

Constraint preconditioning for indefinite linear systems

Carsten Keller^{1,2}, Nicholas I. M. Gould^{3,4} and Andrew J. Wathen^{1,2}

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 which 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.

¹ Oxford University Computing Laboratory,
Parks Road, Oxford OX1 3QD, England, EU.
Email : ck@comlab.ox.ac.uk and wathen@comlab.ox.ac.uk

² Current reports available from
“<http://www.comlab.ox.ac.uk/oucl/publications/natr.html>” .

³ Department for Computation and Information, Rutherford Appleton Laboratory,
Chilton, Oxfordshire, OX11 0QX, England, EU.
Email : n.gould@rl.ac.uk

⁴ Current reports available from
“<http://www.numerical.rl.ac.uk/reports/reports.html>” .

Department for Computation and Information
Atlas Centre
Rutherford Appleton Laboratory
Oxfordshire OX11 0QX
February 9, 1999.

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

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

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 \mathcal{A} is non-singular, in which case B must be of full rank.

Example 1. Consider the problem of minimizing a function of n variables subject to m linear equality constraints on the variables, i.e.

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

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

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

where the λ_i are referred to as *Lagrangian multipliers*. By differentiating L with respect to x and λ the solution to (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. \square

Example 2 (The Stokes Problem). The Stokes equations in compact form are defined by

$$\left. \begin{aligned} -\nu \Delta u + \text{grad } p &= f \\ \text{div } u &= 0 \end{aligned} \right\} \text{ in } \Omega. \quad (1.3)$$

Discretising equations (1.3) together with the boundary conditions

$$u = 0 \text{ on } \Gamma$$

defines a linear system of equations of the form (1.1), where $b_1 = f$, $b_2 = 0$, $x_1 = u$ and $x_2 = p$. \square

Among the most important iterative methods currently available, Krylov subspace methods 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 generalised minimum residual method (GMRES) for unsymmetric matrices, although many other methods are available—see for example Greenbaum [12].

One common feature of the above methods is that the solution of the linear system (1.1) is found within $n + m$ iterations in exact arithmetic—see Joubert and Manteuffel [14, 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 non-singular preconditioners of the form

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

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), leads one to hope for a more favourable distribution of the eigenvalues of the (left-)preconditioned linear system

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

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 non-linear programming problems—see also Coleman [4], Polyak [18] and Gould et al. [11]. Here we derive arguments that confirm and extend some of the results in [16] and highlight the favourable features of a preconditioner of the form \mathcal{G} . Note that Golub and Wathen [10] recently considered a symmetric preconditioner of the form (1.4) for problems of the form (1.1) where A is non-symmetric.

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 behaviour 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, is small, as 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 non-normal systems convergence as opposed to termination is not so readily described—see Greenbaum [12, p. 5].

2.1 Eigenvalue Distribution

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

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \quad (2.6)$$

Let $QR = \begin{bmatrix} Y & Z \end{bmatrix} \begin{bmatrix} R^T & 0^T \end{bmatrix}^T$ be an orthogonal factorisation 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 (2.6) by the non-singular and square matrix

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

and postmultiplying by its transpose gives

$$\begin{bmatrix} Z^T AZ & Z^T AY & 0 \\ Y^T AZ & Y^T AY & R \\ 0 & R^T & 0 \end{bmatrix} \begin{bmatrix} x_z \\ x_y \\ y \end{bmatrix} = \lambda \begin{bmatrix} Z^T GZ & Z^T GY & 0 \\ Y^T GZ & Y^T GY & R \\ 0 & R^T & 0 \end{bmatrix} \begin{bmatrix} x_z \\ x_y \\ y \end{bmatrix}, \quad (2.7)$$

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 inter-

changes on both matrices in (2.7) 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

$$\mathcal{P} = \tilde{\mathcal{G}}^{-1}\tilde{\mathcal{A}} = \begin{bmatrix} & m & & & & \\ & & (n-m) & & & \\ & & & m & & \\ I & & 0 & & 0 & \\ \Theta & & (Z^T GZ)^{-1}(Z^T AZ) & & 0 & \\ \Upsilon & & \Gamma & & I & \end{bmatrix} \begin{matrix} m \\ (n-m) \\ m \end{matrix}. \quad (2.8)$$

Here the precise forms of Θ , Υ and Γ are irrelevant for the argument that follows; they are in general non-zero. 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 generalised eigenvalue problem $Z^T AZ x_z = \lambda Z^T GZ x_z$.

