

CONSTRAINT PRECONDITIONING FOR INDEFINITE LINEAR SYSTEMS*

CARSTEN KELLER[†], NICHOLAS I. M. GOULD[‡], AND ANDREW J. WATHEN[†]

Abstract. The problem of finding good preconditioners for the numerical solution of indefinite linear systems is considered. Special emphasis is put on preconditioners that have a 2×2 block structure and that incorporate the $(1, 2)$ and $(2, 1)$ blocks of the original matrix. Results concerning the spectrum and form of the eigenvectors of the preconditioned matrix and its minimum polynomial are given. The consequences of these results are considered for a variety of Krylov subspace methods. Numerical experiments validate these conclusions.

Key words. preconditioning, indefinite matrices, Krylov subspace methods

AMS subject classifications. 65F10, 65F15, 65F50

PII. S0895479899351805

1. Introduction. In this paper, we are concerned with investigating a new class of preconditioners for indefinite systems of linear equations of a sort which arise in constrained optimization as well as in least-squares, saddle-point, and Stokes problems. We attempt to solve the indefinite linear system

$$(1.1) \quad \underbrace{\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix}}_A \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \underbrace{\begin{bmatrix} b_1 \\ b_2 \end{bmatrix}}_b,$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and $B \in \mathbb{R}^{m \times n}$. Throughout the paper we shall assume that $m \leq n$ and that A is nonsingular, in which case B must be of full rank.

Example 1.1 (quadratic programming problem). Consider the problem of minimizing a function of n variables subject to m linear equality constraints on the variables, i.e.,

$$(1.2) \quad \begin{aligned} \underset{x}{\text{minimize}} \quad & f(x) = \frac{1}{2}x^T Ax - c^T x \\ \text{subject to} \quad & Bx = d. \end{aligned}$$

Any finite solution to (1.2) is a stationary point of the Lagrangian function

$$L(x, \lambda) = \frac{1}{2}x^T Ax - c^T x + \lambda^T (Bx - d),$$

where the λ_i are referred to as *Lagrangian multipliers*. On differentiating L with respect to x and λ the solution of (1.2) is readily seen to satisfy $n + m$ linear equations of the form (1.1), with $x_1 = x$, $x_2 = \lambda$, $b_1 = c$, and $b_2 = d$. For this application these are known as the Karush–Kuhn–Tucker (KKT) conditions.

Example 1.2 (saddle-point problems). Mixed finite element approximations of variational problems are expressible in the following common form.

*Received by the editors February 1, 1999; accepted for publication (in revised form) by M. L. Overton August 31, 1999; published electronically April 18, 2000.

<http://www.siam.org/journals/simax/21-4/35180.html>

[†]Oxford University Computing Laboratory, Parks Road, Oxford OX1 3QD, UK (ck@comlab.ox.ac.uk, wathen@comlab.ox.ac.uk).

[‡]Rutherford Appleton Laboratory, Chilton, Didcot, OX11 0QX, UK (N.I.M.Gould@rl.ac.uk).

Find u and p in certain spaces such that

$$(1.3) \quad \begin{aligned} a(u, v) + b(v, p) &= \langle f, v \rangle, \\ b(u, q) &= \langle g, q \rangle. \end{aligned}$$

For appropriate v and q (1.3) leads to systems of the form (1.1), where x_1 are the coefficients of the approximation of u and x_2 are the coefficients of the approximation of p with respect to chosen bases. See, for example, Quarteroni and Valli [23, Chapters 7, 9].

Krylov subspace methods, which are among the most important iterative methods currently available, apply techniques that involve orthogonal projections onto subspaces of the form

$$\mathcal{K}(\mathcal{A}, b) \equiv \text{span}\{b, \mathcal{A}b, \mathcal{A}^2b, \dots, \mathcal{A}^{n-1}b, \dots\}.$$

The most common schemes that use this idea are the method of conjugate gradients (CG) for symmetric positive definite matrices, the method of minimum residuals (MINRES) for symmetric and possibly indefinite matrices, and the generalized minimum residual method (GMRES) for unsymmetric matrices, although many other methods are available; see, for example, Greenbaum [13].

One common feature of the above-mentioned methods is that the solution of the linear system (1.1) is found within $n+m$ iterations in exact arithmetic; see Joubert and Manteuffel [15, p. 152]. For very large (and possibly sparse) linear systems this upper limit on the number of iterations is often not practical. The idea of preconditioning attempts to improve on the spectral properties, i.e., the clustering of the eigenvalues, such that the total number of iterations required to solve the system to within some tolerance is decreased substantially.

In this paper, we are specifically concerned with nonsingular preconditioners of the form

$$(1.4) \quad \mathcal{G} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$ approximates, but is not the same as, A . The inclusion of the exact representation of the (1, 2) and (2, 1) matrix blocks in the preconditioner, which are often associated with constraints (see Example 1.1), leads one to hope for a more favorable distribution of the eigenvalues of the (left-)preconditioned linear system

$$(1.5) \quad \mathcal{G}^{-1} \mathcal{A}x = \mathcal{G}^{-1}b.$$

Since these blocks are unchanged from the original system, we shall call \mathcal{G} a *constraint preconditioner*. A preconditioner of the form \mathcal{G} has recently been used by Lukšan and Vlček [16] in the context of constrained nonlinear programming problems; see also Coleman [7], Polyak [21], and Gould, Hribar, and Nocedal [12]. Here, we derive arguments that confirm and extend some of the results in [16] and highlight the favorable features of a preconditioner of the form \mathcal{G} . Note that Golub and Wathen [11] recently considered a symmetric preconditioner of the form (1.4) for problems of the form (1.1), where A is nonsymmetric.

We comment that for certain partial differential equation problems, which give rise to linear systems of the form (1.1), extremely effective positive definite preconditioners have been proposed. These make use of particular structures in the underlying problems but they are not necessarily of the form (1.4). For example, for Stokes problems

describing slow viscous flow of an incompressible fluid see Silvester and Wathen [26]. The preconditioning discussed in this paper does not assume any underlying structure except for the algebraic block structure indicated in (1.1).

In section 2, we determine the eigensolution distribution of the preconditioned system and give lower and upper bounds for the eigenvalues of $\mathcal{G}^{-1}\mathcal{A}$ in the case when the submatrix G is positive definite. Section 3 describes the convergence behavior of a Krylov subspace method such as GMRES, section 4 investigates possible implementation strategies, while in section 5 we give numerical results to support the theory developed in this paper.

2. Preconditioning \mathcal{A} . For symmetric (and in general normal) matrix systems, the convergence of an applicable iterative method is determined by the distribution of the eigenvalues of the coefficient matrix. In particular, it is desirable that the number of distinct eigenvalues, or at least the number of clusters, be small, because in this case convergence will be rapid. To be more precise, if there are only a few distinct eigenvalues, then optimal methods like CG, MINRES, or GMRES will terminate (in exact arithmetic) after a small and precisely defined number of steps. We prove a result of this type below. For nonnormal systems, convergence, as opposed to termination, is not so readily described; see Greenbaum [13, p. 5].

