

40 solution \tilde{x} in a given iteration is of the form:

$$41 \quad \tilde{x} = x_0 + p(A)r_0,$$

42 where p is a polynomial, and its related residual is equal to

$$43 \quad (1.2) \quad \tilde{r} = b - A\tilde{x} = (I - Ap(A))r_0 \equiv \rho(A)r_0.$$

44 Note that the approximate solution is a member of the affine Krylov subspace $x_0 +$
 45 $\mathcal{K}_m(A, r_0)$. The acceleration procedures based on Krylov subspace methods that
 46 have been developed in the literature are all based on polynomial iterations where
 47 the iterates are of the form given above, and the polynomials are obtained using
 48 various criteria. For example, the criterion employed in GMRES [26] is to select the
 49 polynomial p to make the residual norm $\|\tilde{r}\|_2$ as small as possible. The Chebyshev
 50 “semi-iterative” method [12, 13] constructs p so that the residual polynomial $\rho(t)$ is an
 51 appropriately shifted and scaled Chebyshev polynomial of the first kind. The residual
 52 polynomial is built so that it is small in an ellipse that encloses the spectrum of the
 53 matrix A . In these methods, the polynomial p can be either used directly to solve
 54 linear systems approximately in an iterative scheme as in [12, 13] and other works, or
 55 it can be exploited as a preconditioner in combination with an acceleration such as
 56 GMRES for example.

57 Polynomial preconditioners are quite appealing because they are simple to use
 58 and because they can be highly effective for some problems. The construction of the
 59 polynomial preconditioner does not involve matrix factorizations and it is also inde-
 60 pendent of reordering schemes. Moreover, applying the preconditioner relies heavily
 61 on one single operation namely the matrix-vector multiplication associated with the
 62 original coefficient matrix A . This operation has been studied and optimized for over
 63 decades by researchers, see, e.g., [2, 32, 3, 18] and is often extremely efficient for
 64 sparse matrices. In addition, the computations are completely free of inner product
 65 which is communication-intensive and limits the performance in a distributed mem-
 66 ory environment. The paper brings three main contributions which are summarized
 67 below:

- 68 • **Improved numerical stability.** In the past, several polynomial precondition-
 69 tioners have been proposed in the literature [24, 22, 21, 11]. However, all of
 70 these methods suffer from numerical stability issues that hampers their use
 71 for higher degrees. In contrast, the proposed methods build a polynomial basis
 72 via an Arnoldi-like procedure. This procedure represents the polynomial
 73 implicitly and has well-controlled numerical stability. As a result, the pro-
 74 posed polynomial preconditioners can be computed accurately for arbitrary
 75 degree.
- 76 • **Guaranteed effect in spectrum.** The proposed polynomial precondition-
 77 ers are constructed by solving a discrete least-squares problem based on the
 78 spectrum of the coefficient matrix so that the spectrum of the preconditioned
 79 system will be better clustered. In contrast, those based on GMRES poly-
 80 nomials cannot guarantee to yield a good preconditioner as pointed out in
 81 [31, 11].
- 82 • **Efficient construction and application.** The proposed polynomial pre-
 83 conditioners are built in a carefully designed polynomial space which has
 84 much smaller dimension compared to the matrix size. As a result, the cost
 85 of building the polynomial is essentially negligible. In the application phase,
 86 a technique based on short-term recurrence is proposed in Section 3.3 which

87 can significantly accelerate the application of the preconditioner on a vector
 88 and reduce the storage requirement.

89 The rest of this paper is organized as follows. Section 2 introduces a few ways
 90 to derive polynomial preconditioners based on solving minimax problems. Section 3
 91 presents an Arnoldi-like procedure to generate a stable polynomial basis based on
 92 the boundary of the spectrum of the coefficient matrix. Several improvements are
 93 discussed in Section 4 and numerical examples are provided in Section 5. Finally,
 94 concluding remarks are draw in Section 6.

95 **2. Polynomial construction via an explicit basis.** In this section, we will
 96 discuss a few ways to derive a polynomial preconditioner when an explicit basis $\{\phi_i(z)\}$
 97 for the polynomial space is given.

98 **2.1. Classical minimax problem.** In many applications, the boundary of the
 99 spectrum of A is not hard to estimate. For example, this can be done either by
 100 analyzing the physical problem [20] where (1.1) is derived from or approximated by
 101 methods such as the Arnoldi iteration [1, 10]. Assume that all eigenvalues of A
 102 are contained in a simply connected domain $\Omega \subset \mathbb{C}$ and denote by $\Gamma = \partial\Omega$ the boundary
 103 of Ω . Here, we further assume that Ω does not contain the origin and that Γ is
 104 piecewise smooth; see Figure 2.1 for an illustration.

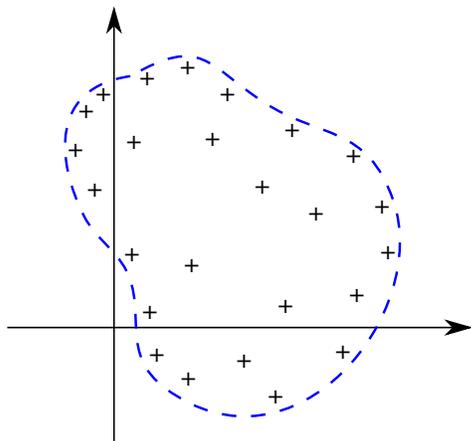


FIG. 2.1. *Eigenvalues of the matrix enclosed by a closed curve.*

105 From (1.2) we have that

$$106 \quad \|\tilde{r}\| \leq \|I - Ap(A)\| \|r_0\|.$$

107 In order to make $\|\tilde{r}\|$ small, we could choose p so that $\|I - Ap(A)\|$ is small. A
 108 straightforward criterion to ensure this is simply to require that $|1 - zp(z)|$ be small
 109 for all $z = \lambda$ where λ is an eigenvalue of A . Unfortunately, this approach involves
 110 all the eigenvalues of A , which is not practically feasible so an alternative is to seek
 111 p so that the maximum of $|1 - zp(z)|$ in the region Ω is small. Since we assume the
 112 eigenvalues of A are enclosed by Γ , and since $1 - zp(z)$ is holomorphic, the maximum
 113 modulus principle [27] tells us that the maximum value of $|1 - zp(z)|$ on Ω is achieved
 114 on the boundary Γ . Thus for a fixed $m > 0$, the sought-after polynomial p can be

115 characterized by the following minimax problem:

$$116 \quad (2.1) \quad \min_{p \in \mathcal{P}_{m-1}} \max_{z \in \Gamma} |1 - zp(z)|,$$

117 where \mathcal{P}_{m-1} denotes the set of all complex polynomials of degree less than m .

118 It is important to note that an approach based on this framework can be viewed as
 119 a heuristic only because in the highly non-normal case the norm of $\|I - Ap(A)\|$ is not
 120 always tightly related to the maximum of $|1 - zp(z)|$ on the contour Γ that contains
 121 the spectrum, see, for example, the articles on the Crouzeix conjecture [6, 7, 8]. For
 122 many practical problems minimizing some norm of $|1 - zp(z)|$ on the contour Γ will
 123 yield good results.

124 Defining the Chebyshev norm on any set $\mathcal{D} \subset \mathbb{C}$ of a function f by $\|f\|_{\mathcal{D}} =$
 125 $\max_{z \in \mathcal{D}} |f(z)|$, the minimax problem (2.1) can be rewritten as

$$126 \quad (2.2) \quad \min_{p \in \mathcal{P}_{m-1}} \|1 - zp(z)\|_{\Gamma}.$$

127 This is a Chebyshev approximation problem in functional form with a domain that is
 128 a continuous subset of the complex plane. The problem can be solved by a Remez-like
 129 algorithm [4, 30, 23] or the Lanczos τ -method [15, 5]. However, when the geometry of
 130 Γ becomes irregular or the degree of the polynomial increases, these methods might
 131 fail. As a result, we will not attempt to solve the minimax problem (2.1) directly.

