The University of Reading
Department of Mathematics

Constraint-style preconditioners for regularized saddle point problems

H. S. Dollar

Numerical Analysis Report 3/06

March 27, 2006

Constraint-style preconditioners for regularized saddle point problems

H. S. Dollar[1] [2]

Abstract

The problem of finding good preconditioners for the numerical solution of an important class of indefinite linear systems is considered. These systems are of a regularized saddle point structure

$$\begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} c \\ d \end{bmatrix},$$

where $A \in \mathbb{R}^{n\times n}$, $C \in \mathbb{R}^{m\times m}$ are symmetric and $B \in \mathbb{R}^{m\times n}$.

In *Constraint preconditioning for indefinite linear systems,* SIAM J. Matrix Anal. Appl., 21 (2000), Keller, Gould and Wathen analyze the idea of using constraint preconditioners that have a specific 2 by 2 block structure for the case of C being zero. We shall extend this idea by allowing the (2,2) block to be symmetric and positive semi-definite. Results concerning the spectrum and form of the eigenvectors are presented, as are numerical results to validate our conclusions.

1 Introduction

Recently, a large amount of work has been devoted to the problem of solving large linear systems in saddle point form. Such systems arise in a wide variety of technical and scientific applications. For example, interior point methods in both linear and nonlinear optimization require the solution of a sequence of systems in saddle point form [26]. Another popular field, which is a major source of saddle point problems, is that of mixed finite element methods in engineering fields, see [10] and [19, Chapters 7,9]. An excellent survey of numerical methods for algebraic saddle point problems has been written by Benzi, Golub and Liesen [3].

We wish to find the solution of block 2×2 linear systems of the form

$$\underbrace{\begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}}_{\mathcal{A}} \begin{bmatrix} x \\ y \end{bmatrix} = \underbrace{\begin{bmatrix} c \\ d \end{bmatrix}}_{b}, \tag{1}$$

where $A \in \mathbb{R}^{n\times n}$, $C \in \mathbb{R}^{m\times m}$ are symmetric and $B \in \mathbb{R}^{m\times n}$. We shall assume that $m \leq n$ and B is of full rank. If A and C are positive definite, then the matrix $\mathcal{A}$ is a permuted quasidefinite matrix, [25]. Vanderbei has shown

[1]Department of Mathematics, University of Reading, Whiteknights, P.O. Box 220, Reading, Berkshire, RG6 6AX, U.K. (H.S.Dollar@reading.ac.uk).

[2]This authors work was supported by an Oxford University Computing Laboratory Doctoral Training Account.

that quasidefinite matrices are strongly factorizable , i.e., a Cholesky-like factorization LDL^T exists for any symmetric row and column permutation of the quasidefinite matrix, [25]. The diagonal matrix has n positive and m negative pivots. However, we shall not confine ourselves to quasidefinite matrices.

It may be attractive to use iterative methods to solve systems such as (1), particularly for large m and n. In particular, Krylov subspace methods might be used. It is often advantageous to use a preconditioner, $\mathcal{P}$, with such iterative methods. The preconditioner should reduce the number of iterations required for convergence but not significantly increase the amount of computation required at each iteration, [24, Chapter 13].

In Section 2 we shall firstly review the well known spectral properties of a technique commonly known as constraint preconditioning when $C = 0$ [15, 17]. For the case of $C = 0$, a constraint preconditioner exactly reproduces the (constraint) blocks B, B^T and the $C = 0$ block. It is restrictive to assume that the matrix C in the saddle point systems is always a zero matrix: a number of situations arise in which $C \neq 0$ [1, 16, 23]. In all these cases, C is positive semi-definite and, hence, we shall consider the idea of extending constraint preconditioners to the case of C being positive semi-definite. In particular, the preconditioner will exactly reproduce the B, B^T and C blocks, whilst the A block will be replaced by a symmetric block which we refer to as G; this is considered in Sections 3 and 4. Such a preconditioner has been considered before, for example, Perugia and Simoncini consider the case of $G = I$ [18], and Siefert and de Sturler assume that G is nonsingular [22], but we show that these assumptions can be relaxed. In Section 5 we shall report numerical results where our preconditioners have been used to solve various test problems.

2 Constraint preconditioners

Let us initially assume that $C = 0$. Keller, Gould and Wathen [15] investigated the spectral properties of the resulting preconditioned system when we use of a preconditioner of the form

$$\mathcal{P} = \begin{bmatrix} G & B^T \\ B & 0 \end{bmatrix}, \tag{2}$$

where G approximates but (in general) is not the same as A. They were able to prove various results about the eigenvalues and eigenvectors for the preconditioned systems $\mathcal{P}^{-1}\mathcal{A}$, where $\mathcal{A}$ and $\mathcal{P}$ are defined in (1) and (2) respectively. $\mathcal{P}$ is called a *constraint preconditioner*. Proof of the following theorem can be found in [15].

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

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

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

1. *an eigenvalue at 1 with multiplicity $2m$;*
2. *$n-m$ eigenvalues λ which are defined by the generalized eigenvalue problem $Z^TAZx_z = \lambda Z^TGZx_z$.*

This accounts for all of the eigenvalues.

If either Z^TAZ or Z^TGZ are positive definite, then the indefinite preconditioner $\mathcal{P}$ applied to the indefinite saddle point matrix $\mathcal{A}$ with $C = 0$ yields a preconditioned matrix $\mathcal{P}^{-1}\mathcal{A}$ which has real eigenvalues [15]. If both Z^TAZ and Z^TGZ are positive definite, then we can use a projected preconditioned conjugate gradient method to find x and y, see [12]. Results about the associated eigenvectors and the Krylov subspace dimension can also be found in [15].

3 Constraint preconditioners for the case of symmetric and positive definite C

In this section we shall assume that the matrix C is symmetric and positive definite. The term *constraint preconditioner* was introduced in [15] because the (1,2) and (2,1) matrix blocks of the preconditioner are exact representations of those in $\mathcal{A}$, where these blocks represent constraints. However, we also observe that the (2,2) matrix block is an exact representation when $C = 0$. This motivates the generalization of the constraint preconditioner to take the form

$$\mathcal{P} = \begin{bmatrix} G & B^T \\ B & -C \end{bmatrix}, \tag{3}$$

where $G \in \mathbb{R}^{n\times n}$ approximates, but is, in general, not the same as A.