2.1. Eigenvalue distribution. The eigenvalues of the preconditioned coefficient matrix $\mathcal{G}^{-1}\mathcal{A}$ may be derived by considering the generalized eigenvalue problem

$$(2.1) \quad \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}.$$

Let $QR = [Y \ Z][R^T \ 0^T]^T$ be an orthogonal factorization of B^T , where $R \in \mathbb{R}^{m \times m}$ is upper triangular, $Y \in \mathbb{R}^{n \times m}$, and $Z \in \mathbb{R}^{n \times (n-m)}$ is a basis for the nullspace of B . Premultiplying (2.1) by the nonsingular and square matrix

$$\begin{bmatrix} Z^T & 0 \\ Y^T & 0 \\ 0 & I \end{bmatrix}$$

and postmultiplying by its transpose gives

$$(2.2) \quad \begin{bmatrix} Z^T AZ & Z^T AY & 0 \\ Y^T AZ & Y^T AY & R \\ 0 & R^T & 0 \end{bmatrix} \begin{bmatrix} x_z \\ x_y \\ y \end{bmatrix} = \lambda \begin{bmatrix} Z^T GZ & Z^T GY & 0 \\ Y^T GZ & Y^T GY & R \\ 0 & R^T & 0 \end{bmatrix} \begin{bmatrix} x_z \\ x_y \\ y \end{bmatrix},$$

with $x = Zx_z + Yx_y$ and where we made use of the equalities $BZ = 0$ and $R = (BY)^T$. Performing a simultaneous sequence of row and column interchanges on both matrices in (2.2) reveals two lower block-triangular matrices

$$\tilde{\mathcal{A}} = \begin{bmatrix} R^T & 0 & 0 \\ Z^T AY & Z^T AZ & 0 \\ Y^T AY & Y^T AZ & R \end{bmatrix}, \quad \tilde{\mathcal{G}} = \begin{bmatrix} R^T & 0 & 0 \\ Z^T GY & Z^T GZ & 0 \\ Y^T GY & Y^T GZ & R \end{bmatrix},$$

and thus the preconditioned coefficient matrix $\mathcal{G}^{-1}\mathcal{A}$ is similar to

$$(2.3) \quad \mathcal{P} = \tilde{\mathcal{G}}^{-1}\tilde{\mathcal{A}} = \begin{bmatrix} m & & & & & \\ & (n-m) & & & & \\ & & m & & & \\ & & & I & & \\ & & & & 0 & \\ & & & & & (Z^T GZ)^{-1}(Z^T AZ) & 0 \\ & & & & & & 0 \\ & & & & & & & I & & \\ & & & & & & & & m & \\ & & & & & & & & & (n-m) \\ & & & & & & & & & & m \end{bmatrix}.$$

Here, the precise forms of Θ , Υ , and Γ are irrelevant for the argument that follows; they are in general nonzero. We just proved the following theorem.

THEOREM 2.1. *Let $\mathcal{A} \in \mathbb{R}^{(n+m) \times (n+m)}$ be a symmetric and indefinite matrix of the form*

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and $B \in \mathbb{R}^{m \times n}$ is of full rank. Assume Z is an $n \times (n - m)$ basis for the nullspace of B . Preconditioning \mathcal{A} by a matrix of the form

$$\mathcal{G} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$ is symmetric, $G \neq A$, and $B \in \mathbb{R}^{m \times n}$ is as above, implies that the matrix $\mathcal{G}^{-1}\mathcal{A}$ has

- (1) an eigenvalue at 1 with multiplicity $2m$, and
- (2) $n - m$ eigenvalues which are defined by the generalized eigenvalue problem $Z^T A Z x_z = \lambda Z^T G Z x_z$.

Note that, if either $Z^T A Z$ is positive definite or we choose G so that $Z^T G Z$ is positive definite, then the indefinite constrained preconditioner applied to the indefinite linear system (1.1) yields the preconditioned matrix \mathcal{P} which has real eigenvalues.

Remark 2.2. In the above argument, we assumed that B has full row rank and consequently applied an orthogonal factorization of B^T , which resulted in an upper triangular matrix $R \in \mathbb{R}^{m \times m}$. If B does not have full row rank, i.e., $\text{rank}(B) = m - k$ for some integer $k \leq m$, then k zero rows and columns can be deleted from both matrices in (2.2), thus giving a reduced system of dimension $(n + m - k) \times (n + m - k)$. This removal of the redundant information does not impose any restriction on the proposed preconditioner, since all mathematical arguments apply equivalently to the reduced system of equations.

2.2. Eigenvector distribution. We mentioned above that the termination of a Krylov subspace method is related to the location of the eigenvalues and the number of corresponding linearly independent eigenvectors. In order to establish the association between eigenvectors and eigenvalues, we expand the general eigenvalue problem (2.2), yielding

$$(2.4) \quad Z^T A Z x_z + Z^T A Y x_y = \lambda [Z^T G Z x_z + Z^T G Y x_y],$$

$$(2.5) \quad Y^T A Z x_z + Y^T A Y x_y + R y = \lambda [Y^T G Z x_z + Y^T G Y x_y + R y],$$

$$(2.6) \quad R^T x_y = \lambda R^T x_y.$$

From (2.6), it may be deduced that either $\lambda = 1$ or $x_y = 0$. In the former case, (2.4) and (2.5) simplify to

$$\begin{aligned} Z^T A Z x_z + Z^T A Y x_y &= Z^T G Z x_z + Z^T G Y x_y, \\ Y^T A Z x_z + Y^T A Y x_y &= Y^T G Z x_z + Y^T G Y x_y, \end{aligned}$$

which can consequently be written as

$$(2.7) \quad Q^T A Q w = Q^T G Q w,$$

where $Q = \begin{bmatrix} Y & Z \end{bmatrix}$ and $w = \begin{bmatrix} x_y^T & x_z^T \end{bmatrix}^T$. Since Q is orthogonal, the general eigenvalue problem (2.7) is equivalent to considering

$$(2.8) \quad Aw = \sigma Gw,$$

where $w \neq 0$ if and only if $\sigma = 1$. There are m linearly independent eigenvectors $\begin{bmatrix} 0^T & 0^T & y^T \end{bmatrix}^T$ corresponding to $w = 0$, and a further i ($1 \leq i \leq n$) linearly independent eigenvectors (corresponding to eigenvalues $\sigma = 1$ of (2.8)).

Now, suppose $\lambda \neq 1$, in which case $x_y = 0$. Equations (2.4) and (2.5) yield

$$(2.9) \quad Z^T AZx_z = \lambda Z^T GZx_z,$$

$$(2.10) \quad Y^T AZx_z + Ry = \lambda \left[Y^T GZx_z + Ry \right].$$

The generalized eigenvalue problem (2.9) defines $n - m$ eigenvalues, where j ($1 \leq j \leq n - m$) of these are not equal to 1 and for which two cases have to be distinguished. If $x_z \neq 0$, y must satisfy

$$\left[Y^T AZ - \lambda Y^T GZ \right] x_z = (\lambda - 1) Ry,$$

from which follows that the corresponding eigenvectors are defined by $\begin{bmatrix} x_z^T & 0^T & y^T \end{bmatrix}^T$. If $x_z = 0$, we deduce from (2.10) that

$$Ry = \lambda Ry,$$

and hence that $y = 0$ since $\lambda \neq 1$. As $\begin{bmatrix} x_z^T & x_y^T & y^T \end{bmatrix}^T = 0$ in this case, no extra eigenvectors arise.

Summarizing the above, it is evident that \mathcal{P} has $m + i + j$ eigenvectors. We now show that, under realistic assumptions, these eigenvectors are in fact linearly independent.