Note that the indefinite constrained preconditioner applied to the indefinite linear system (1.1) yields the preconditioned matrix \mathcal{P} which has real eigenvalues.

Remark 1. In the above argument we assumed that B has full row rank and consequently applied an orthogonal factorisation of B^T which resulted in a

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 (2.7), 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 equivalently apply to the reduced system of equations.

2.2 Eigenvector Distribution

We mentioned above that the termination for 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 (2.7), yielding

$$Z^T AZx_z + Z^T AYx_y = \lambda [Z^T GZx_z + Z^T GYx_y], \quad (2.9)$$

$$Y^T AZx_z + Y^T AYx_y + Ry = \lambda [Y^T GZx_z + Y^T GYx_y + Ry], \quad (2.10)$$

$$R^T x_y = \lambda R^T x_y. \quad (2.11)$$

From (2.11) it may be deduced that either $\lambda = 1$ or $x_y = 0$. In the former case equations (2.9) and (2.10) simplify to

$$Z^T AZx_z + Z^T AYx_y = Z^T GZx_z + Z^T GYx_y,$$

$$Y^T AZx_z + Y^T AYx_y = Y^T GZx_z + Y^T GYx_y,$$

which can consequently be written as

$$Q^T AQw = Q^T GQw, \quad (2.12)$$

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 (2.12) is equivalent to considering

$$Aw = \sigma Gw, \quad (2.13)$$

with $w \neq 0$ 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 (2.13)).

Now suppose $\lambda \neq 1$, in which case $x_y = 0$. Equations (2.9) and (2.10) yield

$$Z^T AZx_z = \lambda Z^T GZx_z \quad (2.14)$$

$$Y^T AZx_z + Ry = \lambda [Y^T GZx_z + Ry]. \quad (2.15)$$

The general eigenvalue problem (2.14) 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

$$\left[x_z^T \quad 0^T \quad y^T \right]^T.$$

If $x_z = 0$, we deduce from (2.15) that

$$Ry = \lambda Ry,$$

and hence that $y = 0$ since $\lambda \neq 1$. As $\left[x_z^T \quad x_y^T \quad y^T \right]^T = 0$ in this case, no extra eigenvectors arise.

Summarising 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 2.2. 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. 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 $\left[0^T \quad 0^T \quad y^T \right]^T$ 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 $\begin{bmatrix} x_z^T & 0^T & y^T \end{bmatrix}^T$ 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

$$\begin{aligned} \begin{bmatrix} 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} + \begin{bmatrix} x_{z_1}^{(2)} & \cdots & x_{z_i}^{(2)} \\ x_{y_1}^{(2)} & \cdots & x_{y_i}^{(2)} \\ y_1^{(2)} & \cdots & y_i^{(2)} \end{bmatrix} \begin{bmatrix} a_1^{(2)} \\ \vdots \\ a_i^{(2)} \end{bmatrix} \\ + \begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} \end{aligned} \quad (2.16)$$

implies that the vectors $a^{(k)}$ ($k = 1, \dots, 3$) are zero vectors. Multiplying (2.16) 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 $\sigma_k = 1$ ($k = 1, \dots, i$), whereas the last matrix arises from $\lambda_k \neq 1$ ($k = 1, \dots, j$), gives

$$\begin{aligned} \begin{bmatrix} 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} + \begin{bmatrix} x_{z_1}^{(2)} & \cdots & x_{z_i}^{(2)} \\ x_{y_1}^{(2)} & \cdots & x_{y_i}^{(2)} \\ y_1^{(2)} & \cdots & y_i^{(2)} \end{bmatrix} \begin{bmatrix} a_1^{(2)} \\ \vdots \\ a_i^{(2)} \end{bmatrix} \\ + \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_j \end{bmatrix} \begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}. \end{aligned} \quad (2.17)$$