For symmetric matrix systems, the convergence of an applicable iterative method is determined by the distribution of the eigenvalues of the coefficient matrix. It is often desirable for the number of distinct eigenvalues to be small so that the rate of convergence is rapid. For non-normal systems the convergence is not so readily described, see [14, page 6].

We shall use the following assumptions in the theorems of this section:

A1 $C \in \mathbb{R}^{m\times m}$ is symmetric and positive definite,

A2 $A \in \mathbb{R}^{n\times n}$ is symmetric,

A3 $B \in \mathbb{R}^{m\times n}$ $(m < n)$ has full rank,

A4 $G \in \mathbb{R}^{n\times n}$ is symmetric,

A5 $\mathcal{A} \in \mathbb{R}^{(n+m)\times(n+m)}$ is as defined in (1),

A6 $\mathcal{P} \in \mathbb{R}^{(n+m)\times(n+m)}$ is as defined in (3).

In the next section **A1** will be relaxed.

Theorem 3.1. *Assume that **A1**-**A6** hold, then the matrix $\mathcal{P}^{-1}\mathcal{A}$ has*

- *an eigenvalue at 1 with multiplicity m, and*
- *n eigenvalues which are defined by the generalized eigenvalue problem*

$$\left(A + B^T C^{-1} B\right) x = \lambda \left(G + B^T C^{-1} B\right) x.$$

This accounts for all of the eigenvalues.

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

$$\begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \begin{bmatrix} G & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \tag{4}$$

Expanding this out we obtain

$$Ax + B^T y = \lambda G x + \lambda B^T y, \tag{5}$$

and

$$Bx - Cy = \lambda B x - \lambda C y. \tag{6}$$

Equation (6) implies that either $\lambda = 1$ or $Bx - Cy = 0$. If the former holds then (5) becomes

$$Ax = Gx. \tag{7}$$

Equation (7) is trivially satisfied by $x = 0$ and, hence, there are m linearly independent eigenvectors of the form $\begin{bmatrix} 0^T & y^T \end{bmatrix}$ associated with the unit eigenvalue. If there exist any $x \neq 0$ which satisfy (7), then there will be a i $(0 \leq i \leq n)$ linearly independent eigenvectors of the form $\begin{bmatrix} x^T & y^T \end{bmatrix}$ where the components x arise from the generalized eigenvalue problem $Ax = Gx$.

If $\lambda \neq 1$, then (6) implies that

$$y = C^{-1} B x.$$

Substituting this into (5) yields the generalized eigenvalue problem

$$\left(A + B^T C^{-1} B\right) x = \lambda \left(G + B^T C^{-1} B\right) x. \tag{8}$$

Thus, the non-unit eigenvalues of $\mathcal{P}^{-1}\mathcal{A}$ are defined as the non-unit eigenvalues of (8). Noting that if (8) has any unit eigenvalues, then the values of $x(\neq 0)$ which satisfy this are exactly those which arise from the generalized eigenvalue problem $Ax = Gx$, we complete our proof. □

If $A + B^T C^{-1} B$ or $G + B^T C^{-1} B$ are positive definite, then the preconditioned system has real eigenvalues. If both $A + B^T C^{-1} B$ and $G + B^T C^{-1} B$ are positive definite, then we can apply a projected preconditioned conjugate gradient method to find x and y [7, 11]. We also observe that if C has a small 2-norm, $\|A\|_2 = \mathcal{O}(1)$ and $\|G\|_2 = \mathcal{O}(1)$, then the $B^T C^{-1} B$ terms will dominate the generalized eigenvalue problem (8) for $Bx \neq 0$ and, hence, there will be at least a further m eigenvalues clustered about 1 for $\|C\|_2 \ll 1$. This additional clustering of part of the spectrum of $\mathcal{P}^{-1}\mathcal{A}$ will often translate into a speeding up of the convergence of a selected Krylov subspace method, [2, Section 1.3].

Theorem 3.2. *Assume that **A1**-**A6** hold and $G + B^T C^{-1} B$ is positive definite, then the matrix $\mathcal{P}^{-1}\mathcal{A}$ has $n + m$ eigenvalues as defined in Theorem 3.1 and $m + i + j$ linearly independent eigenvectors. There are*

- *m eigenvectors of the form $\begin{bmatrix} 0^T & y^T \end{bmatrix}$ that correspond to the case $\lambda = 1$,*
- *i $(0 \leq i \leq n)$ eigenvectors of the form $\begin{bmatrix} x^T & y^T \end{bmatrix}$ arising from $Ax = \sigma Gx$ for which the i vectors x are linearly independent, $\sigma = 1$, and $\lambda = 1$, and*
- *j $(0 \leq j \leq n)$ eigenvectors of the form $\begin{bmatrix} x^T & y^T \end{bmatrix}$ that correspond to the case $\lambda \neq 1$.*

Proof. The form of the eigenvectors follows directly from the proof of Theorem 3.1. It remains for us to show that

$$\begin{bmatrix} 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} + \begin{bmatrix} x_1^{(2)} & \cdots & x_i^{(2)} \\ y_1^{(2)} & \cdots & y_i^{(2)} \end{bmatrix} \begin{bmatrix} a_1^{(2)} \\ \vdots \\ a_i^{(2)} \end{bmatrix} + \begin{bmatrix} x_1^{(3)} & \cdots & x_j^{(3)} \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} a_1^{(3)} \\ \vdots \\ a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} \quad (9)$$

implies that the vectors $a^{(k)}$ $(k = 1, 2, 3)$ are zero vectors. Multiplying (9) by $\mathcal{P}^{-1}\mathcal{A}$, and recalling that in the previous equation the first matrix arises from the case $\lambda_k = 1$ $(k = 1, \ldots, m)$, the second matrix from the case $\lambda_k = 1$ and $\sigma_k = 1$ $(k = 1, \ldots, i)$, and the last matrix from $\lambda_k \neq 1$ $(k = 1, \ldots, j)$, gives

$$\begin{bmatrix} 0 & \cdots & 0 \\ y_1^{(1)} & \cdots & y_m^{(1)} \end{bmatrix} \begin{bmatrix} a_1^{(1)} \\ \vdots \\ a_m^{(1)} \end{bmatrix} + \begin{bmatrix} x_1^{(2)} & \cdots & x_i^{(2)} \\ y_1^{(2)} & \cdots & y_i^{(2)} \end{bmatrix} \begin{bmatrix} a_1^{(2)} \\ \vdots \\ a_i^{(2)} \end{bmatrix} + \begin{bmatrix} x_1^{(3)} & \cdots & x_j^{(3)} \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} \lambda_1 a_1^{(3)} \\ \vdots \\ \lambda_j a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} \quad (10)$$