THEOREM 2.3. *Let $\mathcal{A} \in \mathbb{R}^{(n+m) \times (n+m)}$ be a symmetric and indefinite matrix of the form*

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and $B \in \mathbb{R}^{m \times n}$ is of full rank. Assume the preconditioner \mathcal{G} is defined by a matrix of the form

$$\mathcal{G} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$ is symmetric, $G \neq A$, and $B \in \mathbb{R}^{m \times n}$ is as above. Let Z denote an $n \times (n - m)$ basis for the nullspace of B and suppose that $Z^T GZ$ is positive definite. The preconditioned matrix $\mathcal{G}^{-1} \mathcal{A}$ has $n + m$ eigenvalues as defined by Theorem 2.1 and $m + i + j$ linearly independent eigenvectors. There are

- (1) m eigenvectors of the form $\begin{bmatrix} 0^T & 0^T & y^T \end{bmatrix}^T$ that correspond to the case $\lambda = 1$;
- (2) i ($0 \leq i \leq n$) eigenvectors of the form $\begin{bmatrix} x_z^T & x_y^T & y^T \end{bmatrix}^T$ arising from $Aw = \sigma Gw$ with $w = \begin{bmatrix} x_y^T & x_z^T \end{bmatrix}^T$ linearly independent, $\sigma = 1$, and $\lambda = 1$; and

(3) j ($0 \leq j \leq n-m$) eigenvectors of the form $[x_z^T \ 0^T \ y^T]^T$ that correspond to the case $\lambda \neq 1$.

Proof. To prove that the $m+i+j$ eigenvectors of \mathcal{P} are linearly independent, we need to show that

$$(2.11) \quad \begin{bmatrix} 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} + \begin{bmatrix} x_{z_1}^{(2)} & \cdots & x_{z_i}^{(2)} \\ x_{y_1}^{(2)} & \cdots & x_{y_i}^{(2)} \\ y_1^{(2)} & \cdots & y_i^{(2)} \end{bmatrix} \begin{bmatrix} a_1^{(2)} \\ \vdots \\ a_i^{(2)} \end{bmatrix} + \begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$$

implies that the vectors $a^{(k)}$ ($k = 1, \dots, 3$) are zero vectors. Multiplying (2.11) by \mathcal{A} and \mathcal{G}^{-1} , and recalling that in the previous equation the first matrix arises from the case $\lambda_k = 1$ ($k = 1, \dots, m$), the second matrix from the case $\lambda_k = 1$ and $\sigma_k = 1$ ($k = 1, \dots, i$), and the last matrix from $\lambda_k \neq 1$ ($k = 1, \dots, j$), gives

$$(2.12) \quad \begin{bmatrix} 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} + \begin{bmatrix} x_{z_1}^{(2)} & \cdots & x_{z_i}^{(2)} \\ x_{y_1}^{(2)} & \cdots & x_{y_i}^{(2)} \\ y_1^{(2)} & \cdots & y_i^{(2)} \end{bmatrix} \begin{bmatrix} a_1^{(2)} \\ \vdots \\ a_i^{(2)} \end{bmatrix} + \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_j \end{bmatrix} \begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Subtracting (2.11) from (2.12) we obtain

$$\begin{bmatrix} \lambda_1 - 1 & & \\ & \ddots & \\ & & \lambda_j - 1 \end{bmatrix} \begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix},$$

which simplifies to

$$(2.13) \quad \begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix},$$

since $\lambda_k \neq 1$ ($k = 1, \dots, j$).

The assumption that $Z^T G Z$ is positive definite implies that $x_{z_k}^{(3)}$ ($k = 1, \dots, j$) in (2.13) are linearly independent and thus that $a_k^{(3)} = 0$ ($k = 1, \dots, j$). Similarly, $a_k^{(2)} = 0$ ($k = 1, \dots, i$) follows from the linear independence of $[x_{z_k}^{(2)T} \ x_{y_k}^{(2)T}]^T$ ($k = 1, \dots, i$), and thus (2.11) simplifies to

$$\begin{bmatrix} 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}.$$

However, $y_k^{(1)}$ ($k = 1, \dots, m$) are linearly independent and thus $a_k^{(1)} = 0$ ($k = 1, \dots, m$). \square

Remark 2.4. Note that the result of Theorem 2.3 remains true if $Z^T(\gamma A + \sigma G)Z$ is positive definite for some scalars γ and σ ; see Parlett [20, p. 343] for details.

To show that the eigenvector bounds of Theorem 2.3 can in fact be attained, consider the following two examples.

Example 2.5 (minimum bound). Consider the matrices

$$\mathcal{A} = \begin{bmatrix} 1 & 2 & 0 \\ 2 & 2 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad \mathcal{G} = \begin{bmatrix} 1 & 3 & 0 \\ 3 & 4 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

so that $m = 1$ and $n = 2$. The preconditioned matrix \mathcal{P} has an eigenvalue at 1 with multiplicity 3, but only one eigenvector arising from case (1) in Theorem 2.3. This eigenvector may be taken to be $[0 \ 0 \ 1]^T$.

Example 2.6 (maximum bound). Let $\mathcal{A} \in \mathbb{R}^{3 \times 3}$ be defined as in Example 2.5, but assume $\mathcal{G} = \mathcal{A}$. The preconditioned matrix \mathcal{P} has an eigenvalue at 1 with multiplicity 3 and clearly a complete set of eigenvectors. These may be taken to be the columns of the identity matrix.

2.3. Eigenvalue bounds. It is apparent from the calculations in the previous section that the eigenvalue at 1 with multiplicity $2m$ is independent of the choice of G in the preconditioner. On the contrary, the $n - m$ eigenvalues that are defined by (2.9) are highly sensitive to the choice of G . If G is a close approximation of A , we can expect a more favorable distribution of eigenvalues and consequently may expect faster convergence of an appropriate iterative method. In order to determine a good factorization of A , it will be helpful to find intervals in which the $n - m$ eigenvalues are located. If G is a positive definite matrix, one possible approach is provided by Cauchy’s interlace theorem.

THEOREM 2.7 (Cauchy’s interlace theorem). *Suppose $T \in \mathbb{R}^{n \times n}$ is symmetric and*

$$T = \begin{bmatrix} H & \star \\ \star & \star \end{bmatrix},$$

where $H \in \mathbb{R}^{m \times m}$ with $m < n$. Label the eigenpairs of T and H as

$$\begin{aligned} Tz_i &= \alpha_i z_i, & i &= 1, \dots, n, & \alpha_1 &\leq \alpha_2 \leq \dots \leq \alpha_n, \\ Hy_i &= \lambda_i y_i, & i &= 1, \dots, m, & \lambda_1 &\leq \lambda_2 \leq \dots \leq \lambda_m. \end{aligned}$$

Then

$$\alpha_k \leq \lambda_k \leq \alpha_{k+(n-m)}, \quad k = 1, \dots, m.$$

Proof. See Parlett [20, p. 203]. \square

The applicability of Theorem 2.7 is verified by recalling the definitions of Q and Z given in the previous section and by considering the generalized eigenvalue problems

$$(2.14) \quad Q^T A Q v = \alpha Q^T G Q v$$

and

$$(2.15) \quad Z^T A Z w = \lambda Z^T G Z w.$$

Since G is positive definite, so is $Q^T G Q$, and we may therefore write

$$Q^T G Q = \begin{bmatrix} Z^T G Z & Z^T G Y \\ Y^T G Z & Y^T G Y \end{bmatrix} = \underbrace{\begin{bmatrix} L & 0 \\ R & S \end{bmatrix}}_M \underbrace{\begin{bmatrix} L^T & R^T \\ 0 & S^T \end{bmatrix}}_{M^T},$$

where $LL^T = Z^T G Z$, $R = Y^T G Z L^{-T}$, and $SS^T = Y^T G Y - RR^T$. Rewriting (2.14) and (2.15) gives

$$(2.16) \quad M^{-1} Q^T A Q M^{-T} u = \alpha u$$

and

$$(2.17) \quad L^{-1} Z^T A Z L^{-T} z = \lambda z,$$

where $u = M^T v$ and $z = L^T w$.

