

A Domain Decomposition Rayleigh-Ritz Algorithm for Symmetric Generalized Eigenvalue Problems

Vassilis Kalantzis

September 2020

EPrint ID: 2020.2

IBM Research
Thomas J. Watson Research Center

Preprints available from:

<https://researcher.watson.ibm.com/researcher/view.php?person=ibm-vkal>



A DOMAIN DECOMPOSITION RAYLEIGH–RITZ ALGORITHM FOR SYMMETRIC GENERALIZED EIGENVALUE PROBLEMS*

VASSILIS KALANTZIS[†]

Abstract. This paper proposes a parallel domain decomposition Rayleigh–Ritz projection scheme to compute a selected number of eigenvalues (and, optionally, associated eigenvectors) of large and sparse symmetric pencils. The projection subspace associated with interface variables is built by computing a few of the eigenvectors and associated leading derivatives of a zeroth-order approximation of the nonlinear matrix-valued interface operator. On the other hand, the projection subspace associated with interior variables is built independently in each subdomain by exploiting local eigenmodes and matrix resolvent approximations. The sought eigenpairs are then approximated by a Rayleigh–Ritz projection onto the subspace formed by the union of these two subspaces. Several theoretical and practical details are discussed, and upper bounds of the approximation errors are provided. Our numerical experiments demonstrate the efficiency of the proposed technique on sequential/distributed memory architectures as well as its competitiveness against schemes such as shift-and-invert Lanczos and automated multilevel substructuring combined with p -way vertex-based partitionings.

Key words. symmetric generalized eigenvalue problem, domain decomposition, high-performance computing, spectral Schur complement, Rayleigh–Ritz

AMS subject classifications. 65F15, 15A18, 65F50

DOI. 10.1137/19M1280004

1. Introduction. This paper proposes a Rayleigh–Ritz projection scheme based on algebraic domain decomposition to compute eigenvalues (and, optionally, associated eigenvectors) of large and sparse symmetric matrix pencils. In particular, our focus lies in the computation of a large number of eigenvalues located immediately on the right of some real scalar. Eigenvalue problems of this form appear in applications such as low-frequency response analysis [16, 31] and spectral clustering [39], among others. Extensions to the case where the sought eigenvalues are located immediately on the left of some real scalar are straightforward.

Computing eigenvalues located the closest to a given scalar is typically achieved by enhancing the projection method of choice by shift-and-invert [13, 41]. When the required accuracy in the approximation of the sought eigenvalues is not high, an alternative to Krylov subspace techniques is the automated multilevel substructuring technique (AMLS) [6, 7, 10, 28]. From an algebraic perspective, AMLS performs a Rayleigh–Ritz projection on a large projection subspace while avoiding excessive orthogonalization costs and has been demonstrated as a superior alternative to shift-and-invert Krylov subspace techniques when a very large number of eigenvalues is sought [29]. An analysis of the AMLS algorithm for elliptic PDE problems from a variational viewpoint can be found in [7]. Therein, tools from the mathematical theory of substructuring domain decomposition for elliptic PDEs are considered so as to understand how AMLS scales as the mesh is refined.

*Submitted to the journal’s Software and High-Performance Computing section August 8, 2019; accepted for publication (in revised form) September 29, 2020; published electronically DATE.

<https://doi.org/10.1137/19M1280004>

Funding: This work was supported by the International Business Machines Corporation through its Herman H. Goldstine Postdoctoral Fellowship program and by the Minnesota Supercomputing Institute (MSI) at the University of Minnesota.

[†]Thomas J. Watson Research Center, IBM Research, Yorktown Heights, NY 10598 USA (vkal@ibm.com).

43 This paper focuses on algebraic domain decomposition eigenvalue solvers where
 44 the concept of domain decomposition is applied directly to the eigenvalue equation.
 45 Algebraic domain decomposition eigenvalue solvers start by calling an algebraic graph
 46 partitioner to partition the graph associated with the given matrix pencil into a num-
 47 ber of nonoverlapping subdomains. The variables within each subdomain are then
 48 classified into two different categories: (a) (interior) variables which are coupled with
 49 variables located only in the same subdomain, and (b) (interface) variables which are
 50 both coupled with local variables and variables located in neighboring subdomains.
 51 The sought eigenpairs are then approximated by a Rayleigh–Ritz projection onto a
 52 subspace formed by the combination of subspaces associated with the two different
 53 types of variables. An in-depth analysis of algebraic domain decomposition eigenvalue
 54 solvers can be found in [20].

55 The main challenge of algebraic domain decomposition eigenvalue solvers is the
 56 construction of the projection subspace associated with interface variables due to the
 57 nonlinear nature of the interface matrix operator, also known as “spectral Schur com-
 58 plement” [21, 23]. From a purely algebraic perspective, AMLS builds this subspace
 59 by computing a few of the eigenvectors of a linear generalized eigenvalue problem
 60 involving the Schur complement matrix and its first derivative [6]. As a result, the
 61 accuracy provided by AMLS can deteriorate considerably as we look to compute ei-
 62 genvalues located away from the origin. An alternative suggested recently is to form
 63 the projection subspace associated with interface variables by applying a complex rati-
 64 onal transformation to the original pencil so as to annihilate components associated
 65 with unwanted eigenvalues, e.g., see the RF-DDES algorithm [24] and the discussion
 66 in [22, 27]. While these techniques can indeed lead to enhanced accuracy, multiple
 67 matrix factorizations computed in complex arithmetic are required.


68 In this paper we propose an algorithm which preserves advantages of algebraic
 69 domain decomposition eigenvalue solvers such as reduced orthogonalization costs and
 70 inherent parallelism while, at the same time, increases their accuracy without consid-
 71 ering more than one shift.

72 The key characteristics of the proposed scheme are summarized below.

73 (1) *Zeroth-order truncation of the interface matrix operator.* In contrast to AMLS,
 74 the algorithm proposed in this paper generates the projection subspace associated with
 75 interface variables by solving partially a standard eigenvalue problem with the Schur
 76 complement matrix. This approach avoids the need to compute/apply the derivative
 77 of the spectral Schur complement.

78 (2) *Exploiting Taylor series expansions of interface eigenvectors.* The accuracy of
 79 the projection subspace associated with interface variables is enhanced by expanding
 80 the (analytic) eigenvectors of the spectral Schur complement through their Taylor
 81 series and injecting a few leading eigenvector derivatives into the subspace. We show
 82 theoretically (and verify experimentally) that injecting up to second-order eigenvec-
 83 tor derivatives leads to eigenvalue approximations for which the upper bound of the
 84 absolute error reduces quartically. These eigenvector derivatives are computed inde-
 85 pendently of each other by exploiting deflated Krylov subspace solvers.

86 (3) *Reduced orthogonalization costs and enhanced parallelism.* Similarly to do-
 87 main decomposition schemes such as AMLS and RF-DDES, orthonormalization is
 88 applied only to vectors whose length is equal to the number of interface variables in-
 89 stead of the entire set of domain variables. This becomes especially important when
 90 both a large number of eigenvalues is sought and the number of interface variables
 91 is much smaller compared to the total number of equations/unknowns. In addition,
 92 the projection subspaces associated with interior variables in each subdomain are



106 FIG. 1.1. An illustration of the eigenvalue problem considered in this paper. Our goal is to
 107 compute the $n_{ev} = 5$ (real) eigenvalues located immediately on the right of $\alpha \in \mathbb{R}$. The sought
 108 eigenvalues are denoted by $\lambda_{\kappa_1}, \dots, \lambda_{\kappa_{n_{ev}}}$, $1 \leq \kappa_1 \leq \kappa_{n_{ev}} \leq n$.

93 built independently of each other by computing local eigenmodes and computing up
 94 to second-order resolvent expansions. We report experiments performed on sequen-
 95 tial/distributed memory architectures and demonstrate the suitability of the proposed
 96 technique in high-performance computing settings.

97 **1.1. Notation and outline.** The eigenvalue problem considered in this paper
 98 is of the form $Ax = \lambda Mx$ where the matrices $A \in \mathbb{R}^{n \times n}$ and $M \in \mathbb{R}^{n \times n}$ are assumed
 99 large, sparse, and symmetric, and matrix M is also positive-definite (SPD). Our
 100 goal is to compute the $n_{ev} \ll n$ eigenvalues located immediately on the right of
 101 a user-given scalar α . An illustrative example is shown in Figure 1.1. Extensions
 102 of the proposed technique to the computation of eigenvalues located immediately
 103 on the left of a user-given scalar are straightforward. Throughout the rest of this
 104 paper we will denote the pencil $L - \lambda K$ by (L, K) , and for any such pencil we define
 105 $\Lambda(L, K) := \{\lambda \mid \det[L - \lambda K] = 0\}$.

109 The outline of this paper is as follows. Section 2 gives a brief introduction to
 110 Rayleigh–Ritz projection subspaces from the viewpoint of algebraic domain decom-
 111 position. Section 3 presents a spectral Schur complement approach to approximate a
 112 single eigenpair of the pencil (A, M) . Section 4 extends these results to a Rayleigh–
 113 Ritz projection that approximates all sought eigenpairs from a common projection
 114 subspace. Section 5 presents numerical experiments performed on sequential and dis-
 115 tributed memory environments. Finally, in section 6 we give our concluding remarks.

116 2. Rayleigh–Ritz projections from a domain decomposition viewpoint.

117 The standard approach to compute a few eigenpairs of sparse and symmetric ma-
 118 trix pencils is by applying the Rayleigh–Ritz technique onto a carefully constructed
 119 subspace \mathcal{Z} of \mathbb{R}^n [33]. The goal is to find a subspace \mathcal{Z} that includes an invariant
 120 subspace associated with the n_{ev} sought eigenvalues $\lambda_{\kappa_1}, \dots, \lambda_{\kappa_{n_{ev}}}$. The sought eigen-
 121 pairs can then be recovered (in the absence or roundoff errors) as a subset of the Ritz
 122 pairs of the matrix pencil $(Z^T A Z, Z^T M Z)$, where matrix Z represents a basis of \mathcal{Z} .

123 Let the matrices A and M be partitioned in a 2×2 block form

$$124 \quad A = \begin{pmatrix} B & E \\ E^T & C \end{pmatrix} \quad \text{and} \quad M = \begin{pmatrix} M_B & M_E \\ M_E^T & M_C \end{pmatrix},$$

126 where B and M_B are square matrices of size $d \times d$, E and M_E are rectangular matrices
 127 of size $d \times s$, C and M_C are square matrices of size $s \times s$, and $n = d + s$. Without
 128 loss of generality we assume $n_{ev} \leq s$. Similarly, the eigenvector $x^{(i)}$ associated with
 129 eigenvalue λ_i can be written as

$$130 \quad x^{(i)} = \begin{pmatrix} u^{(i)} \\ y^{(i)} \end{pmatrix}, \quad u^{(i)} \in \mathbb{R}^d, \quad y^{(i)} \in \mathbb{R}^s.$$

132 Rewriting $Ax^{(i)} = \lambda_i Mx^{(i)}$ using the above block form gives

$$(2.1) \quad \begin{pmatrix} B - \lambda_i M_B & E - \lambda_i M_E \\ E^T - \lambda_i M_E^T & C - \lambda_i M_C \end{pmatrix} \begin{pmatrix} u^{(i)} \\ y^{(i)} \end{pmatrix} = 0,$$

while eliminating $u^{(i)}$ from the second row equation in (2.1) leads to the $s \times s$ nonlinear eigenvalue problem

$$(2.2) \quad \left[C - \lambda_i M_C - (E - \lambda_i M_E)^T (B - \lambda_i M_B)^{-1} (E - \lambda_i M_E) \right] y^{(i)} = 0,$$

from which λ_i and $y^{(i)}$ can be determined. The missing $d \times 1$ part of $x^{(i)}$, $u^{(i)}$, is then recovered by solving the linear system

$$(B - \lambda_i M_B) u^{(i)} = -(E - \lambda_i M_E) y^{(i)}.$$

From a domain decomposition perspective, an ideal choice is to set $\mathcal{Z} = \mathcal{U} \oplus \mathcal{Y}$ where

$$(2.3) \quad \mathcal{U} = \text{span} \left(\begin{bmatrix} u^{(\kappa_1)}, \dots, u^{(\kappa_{nev})} \\ 0_{s, nev} \end{bmatrix} \right), \quad \mathcal{Y} = \text{span} \left(\begin{bmatrix} 0_{d, nev} \\ y^{(\kappa_1)}, \dots, y^{(\kappa_{nev})} \end{bmatrix} \right),$$

and $0_{\chi, \psi}$ denotes the zero matrix of size $\chi \times \psi$.

The main goal of algebraic domain decomposition eigenvalue solvers is to build a projection subspace which, ideally, includes the subspace in (2.3).

2.1. Domain decomposition reordering. Practical applications of algebraic domain decomposition eigenvalue solvers rely on relabeling the unknowns/equations of the eigenvalue problem $Ax = \lambda Mx$ such that the matrix $B - \lambda M_B$ is block-diagonal. This can be easily achieved by applying a graph partitioner to the adjacency graph of the matrix $|A| + |M|$, e.g., [25, 34].

In this paper we only consider p -way partitionings, and the partitioner divides the graph into $p > 1$ nonoverlapping subdomains. The rows/columns of matrices A and M are then reordered so that unknowns/equations associated with interior variables (i.e., nodes of the adjacency graph which are connected only to nodes located in the same partition) are listed before those associated with interface variables (i.e., nodes of the adjacency graph which are connected with nodes located in neighboring partitions). The permuted matrices A and M can be written¹ as

$$(2.4) \quad A := PAP^T = \begin{pmatrix} B_1 & & & E_1 \\ & B_2 & & E_2 \\ & & \ddots & \vdots \\ & & & B_p & E_p \\ E_1^T & E_2^T & \dots & E_p^T & C \end{pmatrix}, \quad \text{and}$$

$$M := PMP^T = \begin{pmatrix} M_B^{(1)} & & & M_E^{(1)} \\ & M_B^{(2)} & & M_E^{(2)} \\ & & \ddots & \vdots \\ & & & M_B^{(p)} & M_E^{(p)} \\ \left(M_E^{(1)} \right)^T & \left(M_E^{(2)} \right)^T & \dots & \left(M_E^{(p)} \right)^T & M_C \end{pmatrix}.$$

¹The eigenvalues of the pencil (A, M) in (2.4) are identical to those prior to the symmetric permutation. If eigenvectors are also of interest, these need be postmultiplied by P^T . Throughout the rest of this paper we will work with the reordered matrices shown in (2.4).

