

THE ANTITRIANGULAR FACTORIZATION OF SADDLE POINT MATRICES*

J. PESTANA[†] AND A. J. WATHEN[‡]

Abstract. Mastronardi and Van Dooren [*SIAM J. Matrix Anal. Appl.*, 34 (2013), pp. 173–196] recently introduced the block antitriangular (“Batman”) decomposition for symmetric indefinite matrices. Here we show the simplification of this factorization for saddle point matrices and demonstrate how it represents the common nullspace method. We show that rank-1 updates to the saddle point matrix can be easily incorporated into the factorization and give bounds on the eigenvalues of matrices important in saddle point theory. We show the relation of this factorization to constraint preconditioning and how it transforms but preserves the structure of block diagonal and block triangular preconditioners.

Key words. block triangular preconditioner, convergence, eigenvalues, eigenvectors, iterative method, saddle point system

AMS subject classifications. 65F05, 15A23, 65F15, 65F08

DOI. 10.1137/130934933

1. Introduction. The antitriangular factorization proposed by Mastronardi and Van Dooren [17] converts a symmetric indefinite matrix $H \in \mathbb{R}^{p \times p}$ into a block antitriangular matrix M using orthogonal similarity transforms. The factorization can be performed in a backward stable manner and linear systems with the block antitriangular matrix may be efficiently solved. Moreover, the orthogonal similarity transforms preserve eigenvalues and reveal the inertia of H . Thus, from M one can determine the triple (n_-, n_0, n_+) of H , where n_- is the number of negative eigenvalues, n_0 is the number of zero eigenvalues, and n_+ is the number of positive eigenvalues.

The antitriangular factorization takes the form

$$(1.1) \quad H = QMQ^T, \quad Q^{-1} = Q^T, \quad M = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & Y^T \\ 0 & 0 & X & Z^T \\ 0 & Y & Z & W \end{bmatrix} \begin{matrix} \} n_0 \\ \} n_1 \\ \} n_2 \\ \} n_1 \end{matrix},$$

where $n_1 = \min(n_-, n_+)$, $n_2 = \max(n_-, n_+) - n_1$, $Z \in \mathbb{R}^{n_1 \times n_2}$, $W \in \mathbb{R}^{n_1 \times n_1}$ is symmetric, $X = \epsilon LL^T \in \mathbb{R}^{n_2 \times n_2}$ is symmetric definite whenever $n_2 > 0$, and $Y \in \mathbb{R}^{n_1 \times n_1}$ is nonsingular and antitriangular, so that entries above the main antidiagonal are zero. Additionally,

$$\epsilon = \begin{cases} 1 & \text{if } n_+ > n_-, \\ -1 & \text{if } n_- > n_+. \end{cases}$$

The matrix M is strictly antitriangular whenever $n_2 = 0, 1$, i.e., whenever the number of positive and negative eigenvalues differs by at most one. However, the

*Received by the editors August 29, 2013; accepted for publication (in revised form) by N. Mastronardi January 28, 2014; published electronically April 1, 2014. This publication was based on work supported in part by award KUK-C1-013-04 from the King Abdullah University of Science and Technology (KAUST).

<http://www.siam.org/journals/simax/35-2/93493.html>

[†]School of Mathematics, University of Manchester, Manchester M13 9PL, United Kingdom (jennifer.pestana@manchester.ac.uk).

[‡]Computing Laboratory, Oxford University, Oxford OX1 3QD, United Kingdom (wathen@maths.ox.ac.uk).

“bulge” X increases in dimension as H becomes closer to definite. In the extreme case that H is symmetric positive (or negative) definite $n_0 = n_1 = 0$, i.e., X is itself a $p \times p$ matrix. Accordingly, the antitriangular factorization is perhaps best suited to matrices that have a significant number of both positive and negative eigenvalues. We emphasize, however, its generality for real symmetric matrices.

Saddle point matrices are symmetric and indefinite, so that the antitriangular factorization can be applied. These matrices arise in numerous applications [2, section 2] and have the form

$$(1.2) \quad \mathcal{A} = \left[\begin{array}{cc} A & B^T \\ B & 0 \end{array} \right] \left\{ \begin{array}{l} n \\ m \end{array} \right\},$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric (but not necessarily positive definite) and $B \in \mathbb{R}^{m \times n}$, $m \leq n$. The matrix \mathcal{A} is nonsingular with n positive eigenvalues and m negative eigenvalues when A is positive definite on the nullspace of B and $\text{rank}(B) = m$. We consider only this most common situation here.