Subtracting (9) from (10) we obtain

$$\begin{bmatrix} x_1^{(3)} & \cdots & x_j^{(3)} \\ y_1^{(3)} & \cdots & y_j^{(3)} \end{bmatrix} \begin{bmatrix} (\lambda_1 - 1)a_1^{(3)} \\ \vdots \\ (\lambda_j - 1)a_j^{(3)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}.$$

The assumption that $G + B^T C^{-1} B$ is positive definite implies that $x_k^{(3)}$ $(k = 1, \ldots, j)$ are linearly independent and thus that $(\lambda_k - 1)a_1^{(3)} = 0$, $(k = 1, \ldots, j)$. The eigenvalues λ_k $(k = 1, \ldots, j)$ are non-unit which implies that $a_k^{(3)} = 0$ $(k = 1, \ldots, j)$. We also have linear independence of $x_k^{(2)}$ $(k = 1, \ldots, i)$ and thus $a_k^{(2)} = 0$ $(k = 1, \ldots, i)$. Equation (9) simplifies to

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

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

Krylov subspace theory states that iteration with any method with an optimality property, e.g. GMRES [21], will terminate when the degree of the minimum polynomial is attained. This is also true of some other (non-optimal) practical iterations such as BiCGTAB as long as failure does not occur. In particular, the degree of the minimum polynomial is equal to the dimension of the corresponding Krylov subspace $\mathcal{K}\left(\mathcal{P}^{-1}\mathcal{A}, b\right)$ (for general b), [20, Proposition 6.1].

Theorem 3.3. *Assume that* ***A1-A6*** *hold and $G + B^T C^{-1} B$ is positive definite, then the dimension of the Krylov subspace $\mathcal{K}\left(\mathcal{P}^{-1}\mathcal{A}, b\right)$ is at most* $\min\{n+2, n+m\}$.

Proof. As in the proof of Theorem 3.1, the generalized eigenvalue problem is

$$\begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \begin{bmatrix} G & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \tag{11}$$

Suppose that the preconditioned matrix $\mathcal{P}^{-1}\mathcal{A}$ takes the form

$$\mathcal{P}^{-1}\mathcal{A} = \begin{bmatrix} \Theta_1 & \Theta_3 \\ \Theta_2 & \Theta_4 \end{bmatrix}, \tag{12}$$

where $\Theta_1 \in \mathbb{R}^{n \times n}$, $\Theta_2 \in \mathbb{R}^{m \times n}$, $\Theta_3 \in \mathbb{R}^{n \times m}$, and $\Theta_4 \in \mathbb{R}^{m \times m}$. Using the facts that $\mathcal{P}\left(\mathcal{P}^{-1}\mathcal{A}\right) = \mathcal{A}$ and B has full row rank, we obtain $\Theta_3 = 0$ and $\Theta_4 = I$. The precise forms of Θ_1 and Θ_2 are irrelevant for the argument that follows.

From the earlier eigenvalue derivation, it is evident that the characteristic polynomial of the preconditioned linear system (12) is

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

In order to prove the upper bound on the Krylov subspace dimension, we need to show that the order of the minimum polynomial is less than or equal to $\min\{n+2, n+m\}$. Expanding the polynomial $\left(\mathcal{P}^{-1}\mathcal{A} - I\right)\prod_{i=1}^{n}\left(\mathcal{P}^{-1}\mathcal{A} - \lambda_i I\right)$ of degree $n+1$, we obtain

$$\begin{bmatrix} (\Theta_1 - I)\prod_{i=1}^{n}(\Theta_1 - \lambda_i I) & 0 \\ \Theta_2 \prod_{i=1}^{n}(\Theta_1 - \lambda_i I) & 0 \end{bmatrix}.$$

Since Θ_1 has a full set of linearly independent eigenvectors, Θ_1 is diagonalizable. Hence,

$$(\Theta_1 - I)\prod_{i=1}^{n}(\Theta_1 - \lambda_i I) = 0.$$

We therefore obtain

$$\left(\mathcal{P}^{-1}\mathcal{A} - I\right)\prod_{i=1}^{n}\left(\mathcal{P}^{-1}\mathcal{A} - \lambda_i I\right) = \begin{bmatrix} 0 & 0 \\ \Theta_2 \prod_{i=1}^{n}(\Theta_1 - \lambda_i I) & 0 \end{bmatrix}. \qquad (13)$$

If $\Theta_2 \prod_{i=1}^{n}(\Theta_1 - \lambda_i I) = 0$, then the order of the minimum polynomial of $\mathcal{P}^{-1}\mathcal{A}$ is less than or equal to $\min\{n+1, n+m\}$. If $\Theta_2 \prod_{i=1}^{n}(\Theta_1 - \lambda_i I) \neq 0$, then the dimension of $\mathcal{K}\left(\mathcal{P}^{-1}\mathcal{A}, b\right)$ is at most $\min\{n+2, n+m\}$ since multiplication of (13) by another factor $\left(\mathcal{P}^{-1}\mathcal{A} - I\right)$ gives the zero matrix. □

Theorem 3.3 tells us that with preconditioner

$$\mathcal{P} = \begin{bmatrix} G & B^T \\ B & -C \end{bmatrix}$$

for

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}$$

the dimension of the Krylov subspace is no greater than $\min\{n+2, n+m\}$ under appropriate assumptions. Hence, termination (in exact arithmetic) is guaranteed in a number of iterations smaller than this.

4 Constraint preconditioners for the case of symmetric and positive semi-definite C

We shall relax assumption **A1** and instead make the following assumptions in the theorems of this section:

B1 $C \in \mathbb{R}^{m\times m}$ is symmetric and positive semi-definite, and has rank p where $0 < p < m$,

B2 C is factored as $C = EDE^T$, where $E \in \mathbb{R}^{m\times p}$, and $D \in \mathbb{R}^{p\times p}$ is non-singular,

B3 The matrix $F \in \mathbb{R}^{m\times(m-p)}$ is such that its columns span the nullspace of C,

B4 $\begin{bmatrix} E & F \end{bmatrix} \in \mathbb{R}^{m\times m}$ is orthogonal,

B5 The columns of $N \in \mathbb{R}^{n\times(n-m+p)}$ span the nullspace of F^TB.