164 Let us denote the number of interior and interface variables residing in the j th
 165 subdomain by d_j and s_j , respectively. Matrices B_j and $M_B^{(j)}$ are of size $d_j \times d_j$, while
 166 E_j and $M_E^{(j)}$ are rectangular matrices of size $d_j \times s$ where $s = \sum_{j=1}^p s_j$. In particular,
 167 E_j and $M_E^{(j)}$ have a special nonzero pattern of the form $E_j = [0_{d_j, \ell_j}, \hat{E}_j, 0_{d_j, \xi_j}]$, and
 168 $M_E^{(j)} = [0_{d_j, \ell_j}, \hat{M}_E^{(j)}, 0_{d_j, \xi_j}]$, where $\ell_j = \sum_{k=1}^{k < j} s_k$, and $\xi_j = \sum_{k > j}^{k=p} s_k$. Note that
 169 typically the value of p is chosen such that $s \ll d$.

170 Taking advantage of the Schur complements framework, the vectors $u^{(i)}$ and $y^{(i)}$
 171 can be further partitioned as

$$172 \quad u^{(i)} = \begin{pmatrix} u_1^{(i)} \\ \vdots \\ u_p^{(i)} \end{pmatrix} \quad \text{and} \quad y^{(i)} = \begin{pmatrix} y_1^{(i)} \\ \vdots \\ y_p^{(i)} \end{pmatrix},$$

173 where $u_j^{(i)} \in \mathbb{R}^{d_j}$ and $y_j^{(i)} \in \mathbb{R}^{s_j}$ denote the components of vectors $u^{(i)}$ and $y^{(i)}$ which
 174 are associated with the j th subdomain, respectively. Once the subvector $y_j^{(i)}$ becomes
 175 available, $u_j^{(i)}$ is computed independently of the rest of the subvectors of $u^{(i)}$ by solving
 176 the local linear system
 177

$$178 \quad (B_j - \lambda_i M_B^{(j)}) u_j^{(i)} = -(\hat{E}_j - \lambda_i \hat{M}_E^{(j)}) y_j^{(i)}.$$

180 **3. A scheme to approximate a single eigenpair.** This section considers the
 181 approximation of a single eigenpair (λ, x) by $(\hat{\lambda}, \hat{x})$, where $\hat{\lambda} = \frac{\hat{x}^T A \hat{x}}{\hat{x}^T M \hat{x}}$ is the Rayleigh
 182 quotient associated with the approximate eigenvector \hat{x} .

183 **3.1. Spectral Schur complements.** We begin by defining the univariate, non-
 184 linear, matrix-valued function

$$185 \quad S : \zeta \in \mathbb{R} \rightarrow S(\zeta) \in \mathbb{R}^{s \times s}, \quad S(\zeta) = C - \zeta M_C - (E - \zeta M_E)^T (B - \zeta M_B)^{-1} (E - \zeta M_E).$$

187 For each $\zeta \in \mathbb{R} \setminus \Lambda(B, M_B)$, there exist s real eigenvalues and corresponding orthogonal
 188 eigenvectors of $S(\zeta)$. When $\zeta \in \Lambda(B, M_B)$, there exist at least $s - d$ (if $s \geq d$)
 189 eigenvalues of $S(\zeta)$ which are well-defined (see also section 3.2).

190 **DEFINITION 3.1.** For $j = 1, \dots, s$, we define the scalar-vector pairs

$$191 \quad (\mu_j, y_j) : \zeta \in \mathbb{R} \setminus \mathcal{D}_j \rightarrow (\mu_j(\zeta), y_j(\zeta)) \in \{\mathbb{R}, \mathbb{R}^s\}, \quad \text{where } \mathcal{D}_j \subseteq \Lambda(B, M_B),$$

193 such that for any $\zeta \notin \mathcal{D}_j$, the pair $(\mu_j(\zeta), y_j(\zeta))$ satisfies

$$194 \quad S(\zeta) y_j(\zeta) = \mu_j(\zeta) y_j(\zeta).$$

196 Each function $\mu_j(\zeta)$ (from now on referred to as ‘‘eigenvalue curve’’) has either no
 197 or a finite number of poles \mathcal{D}_j , and these poles are eigenvalues of the pencil (B, M_B) .
 198 Moreover, the corresponding eigenvector $y_j(\zeta)$ is uniquely defined (up to a normalizing
 199 factor) and the associated spectral projector is analytic [4, 26]. Throughout the rest
 200 of this paper we assume that the spectrum of $S(\zeta)$ is implicitly arranged so that
 201 the eigenvalue curves $\mu_1(\zeta), \dots, \mu_s(\zeta)$ are analytic functions of ζ everywhere within
 202 their domain of definition.² This assumption corresponds to a mere reordering of
 203 the eigenpairs of $S(\zeta)$ and does not affect the practicality of the algorithm proposed
 204 throughout this paper nor require any additional work.

²See also the discussion in [38, section 4].

205 *Remark 1.* One idea to order the subscripts of the eigenvalue curves is to denote
 206 by $\mu_j(\zeta)$ the eigenvalue curve for which $\mu_j(\zeta_0)$, $\zeta_0 \in \mathbb{R} \setminus \Lambda(B, M_B)$ is equal to the j th
 207 algebraically smallest eigenvalue of matrix $S(\zeta_0)$. Throughout this paper we denote
 208 by $\mu_j(\zeta)$ the eigenvalue curve for which $\mu_j(0)$ is equal to the j th algebraically smallest
 209 eigenvalue of matrix $S(0)$.

210 **3.2. Behavior of eigenvalue curves at their poles.** In general, it is not
 211 possible to determine what eigenvalues (if any) of (B, M_B) are poles of $\mu_j(\zeta)$, nor
 212 does the algorithm proposed in this paper require such information. Nonetheless, we
 213 can determine the number of eigenvalue curves that remain analytic as ζ approaches
 214 an eigenvalue of (B, M_B) .

215 **DEFINITION 3.2.** *The eigenpairs of the matrix pencil (B, M_B) will be denoted by*
 216 $(\delta_\ell, v^{(\ell)})$, $\ell = 1, \dots, d$, where $V = [v^{(1)}, v^{(2)}, \dots, v^{(d)}]$ is scaled so that $V^T M_B V = I$.

217 Let $w_\zeta^{(\ell)} = (E - \zeta M_E)^T v^{(\ell)}$, $\ell = 1, \dots, d$, and assume that δ_k is an eigenvalue
 218 of (B, M_B) with multiplicity $\rho_\delta \leq s$, and corresponding eigenvectors $v_k^{(1)}, \dots, v_k^{(\rho_\delta)}$.
 219 Then, if $\mathbf{rank}(\lim_{\zeta \rightarrow \delta_k} (E - \zeta M_E)^T [v_k^{(1)}, \dots, v_k^{(\rho_\delta)}]) = \theta \leq \rho_\delta$, there exist integers
 220 $j_1, \dots, j_\theta \in \{1, \dots, s\}$ such that

$$221 \quad \begin{cases} \lim_{\zeta \rightarrow \delta_k} \mu_{\{j_1, \dots, j_\theta\}}(\zeta) = -\infty, & \text{when } \zeta < \delta_k, \text{ and} \\ \lim_{\zeta \rightarrow \delta_k} \mu_{\{j_1, \dots, j_\theta\}}(\zeta) = +\infty, & \text{when } \zeta > \delta_k. \end{cases}$$

223 As $\zeta \rightarrow \delta_k$, all but the above eigenvalue curves cross δ_k in an analytical manner. The
 224 eigenvalue curves are strictly decreasing in their entire domain of definition. Details
 225 on the above discussion can be found in [23, Theorem 4.1] and [30, Theorem 2.1] when
 226 $M = I$. Extensions to the case $M \neq I$ are straightforward.

227 **3.3. Taylor series expansion of $y_j(\lambda)$.** Following (2.2), a scalar $\lambda \notin \Lambda(B, M_B)$
 228 is an eigenvalue³ of the matrix pencil (A, M) if and only if there exists an integer
 229 $1 \leq j \leq s$ such that $\mu_j(\lambda) = 0$ (e.g., see Figure 3.1). The eigenvector $y_j(\lambda)$ associated
 230 with the eigenvalue $\mu_j(\lambda)$ is then equal to the bottom $s \times 1$ subvector of eigenvector x
 231 associated with λ . Therefore, computing $y_j(\lambda)$ is the first step toward approximating
 232 the eigenvector x .

233 Let now $\sigma \in [\lambda^-, \lambda^+]$ be an approximation of λ , where $[\lambda^-, \lambda^+]$ is located between
 234 two consecutive poles of $\mu_j(\zeta)$ (if such poles exist).⁴ In the ideal scenario, we have
 235 $\sigma \equiv \lambda$, which leads to $y_j(\sigma) = y_j(\lambda)$. In practice, we can only hope that $\sigma \approx \lambda$,
 236 and thus $y_j(\sigma)$ is only an approximation of $y_j(\lambda)$. To improve this approximation, we
 237 exploit the analyticity of the eigenpair $(\mu_j(\zeta), y_j(\zeta))$.


240 Let $\frac{d^i y_j(\zeta)}{d\zeta^i}$ denote the i th derivative of the univariate vector-valued function $y_j(\zeta)$.
 241 Expanding $y_j(\lambda)$ through its Taylor series around σ gives

$$242 \quad (3.1) \quad y_j(\lambda) = \sum_{i=0}^{\infty} \frac{(\lambda - \sigma)^i}{i!} \left(\frac{d^i y_j(\zeta)}{d\zeta^i} \right)_{\zeta=\sigma}.$$

244 The above expression suggests that even when σ is not very close to λ , we can improve
 245 the approximation of $y_j(\lambda)$ by considering higher-order derivatives of $y_j(\zeta)$ evaluated
 246 at σ .

³The case where $\lambda \in \Lambda(B, M_B)$ is more involved as it is possible that $\mathbf{det}[S(\lambda)] \neq 0$. Nonetheless, such scenarios can be easily detected, e.g., see [30].

⁴In practice a pole of $\mu_j(\zeta)$ located inside the interval $[\lambda^-, \lambda^+]$ poses no threat as long as σ is chosen carefully.



238 FIG. 3.1. Illustration of the concept of eigenvalue curves. We assume that $\lambda \equiv \lambda_1$ is a root of
 239 the eigenvalue curve $\mu_1(\zeta)$. The figure shows two potential choices of the variable $\sigma \in [\lambda^-, \lambda^+]$.

247 Following (3.1), the eigenvector x can be written as

$$\begin{aligned}
 x &= \begin{pmatrix} -(B - \lambda M_B)^{-1}(E - \lambda M_E)y_j(\lambda) \\ y_j(\lambda) \end{pmatrix} \\
 &= \begin{pmatrix} -(B - \lambda M_B)^{-1}(E - \lambda M_E) \left[\sum_{i=0}^{\infty} \frac{(\lambda - \sigma)^i}{i!} \left(\frac{d^i y_j(\zeta)}{d\zeta^i} \right)_{\zeta=\sigma} \right] \\ \sum_{i=0}^{\infty} \frac{(\lambda - \sigma)^i}{i!} \left(\frac{d^i y_j(\zeta)}{d\zeta^i} \right)_{\zeta=\sigma} \end{pmatrix} \\
 &= \begin{pmatrix} -(B - \lambda M_B)^{-1}(E - \lambda M_E) \left[y_j(\sigma), \left(\frac{dy_j(\zeta)}{d\zeta} \right)_{\zeta=\sigma}, \left(\frac{d^2 y_j(\zeta)}{d\zeta^2} \right)_{\zeta=\sigma}, \dots \right] \\ \left[y_j(\sigma), \left(\frac{dy_j(\zeta)}{d\zeta} \right)_{\zeta=\sigma}, \left(\frac{d^2 y_j(\zeta)}{d\zeta^2} \right)_{\zeta=\sigma}, \dots \right] \end{pmatrix} \begin{pmatrix} 1 \\ (\lambda - \sigma)/1! \\ (\lambda - \sigma)^2/2! \\ \vdots \end{pmatrix}.
 \end{aligned}$$

249
 250

251 3.4. Rayleigh quotient approximation of λ .

252 DEFINITION 3.3. We define the following matrix-valued functions of $\zeta \in \mathbb{R}$:

253
 254

$$B_\zeta = B - \zeta M_B, \quad E_\zeta = E - \zeta M_E, \quad \text{and} \quad C_\zeta = C - \zeta M_C.$$

255
 256

DEFINITION 3.4. We define the M -norm of an SPD $n \times n$ matrix M and a nonzero vector $x \in \mathbb{R}^n$ to be equal to $\|x\|_M = \sqrt{x^T M x}$.

257

3.4.1. A basic approximation.

258

PROPOSITION 3.5. Let $\lambda \in \Lambda(A, M)$ satisfy $\mu_j(\lambda) = 0$, and $\lambda^- \leq \sigma \leq \lambda \leq \lambda^+$, where λ is the only root of $\mu_j(\zeta)$ located inside the interval $[\lambda^-, \lambda^+]$. Additionally, let $\tau \in \mathbb{Z}$, and define the vectors

260

$$(3.2) \quad \hat{y} = \left[y_j(\sigma), \left(\frac{dy_j(\zeta)}{d\zeta} \right)_{\zeta=\sigma}, \dots, \left(\frac{d^\tau y_j(\zeta)}{d\zeta^\tau} \right)_{\zeta=\sigma} \right] \begin{pmatrix} \lambda^{-\sigma/1!} \\ \vdots \\ (\lambda-\sigma)^\tau / \tau! \end{pmatrix}$$

and

$$\hat{x} = \begin{pmatrix} -B_\sigma^{-1} E_\sigma \hat{y} \\ \hat{y} \end{pmatrix}.$$

Then, if $\hat{\lambda} = \frac{\hat{x}^T A \hat{x}}{\hat{x}^T M \hat{x}}$,

$$|\lambda - \hat{\lambda}| = O((\lambda - \sigma)^2),$$

where the big- O symbol is to be interpreted in the limit as $\sigma \rightarrow \lambda$.