Now, since the matrix $M^{-1} Q^T A Q M^{-T}$ is similar to $G^{-1} A$, (2.16) defines the same eigenvalues α_i ($i = 1, \dots, n$) as $G^{-1} A$. We may therefore apply Theorem 2.7 directly. The result is that the $n - m$ eigenvalues λ_i of (2.9) satisfy $\alpha_k \leq \lambda_k \leq \alpha_{k+m}$ ($k = 1, \dots, n - m$). In particular, the λ_i are bounded by the extreme eigenvalues of $G^{-1} A$ so that the λ_i will necessarily be clustered if G is a good approximation of A . Furthermore, a good preconditioner G for A implies that $Z^T G Z$ is at least as good a preconditioner for $Z^T A Z$. To show that the preconditioner $Z^T G Z$ can in fact be much better, consider the following example, taken from the CUTE collection [5].

Example 2.8. Consider the convex quadratic programming problem *BLOWEYC*, which may be formulated as

$$\begin{aligned} &\text{minimize } u(s)^T A u(s) + u(s)^T w(s) - v(s)^T A u(s) - 2.0 v(s)^T w(s) - u(s)^T v(s) \\ &\text{subject to } A w(s) = u(s), \quad u(s) \in [-1, 1], \quad \text{and} \quad \int_0^1 u(s) ds = 0.4. \end{aligned}$$

Selecting a size parameter of 500 discretization intervals defines a set of linear equations of the form (1.1), where $n = 1002$ and $m = 502$. Letting G be the diagonal of A , we may deduce from the above theory that the extreme eigenvalues of $G^{-1} A$ give lower and upper bounds for the $n - m$ eigenvalues defined by the general eigenvalue problem (2.9). In Figure 2.1(a), the 1002 eigenvalues of $G^{-1} A$ are drawn as vertical lines, whereas Figure 2.1(b) displays the 500 eigenvalues of $(Z^T G Z)^{-1} Z^T A Z$.

The spectrum of Figure 2.1(a) is equivalent to a graph of the entire spectrum of \mathcal{P} , but with an eigenvalue at 1 of multiplicity 502 removed. Rounded to two decimal places, the numerical values of the two extreme eigenvalues of $G^{-1} A$ are 0.02 and 1.98, whereas the extreme eigenvalues of $(Z^T G Z)^{-1} Z^T A Z$ are given by 0.71 and 1. Note that for this example a large number of eigenvalues of $G^{-1} A$ are clustered in the approximate intervals $[0.02, 0.38]$ and $[1.65, 1.97]$. The eigenvalue distribution in Figure 2.1(b) reveals that there is one eigenvalue near 0.71 and a group of eigenvalues near 1. It follows that any appropriate iterative method which solves (1.5) can be expected to converge in a very small number of steps; this is verified by the numerical results presented in section 5.

It is readily seen from Example 2.8 that in this case the bounds provided by Theorem 2.7 are not descriptive in that there is significantly more clustering of the eigenvalues than implied by the theorem.




FIG. 2.1. Continuous vertical lines represent the eigenvalues of (a) $G^{-1}A$ and (b) $(Z^T G Z)^{-1} Z^T A Z$.

3. Convergence. In the context of this paper, the convergence of an iterative method under preconditioning is influenced not only by the spectral properties of the coefficient matrix, but also by the relationship between the dimensions n and m . In particular, it follows from Theorem 2.1 that in the special case when $n = m$ the preconditioned linear system (1.5) has only one eigenvalue at 1 with multiplicity $2n$. For $m < n$, matrix (2.3) gives an eigenvalue at 1 with multiplicity $2m$ and $n - m$ (generally distinct) eigenvalues whose value may or may not be equal to 1. Before we examine how these results determine upper bounds on the number of iterations of an appropriate Krylov subspace method, we recall the definition of the *minimum polynomial* of a matrix.

DEFINITION 3.1. Let $\mathcal{A} \in \mathbb{R}^{(n+m) \times (n+m)}$. The monic polynomial f of minimum degree such that $f(\mathcal{A}) = 0$ is called the *minimum polynomial* of \mathcal{A} .

The importance of this definition becomes apparent when considering subsequent results and by recalling that similar matrices have the same minimum polynomial.

Krylov subspace theory states that iteration with any method with an optimality property such as GMRES will terminate when the degree of the minimum polynomial is attained; see Saad and Schultz [25, Proposition 2] and also Campbell et al. [6] in the case of GMRES. (To be precise, the number may be smaller in special cases where b is a combination of a few eigenvectors that affect the “grade” of \mathcal{A} with respect to b .) In particular, the degree of the minimum polynomial is equal to the dimension of the corresponding Krylov subspace (for general b) (see Saad [24, Proposition 6.1]), and so the following theorems are relevant.

THEOREM 3.2. Let $\mathcal{A} \in \mathbb{R}^{(n+m) \times (n+m)}$ be a symmetric and indefinite matrix of the form

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and $B \in \mathbb{R}^{m \times n}$ is of full rank. Let $m = n$. If \mathcal{A} is preconditioned by a matrix of the form

$$\mathcal{G} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$, $G \neq A$, and $B \in \mathbb{R}^{m \times n}$ is as above, then the Krylov subspace $\mathcal{K}(\mathcal{P}, b)$ is of dimension at most 2 for any b .

Proof. Writing the preconditioned system (2.3) in its explicit form, we observe that \mathcal{P} is in fact given by

$$(3.1) \quad \begin{bmatrix} I & 0 \\ \Upsilon & I \end{bmatrix},$$

where Υ is nonzero if and only if $A \neq G$. To show that the dimension of the corresponding Krylov subspace is at most 2, we need to determine the minimum polynomial of the system. It is evident from (3.1) that the eigenvalues of \mathcal{P} are all 1 and $\mathcal{P} - I \neq 0$. However, $(\mathcal{P} - I)^2 = 0$ and so the minimum polynomial is of order 2. \square

Remark 3.3. It is of course possible in the case $n = m$ to solve the (square) constrained equation $Bx_1 = b_2$ and then to obtain $x_2 = B^{-T}(b_1 - Ax_1)$. This gives motivation for why the result of Theorem 3.2 is independent of G .

Remark 3.4. The important consequence of Theorem 3.2 is that termination of an iteration method such as GMRES will occur in at most two steps for any choice of b , even though the preconditioned matrix is not diagonalizable (unless $A = G$).