The algorithm for computing an antitriangular factorization proposed by Mastronardi and Van Dooren is designed to be applicable to *all* symmetric indefinite matrices. In this note we show that their algorithm simplifies when applied to saddle point matrices. An alternative based on a QR factorization of B^T , which is like the approach applied by Mastronardi and Van Dooren to specific saddle point problems arising in constrained indefinite least squares [16], gives a different but related antitriangular form. Both algorithms are shown to be strongly backward stable but the optimal algorithm in terms of cost depends on the sizes of m and n .

Low-rank updates of A and B in \mathcal{A} , such as those used in quasi-Newton methods [8], interior point methods [1], or the augmented Lagrangian method [2, section 3.5], can be efficiently incorporated into an antitriangular factorization of \mathcal{A} . Additionally, bounds on the eigenvalues of A and the (negative) Schur complement $BA^{-1}B^T$ that depend only on the smaller blocks W , X , and Y of the antitriangular matrix can be obtained.

We show that solving a saddle point system in antitriangular form is equivalent to applying the nullspace method [2, section 6], [21, section 15.2]. In other words, the antitriangular factorization allows the nullspace method to be represented not just as a procedure but also as a matrix decomposition, similar to other well-known methods for solving linear systems like Gaussian elimination.

If the matrix \mathcal{A} is large, we may solve the saddle point system by an iterative method rather than a direct method like the antitriangular factorization. When preconditioning is required, block preconditioners, such as block diagonal, block triangular, and constraint preconditioners, are popular choices for saddle point systems. We show that the same orthogonal transformation matrix that converts \mathcal{A} into an antitriangular matrix can be applied to these preconditioners and that relevant structures are preserved.

The outline of our manuscript is as follows. The two algorithms are given in section 2, where their complexities are also compared. Stability, extensions, and low-rank updates are discussed in section 3, while the connection to the nullspace method is outlined in section 4. We state our eigenvalue bounds in section 5 and discuss preconditioners in section 6. Finally, section 7 contains our conclusions.

Throughout, we use MATLAB notation to denote submatrices. Thus $K(q : r, s : t)$ is the submatrix of K comprising the intersection of rows q to r with columns s to t . Also, $K(r : -1 : q, s : t)$ (or $K(q : r, t : -1 : s)$) represents the submatrix

$K(q : r, s : t)$ with its rows (or columns) in reverse order. The nullspace and range of a matrix K are denoted by $\text{null}(K)$ and $\text{range}(K)$, respectively.

2. An antitriangular factorization of saddle point matrices. We are interested in applying orthogonal transformations to the saddle point matrix \mathcal{A} in (1.2) to obtain the antitriangular matrix (1.1). Since \mathcal{A} is nonsingular with n positive eigenvalues and m negative eigenvalues, in this case the antitriangular matrix has the specific form

$$\mathcal{M} = \left[\begin{array}{ccc} 0 & 0 & Y^T \\ 0 & X & Z^T \\ Y & Z & W \end{array} \right] \left. \begin{array}{l} \} m \\ \} n - m \\ \} m \end{array} \right\},$$

where $Y \in \mathbb{R}^{m \times m}$ is antitriangular, $X \in \mathbb{R}^{(n-m) \times (n-m)}$ is symmetric positive definite, and $W \in \mathbb{R}^{m \times m}$ is symmetric. We note that linear systems with \mathcal{M} can be solved with the obvious “antitriangular” substitution (finding the last variable from the first equation, the second-last variable from the second equation, and so forth) twice, with a solve with the positive definite matrix X (using, for example, a Cholesky factorization) in between.

2.1. The algorithm of Mastronardi and Van Dooren. Although it is possible to compute an antitriangular factorization of (1.2) by the algorithm of Mastronardi and Van Dooren, the result is somewhat more involved than necessary since the algorithm simplifies if we first first permute \mathcal{A} to

$$(2.1) \quad \tilde{\mathcal{A}} = Q_1^T \mathcal{A} Q_1 = \left[\begin{array}{cc} 0 & B \\ B^T & A \end{array} \right] \left. \begin{array}{l} \} m \\ \} n \end{array} \right\}, \quad Q_1 = \begin{bmatrix} 0 & I_n \\ I_m & 0 \end{bmatrix}.$$

Given the permuted matrix (2.1), the algorithm of Mastronardi and Van Dooren proceeds outward from the (1,1) entry of $\tilde{\mathcal{A}}$, at each stage updating the antitriangular factorization of $\tilde{\mathcal{A}}(1 : k, 1 : k)$ to give a factorization of $\tilde{\mathcal{A}}(1 : k+1, 1 : k+1)$. Accordingly, the first m steps leave $\tilde{\mathcal{A}} = Q_1^T \mathcal{A} Q_1$ unchanged and the inertia of $\tilde{\mathcal{A}}(1 : m, 1 : m)$ is $(0, m, 0)$.