The exact form of the factorization of C in **B2** is clearly not relevant and, also, clearly not unique – a spectral decomposition is a possibility.

Theorem 4.1. *Assume that **A2**-**A6** and **B1**-**B5** hold, then the matrix $\mathcal{P}^{-1}\mathcal{A}$ has*

- *an eigenvalue at 1 with multiplicity $2m-p$, and*
- *$n-m+p$ eigenvalues which are defined by the generalized eigenvalue problem*

$$N^T\left(A + B^TED^{-1}E^TB\right)Nz = \lambda N\left(G + B^TED^{-1}E^TB\right)Nz.$$

This accounts for all of the eigenvalues.

Proof. Any $y \in \mathbb{R}^m$ can be written as $y = Ey_e + Fy_f$. Substituting this into the generalized eigenvalue problem (4) and premultiplying by $\begin{bmatrix} I & 0 \\ 0 & E^T \\ 0 & F^T \end{bmatrix}$ we obtain

$$\left[\begin{array}{cc|c} A & B^TE & B^TF \\ E^TB & -D & 0 \\ \hline F^TB & 0 & 0 \end{array}\right]\begin{bmatrix} x \\ y_e \\ y_f \end{bmatrix} = \lambda\left[\begin{array}{cc|c} G & B^TE & B^TF \\ E^TB & -D & 0 \\ \hline F^TB & 0 & 0 \end{array}\right]\begin{bmatrix} x \\ y_e \\ y_f \end{bmatrix}. \quad (14)$$

Noting that the (3,3) block has dimension $(m-p)\times(m-p)$ and is a zero matrix in both coefficient matrices, we can apply Theorem 2.1 from [15] to obtain that $\mathcal{P}^{-1}\mathcal{A}$ has

- an eigenvalue at 1 with multiplicity $2(m-p)$, and
- $n-m+2p$ eigenvalues which are defined by the generalized eigenvalue problem

$$\overline{N}^T\begin{bmatrix} A & B^TE \\ E^TB & -D \end{bmatrix}\overline{N}w_n = \lambda\overline{N}^T\begin{bmatrix} G & B^TE \\ E^TB & -D \end{bmatrix}\overline{N}w_n, \quad (15)$$

where $\overline{N}$ is an $(n+p) \times (n-m+2p)$ basis for the nullspace of $\begin{bmatrix} F^T B & 0 \end{bmatrix} \in \mathbb{R}^{(m-p)\times(n+p)}$, and

$$\begin{bmatrix} x \\ y_e \end{bmatrix}^T = \overline{N} w_n + \begin{bmatrix} B^T F \\ 0 \end{bmatrix} w_b.$$

Letting $\overline{N} = \begin{bmatrix} N & 0 \\ 0 & I \end{bmatrix}$, then (15) becomes

$$\begin{bmatrix} N^T A N & N^T B^T E \\ E^T B N & -D \end{bmatrix} \begin{bmatrix} w_{n1} \\ w_{n2} \end{bmatrix} = \lambda \begin{bmatrix} N^T G N & N^T B^T E \\ E^T B N & -D \end{bmatrix} \begin{bmatrix} w_{n1} \\ w_{n2} \end{bmatrix}. \quad (16)$$

This generalized eigenvalue problem is exactly that of the form considered in Theorem 3.1, so (16) has an eigenvalue at 1 with multiplicity p and the remaining eigenvalues are defined by the generalized eigenvalue problem

$$N^T \left(A + B^T E D^{-1} E^T B\right) N w_{n1} = \lambda N^T \left(G + B^T E D^{-1} E^T B\right) N w_{n1}. \quad (17)$$

Hence, $\mathcal{P}^{-1}\mathcal{A}$ has an eigenvalue at one with multiplicity $2m - p$ and the other eigenvalues are defined by the generalized eigenvalue problem (17). □

As for the cases $C = 0$ and C nonsingular, we are able obtain conditions which guarantee that the eigenvalues are real and for which a projected preconditioned conjugate gradient method could be applied to find x and y; respectively, these conditions are:

- either $N^T \left(A + B^T E D^{-1} E^T B\right) N$ or $N^T \left(G + B^T E D^{-1} E^T B\right) N$ are positive definite,
- both $N^T \left(A + B^T E D^{-1} E^T B\right) N$ and $N^T \left(G + B^T E D^{-1} E^T B\right) N$ are positive definite.

Interestingly, the projected preconditioned conjugate gradient method is also derived by the use of a factorization of C as in the assumption **B2**; transformations are then used to remove the requirement of needing to factorize C [7].

Similarly to the case $p = m$, if C has a small 2-norm, $\|A\| = \mathcal{O}(1)$ and $\|G\| = \mathcal{O}(1)$, then the $N^T B^T E D^{-1} E^T B N$ terms will dominate the generalized eigenvalue problem (17) for $E^T B N w_{n1} \neq 0$ and, hence, there will be at least p further eigenvalues clustered about 1 for $\|C\|_2 \ll 1$.

Theorem 4.2. *Assume that **A2-A6**, **B1-B5** hold and $G + B^T E D^{-1} E^T B$ is positive definite, then the matrix $\mathcal{P}^{-1}\mathcal{A}$ has $n + m$ eigenvalues as defined in Theorem 3.1 and $m + i + j$ linearly independent eigenvectors. There are*

- *m eigenvectors of the form $\begin{bmatrix} 0^T & y^T \end{bmatrix}$ that correspond to the case $\lambda = 1$,*
- *i $(0 \leq i \leq n)$ eigenvectors of the form $\begin{bmatrix} x^T & y^T \end{bmatrix}$ arising from $Ax = \sigma Gx$ for which the i vectors x are linearly independent, $\sigma = 1$, and $\lambda = 1$, and*

- j $(0 \leq j \leq n)$ *eigenvectors of the form* $\begin{bmatrix} x^T & y^T \end{bmatrix}$ *that correspond to the case* $\lambda \neq 1$.