THEOREM 3.5. *Let $\mathcal{A} \in \mathbb{R}^{(n+m) \times (n+m)}$ be a symmetric and indefinite matrix of the form*

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and $B \in \mathbb{R}^{m \times n}$ is of full rank. Assume $m < n$ and \mathcal{A} is nonsingular. Furthermore, assume \mathcal{A} is preconditioned by a matrix of the form

$$\mathcal{G} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$ is symmetric, $G \neq A$, and $B \in \mathbb{R}^{m \times n}$ is as above. If $Z^T G Z$ is positive definite, where Z is an $n \times (n - m)$ basis for the nullspace of B , then the dimension of the Krylov subspace $\mathcal{K}(\mathcal{P}, b)$ is at most $n - m + 2$.

Proof. From the eigenvalue derivation in section 2.1, it is evident that the characteristic polynomial of the preconditioned linear system (1.5) is

$$(\mathcal{P} - I)^{2m} \prod_{i=1}^{n-m} (\mathcal{P} - \lambda_i I).$$

To prove the upper bound on the dimension of the Krylov subspace, we need to show that the order of the minimum polynomial is less than or equal to $n - m + 2$.

Expanding the polynomial $(\mathcal{P} - I) \prod_{i=1}^{n-m} (\mathcal{P} - \lambda_i I)$ of degree $n - m + 1$, we obtain a matrix of the form

$$(3.2) \quad \begin{bmatrix} 0 & 0 & 0 \\ \left[\prod_{i=1}^{n-m} (S - \lambda_i I) \right] \Theta & (S - I) \prod_{i=1}^{n-m} (S - \lambda_i I) & 0 \\ \Phi_{n-m} & \Gamma \prod_{i=1}^{n-m} (S - \lambda_i I) & 0 \end{bmatrix},$$

where $S = (Z^T G Z)^{-1} Z^T A Z$ and

$$\Phi_{n-m} = \Upsilon \prod_{i=1}^{n-m} (I - \lambda_i I) + \Gamma \left[\prod_{i=1}^{n-m-1} (S - \lambda_i I) \right] \Theta + \Gamma \Psi_{n-m}.$$

Here, Ψ_{n-m} is defined by the recursive formula

$$\Psi_{n-m} = \left[\Psi_{n-m-1} + \left[\prod_{i=1}^{n-m-2} (S - \lambda_i I) \right] \Theta \right] (I - \lambda_{n-m} I) \quad (n - m > 2),$$

with base cases $\Psi_1 = 0$ and $\Psi_2 = \Theta(I - \lambda_2 I)$.

Note that the (2, 1), (2, 2), and (3, 2) entries of matrix (3.2) are in fact zero, since the λ_i ($i = 1, \dots, n - m$) are the eigenvalues of S , which is similar to a symmetric matrix and is thus diagonalizable. Thus, (3.2) may be written as

$$(3.3) \quad \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \Phi_{n-m} & 0 & 0 \end{bmatrix},$$

and what remains is to distinguish two different cases for the value of Φ_{n-m} , that is, $\Phi_{n-m} = 0$ and $\Phi_{n-m} \neq 0$. In the former case, the order of the minimum polynomial of \mathcal{P} is less than or equal to $n - m + 1$, and thus the dimension of the Krylov subspace $\mathcal{K}(\mathcal{P}, b)$ is of the same order. In the latter case, the dimension of $\mathcal{K}(\mathcal{P}, b)$ is less than or equal to $n - m + 2$ since multiplication of (3.3) by another factor $(\mathcal{P} - I)$ gives the zero matrix. \square

The upper bound on the dimension of the Krylov subspace, as stated in Theorem 3.5, can be reduced in the special case when $(Z^T G Z)^{-1} (Z^T A Z)$ has repeated eigenvalues. This result is stated in Theorem 3.7. The following (randomly generated) example shows that the bound in Theorem 3.5 is attainable.

Example 3.6. Let $A \in \mathbb{R}^{6 \times 6}$ and $B^T \in \mathbb{R}^{6 \times 2}$ be given by

$$A = \begin{bmatrix} 2.69 & 1.62 & 1.16 & 1.60 & 0.81 & -1.97 \\ 1.62 & 6.23 & -1.90 & 1.89 & 0.90 & 0.05 \\ 1.16 & -1.90 & 4.01 & -0.16 & -0.16 & -1.60 \\ 1.60 & 1.89 & -0.16 & 1.45 & 0.01 & -0.89 \\ 0.81 & 0.90 & -0.16 & 0.01 & 1.94 & 0.38 \\ -1.97 & 0.05 & -1.60 & -0.89 & 0.38 & 5.38 \end{bmatrix}, \quad B^T = \begin{bmatrix} 0 & -0.59 \\ -0.59 & 0 \\ 0 & 2.00 \\ 0 & 0 \\ -0.02 & 0 \\ 0.33 & 0.17 \end{bmatrix},$$

and assume that $G = \text{diag}(A)$. For the above matrices the (3, 1) entry of (3.3) is

$$\Phi_4 = \begin{bmatrix} 0 & -0.07 \\ -0.22 & -0.02 \end{bmatrix}.$$

It follows that the minimum polynomial is of order 6, and thus the bound given in Theorem 3.5 is sharp.

THEOREM 3.7. *Let $\mathcal{A} \in \mathbb{R}^{(n+m) \times (n+m)}$ be a symmetric and indefinite matrix of the form*

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and $B \in \mathbb{R}^{m \times n}$ is of full rank. Assume $m < n$, A is nonsingular, and \mathcal{A} is preconditioned by a matrix of the form

$$\mathcal{G} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$ is symmetric, $G \neq A$, and $B \in \mathbb{R}^{m \times n}$ is as above. Furthermore, let Z be an $n \times (n - m)$ basis for the nullspace of B and assume $(Z^T G Z)^{-1}(Z^T A Z)$ has k ($1 \leq k \leq n - m$) distinct eigenvalues λ_i ($1 \leq i \leq k$) of respective multiplicity μ_i , where $\sum_{i=1}^k \mu_i = n - m$. Then the dimension of the Krylov subspace $\mathcal{K}(\mathcal{P}, b)$ is at most $k + 2$.

Proof. The proof is similar to the one for Theorem 3.5. In the case when $(Z^T G Z)^{-1}(Z^T A Z)$ has k distinct eigenvalues of multiplicity μ_i , we may, without loss of generality, write the characteristic polynomial of \mathcal{P} as

$$(\mathcal{P} - I)^{2m-2} \left[\prod_{i=1}^k (\mathcal{P} - \lambda_i I)^{\mu_i-1} \right] (\mathcal{P} - I) \underbrace{(\mathcal{P} - I) \left[\prod_{i=1}^k (\mathcal{P} - \lambda_i I) \right]}_{(\dagger)}.$$

Expanding (\dagger) , we obtain the matrix

$$(3.4) \quad \begin{bmatrix} 0 & 0 & 0 \\ \left[\prod_{i=1}^k (S - \lambda_i I) \right] \Theta & (S - I) \prod_{i=1}^k (S - \lambda_i I) & 0 \\ \Phi_k & \Gamma \prod_{i=1}^k (S - \lambda_i I) & 0 \end{bmatrix},$$

where $S = (Z^T G Z)^{-1} Z^T A Z$ and

$$\Phi_k = \Upsilon \prod_{i=1}^k (I - \lambda_i I) + \Gamma \left[\prod_{i=1}^{k-1} (S - \lambda_i I) \right] \Theta + \Gamma \Psi_k.$$

Here, Ψ_k is given by the recursive formula

$$\Psi_k = \left[\Psi_{k-1} + \left[\prod_{i=1}^{k-2} (S - \lambda_i I) \right] \Theta \right] (I - \lambda_k I) \quad (k > 2),$$

with base cases $\Psi_1 = 0$ and $\Psi_2 = \Theta(I - \lambda_2 I)$.

