

PRECONDITIONING SADDLE-POINT SYSTEMS WITH APPLICATIONS IN OPTIMIZATION*

H. SUE DOLLAR[†], NICHOLAS I. M. GOULD[†], MARTIN STOLL[‡], AND
ANDREW J. WATHEN[‡]

Abstract. Saddle-point systems arise in many applications areas, in fact in any situation where an extremum principle arises with constraints. The Stokes problem describing slow viscous flow of an incompressible fluid is a classic example coming from PDEs and in the area of optimization such problems are ubiquitous. In this paper we present a framework into which many well-known methods for solving saddle-point systems fit. Based on this description we show how new approaches for the solution of saddle-point systems arising in optimization can be derived from the Bramble–Pasciak conjugate gradient approach widely used in PDEs and more recent generalizations thereof. In particular we derive a class of new solution methods based on the use of preconditioned conjugate gradients in nonstandard inner products and demonstrate how these can be understood through more standard machinery. We show connections to constraint preconditioning and give the results of numerical computations on a number of standard optimization test examples.

Key words. saddle-point problems, iterative solver, preconditioning, optimization

AMS subject classifications. 65F10, 65F50

DOI. 10.1137/080727129

1. Introduction. Saddle-point systems arise in many application areas, in fact in any situation where an extremum principle arises with constraints. The Stokes problem describing slow viscous flow of an incompressible fluid is a classic example coming from PDEs; here minimization of viscous energy (or flow induced by a body force or boundary forcing) is constrained by conservation of mass (see, for example, [8]). In the area of optimization, such problems of finding optima in the presence of constraints are ubiquitous.

We consider saddle-point problems of the general form

$$(1) \quad Kz \equiv \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ g \end{bmatrix} \equiv d,$$

where $A \in \mathbb{R}^{n \times n}$ and $C \in \mathbb{R}^{m \times m}$ are symmetric matrices and $B \in \mathbb{R}^{m \times n}$. We assume that K is nonsingular, sufficient conditions for which are that B is of full rank, C is positive semidefinite, and A is positive definite on the kernel of B (see [3]). In practice, the properties for the blocks A and C usually vary with the underlying application. In the context of PDEs and mixed finite element methods, we can usually assume that A is positive definite and C is positive semidefinite, whereas in many problems arising in optimization, A can be indefinite (for example, when directions of negative curvature arise in sequential quadratic programming). We refer the reader to [3] for information about further applications that result in saddle-point problems of the form (1). The notation used in this paper is reasonably standard, but note that in

*Received by the editors June 13, 2008; accepted for publication (in revised form) May 6, 2009; published electronically February 5, 2010.

<http://www.siam.org/journals/sisc/32-1/72712.html>

[†]Computational Science and Engineering Department, Rutherford Appleton Laboratory, Chilton, Oxon OX11 0QX, UK (sue.dollar@stfc.ac.uk, nick.gould@stfc.ac.uk).

[‡]Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK (martin.stoll@comlab.ox.ac.uk, andy.wathen@comlab.ox.ac.uk).

line with the possibility of indefinite blocks, $\lambda_{\min}(M)$ denotes the leftmost eigenvalue and $\lambda_{\max}(M)$ the rightmost eigenvalue of a symmetric matrix M .

There are many methods for solving saddle-point problems (see [3] for a survey). In this paper we give a formulation that represents a framework for many solvers, some already known and some new. In particular, we introduce a new preconditioning strategy based on that of Bramble and Pasciak [5] and a variant of a strategy recently introduced by Forsgren, Gill, and Griffin [11] which extends the idea of constraint preconditioning [21]. An important feature of our reformulation is that the various strategies can be used to formulate preconditioned conjugate gradient (PCG) methods in nonstandard inner products (see [35]). However, we also show how such methods can be thought of as acting in the standard ℓ_2 inner product with different preconditioners.

We derive a method similar to the Bramble–Pasciak method, which is very similar to the method of Forsgren, Gill, and Griffin in the case that C is positive definite but is well defined and an effective method even if C is semidefinite (including the extreme case that it is zero). We present the results of numerical computations on optimization examples from the CUTer test set [16].

2. Reformulation. It follows directly that any solution z to (1) also satisfies

$$(2) \quad \left(\sigma \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} + \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} D & F^T \\ F & E \end{bmatrix} \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \right) \begin{bmatrix} x \\ y \end{bmatrix} \\ = \sigma \begin{bmatrix} b \\ g \end{bmatrix} + \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} D & F^T \\ F & E \end{bmatrix} \begin{bmatrix} b \\ g \end{bmatrix}$$

for arbitrary σ , symmetric matrices $D \in \mathbb{R}^{n \times n}$ and $E \in \mathbb{R}^{m \times m}$, and a matrix $F \in \mathbb{R}^{m \times n}$. We will be especially interested in the case where σ , D , E , and F ensure that the coefficient matrix in (2) is nonsingular. We denote the coefficient matrix and right-hand side of (2) as $K(\sigma, D, E, F)$ and $d(\sigma, D, E, F)$, respectively, and note that $K = K(1, 0, 0, 0)$ and $d = d(1, 0, 0, 0)$. Many well-known methods can be represented using this reformulation. For example,

- $K(0, I, I, 0)$ gives the normal equations for (1);
- $K(-1, A^{-1}, 0, 0)$ gives the Schur-complement method for finding y when A is nonsingular;
- $K(0, A^{-1}, C^{-1}, 0)$ gives the primal-dual Schur complement method for finding x and y simultaneously when both A and C are nonsingular; and
- $K(1, 0, (1 + \nu)C^{-1}, 0)$ for a given ν (in particular $\nu = 1$) gives the system to which Forsgren, Gill, and Griffin [11] apply the PCG method. The matrices C and $A + B^T C^{-1} B$ are assumed to be positive definite.

There are also a variety of methods that solve (1) by applying the conjugate gradient (CG) method within a nonstandard inner product. The general framework (2) may also be used to represent such methods:

- PCG applied to $K(-1, A_0^{-1}, 0, 0)$ with preconditioner

$$R = \begin{bmatrix} A - A_0 & 0 \\ 0 & C_0 \end{bmatrix}$$

gives the well-respected Bramble–Pasciak configuration for a given A_0 and C_0 ; see [5]. The matrices A , C_0 , and $BA^{-1}B^T + C$ are assumed to be positive

definite, and A_0 is such that $A - A_0$ and $A_0^{-1} - A^{-1}$ are also symmetric and positive definite.

- PCG applied to $K(-\gamma, I, -I, 0)$ with preconditioner

$$R = \begin{bmatrix} A - \gamma I & B^T \\ B & \gamma I - C \end{bmatrix}$$

gives Liesen and Parlett's method [23] for a given γ . The matrix A is assumed to be positive definite, γ lies in the interval $[\lambda_{\max}(C), \lambda_{\min}(A)]$, and $\|(\gamma I - C)^{-\frac{1}{2}} B (A - \gamma I)^{-\frac{1}{2}}\| < 1$. This method extends that of Benzi and Simoncini [4] to the case where $C \neq 0$.

- PCG applied to $K(-(\alpha + \beta\gamma), \alpha A_0^{-1} + \beta I, -\beta I, 0)$ with preconditioner

$$R = \begin{bmatrix} A - (\alpha + \beta\gamma)(\alpha A_0^{-1} + \beta I)^{-1} & B^T \\ B & \frac{\alpha + \beta\gamma}{\beta} I \end{bmatrix}$$

gives an example of the combination preconditioning method of Stoll and Wathen [35]. The assumptions of both Bramble–Pasciak and Liesen and Parlett must hold, and α , β , and γ must be chosen such that $K(-(\alpha + \beta\gamma), \alpha A_0^{-1} + \beta I, -\beta I, 0) = \alpha K(-1, A_0^{-1}, 0, 0) + \beta K(-\gamma, I, -I, 0)$ is positive definite.

- PCG applied to $K(1, A_0^{-1}(B^T C_0^{-1} B - A_0^{-1}) A_0^{-1}, C_0^{-1}, -C_0^{-1} B A_0^{-1})$ with preconditioner

$$R = \begin{bmatrix} A_0 - A & 0 \\ 0 & B^T A_0^{-1} B - C_0 \end{bmatrix}$$

represents the method presented by Schöberl and Zulehner for the case $C = 0$ [32] for given A_0 and C_0 . The matrix A is assumed to be positive definite on the kernel of B , A_0 is such that $A_0 - A$ is symmetric and positive definite, and C_0 is such that $B A_0^{-1} B^T - C_0$ is symmetric and positive definite.

