Citation for published version:
Robbé, M, Sadkane, M & Spence, A 2009, 'Inexact inverse subspace iteration with preconditioning applied to non-Hermitian eigenvalue problems', *SIAM Journal On Matrix Analysis and Applications (SIMAX)*, vol. 31, no. 1, pp. 92-113. https://doi.org/10.1137/060673795

DOI:
10.1137/060673795

Publication date:
2009

Link to publication

University of Bath

Alternative formats
If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 25. Aug. 2022

INEXACT INVERSE SUBSPACE ITERATION WITH PRECONDITIONING APPLIED TO NON-HERMITIAN EIGENVALUE PROBLEMS

MICKAËL ROBBÉ*, MILOUD SADKANE* , AND ALASTAIR SPENCE †

Abstract. Convergence results are provided for inexact inverse subspace iteration applied to the problem of finding the invariant subspace associated with a small number of eigenvalues of a large sparse matrix. These results are illustrated by the use of block-GMRES as the iterative solver. The costs of the inexact solves are measured by the number of inner iterations needed by the iterative solver at each outer step of the algorithm. It is shown that for a decreasing tolerance the number of inner iterations should not increase as the outer iteration proceeds, but it may increase for preconditioned iterative solves. However, it is also shown that an appropriate small rank change to the preconditioner can produce significant savings in costs, and in particular, can produce a situation where there is no increase in the costs of the iterative solves even though the solve tolerances are reducing. Numerical examples are provided to illustrate the theory.

Key words. Eigenvalue approximation, inverse subspace iteration, iterative methods, preconditioning.

AMS subject classifications. 65F10, 65F15

1. Introduction. Inverse subspace iteration is a block version of the inverse iteration. It computes an approximation of the invariant subspace of a large matrix $A \in \mathbb{C}^{n\times n}$ corresponding to the eigenvalues in an isolated cluster around a given shift σ. The corresponding algorithm is very simple and can formally be written as

$$X_i = (A - \sigma I)^{-1} X_{i-1},\ i = 1, 2, \ldots, \tag{1.1}$$

where $X_0 \in \mathbb{C}^{n\times p}$ is full rank with $p \ll n$. As the iterations unfold, the invariant subspace and hence the eigenvectors corresponding to eigenvalues near σ eventually dominates X_i. The method is known to be reliable [17, 26, 20, 28] and, although its convergence is linear, only a few iterations are needed to converge provided that the target eigenvalues lie in a cluster well separated from the rest of the spectrum and p is chosen as large as the number of eigenvalues in the cluster. The drawback of this method is that each iteration necessitates the exact solution of a block linear system, that is, a linear system with multiple right-hand sides of the form

$$(A - \sigma I)Y = X, \quad Y, X \in \mathbb{C}^{n\times p}, \tag{1.2}$$

which is a challenge when n is large. The first aim of this paper is to analyze the convergence of (1.1) when the underlying block linear systems (1.2) are solved inexactly by an iterative method. The method obtained this way belongs to the wide class of "inner-outer" iterative methods. The outer iteration is the inverse subspace iteration and the inner iteration is the iterative solution of the block linear system (1.2). The results in this paper extend the results in [14] and [11] on inexact inverse iteration to inexact inverse subspace iteration.
The second aim of this paper is to discuss the performance of unpreconditioned and

*Université de Bretagne Occidentale. Département de Mathématiques. 6, Av. Le Gorgeu. CS 93837. 29238 Brest Cedex 3, France. (robbe@univ-brest.fr, sadkane@univ-brest.fr).

†Department of Mathematical Sciences, University of Bath, Claverton Down, BA2 7AY, United Kingdom. (a.spence@maths.bath.ac.uk).

preconditioned block-GMRES as the inexact solver. If P denotes a preconditioner for $(A-\sigma I)$, the (right) preconditioned form of (1.2) is

$$(A-\sigma I)P^{-1}\widetilde{Y} = X, \quad Y = P^{-1}\widetilde{Y}, \tag{1.3}$$

with the aim that (1.3) is solved more efficiently than (1.2). For inexact inverse iteration, [4, 6] consider the costs of the inner solves for Krylov solvers and analyse cases where the number of inner iterations may remain approximately constant or may increase as the outer iteration proceeds. In this paper we extend these results to the block case. Moreover, we show how a rank-p modification of P gives a "tuned" preconditioner which eliminates the increase in the number of inner iterations as the outer iteration proceeds.

Recently, inexact inverse iteration has been discussed by [14], [11], and [5], and for the symmetric case by [25] and [4]. The idea of tuning the preconditioner for eigenvalue problems was introduced in [7, 8] for inexact inverse iteration. There is considerable interest in inexact solves for subspace based methods, especially in relation to the Jacobi-Davidson method (JD) [24, 2], [16] and the Riccati-based methods as developed in [18], [3], the latter may be viewed as the block analogue of JD and are useful for computing invariant subspaces. Other useful methods which use inexact solves within inner outer iterations include the trace minimization [22] and the inexact Raleigh quotient (IRQ) iteration [23]. A link between IRQ and the simplified JD [15], i.e. the correction equation in JD without expanding the search space , has been established in [23]. Other methods which use preconditioned iterative solves on subspaces are LOBPCG [13] and the truncated-CG-base trust-region [1] which are particularly successful for finding extremal eigenvalues of symmetric matrices. The latter is also related to the simplified JD [1].

The tuned preconditioner developed in this paper is effectively suited for inexact inverse subspace iteration. As such, it does not apply for example to JD. Tuning the preconditioner for JD, not only for the simplified JD, would certainly produce a very efficient eigensolver. This point will hopefully be treated in a future work.

In Section 2 we present the inexact inverse subspace iteration algorithm and some preliminary results. In particular, we discuss some tools for measuring the closeness between subspaces. Section 3 presents a convergence theory for the inexact (and exact) inverse subspace iteration. We shall show that provided these linear systems are solved to an appropriately chosen decreasing tolerance then the method attains a linear rate of convergence just as in the case of exact solves (Theorem 3.1). In Section 3.1 we consider the use of block-GMRES as the (unpreconditioned) solver. We show that for a decreasing tolerance the number of inner iterations should not increase as the outer iteration converges. The case of preconditioned solves is discussed in Section 4. Our main result, presented in Section 4.2, is that if a standard preconditioner is modified by a small rank change then there is again no increase in the number of inner iterations as the outer iteration proceeds. We call the process of modifying the preconditioner in this way "tuning". In Section 5, numerical tests are given to illustrate the theory. In particular, it is shown that significant savings are obtained when a tuned preconditoner is used.

The main difference with the works in [14, 11, 4, 6, 7, 8] is that besides the extension to block case, our convergence theory for inexact subspace iteration and for bloc-GMRES works under rather weak assumptions; for example the matrix A can be defective. Furthermore, we provide a rigorous proof that the tuned preconditioner removes the dependence on the number of inner iterations.

2. Inexact inverse subspace iteration. In this section we describe the inexact inverse subspace iteration algorithm, and in Section 2.1 revise some background material, especially relating to the angle between two subspaces.

In many applications interest centres on the invariant subspace corresponding to the eigenvalues nearest zero, and from now on we shall choose the shift in (1.2) to be zero. Much of what we say extends to the case of a nonzero fixed shift.

Inexact inverse subspace iteration is described in the following inner-outer algorithm:

ALGORITHM 1 (Inexact Subspace Iteration).
Given $\delta \geq 0$ *and* $X_0 \in \mathbb{C}^{n\times p}$ *with* $X_0^* X_0 = I$,
For $i = 0, 1, \ldots$

1. *Compute* $L_i = X_i^* A X_i$,
2. *Set* $R_i = AX_i - X_i L_i$ *and test the convergence,*
3. *Solve* $AY_i = X_i$ *inexactly, that is, compute* Y_i *such that*

$$X_i - AY_i = \Delta_i \quad with \quad \|\Delta_i\| \leq \tau_i = \delta \|R_i\|,$$

4. *Orthonormalize the columns of* Y_i *into* X_{i+1}.

End For i

In Algorithm 1 and throughout this paper, the symbol $\|\ \|$ denotes the Euclidean norm or its induced matrix norm.

In Section 3 we first analyse the convergence of Algorithm 1 with no particular solver in mind, and in Section 3.1 we discuss the case when the block linear systems in step 3 of Algorithm 1 are solved by block-GMRES. Note that if the block systems are solved exactly, then the (exact) inverse subspace iteration (1.1) is recovered.

The next section gathers some technical details which will be used throughout this paper.

2.1. Notation and preliminaries. We assume that the eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$ of A are such that

$$0 < |\lambda_1| \leq \ldots \leq |\lambda_p| \ < \ |\lambda_{p+1}| \leq \ldots \leq |\lambda_n|. \tag{2.1}$$

By Schur's theorem we may decompose the matrix A to upper triangular form by a unitary matrix $\begin{pmatrix} V_1 & V_1^\perp \end{pmatrix}$:

$$A = \begin{pmatrix} V_1 & V_1^\perp \end{pmatrix} \begin{pmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{pmatrix} \begin{pmatrix} V_1 & V_1^\perp \end{pmatrix}^* \tag{2.2}$$

where $V_1 \in \mathbf{C}^{n\times p}$, $V_1^\perp \in \mathbf{C}^{n\times(n-p)}$, $T_{11} \in \mathbf{C}^{p\times p}$ and $T_{22} \in \mathbf{C}^{(n-p)\times(n-p)}$. The spectra of T_{11} and T_{22} are respectively $\lambda_1, \ldots \lambda_p$ and $\lambda_{p+1}, \ldots \lambda_n$. Let $Q \in \mathbf{C}^{p\times(n-p)}$ be the unique solution of the Sylvester equation

$$QT_{22} - T_{11}Q = T_{12}. \tag{2.3}$$

Then A can be block-diagonalized as follows (see e.g. [10]):

$$A = \begin{pmatrix} V_1 & V_1^\perp \end{pmatrix} \begin{pmatrix} I & Q \\ 0 & I \end{pmatrix} \begin{pmatrix} T_{11} & 0 \\ 0 & T_{22} \end{pmatrix} \begin{pmatrix} I & -Q \\ 0 & I \end{pmatrix} \begin{pmatrix} V_1 & V_1^\perp \end{pmatrix}^*. \tag{2.4}$$

Let

$$V_2 = \left(V_1 Q + V_1^\perp\right) \left(I + Q^* Q\right)^{-\frac{1}{2}},$$

$$L = T_{11}, \quad M = (I + Q^*Q)^{-\frac{1}{2}} T_{22} (I + Q^*Q)^{\frac{1}{2}} .$$

Then the block-diagonalization in (2.4) can be written

$$A = \begin{pmatrix} V_1 & V_2 \end{pmatrix} \begin{pmatrix} L & 0 \\ 0 & M \end{pmatrix} \begin{pmatrix} V_1 & V_2 \end{pmatrix}^{-1} . \tag{2.5}$$

Note that M and T_{22} have the same spectra and that V_1 and V_2 have orthonormal columns. The subspaces $\mathcal{V}_1 = \mathcal{R}(V_1)$ and $\mathcal{V}_2 = \mathcal{R}(V_2)$ spanned by the columns of V_1 and V_2 are complementary invariant subspaces of A associated respectively with the eigenvalues $\lambda_1, \cdots, \lambda_p$ of L and $\lambda_{p+1}, \cdots, \lambda_n$ of M. Our main task in this paper is to compute the invariant subspace $\widetilde{\mathcal{V}}_1 \subset \mathcal{V}_1$ associated with the $q \leq p$ smallest (in modulus) eigenvalues of A by Algorithm 1.