Note that the $(2, 1)$, $(2, 2)$, and $(3, 2)$ blocks of matrix (3.4) are in fact zero. It follows that, for $\Phi_k \neq 0$, a further multiplication of (3.4) by $(\mathcal{P} - I)$ gives the zero matrix and thus the dimension of Krylov subspace $\mathcal{K}(\mathcal{P}, b)$ is less than or equal to $k + 2$. \square

To verify that the bound in Theorem 3.7 is attainable, consider the following example.

Example 3.8. Let $A \in \mathbb{R}^{4 \times 4}$, $G \in \mathbb{R}^{4 \times 4}$, and $B^T \in \mathbb{R}^{4 \times 1}$ be given by

$$A = \begin{bmatrix} 6 & 0 & 0 & 0 \\ 0 & 6 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{bmatrix}, \quad G = \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & \frac{1}{2} \end{bmatrix},$$

and $B = [0 \ 0 \ 10^{-3} \ 10^{-3}]$. Then two of the $n - m$ eigenvalues that are defined by the generalized eigenvalue problem (2.9) are distinct and given by $[2, 4]$. It follows that the $(3, 1)$ entry of (3.4) is nonzero with

$$\Phi_3 = [-1071.41],$$

and so the minimum polynomial is of order 4.

4. Implementation. There are various strategies that can be used to implement the proposed preconditioner, two of which are used in the numerical results in section 5. The first strategy applies the standard (preconditioned) GMRES algorithm [25], where the preconditioner step is implemented by means of a symmetric indefinite factorization of (1.4). Such a factorization of the preconditioner may be much less demanding than the factorization of the initial coefficient matrix if G is a considerably simpler matrix than A . The second approach, discussed in the next section, is based on an algorithm that solves a reduced linear system [12].

4.1. CG on a reduced linear system. In [12], Gould, Hribar, and Nocedal propose a CG-like algorithm to solve equality constrained quadratic programming problems such as the one described in Example 1.1. The algorithm is based on the idea of computing an implicit basis Z which spans the nullspace of B . The nullspace basis is then used to remove the constraints from the system of equations, thus allowing the application of the CG method to the (positive definite) reduced system.

Assume that $W_{zz} = Z^T G Z$ is a symmetric and positive definite preconditioner matrix of dimension $(n-m) \times (n-m)$ and Z is an $n \times (n-m)$ matrix. The algorithm can then be stated as follows.

ALGORITHM 4.1. PRECONDITIONED CG FOR A REDUCED SYSTEM.

- (1) Choose an initial point x_1 satisfying $Bx_1 = b_2$.
- (2) Compute

$$(4.1) \quad \begin{aligned} r &= Ax_1 - b_1, \\ r &\leftarrow r - B^T y, \end{aligned}$$

$$(4.2) \quad \begin{aligned} g &= ZW_{zz}^{-1}Z^T r, \\ p &= -g. \end{aligned}$$

- (3) Repeat the following steps until $|(r^+)^T g^+| \rightarrow 0$ is satisfied:

$$(4.3) \quad \begin{aligned} \alpha &= r^T g / p^T A p, \\ x_1 &\leftarrow x_1 + \alpha p, \\ r^+ &= r + \alpha A p, \\ r^+ &\leftarrow r^+ - B^T y, \end{aligned}$$

$$(4.4) \quad \begin{aligned} g^+ &= ZW_{zz}^{-1}Z^T r^+, \\ \beta &= (r^+)^T g^+ / r^T g, \\ p &\leftarrow -g^+ + \beta p, \\ g &\leftarrow g^+, \\ r &\leftarrow r^+. \end{aligned}$$

The vectors y in (4.1) and (4.3) are arbitrary, but in practice they are chosen to try to make r and r^+ close to $B^T y$. The computation of the preconditioned residual in (4.4) is often the most expensive computational factor in the algorithm. Gould, Hribar, and Nocedal [12] suggest avoiding the explicit use of the nullspace Z , but instead computing g^+ by applying a symmetric indefinite factorization of

$$(4.5) \quad \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} g^+ \\ v^+ \end{bmatrix} = \begin{bmatrix} r^+ \\ 0 \end{bmatrix}.$$

This factorization also determines the vector y in (4.3), where $y = v^+$. The vector y in (4.1) and the preconditioned residual (4.2) are determined in a similar manner.

In practice, (4.5) can often be factored efficiently by using the MA27 package of the Harwell Subroutine Library when G is a simple matrix block, whereas the direct application of MA27 to the original system (1.1) is limited by space requirements as well as time for large enough systems [8]. In this context, the factorization consists of three separate routines, the first two of which analyze and factorize the matrix in (4.5). They need to be executed only once in step (1) of Algorithm 4.1. Repeated calls to the third routine within MA27 apply forward- and backward-substitutions to find the initial point x_1 in step (1), solve for y in (4.1) and g in (4.2), and also find y in (4.3) and g^+ in (4.4).

Remark 4.2. The computation of the projected residual g^+ is often accompanied by significant roundoff errors if this vector is much smaller than the residual r^+ . *Fixed precision iterative refinement* is used in (4.3) to redefine r^+ so that its norm is closer to that of g^+ . In the same manner, (4.1) is used to refine the initial residual r . The result is a dramatic reduction of the roundoff errors in the projection operation; see Gould, Hribar, and Nocedal [12].

Remark 4.3. The current implementations, which apply the proposed preconditioner, require an explicit factorization of the indefinite matrix \mathcal{G} . For larger systems, the complexity of the two methods introduced in this section can grow substantially for increasing problem dimensions and may not be feasible in comparison to the complexity of methods such as (block-) diagonal and triangular preconditioners. The authors are currently investigating a new implementation approach which is mainly based on an adaptation of the direct solver MA27. Preliminary numerical results show the beneficial features of the proposed approach.

5. Numerical results. We now present the results of numerical experiments that reinforce the analysis given in previous sections. The test problems employed are partly matrices that arise in linear and nonlinear optimization (see, for instance, Example 1.1) and partly randomized sparse matrices which arise in areas such as those described by Example 1.2. In particular, the problems in Table 5.1 are of the form (1.3), with A representing a *stiffness matrix*, whereas the examples in Table 5.2 are of the same form, but with A being a *mass matrix* or a *discrete Laplacian*. The four test examples of Table 5.3 are taken from the *constrained and unconstrained testing environment (CUTE)* [5], where the individual problems have the following characteristics.

BLOWEYC is the convex quadratic programming problem

$$\begin{aligned} &\text{minimize } u(s)^T A u(s) + u(s)^T w(s) - v(s)^T A u(s) - 2.0v(s)^T w(s) - u(s)^T v(s) \\ &\text{subject to } A w(s) = u(s), \quad u(s) \in [-1, 1], \quad \text{and} \quad \int_0^1 u(s) ds = 0.4, \end{aligned}$$

which is a simplification of a problem given in [4]. It is part of a modeling exercise, and the actual value of the solution is not used in a genuine practical application.

CVXQP1 is a convex quadratic program whose constraints are linear and which has continuous first and second order derivatives. It is a purely academic problem and has been constructed specifically for test purposes.

MOSARQP2 is also a convex quadratic problem which was first proposed by Morales-Pérez and Sargent [17]. Once again, it is a purely academic problem which has been constructed specifically for test purposes.