It may not be obvious why the above formulations produce algorithms that (in exact arithmetic) produce iterates which are equivalent to those produced by the CG methods within a nonstandard inner product. In section 3, we will reveal why we can reformulate the methods as above. We intend that these reformulations will provide the reader with an insight into the properties of the nonstandard inner-product CG methods without having to use the nonstandard inner-product. Of course, simply reformulating (1) as (2) in itself offers no immediate advantage. However, if $K(\sigma, D, E, F)$ possesses one or more desirable properties, (2) may be preferable to (1).

Before considering the nonstandard inner-product CG methods, we will consider what properties need to hold to guarantee that $K(\sigma, D, E, F)$ is symmetric and positive definite (and thus one may use methods such as CG rather than MINRES). Clearly, D and E need both be symmetric. Furthermore, we may factorize $K(\sigma, D, E, F)$ as

$$\begin{aligned} K(\sigma, D, E, F) &= \begin{bmatrix} \Theta_1 & \Theta_2^T \\ \Theta_2 & \Theta_3 \end{bmatrix} \\ (3) \quad &= \begin{bmatrix} I & \Theta_2^T \Theta_3^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} \Theta_1 - \Theta_2^T \Theta_3^{-1} \Theta_2 & 0 \\ 0 & \Theta_3 \end{bmatrix} \begin{bmatrix} I & 0 \\ \Theta_3^{-1} \Theta_2 & I \end{bmatrix} \end{aligned}$$

or

$$(4) \quad K(\sigma, D, E, F) = \begin{bmatrix} I & 0 \\ \Theta_2 \Theta_1^{-1} & I \end{bmatrix} \begin{bmatrix} \Theta_1 & 0 \\ 0 & \Theta_3 - \Theta_2 \Theta_1^{-1} \Theta_2^T \end{bmatrix} \begin{bmatrix} I & \Theta_1^{-1} \Theta_2^T \\ 0 & I \end{bmatrix},$$

where

$$(5) \quad \Theta_1 = \sigma A + ADA + B^T F A + A F^T B + B^T E B,$$

$$(6) \quad \Theta_2 = \sigma B + BDA - CFA + B F^T B - C E B,$$

$$(7) \quad \Theta_3 = BDB^T - CFB^T - B F^T C + C E C - \sigma C.$$

Using Sylvester's law of inertia [15], we obtain the following theorem.

THEOREM 2.1. *Let Θ_1 , Θ_2 , and Θ_3 be as defined in (5)–(7). $K(\sigma, D, E, F)$ is symmetric and positive definite if and only if*

- D and E are symmetric,
- Θ_3 is positive definite, and
- $\Theta_1 - \Theta_2^T \Theta_3^{-1} \Theta_2$ is positive definite.

Equivalently, $K(\sigma, D, E, F)$ is symmetric and positive definite if and only if

- D and E are symmetric,
- Θ_1 is positive definite, and
- $\Theta_3 - \Theta_2 \Theta_1^{-1} \Theta_2^T$ is positive definite.

Clearly, $K(\sigma, D, E, F)$ is symmetric and positive definite if and only if $K^{-1}K(\sigma, D, E, F)K^{-1}$ is symmetric and positive definite. This is equivalent to requiring that

$$\sigma K^{-1} + \begin{bmatrix} D & F^T \\ F & E \end{bmatrix}$$

be symmetric and positive definite. We will consider different cases for A and C separately. Proofs for the following corollaries may be found in Appendix A.

COROLLARY 2.2. *If A is symmetric and nonsingular, and*

$$\begin{aligned} S_A &= C + BA^{-1}B^T, \\ \Upsilon_1 &= D + \sigma A^{-1} - \sigma A^{-1}B^T S_A^{-1}BA^{-1}, \\ \Upsilon_2 &= \sigma F + S_A^{-1}BA^{-1}, \\ \Upsilon_3 &= E - \sigma S_A^{-1}, \end{aligned}$$

then $K(\sigma, D, E, F)$ is symmetric and positive definite if and only if

- D and E are symmetric,
- Υ_3 is positive definite, and
- $\Upsilon_1 - \Upsilon_2^T \Upsilon_3^{-1} \Upsilon_2$ is positive definite.

COROLLARY 2.3. *If C is symmetric and nonsingular, and*

$$\begin{aligned} S_C &= A + B^T C^{-1}B, \\ \Delta_1 &= D + \sigma S_C^{-1}, \\ \Delta_2 &= F + \sigma C^{-1}BS_C^{-1}, \\ \Delta_3 &= E + \sigma C^{-1}BS_C^{-1}B^T C^{-1} - C^{-1}, \end{aligned}$$

then $K(\sigma, D, E, F)$ is symmetric and positive definite if and only if

- D and E are symmetric,
- Δ_1 is positive definite, and
- $\Delta_3 - \Delta_2 \Delta_1^{-1} \Delta_2^T$ is positive definite.

COROLLARY 2.4. If $C = 0$, the columns of $Z \in \mathbb{R}^{n \times (n-m)}$ span the nullspace of B , $B^\dagger = B^T (BB^T)^{-1}$ is the Moore–Penrose inverse of B [15], and

$$\begin{aligned} S_Z &= Z^T A Z, \\ \Gamma_1 &= D + \sigma Z S_Z^{-1} Z^T, \\ \Gamma_2 &= F + \sigma B^{\dagger T} (I - A Z S_Z^{-1} Z^T), \\ \Gamma_3 &= E + \sigma B^{\dagger T} (A Z S_Z Z^T A - A) B^\dagger. \end{aligned}$$

$K(\sigma, D, E, F)$ is symmetric and positive definite if and only if

- D and E are symmetric,
- Γ_1 is positive definite, and
- $\Gamma_3 - \Gamma_2 \Gamma_1^{-1} \Gamma_2^T$ is positive definite.

Conditions for the case where C is rank-deficient but nonzero may be derived by factoring C as

$$C = U^T \begin{bmatrix} \widehat{C} & 0 \\ 0 & 0 \end{bmatrix} U,$$

where \widehat{C} is nonsingular and U is unitary. Premultiplying K by $\begin{bmatrix} I & 0 \\ 0 & U \end{bmatrix}$ and postmultiplying by the inverse of this matrix reveals a saddle-point system to which either Corollary 2.3 or Corollary 2.4 could be applied.

3. Equivalence of nonstandard inner-product CG methods and standard PCG methods. In this section, we illustrate the equivalence of the reformulation and a class of methods that apply CG with a nonstandard inner-product. Such examples can be found in [4, 5, 10, 23, 32, 35]. These methods all have a common framework. That is, matrices P and H are formed such that P is nonsingular, H is symmetric and positive definite, and $P^{-1}K$ is self-adjoint in the inner-product $\langle \cdot, \cdot \rangle$ defined by $\langle x, y \rangle_H = x^T H y$, i.e., $HP^{-1}K$ is symmetric. Moreover, P and H are chosen such that $P^{-1}K$ is positive definite in the inner-product $\langle \cdot, \cdot \rangle_H$. The methods then use this nonstandard inner product within the CG method, as illustrated in Algorithm 3.1. At iteration k of Algorithm 3.1, $\text{span} \{p^{(0)}, p^{(1)}, \dots, p^{(k-1)}\} = \text{span} \{r^{(0)}, r^{(1)}, \dots, r^{(k-1)}\}$, $r^{(k)T} H r^{(j)} = 0$, and $p^{(k)T} H P^{-1} K p^{(j)} = 0$ for all $j < k$; see [23, Theorem 3.2]. Hence, Algorithm 3.1 may be reformulated as Algorithm 3.2.

We observe that $P^{-1}K$ is self-adjoint and positive definite in the inner-product $\langle \cdot, \cdot \rangle_H$ if and only if $HP^{-1}K$ is symmetric and positive definite. Thus, an alternative method for solving (1) is to apply the PCG method (with a preconditioner L) to the equivalent symmetric and positive definite system

$$(8) \quad HP^{-1}Kz = HP^{-1}d.$$

Such a method is given by Algorithm 3.3. Eliminating $s^{(k)}$ from Algorithm 3.3 we obtain Algorithm 3.4.