Proof. Let $\rho(z) = \frac{z^T A z}{z^T M z}$ be the Rayleigh quotient of any nonzero vector $z \in \mathbb{R}^n$. Expanding $\rho(z)$ through its Taylor series expansion around the eigenvector x gives

$$\rho(z) = \rho(x) + (z - x)^T \nabla \rho(x) + O(\|z - x\|_M^2) \quad \text{as } \hat{z} \rightarrow x.$$

The gradient of the Rayleigh quotient is equal to $\nabla \rho(z) = 2 \frac{Az(z^T M z) - Mz(z^T A z)}{(z^T M z)^2}$, and thus $\nabla \rho(x) = 0$. It follows that

$$|\lambda - \hat{\lambda}| = |\rho(x) - \rho(\hat{x})| = O(\|\hat{x} - x\|_M^2) \quad \text{as } \hat{x} \rightarrow x.$$

Write now $E_\lambda = E_\sigma - (\lambda - \sigma)M_E$ and define the vector

$$r = \begin{pmatrix} -B_\lambda^{-1} E_\lambda \left[\left(\frac{d^{\tau+1} y_j(\zeta)}{d\zeta^{\tau+1}} \right)_{\zeta=\sigma}, \left(\frac{d^{\tau+2} y_j(\zeta)}{d\zeta^{\tau+2}} \right)_{\zeta=\sigma}, \dots \right] \\ \left[\left(\frac{d^{\tau+1} y_j(\zeta)}{d\zeta^{\tau+1}} \right)_{\zeta=\sigma}, \left(\frac{d^{\tau+2} y_j(\zeta)}{d\zeta^{\tau+2}} \right)_{\zeta=\sigma}, \dots \right] \end{pmatrix} \begin{pmatrix} (\lambda - \sigma)^{\tau+1} / (\tau + 1)! \\ (\lambda - \sigma)^{\tau+2} / (\tau + 2)! \\ \vdots \end{pmatrix}.$$

The difference $x - \hat{x}$ can then be written as

$$x - \hat{x} = r - \begin{pmatrix} [B_\lambda^{-1} - B_\sigma^{-1}] E_\sigma \hat{y} \\ 0 \end{pmatrix} + \begin{pmatrix} (\lambda - \sigma) B_\lambda^{-1} M_E \hat{y} \\ 0 \end{pmatrix}.$$

Let now $E_\sigma \hat{y}$ and $M_E \hat{y}$ be expanded in the basis $\{M_B v^{(\ell)}\}_{\ell=1, \dots, d}$:

$$E_\sigma \hat{y} = M_B \sum_{\ell=1}^d \epsilon_\ell v^{(\ell)}, \quad M_E \hat{y} = M_B \sum_{\ell=1}^d \gamma_\ell v^{(\ell)},$$

where $[\epsilon_\ell, \gamma_\ell]^T \in \mathbb{R}^2$ are the expansion coefficients associated with the direction $v^{(\ell)}$.

Taking advantage of the identity $B_\lambda^{-1} - B_\sigma^{-1} = (\lambda - \sigma) B_\lambda^{-1} M_B B_\sigma^{-1}$ and noticing that $\|r\|_M = O((\lambda - \sigma)^{\tau+1})$ leads to

(3.3)

$$\begin{aligned} \|x - \hat{x}\|_M^2 &= \left\| r - \begin{pmatrix} (\lambda - \sigma) B_\lambda^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y} \\ 0 \end{pmatrix} \right\|_M^2 \\ &= \|r\|_M^2 + \|(\lambda - \sigma) B_\lambda^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y}\|_{M_B}^2 \\ &\quad - 2(\lambda - \sigma) r^T M \begin{pmatrix} B_\lambda^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y} \\ 0 \end{pmatrix} \\ &= O((\lambda - \sigma)^{2\tau+2}) + (\lambda - \sigma)^2 \sum_{\ell=1}^d \left(\frac{\epsilon_\ell - \gamma_\ell (\delta_\ell - \sigma)}{(\delta_\ell - \sigma)(\delta_\ell - \lambda)} \right)^2 + O((\lambda - \sigma)^{\tau+2}), \end{aligned}$$

295 where the big- O symbol is to be interpreted in the limit as $\sigma \rightarrow \lambda$, and we made use
296 of the relations

$$297 \quad (\lambda - \sigma)B_\lambda^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y} = (\lambda - \sigma) \sum_{\ell=1}^d \left(\frac{\epsilon_\ell - \gamma_\ell (\delta_\ell - \sigma)}{(\delta_\ell - \sigma)(\delta_\ell - \lambda)} \right) v^{(\ell)}$$

299 and

$$300 \quad \left\| \begin{pmatrix} (\lambda - \sigma)B_\lambda^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y} \\ 0 \end{pmatrix} \right\|_M^2 = \|(\lambda - \sigma)B_\lambda^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y}\|_{M_B}^2$$

$$301 \quad = (\lambda - \sigma)^2 \sum_{\ell=1}^d \left(\frac{\epsilon_\ell - \gamma_\ell (\delta_\ell - \sigma)}{(\delta_\ell - \sigma)(\delta_\ell - \lambda)} \right)^2. \quad \square$$

303 Proposition 3.5 remains valid even when the interval $[\lambda^-, \lambda^+]$ includes more than
304 one root of $\mu_j(\zeta)$ (including multiple eigenvalues). However, only one of these ei-
305 genvalues can be approximated. Moreover, the interval $[\lambda^-, \lambda^+]$ can include poles of
306 $\mu_j(\zeta)$ but σ needs to be algebraically smaller than these poles.

307 **3.4.2. Improving accuracy by deflation.** The bound on $\|x - \hat{x}\|_M^2$ appearing
308 in Proposition 3.5 can be improved (reduced) by explicitly removing those directions
309 in which $\|x - \hat{x}\|_M^2$ is large, i.e., the directions corresponding to the eigenvectors associ-
310 ated with the few smallest (in magnitude) eigenvalues of the matrix pencil (B_σ, M_B) .
311 This is especially true for the second and third terms shown in (3.3), both of which
312 depend on the distance of δ_ℓ , $\ell = 1, \dots, d$, from both σ and λ .

313 More specifically, let the approximate eigenvector \hat{x} set as

$$314 \quad (3.4) \quad \hat{x} = \begin{pmatrix} -B_\sigma^{-1} E_\sigma \hat{y} \\ \hat{y} \end{pmatrix} - \begin{pmatrix} [v^{(1)}, \dots, v^{(\hat{\kappa})}] \\ 0 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \vdots \\ \nu_{\hat{\kappa}} \end{pmatrix},$$

316 where \hat{y} is defined in (3.2), and $\nu_j = (\lambda - \sigma) \left(\frac{\epsilon_j - \gamma_j (\delta_j - \sigma)}{(\delta_j - \lambda)(\delta_j - \sigma)} \right)$, $j = 1, \dots, \hat{\kappa} \leq d$. Following
317 the reasoning in Proposition 3.5 (we omit the intermediate steps), we can show

$$318 \quad \|x - \hat{x}\|_M^2 = O((\lambda - \sigma)^{2\tau+2}) + (\lambda - \sigma)^2 \sum_{\ell=\hat{\kappa}+1}^d \left(\frac{\epsilon_\ell - \gamma_\ell (\delta_\ell - \sigma)}{(\delta_\ell - \sigma)(\delta_\ell - \lambda)} \right)^2 + O((\lambda - \sigma)^{\tau+2}).$$

320 The eigenvector approximation \hat{x} shown in (3.4) becomes appealing when $\hat{\kappa}$ is set
321 such that the eigenvalues $\delta_{\hat{\kappa}+1}, \dots, \delta_d$ lie far away from both λ and σ .

322 **3.4.3. Reducing the asymptotic order of the upper bound.** The analysis
323 in Proposition 3.5 suggests that regardless of the value of τ , at the limit $\sigma \rightarrow \lambda$ the
324 term $|\lambda - \hat{\lambda}|$ will be of the order $O((\lambda - \sigma)^2)$ due to the approximation of the term
325 $B_\lambda^{-1} E_\lambda$ by $B_\sigma^{-1} E_\sigma$.

326 Let us write $B_\lambda^{-1} = B_{\sigma+\epsilon}^{-1}$ where $\epsilon < \min_{1 \leq \ell \leq d} |\delta_\ell - \sigma|$. The matrix resolvent
327 can then be written as $B_\lambda^{-1} = ((I - \epsilon M_B B_\sigma^{-1}) B_\sigma)^{-1} = B_\sigma^{-1} \sum_{k=0}^{\infty} [\epsilon M_B B_\sigma^{-1}]^k$. The
328 approximation $u \approx -B_\sigma^{-1} E_\sigma \hat{y}$ can be replaced by $u \approx -(B_\sigma^{-1} + \epsilon B_\sigma^{-1} M_B B_\sigma^{-1}) E_\sigma \hat{y}$.
329 When $M_E \neq 0$, the error term associated with $B_\lambda^{-1} M_E \hat{y}$ can also be smoothed out
330 by adding the vector $B_\sigma^{-1} M_E \hat{y}$, i.e., $u \approx -(B_\sigma^{-1} E_\sigma + \epsilon B_\sigma^{-1} M_B B_\sigma^{-1} E_\sigma - B_\sigma^{-1} M_E) \hat{y}$.

331 PROPOSITION 3.6. Let $\lambda \in \Lambda(A, M)$ satisfy $\mu_j(\lambda) = 0$, and $\lambda^- \leq \sigma \leq \lambda \leq \lambda^+$,
 332 where λ is the only root of $\mu_j(\zeta)$ located inside the interval $[\lambda^-, \lambda^+]$. Additionally, let
 333 $\tau \in \mathbb{Z}$ and define the vectors

$$334 \quad \hat{y} = \left[y_j(\sigma), \left(\frac{dy_j(\zeta)}{d\zeta} \right)_{\zeta=\sigma}, \dots, \left(\frac{d^\tau y_j(\zeta)}{d\zeta^\tau} \right)_{\zeta=\sigma} \right] \begin{pmatrix} \lambda^{-\sigma/1!} \\ \vdots \\ (\lambda-\sigma)^\tau / \tau! \end{pmatrix}$$

335 and
 336

$$337 \quad \hat{x} = \begin{pmatrix} -[B_\sigma^{-1} + (\lambda - \sigma)B_\sigma^{-1}M_B B_\sigma^{-1}] E_\sigma \hat{y} \\ \hat{y} \end{pmatrix} + \begin{pmatrix} B_\sigma^{-1} M_E \hat{y} \\ 0 \end{pmatrix} - \begin{pmatrix} [v^{(1)}, \dots, v^{(\hat{\kappa})}] \\ 0 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \vdots \\ \nu_{\hat{\kappa}} \end{pmatrix},$$

338 where $\nu_j = (\lambda - \sigma)^2 \left(\frac{\epsilon_j - \gamma_j (\delta_j - \sigma)}{(\delta_j - \lambda)(\delta_j - \sigma)^2} \right)$, $j = 1, \dots, \hat{\kappa} \leq d$.

340 Then, if $\hat{\lambda} = \frac{\hat{x}^T A \hat{x}}{\hat{x}^T M \hat{x}}$,

$$341 \quad |\lambda - \hat{\lambda}| = O((\lambda - \sigma)^\chi), \quad \text{where} \quad \begin{cases} \chi = 2, & \text{when } \tau = 0, \text{ and} \\ \chi = 4, & \text{when } \tau \geq 1, \end{cases}$$

342 and the big-O symbol is to be interpreted in the limit as $\sigma \rightarrow \lambda$.

