scaled so that $\|\mathbf{F}\| = 1$, as is the case when \mathbf{F} is derived from a reduced unitary diagonalization of the skew-Hermitian part of \mathbf{A} . Therefore $\|\mathbf{A} - \mathbf{A}^*\|/2 = \|\mathbf{FCF}^*\| \leq \|\mathbf{C}\|$. Let $\widetilde{\mathbf{W}}$ and $\widetilde{\mathbf{u}}$ denote approximate solutions to $\mathbf{HW} = \mathbf{F}$ and $\mathbf{Hu} = \mathbf{b}$ derived, e.g., from MINRES. With these approximations in hand, one would replace (4.3) with the perturbed s -dimensional system

$$(4.6) \quad (\mathbf{F}^*\widetilde{\mathbf{W}} + \mathbf{C}^{-1})\widetilde{\mathbf{y}} = \mathbf{F}^*\widetilde{\mathbf{u}},$$

which can be solved directly with Gaussian elimination to yield the approximation

$$\widetilde{\mathbf{x}} = \widetilde{\mathbf{u}} - \widetilde{\mathbf{W}}\widetilde{\mathbf{y}}$$

to the desired solution \mathbf{x} . (For purposes of this analysis, we implicitly assume that this direct solve is computed exactly.) The residual of this approximation can be expressed in terms of the other approximations:

$$\begin{aligned} \tilde{\mathbf{r}} &:= \mathbf{b} - \mathbf{A}\tilde{\mathbf{x}} \\ &= \mathbf{b} - (\mathbf{H} + \mathbf{F}\mathbf{C}\mathbf{F}^*)(\tilde{\mathbf{u}} - \widetilde{\mathbf{W}}\tilde{\mathbf{y}}) \\ &= \mathbf{b} - \mathbf{H}\tilde{\mathbf{u}} + \mathbf{H}\widetilde{\mathbf{W}}\tilde{\mathbf{y}} - \mathbf{F}\mathbf{C}(\mathbf{F}^*\tilde{\mathbf{u}} - \mathbf{F}^*\widetilde{\mathbf{W}}\tilde{\mathbf{y}}) \\ &= \mathbf{b} - \mathbf{H}\tilde{\mathbf{u}} + (\mathbf{H}\widetilde{\mathbf{W}} - \mathbf{F})\tilde{\mathbf{y}}. \end{aligned}$$

Basic norm inequalities yield a simple, dynamic stopping criterion.

THEOREM 4.1. *Suppose $\mathbf{A} \in \mathbb{C}^{n \times n}$ is a nonsingular matrix of the form (4.1) with $\|\mathbf{F}\| \leq 1$ and nonsingular Hermitian part \mathbf{H} . Let $\widetilde{\mathbf{W}}$ and $\tilde{\mathbf{u}}$ be approximate solutions to $\mathbf{H}\mathbf{W} = \mathbf{F}$ and $\mathbf{H}\mathbf{u} = \mathbf{b}$ with residuals $\mathbf{R}_{\mathbf{W}} := \mathbf{F} - \mathbf{H}\widetilde{\mathbf{W}}$ and $\mathbf{r}_{\mathbf{u}} := \mathbf{b} - \mathbf{H}\tilde{\mathbf{u}}$, and suppose further that $\mathbf{F}^*\widetilde{\mathbf{W}} + \mathbf{C}^{-1}$ is nonsingular. Then provided*

$$\frac{\|\mathbf{R}_{\mathbf{W}}\|\|\tilde{\mathbf{y}}\|}{\|\mathbf{b}\|} < \varepsilon/2 \quad \text{and} \quad \frac{\|\mathbf{r}_{\mathbf{u}}\|}{\|\mathbf{b}\|} < \varepsilon/2,$$

and $\tilde{\mathbf{y}}$ exactly solves (4.6), then the approximate solution $\tilde{\mathbf{x}} := \tilde{\mathbf{u}} - \widetilde{\mathbf{W}}\tilde{\mathbf{y}}$ satisfies

$$(4.7) \quad \frac{\|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}\|}{\|\mathbf{b}\|} < \varepsilon.$$

Since $\tilde{\mathbf{y}}$ depends on $\widetilde{\mathbf{W}}$, the stopping criterion for $\widetilde{\mathbf{W}}$ (i.e., $\|\mathbf{R}_{\mathbf{W}}\|\|\tilde{\mathbf{y}}\|/\|\mathbf{b}\| < \varepsilon/2$) in Theorem 4.1 cannot be expressed a priori. Once a candidate value for $\widetilde{\mathbf{W}}$ has been found, one can solve the small $s \times s$ system (4.6) for $\tilde{\mathbf{y}} \in \mathbb{C}^s$, where $s \ll n$.³ With $\tilde{\mathbf{y}}$ in hand, one can check if $\widetilde{\mathbf{W}}$ satisfies the stopping criterion; if not, conduct further MINRES iterations to refine $\widetilde{\mathbf{W}}$, and test the criterion again with the updated $\tilde{\mathbf{y}}$.