Observe that when $L = H$, Algorithms 3.2 and 3.4 are equivalent (in exact arithmetic). Thus, application of a nonstandard inner-product CG method with matrices

ALGORITHM 3.1. *Nonstandard inner-product CG (variant 1).*

Given $z^{(0)} = 0$, set $r^{(0)} = P^{-1}(d - Kz^{(0)})$
 and $p^{(0)} = r^{(0)}$
for $k = 0, 1, \dots$ **do**
 $\alpha = \frac{\langle r^{(k)}, p^{(k)} \rangle_H}{\langle P^{-1}Kp^{(k)}, p^{(k)} \rangle_H}$
 $z^{(k+1)} = z^{(k)} + \alpha p^{(k)}$
 $r^{(k+1)} = r^{(k)} - \alpha P^{-1}Kp^{(k)}$
 $\beta = \frac{\langle P^{-1}Kr^{(k+1)}, p^{(k)} \rangle_H}{\langle P^{-1}Kp^{(k)}, p^{(k)} \rangle_H}$
 $p^{(k+1)} = r^{(k+1)} - \beta p^{(k)}$
end for

ALGORITHM 3.2. *Nonstandard inner-product CG (variant 2).*

Given $z^{(0)} = 0$, set $r^{(0)} = P^{-1}(d - Kz^{(0)})$
 and $p^{(0)} = r^{(0)}$
for $k = 0, 1, \dots$ **do**
 $\alpha = \frac{r^{(k)T} H r^{(k)}}{p^{(k)T} H P^{-1} K p^{(k)}}$
 $z^{(k+1)} = z^{(k)} + \alpha p^{(k)}$
 $r^{(k+1)} = r^{(k)} - \alpha P^{-1} K p^{(k)}$
 $\beta = \frac{r^{(k+1)T} H r^{(k+1)}}{r^{(k)T} H r^{(k)}}$
 $p^{(k+1)} = r^{(k+1)} + \beta p^{(k)}$
end for

ALGORITHM 3.3. *Preconditioned CG method for solving $HP^{-1}Kx = HP^{-1}b$ with symmetric and positive definite preconditioner L (variant 1).*

Given $z^{(0)} = 0$, set $s^{(0)} = HP^{-1}(d - Kz^{(0)})$
 Solve $Lq^{(0)} = s^{(0)}$ and set $p^{(0)} = q^{(0)}$
for $k = 0, 1, \dots$ **do**
 $\alpha = \frac{s^{(k)T} q^{(k)}}{p^{(k)T} H P^{-1} K p^{(k)}}$
 $z^{(k+1)} = z^{(k)} + \alpha p^{(k)}$
 $s^{(k+1)} = s^{(k)} - \alpha HP^{-1}Kp^{(k)}$
 Solve $Lq^{(k+1)} = s^{(k+1)}$
 $\beta = \frac{s^{(k+1)T} q^{(k+1)}}{s^{(k)T} q^{(k)}}$
 $p^{(k+1)} = z^{(k+1)} + \beta p^{(k)}$
end for

ALGORITHM 3.4. *Preconditioned CG method for solving $HP^{-1}Kx = HP^{-1}b$ with symmetric and positive definite preconditioner L (variant 2).*

Given $z^{(0)} = 0$, set $q^{(0)} = L^{-1}HP^{-1}(d - Kz^{(0)})$ and $p^{(0)} = q^{(0)}$
for $k = 0, 1, \dots$ **do**
 $\alpha = \frac{q^{(k)T} L q^{(k)}}{p^{(k)T} H P^{-1} K p^{(k)}}$
 $z^{(k+1)} = z^{(k)} + \alpha p^{(k)}$
 $q^{(k+1)} = q^{(k)} - \alpha L^{-1}HP^{-1}Kp^{(k)}$
 $\beta = \frac{q^{(k+1)T} L q^{(k+1)}}{q^{(k)T} L q^{(k)}}$
 $p^{(k+1)} = q^{(k+1)} + \beta p^{(k)}$
end for

H and P produces iterates that are equivalent to those formed by applying the standard PCG method to the symmetric and positive definite problem (8) with symmetric and positive definite preconditioner H . The convergence of Algorithms 3.1 and 3.2 can, hence, be described by the eigenvalues of $P^{-1}K$, and that of Algorithms 3.3 and 3.4 are described by the eigenvalues of $L^{-1}HP^{-1}K$. For further details about the connection between CG in the nonstandard inner-product and in the standard inner product, we refer the reader to [1, 2, 20].

We note that, in our framework (2), the matrix HP^{-1} in (8) corresponds to

$$\sigma I + \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} D & F^T \\ F & E \end{bmatrix}.$$

We will illustrate the above observations by considering the Bramble–Pasciak CG method. This method assumes that A is symmetric and positive definite and sets

$$H = \begin{bmatrix} A - A_0 & 0 \\ 0 & S_0 \end{bmatrix} \quad \text{and} \quad P = \begin{bmatrix} A_0 & 0 \\ B & -S_0 \end{bmatrix},$$

where A_0 is an approximation to A , $A - A_0$ is symmetric and positive definite, and S_0 is a symmetric and positive definite approximation to $C + BA^{-1}B^T$. We note that Bramble and Pasciak [5] only considered $S_0 = I$ but this was extended to more general

S_0 in [22, 24, 33]. Setting $\sigma = -1$, $D = A_0^{-1}$, $E = 0$, and $F = 0$, we can confirm that

$$HP^{-1} = \sigma I + \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} D & F^T \\ F & E \end{bmatrix}.$$

The entries in the matrix $HP^{-1}K$ are independent of the choice of S_0 .

Corollary 2.2 implies that $HP^{-1}K$ will be symmetric and positive definite if and only if both $BA^{-1}B^T + C$ and $A_0^{-1} - A^{-1}$ are positive definite. We note that these are exactly the same conditions as those derived by Klawonn [22].

Finally, in Figure 1, we plot the convergence history of the Bramble–Pasciak CG method and the standard PCG method without preconditioning CG and with preconditioner $L = H$ PCG when applied to a Stokes problem of dimension 59 that was generated by IFISS [7]. We set $A_0 = 0.5A$ and $S_0 = I$. The matrix P was constructed such that $P^{-1}K$ has good convergence properties but we have no reason to expect that $HP^{-1}K$ will also have good convergence properties. For this example, we would therefore expect that the Bramble–Pasciak and PCG methods will outperform the standard CG method: Figure 1 confirms our prediction. As expected, when the preconditioner $L = H$ is used within PCG the convergence curve is almost identical to that of the Bramble–Pasciak CG method (the slight deviation is due to round-off error). Of course, different choices of H and P could result in the standard CG method outperforming the nonstandard inner-product CG method: the impractical choice $H = K^2$ and $P = K^{-3}$ is one such example.




FIG. 1. *Bramble–Pasciak CG* *unpreconditioned CG for* $HP^{-1}Kz = HP^{-1}d$, *and PCG for* $HP^{-1}Kz = HP^{-1}d$ *with preconditioner* H *for a Stokes problem of dimension 59.*

4. Using the reformulation. In sections 2 and 3, we illustrated that different methods for solving saddle-point problems can be presented within the same framework; see (2). Furthermore, we showed that nonstandard inner-product CG methods for solving saddle-point systems can be reformulated as standard PCG methods. In this section, we will review the properties of the Forsgren, Gill, and Griffin (FGG)

method and derive a Bramble–Pasciak-style method that may have similar convergence properties. Our new method allows us to relax the assumption that C is symmetric and positive definite, which is required in the FGG approach.