132 We can instead solve a discrete version of the problem, i.e., we can simplify (2.1)
 133 by replacing the continuous contour Γ by a discrete one. Let $\Gamma_n = \{z_1, z_2, \dots, z_n\}$
 134 be an n -point discretization of the boundary Γ . This discretization should capture
 135 the geometric characteristics of Γ , a uniform discretization of Γ usually suffices in
 136 practice. In certain cases when Γ contains a high curvature or discontinuous part,
 137 we can either add additional points to refine the discretization in this area or simply
 138 replace this part by a smoother curve before the discretization. We then consider
 139 the Chebyshev norm on the discrete set Γ_n and define the following *discrete minimax*
 140 *problem*:

$$141 \quad (2.3) \quad \min_{p \in \mathcal{P}_{m-1}} \|1 - zp(z)\|_{\Gamma_n}.$$

142 Write $p(z) = \sum_{i=1}^m \alpha_i \phi_i(z)$ and denoted by $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_m]^T \in \mathbb{C}^m$ the
 143 column vector of all the coefficients, (2.3) becomes

$$144 \quad \min_{\alpha \in \mathbb{C}^m} \max_{1 \leq i \leq n} \left| 1 - z_i \sum_{j=1}^m \alpha_j \phi_j(z_i) \right|.$$

145 Define an $n \times m$ matrix F with entries given by

$$146 \quad f_{ij} = z_i \phi_j(z_i), \quad 1 \leq i \leq n, \quad 1 \leq j \leq m,$$

147 and $e \in \mathbb{C}^n$ the column vector of all ones, (2.3) can be reformulated in the matrix
 148 form as

$$149 \quad \min_{\alpha \in \mathbb{C}^m} \|e - F\alpha\|_{\infty}.$$

150 We refer the readers to [29, 28, 33, 16] for some discussions on algorithms for
 151 solving the above complex linear programming problem. This problem uses the in-
 152 finity norm in \mathbb{C}^n . We will not consider this approach in the remainder of the paper.

153 Instead we will replace the infinity norm by the 2-norm in \mathbb{C}^n . The least-squares
 154 polynomial will be computed by a GMRES-like procedure in polynomial space which
 155 is described next.

156 **3. Polynomial construction via an Arnoldi process.** Define an inner prod-
 157 uct for the polynomial space as

$$158 \quad (3.1) \quad \langle p_1, p_2 \rangle = \sum_{i=1}^n p_1(z_i) \overline{p_2(z_i)}.$$

159 This sesqui-linear form is a valid inner product of the space of polynomials \mathcal{P}_m as long
 160 as m does not exceed the number of points n . We will denote by $\|\cdot\|_\omega$ the related
 161 norm. Then we would like to solve the following discrete least-squares problem instead
 162 of (2.2) :

$$163 \quad (3.2) \quad \min_{p \in \mathcal{P}_{m-1}} \|1 - zp(z)\|_\omega^2.$$

164 Instead of specifying a basis $\{\phi_i(z)\}_{i=1}^m$ in advance as in Section 2, we will actually
 165 build the polynomial basis dynamically in an Arnoldi-like a process.

166 **3.1. GMRES in polynomial space.** The construction procedure for the opti-
 167 mal polynomial is similar to GMRES in vector space and is described in Algorithm 3.1.
 168 For the sake of conformity with the notation used in the standard Arnoldi process,
 169 the polynomial basis of degree l is represented by q_{l+1} , instead of q_l .

170 It is easy to see that the Arnoldi-like process Algorithm 3.1 will indeed generate
 171 a set of orthonormal polynomial basis $\{q_i\}_{i=1}^m$ with respect to the inner product (3.1),
 172 there will be no stability issue even for high degrees due to the full orthogonalization.
 173 The question now is how to represent the polynomials and how to carry out the actual
 174 computations that are involved in Algorithm 3.1. In fact we have a number of choices
 175 of which we will only retain one. The simplest choice, a poor one for obvious reasons
 176 of stability, is to use the power series representation. In this case, a polynomial $p(z) =$
 177 $\alpha_0 + \alpha_1 z + \dots + \alpha_{m-1} z^{m-1}$ will be represented by the vector $[\alpha_0, \alpha_1, \dots, \alpha_{m-1}]^T \in \mathbb{C}^m$.
 178 For example, the polynomial multiplication $q := zq_j$ in Step 3 amounts to shifting
 179 all components of the representing vector down by one position and putting a zero
 180 in the first position; addition, subtraction and scalar multiplication all translate to
 181 the corresponding operation on the vector; inner products are also easy to compute
 182 efficiently once the Gram matrix of the power series basis is computed.

183 However, we will not use any explicit representations because, as we will show
 184 later, we are more interested in the coefficients h_{ij} than the polynomials themselves.
 185 Therefore we will represent the polynomials implicitly by the evaluations on the points
 186 $\{z_i\}_{i=1}^n$, i.e., a polynomial p is represented by a vector $[p(z_1), p(z_2), \dots, p(z_n)]^T \in \mathbb{C}^n$.
 187 Under this representation, the polynomial multiplication $q := zq_j$ in Step 3 will be
 188 translated simply into the entry-wise multiplication of two vectors of length n , and
 189 the inner products in Steps 5 and 8 become standard inner products in vector space
 190 \mathbb{C}^n .

191 We now address the solution of the discrete least-squares problem (3.2). Define
 192 the $(m+1) \times m$ matrix H_m where $(H_m)_{ij} = h_{ij}$, for $i \leq j+1$ and $(H_m)_{ij} = 0$,
 193 for $i > j+1$, so that H_m is an upper-Hessenberg matrix. If we abuse the notation
 194 and replace all polynomials by their vector representations in Algorithm 3.1, then the
 195 constant 1 in (3.2) becomes βq_1 where $\beta = \|\mathbb{1}\|_\omega = \sqrt{n}$. Define $Q_l = [q_1, q_2, \dots, q_l]$
 196 to be the column concatenation of the first l basis vectors, then each Q_l , for all

Algorithm 3.1 *The Arnoldi-like process in polynomial space**Input:* Discretization points $\{z_i\}_{i=1}^n$ on Γ and degree m *Output:* Orthogonal polynomial basis $\{q_i\}_{i=1}^{m+1}$

-
- 1: Set $q_1 = \mathbb{1}/\|\mathbb{1}\|_\omega$ $\triangleright q_1$ is of degree 0 and norm 1
 - 2: **for** $j = 1, 2, \dots, m$ **do**
 - 3: Compute $q := zq_j$ \triangleright Increase degree
 - 4: **for** $i = 1, 2, \dots, j$ **do**
 - 5: Compute $h_{ij} = \langle q, q_i \rangle$
 - 6: Compute $q = q - h_{ij}q_i$
 - 7: **end for** \triangleright Full orthogonalization
 - 8: Compute $h_{j+1,j} = \|q\|_\omega$
 - 9: Compute $q_{j+1} = q/h_{j+1,j}$ \triangleright Normalize the new basis
 - 10: **end for**
-

197 $1 \leq l \leq m + 1$, is unitary. If p is expressed linearly in the basis $\{q_1, q_2, \dots, q_m\}$ as

198
$$p = \sum_{i=1}^m \alpha_i q_i = Q_m \alpha$$

199 where $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_m]^T$, then the polynomial zp in (3.2) becomes

200
$$\begin{aligned} zp &= \sum_{j=1}^m \alpha_j (zq_j) = \sum_{j=1}^m \alpha_j \sum_{i=1}^{j+1} h_{ij} q_i = \sum_{j=1}^m \alpha_j \sum_{i=1}^{m+1} h_{ij} q_i \quad (h_{ij} = 0 \text{ when } i - j > 1) \\ &= \sum_{i=1}^{m+1} \sum_{j=1}^m q_i h_{ij} \alpha_j = Q_{m+1} H_m \alpha. \end{aligned}$$

201 In the end we observe that solving (3.2) amounts to minimizing with respect to $\alpha \in \mathbb{C}^n$
 202 the objective function

203
$$J(\alpha) = \|\beta q_1 - Q_{m+1} H_m \alpha\|_2^2.$$

204 Since Q_{m+1} is unitary and $q_1 = Q_{m+1} e_1$ where $e_1 = [1, 0, \dots, 0]^T$ is a vector of length
 205 $m + 1$, this can be further reduced to

206 (3.3)
$$J(\alpha) = \|\beta e_1 - H_m \alpha\|_2^2.$$

207 Note that this is a standard least-squares problem, which is identical to the one
 208 encountered in the GMRES process.