Adapting notation slightly, let $\widehat{\mathbf{W}}$ and $\widetilde{\mathbf{W}}$ denote two approximate solutions to $\mathbf{H}\mathbf{W} = \mathbf{F}$ with corresponding solutions $\widehat{\mathbf{y}}$ and $\tilde{\mathbf{y}}$ to (4.6). The following result quantifies the rate at which $\widehat{\mathbf{y}} \rightarrow \tilde{\mathbf{y}}$ as $\widehat{\mathbf{W}} \rightarrow \widetilde{\mathbf{W}}$, thus emphasizing that the dynamic stopping criterion supplied by Theorem 4.1 is stable with respect to refinements to $\widetilde{\mathbf{W}}$.

THEOREM 4.2. *Let $\tilde{\mathbf{y}}, \widehat{\mathbf{y}} \in \mathbb{C}^s$ solve the nonsingular linear systems*

$$\begin{aligned} (\mathbf{F}^*\widetilde{\mathbf{W}} + \mathbf{C}^{-1})\tilde{\mathbf{y}} &= \mathbf{F}^*\tilde{\mathbf{u}}, \\ (\mathbf{F}^*\widehat{\mathbf{W}} + \mathbf{C}^{-1})\widehat{\mathbf{y}} &= \mathbf{F}^*\tilde{\mathbf{u}}, \end{aligned}$$

with $\|\mathbf{F}\| \leq 1$. Then for sufficiently small $\|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\|$ and $\widetilde{\mathbf{\Omega}} := \mathbf{F}^*\widetilde{\mathbf{W}} + \mathbf{C}^{-1}$,

$$\|\tilde{\mathbf{y}} - \widehat{\mathbf{y}}\| \leq \frac{\|\widetilde{\mathbf{\Omega}}^{-1}\| \|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\| \|\tilde{\mathbf{y}}\|}{1 - \|\widetilde{\mathbf{\Omega}}^{-1}\| \|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\|}$$

and

$$\|\widehat{\mathbf{y}}\| \leq \|\tilde{\mathbf{y}}\| + \frac{\|\widetilde{\mathbf{\Omega}}^{-1}\| \|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\| \|\tilde{\mathbf{y}}\|}{1 - \|\widetilde{\mathbf{\Omega}}^{-1}\| \|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\|}.$$

Proof. Since the difference between the coefficient matrix in solving for $\tilde{\mathbf{y}}$ and $\widehat{\mathbf{y}}$ is bounded in norm by $\|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\|$, the result follows directly from basic perturbation theory for linear systems [14, Thm. 7.2], provided $\|\widetilde{\mathbf{\Omega}}^{-1}\| \|\widetilde{\mathbf{W}} - \widehat{\mathbf{W}}\| < 1$. \square

³In the case of the one-dimensional scattering problem in the last section, $s = 2$.

Recasting $\mathbf{Ax} = \mathbf{b}$ into $\mathbf{Hx} + \mathbf{Fy} = \mathbf{b}$ and $\mathbf{CF}^*\mathbf{x} = \mathbf{y}$ expresses the right-hand side vector \mathbf{b} as the sum of vectors in $\text{Ran}(\mathbf{H})$ and $\text{Ran}(\mathbf{F})$. An a priori bound for $\|\mathbf{y}\|$ or $\|\tilde{\mathbf{y}}\|$ would thus require knowledge of quantities such as $\|\mathbf{A}^{-1}\|$ or $\|\mathbf{H}^{-1}\|$, both of which are computationally prohibitive, particularly compared to the cost of evaluating the dynamic bound in Theorem 4.1. We provide such a bound mainly for theoretical interest.

THEOREM 4.3. *Suppose $\|\mathbf{F}\| \leq 1$, and that \mathbf{y} and $\tilde{\mathbf{y}}$ solve*

$$\mathbf{\Omega y} = \mathbf{F}^*\mathbf{u}, \quad \tilde{\mathbf{\Omega}}\tilde{\mathbf{y}} = \mathbf{F}^*\tilde{\mathbf{u}}$$

for nonsingular $\mathbf{\Omega} := \mathbf{F}^*\mathbf{W} + \mathbf{C}^{-1}$ and $\tilde{\mathbf{\Omega}} := \mathbf{F}^*\tilde{\mathbf{W}} + \mathbf{C}^{-1}$. If $\|\mathbf{H}^{-1}\| \|\mathbf{\Omega}^{-1}\| \|\mathbf{R}_\mathbf{W}\| < 1$, then

$$\|\tilde{\mathbf{y}}\| \leq \|\mathbf{\Omega}^{-1}\| \|\mathbf{H}^{-1}\| \left(\frac{\|\mathbf{b}\| + \|\mathbf{r}_\mathbf{u}\|}{1 - \|\mathbf{H}^{-1}\| \|\mathbf{\Omega}^{-1}\| \|\mathbf{R}_\mathbf{W}\|} \right),$$

where $\mathbf{R}_\mathbf{W} := \mathbf{F} - \mathbf{H}\tilde{\mathbf{W}}$ and $\mathbf{r}_\mathbf{u} := \mathbf{b} - \mathbf{H}\tilde{\mathbf{u}}$.