4.1. The method of Forsgren, Griffin, and Gill (FGG). Forsgren, Gill, and Griffin [11] worked with a saddle-point problem of the general form

$$(9) \quad K(\nu) \begin{bmatrix} x \\ y \end{bmatrix} \equiv \begin{bmatrix} A + (1 + \nu)B^T C^{-1}B & -\nu B^T \\ -\nu B & \nu C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b + (1 + \nu)B^T C^{-1}g \\ -\nu g \end{bmatrix},$$

where $\nu \in \mathbb{R}$ and, as we have already observed, $K(\nu) \equiv K(1, 0, (1 + \nu)C^{-1}, 0)$ in our general setting. We emphasize the fact that C must be nonsingular in this formulation. The case $\nu = -1$ gives the standard saddle-point formulation, $\nu = 0$ a condensed system which is equivalent to the Schur-complement method for finding the solution, and $\nu = 1$ the doubly augmented system with matrix

$$(10) \quad K(1) = \begin{bmatrix} A + 2B^T C^{-1}B & -B^T \\ -B & C \end{bmatrix}.$$

Note that $K(\nu)$ is positive definite if $A + B^T C^{-1}B > 0$, $C > 0$, and $\nu > 0$. In addition, a general preconditioner

$$(11) \quad M(\nu) = \begin{bmatrix} G + (1 + \nu)B^T C^{-1}B & -\nu B^T \\ -\nu B & \nu C \end{bmatrix}$$

is introduced, where G is an approximation to A and $G + B^T C^{-1}B > 0$. Again, $M(\nu)$ represents different preconditioners for different instances of ν . In practice it is often useful to use the decomposition

$$(12) \quad \begin{bmatrix} G + (1 + \nu)B^T C^{-1}B & -\nu B^T \\ -\nu B & \nu C \end{bmatrix} = \begin{bmatrix} I & (1 + \nu)B^T C^{-1} \\ 0 & -\nu I \end{bmatrix} \begin{bmatrix} G & B^T \\ B & -C \end{bmatrix}$$

to solve a system with the preconditioner $M(\nu)$. Note that the first factor here is easily inverted and the second factor is in the form of a constraint preconditioner [21]. The eigenvalues of the preconditioned system $M(\nu)^{-1}K(\nu)$ are independent of ν and are given by the eigenvalues of

$$(13) \quad (G + B^T C^{-1}B)^{-1}(A + B^T C^{-1}B)$$

together with m unit eigenvalues. Therefore, in exact arithmetic, convergence is given in at most $n + 1$ steps.

4.2. A Bramble–Pasciak-like approach. In this section, we show the equivalence of the FGG method and a Bramble–Pasciak-like method. Multiplying K by -1 and block-symmetrically permuting its rows (and columns) we obtain a matrix of the form

$$\overline{K} = \begin{bmatrix} C & -B \\ -B^T & -A \end{bmatrix}.$$

Applying Bramble–Pasciak to this matrix we obtain

$$\overline{H} = \begin{bmatrix} C - C_0 & 0 \\ 0 & A_0 \end{bmatrix} \quad \text{and} \quad \overline{P} = \begin{bmatrix} C_0 & 0 \\ -B^T & -A_0 \end{bmatrix},$$

where A_0 and C_0 are symmetric and nonsingular. Removing the permutation and multiplication we obtain the preconditioner

$$(14) \quad P = \begin{bmatrix} A_0 & B^T \\ 0 & -C_0 \end{bmatrix} \quad \text{with} \quad P^{-1} = \begin{bmatrix} A_0^{-1} & A_0^{-1}B^TC_0^{-1} \\ 0 & -C_0^{-1} \end{bmatrix}$$

and the bilinear form matrix

$$(15) \quad H = \begin{bmatrix} A_0 & 0 \\ 0 & C - C_0 \end{bmatrix}.$$

It is easy to show that $HP^{-1}K$ is symmetric (i.e., $P^{-1}K$ self-adjoint in the bilinear form given by H) and that

$$HP^{-1}K = K(1, 0, C_0^{-1}, 0).$$

Thus, if C is positive definite and $C_0 = \frac{1}{1+\nu}C$, we obtain the FGG reformulation (see section 2).

We stress the fact that the FGG method assumes that the matrix C is nonsingular, whereas this assumption may be relaxed in (14)–(15). Instead, we require that C_0 be a nonsingular matrix which is chosen in such a manner that $C - C_0$ is nonsingular: clearly we can always find such a C_0 for any choice of C .

Typically, the properties of A and C depend on the underlying application. We will now discuss some of the common cases. Since, as we noted in section 3, the iteration is applied implicitly to a system with matrix $P^{-1}K$, and as the eigenvalues of such a matrix influence convergence, we analyze the eigenvalues of the matrix $P^{-1}K$ for various choices of A_0 and C_0 . If A_0 and C_0 can be chosen such that H and $HP^{-1}K$ are positive definite, then this would enable us to solve (1) by applying PCG (with preconditioner H) to the equivalent system (8). Using H as the preconditioner enables us to rewrite the method in the style of a nonstandard inner-product CG method (Algorithms 3.1 and 3.2 in section 3).

If H is symmetric and positive definite, but $HP^{-1}K$ is symmetric but not positive definite, then we cannot reliably apply the CG method. However, we may solve (8) by applying H -MINRES [34,35] based on the fact that the preconditioned matrix $P^{-1}H$ is self-adjoint in the inner product defined by H . Note that H always defines an inner product if we set $C_0 := -C_0$ for positive definite C_0 in the above setup. Another possibility is to use the simplified QMR (sQMR) method of Freund [12,34,35] where a simplified version of the nonsymmetric Lanczos process is used based on the identity $HP^{-1}K = KP^{-T}H$. Note that sQMR does not need the positivity of H , only the self-adjointness of $P^{-1}K$ in the bilinear form defined by H . Note that for the unsymmetric matrix $P^{-1}K$, we can also use methods such as BICG or BICGSTAB, but that these methods do not necessarily benefit from the self-adjointness of the preconditioned matrix at the cost of a more expensive implementation.

4.2.1. C positive definite. The case where C is positive definite can sometimes be found in optimization [11] (as well as other areas [3]) and usually occurs because of some explicit regularization [29]. Optimality conditions imply that $A + B^T C^{-1} B$ should be positive definite. Suppose that we set $C_0 = C$; then the eigenvalues of $P^{-1}K$ are given by the following theorem.

THEOREM 4.1. *Let*

$$K = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \quad \text{and} \quad P = \begin{bmatrix} A_0 & B^T \\ 0 & -C \end{bmatrix}$$

with nonsingular A_0 and C . Then $P^{-1}K$ has

- m eigenvalues at 1,
- that the remaining n eigenvalues are defined by the generalized eigenvalue problem

$$(A + B^T C^{-1} B) x = \lambda A_0 x.$$

Proof. See [33, Proposition 1]. \square

Note that if $A_0 = G + B^T C^{-1} B$, then $P^{-1}K$ will have the same eigenvalues as $M(\nu)^{-1}K(\nu)$, where $K(\nu)$ and $M(\nu)$ are defined by (9) and (11), respectively.

However, if $C_0 = C$, then H in (15) will be singular and, therefore, Algorithms 3.1 and 3.2 may break down. Suppose that we instead choose $C_0 = (1 + \nu)^{-1}C$, where $\nu \neq -1$. If A_0 is chosen to be a symmetric and positive definite matrix, H will be symmetric and positive definite if and only if $\nu > 0$ or $\nu < -1$. Applying Corollary 2.3 with $\sigma = 1$, $D = 0$, $E = C_0^{-1}$, and $F = 0$, we find that $HP^{-1}K$ is positive definite if and only if $A + B^T C^{-1} B$ and $C_0^{-1} - C^{-1}$ are both positive definite. If $C_0 = (1 + \nu)^{-1}C$, then $C_0^{-1} - C^{-1}$ is positive definite if and only if $\nu > 0$. This confirms the result by Forsgren, Griffin, and Gill that $K(\nu)$ is positive definite if $A + B^T C^{-1} B > 0$ and $\nu > 0$; see section 4.1. Note that C_0 has to be scaled in order to guarantee the positivity of $C - C_0$. This is a trivial task if C_0 is a multiple of C , but for general C_0 this usually involves an eigenvalue estimation problem for $C_0^{-1}C$. Theorem 4.2 provides results on the eigenvalues of the resulting matrix $P^{-1}K$.

THEOREM 4.2. *Let B have rank $r > 0$ and $Z \in \mathbb{R}^{n \times (n-r)}$ be such that its columns span the nullspace of B . Additionally, let*

$$K = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \quad \text{and} \quad P = \begin{bmatrix} A_0 & B^T \\ 0 & -(1 + \nu)^{-1}C \end{bmatrix}$$

with nonsingular A_0 and C , where $\nu \neq 0$ and $\nu \neq -1$. Suppose that the generalized eigenvalue problem $Z^T A Z x_z = \lambda Z^T A_0 Z x_z$ has j ($0 \leq j \leq n - r$) eigenvalues equal to $1 + \nu$. Then $P^{-1}K$ has

- j eigenvalues at $1 + \nu$,
- that the remaining eigenvalues satisfy the quadratic eigenvalue problem

$$(16) \quad \lambda^2 A_0 x - \lambda (A + (1 + \nu)(A_0 + B^T C^{-1} B)) x + (1 + \nu)(A + B^T C^{-1} B) x = 0$$

subject to $\lambda \neq 1 + \nu$.