The smallest (largest) singular value of a matrix B is denoted by $\sigma_{min}(B) = \min_{\|x\|=1} \|Bx\|$ ($\sigma_{max}(B) = \|B\| = \max_{\|x\|=1} \|Bx\|$). The separation $\mathrm{sep}(E, F)$ between two matrices $E \in \mathbb{C}^{p\times p}$ and $F \in \mathbb{C}^{q\times q}$ is defined as (see [27]):

$$\mathrm{sep}(E, F) = \min_{\|X\|=1} \|EX - XF\| .$$

It is known that $\mathrm{sep}(E, F) > 0$ if and only if E and F have disjoint spectra. Our analysis will lead us to use either $\mathrm{sep}(T_{22}, L)$ or $\mathrm{sep}(M, L)$. These quantities are equivalent since (see [27]):

$$\mathrm{sep}(T_{22}, L)/\kappa \leq \mathrm{sep}(M, L) \leq \kappa\, \mathrm{sep}(T_{22}, L)$$

where $\kappa = \sqrt{\frac{1+\sigma^2_{\max}(Q)}{1+\sigma^2_{\min}(Q)}}$, with Q defined by (2.3).

Let

$$W_1 = V_1 - V_1^{\perp} Q^* \quad \text{and} \quad W_2 = V_1^{\perp} (I + Q^*Q)^{\frac{1}{2}} . \tag{2.6}$$

Then $\mathcal{R}(W_1)$ and $\mathcal{R}(W_2)$ are complementary invariant subspaces of A^* corresponding to the eigenvalues $\bar{\lambda}_1, \cdots, \bar{\lambda}_p$ of L^* and $\bar{\lambda}_{p+1}, \cdots, \bar{\lambda}_n$ of M^* and we have

$$\begin{pmatrix} W_1 & W_2 \end{pmatrix}^* = \begin{pmatrix} V_1 & V_2 \end{pmatrix}^{-1} . \tag{2.7}$$

The spectral projection on $\mathcal{V}_1$ is defined by

$$\mathcal{P} = V_1 W_1^* . \tag{2.8}$$

Note that

$$\|\mathcal{P}\| = \|W_1\| = \sqrt{1 + \|Q\|^2} . \tag{2.9}$$

To understand the performance of Algorithm 1 we need to measure the deviation of X_i from V_1. This can be done by monitoring the angle between the subspaces $\mathcal{V}_1$ and $\mathcal{X}_i = \mathcal{R}(X_i)$. One tool is the sine of the largest canonical angle between $\mathcal{V}_1$ and $\mathcal{X}_i$ defined by (see [10, p.584]):

$$\sin \angle(\mathcal{X}_i, \mathcal{V}_1) = \left\| (V_1^{\perp})^* X_i \right\| . \tag{2.10}$$

We assume that the subspaces $\mathcal{X}_i$ and $\mathcal{V}_1$ have the same dimension. Then (see [10, p.76])

$$\sin \angle(\mathcal{X}_i, \mathcal{V}_1) = \sin \angle(\mathcal{V}_1, \mathcal{X}_i) = \|X_i X_i^* - V_1 V_1^*\| \tag{2.11}$$

$$= \min_{Z\in\mathbb{C}^{p\times p}} \|X_i - V_1 Z\| = \min_{Z\in\mathbb{C}^{p\times p}} \|V_1 - X_i Z\| . \tag{2.12}$$

We also assume that the matrix X_i can be decomposed as

$$X_i = V_1 C_i + V_2 S_i \quad \text{with } \|S_i\| < 1. \tag{2.13}$$

Using (2.7), we see that the matrices C_i and S_i are given by

$$C_i = W_1^* X_i \in \mathbb{C}^{p\times p}, \quad S_i = W_2^* X_i \in \mathbb{C}^{(n-p)\times p}. \tag{2.14}$$

From (2.13), formula (2.10) becomes

$$\sin\angle(\mathcal{X}_i, \mathcal{V}_1) = \|(V_1^\perp)^* V_2 S_i\|, \tag{2.15}$$

which shows that $\|S_i\|$ can also be used to measure the deviation between $\mathcal{V}_1$ and $\mathcal{X}_i$. In fact, we will cast our results in terms of the quantities

$$\sin\angle(\mathcal{X}_i, \mathcal{V}_1),\ t_i = \left\|S_i C_i^{-1}\right\| \quad \text{or} \quad s_i = t_i \left\|C_i\right\|.$$

Note that in the case when A is Hermitian, then t_i and $\|S_i\|$ represent respectively the tangent and the sine of the largest angle between $\mathcal{X}_i$ and $\mathcal{V}_i$. The following proposition shows that all these quantities behave like $\|S_i\|$. So $\mathcal{X}_i \to \mathcal{V}_1$ if and only if $t_i \to 0$ or $s_i \to 0$ or $\|S_i\| \to 0$.

PROPOSITION 2.1. *Let X_i be partitioned as in (2.13). Then*

1) C_i is nonsingular and therefore t_i is well defined. Moreover, the singular values of C_i satisfy

$$0 < 1 - \|S_i\| \le \sigma_k(C_i) \le 1 + \|S_i\|,\ k = 1, \ldots, p, \tag{2.16}$$

and C_i can be written

$$C_i = U_i + \Upsilon_i \tag{2.17}$$

where U_i is unitary and $\|\Upsilon_i\| \le \|S_i\| < 1$.

2) $\sin\angle(\mathcal{X}_i, \mathcal{V}_1) \le \|S_i\| \le s_i \le \|S_i\| \frac{1+\|S_i\|}{1-\|S_i\|}$.

3) $\sin\angle(\mathcal{X}_i, \mathcal{V}_1) \le t_i \le \frac{\|S_i\|}{1-\|S_i\|}$.

4) $\|S_i\| \le \|\mathcal{P}\| \sin\angle(\mathcal{X}_i, \mathcal{V}_1)$.

Proof.

1) Assume C_i is singular and let u be a nonzero vector such that $C_i u = 0$. Then

$$\|u\| = \|X_i u\| = \|V_2 S_i u\| \le \|S_i\|\|u\| < \|u\|,$$

a contradiction. Hence C_i is nonsingular. The k^{th} singular values of X_i and $V_1 C_i$ satisfy (see [10, p. 428])

$$|\sigma_k(X_i) - \sigma_k(V_1 C_i)| \le \|X_i - V_1 C_i\| \le \|S_i\|,$$

and hence

$$|1 - \sigma_k(C_i)| \le \|S_i\|.$$

Let $C_i = W_i^{(l)} \Sigma_i W_i^{(r)}$ be the singular value decomposition of C_i. Then C_i can be written as in (2.17) with $U_i = W_i^{(l)} W_i^{(r)}$ and $\Upsilon_i = W_i^{(l)} (\Sigma_i - I) W_i^{(r)}$.

2) The first bound follows from (2.15) and the other ones from the definition of s_i and (2.16).

3)

$$\begin{aligned}\sin\angle(\mathcal{X}_i,\mathcal{V}_1) &= \|(X_i^\perp)^* V_1\| = \|(X_i^\perp)^*(V_1 - X_i C_i^{-1})\| \\ &\le \|V_1 - X_i C_i^{-1}\| = \|V_2 S_i C_i^{-1}\| = t_i.\\ t_i = \|S_i C_i^{-1}\| &\le \frac{\|S_i\|}{\sigma_{\min}(C_i)} \le \frac{\|S_i\|}{1-\|S_i\|}.\end{aligned}$$

4)

$$\begin{aligned}\|S_i\| = \|W_2^* X_i\| &= \left\|(I+Q^*Q)^{\frac{1}{2}}\ (V_1^\perp)^* X_i\right\| \\ &\le \|\mathcal{P}\|\sin\angle(\mathcal{X}_i,\mathcal{V}_1).\end{aligned}$$

□

The following proposition gives bounds on the residual norm.

PROPOSITION 2.2. *The following inequalities hold:*

$$sep(T_{22}, L_i)\ \sin\angle(\mathcal{X}_i,\mathcal{V}_1) \le \|R_i\| \le ||\mathcal{S}||s_i.$$

where $\mathcal{S}$ is the Sylvester operator $X \to \mathcal{S}(X) = MX - XL$ and

$$\|\mathcal{S}|| = \sup_{\|X\|=1} \|\mathcal{S}(X)\|.$$

Proof.

$$\begin{aligned}\|R_i\| &\ge \|(V_1^\perp)^* R_i\| \\ &= \|(V_1^\perp)^* A X_i - (V_1^\perp)^* X_i L_i\|.\end{aligned}$$

From (2.2) we have $(V_1^\perp)^* A = T_{22}(V_1^\perp)^*$. Then

$$\|R_i\| \ge \|T_{22}(V_1^\perp)^* X_i - (V_1^\perp)^* X_i L_i\| \ge \mathrm{sep}(T_{22}, L_i)\ \|(V_1^\perp)^* X_i\|.$$

Also,

$$\begin{aligned}\|R_i\| &= \min_{Z\in\mathbb{C}^{p\times p}} \|AX_i - X_i Z\| \quad (\text{see } [27, \text{Thm.1.15}]) \\ &\le \|AX_i - X_i C_i^{-1} L C_i\| \\ &= \|MS_i - S_i C_i^{-1} L C_i\| \\ &= \|\left(MS_iC_i^{-1} - S_iC_i^{-1}L\right)C_i\| \le \|\mathcal{S}\|\ \|S_iC_i^{-1}\|\|C_i\|.\end{aligned}$$

□

3. Convergence analysis of Algorithm 1. In this section we analyze the convergence of Algorithm 1 when the inner iterations are solved inexactly. First, we make no assumption on the inexact solver except that step 3 in Algorithm 1 is satisfied. Then, in Section 3.1, we assume that a block-GMRES method is the inexact solver.