209 Once α is found from (3.3), we obtain a polynomial p of degree $m - 1$ and the
 210 matrix M defined by $M^{-1} = p(A)$ can be used as a preconditioner for solving the
 211 linear system $Ax = b$.

212 To apply M^{-1} to a vector v , note that

213 (3.4)
$$M^{-1}v = p(A)v = \sum_{i=1}^m \alpha_i q_i(A)v := \sum_{i=1}^m \alpha_i v_i$$

214 where we define $v_i \equiv q_i(A)v$ for $1 \leq i \leq m$. Since $q_1 = \mathbb{1}/\sqrt{n}$ so

215 (3.5)
$$v_1 = q_1(A)v = Iv/\sqrt{n} = v/\sqrt{n}.$$

216 From the Arnoldi-like process [Algorithm 3.1](#) we have that $zq_i = \sum_{j=1}^{i+1} h_{ji}q_j$, thus

$$217 \quad Av_i = Aq_i(A)v = \left[\sum_{j=1}^{i+1} h_{ji}q_j(A) \right] v = h_{i+1,i}v_{i+1} + \sum_{j=1}^i h_{ji}v_j, \quad 1 \leq i \leq m-1,$$

218 and hence

$$219 \quad (3.6) \quad v_{i+1} = \frac{1}{h_{i+1,i}} \left(Av_i - \sum_{j=1}^i h_{ji}v_j \right), \quad 1 \leq i \leq m-1.$$

220 The v_i 's can be computed recursively from [\(3.5\)](#) and [\(3.6\)](#) and the final result is just a
 221 linear combination of v_i 's with the coefficient α . Note that the only information needed
 222 is the pre-calculated entries available in H_m , the basis Q_{m+1} is not involved directly
 223 in the least-squares problem [\(3.3\)](#) for finding α or in applying the preconditioner
 224 [\(3.4\)–\(3.6\)](#).

225 We note that the idea of using an Arnoldi-like procedure to generate orthogonal
 226 polynomials is not completely new. In fact, the framework is similar to what was
 227 discussed in [\[24\]](#) where the Chebyshev polynomial basis is used to construct a poly-
 228 nomial p that minimizes the residual polynomial $1 - zp$ under some specially defined
 229 norm. But, as mentioned in [\[24\]](#), this algorithm suffers from numerical stability issues
 230 and the polynomial degree has to be kept low. The main reason is that the two norms
 231 used in [\[24\]](#) are completely different, namely, the norm used to form the Chebyshev
 232 polynomial basis and the one used to characterize the solution are not compatible. As
 233 a result, the orthogonal polynomial basis is no longer orthogonal in the inner product
 234 space associated with the optimization problem that generates the solution. The same
 235 argument holds true for other methods that try to construct a polynomial from the
 236 span of a given basis. Since the algorithm proposed in this manuscript uses only the
 237 inner product [\(3.1\)](#) and implicitly constructs the polynomial p , p will be accurately
 238 computed for high degrees (as long as $m < n$). Another class of method construct the
 239 polynomial by finding all of its roots and represents the polynomial by the product
 240 of a series of degree one polynomials, e.g., in [\[21, 11\]](#). These methods also suffer from
 241 stability issues when the degree is high mainly due to numerical cancellation.

242 **3.2. Connection to GMRES.** In comparing the proposed approach to the
 243 standard GMRES approach, one can observe that [\(3.3\)](#) is exactly the same least-
 244 squares problem that we solve in standard GMRES except that the coefficients of H_m
 245 are generated in a vector space of dimension n , the number of points on the contour.
 246 Looking more carefully at the algorithm, it is also possible to show that in fact *it*
 247 *is equivalent to the standard GMRES algorithm applied to the diagonal matrix whose*
 248 *entries are the discretization points z_1, z_2, \dots, z_n .* They are equivalent in the sense
 249 that they would generate the same Hessenberg matrix H_m and in the end also the
 250 same polynomial. For this reason we may refer to this approach as a *proxy-GMRES*
 251 algorithm since that the original matrix is replaced by a small (“proxy”) diagonal
 252 matrix whose spectrum captures the original spectrum well.

253 **3.3. Short-term recurrence.** Because the matrix H_m in [Algorithm 3.1](#) is an
 254 upper Hessenberg matrix, computing $p(A)v$ for a degree $m-1$ polynomial p costs
 255 $\mathcal{O}(m^2N)$ operations and requires $\mathcal{O}(mN)$ storage. This implies that despite the good
 256 numerical stability of the algorithm, its computation cost and storage quickly become
 257 unacceptably high as m increases. Motivated by the three-term recurrence for Cheby-

258 shev polynomials, we will show in this section that a short-term recurrence can be
 259 exploited to significantly reduce these costs.

260 The basic idea is to replace the full orthogonalization in Steps 4 to 7 in Algo-
 261 rithm 3.1 by a partial orthogonalization. That is, the newly generated polynomial q
 262 in Step 3 is only orthogonalized against the most recent k basis, which leads to the
 263 following short-term recurrence relation

$$264 \quad t_{j+1,j}\hat{q}_{j+1} = z\hat{q}_j - \sum_{i=j-k+1}^j t_{ij}\hat{q}_i, \quad 1 \leq j \leq m,$$

265 where t_{ij} ($1 \leq i \leq j$), $t_{j+1,j}$ and \hat{q}_{j+1} are generated in the same way as in Steps 5, 8
 266 and 9 in Algorithm 3.1, respectively. The computed basis $\{\hat{q}_i\}_{i=1}^{m+1}$ form the columns
 267 of \hat{Q}_{m+1} and t_{ij} 's form an $(m+1) \times m$ matrix T_m . Notice that \hat{Q}_{m+1} is not unitary
 268 anymore and T_m is a banded matrix with one subdiagonal and $k-1$ superdiagonals.
 269 For example, when $k=2$, we have the three-term recurrence for the computed basis
 270 \hat{q}_i and T_m is a tridiagonal matrix. This is similar to the Chebyshev polynomial case.
 271 In the extreme case when $k=m$, the partial reorthogonalization becomes equivalent
 272 to the full orthogonalization and all results in Section 3.1 are recovered.

273 Similar to Section 3.1, with the new basis $\{\hat{q}_j\}_{j=1}^{m+1}$ from the short-term recurrence
 274 we can rewrite (3.2) into minimizing with respect to $\hat{\alpha} \in \mathbb{C}^n$ a new objective function

$$275 \quad (3.7) \quad \hat{J}(\hat{\alpha}) = \|\beta\hat{q}_1 - \hat{Q}_{m+1}T_m\hat{\alpha}\|_2^2.$$

276 Solving for $\hat{\alpha}$ in this problem typically needs to compute an orthogonal factorization
 277 of the matrix $\hat{Q}_{m+1}T_m$, which requires some additional computation cost and storage
 278 (since both \hat{Q}_{m+1} and T_m need to be stored) compared to computing α in (3.3).
 279 However, recall that all these computations are still within a vector space of dimension
 280 n which is typically much smaller than N .

281 On the other hand, applying the preconditioner $M^{-1} = \hat{p}(A)$ where \hat{p} is repre-
 282 sented in the new basis \hat{Q}_m is slightly different. More specifically, (3.6) is replaced by
 283 the corresponding short-term version

$$284 \quad (3.8) \quad v_{i+1} = \frac{1}{t_{i+1,i}} \left(Av_i - \sum_{j=i-k+1}^i t_{ij}v_j \right).$$

285 Due to the above short-term recurrence, the application of the preconditioner $M^{-1} =$
 286 $\hat{p}(A)$ on a vector only requires $\mathcal{O}(mkN)$ operations and $\mathcal{O}(kN)$ storage.