The next stage of the algorithm uses Householder matrices to convert B to an antitrapezoidal matrix

$$(2.2) \quad HB = V = \begin{bmatrix} V^{(1)} & V^{(2)} \end{bmatrix},$$

where $H \in \mathbb{R}^{m \times m}$ is orthogonal, $V^{(1)} \in \mathbb{R}^{m \times m}$ is antitriangular, and $V^{(2)} \in \mathbb{R}^{m \times (n-m)}$. Accordingly, after a further m steps, we obtain

$$Q_2^T Q_1^T \mathcal{A} Q_1 Q_2 = \begin{bmatrix} 0 & V \\ V^T & A \end{bmatrix}, \quad Q_2 = \begin{bmatrix} H^T & \\ & I_n \end{bmatrix},$$

and the inertia of the submatrix formed from the first $2m$ rows and columns of $Q_2^T Q_1^T \mathcal{A} Q_1 Q_2$ is $(m, 0, m)$.

If $n = m$ we are finished, although this situation is rare in practice. Otherwise, we must reduce V to antitriangular form by Givens rotations from the right. Thus,

$$(2.3) \quad VG = V[G^{(1)} \quad G^{(2)}] = \begin{bmatrix} 0 & Y^T \end{bmatrix},$$

where $G \in \mathbb{R}^{n \times n}$, $G^{(1)} \in \mathbb{R}^{n \times (n-m)}$, $Y \in \mathbb{R}^{m \times m}$. It follows from (2.2) that

$$B \begin{bmatrix} G^{(1)} & G^{(2)} \end{bmatrix} = \begin{bmatrix} 0 & (YH)^T \end{bmatrix}$$

and that $G^{(1)}$ and $G^{(2)}$ are bases for $\text{null}(B)$ and $\text{range}(B^T)$, respectively. Thus, applying

$$Q_3 = \begin{bmatrix} I_m & \\ & G \end{bmatrix}$$

gives the antitriangular form

$$\mathcal{M}_1 = Q_1^T A Q_1 = \begin{bmatrix} 0 & 0 & Y_1^T \\ 0 & X_1 & Z_1^T \\ Y_1 & Z_1 & W_1 \end{bmatrix}, \quad Q_1 = Q_1 Q_2 Q_3 = \begin{bmatrix} 0 & G^{(1)} & G^{(2)} \\ H^T & 0 & 0 \end{bmatrix},$$

where $Z_1 = (G^{(2)})^T A G^{(1)} \in \mathbb{R}^{m \times (n-m)}$, $W_1 = (G^{(2)})^T A G^{(2)} \in \mathbb{R}^{m \times m}$ is symmetric and $X_1 = (G^{(1)})^T A G^{(1)} \in \mathbb{R}^{(n-m) \times (n-m)}$ is symmetric positive definite, since A is positive definite on the nullspace of B . An algorithm for this procedure is given in Algorithm 1.

ALGORITHM 1. Antitriangular factorization of a saddle point matrix by the algorithm of Mastronardi and Van Dooren.

Input: Saddle point matrix \mathcal{A} from (1.2)

Output: Antitriangular matrix \mathcal{M}_1 and orthogonal matrix Q_1 such that $\mathcal{A} = Q_1 \mathcal{M}_1 Q_1^T$

Permute the rows and columns of \mathcal{A} as in (2.1)

Compute an upper trapezoidal factorization of B as in (2.2) by Householder matrices

Compute the Givens rotations that transform V to antitriangular form

$VG = \begin{bmatrix} 0 & Y_1^T \end{bmatrix}$ as in (2.3)

Set $G^{(1)} = G(1:n, 1:n-m)$ and $G^{(2)} = G(1:n, n-m+1:n)$

Set $X_1 = (G^{(1)})^T A G^{(1)}$, $Z_1 = (G^{(2)})^T A G^{(1)}$ and $W_1 = (G^{(2)})^T A G^{(2)}$

Set $\mathcal{M}_1 = \begin{bmatrix} 0 & 0 & Y_1^T \\ 0 & X_1 & Z_1^T \\ Y_1 & Z_1 & W_1 \end{bmatrix}$ and $Q_1 = \begin{bmatrix} 0 & G^{(1)} & G^{(2)} \\ H^T & 0 & 0 \end{bmatrix}$

Note that in this case we avoid the more complex case c in the Mastronardi and Van Dooren algorithm since, although A may not be definite, the positive definite part of A is automatically obtained in the process of antitriangularizing B and B^T . In contrast, applying the algorithm to (1.2) would involve this third case.