THEOREM 3.1. *(Convergence of Algorithm 1).*
Assume that X_0 is close to V_1. If a tolerance of the form $\tau_i = \delta\|R_i\|$ is chosen with small enough δ (e.g. when $\delta < \left(\|(C_i^{-1}\|\|\mathcal{P}\|\|R_i\|\right)^{-1}$), then we have

$$t_{i+1} \le \|M^{-1}\|\|L\|\frac{t_i + \alpha_i\tau_i}{1-\alpha_i\tau_i}, \tag{3.1}$$

with $\alpha_i = \|C_i^{-1}\|\|\mathcal{P}\| \leq \frac{\|\mathcal{P}\|}{1-\|S_i\|}$, *and where* M *and* L *are defined in (2.5).*
If, in addition, $\|M^{-1}\|\|L\| < 1$, *then Algorithm 1 converges linearly.*

Proof. Note first that

$$t_{i+1} = \left\|S_{i+1}C_{i+1}^{-1}\right\| = \left\|(W_2^*X_{i+1})(W_1^*X_{i+1})^{-1}\right\| = \left\|(W_2^*X_{i+1}K)(W_1^*X_{i+1}K)^{-1}\right\|$$

where $K \in \mathbb{C}^{p\times p}$ is an arbitrary nonsingular matrix. In particular $t_{i+1} = \|(W_2^*Y_i)(W_1^*Y_i)^{-1}\|$ and therefore

$$\begin{aligned}
t_{i+1} &= \left\|W_2^*A^{-1}(X_i-\Delta_i)\left(W_1^*A^{-1}(X_i-\Delta_i)\right)^{-1}\right\| \\
&= \left\|M^{-1}W_2^*(X_i-\Delta_i)\left(L^{-1}W_1^*(X_i-\Delta_i)\right)^{-1}\right\| \\
&\leq \|M^{-1}\|\|L\|\left\|W_2^*(X_i-\Delta_i)\left(W_1^*(X_i-\Delta_i)\right)^{-1}\right\| \\
&\leq \|M^{-1}\|\|L\|\left\|W_2^*(X_i-\Delta_i)C_i^{-1}\left(I-(W_1^*\Delta_i)C_i^{-1}\right)^{-1}\right\| \\
&\leq \|M^{-1}\|\|L\|\frac{t_i+\|W_2^*\Delta_i\|\|C_i^{-1}\|}{1-\|W_1^*\Delta_i\|\|C_i^{-1}\|},
\end{aligned}$$

The condition on δ ensures that $\|(W_1^*\Delta_i)C_i^{-1}\| < 1$. Thus

$$t_{i+1} \leq \|M^{-1}\|\|L\|\frac{t_i+\|W_2^*\Delta_i\|\|C_i^{-1}\|}{1-\|W_1^*\Delta_i\|\|C_i^{-1}\|}.$$

Note that $\|W_2^*\Delta_i\| \leq \|W_2\|\tau_i$. From (2.6) and (2.9) we have $\|W_2\| = \|\mathcal{P}\|$, and from Proposition 2.2, $\|C_i^{-1}\| \leq 1/(1-\|S_i\|)$. The expression $\|W_1^*\Delta_i\|$ is bounded in a similar way, and (3.1) is obtained. Since $\tau_i \leq \delta\|R_i\| \leq \delta\|\mathcal{S}\|\|C_i\|t_i$, linear convergence is established for small δ and $\|M^{-1}\|\|L\| < 1$. ∎

When A is Hermitian, $\|M^{-1}\|\|L\| = |\lambda_p|/|\lambda_{p+1}| < 1$, so the condition $\|M^{-1}\|\|L\| < 1$ is automatically satisfied, moreover, in this case $Q = 0$ in (2.4) and $\|\mathcal{P}\| = 1$ (see (2.9)), $\|C_i^{-1}\| = 1/\cos\angle(\mathcal{X}_i,\mathcal{V}_1)$ and $\|W_1^*\Delta_i\| = \|W_2^*\Delta_i\| = \|\Delta_i\|$. Thus (3.1) becomes

$$\tan\angle(\mathcal{X}_{i+1},\mathcal{V}_1) \leq \frac{|\lambda_p|}{|\lambda_{p+1}|}\,\frac{\sin\angle(\mathcal{X}_i,\mathcal{V}_1)+\tau_i}{\cos\angle(\mathcal{X}_i,\mathcal{V}_1)-\tau_i}. \tag{3.2}$$

If we now take $\delta = 0$ in Theorems 3.1, then we recover two convergence results for the exact solves case (see [20, Thm. 5.2] and [28, p. 383]). This point is clarified in the next corollaries.

COROLLARY 3.2. *If the block systems in step 3 of Algorithm 1 are solved exactly, i.e.* $\tau_i = 0$, *then*

$$t_i \leq \|M^{-1}\|\|L\|t_{i-1},$$

and

$$t_i \leq \|M^{-i}\|\|L^i\|t_0 \leq \left(\frac{|\lambda_p|}{|\lambda_{p+1}|}+\eta_i\right)^i t_0,$$

with $\lim_{i\to\infty}\eta_i = 0$.

Proof. The first inequality follows directly from Theorem 3.1. For the second one, we have

$$\begin{aligned} t_i &= \left\| \left(W_2^* A^{-1} X_i\right) \left(W_1^* A^{-1} X_i\right)^{-1} \right\| \\ &= \left\| \left(W_2^* A^{-i} X_0\right) \left(W_1^* A^{-i} X_0\right)^{-1} \right\| \\ &= \| \left(M^{-i} S_0\right) \left(L^{-i} C_0\right)^{-1} \| \leq \|M^{-i}\| \|L^i\| t_0. \end{aligned}$$

Now, using the fact that for any square matrix E, $\|E^i\| \leq \left(\rho(E) + \eta_E^{(i)}\right)^i$, where $\rho(E)$ is the spectral radius of E and $\lim_{i\to\infty} \eta_E^{(i)} = 0$, we obtain with obvious notation

$$\|M^{-i}\| \|L^i\| \leq \left(\rho(M^{-1}) + \eta_{M^{-1}}^{(i)}\right)^i \left(\rho(L) + \eta_L^{(i)}\right)^i = \left(\frac{|\lambda_p|}{|\lambda_{p+1}|} + \eta_i\right)^i$$

with

$$\eta_i = \eta_{M^{-1}}^{(i)} |\lambda_p| + \eta_L^{(i)} / |\lambda_{p+1}| + \eta_{M^{-1}}^{(i)} \eta_L^{(i)} \to 0 \quad \text{as} \quad i \to \infty.$$

□

In a practice, the block size p is chosen to enable the computation of invariant subspaces corresponding to close/multiple/complex pairs of eigenvalues. Therefore, to speed up the convergence, it is desirable to choose p larger than the dimension of the sought invariant subspace. Thus an estimate of the angle between $\mathcal{X}_i$ and a subspace $\widetilde{\mathcal{V}}_1 \subset \mathcal{V}_1$ is needed. Corollary 3.2 does not give such an estimate because t_i relates $\mathcal{X}_i$ to $\mathcal{V}_1$ not to a subspace $\widetilde{\mathcal{V}}_1 \subset \mathcal{V}_1$. The following corollary treats this point.

COROLLARY 3.3. *Assume that the block systems in step 3 of Algorithm 1 are solved exactly and that* $V_1 = \begin{pmatrix} \widetilde{V}_1 & \widetilde{\widetilde{V}}_1 \end{pmatrix}$ *with* $\widetilde{V}_1 \in \mathbb{C}^{n\times q}$, $q \leq p$ *and* $\widetilde{\mathcal{V}}_1 := \mathcal{R}(\widetilde{V}_1)$ *span an invariant subspace of* A *associated with the eigenvalues* $\lambda_1, \ldots, \lambda_q$. *Then*

$$\sin\angle(\widetilde{\mathcal{V}}_1, \mathcal{X}_i) \leq \left(\frac{|\lambda_q|}{|\lambda_{p+1}|} + \eta_i\right)^i \left\|(I - \mathcal{P})\widetilde{X}_0\right\|$$

with $\lim_{i\to\infty} \eta_i = 0$ *and* $\widetilde{X}_0 = X_0 C_0^{-1} \begin{pmatrix} I_q \\ 0 \end{pmatrix}$.

Proof. From the proof of Corollary 3.2, we have

$$\begin{aligned} \left\| S_i C_i^{-1} \begin{pmatrix} I_q \\ 0 \end{pmatrix} \right\| &= \left\| \left(M^{-i} S_0\right) \left(L^{-i} C_0\right)^{-1} \begin{pmatrix} I_q \\ 0 \end{pmatrix} \right\| \\ &= \left\| M^{-i} \left(S_0 C_0^{-1}\right) L^i \begin{pmatrix} I_q \\ 0 \end{pmatrix} \right\|. \end{aligned}$$

Since L is upper triangular, $L^i \begin{pmatrix} I_q \\ 0 \end{pmatrix} = \begin{pmatrix} I_q \\ 0 \end{pmatrix} \left(L_{1:q,1:q}\right)^i$. Then

$$\begin{aligned} \left\| S_i C_i^{-1} \begin{pmatrix} I_q \\ 0 \end{pmatrix} \right\| &\leq \|M^{-i}\| \| \left(L_{1:q,1:q}\right)^i \| \left\| S_0 C_0^{-1} \begin{pmatrix} I_q \\ 0 \end{pmatrix} \right\| \\ &\leq \left(\frac{|\lambda_q|}{|\lambda_{p+1}|} + \eta_i\right)^i \left\| S_0 C_0^{-1} \begin{pmatrix} I_q \\ 0 \end{pmatrix} \right\| \quad \text{as in Corollary 3.2.} \end{aligned}$$

The proof is completed by noting that

$$\begin{aligned}\sin\angle(\widetilde{\mathcal{V}}_1,\mathcal{X}_i) &= \left\|(I-X_iX_i^*)\widetilde{V}_1\right\| \\ &\le \left\|\widetilde{V}_1 - X_iC_i^{-1}\begin{pmatrix} I_q \\ 0\end{pmatrix}\right\| = \left\|S_iC_i^{-1}\begin{pmatrix} I_q \\ 0\end{pmatrix}\right\|,\end{aligned}$$

and that

$$\left\|S_0C_0^{-1}\begin{pmatrix} I_q \\ 0\end{pmatrix}\right\| = \left\|X_0C_0^{-1}\begin{pmatrix} I_q \\ 0\end{pmatrix} - \widetilde{V}_1\right\| = \left\|(I-\mathcal{P})\widetilde{X}_0\right\|.$$

□

Note that this corollary generalizes [21, Thm. 5.2] in the sense that the estimate on $\sin\angle(\widetilde{\mathcal{V}}_1,\mathcal{X}_i)$ deals with invariant subspaces rather than eigenvectors.