Proof. The perturbation bound [14, Thm. 7.2] implies that

$$\begin{aligned} \|\tilde{\mathbf{y}}\| &\leq \|\mathbf{y}\| + \|\tilde{\mathbf{y}} - \mathbf{y}\| \\ (4.8) \quad &\leq \|\mathbf{y}\| + \frac{\|\mathbf{H}^{-1}\| \|\mathbf{\Omega}^{-1}\|}{1 - \|\mathbf{H}^{-1}\| \|\mathbf{\Omega}^{-1}\| \|\mathbf{R}_\mathbf{W}\|} (\|\mathbf{r}_\mathbf{u}\| + \|\mathbf{R}_\mathbf{W}\| \|\mathbf{y}\|), \end{aligned}$$

since $\|\mathbf{\Omega} - \tilde{\mathbf{\Omega}}\| \leq \|\mathbf{H}^{-1}\| \|\mathbf{R}_\mathbf{W}\|$ and $\|\mathbf{F}^*\mathbf{u} - \mathbf{F}^*\tilde{\mathbf{u}}\| \leq \|\mathbf{H}^{-1}\| \|\mathbf{r}_\mathbf{u}\|$. The fact that $\mathbf{y} = \mathbf{\Omega}^{-1}\mathbf{F}^*\mathbf{u} = \mathbf{\Omega}^{-1}\mathbf{F}^*\mathbf{H}^{-1}\mathbf{b}$ implies $\|\mathbf{y}\| \leq \|\mathbf{\Omega}^{-1}\| \|\mathbf{H}^{-1}\| \|\mathbf{b}\|$, from which follows

$$\begin{aligned} \|\tilde{\mathbf{y}}\| &\leq \|\mathbf{\Omega}^{-1}\| \|\mathbf{H}^{-1}\| \left(\|\mathbf{b}\| + \frac{\|\mathbf{r}_\mathbf{u}\| + \|\mathbf{\Omega}^{-1}\| \|\mathbf{R}_\mathbf{W}\| \|\mathbf{H}^{-1}\| \|\mathbf{b}\|}{1 - \|\mathbf{H}^{-1}\| \|\mathbf{\Omega}^{-1}\| \|\mathbf{R}_\mathbf{W}\|} \right) \\ &= \|\mathbf{\Omega}^{-1}\| \|\mathbf{H}^{-1}\| \left(\frac{\|\mathbf{b}\| + \|\mathbf{r}_\mathbf{u}\|}{1 - \|\mathbf{H}^{-1}\| \|\mathbf{\Omega}^{-1}\| \|\mathbf{R}_\mathbf{W}\|} \right). \quad \square \end{aligned}$$

COROLLARY 4.4. *Provided the approximate solutions $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{W}}$ are sufficiently accurate that*

$$\frac{\|\mathbf{r}_\mathbf{u}\|}{\|\mathbf{b}\|} \leq \frac{\varepsilon}{2} \quad \text{and} \quad \|\mathbf{R}_\mathbf{W}\| \leq \frac{\varepsilon}{\|\mathbf{\Omega}^{-1}\| \|\mathbf{H}^{-1}\| (2 + 2\varepsilon)},$$

the relative residual satisfies $\|\mathbf{r}\|/\|\mathbf{b}\| < \varepsilon$.

We emphasize that the quantity $\|\mathbf{\Omega}^{-1}\|$ in the hypothesis and bound renders this result inapplicable a priori, since $\mathbf{\Omega}$ is a function of \mathbf{W} . Approximating the value $\|\mathbf{\Omega}^{-1}\|$ from $\tilde{\mathbf{W}}$ could give a dynamic estimate. Note that given an approximation $\tilde{\mathbf{W}}$, one can directly compute $\|\tilde{\mathbf{y}}\|$.

The Schur complement method for solving nearly Hermitian linear systems is summarized in Algorithm 4.1, which uses the dynamic stopping criterion. Since this convergence test involves $\tilde{\mathbf{y}}$, which requires an approximation to \mathbf{W} , MINRES is applied to the problems $\mathbf{H}\mathbf{w}_j = \mathbf{f}_j$ for $j = 1, \dots, s$ concurrently. Lines 8 through 10 of the algorithm can be replaced by a step-by-step block MINRES [17]. (Block methods for multiple right-hand sides can potentially give more rapid convergence than the aggregate cost of solving the systems one at a time [11, 17].) Also observe that it is not necessary to build the matrix \mathbf{H} ; one need only compute the matrix-vector product \mathbf{Hx} for a vector \mathbf{x} .

ALGORITHM 4.1: THE SCHUR COMPLEMENT METHOD

Input : $\mathbf{A} \in \mathbb{C}^{n \times n}$, $\mathbf{F}, \mathbf{G} \in \mathbb{C}^{n \times s}$, $\mathbf{b} \in \mathbb{C}^n$, $tol > 0$, $m =$ maximum number of MINRES (or CG) iterations

Output: $\tilde{\mathbf{x}}$, an approximate solution to $\mathbf{Ax} = \mathbf{b}$

$\mathbf{H} \leftarrow (\mathbf{A} + \mathbf{A}^*)/2$

$\tilde{\mathbf{u}} \leftarrow$ apply MINRES to $\mathbf{Hu} = \mathbf{b}$ with relative residual tolerance $tol/2$ and initial iterate $\tilde{\mathbf{u}}_0 = \mathbf{0}$ (or CG if \mathbf{H} is positive definite)

for $k = 1$ to m **do**

for $j = 1$ to s **do**

$\mathbf{f} \leftarrow$ j th column of \mathbf{F}

$\tilde{\mathbf{w}}_j \leftarrow \operatorname{argmin}_{\tilde{\mathbf{w}} \in \mathcal{K}_k(\mathbf{H}, \mathbf{f})} \|\mathbf{f} - \mathbf{H}\tilde{\mathbf{w}}\|$ via MINRES

if $\|\mathbf{F} - \mathbf{H}\tilde{\mathbf{W}}\| < tol/2$ **then**

$\tilde{\mathbf{y}} \leftarrow (\mathbf{F}^*\tilde{\mathbf{W}} + \mathbf{C}^{-1})^{-1}\mathbf{F}^*\mathbf{u}$

if $\|\mathbf{F} - \mathbf{H}\tilde{\mathbf{W}}\| \|\tilde{\mathbf{y}}\| / \|\mathbf{b}\| < tol/2$ **then**