Proof. Assume that $(\lambda, [x^T \ y^T]^T)$ represents an eigenpair of $P^{-1}K$. Then

$$(17) \quad Ax + B^T y = \lambda (A_0 x + B^T y),$$

$$(18) \quad Bx - Cy = -\frac{\lambda}{1 + \nu} Cy.$$

Let $\lambda = 1 + \nu$. Equation (18) implies that $Bx = 0$. Let $Z \in \mathbb{R}^{n \times (n-r)}$ be such that its columns span the nullspace of B , and let $Y \in \mathbb{R}^{n \times r}$ be such that its columns span the range of the columns of B^T . If $x = Yx_y + Zx_z$, then $Bx = 0$ implies that $x_y = 0$. Premultiplying (17) by $[Y \ Z]^T$ and substituting in $x = Zx_z$ we obtain

$$(19) \quad Y^T AZx_z + (BY)^T y = (1 + \nu) \left(Y^T A_0 Zx_z + (BY)^T y \right),$$

$$Z^T AZx_z = (1 + \nu) Z^T A_0 Zx_z.$$

Hence, $x_z \neq 0$ if and only if $1 + \nu$ is an eigenvalue of the generalized eigenvalue problem $Z^T AZx_z = \lambda Z^T A_0 Zx_z$. Given such an x_z , y can be defined using (19).

Let $\lambda \neq 1 + \nu$. Equation (18) implies that

$$y = \frac{1 + \nu}{1 + \nu - \lambda} C^{-1} Bx.$$

Substituting this into (17) and rearranging we obtain the quadratic eigenvalue problem

$$\lambda^2 A_0 x - \lambda (A + (1 + \nu) (A_0 + B^T C^{-1} B)) x + (1 + \nu) (A + B^T C^{-1} B) x = 0.$$

This completes the proof. \square

In practice, there are unlikely to be many (if any) eigenvalues equal to $1 + \nu$. However, in certain situations, the quadratic eigenvalue problem (16) will give us useful information about the distribution of the eigenvalues. Consider the case of convex nonlinear optimization problems of the form

$$\min f(x) \quad \text{such that} \quad c(x) \geq 0,$$

where $x \in \mathbb{R}^n$, and $f : \mathbb{R}^n \mapsto \mathbb{R}$ and $-c : \mathbb{R}^n \mapsto \mathbb{R}^{m_1}$ are convex and twice differentiable. The use of primal-dual interior point methods [36] for this problem will result in a sequence of saddle-point problems of the form (1) where C is diagonal and positive definite. If a partitioning method is applied to the constraints [19], then the entries of the resulting matrix C become small as the iterates of the interior point method approach optimality. In the following corollary we show that as the entries in C become small, many of the eigenvalues defined by (16) will be clustered around two points.

COROLLARY 4.3. *Assume that $A + B^T C^{-1} B$ and C are symmetric and positive definite. Let $A_0 = G + (1 + \nu) B^T C^{-1} B$ and assume that A_0 is symmetric and positive definite. Consider the quadratic eigenvalue problem (16). If $Bx \neq 0$, the eigenvalues will converge to either $\frac{2 + \nu + \sqrt{\nu^2 + 4\nu}}{2}$ or $\frac{2 + \nu - \sqrt{\nu^2 + 4\nu}}{2}$ as $\|C\| \rightarrow 0$. If $Bx = 0$ and $\lambda \neq 1 + \nu$, the eigenvalues satisfy the generalized eigenvalue problem $Ax = \lambda Gx$ (subject to $\lambda \neq 1 + \nu$).*

Proof. Consider the case $Bx \neq 0$. Let

$$\alpha = x^* A x, \quad \gamma = x^* G x, \quad \text{and} \quad \xi = x^* B^T C^{-1} B x.$$

Premultiplying (16) by x^* we obtain

$$\lambda^2 (\gamma + (1 + \nu) \xi) - \lambda (\alpha - \gamma + (2 + \nu) (\gamma + (1 + \nu) \xi)) + (1 + \nu) (\alpha + \xi).$$

Hence,

$$\begin{aligned} \lambda &= \frac{\alpha - \gamma + (2 + \nu) (\gamma + (1 + \nu) \xi) \pm \sqrt{(\alpha - \gamma + (2 + \nu) (\gamma + (1 + \nu) \xi))^2 - 4(1 + \nu) (\alpha + \xi) (\gamma + (1 + \nu) \xi)}}{2(\gamma + (1 + \nu) \xi)} \\ &= \frac{\alpha - \gamma}{\gamma + (1 + \nu) \xi} + \frac{2 + \nu}{2} \pm \sqrt{\frac{(\alpha - \gamma)^2}{4(\gamma + (1 + \nu) \xi)^2} - \frac{\nu(\alpha + \gamma)}{2(\gamma + (1 + \nu) \xi)} + \frac{\nu^2 - 4\nu}{4}}. \end{aligned}$$

Assume that $\|Bx\| = 1$. Clearly α and γ are bounded from above and below independently of C for all x satisfying $\|Bx\| = 1$. Let $z = Bx$, $c_l = \lambda_{\min}(C)$, and $c_u = \lambda_{\max}(C)$. Now $\xi = z^* C^{-1} z$ and, hence,

$$c_u^{-1} \leq \xi \leq c_l^{-1}.$$

Therefore

$$\frac{c_l}{1 + \nu + c_l \gamma_u} \leq \frac{1}{\gamma + (1 + \nu)\xi} \leq \frac{c_u}{1 + \nu + c_u \gamma_l},$$

where $\gamma_l \leq x^* G x \leq \gamma_u$, and

$$\frac{1}{\gamma + (1 + \nu)\xi} \rightarrow 0 \quad \text{as} \quad \|C\| \rightarrow 0.$$

This completes the case $Bx \neq 0$.

If $Bx = 0$, the quadratic eigenvalue problem (16) becomes

$$\lambda^2 G x - \lambda(A + (1 + \nu)G)x + (1 + \nu)Ax = 0.$$

Noting that

$$\lambda^2 G - \lambda(A + (1 + \nu)G) + (1 + \nu)A = (\lambda - 1 - \nu)(\lambda G - A),$$

we obtain the desired result. \square

Figure 2 shows the eigenvalue distribution of $P^{-1}K$ for different choices of C , where $\nu = 0.1$, $C_0 = (1 + \nu)^{-1}C$, and $A_0 = \text{diag}(A) + B^T C^{-1}B$. As before, we consider the matrix CVXQP3.S. As predicted by Corollary 4.3, as $\|C\| \rightarrow 0$, a large number of the eigenvalues converges to either $\frac{2 + \nu + \sqrt{\nu^2 + 4\nu}}{2}$ or $\frac{2 + \nu - \sqrt{\nu^2 + 4\nu}}{2}$.




FIG. 2. Eigenvalue distribution of $P^{-1}K$ for different choices of C , where $\nu = 0.1$, $C_0 = (1 + \nu)^{-1}C$, and $A_0 = \text{diag}(A) + B^T C^{-1}B$.

In the case where a factorization of $A_0 = G + B^T C_0^{-1} B$ should be avoided, it may be helpful to decompose the matrix P as

$$(20) \quad P = \begin{bmatrix} G + B^T C_0^{-1} B & B^T \\ 0 & -C_0 \end{bmatrix} = \begin{bmatrix} G & B^T \\ B & -C_0 \end{bmatrix} \begin{bmatrix} I & 0 \\ C_0^{-1} B & I \end{bmatrix}.$$

Forsgren, Griffin, and Gill recommend a similar trick for their method; see section 4.1.

4.2.2. A positive definite and C positive semidefinite. If A is positive definite, then we may let $A_0 = A$; the analysis presented here is not based on the assumption that C is positive definite. The eigenvalues of $P^{-1}K$ are defined by Theorem 4.4.