343 *Proof.* First notice that

$$344 \quad (\lambda - \sigma) (B_\lambda^{-1} - B_\sigma^{-1}) = (\lambda - \sigma)^2 B_\lambda^{-1} M_B B_\sigma^{-1},$$

$$345 \quad (B_\lambda^{-1} - B_\sigma^{-1} - (\lambda - \sigma) B_\sigma^{-1} M_B B_\sigma^{-1}) = (\lambda - \sigma)^2 B_\lambda^{-1} M_B B_\sigma^{-1} M_B B_\sigma^{-1}.$$

346 We can then write the difference $x - \hat{x}$ as

$$347 \quad x - \hat{x} = r - (\lambda - \sigma)^2 \left[\begin{pmatrix} B_\lambda^{-1} M_B B_\sigma^{-1} M_B B_\sigma^{-1} E_\sigma \hat{y} \\ 0 \end{pmatrix} - \begin{pmatrix} B_\lambda^{-1} M_B B_\sigma^{-1} M_E \hat{y} \\ 0 \end{pmatrix} \right]$$

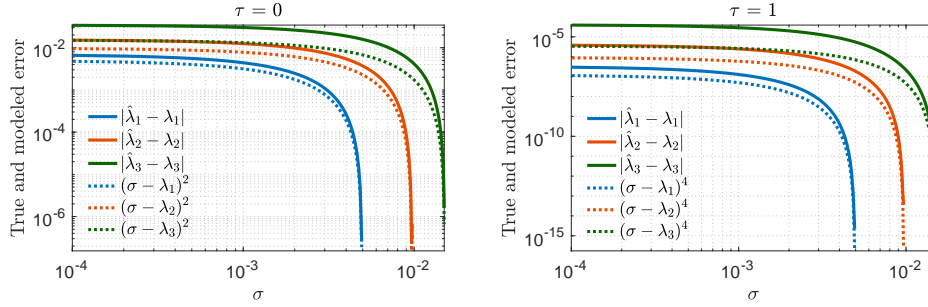
$$348 \quad + \begin{pmatrix} [v^{(1)}, \dots, v^{(\hat{\kappa})}] \\ 0 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \vdots \\ \nu_{\hat{\kappa}} \end{pmatrix}$$

$$349 \quad = r - (\lambda - \sigma)^2 \begin{pmatrix} B_\lambda^{-1} M_B B_\sigma^{-1} (M_B B_\sigma^{-1} E_\sigma - M_E) \hat{y} \\ 0 \end{pmatrix} + \begin{pmatrix} [v^{(1)}, \dots, v^{(\hat{\kappa})}] \\ 0 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \vdots \\ \nu_{\hat{\kappa}} \end{pmatrix}.$$


350 Following the same reasoning as in Proposition 3.5 (we omit the intermediate steps),
 351 we get

$$352 \quad \|x - \hat{x}\|_M^2 = O((\lambda - \sigma)^{2(\tau+1)}) + (\lambda - \sigma)^4 \sum_{\ell=\hat{\kappa}+1}^d \left(\frac{\epsilon_\ell - \gamma_\ell (\delta_\ell - \sigma)}{(\delta_\ell - \sigma)^2 (\delta_\ell - \lambda)} \right)^2 + O((\lambda - \sigma)^{\tau+3}). \square$$

353 Proposition 3.6 tells us that a first-order approximation of the resolvent B_λ^{-1}
 354 combined with $\tau = 1$ leads to eigenvalue approximation errors of the order $O((\lambda - \sigma)^4)$
 355 as $\sigma \rightarrow \lambda$. Figure 3.2 plots the approximation error of eigenvalues λ_1 , λ_2 , and λ_3 ,
 356 obtained by Proposition 3.6 for some Dirichlet discretization of the Laplacian operator
 357 in the two-dimensional space. In agreement with Proposition 3.6, the true error curves
 358 follow nicely those of $(\sigma - \lambda_j)^2$ (when $\tau = 0$) and $(\sigma - \lambda_j)^4$ (when $\tau = 1$), respectively.
 359 The error reduction remains quartic even when $\tau \geq 2$. More generally, an increase by
 360 one in the value of τ should be accompanied by the addition of one more term in the
 361 Neumann series approximation of the resolvent B_λ^{-1} if the order of the upper bound
 362 is to be decreased.



357 FIG. 3.2. Approximation error of eigenvalues λ_1 , λ_2 , and λ_3 , obtained by Proposition 3.6 as
 358 $\sigma \in [0, \lambda_3)$. Matrix A is formed by a regular Dirichlet discretization of the Laplacian operator over
 359 a square domain where the grid is partitioned into $p = 4$ subdomains ($M = I$). Left: $\tau = 0$. Right:
 360 $\tau = 1$.



381 FIG. 4.1. The eigenpairs of the matrix $S(\sigma)$ can be exploited for the approximation of more
 382 than one eigenpair of the matrix pencil (A, M) . In this example we have $n_{ev} = 3$ and $\psi_i = \kappa_i = i$.
 383 Two different choices of σ are shown.

371 **4. Computing a large number of eigenpairs.** The technique discussed in
 372 section 3 can be extended to the simultaneous approximation of all n_{ev} sought eigen-
 373 pairs. More specifically, denote the n_{ev} sought eigenvalues located immediately on
 374 the right of a real scalar α by $\lambda_{\kappa_1} \leq \dots \leq \lambda_{\kappa_{n_{ev}}}$, and let eigenvalue λ_{κ_i} be a root of
 375 the eigenvalue curve $\mu_{\psi_i}(\zeta)$, i.e., $\psi_i = \mathbf{arg}\{j \mid \mu_j(\lambda_{\kappa_i}) = 0\}$. The eigenpair associated
 376 with eigenvalue λ_{κ_i} can then be approximated independently of the rest by comput-
 377 ing the eigenpair $(\mu_{\psi_i}(\sigma), y_{\psi_i}(\sigma))$ and considering eigenvector $y_{\psi_i}(\sigma)$ as a zeroth-order
 378 approximation of the eigenvector $y_{\psi_i}(\lambda_{\kappa_i}) \equiv y^{(\kappa_i)}$. Figure 4.1 illustrates an example
 379 where $n_{ev} = 3$ and $\psi_i = \kappa_i = i$. Eigenvalue λ_i is a root of eigenvalue curve $\mu_i(\zeta)$, and
 380 the theory presented in section 3 applies to each eigenvalue curve independently.

384 Assume for the moment that no eigenvalue curves cross each other and let $\sigma = \alpha$.
 385 Since the eigenvalue curves are strictly decreasing, we can infer that the eigenvalues

386 $\lambda_{\kappa_1}, \dots, \lambda_{\kappa_{n_{ev}}}$ are roots of consecutive eigenvalue curves, i.e., $\psi_i = \psi_1 + (i - 1)$. Ad-
 387 ditionally, since $\alpha \leq \lambda_{\kappa_1}$, $\mu_{\psi_1}(\zeta)$ is precisely the eigenvalue curve which crosses the
 388 algebraically smallest nonnegative eigenvalue of $S(\alpha)$. The two above observations
 389 tell us that the eigenvectors $y_{\psi_1}(\alpha), \dots, y_{\psi_{n_{ev}}}(\alpha)$ associated with the n_{ev} algebraically
 390 smallest nonnegative eigenvalues of the matrix $S(\alpha)$ form a zeroth-order approxima-
 391 tion of the vectors $y_{\psi_1}(\lambda_{\kappa_1}) \equiv y^{(\lambda_{\kappa_1})}, \dots, y_{\psi_{n_{ev}}}(\lambda_{\kappa_{n_{ev}}}) \equiv y^{(\lambda_{\kappa_{n_{ev}}})}$. Note that the value
 392 of ψ_1 itself is not needed.

393 In practice, eigenvalue curves which intersect each other pose no threat as long
 394 as each one of the n_{ev} algebraically smallest nonnegative eigenvalues of the matrix
 395 $S(\alpha)$ coincides with one of the values $\mu_{\psi_1}(\alpha), \dots, \mu_{\psi_{n_{ev}}}(\alpha)$. Our default strategy
 396 throughout the rest of this paper is to set $\sigma = \alpha$.

397 **4.1. A Rayleigh–Ritz algorithm.** The theoretical results presented in section
 398 3 focused on the eigenvalue approximation error resulting by a Rayleigh quotient
 399 approximation with an approximate eigenvector \hat{x} . In practice, we cannot form \hat{x} since
 400 we do not know the quantity $\lambda - \sigma$. Nonetheless, we can overcome this drawback by
 401 approximating all n_{ev} eigenvalues from a single subspace by means of a Rayleigh–Ritz
 402 projection.

ALGORITHM 4.1. *The complete procedure*

0. *Input: $A, M, p, \alpha, n_{ev}, \kappa$*
1. *Reorder A and M as in (2.4)*
2. *Call Algorithm 4.2 (subspace associated with interface variables)*
3. *Call Algorithm 4.3 (subspace associated with interior variables)*
4. *Build the projection matrix Z as described in section 4.3.1*
5. *Solve $(Z^T A Z)\tilde{x} = \hat{\lambda}(Z^T M Z)\tilde{x}$ for the n_{ev} sought $(\tilde{\lambda}, \tilde{x})$*
6. *Form the Ritz vectors $\hat{x} = Z\tilde{x}$ associated with the n_{ev} computed eigenpairs from step 5*

404 Algorithm 4.1 describes the complete algorithmic procedure to compute the n_{ev}
 405 eigenvalues located immediately on the right of the user-given scalar $\alpha \in \mathbb{R}$. The
 406 Rayleigh–Ritz eigenvalue problem shown in step 5 must be solved only for the eigen-
 407 pairs associated with the n_{ev} smallest eigenvalues that are greater than or equal to
 408 α . Except for the matrices A and M , and the scalars n_{ev} and α , Algorithm 4.1 also
 409 requires a number of partitions p and a nonzero integer κ denoting the number of
 410 eigenvectors computed from each matrix pencil $(B_\sigma^{(j)}, M_B^{(j)})$, $j = 1, \dots, p$.

411 **4.2. Building the projection subspace associated with interface vari-**
 412 **ables.** Algorithm 4.2 begins by computing the eigenvectors associated with the n_{ev}
 413 smallest nonnegative eigenvalues of the matrix $S(\alpha)$. These can be computed by
 414 any appropriate sparse eigenvalue solver, e.g., the (block) Lanczos method com-
 415 bined with shift-and-invert [13] or (generalized) Davidson [36]. The algorithm pro-
 416 ceeds by computing the derivatives $y'_{\psi_i}(\alpha) \equiv (\frac{dy_j(\zeta)}{d\zeta})_{\zeta=\alpha}, y''_{\psi_i}(\alpha) \equiv (\frac{d^2 y_j(\zeta)}{d\zeta^2})_{\zeta=\alpha}, \dots$,
 417 $i = 1, 2, \dots, n_{ev}$. Details on the practical computation of these derivatives up to a
 418 second order are provided in the appendix.

ALGORITHM 4.2. *Projection subspace associated with interface variables*

0. *Input: α, n_{ev}*
1. *Solve $S(\alpha)y(\alpha) = \mu(\alpha)y(\alpha)$ for the n_{ev} smallest nonnegative eigenvalues $\mu(\alpha)$ and associated eigenvectors $y(\alpha)$*
- . *Denote these eigenpairs as $(\mu_{\psi_i}(\alpha), y_{\psi_i}(\alpha))$, $i = 1, 2, \dots, n_{ev}$*
2. *Form $Y = [y_{\psi_1}(\alpha), y'_{\psi_1}(\alpha), y''_{\psi_1}(\alpha), \dots, y_{\psi_2}(\alpha), y'_{\psi_2}(\alpha), y''_{\psi_2}(\alpha), \dots]$*

420 **4.3. Building the projection subspace associated with interior vari-**
 421 **ables.** Algorithm 4.3 provides a formal description of the procedure followed for
 422 the construction of the projection subspace associated with interior variables. The
 423 routine “`eigs`($B_\alpha^{(j)}, M_B^{(j)}, \kappa, \text{sm}$)” listed in step 2 denotes the computation of the eigen-
 424 vectors associated with the κ smallest (in magnitude) eigenvalues of each matrix pencil
 425 ($B_\alpha^{(j)}, M_B^{(j)}$), $j = 1, \dots, p$, and can be performed by any appropriate sparse eigenvalue
 426 solver. These eigenvectors form the columns of the $d_j \times \kappa$ matrix V_j . The value of κ
 427 can be set either a priori or adaptively, e.g., see the related discussion in [40].

ALGORITHM 4.3. *Projection subspace associated with interior variables*

0. *Input:* α, p, κ, Y := orthonormal basis returned by Algorithm 4.2

1a. *For* $j = 1, \dots, p$

2. Compute $V_j = \text{eigs}(B_\alpha^{(j)}, M_B^{(j)}, \kappa, \text{sm})$

1b. *End*

3. *Form* $U = [\bar{U}^{(1)}, \dots, \bar{U}^{(p)}, -B_\alpha^{-1}E_\alpha Y, -B_\alpha^{-1}M_B B_\alpha^{-1}E_\alpha Y, B_\alpha^{-1}M_E Y]$

where $\bar{U}^{(j)} = \begin{pmatrix} 0_{\ell_j, \kappa} \\ V_j \\ 0_{\xi_j, \kappa} \end{pmatrix}$ and we recall $\ell_j = \sum_{k=1}^{k < j} s_k$, and $\xi_j = \sum_{k > j} s_k$

429 **4.3.1. The Rayleigh–Ritz eigenvalue problem.** The matrix Y returned by
 430 Algorithm 4.2 is distributed among the p subdomains and can be written as

$$(4.1) \quad Y = \begin{bmatrix} Y_1^T & Y_2^T & \cdots & Y_p^T \end{bmatrix}^T,$$

433 where $Y_j \in \mathbb{R}^{s_j \times \eta}$ is the row block of matrix Y associated with the j th subdomain
 434 and $\eta \in \mathbb{Z}^*$ denotes the column dimension of matrix Y . By definition η is equal to an
 435 integer multiple of n_{ev} .

436 Define now the matrices

$$\left(\hat{B}_\alpha^{(j)}\right)^{-1} = \left(B_\alpha^{(j)}\right)^{-1} \left(I - V_j V_j^T M_B^{(j)}\right) \text{ and}$$

$$P_j = \begin{bmatrix} E_\alpha^{(j)}, -M_E^{(j)} \end{bmatrix} \begin{pmatrix} Y_j & 0 \\ 0 & Y_j \end{pmatrix}.$$

439 The projection matrix Z can be then written as

$$Z = \begin{bmatrix} V_1 & & -\left(\hat{B}_\alpha^{(1)}\right)^{-1} P_1 & -\left(B_\alpha^{(1)} \left(M_B^{(1)}\right)^{-1} \hat{B}_\alpha^{(1)}\right)^{-1} E_\alpha^{(1)} Y_1 \\ & V_2 & -\left(\hat{B}_\alpha^{(2)}\right)^{-1} P_2 & -\left(B_\alpha^{(2)} \left(M_B^{(2)}\right)^{-1} \hat{B}_\alpha^{(2)}\right)^{-1} E_\alpha^{(2)} Y_2 \\ & & \vdots & \vdots \\ & & & -\left(\hat{B}_\alpha^{(p)}\right)^{-1} P_p & -\left(B_\alpha^{(p)} \left(M_B^{(p)}\right)^{-1} \hat{B}_\alpha^{(p)}\right)^{-1} E_\alpha^{(p)} Y_p \\ & & & & [Y, 0_{s, \eta}] \end{bmatrix}.$$

443 The total memory overhead associated with the j th subdomain is equal to that of
 444 storing $\kappa d_j + (3d_j + s_j)\eta$ floating-point scalars. The dimension of the Rayleigh–Ritz
 445 pencil ($Z^T A Z, Z^T M Z$) is equal to $\kappa p + 3\eta$ and can be solved by the appropriate routine
 446 in LAPACK, e.g., `dsygv` [3]. As a sidenote, when $M_E^{(j)} = 0$ we have $P_j = E_\alpha^{(j)} Y_j$ and
 447 the bottom row block of matrix Z becomes $[0_{s, p\kappa}, Y, 0_{s, \eta}]$. The dimension of the
 448 Rayleigh–Ritz pencil then reduces to $\kappa p + 2\eta$.