3.1. Use of block-GMRES as inner iteration. In this section we restrict attention to the use of block-GMRES as inner solver in Algorithm 1. Block-GMRES belongs to the family of block Krylov subspace methods (see [21]), and it is attractive for large (sparse) linear systems with multiple right-hand sides, as in the case of interest.

Assume that block-GMRES is used to solve a linear system with multiple right-hand sides of the form

$$AZ = B, \quad B \in \mathbb{C}^{n\times p}, \tag{3.3}$$

and that B can be decomposed as

$$B = V_1B_1 + V_2B_2, \tag{3.4}$$

with V_1 and V_2 given in (2.5), $B_1 \in \mathbb{C}^{p\times p}$ nonsingular and $B_2 \in \mathbb{C}^{(n-p)\times p}$. Then we have the following theorem:

THEOREM 3.4. *The residual $B - AZ_k$ associated with the approximate solution Z_k of (3.3) obtained with k iterations of block-GMRES starting with $Z_0 = 0$ is such that*

$$\|B - AZ_k\| \le \min_{p\in\bar{\mathbf{P}}_{k-1}} \|p(M)\|\|\mathcal{S}\|\|L^{-1}\|\|B_2B_1^{-1}\|\|B_1\|, \tag{3.5}$$

where $\bar{\mathbf{P}}_l$ is the set of complex polynomials p_l of degree at most l such that $p_l(0) = 1$ and $\mathcal{S}$ is the Sylvester operator defined in Proposition 2.2.

Proof. Since $Z_k \in \mathcal{R}\left(B \;\; AB \;\; \ldots \;\; A^{k-1}B\right)$ and block-GMRES minimizes the residual, we have

$$\|B - AZ_k\| = \min_{G_1,\ldots,G_k\in\mathbb{C}^{p\times p}} \left\|B + \sum_{i=1}^{k} A^iBG_i\right\|.$$

Let $f_1,\ldots,f_{k-1} \in \mathbb{C}$. Set $F = B_1^{-1}L^{-1}B_1$ and choose

$$G_1 = f_1I - F, \;\; G_i = f_iI - f_{i-1}F, \; i = 2,\ldots,k-1 \;\text{ and }\; G_k = -f_{k-1}F.$$

Then

$$\|B - AZ_k\| \le \min_{f_1,\ldots,f_{k-1}\in\mathbb{C}} \left\|B - ABF + \sum_{i=1}^{k-1} f_iA^i(B - ABF)\right\|.$$

Now observe that the decomposition (3.4) yields

$$\begin{aligned} B - ABF &= V_2 \left(B_2 - MB_2F\right) \\ &= V_2 \left(B_2B_1^{-1}L - MB_2B_1^{-1}\right) L^{-1}B_1 \\ &= -V_2\mathcal{S}(B_2B_1^{-1})\, L^{-1}B_1 \end{aligned}$$

and therefore

$$\|B - AZ_k\| \leq \min_{f_1,\ldots,f_{k-1}\in\mathbb{C}} \left\| \left(I + \sum_{i=1}^{k-1} f_i M^i \right) \mathcal{S}(B_2B_1^{-1})\, L^{-1}B_1 \right\|.$$

The proof is completed by noting that

$$\min_{f_1,\ldots,f_p\in\mathbb{C}} \left\| I + \sum_{i=1}^{k-1} f_i M^i \right\| = \min_{p\in\bar{\mathbf{P}}_{k-1}} \|p(M)\|$$

and that

$$\left\|\mathcal{S}(B_2B_1^{-1})\, L^{-1}B_1\right\| \leq \|\mathcal{S}\|\|B_2B_1^{-1}\|\|L^{-1}\|\|B_1\|.$$

□

Note that the minimum in (3.5) is taken with respect to the matrix M and not A as in the usual theory. Also note that according to Proposition 2.1 the quantity $\|B_2B_1^{-1}\|\|B_1\|$ behaves like the sine of the largest canonical angle between $\mathcal{V}_1$ and $\mathcal{R}(B)$.

To estimate $\min_{p\in\bar{\mathbf{P}}_{k-1}} \|p(M)\|$ we use the following lemma whose proof can be, for example, read off from that of [12, Lemma 1].

PROPOSITION 3.5. *Let E be a convex, closed bounded set in the complex plane, and let ϕ be the conformal mapping from the exterior of E onto the exterior of the unit disk such that $\phi(\infty) = \infty$ and $\phi'(\infty) > 0$. If the numerical range of M is contained in E and $0 \notin E$, then*

$$\min_{p\in\bar{\mathbf{P}}_{k-1}} \|p(M)\| \leq N \left(\frac{1}{|\phi(0)|} \right)^{k-1}, \tag{3.6}$$

with $N = \frac{3\, l(\partial E)}{2\pi d(\partial E)}$ where $l(\partial E)$ is the length of the boundary curve ∂E of E and $d(\partial E)$ is the minimal distance between the numerical range of M and ∂E.

An immediate corollary is as follows.

COROLLARY 3.6. *Let M be perturbed to $M + \delta M$, where $\|\delta M\| < d(\partial E)$, then*

$$\min_{p\in\bar{\mathbf{P}}_{k-1}} \|p(M + \delta M)\| \leq N_\delta \left(\frac{1}{|\phi(0)|} \right)^{k-1}, \tag{3.7}$$

with $N_\delta = \frac{3\, l(\partial E)}{2\pi(d(\partial E) - \|\delta M\|)}$.

A favorable situation for the bound in Proposition 3.5 is when the numerical range of M is well separated from 0. Then $|\phi(0)| \gg 1$ and $\min_{p\in\bar{\mathbf{P}}_{k-1}} \|p(M)\|$ goes to 0 quickly as k increases.

Proposition 3.5 remains valid if the numerical range of M is replaced by the $\epsilon-$pseudospectrum of M defined, for $\epsilon > 0$, by

$$\Lambda_\epsilon(M) = \{\lambda \in \mathbb{C} :\ \sigma_{\min}(\lambda I - M) \leq \epsilon\}.$$

Then the constant N in (3.6) should be replaced by the larger one $N = \frac{3\, l(\partial E)}{2\pi\, \epsilon}$ (see [12]). The advantage here is that the set $\Lambda_\epsilon(M)$ is generally smaller than the numerical range of M (see [29]). Thus the set E can be chosen far from 0 which leads to the favorable condition $|\phi(0)| \gg 1$. For the perturbed case a similar change is needed in (3.7).

A combination of Theorem 3.4 and Proposition 3.5 gives the following result.

THEOREM 3.7. *Let Z_k be the approximate solution of (3.3) obtained with k iterations of block-GMRES starting with $Z_0 = 0$. Under the assumptions of Proposition 3.5, if*

$$k \geq 1 + \frac{1}{\log|\phi(0)|}\left(\log\left(N\|\mathcal{S}\|\|L^{-1}\|\right) + \log\frac{\|B_2B_1^{-1}\|\|B_1\|}{\tau}\right) \tag{3.8}$$

then $\|B - AZ_k\| \leq \tau$.

Note that the bound in (3.8) is only a sufficient condition which guarantees that the norm of the residual be less than τ. It is clear that the required accuracy may be reached for k smaller than the bound (3.8) suggests.

In step 3 of Algorithm 1, the system to be solved by block-GMRES is $AY_i = X_i$. The right-hand side X_i decomposes as in (2.13) which is of the same form as (3.4). In this context, Theorem 3.7 tells us that the residual obtained with k_i iterations of block-GMRES starting with 0 is less than $\tau_i = \delta\|R_i\|$ if

$$\begin{aligned} k_i &\geq 1 + \frac{1}{\log|\phi(0)|}\left(\log\left(N\|\mathcal{S}\|\|L^{-1}\|\right) + \log\frac{\|S_iC_i^{-1}\|\|C_i\|}{\tau_i}\right) \\ &= 1 + \frac{1}{\log|\phi(0)|}\left(\log\left(N\|\mathcal{S}\|\|L^{-1}\|\right) + \log\frac{s_i}{\delta\|R_i\|}\right). \end{aligned} \tag{3.9}$$

The next proposition shows that as $\mathcal{X}_i$ starts to approximate $\mathcal{V}_1$, the ratio $s_i/\|R_i\|$ is bounded independent of i, and thus, the number of inner iterations needed by block-GMRES is bounded independent of i.

PROPOSITION 3.8. *Let X_i be decomposed as*

$$X_i = V_1C_i + V_2S_i \quad \textit{with} \quad \|S_i\| < \epsilon. \tag{3.10}$$

If $0 \leq \epsilon < \min\left(1, -\frac{1}{2} + \frac{1}{2}\sqrt{1 + sep(T_{22}, L)/\|A\|}\right)$, then

$$\frac{s_i}{\|R_i\|} \leq \frac{1+\epsilon}{1-\epsilon}\,\frac{\|\mathcal{P}\|}{sep(T_{22}, L) - 4\|A\|\epsilon(\epsilon+1)}. \tag{3.11}$$

Hence the number of inner iterations needed by block-GMRES to satisfy the tolerance in step 3 of Algorithm 1 is bounded independent of i.

Proof. Note first that the condition on ϵ ensures that $\mathrm{sep}(T_{22}, L) - 4\|A\|\epsilon(\epsilon+1) > 0$.

Using Propositions 2.1 and 2.2, we have

$$s_i \leq \|S_i\|\,\frac{1+\|S_i\|}{1-\|S_i\|} \leq \|\mathcal{P}\|\sin\angle(\mathcal{X}_i, \mathcal{V}_1)\,\frac{1+\epsilon}{1-\epsilon},$$

$$\|R_i\| \geq \mathrm{sep}\,(T_{22}, L_i)\sin\angle(\mathcal{X}_i, \mathcal{V}_1).$$

Then

$$\frac{s_i}{\|R_i\|} \leq \frac{1+\epsilon}{1-\epsilon} \frac{\|\mathcal{P}\|}{\mathrm{sep}(T_{22}, L_i)}.$$

As in Proposition 2.1, the decomposition (3.10) allows us to write $C_i = U_i + \Upsilon_i$ with U_i unitary and $\|\Upsilon_i\| < \epsilon$. Then

$$\begin{aligned}\mathrm{sep}(T_{22}, L_i) &= \mathrm{sep}\,(T_{22}, U_i^* L_i U_i) \\ &\geq \mathrm{sep}(T_{22}, L) - \|L - U_i^* L_i U_i\| \quad (\text{see } [27, p.234]).\end{aligned}$$

Using the decomposition of X_i in $L_i = X_i^* A X_i$ and the expression of C_i given above, we obtain the bound

$$\|L - U_i^* L_i U_i\| \leq 4\|A\|\epsilon(1+\epsilon),$$

from which (3.11) is obtained. □

This proposition is illustrated in Figure 3.1 on a convection diffusion problem (see Example 1 in Section 5 for the details) where after an initial increase in k_i, the number of inner iterations needed at each outer iteration settles down to an approximately constant value.