THEOREM 4.4. *Let*

$$K = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix} \quad \text{and} \quad P = \begin{bmatrix} A & B^T \\ 0 & -C_0 \end{bmatrix},$$

where A is positive definite, C_0 is symmetric and (positive or negative) definite, and $C - C_0$ is nonsingular. Then $P^{-1}K$ has

- n eigenvalues at 1,
- that the remaining m eigenvalues are defined by the generalized eigenvalue problem

$$(C + BA^{-1}B^T)y = \lambda C_0 y.$$

Proof. If $P^{-1}Kz = \lambda z$, then $KP^{-1}\tilde{z} = \lambda\tilde{z}$, where $\tilde{z} = Pz$. Now

$$KP^{-1} = \begin{bmatrix} I & 0 \\ BA^{-1} & (C + BA^{-1}B^T)C_0^{-1} \end{bmatrix}.$$

Hence, $P^{-1}K$ has n eigenvalues at 1, and the remaining m eigenvalues are defined by the generalized eigenvalue problem

$$(C + BA^{-1}B^T)y = \lambda C_0 y. \quad \square$$

As a result, the convergence of the Bramble–Pasciak-like setup with $A_0 = A$ is given in at most $m + 1$ steps. If $C + BA^{-1}B^T$ and C_0 are both positive definite, then all of the eigenvalues of $P^{-1}K$ will be positive; however, C_0 must be chosen such that $C - C_0$ is positive definite in order to guarantee that H is positive definite. If $C + BA^{-1}B^T$ is positive definite and C_0 is negative definite, then $P^{-1}K$ will have m negative eigenvalues. Hence, CG cannot be reliably applied, but with the replacement of $C_0 := -C_0$ for positive definite C_0 we can always use the H -MINRES method.

The case of A definite and C semidefinite typically occurs when working with the mixed finite element formulation of the Stokes problem; see [8]. Such examples can be easily generated using the IFISS package (cf. [7]). Instead of setting $A_0 = A$, A_0 is generally chosen to be a symmetric and positive definite approximation to A , e.g., an incomplete Cholesky decomposition, and C_0 is chosen to be an approximation to the positive or negative Schur-complement. A more general analysis for arbitrary choices of A_0 and C_0 could be performed in the same manner as that in [35]. However, this does not appear to produce useful bounds for the eigenvalues.

4.2.3. Neither A nor C is positive definite. The case where neither A nor C is positive definite is a more severe case since we cannot set $A_0 = A$ or $C_0 = \frac{1}{1+\nu}C$ and expect to obtain a positive definite matrix H . One remedy is to modify A so that the result is positive definite. This may be achieved either during an attempted sparse factorization of A by suitable modifications to its diagonal entries [9, 13], or by modifying the 1×1 and 2×2 diagonal blocks of a computed sparse indefinite factorization; see [6, 30, 31] and [14, section 4.4.2.2].

For a better understanding of the convergence behavior for the case when A is indefinite and nonsingular, we want to present the analysis of the eigenvalues for the simplified case when $A_0 = A$ and $C_0 = C + BA^{-1}B^T$. Recently, Gould and Simoncini [18] analyzed the eigenvalues of saddle-point problems with indefinite leading block both in the unpreconditioned and when preconditioned by

$$P = \begin{bmatrix} A_0 & 0 \\ 0 & C_0 \end{bmatrix},$$

where A_0 and C_0 are as defined above. For the case $C = 0$, a result given in [25] is reproduced, i.e.,

$$\sigma(P^{-1}K) \subset \left\{ 1, \frac{1}{2} \pm \frac{\sqrt{5}}{2} \right\},$$

and for C being symmetric and positive semidefinite, the following proposition is given (see Proposition 4.2 in [18]).

PROPOSITION 4.5. *We assume that C is symmetric and positive semidefinite, A is indefinite, and B is of full rank. When using the preconditioner*

$$P = \begin{bmatrix} A_0 & 0 \\ 0 & C_0 \end{bmatrix}$$

with $A_0 = A$ and $C_0 = C + BA^{-1}B^T$ and assuming that θ is a finite eigenvalue of the pair $(C + BA^{-1}B^T, C)$, the eigenvalues of the preconditioned matrix $P^{-1}K$ are given by

$$(21) \quad \sigma(P^{-1}K) \subset \left\{ 1, \frac{1}{2}(1 \pm \sqrt{5}), \frac{1}{2\theta} \left(\theta - 1 \pm \sqrt{(1 - \theta)^2 + 4\theta^2} \right) \right\}.$$

As we later compare with the block-diagonal preconditioned MINRES method, we analyze the eigenvalues of K preconditioned by

$$P = \begin{bmatrix} A & B^T \\ 0 & C_0 \end{bmatrix}.$$

Similar to Theorem 4.4, it can be seen that the eigenvalues of

$$KP^{-1} = \begin{bmatrix} I & 0 \\ BA^{-1} & -(C + BA^{-1}B^T)C_0^{-1} \end{bmatrix}$$

are either at 1 or given by the eigenvalues of the pencil $-(C + BA^{-1}B^T, C_0)$. Using the setup given in Proposition 4.5, i.e., $A_0 = A$ and $C_0 = C + BA^{-1}B^T$, the preconditioned system has only two eigenvalues at 1 and -1 . This is in contrast to the block-diagonal preconditioner, as there the preconditioned system even in the idealized case depends on the eigenvalues of the pencil $(C + BA^{-1}B^T, C)$. The numerical results given in section 5 illustrate this property.




FIG. 3. Comparison of the FGG method and the Bramble–Pasciak-like CG method for the matrix $CVXQP1_M$ with $C = I$.

5. Numerical experiments. In this section, we provide examples to show how the Bramble–Pasciak-like method can be applied to different problems. The examples in this section either are taken from the CUTer [16] test set or are generated using the IFISS software package [7]. We will again use the structure presented in section 4.2 where different setups of the original matrix are analyzed. The methods we compare in this section are the CG of Forsgren, Gill, and Griffin (when applicable; see section 4.1) and the Bramble–Pasciak-like CG. We will also compare the aforementioned methods with MINRES [8, 26] where the preconditioner is defined as the block-diagonal matrix

$$P = \begin{bmatrix} A_0 & 0 \\ 0 & M_0 \end{bmatrix},$$

where M_0 is a given matrix, and H –MINRES (section 4.2).

C positive definite. In this example, we consider the matrix $CVXQP1_M$ from CUTer, which is of size 1500×1500 . C will either be the identity matrix or a diagonal matrix with entries of the form 10^{-k} on the diagonal where $2 \leq k \leq 10$. We set $M_0 = 0.9C$, $C_0 = M_0$, and $A_0 = \text{diag}(A) + B^T C_0^{-1} B$. The right-hand side is such that z is the vector of all ones. The results for the Bramble–Pasciak-like setup and the FGG method are shown in Figures 3 and 4. Throughout this section, we compare the relative residuals, where the size of the residual is measured using the Euclidean norm. Note that with the presented setup the nonstandard inner-product CG is guaranteed to work since $C - C_0$ will always be positive definite. It can be seen that, for these two choices of C , the performance of the Bramble–Pasciak-like method is very similar to that of the FGG method. It should be noted that if the decomposed form of the Bramble–Pasciak-like preconditioner (see section 4.2) and the FGG preconditioner (see section 4.1) are used, the computational cost of each iteration will be similar for both methods.

A positive definite. The examples we consider in this section are generated using IFISS to discretize the Stokes problem with mixed finite elements: this provides a configuration where A is positive definite and C is positive semidefinite; see [8].




FIG. 4. Comparison of the FGG method and the Bramble–Pasciak-like method for the matrix $CVXQP1_M$ with a random diagonal and positive definite matrix C .

The first test matrix is of size 6659×6659 and describes the flow over a backward-facing step. The matrix A_0 is taken to be the incomplete Cholesky factorization with zero fill-in [28]. A matrix M_0 is generated by IFISS to be the positive definite pressure mass matrix. Note that the Bramble–Pasciak-like method cannot be used reliably with CG. Nevertheless, it very frequently works in practice (see [27] for more details). We also want to mention that methods tailored for this setup, such as the classical Bramble–Pasciak method, require a scaling of the preconditioners. It can be seen from the results in Figure 5 that the Bramble–Pasciak-like CG method (with $C_0 = M_0$) is initially outperformed by the H -MINRES method (with $C_0 = -M_0$) and the standard preconditioned MINRES method. However, the latter two methods then almost stagnate for a large number of iterations (in terms of the Euclidean norm applied to the residual) and, hence, the Bramble–Pasciak-like CG method reaches the desired relative tolerance of 10^{-6} in significantly fewer iterations. We note that the Bramble–Pasciak-like method might break down but we observe good behavior for this problem.