449 **4.4. Comparing Algorithm 4.1 with shift-and-invert Lanczos.** A natural
 450 question is how Algorithm 4.1 compares against shift-and-invert Lanczos when the
 451 latter is applied directly to the pencil $(A - \alpha M, M)$.

452 The first key difference between these two techniques is *orthogonalization cost*.
 453 Applying k steps of shift-and-invert Lanczos to matrix pencils $(S(\alpha), I)$ and $(A -$
 454 $\alpha M, M)$ leads to a total orthogonalization cost of $O(k^2 s)$ and $O(k^2 n)$, respectively.
 455 Thus, Algorithm 4.1 reduces orthogonalization costs by a factor of n/s , and this
 456 difference becomes more pronounced as n_{ev} increases (since $k \geq n_{ev}$).

457 The second key difference between Algorithm 4.1 and shift-and-invert Lanczos
 458 is *the number of linear system solutions with B_α as the coefficient matrix*. It is
 459 straightforward to verify that applying k steps of shift-and-invert Lanczos to the
 460 pencil $(A - \alpha M, M)$ requires $2k$ such linear system solutions. In contrast, Algorithm
 461 4.1 requires 3η (2η if $M_E = 0$) linear solves with B_α . Nonetheless, these 3η linear
 462 solves can be performed simultaneously, since all right-hand sides are available at the
 463 same time. Thus, the associated cost can be marginally higher than that of solving
 464 for a few right-hand sides. To the above cost we also need to add the computational
 465 cost required to compute the eigenvectors of the block-diagonal pencil (B_α, M_B) in
 466 Algorithm 4.3.

467 On the other hand, the accuracy achieved by shift-and-invert Lanczos can be
 468 significantly higher than that of Algorithm 4.1. Nonetheless, in many applications
 469 the sought eigenpairs need not be computed to high accuracy, e.g., eigenpairs of
 470 problems originating from discretizations need be approximated up to the associated
 471 discretization error.

472 **4.5. Parallel computing.** Algorithm 4.1 is based on domain decomposition
 473 and is well-suited for execution in modern distributed memory environments. For
 474 example, each one of the p subdomains can be mapped to a separate MPI process.
 475 Performing any type of operation with the matrices $B_\sigma^{(j)}$, $E_\sigma^{(j)}$, and $M_E^{(j)}$ is then
 476 an entirely local process in the j th subdomain (i.e., Algorithm 4.3 is embarrassingly
 477 parallel). An additional layer of parallelism can be realized in each subdomain by
 478 further exploiting shared memory parallelism to perform the required computations
 479 with the aforementioned matrices; e.g., see [24] for a similar discussion in the context
 480 of domain decomposition eigenvalue solvers and [15] for a general discussion on parallel
 481 computing and domain decomposition.

482 In contrast to Algorithm 4.3, Algorithm 4.2 involves computations with the Schur
 483 complement matrix $S(\alpha)$, which is distributed among the p subdomains. In this case,
 484 point-to-point communication among neighboring subdomains is necessary. Finally,
 485 the Rayleigh–Ritz eigenvalue problem is typically small enough so that it can be
 486 replicated in all MPI processes and solved redundantly.

487 **5. Numerical experiments.** Our sequential experiments are conducted in a
 488 MATLAB environment (version R2018b), using 64-bit arithmetic, on a single core of
 489 a computing system equipped with an Intel Haswell E5-2680v3 processor and 32 GB
 490 of system memory.

491 The eigenvalues of the pencil (A, M) are ordered as $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, and
 492 throughout the rest of this section we focus in computing the n_{ev} algebraically smallest
 493 eigenvalues $\lambda_1, \dots, \lambda_{n_{ev}}$. Unless mentioned otherwise, the default values in Algorithm
 494 4.1 will be set as $p = 8$, $\kappa = 5$, and $\alpha = 0$. For indefinite pencils we first shift the
 495 spectrum so that all eigenvalues become positive, and then apply Algorithm 4.1 with
 496 $\alpha = 0$. Throughout the rest of this paper we assume $\psi_i = i$, $i = 1, \dots, n_{ev}$.

We consider three different choices to set $\mathbf{span}(Y)$ in Algorithm 4.2:

$$\{y\} := \mathbf{span} \left([y_i(\alpha)]_{i=1, \dots, 3n_{ev}} \right),$$

$$\{y, dy\} := \mathbf{span} \left(\left[y_i(\alpha), \left(\frac{dy_i(\zeta)}{d\zeta} \right)_{\zeta=\alpha} \right]_{i=1, \dots, n_{ev}} \right),$$

$$\{y, dy, d^2y\} := \mathbf{span} \left(\left[y_i(\alpha), \left(\frac{dy_i(\zeta)}{d\zeta} \right)_{\zeta=\alpha}, \left(\frac{d^2y_i(\zeta)}{d\zeta^2} \right)_{\zeta=\alpha} \right]_{i=1, \dots, n_{ev}} \right).$$

The eigenpairs of matrix $S(\alpha)$ are computed up to a residual norm of 1.0×10^{-6} .

5.1. A model problem. Our first test case consists of a five-point stencil finite difference discretization of the Dirichlet eigenvalue problem

$$(5.1) \quad \Delta u + \lambda u = 0 \text{ in } \Omega := (0, 1) \times (0, 1), \quad u|_{\partial\Omega} = 0,$$

where Δ denotes the Laplace operator. Our goal is not to compute the actual eigenvalues of the Laplace operator but rather to assess the performance of Algorithm 4.1. To this end, the Dirichlet eigenvalue problem is discretized on a 250×250 mesh, and we set $n_{ev} = 150$. Note that A has many eigenvalues of multiplicity $\rho_\lambda = 2$. The proposed method can naturally capture multiple eigenvalues in contrast to (nonblock) Krylov-based approaches such as shift-and-invert Lanczos.

Figure 5.1 plots the relative eigenvalue errors and associated residual norms of the approximate eigenpairs returned by Algorithm 4.1 when $U = [u^{(1)}, \dots, u^{(n_{ev})}]$, i.e., there is no error associated with interior variables. In agreement with the discussion in section 3, adding more eigenvector derivatives leads to higher accuracy, especially for those eigenvalues located closest to α . Table 5.1 lists the maximum/minimum

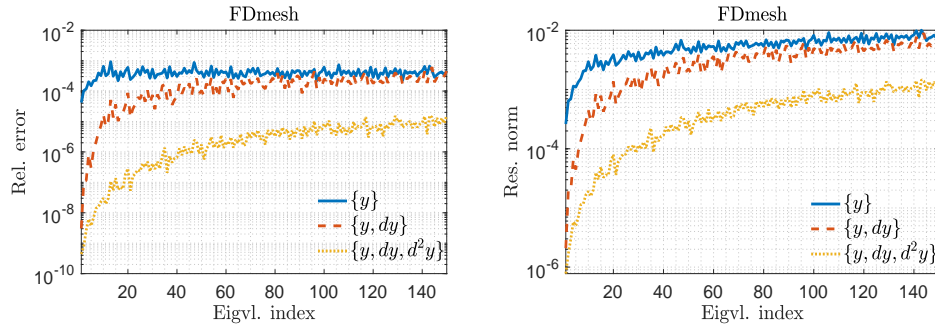



FIG. 5.1. Relative eigenvalue errors (left) and residual norms (right) returned by Algorithm 4.1 in the approximation of the $n_{ev} = 150$ algebraically smallest eigenvalues of the 250×250 finite difference discretization of the Dirichlet eigenvalue problem when $U = [u^{(1)}, \dots, u^{(n_{ev})}]$.

TABLE 5.1

Maximum/minimum relative eigenvalue error achieved by Algorithm 4.1 in the approximation of the $n_{ev} = 50$ algebraically smallest eigenvalues of the 250×250 finite difference discretization of the Dirichlet eigenvalue problem for $U = [u^{(1)}, \dots, u^{(n_{ev})}]$ and various values of p .

	$\{y\}$		$\{y, dy\}$		$\{y, dy, d^2y\}$	
	max	min	max	min	max	min
$p = 2$	1.1×10^{-6}	9.0×10^{-9}	5.7×10^{-7}	5.2×10^{-12}	2.9×10^{-9}	1.7×10^{-15}
$p = 4$	7.5×10^{-6}	2.8×10^{-8}	8.0×10^{-6}	1.9×10^{-11}	4.5×10^{-7}	1.6×10^{-13}
$p = 8$	4.7×10^{-5}	7.2×10^{-8}	8.4×10^{-5}	2.0×10^{-11}	6.7×10^{-6}	2.8×10^{-12}
$p = 16$	1.5×10^{-4}	1.1×10^{-7}	3.4×10^{-4}	5.8×10^{-11}	5.3×10^{-5}	9.4×10^{-12}



532 FIG. 5.2. Plots of the relative eigenvalue errors (left) and residual norms (right) returned by
 533 Algorithm 4.1 in the approximation of the $n_{ev} = 150$ algebraically smallest eigenvalues of the $250 \times$
 534 250 finite difference discretization of the Dirichlet eigenvalue problem when $Y = [y^{(1)}, \dots, y^{(n_{ev})}]$,
 535 and κ varies.

TABLE 5.2

536 n : size of (A, M) ; s : number of interface variables ($p = 8$); $nnz(\cdot)$: number of nonzero entries.
 537

#	Matrix pencil	n	$nnz(A)/n$	$nnz(M)/n$	s	Application	Source
1.	Si2	769	23.1	1.0	547	Quantum Chemistry	[9]
2.	nos3	960	16.5	1.0	170	Structural	[9]
3.	FEmesh	2,689	6.9	6.9	190	Finite Element	-
4.	VCNT_4000	4,000	40.0	1.0	640	Quantum Mechanics	[19]
5.	Kuu/Muu	7,102	47.9	24.0	488	Structural	[9]
6.	fv1	9,604	8.9	1.0	453	2D/3D	[9]
7.	FDmesh	62,500	5.0	1.0	1,032	2D	-
8.	qa8fk/qa8fm	66,172	25.1	25.1	5,270	3D Acoustic	[9]

525 relative eigenvalue error for the same problem where $n_{ev} = 50$ and p varies. In
 526 summary, increasing the number of interface variables leads to lower accuracy unless
 527 more eigenvector derivatives are computed.

528 Figure 5.2 considers the opposite scenario where U is set as in Algorithm 4.3
 529 and $Y = [y^{(1)}, \dots, y^{(n_{ev})}]$. Here, the asymptotic order of the approximation error is
 530 fixed (i.e., quartic), and increasing the value of κ reduces the upper bound of the
 531 approximation error.

538 **5.2. General pencils.** Throughout the rest of this section we assume that ma-
 539 trices U and Y are set as described in Algorithm 4.1. Details on the numerical
 540 approximation of the first and second derivatives of each computed eigenvector of
 541 matrix $S(\alpha)$ are given in the appendix.⁵

542 We consider the application of Algorithm 4.1 on a set of model problems and
 543 matrix pencils obtained by the SuiteSparse⁶ Matrix Collection [9], and the Elses⁷
 544 matrix library [19]. Details can be found in Table 5.2. The matrix pencil FEmesh
 545 represents a finite elements discretization of the Dirichlet eigenvalue problem on a
 546 $[-1, 1] \times [-1, 1]$ plane where the Laplacian is discretized using linear elements with
 547 target maximum mesh edge length of $h = 0.05$. The matrix FDmesh represents the

⁵The linear system solver we choose is preconditioned MINRES with a stopping tolerance of 1.0×10^{-3} and a maximum number of five preconditioned iterations, whichever occurs first. The convergence criterion is applied on the unpreconditioned residuals. For preconditioning we use matrix $S(\alpha)$ combined with deflation.

⁶<https://sparse.tamu.edu/>.

⁷<http://www.elses.jp/matrix/>.

551
552
553

TABLE 5.3

Maximum relative error of the approximation of the $n_{ev} = 50$ algebraically smallest eigenvalues returned by Algorithm 4.1 as κ varies.

	$\{y\}$		$\{y, dy\}$		$\{y, dy, d^2y\}$		
	$\kappa = 20$	$\kappa = 10$	$\kappa = 20$	$\kappa = 40$	$\kappa = 10$	$\kappa = 20$	$\kappa = 40$
Si2	5.0×10^{-4}	1.4×10^{-3}	9.4×10^{-4}	6.1×10^{-4}	7.7×10^{-4}	1.2×10^{-4}	2.5×10^{-5}
nos3	3.4×10^{-3}	1.4×10^{-3}	5.4×10^{-4}	2.7×10^{-4}	4.8×10^{-3}	3.6×10^{-4}	9.9×10^{-5}
FEmesh	1.7×10^{-3}	2.1×10^{-3}	1.4×10^{-3}	8.9×10^{-4}	9.1×10^{-4}	2.4×10^{-4}	1.1×10^{-4}
VCNT_4000	4.2×10^{-3}	3.6×10^{-3}	2.6×10^{-3}	9.0×10^{-4}	6.7×10^{-3}	9.4×10^{-4}	2.7×10^{-4}
{K,M}uu	7.3×10^{-3}	1.4×10^{-2}	9.6×10^{-3}	7.1×10^{-3}	2.6×10^{-3}	4.8×10^{-4}	1.5×10^{-4}
fv1	3.4×10^{-3}	1.4×10^{-2}	7.4×10^{-4}	3.0×10^{-4}	6.7×10^{-3}	4.1×10^{-4}	3.1×10^{-5}
FDmesh	1.9×10^{-3}	1.9×10^{-3}	1.7×10^{-3}	9.4×10^{-4}	2.4×10^{-3}	9.1×10^{-3}	2.4×10^{-4}
qa8f{k,m}	4.7×10^{-3}	9.0×10^{-3}	5.5×10^{-3}	4.2×10^{-3}	6.1×10^{-3}	2.2×10^{-3}	9.5×10^{-4}

548 250×250 finite difference discretization of the Dirichlet eigenvalue problem discussed
549 in the previous section. With the exception of the matrix pencils FEmesh, Kuu/Muu,
550 and qa8fk/qa8fm, all other pencils are of standard form.