Subtracting Equation (2.16) from (2.17) 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

$$\begin{bmatrix} x_{z_1}^{(3)} & \cdots & x_{z_j}^{(3)} \\ 0 & \cdots & 0 \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (2.18)$$

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 (2.18) are linearly independent and thus $a_k^{(3)} = 0$ ($k = 1, \dots, j$). Similarly, $a_k^{(2)} = 0$ ($k = 1, \dots, i$) follows from the linear independence of $\begin{bmatrix} x_{z_k}^{(2)T} & x_{y_k}^{(2)T} \end{bmatrix}^T$ ($k = 1, \dots, i$), and thus (2.16) simplifies to

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

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

Remark 2. Note that the result of Theorem 2.2 remains true if $Z^T(\gamma A + \sigma G)Z$ is positive definite for some scalars γ and σ —see Parlett [17, p. 343] for details.

To show that the eigenvector bounds of Theorem 2.2 can in fact be attained, consider the following two examples.

Example 3 (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 2.2. This eigenvector may be taken to be $\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$. \square

Example 4 (Maximum bound). Let $\mathcal{A} \in \mathbb{R}^{3 \times 3}$ be defined as in Example 3, 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 $\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T$, $\begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^T$, and $\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$. \square

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 (2.14) are highly sensitive to the choice of G . If G is a close approximation of A , we can expect a more favourable distribution of eigenvalues and consequently may expect faster convergence of an appropriate iterative method. In order to determine a good factorisation 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 2.3 (Cauchy's Interlace Theorem, see [17, Theorem 10.1.1]).

Suppose $\mathcal{T} \in \mathbb{R}^{n \times n}$ is symmetric, and that

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

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

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

Then

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

Proof. See Parlett [17, p. 203]. □

The applicability of Theorem 2.3 is verified by recalling the definitions of Q and Z given in the previous section, and by considering the generalised eigenvalue problems

$$Q^T A Q v = \alpha Q^T G Q v \tag{2.19}$$

and

$$Z^T A Z w = \lambda Z^T G Z w. \tag{2.20}$$

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 (2.19) and (2.20) gives

$$M^{-1} Q^T A Q M^{-T} u = \alpha u \quad (2.21)$$

and

$$L^{-1} Z^T A Z L^{-T} z = \lambda z, \quad (2.22)$$

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 $G^{-1} A$, equation (2.21) defines the same eigenvalues α_i ($i = 1, \dots, n$) as $G^{-1} A$. We may therefore apply Theorem 2.3 directly. The result is that the $n - m$ eigenvalues λ_i of (2.14) 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 [3].

Example 5. Consider the convex quadratic programming problem BLOWEYC which may be formulated as

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

Selecting a size parameter of 500 discretisation 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 by the above theory that the extreme eigenvalues of $G^{-1} A$ give a lower and upper bound for the $n - m$ eigenvalues defined by the general eigenvalue problem (2.14). 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 and multiplicity 502 removed. Rounded to two decimal places the numerical values of the two extreme eigenvalues of $G^{-1} A$ are



(a) Eigenvalues of $G^{-1}A$

(b) Eigenvalues of $(Z^T G Z)^{-1} Z^T A Z$

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

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 that 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. \square

It is readily seen from Example 5 that in this case the bounds provided by Theorem 2.3 are not descriptive in that there is significantly more clustering of the eigenvalues than implied by the theorem.

3 Convergence

In the context of this paper, the convergence of an iterative method under preconditioning is not only influenced 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 (2.8) 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 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.

The Krylov subspace theory states that the iteration with any method with an optimality property such as GMRES will terminate when the degree of the minimum polynomial is attained—see Axelsson [1, p. 463] (To be precise, the number may be less 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) and so the following theorems are relevant.

Theorem 3.1. 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. 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 (2.8) in its explicit form we observe that \mathcal{P} is in fact given by

$$\begin{bmatrix} I & 0 \\ \Upsilon & I \end{bmatrix}, \quad (3.23)$$

where Υ is non-zero 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 (3.23) 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. 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 3.1 is independent of G .