FIT2P is a dual version of a problem that originates from a model for fitting linear inequalities to data by minimization of a sum of piecewise linear penalties.

TABLE 5.1
Random test problems with A being a stiffness matrix ($G = \text{diag}(A)$).

	Stiff I	Stiff II	Stiff III	Stiff IV	Stiff V	Stiff VI
n	578	2178	8450	8450	33282	33282
m	32	128	236	46	128	512
Nonzero entries in A	2316	9740	39948	39948	161804	161804
Nonzero entries in B	427	1871	3600	685	7726	15830
MINRES						
# of iterations	174	387	639	515	1579	1571
Time in seconds	0.4	3.0	18.1	12.7	170.4	175.5
PGMRES						
# of iterations	46	87	228	242	313	319
Time in seconds	0.2	3.5	92.4	104.2	695.7	730.6
RCG						
# of iterations	35	72	197	214	294	295
Time in seconds	0.1	1.0	7.7	7.6	44.8	48.7
MA27						
Time in seconds	0.1	0.3	1.7	2.1	7.2	11.8

As indicated throughout, all matrices are of the form

$$(5.1) \quad \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric, $B \in \mathbb{R}^{m \times n}$ has full rank, and $m \leq n$. Four different approaches to finding solutions to (1.1) are compared: three iterative algorithms based on Krylov subspaces, and the direct solver MA27, which applies a sparse variant of Gaussian elimination; see Duff and Reid [8]. To investigate possible favorable aspects of preconditioning, it makes sense to compare unpreconditioned with preconditioned solution strategies. The indefinite nature of matrix (5.1) suggests the use of MINRES in the unpreconditioned case. Of course, positive definite preconditioning could be employed with MINRES; see, for example, Murphy, Golub, and Wathen [18], or Silvester and Wathen [26], and Wathen and Silvester [27] in case of the Stokes problem. However, this is not done here. We refer to Battermann and Heinkenschloss [3] for some numerical results of this type.

As outlined in section 4, we employ two slightly different strategies in order to implement the preconditioner \mathcal{G} . The first method applies a standard (full) GMRES(\mathcal{A}) code (PGMRES in Tables 5.1, 5.2, and 5.3), which is mathematically equivalent to MINRES(\mathcal{A}) for symmetric matrices \mathcal{A} , whereas the second approach implements Algorithm 4.1 (RCG in Tables 5.1, 5.2 and 5.3). The choice $G = \text{diag}(A)$ in the preconditioner is made for both PGMRES and RCG.

All tests were performed on a SUN Ultra SPARCII-300MHz (ULTRA-30) workstation with 245 MB physical RAM and running SunOS release 5.5.1. Programs were written in standard Fortran 77 using the SUN WorkShop *f77* compiler (version 4.2) with the *-O* optimization flag set. In order to deal with large sparse matrices, we implemented an index storage format that only stores nonzero matrix elements; see Press, Teukolsky, and Vetterling [22]. The termination criterion for all iterative methods was taken to be a residual vector of order less than 10^{-6} in the 2-norm.

As part of its analysis procedure, MA27 accepts the pattern of some coefficient matrix and chooses pivots for the factorization and solution phases of subsequent routines. The amount of pivoting is controlled by the special threshold parameter u ($-1/2 \leq u \leq 1/2$). Modifying u within its positive range influences the accuracy of the resulting solution, whereas a negative value prevents any pivoting; see Duff and Reid [8]. In this context, the early construction of some of the test examples with the default value $u = 0.1$ was accompanied by difficulties in the form of memory limitations. We met the trade-off between less use of memory and solutions of high

TABLE 5.2
Random test problems with A being a mass matrix ($G = \text{diag}(A)$).

	Mass I	Mass II	Mass III	Mass IV	Mass V	Mass VI
n	578	2178	8450	8450	33282	33282
m	32	128	236	46	128	512
Nonzero entries in A	3128	13400	55448	55448	225560	225560
Nonzero entries in B	427	1871	3600	686	7726	15830
MINRES						
# of iterations	135	312	162	67	473	845
Time in seconds	0.3	2.9	6.8	3.5	65.1	117.9
PGMRES						
# of iterations	11	11	11	11	12	12
Time in seconds	0.1	0.7	2.6	2.4	21.0	31.3
RCG						
# of iterations	9	10	11	11	12	12
Time in seconds	0.1	0.7	2.6	2.3	20.3	26.6
MA27						
Time in seconds	0.1	0.4	6.5	4.5	–	24.1

TABLE 5.3
CUTE test problems ($G = \text{diag}(A)$).

	BLOWEYC	CVXQP1	MOSARQP2	FIT2P
n	1002	100	930	13525
m	502	100	30	3000
Nonzero entries in A	3004	672	1020	13525
Nonzero entries in B	2503	295	148	50284
MINRES				
# of iterations	363	> 200	51	180
Time in seconds	1.8	–	0.1	14.6
PGMRES				
# of iterations	2	2	6	1
Time in seconds	0.4	0.1	0.2	13.2
RCG				
# of iterations	2	2	4	1
Time in seconds	0.3	0.1	0.1	12.1
MA27				
Time in seconds	0.2	0.1	0.1	11.5

enough accuracy by choosing the parameter value $u = 0.01$ for the results in Tables 5.1, 5.2, and 5.3. However, this resulted in a floating point exception when using MA27 with one of the larger examples in Table 5.2, which is indicated by a dash.

The numerical results suggest that the inclusion of the (1, 2) and (2, 1) block of A into the preconditioner, together with $G = \text{diag}(A)$, results in a considerable reduction of iterations, where the appropriate bounds of Theorems 3.2, 3.5, and 3.7 are attained in all cases. Specifically, Theorem 3.2 applies in the context of problem CVXQP1.

Test problems Stiff III to Stiff VI in Table 5.1 emphasize the storage problems that are associated with the use of long recurrences in the PGMRES algorithm. The time required to find solutions to Stiff III to Stiff VI via the PGMRES algorithm is not comparable to any of the other methods, which is due to the increased storage requirements and the data trafficking involved. A solution to the memory problems is to restart PGMRES after a prescribed number of iterations, but the iteration counts for such restarts would not be comparable with those of full PGMRES.

The examples in Table 5.2 indicate the beneficial strategy of preconditioning the mass matrix A with its diagonal ($G = \text{diag}(A)$); see Wathen [28]. As expected from the theory described in [28], PGMRES, as well as RCG, converge in virtually all cases in only a few number of steps and also outperform the solutions found by the direct solver MA27 as far as the time measurements are concerned.

These measurements for MA27 suggest that the preconditioned CG algorithm, discussed in section 4.1, is a suitable alternative to the direct solver. Whereas both

MINRES and especially PGMRES are in most cases considerably slower than MA27, the timings for RCG are in virtually all cases comparable. For problems of large enough dimension or bandwidth, the resources required by MA27 must become prohibitive, in which case RCG becomes even more competitive.

6. Conclusion. In this paper, we investigated a new class of preconditioner for indefinite linear systems that incorporate the $(1, 2)$ and $(2, 1)$ blocks of the original matrix. These blocks are often associated with constraints. In our numerical results, we used a simple diagonal matrix G to approximate the $(1, 1)$ block of \mathcal{A} , even though other approximations, such as an incomplete factorization of A , are possible. We first showed that the inclusion of the constraints into the preconditioner clusters at least $2m$ eigenvalues at 1, regardless of the structure of G . However, unless G represents A exactly, \mathcal{P} does not have a complete set of linearly independent eigenvectors and thus the standard convergence theory for Krylov subspace methods is not readily applicable.