End For Loop

$\tilde{\mathbf{x}} \leftarrow \tilde{\mathbf{u}} - \tilde{\mathbf{W}}\tilde{\mathbf{y}}$

There is no guarantee that \mathbf{H} is well conditioned, or even invertible (even when \mathbf{A} is well conditioned: indeed, \mathbf{H} is singular for the examples constructed in section 3.2). Furthermore, $\mathbf{F}^*\tilde{\mathbf{W}} + \mathbf{C}^{-1}$ can be singular, in which case accurate approximations $\tilde{\mathbf{W}}$ give matrices $\mathbf{F}^*\tilde{\mathbf{W}} + \mathbf{C}^{-1}$ that are highly ill conditioned. For such cases, Algorithm 4.1 would likely be unsuitable. While Algorithm 4.1 does offer a method for approximating the solution to (1.1) that requires only three-term recurrences, it is not more efficient than the PGMRES method, provided the latter method does not suffer from the instabilities identified in section 3. When such instabilities occur, Algorithm 4.1 provides an appealing alternative to full GMRES (with its long recurrences) or other methods based on short recurrences that do not satisfy any natural optimality properties. In the absence of such instabilities, PGMRES should generally be preferred to the Schur complement method, as PGMRES requires only that \mathbf{A} be nonsingular, and imposes no requirements on \mathbf{H} .

4.1. Use of the Schur complement method as a preconditioner. Given the utility of Algorithm 4.1 for solving linear systems for which the coefficient matrix has the form (1.2), one naturally wonders whether this Schur complement strategy can be used as a preconditioner for matrices that are close to the form (1.2). Apply \mathbf{M} as a right preconditioner by modifying (1.1) to have the form

$$(4.9) \quad \mathbf{AM}^{-1}\mathbf{z} = \mathbf{b}, \quad \mathbf{x} = \mathbf{M}^{-1}\mathbf{z},$$

where \mathbf{M} is a matrix that is cheap to compute, and makes \mathbf{AM}^{-1} more favorable for GMRES convergence than \mathbf{A} on its own. A right-preconditioned Krylov subspace method selects the j th approximation from the affine subspace

$$\mathbf{x}_j \in \mathbf{x}_0 + \operatorname{span}\{\mathbf{r}_0, \mathbf{AM}^{-1}\mathbf{r}_0, \dots, (\mathbf{AM}^{-1})^{j-1}\mathbf{r}_0\};$$

see, e.g., [25], [22, Chap. 9], [31].

Suppose $\mathbf{A} \in \mathbb{C}^{n \times n}$ has the decomposition $\mathbf{A} = \mathbf{H} + \mathbf{K}$, where \mathbf{H} is the Hermitian part, and the skew-Hermitian part \mathbf{K} can be well approximated by \mathbf{FG}^* for some

$\mathbf{F}, \mathbf{G} \in \mathbb{C}^{n \times s}$ with $s \ll n$, e.g., $\mathbf{K} = \mathbf{F}\mathbf{G}^* + \mathbf{E}$ with $\|\mathbf{E}\| \ll 1$. Systems such as this arise, for example, in discretizations of PDEs with certain Neuman boundary conditions, or more generically when \mathbf{K} has a small number of dominant singular values. In such instances, we expect $\mathbf{M} = \mathbf{H} + \mathbf{F}\mathbf{G}^*$ to be an effective preconditioner for \mathbf{A} , while allowing for rapid application through Algorithm 4.1.

Recall that Algorithm 4.1 requires $s + 1$ applications of MINRES, which could be prohibitive to apply at each iteration of GMRES applied to the preconditioned system. However, s of these applications are needed to solve $\mathbf{H}\mathbf{W} = \mathbf{F}$; the solution of this system can be computed once and reused at each GMRES iteration. Using \mathbf{v}_j here to denote the j th Arnoldi vector in the (outer) preconditioned GMRES iteration, each preconditioner application $\mathbf{M}^{-1}\mathbf{v}_j$ will require only the solution of one Hermitian system, $\mathbf{H}\hat{\mathbf{x}} = \mathbf{v}_j$, and the direct solution of one $s \times s$ system. Thus this preconditioner is relatively cheap to compute at each step, after the up-front cost of solving $\mathbf{H}\mathbf{W} = \mathbf{F}$.

5. Numerical results. To demonstrate the effectiveness of the proposed Schur complement approach, we apply Algorithm 4.1 to several problems. To put this new algorithm in context, we also solve our linear system using MATLAB's full GMRES algorithm [16] and the recently proposed IDR(2) method [32, 33], as implemented by van Gijzen [36]. The IDR(2) method is a modern short recurrence method that uses right and left Krylov subspaces. In the numerical examples in this section, the Schur complement method is implemented using the MINRES algorithm described in [7, sect. 6.5] (which has a similar level of overhead as the PGMRES and IDR implementations we use). In most of our experiments we apply the MINRES runs in series, but the Schur complement algorithm allows for easy parallelization of these calls, since each MINRES run can be done independently. In fact, we show one parallel numerical experiment in section 5.3. When PGMRES is successful, it is often the most efficient algorithm. (Except for the parallel computations in section 5.3, all experiments were run on a MacBook Pro with a 2.66 GHz Intel Core i7 processor and 4GB of DDR3 dynamic RAM.)