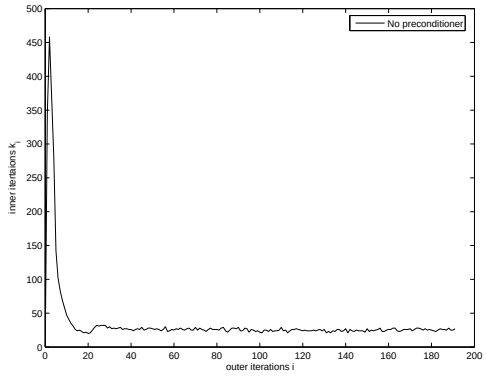

FIG. 3.1. *Outer iterations against inner iterations (Example 1)*

Our aim in the next section is to see if the nice property that k_i is bounded independent of i holds when the system (3.3) is preconditioned.

4. Preconditioning the inexact inverse subspace iteration. A good preconditioner helps to accelerate the computations in step 3 of Algorithm 1 and hence the convergence of this algorithm. A standard way to accomplish this task is to find an approximation P of A such that the systems with the matrix P are cheap to solve. Then the matrix Y_i in step 3 of Algorithm 1 is obtained by applying block-GMRES to the preconditioned block system

$$AP^{-1}Z_i = X_i, \quad Y_i = P^{-1}Z_i. \tag{4.1}$$

Let us denote by Z_{k_i} the approximation of Z_i obtained at iteration k_i of block-GMRES and satisfying

$$\left\|X_i - AP^{-1}Z_{k_i}\right\| \leq \tau_i = \delta \left\|R_i\right\|, \tag{4.2}$$

and so, with $Y_{k_i} = P^{-1}Z_{k_i}$, step 3 in Algorithm 1 is satisfied. The natural question is whether k_i can be bounded independent of i as $\mathcal{X}_i$ approaches $\mathcal{V}_1$, as for the unpreconditioned case in Proposition 3.8.

To answer this question, we attempt to repeat the analysis in the previous section. So, analogously to (2.5), assume that AP^{-1} is block-diagonalized as

$$AP^{-1} = \begin{pmatrix} U_1 & U_2 \end{pmatrix} \begin{pmatrix} K_1 & 0 \\ 0 & K_2 \end{pmatrix} \begin{pmatrix} U_1 & U_2 \end{pmatrix}^{-1}, \quad \text{with } U_i^* U_i = I,\ i = 1, 2 \tag{4.3}$$

where K_1 and K_2 have disjoint spectra. Assume further that X_i is decomposed as

$$X_i = U_1 \widetilde{C}_i + U_2 \widetilde{S}_i \quad \text{with } \|\widetilde{S}_i\| < 1, \tag{4.4}$$

and that analogous hypotheses of Proposition 3.5 and Theorem 3.7 hold. Let

$$\tilde{s}_i = \|\widetilde{S}_i \widetilde{C}_i^{-1}\| \|\widetilde{S}_i\|.$$

The question now is: can the ratio $\tilde{s}_i / \|R_i\|$ be bounded independent of i as $\mathcal{X}_i$ approaches $\mathcal{V}_1$? The a priori answer is 'no', as the following analysis shows.

From Proposition 2.1 and (2.13) and (4.4) we have

$$\tilde{s}_i + s_i \geq \|\widetilde{S}_i\| + \|S_i\| \geq \|U_2 \widetilde{S}_i - V_2 S_i\| = \|U_1 \widetilde{C}_i - V_1 C_i\|$$

and

$$\|U_1 \widetilde{C}_i - V_1 C_i\| \geq \|U_1 \widetilde{C}_i C_i^{-1} - V_1\| \, \sigma_{\min}(C_i).$$

Denoting $\mathcal{U}_1 = \mathcal{R}(U_1)$ and using (2.12) and (2.16), we obtain

$$\|U_1 \widetilde{C}_i - V_1 C_i\| \geq (1 - s_i) \sin \angle(\mathcal{U}_1, \mathcal{V}_1).$$

Therefore

$$\tilde{s}_i \geq \sin \angle(\mathcal{U}_1, \mathcal{V}_1) - s_i \left(1 + \sin \angle(\mathcal{U}_1, \mathcal{V}_1)\right).$$

As $\mathcal{X}_i \to \mathcal{V}_1$, $s_i \to 0$, but there is no reason why $\tilde{s}_i \to 0$. In fact, $\tilde{s}_i / \|R_i\|$ may increase as $\sin \angle(\mathcal{U}_1, \mathcal{V}_1) / \|R_i\|$ leading to a corresponding increase in k_i given by (3.9). Such an increase is shown in Figures 5.1 and 5.5, where an ILU preconditioner is applied to two different examples. The above analysis shows that we do not have a result like (3.11) for preconditioned solves. It also shows that as $\mathcal{X}_i \to \mathcal{V}_1$, a necessary (but not sufficient) condition for a bound similar to (3.11) to hold for preconditioned solves is $\sin \angle(\mathcal{U}_1, \mathcal{V}_1) \approx 0$, that is, $\mathcal{V}_1$ is almost an invariant subspace of AP^{-1}, or equivalently, that $\mathcal{V}_1$ is almost an invariant subspace of P.

4.1. An 'Ideal' Preconditioner. In this subsection we discuss the theoretical case of $\mathcal{U}_1 = \mathcal{V}_1$. We shall see that a preconditioner which satisfies $\mathcal{U}_1 = \mathcal{V}_1$ is

$$\mathbf{P} = AV_1V_1^* + P(I - V_1V_1^*), \tag{4.5}$$

which we call an 'ideal' preconditioner. First, it is easy to see that $\mathbf{P}V_1 = AV_1 = V_1L$. Thus, $\mathcal{V}_1$ is an invariant subspace of both A and $\mathbf{P}$. Moreover, the following proposition shows that if P is a good approximation of A then the spectrum of $A\mathbf{P}^{-1}$ should be clustered near 1.

PROPOSITION 4.1. *Let* $\mathbf{P}$ *be given by (4.5) and assume* A *has the Schur decomposition (2.2). Then the matrix* $A\mathbf{P}^{-1}$ *has the same eigenvalues as the matrix*

$$\begin{pmatrix} I & V_1^* A\mathbf{P}^{-1}V_1^\perp \\ 0 & T_{22}((V_1^\perp)^* P V_1^\perp)^{-1} \end{pmatrix}.$$

Proof. We have

$$\begin{pmatrix} V_1 & V_1^\perp \end{pmatrix}^* A\mathbf{P}^{-1} \begin{pmatrix} V_1 & V_1^\perp \end{pmatrix} = \begin{pmatrix} V_1^* A\mathbf{P}^{-1}V_1 & V_1^* A\mathbf{P}^{-1}V_1^\perp \\ (V_1^\perp)^* A\mathbf{P}^{-1}V_1 & (V_1^\perp)^* A\mathbf{P}^{-1}V_1^\perp \end{pmatrix}.$$

Now observe that $A\mathbf{P}^{-1}V_1 = V_1$ and $(V_1^\perp)^* A = T_{22}(V_1^\perp)^*$. Then

$$\begin{pmatrix} V_1 & V_1^\perp \end{pmatrix}^* A\mathbf{P}^{-1} \begin{pmatrix} V_1 & V_1^\perp \end{pmatrix} = \begin{pmatrix} I & V_1^* A\mathbf{P}^{-1}V_1^\perp \\ 0 & T_{22}(V_1^\perp)^* \mathbf{P}^{-1} V_1^\perp \end{pmatrix}.$$

Finally, since $\mathbf{P}V_1^\perp = PV_1^\perp$, we have

$$\begin{aligned} \left((V_1^\perp)^*\mathbf{P}^{-1}V_1^\perp\right)\left((V_1^\perp)^* P V_1^\perp\right) &= \left((V_1^\perp)^*\mathbf{P}^{-1}V_1^\perp\right)\left((V_1^\perp)^*\mathbf{P}V_1^\perp\right) \\ &= (V_1^\perp)^*\mathbf{P}^{-1}(I - V_1V_1^*)\mathbf{P}V_1^\perp = I. \end{aligned}$$

Hence, $\left(V_1^\perp)^*\mathbf{P}^{-1}V_1^\perp\right) = \left((V_1^\perp)^* P V_1^\perp\right)^{-1}$. □

If P is a good approximation of A, then $(V_1^\perp)^* P V_1^\perp$ will be a good approximation of T_{22}, and hence the eigenvalues of $A\mathbf{P}^{-1}$ should be clustered around 1.

Now, assume that $\mathcal{V}_1$ is a simple invariant subspace of $A\mathbf{P}^{-1}$. This ensures the existence of a block-diagonalization of the form

$$A\mathbf{P}^{-1} = \begin{pmatrix} V_1 & U \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & K \end{pmatrix} \begin{pmatrix} V_1 & U \end{pmatrix}^{-1} \text{ with } U^*U = I, \ \operatorname{sep}(I, K) > 0. \quad (4.6)$$

Assume also that X_i can decomposed, for all $i \geq 0$, in the form

$$X_i = V_1\widetilde{C}_i + U\widetilde{S}_i \text{ with } \widetilde{\epsilon}_i := \|\widetilde{S}_i\| < 1. \quad (4.7)$$

Multiplying (2.13) and (4.7) on the left by W_2^* gives $S_i = W_2^* U \widetilde{S}_i$. It is easy to see that W_2^*U is nonsingular and therefore that $\widetilde{S}_i = (W_2^*U)^{-1} S_i$. Assume $\mathcal{X}_i \to \mathcal{V}_1$, so that $\widetilde{\epsilon}_i \to 0$ and there exists $\widetilde{\epsilon} < 1$ such that $\widetilde{\epsilon}_i \leq \widetilde{\epsilon}$, $\forall i \geq 0$. Then, from Proposition 2.1, we have

$$\begin{aligned} \tilde{s}_i &\leq \frac{1 + \|\widetilde{S}_i\|}{1 - \|\widetilde{S}_i\|} \|\widetilde{S}_i\| \\ &\leq \frac{1 + \widetilde{\epsilon}}{1 - \widetilde{\epsilon}} \|(W_2^*U)^{-1}\| \, s_i. \end{aligned}$$

Now a proof similar to that of Proposition 3.8 shows that $s_i/\|R_i\|$, and therefore that $\tilde{s}_i/\|R_i\|$, is bounded independent of i. This analysis shows that if the ideal preconditioner were available, then we would be able to show that the iterations used by block-GMRES should be independent of i as in Proposition 3.8.