Proof. Proof of the form and linear independence of the $m + i + j$ eigenvalues obtained in a similar manner to the proof Theorem 3.2. □

To show that both the lower and upper bounds on the number of linearly independent eigenvectors can be attained we need only consider variations on Examples 2.5 and 2.6 from [15].

Example 4.1 (minimum bound). *Consider the matrices*

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

such that $m = 2$, $n = 2$ *and* $p = 1$. *The preconditioned matrix* $\mathcal{P}^{-1}\mathcal{A}$ *has an eigenvalue at* 1 *with multiplicity* 4, *but only two linearly independent eigenvectors which arise from case* (1) *of Theorem* 4.2. *These eigenvectors may be taken to be* $\begin{bmatrix} 0 & 0 & 1 & 0 \end{bmatrix}^T$ *and* $\begin{bmatrix} 0 & 0 & 0 & 1 \end{bmatrix}^T$.

Example 4.2 (maximum bound). *Let* $\mathcal{A} \in \mathbb{R}^{4\times 4}$ *be as defined in Example* 4.1, *but assume that* $G = A$. *The preconditioned matrix* $\mathcal{P}^{-1}\mathcal{A}$ *has an eigenvalue at* 1 *with multiplicity* 4 *and clearly a complete set of eigenvectors. These may be taken to be the columns of the identity matrix.*

The linear independence of the $m+i+j$ eigenvectors allows us to obtain an upper bound on the Krylov subspace $\mathcal{K}\left(\mathcal{P}^{-1}\mathcal{A}, b\right)$:

Theorem 4.3. *Assume that* ***A2-A6***, ***B1-B5*** *hold and* $G + B^T E D^{-1} E^T B$ *is positive definite, then the dimension of the Krylov subspace* $\mathcal{K}\left(\mathcal{P}^{-1}\mathcal{A}, b\right)$ *is at most* $\min\{n - m + p + 2, n + m\}$.

Proof. As in the proof of Theorem 3.3, the preconditioned matrix $\mathcal{P}^{-1}\mathcal{A}$ takes the form

$$\mathcal{P}^{-1}\mathcal{A} = \begin{bmatrix} \Theta_1 & 0 \\ \Theta_2 & I \end{bmatrix}, \tag{18}$$

where $\Theta_1 \in \mathbb{R}^{n\times n}$, and $\Theta_2 \in \mathbb{R}^{m\times n}$. The precise forms of Θ_1 and Θ_2 are irrelevant for the argument that follows.

From the earlier eigenvalue derivation, it is evident that the characteristic polynomial of the preconditioned linear system (18) is

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

In order to prove the upper bound on the Krylov subspace dimension, we need to show that the order of the minimum polynomial is less than or equal to $\min\{n-m+p+2, n+m\}$. Expanding the polynomial $\left(\mathcal{P}^{-1}\mathcal{A} - I\right)\prod_{i=1}^{n-m+p}\left(\mathcal{P}^{-1}\mathcal{A} - \lambda_i I\right)$ of degree $n+1$, we obtain

$$\begin{bmatrix} (\Theta_1 - I)\prod_{i=1}^{n-m+p}(\Theta_1 - \lambda_i I) & 0 \\ \Theta_2 \prod_{i=1}^{n-m+p}(\Theta_1 - \lambda_i I) & 0 \end{bmatrix}.$$

Since $G + B^T E D^{-1} E^T B$ is positive definite, Θ_1 has a full set of linearly independent eigenvectors and is diagonalizable. Hence, $(\Theta_1 - I)\prod_{i=1}^{n-m+p}(\Theta_1 - \lambda_i I) = 0$. We therefore obtain

$$\left(\mathcal{P}^{-1}\mathcal{A} - I\right)\prod_{i=1}^{n-m+p}\left(\mathcal{P}^{-1}\mathcal{A} - \lambda_i I\right) = \begin{bmatrix} 0 & 0 \\ \Theta_2 \prod_{i=1}^{n-m+p}(\Theta_1 - \lambda_i I) & 0 \end{bmatrix}. \quad (19)$$

If $\Theta_2 \prod_{i=1}^{n-m+p}(\Theta_1 - \lambda_i I) = 0$, then the order of the minimum polynomial of $\mathcal{P}^{-1}\mathcal{A}$ is less than or equal to $\min\{n-m+p+1, n+m\}$. If $\Theta_2 \prod_{i=1}^{n-m+p}(\Theta_1 - \lambda_i I) = 0$, then the dimension of $\mathcal{K}\left(\mathcal{P}^{-1}\mathcal{A}, b\right)$ is at most $\min\{n-m+p+2, n+m\}$ since multiplication of (19) by another factor $\left(\mathcal{P}^{-1}\mathcal{A} - I\right)$ gives the zero matrix. □

Thus, in exact arithmetic, iteration with any method with an optimality condition will terminate in at most $\min\{n-m+p+2, n+m\}$ iterations. We observe that if $p = m$, then Theorem 4.3 gives the same bound on the Krylov subspace dimension as that in Theorem 3.3 and if $p = 0$, then we obtain the results of [15].

5 Numerical results

The CUTEr test set [13] provides a set of quadratic programming problems. We shall firstly use a problem from this set to illustrate how changing the rank of C affects the multiplicity of the unit eigenvalues and the termination of GMRES, and then present results of numerical tests which compare the four different approaches to finding solutions to (1). All tests were performed in MATLAB® 7.01.

The CVXQP1_S problem from the CUTEr test set is small with $n = 100$ and $m = 50$. It is a convex quadratic program whose constraints are linear; it is a purely academic problem which has been constructed specifically for test problems. "Barrier" penalty terms (in this case 1.1) are added to the diagonal of A to simulate systems that might arise during an iteration of an