5.1. Simple example, revisited. In section 3.2, we introduced a simple, well-conditioned matrix (3.3) with a skew-Hermitian part of low rank. The analysis presented in section 3.3 suggests that as the norm of the skew-Hermitian part grows, the numerical instability of PGMRES causes the residual to stagnate earlier. To apply Algorithm 4.1, however, we require that the Hermitian part of \mathbf{A} be nonsingular; thus for the tests in this section, we modify (3.3) by adding a 2×2 identity to the \mathbf{Z} block (thus shifting the zero eigenvalues of the Hermitian part to one). As is clear from Figure 5.1, this does not change the convergence behavior discussed in section 3.2. It follows from Theorem 4.1 that

$$(5.1) \quad \gamma_k := \|\mathbf{r}_{\mathbf{b},k}\| + \sqrt{2}\|\mathbf{y}\|(\|\mathbf{r}_{\mathbf{f}_1,k}\| + \|\mathbf{r}_{\mathbf{f}_2,k}\|)$$

is an upper bound on the norm of the residual produced at each iteration of the Schur complement method. The vectors $\mathbf{r}_{\mathbf{b},k}$, $\mathbf{r}_{\mathbf{f}_1,k}$, and $\mathbf{r}_{\mathbf{f}_2,k}$ are the residuals produced by MINRES with coefficient matrix $\mathbf{H} = (\mathbf{A} + \mathbf{A}^*)/2$ and right-hand sides \mathbf{b} , \mathbf{f}_1 , and \mathbf{f}_2 . To simplify the illustration of our numerical results in Figure 5.1, we plot γ_k versus k for the Schur complement approach. Observe that this method is superior in run time to full GMRES, PGMRES, and IDR(2), and competitive or superior in iteration count as well. (Note that γ_k can be smaller than the full GMRES residual norm, since the Schur complement method does not draw its approximations from the same Krylov subspace from which GMRES draws its optimal iterates.) We chose IDR(2)

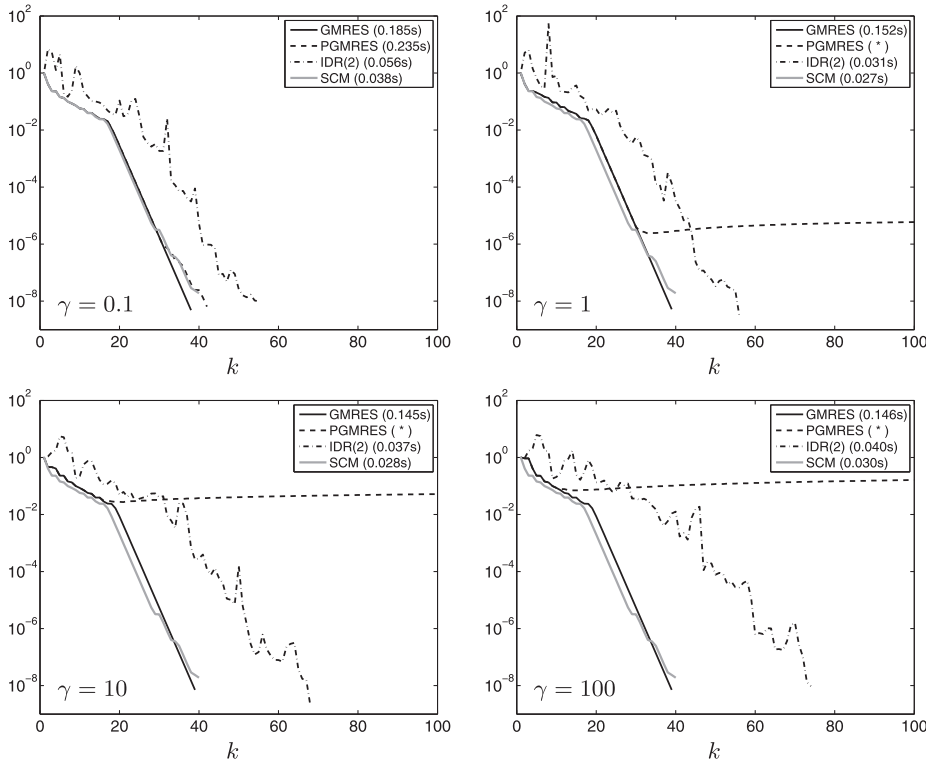


FIG. 5.1. Convergence of the Schur complement method for $\mathbf{Ax} = \mathbf{b}$, with coefficient matrix of dimension $n = 10^4$ given by the simple example (3.3) with a 2×2 identity added to the \mathbf{Z} block, and parameters $\alpha = 0.125$, $\beta = 1$, and $p = 6$; the vector \mathbf{b} is equal to $1/\sqrt{n}$ in all components. The gray line shows the upper bound γ_k in (5.1) on the residual norm. For context, the residual norms for full GMRES, PGMRES, and IDR(2) are also shown, along with timings. For the three largest values of γ , PGMRES(*) indicates that PGMRES does not converge in the given iterations.