287 Now we discuss the stability issue associated with this approach. In exact arith-
 288 metics, it is easy to see that (3.3) and (3.7) are equivalent and the polynomials
 289 $p = Q_m\alpha$ and $\hat{p} = \hat{Q}_m\hat{\alpha}$ obtained from both orthogonalization schemes are exactly
 290 the same. This is because enforcing a short-term recurrence is equivalent to a change
 291 of basis and an update to the corresponding coefficients. However, in floating point
 292 arithmetics, \hat{Q}_m becomes increasingly ill-conditioned when m increases and thus the
 293 coefficient $\hat{\alpha}$ becomes increasingly hard to compute accurately.

294 Next we study the relation between the conditioning of the basis matrix \hat{Q}_m and
 295 the number of recurrence terms k . Figure 3.1 shows how the 2-norm condition number
 296 $\kappa_2(\hat{Q}_m)$ grows for multiple values of k where Γ and Γ_n are draw from the numerical
 297 example in Section 5.2. In Figure 3.1a, the condition number plots for $k=2$ and 3
 298 almost coincide, the relative error between them is shown in Figure 3.1b; similarly for

299 $k = 4, 5$ and $k = 6, 7$. When the recurrence is too short, i.e., $k = 2$ or 3 , the condition
 300 number rapidly grows beyond 10^{12} when m passes 60. In that case, the polynomial
 301 \hat{p} solved with this basis becomes inaccurate and the resulting preconditioner may
 302 become useless. By increasing k to 4 or 5, $\kappa_2(\hat{Q}_m)$ quickly drops from 10^{12} to about
 303 10^3 at $m = 60$ and the basis becomes too ill-conditioned again only when m reaches
 304 180. [Figure 3.1a](#) also shows that the numerical stability keeps getting improved when
 305 k increases to 6.

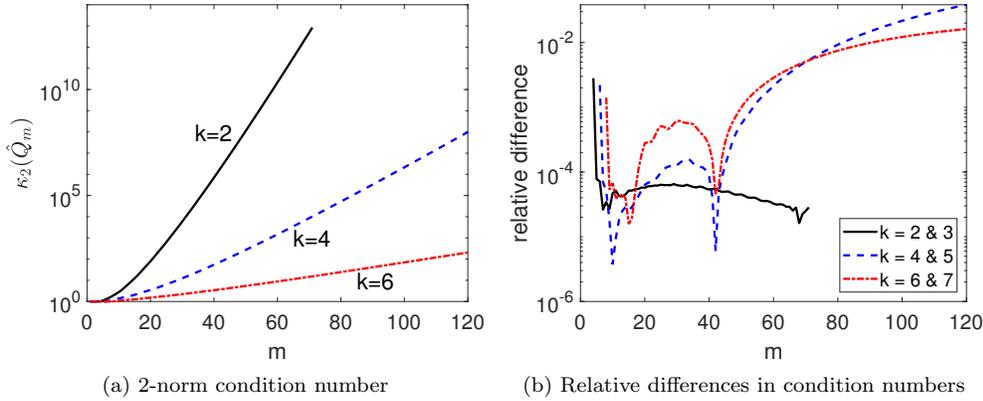


FIG. 3.1. Conditioning of \hat{Q}_m generated with k -term recurrence.

306 The numerical stability of \hat{Q}_m can be monitored inexpensively by its associated
 307 Gram matrix. Denote by \hat{G}_m the $m \times m$ Gram matrix of the basis \hat{Q}_m whose entries
 308 are defined by

309
$$\hat{g}_{ij} = \langle \hat{q}_i, \hat{q}_j \rangle, \quad 1 \leq i, j \leq m$$

310 where the inner product is as defined in (3.1). The matrix \hat{G}_m is Hermitian positive
 311 definite. Let $\hat{G}_m = \hat{L}_m \hat{L}_m^H$ be the Cholesky factorization where \hat{L}_m is lower-triangular
 312 and note that $\kappa_2(\hat{Q}_m) = \sqrt{\kappa_2(\hat{G}_m)} = \kappa_2(\hat{L}_m)$. As the Arnoldi-like process proceeds,
 313 both the Gram matrix \hat{G}_{m+1} and the Cholesky factor \hat{L}_{m+1} can be quickly updated
 314 with \hat{G}_m and \hat{L}_m from the previous step. When a high degree polynomial needs
 315 to be used, the ill-conditioning of \hat{Q}_m can be quickly detected by keeping track of
 316 the condition number of the Cholesky factor \hat{L}_m . Whenever $\kappa_2(\hat{L}_m)$ goes beyond
 317 a certain tolerance, we can stop the process and accept the resulting polynomial
 318 obtained at that point or restart the same process with a longer recurrence relation.
 319 As is indicated in [Figure 3.1](#), we can start from $k = 2$ and increase k by 2 every time
 320 the process restarts. This process is repeated until the desired degree can be reached
 321 while $\kappa_2(\hat{L}_m)$ still remains below the given tolerance. In practice, we find that setting
 322 the tolerance at $\tau = 10^{12}$ usually yields good quality results.

323 **4. Some improvements based on compounding preconditioners.** In the
 324 previous sections, we always assume Γ will exclude the origin. This is because other-
 325 wise, the maximum modulus of the residual polynomial $1 - zp(z)$ will be no less than
 326 1 on Γ and the resulting polynomial preconditioner will not be effective at all. For
 327 ill-conditioned problems, a few eigenvalues of A will stay in a small neighborhood of
 328 the origin and Γ has to be very close to the origin. In that case, a very high degree

329 polynomial becomes mandatory in order to keep the maximum value of $|1 - zp(z)|$
 330 on Γ strictly less than 1. On the other hand, the increased degree will require a
 331 longer recurrence and thus harm the efficiency of the proposed preconditioner. In this
 332 section, we will discuss two compounding techniques to overcome this difficulty.

333 **4.1. Compounding two polynomials.** The first approach is based on com-
 334 pounding two low degree polynomials to mimic the effect of a high degree polynomial.
 335 By doing this, even though the total number of matrix-vector multiplications associ-
 336 ated with A might be slightly increased, the costs associated with vector operations
 337 and storage can be significantly reduced. Also as pointed out in [11], this strategy
 338 can also reduce the number of inner products performed.

339 This can be understood from a simple example. Suppose one high degree polyno-
 340 mial has degree $m - 1$ and the other two low-degree polynomials have degree $m_1 - 1$
 341 and $m_2 - 1$, respectively, with $m = m_1 \times m_2$. For these three polynomials, a k_i -term
 342 recurrence is deployed for a polynomial of degree $m_i - 1$ (for $i = 1, 2$) while a k -term
 343 recurrence is needed for degree $m - 1$. Since m_1 and m_2 are both much smaller
 344 than m , we have $k_1, k_2 \ll k$ when ensuring the numerical stability of the computed
 345 polynomial basis. Thus, applying the preconditioner resulting from compounding the
 346 polynomials will entail fewer vector operations and storage.

347 We now provide more details on how to construct these two low-degree polynomi-
 348 als. First find a contour Γ that encloses all the eigenvalues of A and discretize it as Γ_{n_1} .
 349 Based on Γ_{n_1} , construct the first polynomial p_1 of degree $m_1 - 1$ and select the recur-
 350 rence length k_1 with the procedure discussed in Section 3.3. It can be expected that
 351 most of the eigenvalues of the preconditioned matrix $A_1 := M_1^{-1}A = p_1(A)A$ would
 352 be clustered around $z = 1$. Therefore, a second contour is then selected as a circle \mathcal{C}
 353 centered at $z = 1$ with radius $\theta \in (0, 1)$. Let \mathcal{C}_{n_2} be an n_2 -point discretization of \mathcal{C} and
 354 apply \mathcal{C}_{n_2} to compute the second polynomial p_2 with degree $m_2 - 1$ and recurrence
 355 length k_2 . In the end, the compound polynomial has the form $p(z) := p_1(z)p_2(zp_1(z))$
 356 and the resulting preconditioner is $M^{-1} := p(A) = p_1(A)p_2(Ap_1(A))$.