4.2. The 'Tuned' Preconditioner. Of course, the ideal preconditioner cannot be used in practice since V_1 is unknown, so we replace $\mathbf{P}$ by the "tuned" preconditioner

$$\mathbf{P}_i = AX_iX_i^* + P(I - X_iX_i^*), \tag{4.8}$$

where the V_1 in (4.5) is replaced by X_i computed by Algorithm 1. This preconditioner satisfies the tuning condition

$$\mathbf{P}_iX_i = AX_i.$$

This is a generalisation of the condition proposed in [8] and [7] in the context of inexact inverse iteration, but the motivation given here is different. Note that the tuned preconditioner changes at each iteration i of Algorithm 1 and its quality improves with that of X_i. Since $\mathbf{P}_iX_i = AX_i = X_iL_i + R_i$. We see that as $\mathcal{X}_i \to \mathcal{V}_1$, $R_i \to 0$, and so, in the limit, $\mathbf{P}_i$ has $\mathcal{V}_1$ as a invariant subspace. Also, the tuning condition can be written

$$A\mathbf{P}_i^{-1}(AX_i) = AX_i, \tag{4.9}$$

which means that $A\mathcal{X}_i$ is an invariant subspace of $A\mathbf{P}_i^{-1}$ corresponding to the eigenvalue 1, which is a property shared with the ideal preconditioner given by (4.5). So the tuned preconditioner also has the nice property of clustering around 1 at least a part of the spectrum of $A\mathbf{P}_i^{-1}$. Asymptotically, that is, when $\mathcal{X}_i \to \mathcal{V}_1$, the tuned preconditioner $\mathbf{P}_i$ will behave like the ideal preconditioner.

We now prove a result for the tuned preconditioner corresponding to that given by Proposition 3.8 for unpreconditioned solves, namely, that the number of inner iterations needed to achieve (4.2) will be independent of i. This is to be expected given the closeness of $A\mathbf{P}_i^{-1}$ to $A\mathbf{P}^{-1}$, and is exactly what is observed in the numerical examples discussed in Section 5.

Assume that X_i decomposes as

$$X_i = V_1C_i + V_2S_i \quad \text{with} \ \ \|S_i\| \to 0 \ \ \text{as} \ \ i \to \infty \tag{4.10}$$

and define ϵ_i by

$$\epsilon_i = \max_{j \geq i} \|S_j\|. \tag{4.11}$$

Note that the sequence ϵ_i is decreasing.

In order to prove a result similar to that of Proposition 3.8 for the tuned preconditioned system, i.e. system (4.1) with $\mathbf{P}_i$ as preconditioner, we need the following three lemmas.

LEMMA 4.2. *We have*

$$X_iX_i^* = V_1V_1^* + E_i, \quad \textit{with} \ \ \|E_i\| \leq \|S_i\| \leq \epsilon_i, \tag{4.12}$$

and for ϵ_i small enough, there exist two positive constants α and β independent of i such that

$$\alpha \ \sin\angle(A\mathcal{X}_i, \mathcal{X}_i) \leq \|R_i\|, \tag{4.13}$$

$$\|A\mathbf{P}^{-1} - A\mathbf{P}_i^{-1}\| \leq \beta \ \epsilon_i. \tag{4.14}$$

Proof. We use the same notation as in the proof of Proposition 3.8. The property (4.12) is a consequence of Proposition 2.1 and the fact that $\|E_i\| = \sin\angle(\mathcal{X}_i, \mathcal{V}_1)$ (see (2.11)).

The columns of $AX_i\,(L_i^*L_i + R_i^*R_i)^{-\frac{1}{2}}$ form an orthonormal basis of $A\mathcal{X}_i$. Therefore

$$\begin{aligned}\sin\angle(A\mathcal{X}_i, \mathcal{X}_i) &= \|(I - X_iX_i^*)AX_i\,(L_i^*L_i + R_i^*R_i)^{-\frac{1}{2}}\,\| \\ &= \|R_i\,(L_i^*L_i + R_i^*R_i)^{-\frac{1}{2}}\,\| \\ &\le \|R_i\|/\sigma_{\min}(L_i),\end{aligned}$$

and as in Proposition 3.8,

$$\sigma_{\min}(L_i) \ge \sigma_{\min}(L) - 4\|A\|\epsilon_i(1+\epsilon_i).$$

Then since ϵ_i is decreasing, there exists $\alpha > 0$ independent of i, such that for ϵ_i small enough

$$\sigma_{\min}(L) - 4\|A\|\epsilon_i(1+\epsilon_i) \ge \alpha$$

and then

$$\sin\angle(A\mathcal{X}_i, \mathcal{X}_i) \le \|R_i\|/\alpha.$$

From (4.5), (4.8) and (4.12) we have $\mathbf{P}_i = \mathbf{P} + (A-P)E_i$. Then

$$\begin{aligned}A\mathbf{P}^{-1} - A\mathbf{P}_i^{-1} &= A\mathbf{P}^{-1}\,(\mathbf{P}_i - \mathbf{P})\,\mathbf{P}_i^{-1} \\ &= A\mathbf{P}^{-1}\,(A-P)\,E_i\,(\mathbf{P} + (A-P)E_i)^{-1} \\ \|A\mathbf{P}^{-1} - A\mathbf{P}_i^{-1}\| &\le \|A\mathbf{P}^{-1}\|\|A-P\|\|E_i\|\|\mathbf{P}^{-1}\|\,\left\|\left(I + \mathbf{P}^{-1}(A-P)E_i\right)^{-1}\right\| \\ &\le \frac{\|A\mathbf{P}^{-1}\|\|A-P\|\|\mathbf{P}^{-1}\|}{1 - \|\mathbf{P}^{-1}(A-P)\|\,\|E_i\|}\|E_i\|\end{aligned}$$

and the same argument used for α shows the existence of β independent of i such that

$$\frac{\|A\mathbf{P}^{-1}\|\|A-P\|\|\mathbf{P}^{-1}\|}{1 - \|\mathbf{P}^{-1}(A-P)\|\,\|E_i\|} \le \frac{\|A\mathbf{P}^{-1}\|\|A-P\|\|\mathbf{P}^{-1}\|}{1 - \|\mathbf{P}^{-1}(A-P)\|\,\epsilon_i} \le \beta.$$

□

The following lemma shows that under some natural hypotheses, $A\mathbf{P}_i^{-1}$ will have a block-diagonalization close to that of $A\mathbf{P}^{-1}$ given in (4.6).

LEMMA 4.3. *Assume that $\mathcal{V}_1$ is a simple invariant subspace of $A\mathbf{P}^{-1}$. Then for ϵ_i small enough, $A\mathbf{P}_i^{-1}$ can be block-diagonalized as*

$$A\mathbf{P}_i^{-1} = \begin{pmatrix} \widetilde{U}_i & U_i \end{pmatrix}\begin{pmatrix} I & 0 \\ 0 & K_i \end{pmatrix}\begin{pmatrix} \widetilde{U}_i & U_i \end{pmatrix}^{-1} \tag{4.15}$$

with $\widetilde{U}_i^*\widetilde{U}_i = I$ and $U_i^*U_i = I$ and

$$\begin{aligned}\|K - K_i\| &\le c_1\epsilon_i, \\ \sin\angle(\mathcal{U},\mathcal{U}_i) &\le c_2\epsilon_i \qquad (\mathcal{U} = \mathcal{R}(U) \text{ and } \mathcal{U}_i = \mathcal{R}(U_i)), \\ \sin\angle(\mathcal{V}_1,\widetilde{\mathcal{U}}_i) &\le c_3\epsilon_i \qquad (\widetilde{\mathcal{U}}_i = \mathcal{R}(\widetilde{U}_i) = A\mathcal{X}_i)\end{aligned}$$

where c_1, c_2 and c_3 are positive constants independent of i.

Proof. Since $\mathcal{V}_1$ is a simple invariant subspace, we know that the block-diagonalization (4.6) exists and [27, Thm. 2.8] and Lemma 4.2 can be used to compare the invariant subspaces of $A\mathbf{P}^{-1}$ and $A\mathbf{P}_i^{-1}$. Thus for ϵ_i sufficiently small, there exist matrices U_i and K_i such that

$$A\mathbf{P}_i^{-1}U_i = U_iK_i \quad \text{with} \quad U_i^*U_i = I$$

and positive constants c_1 and c_2 independent of i such that

$$\|K - K_i\| \leq c_1\epsilon_i,$$
$$\sin\angle(\mathcal{U},\mathcal{U}_i) \leq c_2\epsilon_i.$$

From (4.9) it is clear that

$$\widetilde{U}_i = AX_i\left((AX_i)^*(AX_i)\right)^{-1/2}$$

satisfies $A\mathbf{P}_i^{-1}\widetilde{U}_i = \widetilde{U}_i$ and $\widetilde{U}_i^*\widetilde{U}_i = I$. Moreover, there exists c_3 independent of i such that, for ϵ_i sufficiently small,

$$\sin\angle(\mathcal{V}_1,\widetilde{\mathcal{U}}_i) \equiv \sin\angle(\mathcal{V}_1, A\mathcal{X}_i) \leq c_3\epsilon_i,$$

because $\sin\angle(\mathcal{V}_1, A\mathcal{X}_i) \leq \sin\angle(\mathcal{V}_1,\mathcal{X}_i) + \sin\angle(\mathcal{X}_i, A\mathcal{X}_i) \leq \epsilon_i + \|R_i\|/\alpha$, using(4.13), and

$$\|R_i\| = \|(I - X_iX_i^*)AX_i\| \leq \|M\|\|S_i\| + \|E_i\|\|A\| \leq 2\|A\|\epsilon_i.$$

Since 1 is not an eigenvalue of K, then for ϵ_i sufficiently small 1 cannot be an eigenvalue of K_i. This shows the existence of the decomposition (4.15). ∎

The next lemma shows the continuous dependence of a spectral projection on the matrix.