(as opposed to, e.g., IDR(4)) to compare the Schur complement method to a modern short-term recurrence with a small storage requirement.

Whether the Schur complement method is better than preconditioning with the Hermitian part of the matrix depends on the structure and spectral properties of the preconditioned matrix. To illustrate the potential superiority of the Schur complement method, we alter the simple example to have an $s \times s$ skew-Hermitian block, with sub- and superdiagonal entries equal to 1 and -1 . The eigenvalues of the preconditioned matrix fill the complex interval $[1 - 2i, 1 + 2i]$ as s increases, thus keeping the conditioning of the preconditioned problem essentially fixed. The dimension of \mathbf{A} is fixed at $n = 100000$; the right-hand size vector is random. Figure 5.2 compares the CPU time of GMRES with Hermitian preconditioning to the Schur complement method, showing that the superiority of the latter algorithm improves as s grows.

5.2. Lippmann–Schwinger equation. Next, we apply Algorithm 4.1 to the Lippmann–Schwinger integral equation described in section 3.1. Figure 5.3 shows the convergence of each MINRES application required by the Schur complement method, when applied to the integral equation with wave number $\kappa = 10$ and constant refractive index $m = -1$.

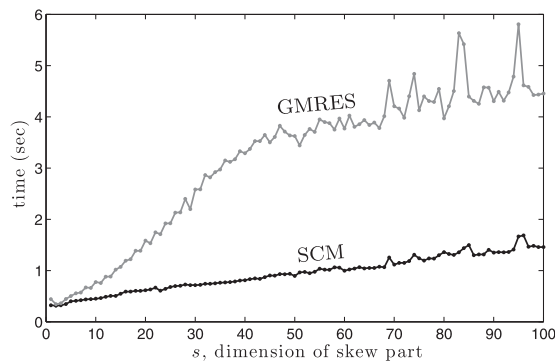


FIG. 5.2. Comparison of the Schur complement method to GMRES (with Hermitian preconditioning) for the simple example (3.3) of dimension $n = 10^5$ with skew-Hermitian part of rank s . The eigenvalues of the Hermitian-preconditioned matrix fill the interval $[1 - 2i, 1 + 2i]$, so the condition number of this matrix is essentially fixed for all these experiments.




FIG. 5.3. Residual norms for the MINRES iterations for the discretized Lippmann–Schwinger integral equation for wave number $\kappa = 10$ on a grid of $n = 10^3$ quadrature points with refractive indices $m(x) \equiv -1$ (left) and $m(x) = -2 - \sin(x)$ (right).

When the refractive index $m(x)$ varies in space, the skew-Hermitian part of the operator no longer has low rank. (Since the model describes wave scattering by a non-homogeneous obstacle, we assume throughout that $m(x) < 0$ for all $x \in (0, 2\pi)$.) In this case, let M denote the multiplication operator defined by $(Mu)(x) = m(x)u(x)$ on $L^2(0, 2\pi)$, and define the integral operator $K_0 : L^2(0, 2\pi) \rightarrow L^2(0, 2\pi)$ based on K in (3.2) but with m removed:

$$(K_0u)(x) = \frac{i\kappa}{2} \int_0^{2\pi} e^{i\kappa|x-y|}u(y) dy.$$

The integral equation (3.1) is then equivalent to

$$(5.2) \quad (M^{-1} + K_0)(Mu)(x) = u^i(x),$$

where $M^{-1} + K_0$ has the special structure required by Algorithm 4.1, since it is the sum of the self-adjoint operator M^{-1} with an operator K_0 that has a rank-2 skew-adjoint part. Indeed, we can view (5.2) as a “right preconditioned” version of (3.1), where the preconditioner is selected not to accelerate convergence, but rather to pose

the problem in a way that allows for solution via a short recurrence. (For another form of such strategic preconditioning, see [8].)

We approximate this problem with the Nyström discretization used in section 3.1, whereby the operator M is represented as a diagonal matrix \mathbf{M} with diagonal entries equal to m evaluated at each of the quadrature points. Figure 5.3 (right) shows the convergence of the MINRES iterations when Algorithm 4.1 is applied to the discretization of (5.2) for $m(x) = -2 - \sin(x)$ with wave number $\kappa = 10$ and $n = 10^3$ quadrature points.

To put this approach in context, we also apply the Schur complement approach to an equation with constant refractive index $m(x) = -1$ for various wave numbers κ , and a discretization of $n = 10^3$ points. We also apply MATLAB’s full GMRES method [16], IDR(2) [36], and PGMRES [1]. The times required to reach a tolerance of 10^{-10} are displayed in Table 5.1.