Remark 4. The important consequence of Theorem 3.1 is that termination of an iteration method such as GMRES will occur in at most 2 steps for any choice of b , even though the preconditioned matrix is not diagonalisable (unless $A = G$).

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. Assume $m < n$ and that \mathcal{A} is non-singular. 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

$$\begin{bmatrix} 0 & 0 & 0 \\ \left[\prod_{i=1}^{n-m} (S - \lambda_i I) \right] \Theta & (S - I) \prod_{i=1}^{n-m} (S - \lambda_i I) & 0 \\ \Phi_{n-m} & \Gamma \prod_{i=1}^{n-m} (S - \lambda_i I) & 0 \end{bmatrix} \quad (3.24)$$

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 (3.24) 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 diagonalisable. Thus (3.24) may be written as

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \Phi_{n-m} & 0 & 0 \end{bmatrix}, \quad (3.25)$$

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 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 (3.25) 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 3.2, 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 3.3. The following (randomly generated) example shows that the bound in Theorem 3.2 is attainable.

Example 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 (3.25) 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 3.2 is sharp. \square

Theorem 3.3. 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 $m < n$, \mathcal{A} is non-singular and that \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 3.2. 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

$$\begin{bmatrix} 0 & 0 & 0 \\ \left[\prod_{i=1}^k (S - \lambda_i I) \right] \Theta & (S - I) \prod_{i=1}^k (S - \lambda_i I) & 0 \\ \Phi_k & \Gamma \prod_{i=1}^k (S - \lambda_i I) & 0 \end{bmatrix} \quad (3.26)$$

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 (3.26) are in fact zero. It follows that, for $\Phi_k \neq 0$, a further multiplication of (3.26) 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 3.3 is attainable consider the following example.

Example 7. 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 $B = \begin{bmatrix} 0 & 0 & 1e^{-3} & 1e^{-3} \end{bmatrix}$. Then two of the $n - m$ eigenvalues that are defined by the generalised eigenvalue problem (2.14) are distinct and given by [2, 4]. It follows that the (3, 1) entry of (3.26) is non-zero with

$$\Phi_3 = \begin{bmatrix} -1071.41 \end{bmatrix},$$

and so the minimum polynomial is of order 4. \square

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 [20], where the preconditioner step is implemented by means of a symmetric indefinite factorisation of (1.4). Such a factorisation of the preconditioner may be much less demanding than the factorisation 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 [11].

4.1 Conjugate Gradients on a Reduced Linear System

In [11] Gould et al. propose a Conjugate Gradient like algorithm to solve equality constrained quadratic programming problems such as the one described in Example 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 Conjugate Gradients 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 satisfying $Bx = b$;

(2) Compute

$$\begin{aligned} r &= Ax + c \\ g &= ZW_{zz}^{-1}Z^T r \\ p &= -g \end{aligned} \tag{4.27}$$

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

$$\begin{aligned} \alpha &= r^T g / p^T A p \\ x &\leftarrow x + \alpha p \\ r^+ &= r + \alpha A p \\ r^+ &\leftarrow r^+ - B^T v \end{aligned} \tag{4.28}$$

$$g^+ = ZW_{zz}^{-1}Z^T r^+ \tag{4.29}$$

$$\begin{aligned} \beta &= (r^+)^T g^+ / r^T g \\ p &\leftarrow -g^+ + \beta p \\ g &\leftarrow g^+ \\ r &\leftarrow r^+ \end{aligned}$$

The computation of the preconditioned residual in (4.29) is often the most expensive computational factor in the algorithm. Gould et al. suggest avoiding the explicit use of the nullspace Z , but instead to compute g^+ by applying a

symmetric indefinite factorisation of

$$\begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} g^+ \\ v^+ \end{bmatrix} = \begin{bmatrix} r^+ \\ 0 \end{bmatrix}. \quad (4.30)$$

In practice (4.30) 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 [6]. In this context the factorisation consists of three separate routines, the first two of which analyse and factorise the matrix in (4.30). 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 in Step (1), solve for g in (4.27) and also to find g^+ in (4.29).