357 It is clear that the preconditioner M^{-1} is a polynomial in A of degree $m_1m_2 - 1$.
 358 Applying M^{-1} on a vector involves two main operations:

- 359 1. Apply $p_1(A)$ to a vector, which follows the formula (3.4), (3.5), and (3.8).
 360 This computation costs $m_1 - 1$ matvecs associated with A and $\mathcal{O}(m_1k_1N)$
 361 from vector operations and $\mathcal{O}(k_1N)$ storage;
- 362 2. Apply $p_2(Ap_1(A))$ to a vector. This operation consists of $m_2 - 1$ matvecs
 363 of $Ap_1(A)$, $\mathcal{O}(m_2k_2N)$ extra costs from vector operations and $\mathcal{O}(k_2N)$ extra
 364 storage.

365 Table 4.1 compares the costs of applying one high degree polynomial preconditioner
 366 verse one compound polynomial preconditioner. It is easy to see that when $m =$
 367 $m_1 \times m_2$, even though both preconditioners perform the same number of *matvecs*
 368 associated with A , the operations and peak storage associated with the compound
 369 polynomial preconditioner can be much less due to the fact that $k_1, k_2 \ll k$.

370 **4.2. Compounding with other preconditioners.** A second approach is to
 371 compound the polynomial preconditioner with other types of preconditioners. For
 372 ill-conditioned problems, it is suggested to perform an approximate factorization on
 373 $A + \sigma I$ for some complex shift σ [35] instead of the original coefficient matrix A . To
 374 simplify the discussion, we assume the incomplete LU factorization (ILU) is explored
 375 here:

TABLE 4.1

The cost and storage of applying the single and compound polynomial preconditioners, the single polynomial is of degree $m - 1$, the compound polynomial is built with two low degree polynomials of degree $m_1 - 1$ and $m_2 - 1$.

	single polynomial	compound polynomial
matvec of A	$m - 1$	$m_1 m_2 - 1$
vector operations	$\mathcal{O}(mkN)$	$\mathcal{O}(m_1 m_2 k_1 N)$
peak storage	$\mathcal{O}(kN)$	$\mathcal{O}((k_1 + k_2)N)$

376
$$A + \sigma I \approx M_1 = LU.$$

377 We will now discuss two different ways to introduce a second level preconditioner.

378 The first option is to compound M_1^{-1} with a polynomial preconditioner of the
 379 form $p(A)$. Define the distance matrix E as $E = I - M_1^{-1}Ap(A)$. An ideal polyno-
 380 mial p should minimize $\|E\|_2$. Although its optimal solution is hard to calculate, an
 381 approximate solution can be computed inexpensively by randomized sampling (see
 382 [14] for details) as follows: first construct a set of polynomial basis $\{q_1, q_2, \dots, q_m\}$
 383 and express the polynomial as a linear combination of the basis $p = \sum_{j=1}^m \alpha_j q_j$,
 384 then pick l random vectors $\omega_1, \omega_2, \dots, \omega_l$ of length N and solve the coefficients $\alpha =$
 385 $[\alpha_1, \alpha_2, \dots, \alpha_m]^T \in \mathbb{C}^m$ from the following problem

386
$$\min_{\alpha \in \mathbb{C}^m} \max_{1 \leq j \leq l} \|E\omega_j\|_2.$$

387 The main drawback of this approach is still numerical stability since it requires a pre-
 388 specified polynomial basis. No suitable norm like (3.1) can be defined in this case,
 389 thus there is no reliable way to generate a good basis set like in Section 3.1.

390 The second option resolves this issue by considering a new linear system

391
$$M_1^{-1}Ax = M_1^{-1}b$$

392 and applying the procedures discussed in Section 3.1 to the new coefficient matrix
 393 $A_1 := M_1^{-1}A$. Note here that the contour Γ for A_1 can be estimated by running a
 394 few steps of the Arnoldi process. After the polynomial p is constructed, the com-
 395 pound preconditioner takes the form of $M^{-1} = p(M_1^{-1}A)$. Suppose the polynomial
 396 p is of degree $m - 1$, then one application of M^{-1} on a vector consists of $m - 1$
 397 matvecs associated with A and $m - 1$ applications of the preconditioner M_1^{-1} . As
 398 mentioned in Section 1, ILU type preconditioners may become the performance bot-
 399 tleneck in a parallel environment due to the sequential nature of the triangular solves.
 400 This framework naturally allows replacing the ILU factorization with more scalable
 401 preconditioners such as SLR, MSLR or GMSLR preconditioners [19, 34, 9].

402 **5. Numerical experiments.** All numerical tests were ran in Matlab on a Desk-
 403 top PC with 3.80 GHz CPU and 8 GB memory. We used flexible GMRES (FGMRES)
 404 with a restart dimension of 50 as the accelerator, the initial guess was set to be the
 405 zero vector and the process was terminated when either the residual was reduced by
 406 a prescribed factor τ or the total number of inner iterations reached 1000.

407 The following notation is used in this section:

- 408 • p-t: the computation time in seconds for constructing the preconditioner,
 409 which includes the time for estimating/discretizing the contour Γ , building

410 the polynomial or/and other preconditioners depending on the specific tests.
 411 F indicates the preconditioner cannot be constructed;
 412 • i-t: iteration time in seconds for FGMRES(50) to converge;
 413 • its: total number of iterations required for FGMRES(50) to converge, F in-
 414 dicates FGMRES(50) does not converge within 1000 inner iterations;
 415 • mv: number of matvecs associated with A performed.

416 **5.1. A diagonal matrix.** In the first example, we generated a 2000×2000
 417 diagonal matrix where all the diagonal entries (eigenvalues) were randomly chosen
 418 from the semiannular region $\Omega = \{z \in \mathbb{C} \mid 0.8 \leq |z| \leq 2, 0 \leq \text{Arg}(z) \leq \pi\}$. The
 419 boundary of this region is shown in Figure 5.1 where the squares are the approximate
 420 eigenvalues (Ritz values) computed by running 60 steps of the Arnoldi algorithm. An
 421 approximate boundary was obtained by running Matlab’s built-in function `boundary`
 422 on the approximate eigenvalues. Figure 5.1 shows that the Ritz values from the
 423 Arnoldi algorithm can characterize the boundary of the spectrum.

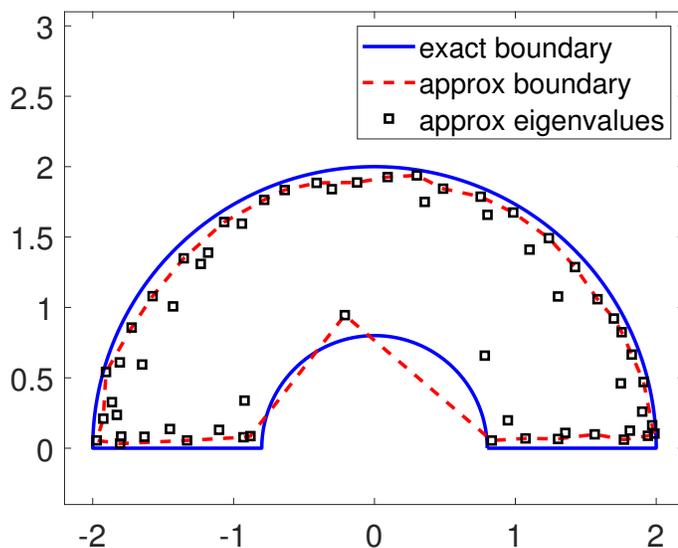


FIG. 5.1. The exact and approximate boundaries of the spectrum and the approximate eigenvalues obtained from 60 steps of the Arnoldi algorithm for the $2,000 \times 2,000$ diagonal matrix in Section 5.1.

424 We first constructed polynomials of degree 29 ($m = 30$) with a recurrence length
 425 $k = 2$. Figure 5.2 shows the contour maps for the function $|1 - zp(z)|$ in log scale
 426 based on both exact (left) and approximate (right) boundaries. Since the estimated
 427 boundary approximates the exact one very well, the two maps look almost identical.
 428 Table 5.1 tabulates the numerical results for solving the linear system with these
 429 constructed polynomial preconditioners, the tolerance was set at $\tau = 10^{-12}$. It took
 430 237 iterations for FGMRES(50) to converge without any preconditioner. On the other
 431 hand, FGMRES(50) with the polynomial preconditioners converged in 8 iterations in
 432 both cases. Although the preconditioned methods performed 3 more matvecs, they
 433 actually took much less time to converge. Similar observations can also be made in
 434 other examples in this section. This is due to the fact that a reduced iteration number
 435 leads to a much smaller subspace for FGMRES and far fewer inner products during
 436 the computation. This performance gap can be expected to become more pronounced

437 when running the experiments on high performance computing architectures.