5.3. A parallel experiment. While our goal is to present the Schur complement method as a viable alternative to PGMRES and other GMRES-based methods, we also want to show that even for large problems, this approach can be competitive. In particular, here we test the performance of the Schur complement method in comparison to both restarted GMRES and IDR(2), the latter requiring 11 vectors of storage. Recall Example 1 in the paper of Beckermann and Reichel [1], a Bratu path-following test problem:

$$\begin{aligned} -\Delta u(\mathbf{x}) - \lambda \exp(u(\mathbf{x})) &= u(\mathbf{x}), & \mathbf{x} \in \Omega, \\ u(\mathbf{x}) &= 0, & \mathbf{x} \in \partial\Omega, \end{aligned}$$

where Ω is the unit square. We chose the same finite difference discretization and parameters as in the tests in [1], but used a grid of 500×500 points. This gives a matrix of order $n = 25 \times 10^4$ that is Hermitian plus a rank-2 skew-Hermitian modification. For this test, we implemented a parallelized version of the Schur complement method in which the three Hermitian solves are done in parallel using the looping control structure `parfor` from MATLAB’s Parallel Computing Toolbox. In Table 5.2, we see

TABLE 5.1

*MATLAB’s implementation of GMRES, PGMRES, van Gijzen’s implementation of IDR(2), and the Schur complement method are applied to the discretized Lippmann–Schwinger integral operator for various wave numbers, using a discretization of $n = 10^3$ points and a right-hand side vector randomly generated by MATLAB’s `randn` function. The timings are in seconds; a * indicates that the method did not reduce the relative residual norm to 10^{-10} within 1000 iterations.*

κ	GMRES	PGMRES	IDR(2)	SCM
1	0.0227	0.0077	0.0067	0.0071
2	0.0252	0.0102	0.0069	0.0085
3	0.0295	0.0159	0.0091	0.0105
4	0.0346	*	0.0113	0.0138
5	0.0399	*	0.0135	0.0154
10	0.0728	*	0.0263	0.0258
20	0.1597	*	0.0618	0.0559
30	0.2739	*	0.1179	0.0942
40	0.4213	*	0.1948	0.1334
50	0.6001	*	*	0.1829
60	0.8141	*	*	0.2464
70	1.0562	*	*	0.3328
80	1.3394	*	*	0.4109
90	1.6781	*	*	0.5176
100	2.0711	*	*	0.6480

TABLE 5.2

Run time and matrix-vector product count for the parallelized Schur complement method, IDR(2), GMRES(11), and GMRES(100) on the two-dimensional Bratu problem from [1], but discretized on a grid of 500×500 points. The convergence tolerance of the relative residual, $\|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}\|$, was 10^{-7} . All methods were preconditioned with an incomplete Cholesky factorization of the Hermitian part with drop tolerance 10^{-2} .

Method	# matvecs	Run time
Parallel SCM	661	38.911
Serial SCM	661	86.966
IDR(2)	340	57.410
GMRES(11)	*	*
GMRES(100)	337	1061.705

* No convergence in a reasonable amount of time.

that the parallel Schur complement method outperforms the other methods in run time, though it requires more matrix-vector products than any other method. This is a result of being able to perform the inner Hermitian iterations in parallel, with an average of 246 matrix-vector products per processor being performed simultaneously.

Tests of the parallel Schur complement method were run on an Oracle Sun Fire X4600 M2 x64 server with eight AMD dual core Opteron 64-bit processors and 128 gigabytes of RAM. The operating system is the 64-bit version of SuSE Linux 11. All computations were done in the 64-bit version of MATLAB R2010a.

6. Discussion and conclusion. PGMRES cleverly exploits the structure of nearly Hermitian matrices to provide a short-term recurrence that is mathematically equivalent to GMRES. Unfortunately, our experiments illustrate that the method suffers from numerical instabilities that can cause the residual to stagnate, even for mildly conditioned coefficient matrices. The analysis presented in section 3.3, as well as the computational experiments in section 3.2, demonstrate that instabilities occur when the skew-Hermitian part is not small in norm. The instability corresponds to a severe loss of orthogonality of the Arnoldi vectors produced by PGMRES, so severe that the “basis” loses linear independence. This is consistent with the analysis presented in [10] for the standard GMRES algorithm.

As an alternative, we show that solving linear systems with a nearly Hermitian coefficient matrix (1.2) is similar in structure to Schur complement problems. A Schur complement approach, with stopping criteria developed in section 4, computes the solution to the original system by solving $s + 1$ Hermitian linear systems and an $s \times s$ system (provided the Hermitian part is invertible). Since one can solve Hermitian linear systems using MINRES or conjugate gradients, this approach requires only three-term recursions (though poor conditioning of the Hermitian part can hamper convergence). The method is easily parallelizable and simple to implement.

Acknowledgements. We gratefully acknowledge fruitful conversations about this work with Haim Avron, Matthias Heinkenschloss, Marlliny Monsalve, and Valeria Simoncini, and the helpful comments of several referees.

REFERENCES

- [1] B. BECKERMANN AND L. REICHEL, *The Arnoldi process and GMRES for nearly symmetric matrices*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 102–120.
- [2] C. BREZINSKI, *Schur complements and applications in numerical analysis*, in The Schur Complement and Its Applications, F. Zhang, ed., Springer, New York, 2005, pp. 227–258.