LEMMA 4.4. *Let B and C be two matrices of the same size and $P_\gamma(B)$ and $P_\gamma(C)$ the spectral projections onto the invariant subspaces of B and C corresponding to the eigenvalues inside a closed contour γ. Assume that $\|B - C\| \leq \xi$ and let $m_\gamma(B) = \max_{\lambda\in\gamma}\|(\lambda I - B)^{-1}\|$. If $\xi m_\gamma(B) < 1$ then*

$$\|P_\gamma(B) - P_\gamma(C)\| \leq \frac{1}{2\pi}l_\gamma\frac{\xi m_\gamma^2(B)}{1 - \xi m_\gamma(B)}$$

where l_γ is the length of γ.

Proof. see e.g. [9, §8.2]. ∎

We are now in a position to state and prove the key result in this paper.

THEOREM 4.5. *Let X_i be decomposed as in (4.10)-(4.11) where ϵ_i is also defined. Assume that $\mathcal{V}_1$ is a simple invariant subspace of $A\mathbf{P}^{-1}$ and that the numerical range of the matrix K in (4.6) satisfies the assumptions of Proposition 3.5.*

Then, for small enough ϵ_i, the number k_i of inner iterations used by block-GMRES to compute Z_{k_i} satisfying the stopping criterion

$$\left\|X_i - A\mathbf{P}_i^{-1}Z_{k_i}\right\| \leq \tau_i = \delta\left\|R_i\right\|, \tag{4.16}$$

is bounded independent of i.

Proof. Let ϕ and E be given by Proposition 3.5 applied to K (instead of M).

For small enough ϵ_i, Lemma 4.3 shows that the decomposition (4.15) holds and Corollary 3.6 can be used with K_i to obtain a constant $\widehat{N}$ independent of i such that

$$\min_{p\in\widetilde{\mathbf{P}}_{k-1}} \|p(K_i)\| \leq \widehat{N}\left(\frac{1}{|\phi(0)|}\right)^{k-1}.$$

Decompose X_i in $\mathcal{R}\left(\widetilde{U}_i \;\; U_i\right)$as

$$X_i = \widetilde{U}_i\widehat{C}_i + U_i\widehat{S}_i$$

and, for ϵ_i small enough, define $\widehat{s}_i$ by $\widehat{s}_i = \|\widehat{S}_i\widehat{C}_i^{-1}\|\|\widehat{C}_i\|$.

It is a simple task to show that the residual obtained with k_i iterations of block-GMRES starting with 0 is less than $\tau_i = \delta\|R_i\|$ if

$$k_i \geq 1 + \frac{1}{\log|\phi(0)|}\left(\log\left(\widehat{N}\|\widehat{\mathcal{S}}_i\|\right) + \log\frac{\widehat{s}_i}{\delta\|R_i\|}\right), \tag{4.17}$$

where $\|\widehat{\mathcal{S}}_i\| = \|I - K_i\| \leq \|I - K\| + c_1\epsilon_i$ can be bounded independent of i since ϵ_i is decreasing.

Now in order to show that k_i can be bounded independent of i for small enough ϵ_i, it only remains to show that the ratio $\widehat{s}_i/\|R_i\|$ possesses this property.

From Proposition 2.1 we have

$$\|\widehat{S}_i\| \leq \|\widehat{\mathcal{P}}_i\|\sin\angle(\widetilde{\mathcal{U}}_i, \mathcal{X}_i) \equiv \|\widehat{\mathcal{P}}_i\|\sin\angle(A\mathcal{X}_i, \mathcal{X}_i)$$

where $\widehat{\mathcal{P}}_i$ is the spectral projection of $A\mathbf{P}_i^{-1}$ onto $\widetilde{\mathcal{U}}_i$.

We have

$$\sin\angle(A\mathcal{X}_i, \mathcal{X}_i) \leq \sin\angle(A\mathcal{X}_i, \mathcal{V}_1) + \sin\angle(\mathcal{V}_1, \mathcal{X}_i) \leq (c_3+1)\epsilon_i.$$

The term $\|\widehat{\mathcal{P}}_i\|$ is bounded as

$$\|\widehat{\mathcal{P}}_i\| \leq \|\widehat{\mathcal{P}} - \widehat{\mathcal{P}}_i\| + \|\widehat{\mathcal{P}}\|$$

where $\widehat{\mathcal{P}}$ is the spectral projection of $A\mathbf{P}^{-1}$ onto $\mathcal{V}_1$. For small enough ϵ_i, (4.14) shows that lemma 4.4 can be applied. Taking, in this lemma, γ as the circle of center 1 and radius ϵ_i, we obtain:

$$\|\widehat{\mathcal{P}} - \widehat{\mathcal{P}}_i\| \leq \frac{\beta m_\gamma^2(A\mathbf{P}^{-1})}{1 - \beta m_\gamma(A\mathbf{P}^{-1})\epsilon_i}\,\epsilon_i.$$

Since ϵ_i is decreasing, we have for ϵ_i small enough

$$\frac{\beta m_\gamma^2(A\mathbf{P}^{-1})}{1 - \beta m_\gamma(A\mathbf{P}^{-1})\epsilon_i} \leq c_4$$

with c_4 independent of i and hence

$$\|\widehat{S}_i\| \leq (c_4\epsilon_i + \|\widehat{\mathcal{P}}\|)(c_3+1)\epsilon_i \leq c_5\epsilon_i \text{ with } c_5 = \left(\|\widehat{\mathcal{P}}\| + c_4\right)(c_3+1).$$

Finally from Proposition 2.1 and Lemma 4.2, we have for ϵ_i small enough

$$\begin{aligned}\frac{\widehat{s}_i}{\|R_i\|} &\leq \frac{1+\|\widehat{S}_i\|}{1-\|\widehat{S}_i\|}\ \frac{\|\widehat{S}_i\|}{\alpha\ \sin\angle(A\mathcal{X}_i,\mathcal{X}_i)}\\ &\leq \frac{1+\|\widehat{S}_i\|}{1-\|\widehat{S}_i\|}\ \frac{\|\widehat{\mathcal{P}}_i\|}{\alpha}\\ &\leq \frac{1+c_5\epsilon_i}{1-c_5\epsilon_i}\ \frac{c_4\epsilon_i+\|\widehat{\mathcal{P}}\|}{\alpha}.\end{aligned}$$

Since ϵ_i is decreasing, the last inequality shows that the ratio $\frac{\widehat{s}_i}{\|R_i\|}$ is bounded independent of i for small enough ϵ_i. □

The numerical results illustrate this theorem, namely that the number of iterations is asymptotically independent of i, see Figures 5.1 and 5.5.

5. Numerical tests. In this section we present some numerical tests to illustrate the performance of Algorithm 1 when step 3 is replaced by the preconditioned block system

$$AP^{-1}Z_i = X_i,\quad Y_i = P^{-1}Z_i, \tag{5.1}$$

solved by block-GMRES with the tolerance $\tau_i = \min(\delta, \delta\|R_i\|)$, $\delta = 10^{-3}$.

Any version of block-GMRES can be used to illustrate the theory. We have chosen to use a new variant of block-GMRES which detects the near-dependance in the corresponding block-Arnoldi basis and then adapts the block sizes accordingly. As a consequence, this variant selects appropriate directions for convergence. See [19] for the details.

We compare two preconditioners:

- ILU preconditioner: P is obtained from the incomplete LU factorization of A with a drop tolerance fixed at 10^{-1}.
- Tuned preconditioner: $\mathbf{P}_i = P + F_i X_i^*$ where $F_i = AX_i - PX_i$ and P is as above. In this case the computation of $\mathbf{P}_i^{-1}$ in $Y_i = \mathbf{P}_i^{-1}Z_i$ uses the Woodbury formula (see [10]). Note that the application of $\mathbf{P}_i^{-1}$ within block-GMRES requires little extra work compared with the application of P^{-1}, with the additional work mainly needed at the outer step.

For each example, we give information on the spectrum of A, the block size p, the dimension q of the computed invariant subspace $\widetilde{\mathcal{V}}_1$ associated to the eigenvalues near 0. We show the inner iterations for the two preconditioners and the norm of the residuals, denoted by Γ_i, associated to the computed invariant subspaces.

Example 1. A is obtained with a five-point stencil and centered difference discretization of the convection diffusion operator (see [11]):

$$\begin{cases}\mathcal{A}u = \Delta u + 10\frac{\partial u}{\partial x} + 10\frac{\partial u}{\partial y}, & \text{on } \Omega = [0,1]\times[0,1],\\ u = 0 & \text{on } \partial\Omega.\end{cases}$$

The matrix A is of order $n = 2025$ and has $nz = 9945$ nonzero elements, $\|A\| = 16152$, $\|A - P\| = 1400$. We use $p = 6$ and look for the invariant subspace $\widetilde{\mathcal{V}}_1$ of dimension $q = 4$. The computations stop when $\|\Gamma_i\| < 10^{-8}$. The spectrum of A and the computed eigenvalues are shown in Figure 5.4.

Figure 5.1 shows the inner iterations k_i for the two preconditioners and figures 5.2 and 5.3 show the behavior of $\|\Gamma_i\|$ during the outer iterations and compared with the

total number of inner iterations. Figure 5.1 illustrates well the theory: it shows that as the outer convergence proceeds, the number of inner iterations becomes independent of i when the tuned preconditioner is used but increases when the standard ILU preconditioner is used. Figure 5.2 illustrates that there is little difference in the performance of the two preconditioners with regard to the residual norms in step 3 of Algorithm 1. Figure 5.3 shows the dramatic improvement in overall cost achieved by the tuned preconditioner, with the required tolerance being achieved at 12.65% of the cost needed for the untuned preconditioner.