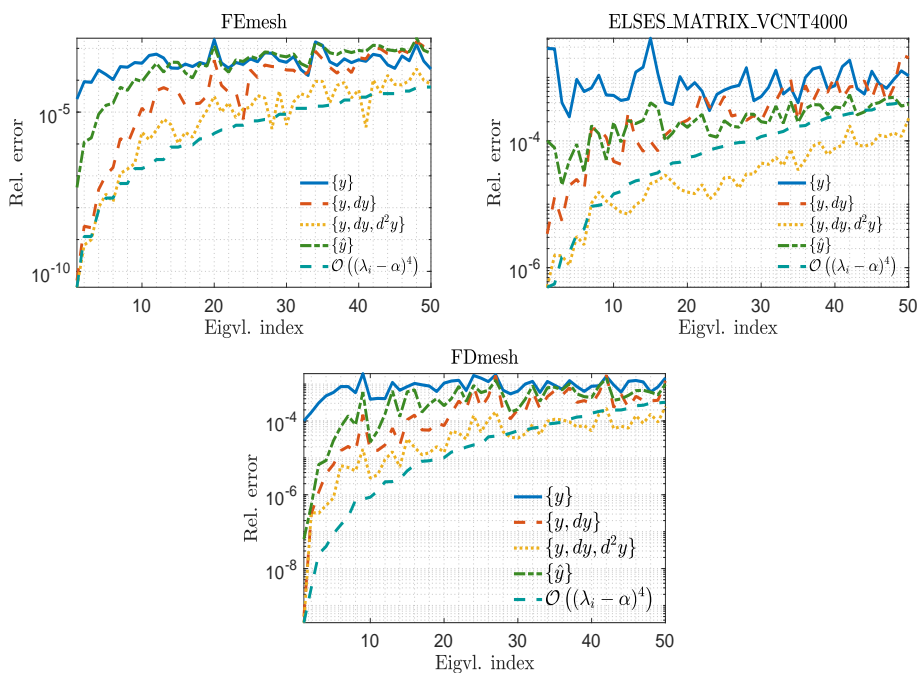
551 Table 5.3 lists the maximum relative error in the approximation of the $n_{ev} = 50$
552 algebraically smallest eigenvalues returned by Algorithm 4.1 for all pencils listed in
553 Table 5.2. Naturally, increasing κ enhances overall accuracy, and this enhancement
554 becomes greater as the interface projection subspace also improves. Note though that
555 the entries associated with the choices $\{y, dy\}$ and $\{y, dy, d^2y\}$ are now closer to each
556 other than what was shown in section 5.1. This is owed to (a) the inexact computation
557 of the first and second derivatives and (b) the fact that the asymptotic order of the
558 error is the same in both cases since only a first-order approximation of the resolvent
559 B_λ^{-1} is considered.

560 Finally, we compare the accuracy achieved by Algorithm 4.1 against that of our
561 own variant of the AMLS method (termed “ p -way AMLS” and denoted by $\{\hat{y}\}$). For
562 reasons of fairness and easiness in the interpretation of our comparisons, our own
563 variant of AMLS is identical to Algorithm 4.1 except that matrix Y is formed as
564 $Y = [\hat{y}^{(1)}, \dots, \hat{y}^{(3n_{ev})}]$ where $\hat{y}^{(i)} \in \mathbb{R}^s$ denotes the eigenvector associated with the i th
565 algebraically smallest eigenvalue of the pencil $(S(\alpha), -S'(\alpha))$. Our variant of AMLS
566 is more accurate (but slower) than standard AMLS described in [7, 10].

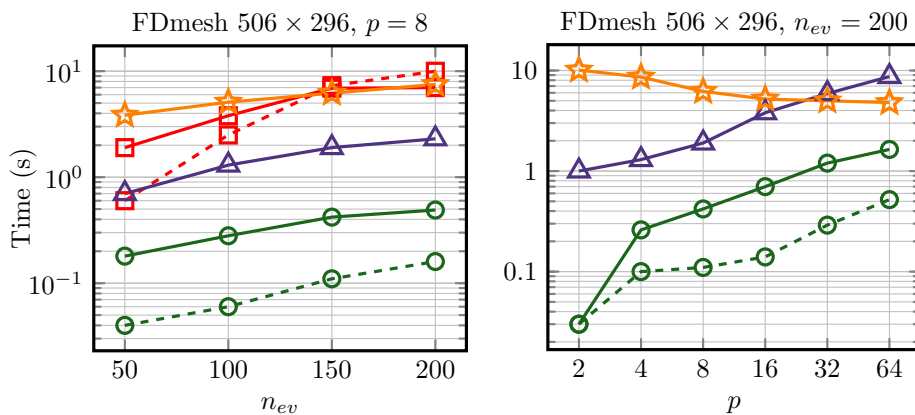
567 Figure 5.3 plots the relative error in the approximation of the $n_{ev} = 50$ algebrai-
568 cally smallest eigenvalues returned by Algorithm 4.1 and p -way AMLS when applied
569 on a subset of the pencils listed in Table 5.2. In summary, p -way AMLS is more accu-
570 rate than the $\{y\}$ variant of Algorithm 4.1 but not as good as the variants $\{y, dy\}$ and
571 $\{y, dy, d^2y\}$, especially for those eigenvalues located closer to α . This performance
572 gap increases favorably for Algorithm 4.1 as the projection subspace associated with
573 interior variables improves.

574 **5.3. Comparisons against shift-and-invert Lanczos.** This section presents
575 wall-clock time comparisons between Algorithm 4.1 (variant $\{y, dy, d^2y\}$) and implic-
576 itly restarted shift-and-invert Lanczos with full orthogonalization applied directly to
577 the pencil (A, M) . We will refer to the latter as IRSIL. The maximum dimension
578 of the Krylov subspace was set to $2n_{ev}$. As our test matrix we choose a five-point
579 506×296 finite difference discretization of the Dirichlet eigenvalue problem.

580 Figure 5.4 plots sequential wall-clock times of the individual steps of IRSIL and
581 Algorithm 4.1 when exploiting both schemes to approximate the n_{ev} algebraically
582 smallest eigenvalues of the discretized Laplacian (left subfigure). In total, IRSIL re-
583 quired 2.5, 6.3, 14.2, and 17.0 seconds to approximate the $n_{ev} = 50, 100, 150,$ and
584



570 FIG. 5.3. Relative error of the approximation of the $n_{ev} = 50$ algebraically smallest eigenvalues
 571 returned by Algorithm 4.1 and p -way AMLS ($\{\hat{y}\}$). We also plot the curve $(\lambda - \alpha)^4$ adjusted so that
 572 its first entry is equal to that of the curve $\{y, dy, d^2 y\}$.



573 FIG. 5.4. Sequential wall-clock times of IRSIL (applied to the pencil $(A - \alpha M, M)$) and Al-
 574 gorithm 4.1 (variant $\{y, dy, d^2 y\}$) for various values of n_{ev} . For IRSIL, we report the amount of
 575 time spent on applying $(A - \alpha M)^{-1}$ (solid) and orthogonalization (dashed) separately (\square). For
 576 Algorithm 4.1 we report the amount of time spent on (a) computing eigenvectors $y_1(\alpha), \dots, y_{n_{ev}}(\alpha)$
 577 (o), (b) approximating the two leading eigenvector derivatives by pseudoblock MINRES (\triangle), and
 578 (c) executing Algorithm 4.3 with $\kappa = 40$ (\star). Steps (a) and (b) form Algorithm 4.2. Notice that for
 579 step (a) we report the amount of time spent on applying the operator (in this case $S(\alpha)^{-1}$) (solid)
 580 separately from that spent on orthogonalization (dashed). Left: fix $p = 8$ and vary n_{ev} . Right: fix
 581 $n_{ev} = 200$ and vary p .

599 $n_{ev} = 200$, algebraically smallest eigenvalues and associated eigenvectors up to a resid-
 600 ual norm 1.0×10^{-8} . On the other hand, Algorithm 4.1 with the default choice $p = 8$
 601 required 4.9, 6.7, 8.7, and 10.5 seconds, respectively. As n_{ev} increases, IRSIL becomes
 602 increasingly slower than Algorithm 4.1 due to its increasing orthogonalization cost.
 603 For example, when $n_{ev} = 200$, IRSIL spent about 10 seconds in orthogonalization,
 604 while Algorithm 4.1 only required about a quarter of a second. The maximum ei-
 605 genvalue error returned by Algorithm 4.1 for any value of n_{ev} tested was $O(10^{-4})$.
 606 Figure 5.4 also plots the sequential execution timing of Algorithm 4.1 as p varies and
 607 $n_{ev} = 200$ is fixed (right subfigure). Increasing the value of p leads to a greater number
 608 of interface variables (and thus increased orthogonalization costs), while solving linear
 609 systems with matrix $S(\alpha)$ also becomes more expensive. On the other hand, larger
 610 values of p generally decrease the amount of time spent in Algorithm 4.3. Note that
 611 for $p = 64$ the maximum error returned by Algorithm 4.1 was of the order $O(10^{-3})$.

612 Table 5.4 lists the total number of iterations required by implicitly restarted
 613 Lanczos to (a) compute the sought eigenvectors of matrix $S(\alpha)$ and (b) solve the
 614 original eigenvalue problem (i.e., compute the n_{ev} sought eigenpairs of (A, M)). For
 615 Algorithm 4.1 we considered setting $p = 4, 16$, and $p = 32$, respectively. In summary,
 616 we observe two patterns. First, as p increases so does the number of iterations required
 617 by Lanczos. In particular, increasing the size of the Schur complement matrix typically
 618 leads to a denser spectrum in matrix $S(\alpha)$ since the eigenvalues of the latter interlace
 619 those of $(A - \alpha M, M)$. This can be verified in Figure 5.5, where we plot the 250
 620 algebraically smallest eigenvalues of matrices FDmesh 506×296 and $S(\alpha \equiv 0)$ as
 621 $p = 4, 16$, and $p = 32$, respectively. Second, working with the Schur complement
 622 matrix becomes the only practical approach if shift-and-invert is not possible, since
 623 applying implicitly restarted Lanczos to (A, M) can lead to very slow convergence
 624 and thus high orthogonalization costs.


625 A preliminary distributed memory implementation of Algorithm 4.1 was built
 626 on top of the PETSc library [5].⁸ Message passing between different processes was
 627 achieved by means of the Message Passing Interface (MPI), a standard application
 628 programming interface for message passing applications [14]. We considered only
 629 one thread per MPI process, and the number of MPI processes will be equal to the
 630 number of subdomains p . To allow for a fair comparison we implemented our own ver-
 631 sion of IRSIL instead of that in the PETSc-based, state-of-the-art eigenvalue package
 632

633 TABLE 5.4

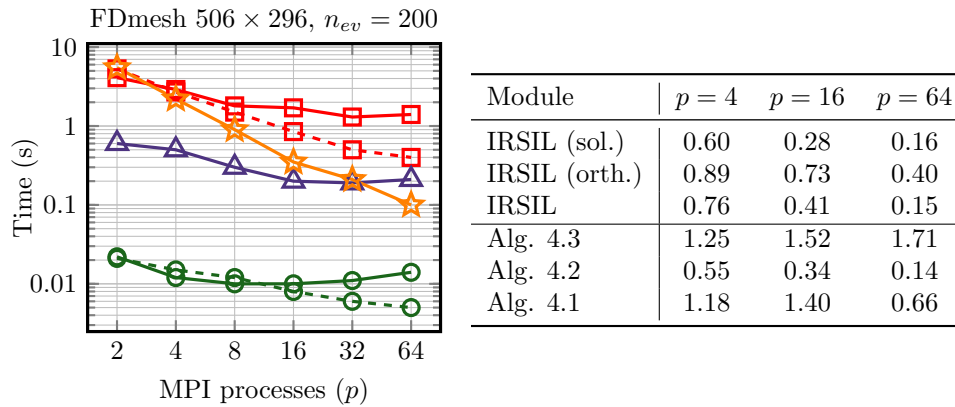
634 *Total number of iterations performed by implicitly restarted Lanczos to compute the sought*
 635 *eigenvectors of (a) matrix $S(\alpha)$ and (b) the original eigenvalue problem. Flag **sa** (**sm**) indicates the*
 636 *absence (presence) of shift-and-invert acceleration. An **F** flag indicates that not all eigenvectors*
 637 *were computed after 20 restarts, where the maximum Krylov subspace dimension was set equal to*
 638 *$2n_{ev}$.*

	$n_{ev} = 100$				$n_{ev} = 200$			
	Alg. 4.1: $p = 4$	$p = 16$	$p = 32$	(A, I)	Alg. 4.1: $p = 4$	$p = 16$	$p = 32$	(A, I)
"sa"	605	1,027	1,289	F	775	1,440	1,723	F
"sm"	304	310	315	317	400	440	472	553

⁸Our code builds on top of the implementation featured in [24] and was compiled using real arithmetic. The source files were compiled with the Intel MPI compiler `mpicc`, using the `-O3` optimization level. The linear system solutions with the distributed matrices $A - \alpha M$ and $S(\alpha)$ were computed by the Multifrontal Massively Parallel Sparse Direct Solver [2] and those with the block-diagonal matrix B_α by MKL PARDISO [1].