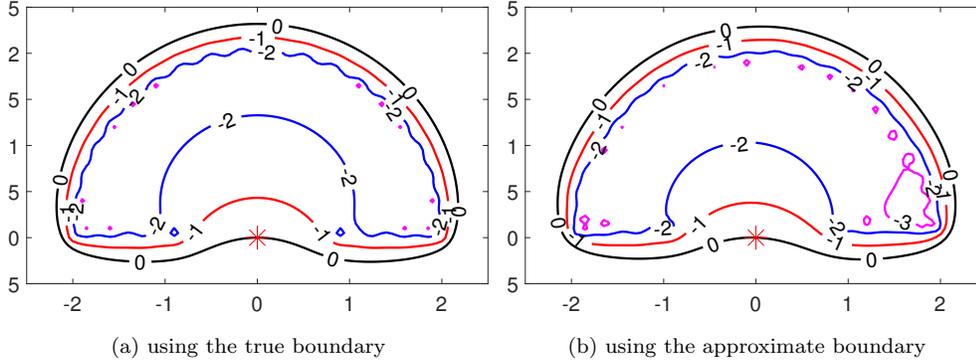


FIG. 5.2. Contour maps of $|1-zp(z)|$ in log scale with different choice of Γ for the $2,000 \times 2,000$ diagonal matrix in Section 5.1. The asterisk marks the origin.

438 We also want to emphasize that the preconditioner construction time was only a
 439 tiny fraction of the iteration time. This is because we used 400 discretization points
 440 for both the exact and approximate boundaries and the corresponding polynomial
 441 space has much smaller dimension compared to the matrix size $N = 2000$.

TABLE 5.1
 Convergence results of FMGRES(50) for the $2,000 \times 2,000$ diagonal matrix test in Section 5.1 with tolerance $\tau = 10^{-12}$.

		p-t	i-t	its	mv
no precondition.		\	0.81	237	237
with precondition.	exact boundary	0.0062	0.61	8	240
	approx. boundary	0.0056	0.61	8	240

442 **5.2. Helmholtz problem.** The second example is the 3D Helmholtz equation

443
$$-\Delta u - \frac{\omega^2}{c^2(x)} u = s$$

444 where ω is the angular frequency and $c(x)$ is the wavespeed. The computational do-
 445 main was the unit cube and the equation was discretized with 7-point stencil finite
 446 difference method. PML boundary conditions were imposed to reduce the artificial re-
 447 flections near the boundaries of the computational domain. We kept 8 grid points per
 448 wavelength. The resulting linear system is sparse complex symmetric with dimension
 449 $N = N_x \times N_y \times N_z$. Moreover, the spectrum of the matrix is contained in a rectangle
 450 area $\{z \in \mathbb{C} \mid \text{real}(z) \in [-1, \rho_1 - 1], \text{imag}(z) \in [-\rho_2, 0]\}$ where the two parameters ρ_1
 451 and ρ_2 are given by [20, Lemma 3.1]. Figure 5.3a shows all the eigenvalues and the
 452 rectangular boundary from [20, Lemma 3.1] for a discretized Helmholtz operator of
 453 size $N = 20^3$. We fixed the tolerance at $\tau = 10^{-3}$ for the Helmholtz equation tests in
 454 this section. Note that this test matrix is non-normal.

455 **5.2.1. Compounding polynomial preconditioners.** Compared with the first
 456 test example, this problem is much harder to solve. First, there are many eigenvalues

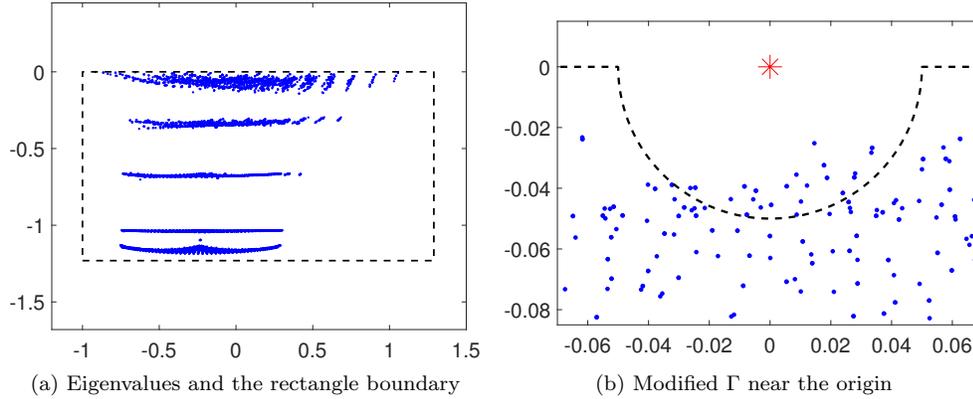


FIG. 5.3. The theoretical rectangular spectrum boundary from [20, Lemma 3.1] and the zoom-in view of the modified Γ near the origin for a discretized Helmholtz operator of size $N = 20^3$.

close to the origin. Second, the theoretical spectrum boundary (the rectangular area) overlaps with the origin. In order to construct an effective polynomial preconditioner, we have to

1. Modify the theoretical rectangular contour [20, Lemma 3.1] to exclude the origin.
2. Use a high degree polynomial.

For this test, the problem size was $N = 100^3 = 1,000,000$ and the boundary Γ was a modified rectangle with the origin excluded. A zoom-in view of Γ near the origin is shown in Figure 5.3b. The boundary Γ was then discretized uniformly on each of its continuous parts with a step size $h = 0.002$ which results to a total of 3551 discretization points. We compared the performance of two polynomial preconditioners on this test matrix. The first one was a single polynomial of degree 599 ($m = 600$). In order to ensure its numerical stability, we used a recurrence length $k = 10$. The second one was a compound polynomial with $m_1 = 60$ and $m_2 = 10$ so that $m_1 \times m_2 = m$. Since m_1 and m_2 are relatively small, the recurrence lengths were set to be $k_1 = k_2 = 2$. The convergence results with these two preconditioners are shown in Table 5.2. In addition, we also ran ILUT-preconditioned FGMRES(50) and CG on the corresponding normal equation with a block Jacobi preconditioner for comparisons.

TABLE 5.2

Convergence results of various preconditioned FGMRES(50) on the Helmholtz equation test with size $N = 100^3$, the tolerance is fixed at $\tau = 10^{-3}$.

Preconditioner type	p-t	i-t	its
no preconditioner	\	\	F
ILUT	F	\	\
CG with diagonal preconditioner	0.49	\	F
single polynomial of degree 600 – 1	5.85	2554.44	16
compound polynomial of degree 60 × 10 – 1	0.08	853.11	18

Due to the ill-conditioning and indefiniteness of the test matrix, the first three methods in Table 5.2 failed to converge. In particular, ILUT even failed to finish

478 the factorization. On the other hand, both polynomial preconditioned methods con-
 479 verged within 18 iterations. Compared to the single polynomial preconditioner, the
 480 compound polynomial preconditioner took much less time to construct and reduced
 481 the iteration time by more than a half even though 1200 more matvecs of A were
 482 performed.

483 The residual histories are plotted in Figure 5.4. It is easy to see that both polyno-
 484 mial preconditioned FGMRES methods converged quickly without encountering any
 485 stagnation while the non-preconditioned FGMRES converged slowly.