The second test matrix is of size 9539×9539 and describes the flow over a channel domain [8]. The matrix A_0 is chosen such that $A_0 = .9A$ and M_0 is again generated by IFISS as the positive-definite pressure mass matrix. The results given in Figure 6 show that the Bramble–Pasciak-like CG method (with $C_0 = M_0$) outperforms the H -MINRES method (with $C_0 = -M_0$) and preconditioned MINRES. We note that the Bramble–Pasciak-like method is not guaranteed to work for this example but, again, we observe good results.

A indefinite and C semidefinite. In this example, we again consider examples from the CUTer test set where the block A is typically indefinite with zero eigenvalues and the matrix C is positive semidefinite. In [11] it is assumed that if the matrix C is semidefinite, it has block-diagonal form and a zero block in the lower corner. Some preprocessing might be required to guarantee this structure in real world examples, this preprocessing is not required here. We present an example where we again




FIG. 5. Comparison of the H -MINRES method, the Bramble–Pasciak-like, and the preconditioned MINRES method for the backward-facing step.




FIG. 6. Comparison of the H -MINRES method, the Bramble–Pasciak-like, and the preconditioned MINRES method for the flow over the channel domain.

consider the CUTer matrix $CVXQP1_M$ and add the block

$$C = \begin{bmatrix} \hat{C} & 0 \\ 0 & \tilde{C} \end{bmatrix} \in \mathbb{R}^{m \times m},$$

where \hat{C} is a matrix with eigenvalues at zero and \tilde{C} is generated using the MATLAB command

$$1e-1 * \text{sprandsym}(p, .1) + 1e1 * \text{speye}(p);$$

with $p = m - 3$. We set A_0 to be the modified Cholesky preconditioner of A , as presented in section 4.2.3, and then create a Schur-complement-type matrix $M_0 = C + BA_0^{-1}B^T$. As a result, M_0 will always be symmetric and positive definite. Note




FIG. 7. Comparison of the H -MINRES method, the Bramble–Pasciak-like, the sQMR method, and the preconditioned MINRES method for the matrix $CVXQP1_M$ with indefinite A and semidefinite C .




FIG. 8. Comparison of the H -MINRES method and the preconditioned MINRES method for the matrix $CVXQP2_M$ with indefinite A and semidefinite C .

that we can always reliably apply H -MINRES when $C_0 = -M_0$. We will also provide results for the Bramble–Pasciak-like method with $C_0 = M_0$ (which is not guaranteed to work in the case of semidefinite C) and results using sQMR for the choice $C_0 = M_0$. From the results given in Figure 7, it can be observed that preconditioned MINRES needs more iterations than the other methods to achieve the given relative tolerance of 10^{-6} . The other methods all perform similarly and converge in a couple of iterations.

The second example in this section is again taken from CUTer. In particular, we use the matrix $CVXQP2_M$, which is of size 1250×1250 . The setup for C is the same as for $CVXQP1_M$, and we again compute a modified Cholesky matrix A_0 for A , which we then use to generate a Schur-complement-type matrix $C_0 = C + BA_0^{-1}B^T$. The results are shown in Figure 8. In both cases, the leading block is indefinite with a small number of very small negative eigenvalues.

The poorer performance for MINRES with block-diagonal preconditioners compared to the methods based on the block-triangular preconditioners can be explained

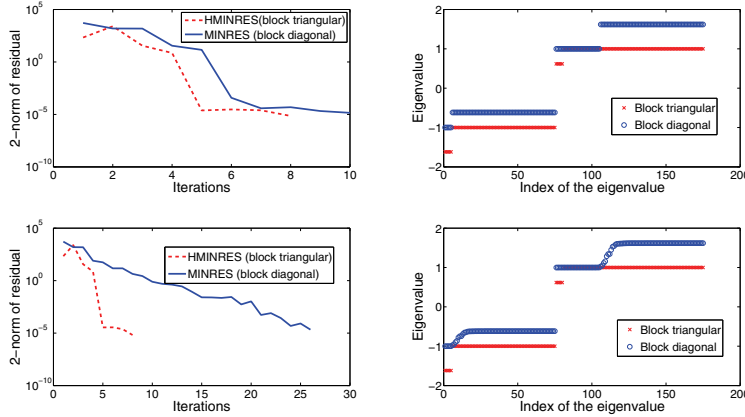


FIG. 9. Eigenvalues of $P^{-1}K$ (right) and convergence (left) for the nonstandard inner-product version of MINRES with BPL^+ block-triangular setup and MINRES with block-diagonal preconditioning for a small 175×175 problem CVXP3-S (see [17]). The top two pictures are for $C = 0$ and the bottom two are for C being positive semidefinite.

by the eigenvalues bounds given in section 4.2.3. Even for the idealized case where $A_0 = A$ and $C_0 = C + BA^{-1}B^T$ the eigenvalues of the preconditioned matrix $\mathbb{P}^{-1}K$ with block-diagonal preconditioning depend on the eigenvalues of the pencil $(C + BA^{-1}B^T, C)$, whereas the eigenvalues of $P^{-1}K$ with block-triangular preconditioner are given by 1 and -1 . This indicates that for the case where $C \neq 0$ the convergence of a method based on the block-diagonal preconditioner is expected to be worse than the convergence of the Bramble–Pasciak-like method. For the case $C = 0$ we expect the performance of MINRES with block-diagonal preconditioning to resemble the convergence behavior of the Bramble–Pasciak-like methods more closely. This is illustrated by the results shown in Figure 9 where the number of eigenvalues clusters is similar for block-diagonal and block-triangular preconditioners in the case $C = 0$. A better clustering is observed for the block-triangular preconditioner if $C \neq 0$.

6. Conclusions. We have presented a reformulation of the saddle-point problem which represents a framework for many well-known solution methods for such problems. We have employed this structure to introduce a Bramble–Pasciak-like method related to a constraint preconditioning technique. The advantage of the presented methods is that we can use block-triangular preconditioners and still employ symmetric methods such as CG or MINRES. Note that even nonsymmetric solvers can be simplified based upon the self-adjointness of the preconditioned matrix in H .

We then considered three cases for which the new methods can be applied. In the case of A being indefinite and C being positive definite, we get competitive results to the FG method, and with the choice C_0 as a multiple of C no preprocessing is required to guarantee the applicability of CG. In other cases, we would have to scale C_0 so that CG is well defined. In the case where A is positive definite and C is positive semidefinite, we illustrate that our method gives competitive results although CG is not well defined. Nevertheless, methods like SQMR still showed good performance in these cases. In the case where both A and C are not positive definite, we showed that the well-defined MINRES (recall we set $C_0 := -C_0$) with nonstandard inner product can outperform MINRES with block-diagonal preconditioning by a large margin.

We have illustrated that competitive results are obtained when the Bramble–Pasciak-like methods are applied to problems arising in optimization. We feel that this is just a first step of understanding the convergence with indefinite leading blocks and see this as a fruitful area for further research.

Appendix A. Proofs. Let

$$\sigma K^{-1} + \begin{bmatrix} D & F^T \\ F & E \end{bmatrix} = \begin{bmatrix} \Omega_1 & \Omega_2^T \\ \Omega_2 & \Omega_3 \end{bmatrix},$$

for given values of Ω_1 , Ω_2 , and Ω_3 . We may factorize this as

$$(22) \quad \begin{bmatrix} \Omega_1 & \Omega_2^T \\ \Omega_2 & \Omega_3 \end{bmatrix} = \begin{bmatrix} I & \Omega_2^T \Omega_3^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} \Omega_1 - \Omega_2^T \Omega_3^{-1} \Omega_2 & 0 \\ 0 & \Omega_3 \end{bmatrix} \begin{bmatrix} I & 0 \\ \Omega_3^{-1} \Omega_2 & I \end{bmatrix}.$$

Using Sylvester's law of inertia, $\sigma K^{-1} + \begin{bmatrix} D & F^T \\ F & E \end{bmatrix}$ is positive definite if and only if Ω_3 and $\Omega_1 - \Omega_2^T \Omega_3^{-1} \Omega_2$ are both positive definite.