2.2. An alternative. By reordering the operations, we obtain an alternative to Algorithm 1 which instead involves only permutations and a QR factorization of B^T ,

$$(2.4) \quad B^T = \underbrace{\begin{bmatrix} U^{(1)} & U^{(2)} \end{bmatrix}}_U \begin{bmatrix} R \\ 0 \end{bmatrix},$$

where $U \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$. Note that now the columns of $U^{(1)} \in \mathbb{R}^{n \times m}$ form an orthonormal basis for $\text{range}(B^T)$, while the columns of $U^{(2)}$ form an orthonormal basis for $\text{null}(B)$.

As in the previous algorithm we start from (2.1) but now apply

$$Q_4 = \begin{bmatrix} I_m & 0 \\ 0 & U \end{bmatrix}$$

to obtain

$$(2.5) \quad Q_4^T Q_1^T \mathcal{A} Q_1 Q_4 = \begin{bmatrix} 0 & R^T & 0 \\ R & \hat{A}_{11} & \hat{A}_{12} \\ 0 & \hat{A}_{12}^T & \hat{A}_{22} \end{bmatrix} \begin{matrix} \} m \\ \} m \\ \} n - m \end{matrix},$$

where $\hat{A}_{ij} = (U^{(i)})^T A U^{(j)}$, $i, j = 1, 2$. Note that $\hat{A}_{22} = (U^{(2)})^T A U^{(2)}$ is positive definite, analogously to the Mastronardi and Van Dooren algorithm.

Then all that remains is to permute the last $n - m$ rows and columns so that R is transformed to an antitriangular matrix that sits in the last m rows. This is achieved by applying

$$(2.6) \quad Q_5 = \begin{bmatrix} I_m & 0 \\ 0 & \hat{S} \end{bmatrix}, \quad \hat{S} = \begin{bmatrix} 0 & S_m \\ I_{n-m} & 0 \end{bmatrix}, \quad S_m = \begin{bmatrix} & & 1 \\ & \ddots & \\ 1 & & \end{bmatrix}.$$

The matrix S_m is the $m \times m$ reverse identity, which satisfies $S_m^{-1} = S_m^T = S_m$.

Combining these steps gives

$$(2.7) \quad \mathcal{M}_2 = Q_2^T \mathcal{A} Q_2 = \begin{bmatrix} 0 & 0 & Y_2^T \\ 0 & X_2 & Z_2^T \\ Y_2 & Z_2 & W_2 \end{bmatrix}, \quad Q_2 = Q_1 Q_4 Q_5 = \begin{bmatrix} 0 & U^{(2)} & U^{(1)} S_m \\ I_m & 0 & 0 \end{bmatrix},$$

where $Y_2 = S_m R \in \mathbb{R}^{m \times m}$, $Z_2 = S_m \hat{A}_{12}$, $W_2 = S_m \hat{A}_{11} S_m$ is symmetric and $X_2 = \hat{A}_{22} \in \mathbb{R}^{(n-m) \times (n-m)}$ is symmetric positive definite. We summarize this method in Algorithm 2.

ALGORITHM 2. Antitriangular factorization of a saddle point matrix using the QR factorization.