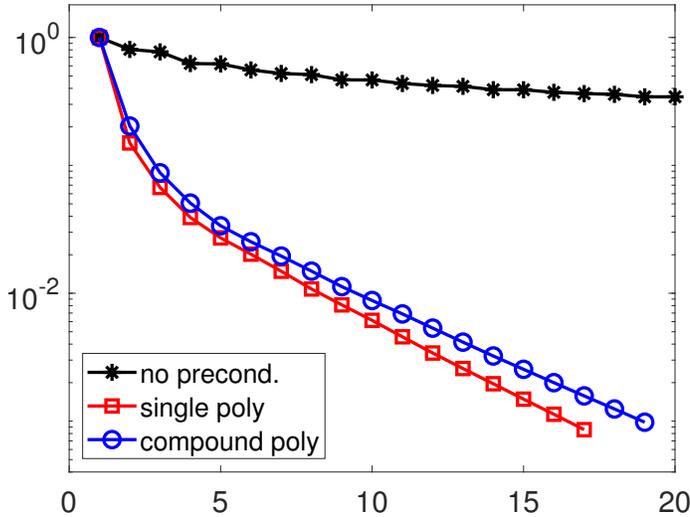


FIG. 5.4. Relative residual histories of three preconditioned FGMRES(50) on the Helmholtz equation test of size $N = 100^3$.

486 **5.2.2. Compounding with SLR.** We also compounded the polynomial pre-
 487 conditioner with the nonsymmetric SLR preconditioner [19] and tested its precondi-
 488 tioning effect on the Helmholtz problem. The problem size was still kept at $N = 100^3$
 489 and the tolerance was set at $\tau = 10^{-3}$. We applied the SLR preconditioner to the
 490 shifted system $M_1 \approx A + \sigma I$ with a complex shift $\sigma = -0.4i$ (pulling the eigenvalues
 491 away from the origin), and then chose a polynomial preconditioner of degree 29 for
 492 the matrix $A_1 = M_1^{-1}A$. The approximate spectrum boundary of A_1 was obtained by
 493 running 80 steps of Arnoldi process, then it was uniformly discretized with 730 points.
 494 The convergence results as well as the comparison with SLR preconditioner are shown
 495 Table 5.3. We see that SLR preconditioned FGMRES(50) failed to converge in 1000
 496 iterations while the SLR compound polynomial preconditioner converged in only 29
 497 iterations and required the least iteration time among all seven methods tested in
 498 Table 5.2 and Table 5.3. Also notice that the construction time of this compound
 499 preconditioner is higher than other methods because it include both the SLR precondi-
 500 tioner construction time and the time to perform 80 steps of Arnoldi process.

501 In order to visualize the spectrum of the preconditioned matrix across different
 502 stages with this compound preconditioner, we also ran the experiment on a smaller
 503 Helmholtz problem of size $N = 20^3$ so that we were able to compute all eigenvalues
 504 of the matrices. Let M_1 denote the SLR preconditioner for the discretized Helmholtz
 505 operator A . The spectrum of $A_1 = M_1^{-1}A$ as well as the approximate eigenvalues
 506 from running 80 steps of the Arnoldi algorithm are shown in Figure 5.5a. A polyno-

TABLE 5.3

Convergence results of the SLR and SLR compound polynomial preconditioned FGMRES(50) on the Helmholtz equation test with size $N = 100^3$, the tolerance was fixed at $\tau = 10^{-3}$.

Preconditioner type	p-t	i-t	its
SLR preconditioner	86.94	\	F
SLR with polynomial of degree 30 – 1	145.85	308.03	29

507 mial preconditioner $p(A_1)$ of degree 29 was constructed and the contour map of the
 508 corresponding residual polynomial $|1 - zp(z)|$ is drawn in Figure 5.5b. Compared with
 509 Figure 5.3b, it is easy to see that the SLR preconditioner pushed the eigenvalues of A_1
 510 further away from the origin. Thus, a low-degree polynomial of p has already led to
 511 an efficient preconditioner, which is supported by both the contour map Figure 5.5b
 512 as well as the spectrum of the preconditioned matrix $A_1p(A_1)$ in Figure 5.5c.

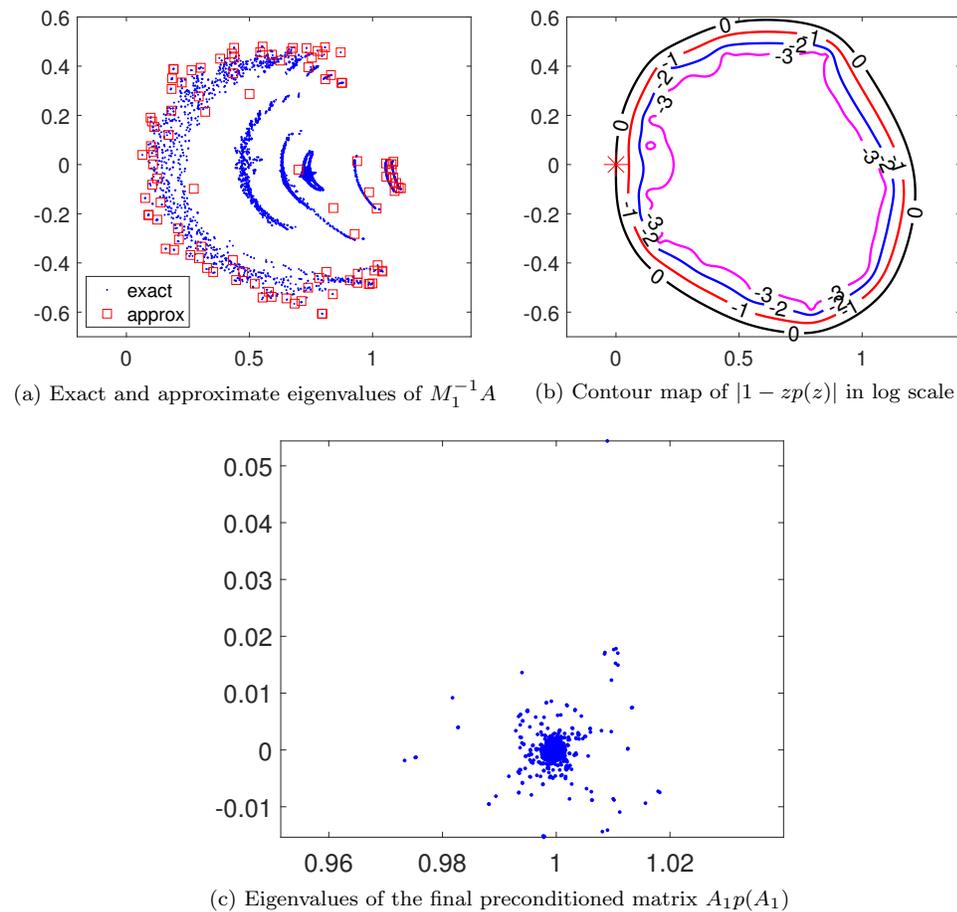


FIG. 5.5. Illustration of the preconditioning effect of both stages of the SLR compound preconditioner $p(M_1^{-1}A)$ on a small discretized Helmholtz equation test of size $N = 20^3$, where M_1 denotes the SLR preconditioner and p has degree 29.

513 **5.3. General sparse matrices.** We also tested the SLR compound poly-
 514 mial preconditioner on several general sparse matrices obtained from the SuiteSparse
 515 Matrix Collection. All of the test problems are nonsymmetric real/non-Hermitian
 516 complex. After the SLR preconditioner M_1 was constructed, we ran 60 steps of
 517 the Arnoldi algorithm to obtain the approximate eigenvalues, then the approximate
 518 spectrum boundary which was a polygon was discretized uniformly on each of its con-
 519 tinuous parts with step size $h = 0.005$. All polynomials were of degree 39 ($m = 40$)
 520 with recurrence length 2 and the tolerance for FGMRES(50) was set at $\tau = 10^{-10}$.
 521 The information of these matrices are listed in Table 5.4, the convergence results
 522 are shown in Table 5.5 together with those from ILUT as comparison. Note that
 523 the preconditioning set-up time for the SLR compound polynomial preconditioner
 524 includes time for SLR preconditioner construction, 60 steps of Arnoldi algorithm on
 525 $A_1 = M_1^{-1}A$ and time for building the polynomial. Despite slightly more expensive
 526 construction costs, SLR compound polynomial preconditioner outperformed ILUT on
 527 all of these 5 tests in the iteration phase.

TABLE 5.4
 Information of the test sparse matrices.