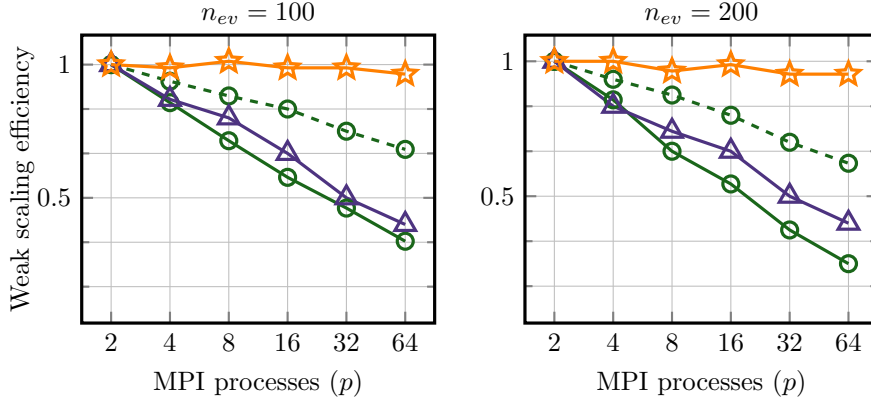
639 FIG. 5.5. Plot of the 250 algebraically smallest eigenvalues of matrix $\text{FDmesh } 506 \times 296$ (left),
 640 and matrix $S(\alpha \equiv 0)$ setting $p = 4, 16, 32$ (right). On the right subfigure, the top, middle, and
 641 bottom curves indicate the choice $p = 4, 16$, and $p = 32$, respectively. The solid part of each curve
 642 indicates the algebraically smallest $n_{ev} = 200$ eigenvalues of each matrix.



643 FIG. 5.6. Left: Distributed memory wall-clock times (strong scaling). Right: Ratio of true
 644 speedup to theoretical speedup (efficiency) computed as $T_s/(pT_p)$, where T_p denotes the wall-clock
 645 time using p MPI processes. For IRSIL we set T_s equal to its sequential wall-clock time shown in
 646 Figure 5.4. For Algorithm 4.1 and its individual steps we set $T_s = 2T_2$.

647 SLEPc [17]. Our experiments were performed at the Mesabi cluster of the Minnesota
 648 Supercomputing Institute (<https://www.msi.umn.edu/>).

649 Figure 5.6 plots the distributed memory wall-clock times of different modules of
 650 IRSIL and Algorithm 4.1 as $n_{ev} = 200$ and $p = 2, 4, 8, 16, 32$, and $p = 64$. The corre-
 651 sponding efficiency for some of these values of p is listed in the accompanying table.
 652 Algorithm 4.3 is the most scalable operation of Algorithm 4.1 and in this example
 653 provides superlinear speedups. As p increases, solving the distributed linear systems
 654 in Lanczos becomes the least scalable operation for both algorithms. Nonetheless, in-
 655 creasing p can still lead to a higher efficiency in Algorithm 4.1 up to a point where the
 656 amount of time spent in Algorithm 4.3 is small and thus the efficiency of the former
 657 is mainly determined by the efficiency of performing computations with the Schur
 658 complement matrix. This is the reason the efficiency of Algorithm 4.1 drops sharply
 659 from $p = 16$ to $p = 64$. This implies that increasing parallelism through increasing p
 660 is nonoptimal, and additional parallel resources should be exploited in a hierarchical
 661 fashion. In total, Algorithm 4.1 is about $3 \times (5 \times)$ faster than IRSIL when 16 (64)
 662 MPI processes are used.



668 FIG. 5.7. Weak scaling efficiency (computed as T_2/T_p where T_p denotes the wall-clock time of
 669 Algorithm 4.1 for $p = 2, 4, 8, 16, 32, 64$) of different modules of Algorithm 4.1 as p varies. The mesh
 670 size of the discretized Laplacian scaled as $\sqrt{p}50 \times \sqrt{p}50$. Left: $n_{ev} = 100$. Right: $n_{ev} = 200$.

663 Finally, Figure 5.7 plots the weak scaling efficiency of different modules of Algo-
 664 rithm 4.1 for a Laplacian discretized on a square mesh with mesh size $\sqrt{p}50 \times \sqrt{p}50$.
 665 The number of sought eigenpairs was set to $n_{ev} = 100$, and $n_{ev} = 200$. Again we ob-
 666 serve that the least scalable parts are those enabling the solution of distributed linear
 667 systems. Moreover, larger values of n_{ev} are more challenging in terms of scalability.

671 **6. Summary and future work.** This paper presented a domain decompo-
 672 sition technique for the computation of a few eigenpairs of symmetric generalized
 673 eigenvalue problems. The proposed technique is based on Rayleigh–Ritz projections
 674 onto subspaces formed by decoupling the original eigenvalue problem into two dis-
 675 tinct subproblems, each one associated with the interface and interior variables of the
 676 domain, respectively. The part of the subspace associated with the interface vari-
 677 ables of the global domain is formed by the eigenvectors and associated derivatives
 678 of a zeroth-order approximation of the nonlinear interface matrix-valued operator.
 679 On the other hand, the part of the subspace associated with the interior variables
 680 is formed in parallel among the different subdomains by exploiting local eigenmodes
 681 and approximations of resolvent expansions.

682 Future work includes a study on the effects which the number of interface variables
 683 has in the accuracy of the technique proposed in this paper. A parallel implementa-
 684 tion with additional levels of parallelism (both distributed and shared memory), possi-
 685 bly combined with harmonic Rayleigh–Ritz projections, is in our short-term plans.

686 **Appendix A. Analytical formulas.** In this section we present analytical
 687 formulas for the computation of the first two derivatives of the eigenpairs $(\mu_i(\sigma), y_i(\sigma))$
 688 of the matrix $S(\sigma)$, $\sigma \in \mathbb{R}$.

689 PROPOSITION A.1. *The first and second derivatives of the matrix $S(\sigma)$, $\sigma \in \mathbb{R}$,
 690 are given by*

$$691 S'(\sigma) = \frac{dS(\sigma)}{d\sigma} = -M_C - E_\sigma^T B_\sigma^{-1} M_B B_\sigma^{-1} E_\sigma + E_\sigma^T B_\sigma^{-1} M_E + M_E^T B_\sigma^{-1} E_\sigma$$

693 and

$$S''(\sigma) = \frac{d^2 S(\sigma)}{d\sigma^2} = 2 \left(E_\sigma^T B_\sigma^{-1} M_B B_\sigma^{-1} M_B B_\sigma^{-1} E_\sigma + M_E^T B_\sigma^{-1} M_E \right) \\ - 2 \left(M_E^T B_\sigma^{-1} M_B B_\sigma^{-1} E_\sigma + E_\sigma^T B_\sigma^{-1} M_B B_\sigma^{-1} M_E \right),$$

respectively. The first and second derivatives of each eigenvalue curve $\mu_i(\sigma)$ are given by

$$\mu'_i(\sigma) = \frac{d\mu_i(\sigma)}{d\sigma} = \frac{y_i^T(\sigma) S'(\sigma) y_i(\sigma)}{y_i^T(\sigma) y_i(\sigma)}$$

and

$$\mu''_i(\sigma) = \frac{d^2 \mu_i(\sigma)}{d\sigma^2} = \frac{y_i^T(\sigma) S''(\sigma) y_i(\sigma) + 2 y_i^T(\sigma) S'(\sigma) y'_i(\sigma)}{y_i^T(\sigma) y_i(\sigma)},$$

respectively. Finally, the first and second derivatives of each eigenvector $y_i(\sigma)$ satisfy the equations

$$(A.1) \quad (S(\sigma) - \mu_i(\sigma)I) y'_i(\sigma) = (\mu'_i(\sigma)I - S'(\sigma)) y_i(\sigma)$$

and

$$(A.2) \quad (S(\sigma) - \mu_i(\sigma)I) y''_i(\sigma) = \left[(\mu''_i(\sigma)I - S''(\sigma)) y_i(\sigma) + 2 (\mu'_i(\sigma)I - S'(\sigma)) y'_i(\sigma) \right],$$

respectively, where $y'_i(\sigma) = \left(\frac{dy_i(\zeta)}{d\zeta} \right)_{\zeta=\sigma}$ and $y''_i(\sigma) = \left(\frac{d^2 y_i(\zeta)}{d\zeta^2} \right)_{\zeta=\sigma}$.

Differentiating the normalization condition $y_i^T(\sigma) y_i(\sigma) = 1$ gives $y_i^T(\sigma) y'_i(\sigma) = 0$, and thus the leading derivative of the eigenvector $y_i(\sigma)$ can be computed by solving the linear system in (A.1). On the other hand, solving the linear system in (A.2) will only provide the second derivative up to the direction $y_i(\sigma)$ (note that $y_i^T(\sigma) y''_i(\sigma) = -\|y'_i(\sigma)\|_2^2$). Nonetheless, the latter eigenvector direction already exists in the subspace $\text{span}(Y)$ (Y is defined in Algorithm 4.2). Throughout the rest of this paper, when we refer to “ $y''_i(\sigma)$ ” it should be understood that we actually refer to the solution of the linear system in (A.2).

Remark 2. When M is equal to the identity matrix, the first and second derivatives of the matrix-valued function $S(\zeta)$ evaluated at σ simplify to the block-diagonal matrices $S'(\sigma) = -I - E_\sigma^T B_\sigma^{-2} E_\sigma$, and $S''(\sigma) = -E_\sigma^T B_\sigma^{-3} E_\sigma$.

Appendix B. Computation of eigenvector derivatives. The computation of the first and second derivatives of each eigenvector $y_i(\sigma)$, $i = 1, 2, \dots, s$, requires the application of an iterative solver, e.g., the minimum residual (MINRES) Krylov subspace method [8, 32], to the solution of the singular linear system

$$(B.1) \quad (S(\sigma) - \mu_i(\sigma)I) x = b^{(i)},$$

where $b^{(i)} \in \mathbb{R}^s$ is defined as

$$b^{(i)} := \begin{cases} (\mu'_i(\sigma)I - S'(\sigma)) y_i(\sigma) & \text{(to compute } y'_i(\sigma)\text{),} \\ (\mu''_i(\sigma)I - S''(\sigma)) y_i(\sigma) + 2 (\mu'_i(\sigma)I - S'(\sigma)) y'_i(\sigma) & \text{(to compute } y''_i(\sigma)\text{).} \end{cases}$$

The eigenvalues of the matrix $S(\sigma) - \mu_i(\sigma)I$ are equal to $\{\mu_k(\sigma) - \mu_i(\sigma)\}_{k=1,2,\dots,s}$.

Let the n_{ev} computed eigenvalues of $S(\sigma)$ be indexed as $\mu_{\psi_1}(\sigma), \dots, \mu_{\psi_{n_{ev}}}(\sigma)$, where $\psi_1 \leq i \leq \psi_{n_{ev}}$. We can enhance the convergence rate of MINRES applied to

734 (B.1) by explicitly removing the components along the directions associated with the
 735 computed eigenvectors of matrix $S(\sigma)$. In particular, the solution of the linear system
 736 with each matrix $S(\sigma) - \mu_i(\sigma)I$ is split into two phases. During the first phase we
 737 apply MINRES to the deflated linear equation

$$738 \quad (B.2) \quad \mathcal{P}(S(\sigma) - \mu_i(\sigma)I)\bar{x} = \mathcal{P}b^{(i)},$$

740 where $\mathcal{P} = I - W(W^TW)^{-1}W^T$, $C = [y_{\psi_1}(\sigma), \dots, y_{\psi_{i-1}}(\sigma), y_{\psi_{i+1}}(\sigma), \dots, y_{\psi_{n_{ev}}}(\sigma)]$,
 741 and $W = (S(\sigma) - \mu_i(\sigma)I)C$. As soon as the deflated linear equation in (B.2) is solved,
 742 the solution of the original linear system is formed as

$$743 \quad (B.3) \quad x = \mathcal{Q}\bar{x} + (I - \mathcal{Q})b^{(i)},$$

745 where $\mathcal{Q} = I - C(W^TW)^{-1}W^T$. Details on deflated MINRES can be found in [11, 12].

746 In case a symmetric factorization of $S(\sigma)$ is already at hand, this can be further
 747 exploited to compute the solution of the linear system $(S(\sigma) - \mu_i(\sigma)I)x = b^{(i)}$ by
 748 solving the linear system⁹

$$749 \quad S(\sigma)^{-1}(S(\sigma) - \mu_i(\sigma)I)x = S(\sigma)^{-1}b^{(i)}.$$

751 The matrix $S(\sigma)^{-1}(S(\sigma) - \mu_i(\sigma)I)$ is symmetric, and its eigenvalues are equal to
 752 $\{\frac{\mu_k(\sigma) - \mu_i(\sigma)}{\mu_k(\sigma)}\}_{k=1, \dots, s}$. This preconditioned linear system can also be combined with
 753 deflation. In the latter case MINRES is applied to the equation $\mathcal{P}S(\sigma)^{-1}(S(\sigma) -$
 754 $\mu_i(\sigma)I)x = \mathcal{P}S(\sigma)^{-1}b^{(i)}$ and the original linear system solution is obtained exactly as
 755 in (B.3).