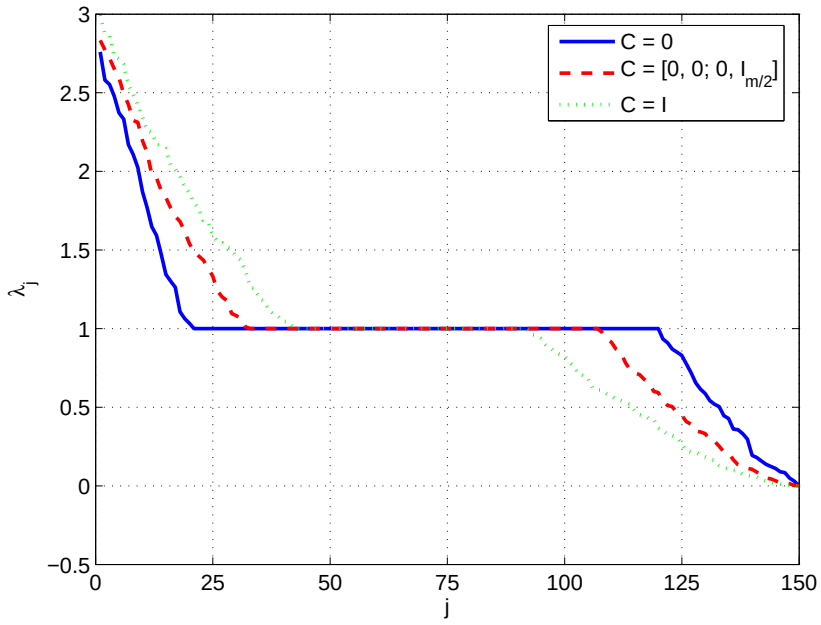

Figure 1: Distribution of the eigenvalues of $\mathcal{P}^{-1}\mathcal{A}$ for various choices of C.

interior-point method for such problems. We shall set $G = diag(A)$, $C = diag(0, \ldots, 0, 1, \ldots, 1)$ and vary the number of zeros on the diagonal of C so as to change its rank.

In Figure 1, we illustrate the change in the eigenvalues of the preconditioned system $\mathcal{P}^{-1}\mathcal{A}$ for three different choices of C. The eigenvalues are sorted so that

$$\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_{n+m}.$$

When $C = 0$, we expect there to be at least $2m$ unit eigenvalues [15]. We observe that our example has exactly $2m$ eigenvalues at 1. From Theorem 3.1, when $C = I$ there will be at least m unit eigenvalues. Our example has exactly m unit eigenvalues, Figure 1.

When C has rank $\frac{m}{2}$, then the preconditioned system $\mathcal{P}^{-1}\mathcal{A}$ has at least $\frac{3m}{2}$ unit eigenvalues, Theorem 4.1. Once again the number of unit eigenvalues for our example is exactly the lower bound given by the theorem.

Now suppose that we use (full) GMRES preconditioned by our extended constraint preconditioner with $G = diag(A)$ and vary the rank of C by changing the number of 1s along the diagonal of C (all other entries are zero). Figure 2 shows that with this choice of G there is a strong correlation between the upper bound on the Krylov subspace dimension and the number of iterations required to reduce the residual by at least a factor of 10^{-12}. This is an extreme example

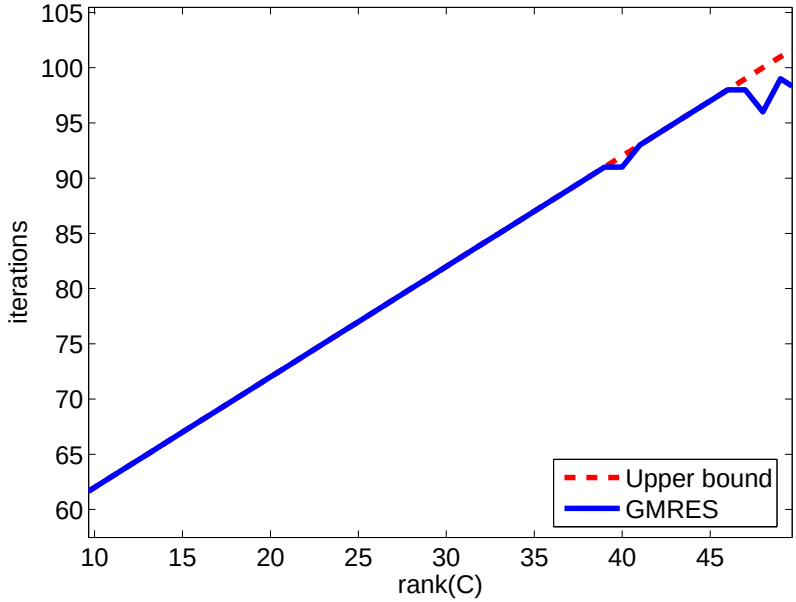

Figure 2: Comparison of upper bound on the Krylov subspace dimension and the number of iterations required to reduce the residual by 10^{-12}.

and, as we will see in the following results, the number of iterations is often a lot lower than the upper bound on the Krylov subspace dimension.

Let us compare five different approaches for solving problems of the form (1). The matrix C is set to have rank $\lceil m/2 \rceil$ and to be diagonal with just entries of 0 and 1, as above. The indefiniteness of the matrix suggests the use of MINRES; we shall use the unpreconditioned version, although positive definite preconditioning could be employed, see [23]. We note that unpreconditioned MINRES is equivalent (in exact arithmetic) to unpreconditioned GMRES for these examples because $\mathcal{A}$ is symmetric. The next three methods apply a standard (full) GMRES method with preconditioners of form $\mathcal{P}$ as considered in Section 4: in the first case we shall set $G = I$ (denoted in our table by GMRES(I)), in the second $G = diag(A)$ (denoted by GMRES(D)), and in the third case

$$G = \begin{bmatrix} 0 & 0 \\ 0 & A_{m+1:n,m+1:n} \end{bmatrix} \tag{20}$$

which we denote by GMRES($\widetilde{A}$). Finally, we shall apply the projected preconditioned conjugate gradient method (variant 2) of Dollar, Gould, Schilders and Wathen [7] with $G = diag(A)$ (this we shall denote by PPCG(D)), and G as in

Problem	m	n	MINRES	GMRES(I)	GMRES(D)	GMRES($\widetilde{A}$)	PPCG(D)	PPCG($\widetilde{A}$)
CVXQP1_M	500	1000	—	547	251	396	95	90
CVXQP2_M	750	1000	—	623	240	192	82	31
GOULDQP2_S	349	699	108	23	20	76	10	1
KSIP	1001	1021	41	9	1	13	1	1
MOSARQP1	700	3200	147	57	10	30	8	3

Table 1: Comparison of different Krylov subspace methods and preconditioners for some of the CUTEr test problems

(20) which we denote by PPCG($\widetilde{A}$). Dollar, Gould, Schilders and Wathen show that the PPCG($\widetilde{A}$) method will terminate (with exact arithmetic) in at most $\min\{2m, n-m+\lceil \frac{m}{2} \rceil\}+1$ iterations.