Group/matrix name	Order	nnz	Origin
Rajat/rajat09	24, 482	105, 573	circuit simulation problem
Dehghani/light_in_tissue	29, 282	406, 084	electromagnetics problem
Goodwin/Goodwin.127	178, 437	5, 778, 545	CFD problem
Kim/kim2	456, 976	11, 330, 020	3D problem
Bourchtein/atmosmodd	1, 270, 432	8, 814, 880	CFD problem

TABLE 5.5
 Convergence results of general sparse matrices with FGMRES(50) and tolerance $\tau = 10^{-10}$,
 all polynomials were of degree 39, column n shows the number of discretization points used on the
 spectrum boundary of A_1 .

matrix	ILUT			SLR compound polynomial			
	p-t	i-t	its	n	p-t	i-t	its
rajat09	F	\	\	763	0.46	3.27	44
light_in_tissue	0.071	2.39	213	572	0.54	0.71	6
Goodwin.127	F	\	\	460	2.19	409.74	311
kim2	4.85	20.52	38	1493	18.72	7.69	2
atmosmodd	1.17	454.66	397	557	36.10	42.71	6

528 **6. Conclusions.** The primary distinction between the polynomial precondition-
 529 ing techniques introduced in this paper and existing techniques is their aim at control-
 530 ling the numerical stability of the polynomial construction and the resulting iterative
 531 process, while allowing the degree of the polynomial to be high. Using a high degree
 532 polynomial is very important in order to guarantee a good quality preconditioner that
 533 will produce convergence in a smaller number of outer iterations. The performance of
 534 the method is significantly improved by a process which relies on a short-term recur-
 535 rence. It is clear that a big appeal of the method is its potential for great performance
 536 in a highly parallel, possibly GPU based, environment. The numerical experiments
 537 show that even in a scalar environment, the method can be effective in solving difficult
 538 problems when other techniques fail.

- [1] W. E. ARNOLDI, *The principle of minimized iterations in the solution of the matrix eigenvalue problem*, Quart. Appl. Math., 9 (1951), pp. 17–29.
- [2] M. M. BASKARAN AND R. BORDAWEKAR, *Optimizing sparse matrix-vector multiplication on GPUs using compile-time and run-time strategies*. technical report, 2008.
- [3] N. BELL AND M. GARLAND, *Implementing sparse matrix-vector multiplication on throughput-oriented processors*, in Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, Nov. 2009, pp. 1–11.
- [4] E. W. CHENEY, *Introduction to approximation theory*, AMS, Providence, Rhode Island, 1982.
- [5] J. P. COLEMAN, *Polynomial approximations in the complex plane*, J. Comput. Appl. Math., 18 (1987), pp. 193–211.
- [6] M. CROUZEIX, *Bounds for analytic functions of matrices*, Integr. Equ. Oper. Theory, 48 (2004), pp. 461–477, <https://doi.org/10.1007/s00020-002-1188-6>.
- [7] M. CROUZEIX, *Numerical range and functional calculus in Hilbert space*, J. Funct. Anal., 244 (2007), pp. 668–690.
- [8] M. CROUZEIX AND C. PALENCIA, *The numerical range is a $(1 + \sqrt{2})$ -spectral set*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 649–655.
- [9] G. DILLON, V. KALANTZIS, Y. XI, AND Y. SAAD, *A hierarchical low rank Schur complement preconditioner for indefinite linear systems*, SIAM J. Sci. Comput., 40 (2018), pp. A2234–A2252.
- [10] H. C. ELMAN, Y. SAAD, AND P. E. SAYLOR, *A hybrid Chebyshev Krylov subspace algorithm for solving nonsymmetric systems of linear equations*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 840–855.
- [11] M. EMBREE, J. A. LOE, AND R. B. MORGAN, *Polynomial preconditioned Arnoldi*, June 2018, <https://arxiv.org/abs/1806.08020>.
- [12] G. H. GOLUB AND R. S. VARGA, *Chebyshev semi-iterative methods, successive overrelaxation iterative methods, and second order Richardson iterative methods*, Numer. Math., 3 (1961), pp. 147–156.
- [13] M. H. GUTKNECHT AND S. RÖLLIN, *The Chebyshev iteration revisited*, Parallel Comput., 28 (2002), pp. 263–283.
- [14] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 52 (2011), pp. 217–288.
- [15] C. LANCZOS, *Trigonometric interpolation of empirical and analytical functions*, J. Math. and Phys., 17 (1938), pp. 123–199.
- [16] D. P. LAURIE AND L. M. VENTER, *a two-phase algorithm for the Chebyshev solution of complex linear equations*, SIAM J. Sci. Comput., 15 (1994), pp. 1440–1451.
- [17] R. LI AND Y. SAAD, *Divide and conquer low-rank preconditioners for symmetric matrices*, SIAM J. Sci. Comput., 35 (2013), pp. A2069–A2095.
- [18] R. LI AND Y. SAAD, *GPU-accelerated preconditioned iterative linear solvers*, J. Supercomput., 63 (2013), pp. 443–466.
- [19] R. LI, Y. XI, AND Y. SAAD, *Schur complement-based domain decomposition preconditioners with low-rank corrections*, Numer. Linear Algebra Appl., 23 (2016), pp. 706–729.
- [20] X. LIU, Y. XI, Y. SAAD, AND M. V. DE HOOP, *Solving the 3D high-frequency Helmholtz equation using contour integration and polynomial preconditioning*, Nov. 2018, <https://arxiv.org/abs/1811.12378>.
- [21] T. A. MANTEUFFEL AND G. STARKE, *On hybrid iterative methods for nonsymmetric systems of linear equations*, Numer. Math., 73 (1996), pp. 489–506.
- [22] N. M. NACHTIGAL, L. REICHEL, AND L. N. TREFETHEN, *A hybrid GMRES algorithm for nonsymmetric linear systems*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 796–825.
- [23] R. PACHÓN AND L. N. TREFETHEN, *Barycentric-Remez algorithms for best polynomial approximation in the chebfun system*, BIT, 49 (2009), pp. 721–741.
- [24] Y. SAAD, *Least squares polynomials in the complex plane and their use for solving nonsymmetric linear systems*, SIAM J. Numer. Anal., 24 (1987), pp. 155–169.
- [25] Y. SAAD, *Iterative methods for sparse linear systems*, SIAM, Philadelphia, 2nd ed., 2003.
- [26] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. and Stat. Comput., 7 (1986), pp. 856–869.
- [27] E. M. STEIN AND R. SHAKARCHI, *Complex Analysis, Princeton Lectures in Analysis, II*, Princeton University Press, Princeton, NJ, 2003.
- [28] R. L. STREIT, *Solution of systems of complex linear equations in the l_∞ norm with constraints on the unknowns*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 132–149.

- 600 [29] R. L. STREIT AND A. H. NUTTALL, *A note on the semi-infinite programming approach to*
601 *complex approximation*, Math. Comp., 40 (1983), pp. 599–605.
- 602 [30] P. T. P. TANG, *A fast algorithm for linear complex Chebyshev approximations*, Math. Comp.,
603 51 (1988), pp. 721–739.
- 604 [31] H. K. THORNQUIST, *Fixed-polynomial approximate spectral transformations for preconditioning*
605 *the eigenvalue problem*. PhD thesis, 2006.
- 606 [32] F. VÁZQUEZ, E. M. GARZÓN, J. A. MARTÍNEZ, AND J. J. FERNÁNDEZ, *The sparse matrix vector*
607 *product on GPUs*. technical report, 2008.
- 608 [33] G. A. WATSON, *A method for the Chebyshev solution of an overdetermined system of complex*
609 *linear equations*, IMA J. Numer. Anal., 8 (1988), pp. 461–471.
- 610 [34] Y. XI, R. LI, AND Y. SAAD, *An algebraic multilevel preconditioner with low-rank corrections*
611 *for sparse symmetric matrices*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 235–259.
- 612 [35] Y. XI AND Y. SAAD, *A rational function preconditioner for indefinite sparse linear systems*,
613 SIAM J. Sci. Comput., 39 (2017), pp. A1145–A1167.