756 **COROLLARY B.1.** *Let the eigenvalues of matrix $S(\sigma)$ be ordered as*

$$757 \quad \mu_1(\sigma) \leq \dots \leq \mu_{\psi_1-1}(\sigma) \leq \mu_{\psi_1}(\sigma) \leq \dots \leq \mu_{\psi_{n_{ev}}}(\sigma) \leq \mu_{\psi_{n_{ev}+1}}(\sigma) \leq \dots \leq \mu_s(\sigma).$$

759 *Then, the effective condition number of the matrix $\mathcal{P}(S(\sigma) - \mu_i(\sigma)I)$ is equal to*

$$760 \quad \kappa_{MR,i} = \frac{\max\{|\mu_1(\sigma) - \mu_i(\sigma)|, |\mu_s(\sigma) - \mu_i(\sigma)|\}}{\min\{|\mu_{\psi_1-1}(\sigma) - \mu_i(\sigma)|, |\mu_{\psi_{n_{ev}+1}}(\sigma) - \mu_i(\sigma)|\}}, \quad i = \psi_1, \dots, \psi_{n_{ev}}.$$


762 *Similarly, the effective condition number of the preconditioned matrix $\mathcal{P}S(\sigma)^{-1}(S(\sigma) -$
 763 $\mu_i(\sigma)I)$ is equal to*

$$764 \quad \kappa_{PMR,i} = \frac{\max\left\{\left|1 - \frac{\mu_i(\sigma)}{\mu_1(\sigma)}\right|, \left|1 - \frac{\mu_i(\sigma)}{\mu_s(\sigma)}\right|\right\}}{\min\left\{\left|1 - \frac{\mu_i(\sigma)}{\mu_{\psi_1-1}(\sigma)}\right|, \left|1 - \frac{\mu_i(\sigma)}{\mu_{\psi_{n_{ev}+1}}(\sigma)}\right|\right\}}.$$

766 If we do not deflate the n_{ev} computed eigenvectors of the matrix $S(\sigma)$, the de-
 767 nominators in Corollary B.1 become $\min\{|\mu_{i-1}(\sigma) - \mu_i(\sigma)|, |\mu_{i+1}(\sigma) - \mu_i(\sigma)|\}$ and
 768 $\min\{1 - \frac{\mu_i(\sigma)}{\mu_{i-1}(\sigma)}, 1 - \frac{\mu_i(\sigma)}{\mu_{i+1}(\sigma)}\}$, respectively. When $0 \leq \sigma \leq \lambda_{\min}(A, M)$, deflated
 769 MINRES can be replaced by a deflated variant of the conjugate gradient method [18];
 770 see, for example, [35, 37].

771 Figure B.1 plots the number of iterations required by MINRES to compute the
 772 eigenvector derivatives $y'_i(\sigma)$, $i = 1, \dots, n_{ev}$, up to a tolerance equal to 1.0×10^{-8} ,
 776 for a 253×148 finite difference discretization of the Dirichlet eigenvalue problem.

⁹Recall that MINRES requires an SPD preconditioner. If $S(\sigma)$ is not SPD, then alternative Kyrlov subspace iterative linear system solvers should be considered.



771 FIG. B.1. Number of (deflated) MINRES iterations with/without preconditioning to compute
 772 the eigenvector derivatives $y'_i(\sigma)$, $i = 1, \dots, n_{ev}$. Legend: “-” (deflation without preconditioning),
 773 “.” (preconditioning without deflation), and “-.” (preconditioning with deflation).

777 Combining preconditioning by matrix $S(\sigma)$ while deflating the invariant subspace
 778 associated with the n_{ev} computed eigenvectors of the matrix $S(\sigma)$ proved to be the
 779 fastest scheme in terms of iterations required to achieve convergence. Moreover, linear
 780 systems corresponding to eigenvector derivatives associated with eigenvalues that lie
 781 closer to σ converge faster due to a smaller effective condition number.

782 **Acknowledgments.** The author would like to thank the anonymous referees,
 783 whose detailed comments and suggestions led to considerable improvements in the
 784 quality of the present manuscript. In addition, the author is indebted to Daniel
 785 Kressner, Anthony Austin, and Yousef Saad for their feedback and encouragement to
 786 pursue the work featured in this manuscript.

787 REFERENCES

- 788 [1] *Fortran Compiler XE 14.0 for Linux*, Intel Corporation, 2018.
- 789 [2] P. R. AMESTOY, I. S. DUFF, J. KOSTER, AND J.-Y. L’EXCELLENT, *A fully asynchronous multi-*
 790 *frontal solver using distributed dynamic scheduling*, SIAM J. Matrix Anal. Appl., 23 (2001),
 791 pp. 15–41.
- 792 [3] E. ANDERSON, Z. BAI, C. BISCHOF, S. BLACKFORD, J. DONGARRA, J. DU CROZ, A. GREEN-
 793 BAUM, S. HAMMARLING, A. MCKENNEY, AND D. SORENSEN, *LAPACK Users’ Guide*, Vol. 9,
 794 SIAM, Philadelphia, 1999.
- 795 [4] A. ANDREW AND R. TAN, *Computation of derivatives of repeated eigenvalues and the corre-*
 796 *sponding eigenvectors of symmetric matrix pencils*, SIAM J. Matrix Anal. Appl., 20 (1998),
 797 pp. 78–100, <https://doi.org/10.1137/S0895479896304332>.
- 798 [5] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN,
 799 A. DENER, V. ELJKHOUT, W. D. GROPP, D. KARPEYEV, D. KAUSHIK, M. G. KNEPLEY,
 800 D. A. MAY, L. C. MCINNES, R. T. MILLS, T. MUNSON, K. RUPP, P. SANAN, B. F. SMITH,
 801 S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Users Manual*, Tech. Report ANL-95/11,
 802 Revision 3.11, Argonne National Laboratory, 2019, <https://www.mcs.anl.gov/petsc>.
- 803 [6] C. BEKAS AND Y. SAAD, *Computation of smallest eigenvalues using spectral Schur comple-*
 804 *ments*, SIAM J. Sci. Comput., 27 (2006), pp. 458–481.
- 805 [7] J. K. BENNIGHOF AND R. B. LEHOUCQ, *An automated multilevel substructuring method for*
 806 *eigenspace computation in linear elastodynamics*, SIAM J. Sci. Comput., 25 (2004),
 807 pp. 2084–2106.
- 808 [8] S.-C. T. CHOI, C. C. PAIGE, AND M. A. SAUNDERS, *MINRES-QLP: A Krylov subspace method*
 809 *for indefinite or singular symmetric systems*, SIAM J. Sci. Comput., 33 (2011), pp. 1810–
 810 1836.
- 811 [9] T. A. DAVIS AND Y. HU, *The University of Florida sparse matrix collection*, ACM Trans. Math.
 812 Software, 38 (2011), pp. 1:1–1:25, <https://doi.org/10.1145/2049662.2049663>.

- 813 [10] W. GAO, X. S. LI, C. YANG, AND Z. BAI, *An implementation and evaluation of the AMLS*
814 *method for sparse eigenvalue problems*, ACM Trans. Math. Software, 34 (2008), pp. 20:1–
815 20:28, <https://doi.org/10.1145/1377596.1377600>.
- 816 [11] A. GAUL, M. H. GUTKNECHT, J. LIESEN, AND R. NABBEN, *A framework for deflated and*
817 *augmented Krylov subspace methods*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 495–518.
- 818 [12] A. GAUL AND N. SCHLÖMER, *Preconditioned recycling Krylov subspace methods for self-adjoint*
819 *problems*, Electron. Trans. Numer. Anal., 44 (2015), pp. 522–547.
- 820 [13] R. G. GRIMES, J. G. LEWIS, AND H. D. SIMON, *A shifted block Lanczos algorithm for solv-*
821 *ing sparse symmetric generalized eigenproblems*, SIAM J. Matrix Anal. Appl., 15 (1994),
822 pp. 228–272, <https://doi.org/10.1137/S0895479888151111>.
- 823 [14] W. GROPP, W. D. GROPP, E. LUSK, A. D. F. E. E. LUSK, AND A. SKJELLUM, *Using MPI:*
824 *Portable Parallel Programming with the Message-passing Interface*, Vol. 1, MIT Press,
825 Cambridge, MA, 1999.
- 826 [15] W. D. GROPP, *Parallel computing and domain decomposition*, in Proceedings of the Fifth Inter-
827 national Symposium on Domain Decomposition Methods for Partial Differential Equations,
828 Philadelphia, 1992, pp. 349–361.
- 829 [16] A. HANNUKAINEN, J. MALINEN, AND A. OJALAMMI, *Efficient Solution of Symmetric Eigenvalue*
830 *Problems from Families of Coupled Systems*, preprint, arXiv:1806.07235, 2018.
- 831 [17] V. HERNANDEZ, J. E. ROMAN, AND V. VIDAL, *SLEPC: A scalable and flexible toolkit for the*
832 *solution of eigenvalue problems*, ACM Trans. Math. Software, 31 (2005), pp. 351–362.
- 833 [18] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*,
834 J. Res. National Bureau of Standards, 49 (1952), pp. 409–436.
- 835 [19] T. HOSHI, H. IMACHI, A. KUWATA, K. KAKUDA, T. FUJITA, AND H. MATSUI, *Numerical aspect*
836 *of large-scale electronic state calculation for flexible device material*, Jpn. J. Ind. Appl.
837 Math., 36 (2019), pp. 685–698.
- 838 [20] V. KALANTZIS, *Domain Decomposition Algorithms for the Solution of Sparse Symmetric Gen-*
839 *eralized Eigenvalue Problems*, Ph.D. thesis, University of Minnesota, 2018.
- 840 [21] V. KALANTZIS, *A spectral Newton-Schur algorithm for the solution of symmetric generalized*
841 *eigenvalue problems*, Electron. Trans. Numer. Anal., 52 (2020), pp. 132–153.
- 842 [22] V. KALANTZIS, J. KESTYN, E. POLIZZI, AND Y. SAAD, *Domain decomposition approaches for*
843 *accelerating contour integration eigenvalue solvers for symmetric eigenvalue problems*, Num-
844 er. Linear Algebra Appl., 25 (2018), e2154.
- 845 [23] V. KALANTZIS, R. LI, AND Y. SAAD, *Spectral Schur complement techniques for symmetric*
846 *eigenvalue problems*, Electron. Trans. Numer. Anal., 45 (2016), pp. 305–329.
- 847 [24] V. KALANTZIS, Y. XI, AND Y. SAAD, *Beyond automated multilevel substructuring: Domain*
848 *decomposition with rational filtering*, SIAM J. Sci. Comput., 40 (2018), pp. C477–C502.
- 849 [25] G. KARYPIS AND V. KUMAR, *A fast and high quality multilevel scheme for partitioning ir-*
850 *regular graphs*, SIAM J. Sci. Comput., 20 (1998), pp. 359–392, <https://doi.org/10.1137/S1064827595287997>.
- 851 [26] T. KATO, *Perturbation Theory for Linear Operators*, 2nd ed., Grundlehren Math. Wiss. 132,
852 Springer, Berlin, 1976, <https://cds.cern.ch/record/101545>.
- 853 [27] J. KESTYN, V. KALANTZIS, E. POLIZZI, AND Y. SAAD, *PFEAST: A high performance sparse*
854 *eigenvalue solver using distributed-memory linear solvers*, in SC'16: Proceedings of the In-
855 ternational Conference for High Performance Computing, Networking, Storage and Analy-
856 sis, IEEE, 2016, pp. 178–189.
- 857 [28] J. H. KO AND Z. BAI, *High-frequency response analysis via algebraic substructuring*, Internat.
858 J. Numer. Methods Engrg., 76 (2008), pp. 295–313.
- 859 [29] L. KOMZSIK AND T. ROSE, *Parallel methods on large-scale structural analysis and physics appli-*
860 *cations substructuring in MSC/NASTRAN for large scale parallel applications*, Computing
861 Systems in Engineering, 2 (1991), pp. 167–173, [http://dx.doi.org/10.1016/0956-0521\(91\)](http://dx.doi.org/10.1016/0956-0521(91)90017-Y)
862 90017-Y.
- 863 [30] S. LUI, *Kron's method for symmetric eigenvalue problems*, J. Comput. Appl. Math., 98 (1998),
864 pp. 35–48.
- 865 [31] K. MEERBERGEN AND Z. BAI, *The Lanczos method for parameterized symmetric linear systems*
866 *with multiple right-hand sides*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 1642–1662.
- 867 [32] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*,
868 SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- 869 [33] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, Classics in Appl. Math. 20, SIAM,
870 Philadelphia, 1998, <https://doi.org/10.1137/1.9781611971163>.
- 871 [34] Y. SAAD, *ILUM: A multi-elimination ILU preconditioner for general sparse matrices*, SIAM
872 J. Sci. Comput., 17 (1996), pp. 830–847.
- 873

- 874 [35] Y. SAAD, M. YEUNG, J. ERHEL, AND F. GUYOMARC'H, *A deflated version of the conjugate*
875 *gradient algorithm*, SIAM J. Sci. Comput., 21 (2000), pp. 1909–1926, [https://doi.org/10.](https://doi.org/10.1137/S1064829598339761)
876 [1137/S1064829598339761](https://doi.org/10.1137/S1064829598339761).
- 877 [36] A. STATHOPOULOS AND J. R. MCCOMBS, *PRIMME: Preconditioned iterative multimethod*
878 *eigensolver: Methods and software description*, ACM Trans. Math. Software, 37 (2010),
879 pp. 21:1–21:30, <https://doi.org/10.1145/1731022.1731031>.
- 880 [37] A. STATHOPOULOS AND K. ORGINOS, *Computing and deflating eigenvalues while solving multiple*
881 *right-hand side linear systems with an application to quantum chromodynamics*, SIAM J.
882 Sci. Comput., 32 (2010), pp. 439–462.
- 883 [38] Y. SU, T. LU, AND Z. BAI, *2D Eigenvalue Problems I: Existence and Number of Solutions*,
884 preprint, arXiv:1911.08109, 2019.
- 885 [39] U. VON LUXBURG, *A tutorial on spectral clustering*, Stat. Comput., 17 (2007), pp. 395–416.
- 886 [40] C. YANG, W. GAO, Z. BAI, X. S. LI, L.-Q. LEE, P. HUSBANDS, AND E. NG, *An algebraic*
887 *substructuring method for large-scale eigenvalue calculation*, SIAM J. Sci. Comput., 27
888 (2005), pp. 873–892, <https://doi.org/10.1137/040613767>.
- 889 [41] H. ZHANG, B. SMITH, M. STERNBERG, AND P. ZAPOL, *SIPs: Shift-and-invert parallel spectral*
890 *transformations*, ACM Trans. Math. Software, 33 (2007), pp. 111–127.