The saddle point systems are all preprocessed such that the first m columns of B are linearly independent. This is achieved by carrying out a column permutation to B (and also appropriately applied to A and the right hand side): this permutation is obtained by finding a sparse LU factorization of B^T using the built-in MATLAB® function `lu`. This preprocessing step is necessary for the preconditioners of GMRES($\widetilde{A}$) and PPCG($\widetilde{A}$) to be nonsingular (the resulting $A_{m+1:n,m+1:n}$ is also assumed to be nonsingular). We make no claim that this is the definitive method for finding a permutation such that the first m columns of the resulting matrix B are linearly independent and recognize that this would be an inefficient method when the dimensions of the saddle point problems are large. Further discussions on such permutations can be found in [6, Chapter 8] and is a source of future work.

In Table 1 we give the iteration counts for a small subset of CUTEr problems. As in the previous example, the matrix C is set to have rank $\lceil m/2 \rceil$ and to be diagonal with just entries of 0 and 1, the right hand sides have been set to be equal to the sum of the columns of $\mathcal{A}$ and a tolerance of 10^{-12} is used. If no iteration count is given, then this indicates that the method failed to terminate to the desired tolerance within $n+m$ iterations. The numerical results suggest that the inclusion of the (1,2), (2,1) and (2,2) blocks of $\mathcal{A}$ into the preconditioner, together with $G = diag(A)$ or G as in (20), results in a considerable reduction in the number of iterations. In particular, the projected preconditioned conjugate gradient methods appear to be favorable because they also have lower storage requirements than the GMRES-based methods.

In Figure 3 we compare how the number of PPCG iterations vary as the size of the entries of C vary. We set G as in (20), $C = \alpha I$ and vary $\alpha \in \mathbb{R}^+$. As α approaches 0 we will expect $2m$ of the eigenvalues to cluster around 1 and for the number of iterations to decrease [2]; this is observed in practice. For these examples the right hand sides have been set to be equal to the sum of the columns of $\mathcal{A}$ and a tolerance of 10^{-12} is used. In Figure 4 the same tests were performed but a random right hand side was used. Again we observe that as α approaches 0 the number of PPCG iterations decrease as expected.

We have not compared the execution times for the different methods. Instead of building $\mathcal{P}$ and then factoring it, as has been done in these tests, we suggest

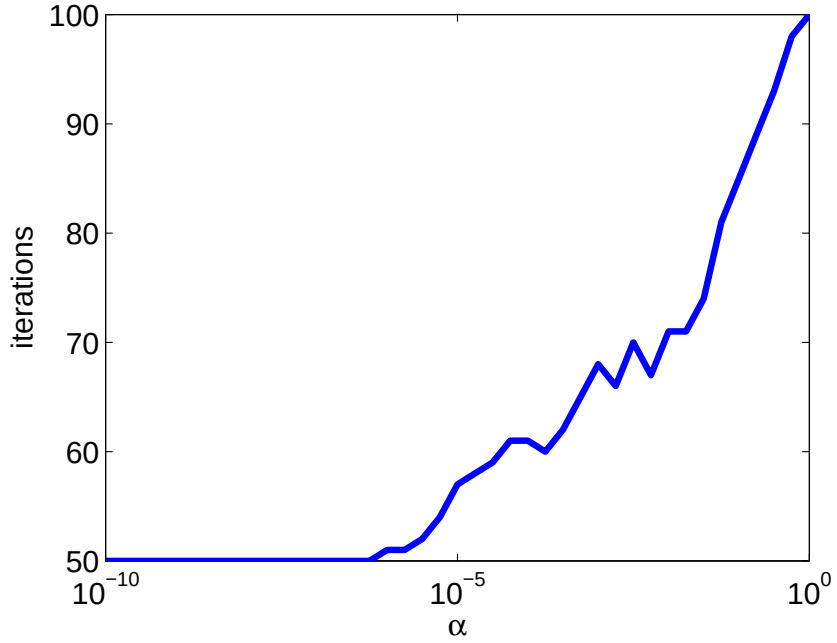

Figure 3: Comparison of the number of PPCG iterations for $C = \alpha I$ and varying α. The right hand sides have been set to be equal to the sum of the columns of $\mathcal{A}$.

the possible use of implicit-factorization constraint preconditioners which only require small factorizations to be carried out [7, 8, 9]. Dollar, Gould, Schilders and Wathen also consider how G might be chosen to further increase the number of eigenvalues at 1 [7].

6 Conclusions

In this paper, we investigated a class of preconditioners for regularized saddle point matrix systems that incorporate the (1,2), (2,1) and (2,2) blocks of the original matrix. We showed that the inclusion of these blocks in the preconditioner clusters at least $2m - p$ eigenvalues at 1, regardless of the structure of G. However, the standard convergence theory for Krylov subspace methods is not readily applicable because, in general, $\mathcal{P}^{-1}\mathcal{A}$ does not have a complete set of linearly independent eigenvectors. Using a minimum polynomial argument, we found a general (sharp) upper bound on the number of iterations required to solve linear systems of the form (1).

To confirm the analytical results of this paper we used a subset of problems

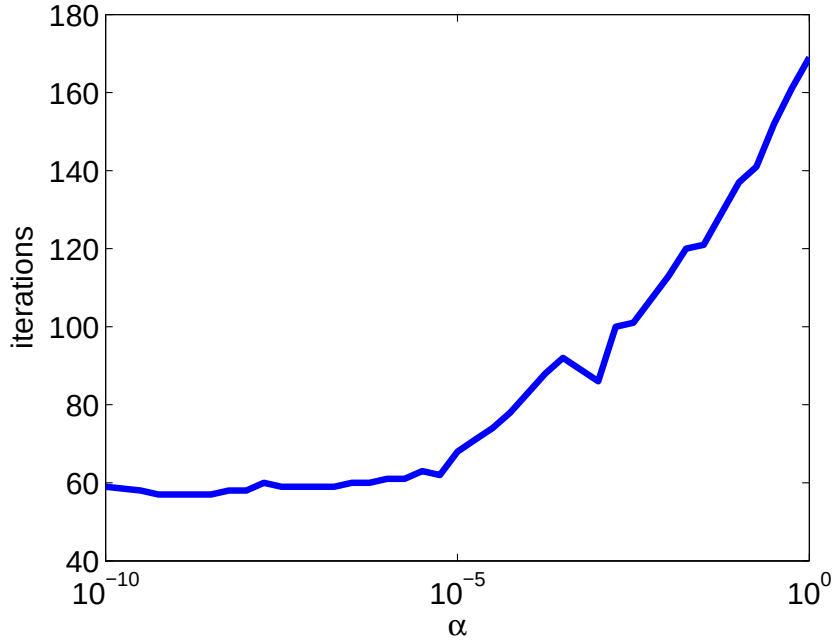

Figure 4: Comparison of the number of PPCG iterations for $C = \alpha I$ and varying α. The right hand side is a random vector.