- [3] Y. CHEN AND V. ROKHLIN, *On the inverse scattering problem for the Helmholtz equation in one dimension*, Inverse Problems, 8 (1992), pp. 365–391.
- [4] P. CONCUS AND G. H. GOLUB, *A generalized conjugate gradient method for nonsymmetric systems of linear equations*, in Computing Methods in Applied Sciences and Engineering, Part I, Lecture Notes in Econom. and Math. Systems, Springer, Berlin, 134 (1976), pp. 56–65.
- [5] R. W. COTTLE, *Manifestations of the Schur complement*, Linear Algebra Appl., 8 (1974), pp. 189–211.
- [6] T. A. DRISCOLL, K.-C. TOH, AND L. N. TREFETHEN, *From potential theory to matrix iterations in six steps*, SIAM Rev., 40 (1998), pp. 547–578.
- [7] B. FISCHER, *Polynomial Based Iteration Methods for Symmetric Linear Systems*, Wiley–Teubner, Chichester, UK, 1996.
- [8] B. FISCHER, A. RAMAGE, D. J. SILVESTER, AND A. J. WATHEN, *Minimum residual methods for augmented systems*, BIT, 38 (1998), pp. 527–543.
- [9] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, 1997.
- [10] A. GREENBAUM, M. ROZLOŽNÍK, AND Z. STRAKOŠ, *Numerical behaviour of the modified Gram–Schmidt GMRES implementation*, BIT, 37 (1997), pp. 706–719.
- [11] M. H. GUTKNECHT, *Block Krylov space methods for linear systems with multiple right-hand sides: an introduction*, in Modern Mathematical Models, Methods and Algorithms for Real World Systems, A. H. Siddiqi, I. S. Duff, and O. Christensen, eds., Anamaya Publishers, New Delhi, 2007, pp. 420–447.
- [12] W. W. HAGER, *Updating the inverse of a matrix*, SIAM Rev., 31 (1989), pp. 221–239.
- [13] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, J. Research Nat. Bur. Standards, 49 (1952), pp. 409–436.
- [14] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, 2002.
- [15] M. HUHTANEN, *A matrix nearness problem related to iterative methods*, SIAM J. Numer. Anal., 39 (2001), pp. 407–422.
- [16] MATHWORKS, *MATLAB 2010a*, Natick, MA, 2010.
- [17] D. P. O’LEARY, *The block conjugate gradient algorithm and related methods*, Linear Algebra Appl., 29 (1980), pp. 293–322.
- [18] C. C. PAIGE, *A useful form of unitary matrix obtained from any sequence of unit 2-norm n -vectors*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 565–583.
- [19] C. C. PAIGE, M. ROZLOŽNÍK, AND Z. STRAKOŠ, *Modified Gram–Schmidt (MGS), least squares, and backward stability of MGS-GMRES*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 264–284.
- [20] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [21] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, SIAM, Philadelphia, 1997.
- [22] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003.
- [23] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Statist. Comput., 7 (1986), pp. 856–869.
- [24] Y. SAAD, *On the Lanczos method for solving symmetric linear systems with several right-hand sides*, Math. Comp., 48 (1987), pp. 651–662.
- [25] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.
- [26] J. SHERMAN AND W. J. MORRISON, *Adjustment of an inverse matrix corresponding to a change in one element of a given matrix*, Ann. Math. Statistics, 21 (1950), pp. 124–127.
- [27] J. A. SIFUENTES, M. EMBREE, AND R. B. MORGAN, *The Stability of GMRES Convergence, with Application to Preconditioning by Approximate Deflation*, Technical Report TR 11-12, Rice University, Department of Computational and Applied Mathematics, Houston, TX, 2011.
- [28] J. SIFUENTES, *Preconditioned Iterative Methods for Inhomogeneous Acoustic Scattering Applications*, Ph.D. thesis, Rice University, Houston, TX, 2010.
- [29] V. SIMONCINI AND E. GALLOPOULOS, *Convergence properties of block GMRES and matrix polynomials*, Linear Algebra Appl., 247 (1996), pp. 97–119.
- [30] V. SIMONCINI AND D. B. SZYLD, *On the occurrence of superlinear convergence of exact and inexact Krylov subspace methods*, SIAM Rev., 47 (2005), pp. 247–272.
- [31] V. SIMONCINI AND D. B. SZYLD, *Recent computational developments in Krylov subspace methods for linear systems*, Numer. Linear Algebra Appl., 14 (2007), pp. 1–59.
- [32] V. SIMONCINI AND D. B. SZYLD, *Interpreting IDR as a Petrov–Galerkin method.*, SIAM J. Sci. Comput., 32 (2010), pp. 1898–1912.

- [33] P. SONNEVELD AND M. B. VAN GIJZEN, *IDR(s): A family of simple and fast algorithms for solving large nonsymmetric systems of linear equations*, SIAM J. Sci. Comput., 31 (2008), pp. 1035–1062.
- [34] D. B. SZYLD AND O. B. WIDLUND, *Variational analysis of some conjugate gradient methods*, East-West J. Numer. Math., 1 (1993), pp. 51–74.
- [35] R. VANDEBRIL AND G. M. DEL CORSO, *An implicit multishift QR-algorithm for Hermitian plus low rank matrices*, SIAM J. Sci. Comput., 32 (2010), pp. 2190–2212.
- [36] M. B. VAN GIJZEN, *The Induced Dimension Reduction Method*, 2010. Available at <http://ta.twi.tudelft.nl/nw/users/gijzen/IDR.html>.
- [37] O. B. WIDLUND, *A Lanczos method for a class of nonsymmetric systems of linear equations*, SIAM J. Numer. Anal., 15 (1978), pp. 801–812.
- [38] M. A. WOODBURY, *Inverting Modified Matrices*, Statistical Research Group, Memo. Rep. no. 42, Princeton University, Princeton, NJ, 1950.
- [39] E. L. YIP, *A note on the stability of solving a rank- p modification of a linear system by the Sherman–Morrison–Woodbury formula*, SIAM J. Sci. Statist. Comput., 7 (1986), pp. 507–513.