Input: Saddle point matrix \mathcal{A} from (1.2)

Output: Antitriangular matrix \mathcal{M}_2 and orthogonal matrix Q_2 such that $\mathcal{A} = Q_2 \mathcal{M}_2 Q_2^T$

Permute the rows and columns of \mathcal{A} as in (2.1)

Compute the QR factorization $B^T = UR$

Set $U^{(1)} = U(1 : n, 1 : m)$ and $U^{(2)} = U(1 : n, m + 1 : n)$

Compute $\hat{A}_{11} = (U^{(1)})^T A U^{(1)}$, $\hat{A}_{12} = (U^{(1)})^T A U^{(2)}$ and $\hat{A}_{22} = (U^{(2)})^T A U^{(2)}$

Set $Y_2 = R(m : -1 : 1, 1 : m)$, $X_2 = \hat{A}_{22}$, $Z_2 = \hat{A}_{12}(m : -1 : 1, 1 : n - m)$ and $W_2 = \hat{A}_{11}(m : -1 : 1, m : -1 : 1)$

Set $\mathcal{M}_2 = \begin{bmatrix} 0 & 0 & Y_2^T \\ 0 & X_2 & Z_2^T \\ Y_2 & Z_2 & W_2 \end{bmatrix}$ and $Q_2 = \begin{bmatrix} 0 & U^{(2)} & U^{(1)}(1 : n, m : -1 : 1) \\ I_m & 0 & 0 \end{bmatrix}$

2.3. Complexity of the antitriangular algorithms. Both Algorithms 1 and 2 start from the permuted matrix $Q_1^T \mathcal{A} Q_1$ in (2.1) and convert B and B^T to antitriangular form. Algorithm 1 achieves this by a two-sided orthogonal transformation $H B G = [0 \ Y_1^T]$, where H gives the intermediate antitrapezoidal form (2.2). Algorithm 2 instead uses the one-sided transformation $B U \hat{S} = [0 \ Y_2^T]$. Thus, the differences between the algorithms are due to the choice of one-sided or two-sided

TABLE 1
Complexity of Algorithms 1 and 2.

Algorithm	Dimensions	Flops
Algorithm 1	$n = m$	$2m^2n - \frac{2}{3}m^3$
	$m < n < 2m$	$3n^3 - 3mn^2 + 8m^2n - \frac{20}{3}m^3$
	$n = 2m$	$6mn^2 + 2m^2n - \frac{20}{3}m^3$
	$n > 2m$	$12n^2m - 10m^2n + \frac{16}{3}m^3$
Algorithm 2	$n \geq m$	$8mn^2 - 2m^2n - \frac{2}{3}m^3$

transformations and the optimal choice in terms of floating point operations depends on the ratio of n to m , as we now show.

The Mastronardi and Van Dooren algorithm first uses Householder transforms to convert B to Y and this requires $2m^2(n - m/3)$ flops [6, section 5.2.1]. If $n > m$ we must apply $n - m$ sequences of m Givens rotations to convert V to the correct antitriangular form. Each sequence annihilates an antidiagonal of V and alters m columns of $Q_2^T Q_1^T A Q_1 Q_2$. The total number of flops required to apply these Givens rotations from the right is

$$6 \sum_{j=1}^{n-m} \sum_{i=1}^m (n + m - i + 1) \approx 6n^2m - 3nm^2 - 3m^3.$$

Although, by exploiting symmetry, we can update V^T and $A(m + 1 : n, 1 : m)$ without additional computations, we must still apply Givens rotations to the rows of $A(m + 1 : n, m + 1 : n)$. The number of operations depends on the size of $n - m$ compared with the size of m . As the first antidiagonal of V is annihilated, we apply $3(n - m)$ operations to the rows of $A(m + 1 : n, m + 1 : n)$, at the second $6(n - m)$, at the third $12(n - m)$, and so on until either we reach the last row of the matrix or we have applied m sequences of Givens rotations. This requires

$$3(n - m) + 6(n - m) \sum_{j=1}^{r-1} i, \quad r = \min\{n - m, m\}$$

flops or, to leading order, $3(n - m)^3$ when $n < 2m$ and $3(n - m)m^2$ otherwise. If $n > 2m$ we must apply additional Givens rotations to make V and V^T antitriangular at a cost of $6m(n - m)(n - 2m)$ flops. The total flop counts for these different cases are given in Table 1.

The cost of Algorithm 2, which involves only the QR factorization of B^T and the formation of $U^T A U$, can also be determined. The QR decomposition of B^T requires $2m^2(n - m/3)$ flops if Householder transformations are used. Then $U^T A$ can be computed in approximately $2mn(2n - m)$ flops [6, section 5.1.6] and similarly for $(U^T A)U$. Thus, the total cost of computing the antitriangular factorization by Algorithm 2 is approximately $8mn^2 - 2m^2(n + m/3)$ flops.

From this comparison it is clear that the optimal algorithm depends on the size of m , the number of constraints, relative to the number of primal variables n . If m is almost as large as n Algorithm 2 is favorable, while Algorithm 1 is better when m is small relative to n .