Equivalently, we may use the factorization

$$(23) \quad \begin{bmatrix} \Omega_1 & \Omega_2^T \\ \Omega_2 & \Omega_3 \end{bmatrix} = \begin{bmatrix} I & 0 \\ \Omega_2 \Omega_1^{-1} & I \end{bmatrix} \begin{bmatrix} \Omega_1 & 0 \\ 0 & \Omega_3 - \Omega_2 \Omega_1^{-1} \Omega_2^T \end{bmatrix} \begin{bmatrix} I & \Omega_1^{-1} \Omega_2^T \\ 0 & I \end{bmatrix}.$$

Using Sylvester's law of inertia, $\sigma K^{-1} + \begin{bmatrix} D & F^T \\ F & E \end{bmatrix}$ is positive definite if and only if Ω_1 and $\Omega_3 - \Omega_2 \Omega_1^{-1} \Omega_2^T$ are both positive definite.

If A is nonsingular, then

$$K^{-1} = \begin{bmatrix} A^{-1} - A^{-1} B^T S^{-1} B A^{-1} & A^{-1} B^T S^{-1} \\ S^{-1} B A^{-1} & -S^{-1} \end{bmatrix},$$

where $S = C + B A^{-1} B^T$. Use of factorization (22) completes the proof of Corollary 2.2.

If C is nonsingular, then

$$K^{-1} = \begin{bmatrix} S^{-1} & S^{-1} B^T C^{-1} \\ C^{-1} B S^{-1} & C^{-1} B S^{-1} B^T C^{-1} - C^{-1} \end{bmatrix},$$

where $S = A + B^T C^{-1} B$. Use of factorization (23) completes the proof of Corollary 2.3.

If $C = 0$, the columns of $Z \in \mathbb{R}^{n \times (n-m)}$ span the nullspace of B , and B^\dagger is the Moore–Penrose inverse of B , then

$$K^{-1} = \begin{bmatrix} Z S^{-1} Z^T & (I - Z S^{-1} Z^T A) B^\dagger \\ B^{\dagger T} (I - A Z S^{-1} Z^T) & B^{\dagger T} (A Z S^{-1} Z^T A - A) B^\dagger \end{bmatrix},$$

where $S = Z^T A Z$. Use of factorization (23) completes the proof of Corollary 2.4.

REFERENCES

- [1] S. F. ASHBY, T. A. MANTEUFFEL, AND P. E. SAYLOR, *A taxonomy for conjugate gradient methods*, SIAM J. Numer. Anal., 27 (1990), pp. 1542–1568.
- [2] S. F. ASHBY, M. J. HOLST, T. A. MANTEUFFEL, AND P. E. SAYLOR, *The role of the inner product in stopping criteria for conjugate gradient iterations*, BIT, 41 (2001), pp. 26–52.

- [3] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.
- [4] M. BENZI AND V. SIMONCINI, *On the eigenvalues of a class of saddle point matrices*, Numer. Math., 103 (2006), pp. 173–196.
- [5] J. H. BRAMBLE AND J. E. PASCIAK, *A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems*, Math. Comp., 50 (1988), pp. 1–17.
- [6] S. H. CHENG AND N. J. HIGHAM, *A modified Cholesky algorithm based on a symmetric indefinite factorization*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 1097–1110.
- [7] H. C. ELMAN, A. RAMAGE, AND D. J. SILVESTER, *Algorithm 886: IFISS, a Matlab toolbox for modelling incompressible flow*, ACM Trans. Math. Software, 33 (2007), Art. 14.
- [8] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Numer. Math. Sci. Comput., Oxford University Press, New York, 2005.
- [9] H. FANG AND D. O’LEARY, *Modified Cholesky algorithms: A catalog with new approaches*, Math. Program. Ser. A, 115 (2008), pp. 319–349.
- [10] B. FISCHER, A. RAMAGE, D. J. SILVESTER, AND A. J. WATHEN, *Minimum residual methods for augmented systems*, BIT, 38 (1998), pp. 527–543.
- [11] A. FORSGREN, P. E. GILL, AND J. D. GRIFFIN, *Iterative solution of augmented systems arising in interior methods*, SIAM J. Optim., 18 (2007), pp. 666–690.
- [12] R. W. FREUND, *Transpose-free quasi-minimal residual methods for non-Hermitian linear systems*, in Recent Advances in Iterative Methods, IMA Vol. Math. Appl. 60, Springer, New York, 1994, pp. 69–94.
- [13] P. E. GILL AND W. MURRAY, *Newton-type methods for unconstrained and linearly constrained optimization*, Math. Programming, 7 (1974), pp. 311–350.
- [14] P. E. GILL, W. MURRAY, AND M. H. WRIGHT, *Practical Optimization*, Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London-New York, 1981.
- [15] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins Stud. Math. Sci., Johns Hopkins University Press, Baltimore, MD, 1996.
- [16] N. GOULD, D. ORBAN, AND P. TOINT, *CUTEr and SifDec: A constrained and unconstrained testing environment, revisited*, ACM Trans. Math. Software, 29 (2003), pp. 373–394.
- [17] N. GOULD, D. ORBAN, AND P. L. TOINT, *CUTEr, a Constrained and Unconstrained Testing Environment, Revisited*, Technical report TR/PA/01/04, CERFACS, Toulouse, France, 2001.
- [18] N. GOULD AND V. SIMONCINI, *Spectral Analysis of Saddle Point Matrices with Indefinite Leading Blocks*, Technical report, Rutherford Appleton Laboratory, Chilton, England, 2008.
- [19] N. I. M. GOULD, *On the accurate determination of search directions for simple differentiable penalty functions*, IMA J. Numer. Anal., 6 (1986), pp. 357–372.
- [20] M. H. GUTKNECHT, *Changing the norm in conjugate gradient type algorithms*, SIAM J. Numer. Anal., 30 (1993), pp. 40–56.
- [21] 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.
- [22] A. KLAOWONN, *Block-triangular preconditioners for saddle point problems with a penalty term*, SIAM J. Sci. Comput., 19 (1998), pp. 172–184.
- [23] J. LIESEN AND B. N. PARLETT, *On nonsymmetric saddle point matrices that allow conjugate gradient iterations*, Numer. Math., 108 (2008), pp. 605–624.
- [24] A. MEYER AND T. STEIDTEN, *Improvements and Experiments on the Bramble–Pasciak Type CG for Mixed Problems in Elasticity*, Technical report SFB393/01-13, TU Chemnitz, Chemnitz, Germany, 2001.
- [25] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972.
- [26] C. C. PAIGE AND M. A. SAUNDERS, *Solutions of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [27] M. ROZLOŽNÍK AND V. SIMONCINI, *Krylov subspace methods for saddle point problems with indefinite preconditioning*, SIAM J. Matrix Anal. Appl., 24 (2002), pp. 368–391.
- [28] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, 2003.
- [29] M. A. SAUNDERS, *Cholesky-based methods for sparse least squares: The benefits of regularization*, in Linear and Nonlinear Conjugate Gradient-Related Methods (Seattle, WA, 1995), Loyce Adams and J. L. Lazareth, eds., SIAM, Philadelphia, 1996, pp. 92–100.
- [30] R. B. SCHNABEL AND E. ESKOW, *A new modified Cholesky factorization*, SIAM J. Sci. Statist. Comput., 11 (1990), pp. 1136–1158.
- [31] R. B. SCHNABEL AND E. ESKOW, *A revised modified Cholesky factorization algorithm*, SIAM J. Optim., 9 (1999), pp. 1135–1148.

- [32] J. SCHÖBERL AND W. ZULEHNER, *Symmetric indefinite preconditioners for saddle point problems with applications to PDE-constrained optimization problems*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 752–773.
- [33] V. SIMONCINI, *Block triangular preconditioners for symmetric saddle-point problems*, Appl. Numer. Math., 49 (2004), pp. 63–80.
- [34] M. STOLL, *Solving Linear Systems Using the Adjoint*, Ph.D. thesis, University of Oxford, Oxford, UK, 2009.
- [35] M. STOLL AND A. WATHEN, *Combination preconditioning and the Bramble–Pasciak⁺ preconditioner*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 582–608.
- [36] S. J. WRIGHT, *Primal-Dual Interior-Point Methods*, SIAM, Philadelphia, 1997.