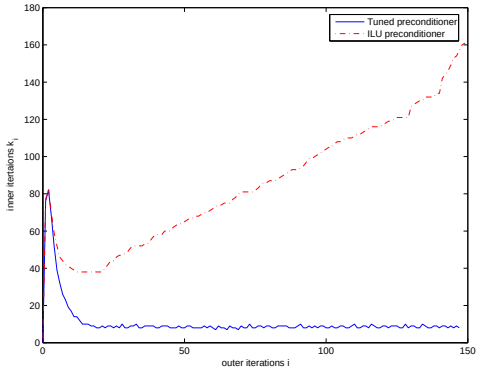

FIG. 5.1. *Outer iterations against inner iterations (Example 1)*

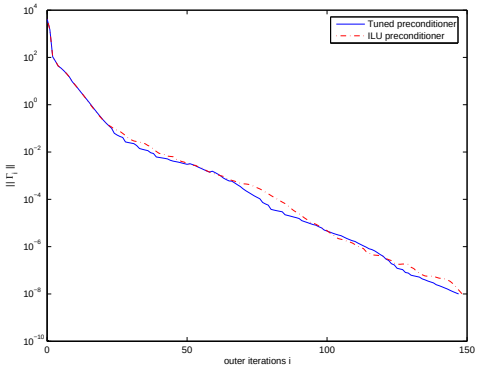

FIG. 5.2. *Residual norms against outer iterations (Example 1)*

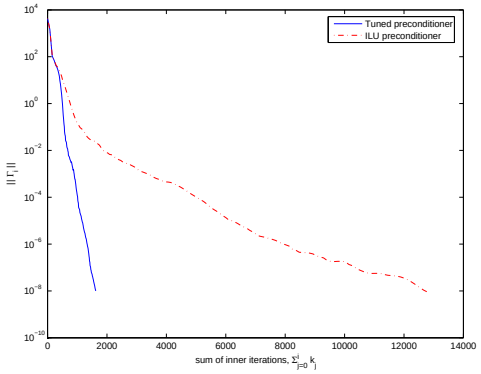

FIG. 5.3. *Residual norms against the total number of inner iterations (Example 1)*

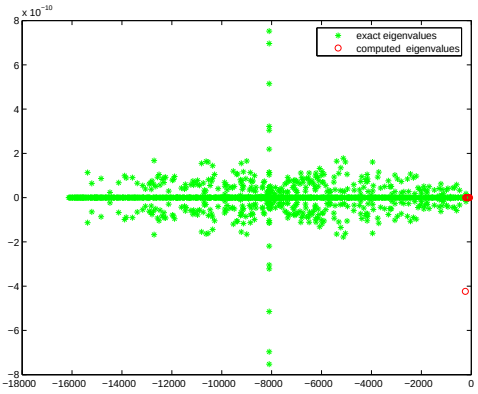

FIG. 5.4. *Eigenvalues A (Example 1)*

Example 2. A is the matrix $QC2534$ from the NEP set [1]. This matrix is complex, symmetric and non-Hermitian. It is of order $n = 2534$ and has $nz = 463360$ nonzero elements, $\|A\| = 3.32$, $\|A - P\| = 0.41$. We use $p = 16$ and look for the invariant subspace $\widetilde{\mathcal{V}}_1$ of dimension $q = 10$. The computations stop when $\|\Gamma_i\| < 10^{-8}$. Figure 5.8 shows the spectrum of A and the computed eigenvalues. Figure 5.5 compares the number of inner iterations for the ILU and tuned preconditioners. Figures 5.6 and 5.7 show the norm of the residual of the invariant subspace associated to the q eigenvalues

[1] see http://math.nist.gov/MatrixMarket/collections/NEP.html

near 0. Similar comments apply as in Example 1. The tuned preconditioner requires a roughly constant number of inner iterations per outer iteration, (see Figure 5.5). Finally, the overall costs for the tuned preconditioned system to achieve $\|\Gamma_i\| < 10^{-8}$ are about 36.18% of the costs for the untuned preconditioner, (see Figure 5.7).

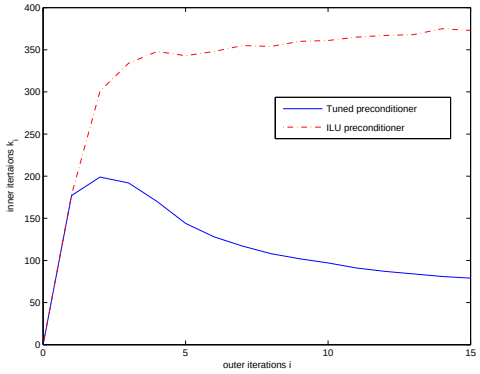

FIG. 5.5. *Outer iterations against inner iterations (Example 2)*

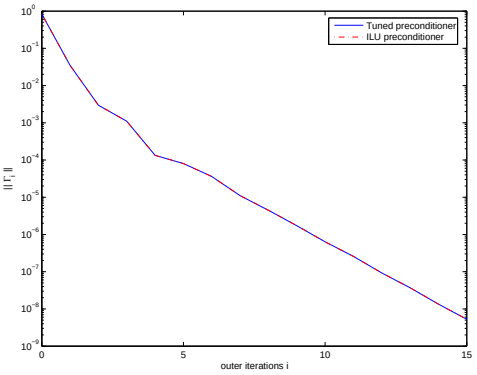

FIG. 5.6. *Residual norms against outer iterations (Example 2)*

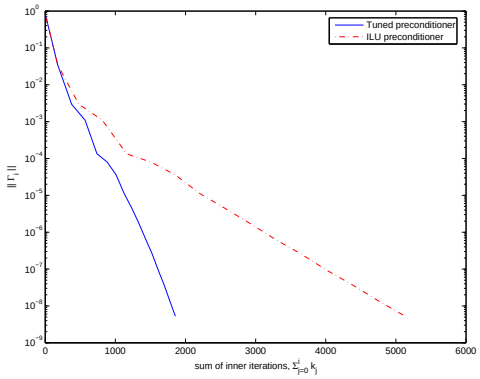

FIG. 5.7. *Residual norms against the total number of inner iterations (Example 2)*

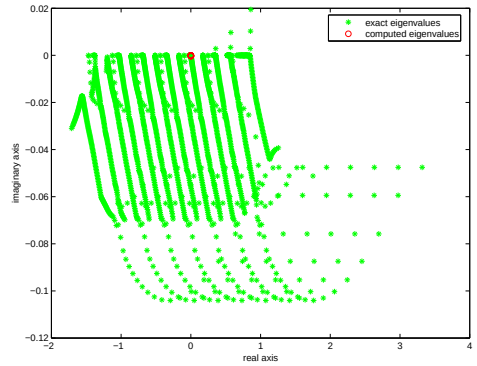

FIG. 5.8. *Eigenvalues of A (Example 2)*

REFERENCES

[1] P.A. ABSIL, C.G. BAKER AND K.A. GALLIVAN, *A truncated-CG style method for symmetric generalized eigenvalue problems,* J. Comput. Appli. Math., 189 (2006), pp. 274–285.

[2] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE AND H. VAN DER VORST, *Templates for the solution of Algebraic Eigenvalue Problems* SIAM, Philadelphia, (2000).

[3] J. BRANDTS, *The Riccati algorithm for eigenvalues and invariant subspaces of matrices with inexpensive action*, Linear Algebra and its Applications, 358 (2003), pp. 335–365.

[4] J. BERNS-MÜLLER, I.G. GRAHAM AND A. SPENCE, *Inexact inverse iteration for symmetric matrices.* Linear Algebra and its Applications, 416 (2006), pp. 389–413.

[5] J. BERNS-MÜLLER AND A. SPENCE, *Inexact inverse iteration with variable shift for nonsymmetric eigenvalue problems.* SIAM J. Matrix Analysis and Applications (2006), to appear.

[6] J. BERNS-MÜLLER AND A. SPENCE, *Inexact inverse iteration and GMRES.* (2006), submitted.

[7] M.A. FREITAG AND A. SPENCE, *A tuned preconditioner for inexact inverse iteration applied to Hermitian eigenvalue problems.* (2006), submitted.

[8] M.A. Freitag and A. Spence, *Convergence of inexact inverse iteration with application to preconditioned inexact solves.* BIT, (2006), accepted.
[9] S. K. Godunov, *Modern Aspects of Linear Algebra*, Translations of Mathematical Monographs, 175. Amer. Math. Soc., Providence, RI, 1998.
[10] G.H. Golub and Ch. Van Loan, *Matrix computations*, The Johns Hopkins University Press, the 2nd edition, 1989.
[11] G.H. Golub and Q. Ye, *Inexact inverse iteration for generalised eigenvalue problems,* BIT, 40 (2000), pp. 671–684.
[12] M. Hochbruck and C. Lubich, *On Krylov subspace approximations to the matrix exponential operator,* SIAM J. Numer. Anal., 34 (1997), pp. 1911–1925.
[13] A.V. Knyazev, *Towards the optimal preconditioned eigensolver: locally optimal block preconditioned conjugate gradient method.* SIAM J. Scientific Computing, 23 (2001), pp.517–541.
[14] Y-L, Lai, K-Y. Lin and L. Wen-Wei, *An inexact inverse iteration for large sparse eigenvalue problems.* Numer. Linear Algebra Applic., 1 (1997), pp. 1–13.
[15] Y. Notay, *Combination of Jacobi-Davidson and conjugate gradients for the partial symmetric eigenproblem,* Numer. Linear Algebra Appl., 9 (2002), pp. 21–44.
[16] Y. Notay, *Inner iterations in eigenvalue solvers,* Report GANMN 05-01, June 2005.
[17] B.N. Parlett, *The symmetric eigenvalue problem*, Classics in Applied Mathematics, SIAM, Philadelphia, 1998.
[18] M. Robbé and M. Sadkane, *Riccati-based preconditioner for computing invariant subspaces of large matrices*, Numer. Math., 92 (2002), pp. 129-159.
[19] M. Robbé and M. Sadkane, *Exact and inexact breakdowns in the block GMRES method*, Linear Algebra Appl., 49 (2006), pp. 265–285.
[20] Y. Saad, *Numerical methods for large eigenvalue problems,* Manchester University Press, Manchester, 1992.
[21] Y. Saad, *Iterative methods for sparse linear systems,* PWS Publishing, Boston, 1996.
[22] A. Sameh and Z. Tong, *The trace minimization method for the symmetric generalized eigenvalue problem,* J. Comput. Appli. Math., 123 (2000), pp. 155–175.
[23] V. Simoncini and Eldén, *Inexact Rayleigh quotient-type methods for eigenvalue computations,* BIT, 42 (2002), pp. 159–182.
[24] G.L.G. Sleijpen and H. van der Vorst, *A Jacobi-Davidson iteration method for linear eigenvalue problems,* SIAM J. Matrix Analysis, 17 (1996), pp.401–425.
[25] P. Smit and M.H.C. Paardekooper, *The effects of inexact solvers in algorithms for symmetric eigenvalue problems,* Linear Algebra Appl., 287 (1999), pp 337–357.
[26] G.W. Stewart, *Simultaneous iteration for computing invariant subspaces of non-Hermitian matrices,* Numer. Math., 25 (1976),pp. 123–136.
[27] G.W. Stewart and J.-G. Sun, *Matrix Perturbation Theory,* Academic Press, San Diego, California, 1990.
[28] G.W. Stewart, *Matrix Algorithms. Volume II: Eigensystems,* SIAM, Philadelphia (2001).
[29] L.N. Trefethen and M. Embree, *Spectra and Pseudospectra : The behavior of nonnormal matrices and operators,* Princeton University Press, 2005.