Remark 5. The computation of the projected residual g^+ is often accompanied by significant roundoff errors if this vector is much smaller than the residual r^+ . *Iterative refinement* is used in (4.28) to redefine r^+ so that its norm is closer to that of g^+ . The result is a dramatic reduction of the roundoff errors in the projection operation—see Gould et al. [11].

5 Numerical Results

We now present the results of numerical experiments that reinforce the analysis given in previous sections. The test problems we use are partly randomised sparse matrices (Table 5.1) and partly matrices that arise in linear and non-linear optimization (Table 5.2)—see Bongartz et al. [3]. As indicated throughout, all matrices are of the form

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix}, \quad (5.31)$$

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 [6]. To investigate possible favourable aspects of preconditioning it makes sense to compare unpreconditioned with preconditioned solution strategies. The indefinite nature of matrix (5.31) suggests the use of MINRES in the unpreconditioned case. 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 and 5.2 below), 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 and 5.2 below). The choice $G = \text{diag}(A)$ in the preconditioner is made for both PGMRES and RCG.

		Random I	Random II	Random III	Random IV
	n	578	2178	8450	8450
	m	32	128	236	46
	non-zero entries in A	2316	9740	39948	39948
	non-zero entries in B	427	1871	3600	686
MINRES	# of iterations	174	387	639	515
	time in seconds	0.4	3.1	17.5	13.1
PGMRES	# of iterations	46	87	228	242
	time in seconds	0.2	3.9	96.0	108.9
RCG	# of iterations	36	67	197	216
	time in seconds	0.1	1.0	5.9	5.9
	MA27 time in seconds	0.1	0.9	5.4	2.9

Table 5.1: Random test problems ($G = \text{diag}(A)$)

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 non-zero matrix elements—see Press et al. [19]. 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 factorisation and solution phases of subsequent routines. The amount of pivoting is controlled by the special 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 [6]. 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 enough accuracy by choosing the parameter value $u = 0.01$ in Tables 5.1 and 5.2.

The time measurements for the eight test examples indicate that the itera-

		BLOWEYC	CVXQP1	MOSARQP2	FIT2P
	n	1002	100	930	13525
	m	502	100	30	3000
	non-zero entries in A	3004	672	1020	13525
	non-zero entries in B	2503	295	148	50284
MINRES	# of iterations	363	no convergence	51	180
	time in seconds	1.8	no convergence	0.1	14.2
PMGREM	# of iterations	2	2	6	1
	time in seconds	0.6	0.1	0.2	13.2
RCG	# of iterations	2	2	6	1
	time in seconds	0.6	0.1	0.2	16.2
	MA27 time in seconds	0.5	0.1	0.1	15.6

Table 5.2: CUTE test problems ($G = \text{diag}(A)$)

tion counts for each of the three proposed iterative methods are comparable as far as operation counts, i.e. work, is concerned. The numerical results suggest that the inclusion of the (1, 2) and (2, 1) block of \mathcal{A} into the preconditioner, together with $G = \text{diag}(A)$, results in a considerable reduction of iterations, where the appropriate bounds of Theorems 3.1, 3.2 and 3.3 are attained in all cases. Specifically, Theorem 3.1 applies in context of problem CVXQP1.

Test problems RANDOM III and RANDOM IV in Table 5.1 emphasise the storage problems that are associated with the use of long recurrences in the PGMRES algorithm. The time required to find a solution to both RANDOM III and RANDOM IV 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 relevance of the time measurements for MA27 are commented on in the next section.

6 Conclusion

In this paper we investigated a new class of preconditioner for indefinite linear systems that incorporate the (2, 1) and (2, 2) 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 factorisation 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 system of the form (1.1) by means of appropriate subspace methods, we used a minimum polynomial argument. Theorem 3.1 considers the special condition $n = m$, in which case termination is guaranteed in two iterations. For $m < n$, Theorem 3.2 gives a general (sharp) upper bound on the dimension of the Krylov subspace, whereas Theorem 3.3 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 an upper and lower bound for the $n - m$ eigenvalues that are defined by the (2,2) block of matrix (2.8).