from the CUTEr test set. We firstly used the CVXQP1_S problem and varied the rank of C to confirm the lower bound on the number of unit eigenvalues and the upper bound on the Krylov subspace dimension. We also compared MINRES for the unpreconditioned matrix system with the GMRES and PPCG methods where the preconditioner incorporate the (1,2), (2,1) and (2,2) blocks of the original matrix. We observed that the preconditioned methods resulted in a considerable reduction in the number of iterations required to reach our desired tolerance. Since GMRES and PPCG minimize different quantities, the number of iterations required may vary although the same preconditioner is used; indeed, we observe this in our results. We also confirmed that as the entries of C approach zero the number of PPCG iterations will decrease because of the additional clustering of eigenvalues around.

We have assumed that the sub-matrices B, B^T and $-C$ in (1) are exactly reproduced in the preconditioner. For truly large-scale problems this will be unrealistic [4, 5] but the theorems in this paper may still be of some interest in the inexact setting as a guide for choosing preconditioners. We wish to investigate this possibility in our future work.

Acknowledgements

The author would like to thank Nick Gould, Wil Schilders, Andy Wathen and the referees for their helpful input during the process of the work.

References

[1] A. Altman and J. Gondzio, *Regularized symmetric indefinite systems in interior point methods for linear and quadratic optimization*, Optim. Methods Softw., 11/12 (1999), pp. 275–302. Interior point methods.

[2] O. Axelsson and V. A. Barker, *Finite element solution of boundary value problems. Theory and computation*, vol. 35 of Classics in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001. Reprint of the 1984 original.

[3] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numerica, 14 (2005), pp. 1–137.

[4] L. Bergamaschi, J. Gondzio, M. Venturin, and G. Zilli, *Inexact constraint preconditioners for linear systems arising in interior point methods*, Tech. Report MS-005-002, University of Edinburgh, 2005.

[5] G. Biros and O. Ghattas, *A Lagrange-Newton-Krylov-Schur method for PDE-constrained optimization*, SIAG/Optimization Views-and-News, 11 (2000), pp. 12–18.

[6] H. S. Dollar, *Iterative Linear Algebra for Constrained Optimization*, Doctor of Philosophy, Oxford University, 2005.

[7] H. S. Dollar, N. I. M. Gould, W. H. A. Schilders, and A. J. Wathen, *Implicit-factorization preconditioning and iterative solvers for regularized saddle-point systems*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 170–189.

[8] H. S. Dollar, N. I. M. Gould, and A. J. Wathen, *On implicit-factorization constraint preconditioners*, in Large-Scale Nonlinear Optimization, Gianni G. Di Pillo and M. Roma, eds., vol. 83 of Nonconvex Optimization and Its Applications, Heidelberg, Berlin, New York, 2006, Springer Verlag, pp. 61–82.

[9] H. S. Dollar and A. J. Wathen, *Approximate factorization constraint preconditioners for saddle-point matrices*, SIAM J. Sci. Comput., 27 (2006), pp. 1555–1572.

[10] H. C. Elman, D. J. Silvester, and A. J. Wathen, *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*, Oxford University Press, Oxford, 2005.

[11] N. I. M. Gould, *Iterative methods for ill-conditioned linear systems from optimization*, in Nonlinear Optimization and Related Topics, G. DiPillo and F. Giannessi, eds., Dordrecht, The Netherlands, 1999, Kluwer Academic Publishers, pp. 123–142.

[12] N. I. M. Gould, M. E. Hribar, and J. Nocedal, *On the solution of equality constrained quadratic programming problems arising in optimization*, SIAM J. Sci. Comput., 23 (2001), pp. 1376–1395.

[13] N. I. M. Gould, D. Orban, and P. L. Toint, *CUTEr (and SifDec), a constrained and unconstrained testing environment, revisited*, ACM Transactions on Mathematical Software, 29 (2003).

[14] A. Greenbaum, *Iterative methods for solving linear systems*, vol. 17 of Frontiers in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.

[15] C. Keller, N. I. M. Gould, and A. J. Wathen, *Constraint preconditioning for indefinite linear systems*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1300–1317.

[16] A. Klawonn, *An optimal preconditioner for a class of saddle point problems with a penalty term*, SIAM J. Sci. Comput., 19 (1998), pp. 540–552.

[17] 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.

[18] I. Perugia and V. Simoncini, *Block-diagonal and indefinite symmetric preconditioners for mixed finite element formulations*, Numer. Linear Algebra Appl., 7 (2000), pp. 585–616.

[19] A. Quarteroni and A. Valli, *Numerical approximation of partial differential equations*, vol. 23 of Springer Series in Computational Mathematics, Springer-Verlag, Berlin, 1994.

[20] Y. Saad, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, second ed., 2003.

[21] 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.

[22] C. Siefert and E. de Sturler, *Preconditioners for generalized saddle-point problems*, Tech. Report UIUCDCS-R-2004-2448, Department of Computer Science, University of Illinois at Urbana-Champaign, 2004. (To appear in SIAM J. Num. Anal.).

[23] D. J. Silvester and A. J. Wathen, *Fast iterative solution of stabilised Stokes systems. II. Using general block preconditioners*, SIAM J. Numer. Anal., 31 (1994), pp. 1352–1367.

[24] H. A. VAN DER VORST, *Iterative Krylov Methods for Large Linear Systems*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, United Kingdom, 1 ed., 2003.

[25] R. J. VANDERBEI, *Symmetric quasidefinite matrices*, SIAM J. Optim., 5 (1995), pp. 100–113.

[26] S. J. WRIGHT, *Primal-dual interior-point methods*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.