To find an upper bound on the number of iterations required to solve linear systems of the form (1.1) by means of appropriate subspace methods, we used a minimum polynomial argument. Theorem 3.2 considers the special condition $n = m$, in which case termination is guaranteed in two iterations. For $m < n$, Theorem 3.5 gives a general (sharp) upper bound on the dimension of the Krylov subspace, whereas Theorem 3.7 defines a considerably stronger result if some of the $n - m$ eigenvalues, defined by $(Z^T G Z)^{-1}(Z^T A Z)$, are repeated.

In the special case when G is a positive definite matrix block, we were able to apply Cauchy's interlacing theorem in order to give upper and lower bounds for the $n - m$ eigenvalues that are defined by the $(2, 2)$ block of matrix (2.3).

To confirm the analytical results in this paper, we used three different subspace methods, MINRES of Paige and Saunders [19] for the unpreconditioned matrix system and RCG of Gould, Hribar, and Nocedal [12] and PGMRES of Saad and Schultz [25] for the preconditioned case. Overall, the results show that the number of iterations is decreased substantially if preconditioning is applied. The Krylov subspaces, which are built during the execution of the two preconditioned implementations, are in theory of equal dimension for any one of the 16 test examples, and thus PGMRES and RCG can be expected to terminate in the same number of steps. However, convergence to any prescribed tolerance may occur for a different number of steps since PGMRES and RCG minimize different quantities. This can be seen in some of the examples. Nevertheless, we note that convergence for both methods is attained much earlier than suggested by the bounds in Theorems 3.2, 3.5, and 3.7.

We have not compared our results with positive definite preconditioning in connection with MINRES—this is known to be a very competitive approach in various applications.

Acknowledgments. The authors would like to thank Gene H. Golub for his insightful comments during the process of the work and an anonymous referee for helpful suggestions.

REFERENCES

- [1] O. AXELSSON, *Iterative Solution Methods*, Cambridge University Press, Cambridge, UK, 1996.
- [2] R. BARRETT, M. W. BERRY, T. F. CHAN, J. DEMMEL, J. DONATO, J. DONGARRA, V. EIJKHOUT, R. POZO, C. ROMINE, AND H. VAN DER VORST, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, SIAM, Philadelphia, 1993.

- [3] A. BATTERMANN AND M. HEINKENSCHLOSS, *Preconditioners for Karush-Kuhn-Tucker matrices arising in the optimal control of distributed systems*, in *Optimal Control of Partial Differential Equations*, Internat. Ser. Numer. Math. 126, W. Desch, F. Kappel, and K. Kunisch, eds., Birkhäuser, Berlin, 1998, pp. 15–32.
- [4] J. F. BLOWEY AND C. M. ELLIOT, *The Cahn-Hilliard gradient theory for phase separation with non-smooth free energy. Part II: Numerical analysis*, *European J. Appl. Math.*, 3 (1992), pp. 147–179.
- [5] I. BONGARTZ, A. CONN, N. GOULD, AND PH. TOINT, *CUTE: Constrained and unconstrained testing environment*, *ACM Trans. Math. Software*, 21 (1995), pp. 123–160.
- [6] S. L. CAMPBELL, I. C. F. IPSEN, C. T. KELLEY, AND C. D. MEYER, *GMRES and the minimum polynomial*, *BIT*, 36 (1996), pp. 664–675.
- [7] T. F. COLEMAN, *Linearly constrained optimization and projected preconditioned conjugate gradients*, in *Proceedings of the Fifth SIAM Conference on Applied Linear Algebra*, *Proc. Appl. Math.* 72, J. Lewis, ed., SIAM, Philadelphia, 1994, pp. 188–122.
- [8] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear equations*, *ACM Trans. Math. Software*, 9 (1983), pp. 302–325.
- [9] H. C. ELMAN, *Perturbation of eigenvalues of preconditioned Navier-Stokes operators*, *SIAM J. Matrix Anal. Appl.*, 18 (1997), pp. 733–751.
- [10] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD, 1996.
- [11] G. H. GOLUB AND A. J. WATHEN, *An iteration for indefinite systems and its application to the Navier-Stokes equations*, *SIAM J. Sci. Comput.*, 19 (1998), pp. 530–539.
- [12] N. I. M. GOULD, M. E. HRIBAR, AND J. NOCEDAL, *On the Solution of Equality Constrained Quadratic Programming Problems Arising in Optimization*, Technical Report RAL-TR-1998-069, Rutherford Appleton Laboratory, Oxfordshire, UK, 1998.
- [13] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, *Frontiers in Appl. Math.* 17, SIAM, Philadelphia, 1997.
- [14] A. GREENBAUM, AND L. GURVITS, *Max-min properties of matrix factor norms*, *SIAM J. Sci. Comput.*, 15 (1994), pp. 348–358.
- [15] W. D. JOUBERT AND T. A. MANTEUFFEL, *Iterative methods for non-symmetric linear systems*, in *Iterative Methods for Large Linear Systems*, D. R. Kincaid and L. J. Hayes, eds., Academic Press, New York, 1990, pp. 149–171.
- [16] L. LUKŠAN AND J. VLČEK, *Indefinitely preconditioned inexact Newton method for large sparse equality constrained non-linear programming problems*, *Numer. Linear Algebra Appl.*, 5 (1998), pp. 219–247.
- [17] J. L. MORALES-PÉREZ AND R. W. H. SARGENT, *On the implementation and performance of an interior point method for large sparse convex quadratic programming*, *Optim. Methods Softw.*, 1 (1992), pp. 153–168.
- [18] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite linear systems*, Technical Report NA99107, Numerical Analysis Group, University of Oxford, UK, 1999.
- [19] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, *SIAM J. Numer. Anal.*, 12 (1975), pp. 617–629.
- [20] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, *Classics Appl. Math.* 20, SIAM, Philadelphia, 1998.
- [21] B. T. POLYAK, *The conjugate gradient method in extremal problems*, *USSR Comput. Math. Math. Phys.*, 9 (1969), pp. 94–112.
- [22] W. H. PRESS, S. A. TEUKOLSKY, AND W. T. VETTERLING, *Numerical Recipes in Fortran: The Art of Scientific Computing*, Cambridge University Press, Cambridge, UK, 1992.
- [23] A. QUATERONI AND A. VALLI, *Numerical Approximation of Partial Differential Equations*, Springer-Verlag, Berlin, 1997.
- [24] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, Boston, 1996.
- [25] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, *SIAM J. Sci. Statist. Comput.*, 7 (1986), pp. 856–869.
- [26] D. J. SILVESTER AND A. J. WATHEN, *Fast iterative solution of stabilised Stokes systems. Part II: Using general block preconditioners*, *SIAM J. Numer. Anal.*, 31 (1994), pp. 1352–1367.
- [27] A. J. WATHEN AND D. J. SILVESTER, *Fast iterative solution of stabilised Stokes systems. Part I: Using simple diagonal preconditioners*, *SIAM J. Numer. Anal.*, 30 (1993), pp. 630–649.
- [28] A. J. WATHEN, *Realistic eigenvalue bounds for the Galerkin mass matrix*, *IMA J. Numer. Anal.*, 7 (1987), pp. 449–457.