To confirm the analytical results in this paper we used three different subspace methods, MINRES of Paige and Saunders for the unpreconditioned matrix system and RCG of Gould et al. and also PGMRES of Saad and Schultz for the preconditioned case. Overall, the results show that the number of iterations is decreased substantially if preconditioning is applied. The Krylov subspaces that are built during the execution of the two preconditioned implementations are in theory of equal dimension for any of the eight 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 3.1, 3.2 and 3.3.

The time measurements for MA27 in the last section suggest that the preconditioned conjugate gradients algorithm, discussed in Section 4.1, is a suitable alternative to the direct solver. Whereas both MINRES and especially PGMRES are 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.

Acknowledgments

The authors would like to thank Gene. H. Golub for his insightful comments during the process of the work.

References

- [1] O. AXELSSON, *Iterative Solution Methods*, Cambridge University Press, Cambridge, 1996.
- [2] R. BARRETT, M. BERRY, T. F. CHAN, J. DEMMEL, J. M. 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, PA, 1993.
- [3] I. BONGARTZ, A. CONN, N. GOULD, AND P. TOINT, *CUTE: Constrained and Unconstrained Testing Environment*, ACM Transactions on Mathematical Software, 21 (1995), pp. 123-160.
- [4] T. F. COLEMAN, *Linearly constrained optimization and projected preconditioned conjugate gradients*, In J. Lewis, editor, Proceedings of the Fifth SIAM Conference on Applied Linear Algebra, SIAM, Philadelphia, PA, 1994, pp. 188-122.
- [5] I. S. DUFF, A. M. ERISMAN, AND J. K. REID, *Direct Methods for Sparse Matrices*, Clarendon Press, Oxford, 1989.
- [6] I.S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear equations*, ACM Transactions on Mathematical Software, 9 (3) (1983), pp. 302-325.
- [7] H. C. ELMAN, *Perturbation of eigenvalues of preconditioned Navier-Stokes operators*, SIAM Journal on Matrix Analysis and Applications, 18 (3) (1997), pp. 733-751.
- [8] B. FISCHER, *Polynomial Based Iteration Methods for Symmetric Linear Systems*, Wiley, New York, 1996
- [9] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, 1996.
- [10] G. H. GOLUB AND A. J. WATHEN, *An iteration for indefinite system and its application to the Navier-Stokes equations*, SIAM Journal on Scientific Computing, 19 (2) (1998), pp. 530-539.

- [11] 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, 1998.
- [12] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, PA, 1997.
- [13] A. GREENBAUM, AND L. GURVITS, *Max-min properties of matrix factor norms*, SIAM Journal on Scientific Computation, 15 (1994), pp. 348-358.
- [14] W. D. JOUBERT AND T. A. MANTEUFFEL, *Iterative methods for non-symmetric linear systems*, In D. R. Kincaid and L. J. Hayes, editor, *Iterative Methods for Large Linear Systems*, Academic Press, 1990, pp. 149-171.
- [15] C. T. KELLEY, *Iterative Methods for Linear and Nonlinear Equations*, SIAM, Philadelphia, PA, 1995.
- [16] L. LUKŠAN AND J. VLČEK, *Indefinitely preconditioned inexact newton method for large sparse equality constrained non-linear programming problems*, Numerical Linear Algebra with Applications, 5 (1998), pp. 219-247.
- [17] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, SIAM Classics in Applied Mathematics, Philadelphia, PA, 1998.
- [18] B. T. POLYAK *The conjugate gradient method in extremal problems*, U.S.S.R. Computational Mathematics and Mathematical Physics, 9 (1969), pp. 94-112.
- [19] W. H. PRESS, S. A. TEUKOLSKY, AND W. T. VETTERLING, *Numerical Recipes in Fortran: The Art of Scientific Computing*, Cambridge University Press, Cambridge, 1992.
- [20] Y. SAAD AND M. H. SCHULTZ, *GMRES: a generalised minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (3) (1986), pp. 856-869.