Unless otherwise stated, we concentrate on the QR variant (Algorithm 2) in the remainder of this manuscript for ease of exposition, but the same analysis could easily be applied to the antitriangular matrix from Algorithm 1. We additionally drop subscripts on the matrices \mathcal{M} , \mathcal{Q} , W , X , Y , and Z .

3. Properties of the antitriangular decomposition. In this section we discuss properties of the antitriangular decomposition of saddle point matrices, including stability, extensions, and low-rank modifications.

3.1. Stability. Algorithms 1 and 2 are not only backward stable (provided the QR decomposition in Algorithm 2 is computed in a backward stable manner) but are strongly stable, in the sense of Sun [25], i.e., the computed matrices \mathcal{Q} and \mathcal{M} satisfy

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

where $\|\Delta A\|_2/\|A\|_2 = O(\epsilon_m)$, $\|\Delta B\|_2/\|B\|_2 = O(\epsilon_m)$, $\|\Delta B^T\|_2/\|B^T\|_2 = O(\epsilon_m)$, and ϵ_m is machine precision. To prove this we first note that the antitriangular decomposition comprises two parts: the antitriangular factorization of B and B^T and the multiplication of A by orthogonal matrices. The factorization of B and B^T by either Algorithm 1 or 2 is backward stable and the computed matrices \bar{G} and \bar{U} satisfy $\|\bar{G} - G\|_2 = O(\epsilon_m)$ and $\|\bar{U} - U\|_2 = O(\epsilon_m)$ for some orthogonal G and U ; this is proved in a similar way to the QR factorization results in the book by Higham [9, Chapter 19]. Additionally, multiplication by an orthogonal matrix is backward stable [9, section 3.5]. Consequently, neither algorithm should have problems with breakdown, but if B is numerically rank deficient, then Y will be as well, as expected.

3.2. Extensions. Although we consider only real matrices, since this is the most prevalent case in practice, the extension of Algorithms 1 and 2 to complex Hermitian matrices is trivial if we apply unitary matrices instead of orthogonal ones.

We can also find a factorization of non-Hermitian matrices, such as block complex symmetric matrices, i.e., matrices of the form (1.2) but with $A \in \mathbb{C}^{n \times n}$, $A = A^T$, and $B \in \mathbb{C}^{m \times n}$. Such matrices arise in, for example, electrical networks [2, p. 5], [11], [15]. In this setting \mathcal{Q} is complex, but the complex symmetry is preserved, i.e., the resulting antitriangular matrix is complex symmetric. It no longer makes sense to discuss inertia, since the eigenvalues of \mathcal{A} may be complex. Moreover, these eigenvalues are not preserved by \mathcal{M} , since $\mathcal{Q}^T \neq \mathcal{Q}^{-1}$. However, solving systems with this matrix is straightforward and the process is equivalent to a nullspace method, such as that employed by Mahawar and Sarin [15].

3.3. Updating the antitriangular factorization. Mastronardi and Van Dooren showed that the antitriangular factorization can be efficiently updated when a rank-1 modification is applied. These updates can be somewhat involved when applied to a general symmetric matrix but the procedure simplifies for saddle point matrices. We discuss some relevant modifications here.

If A is ill-conditioned or singular it may be desirable to apply the augmented Lagrangian approach in which we replace (4.1) by [2, section 3.5]

$$(3.1) \quad \mathcal{A}_{AL}x = \begin{bmatrix} A + B^T E B & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f + B^T E g \\ g \end{bmatrix},$$

where $E \in \mathbb{R}^{m \times m}$ is symmetric positive definite. Updating the antitriangular factorization in this case is straightforward, since $B^T E B$ is orthogonal to $\text{null}(B)$. Thus, given the antitriangular factorization (2.7) of $\mathcal{A} = \mathcal{Q}\mathcal{M}\mathcal{Q}^T$, the antitriangular factorization of \mathcal{A}_{AL} is $\mathcal{A}_{AL} = \mathcal{Q}\mathcal{M}_{AL}\mathcal{Q}^T$, where

$$\mathcal{M}_{AL} = \begin{bmatrix} 0 & 0 & Y^T \\ 0 & X & Z^T \\ Y & Z & W + Y E Y^T \end{bmatrix}.$$

The idea can be extended to more general symmetric positive semidefinite updates $F \in \mathbb{R}^{n \times n}$ to A . If

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

then the antitriangular factorization of \mathcal{A}_F is $\mathcal{A}_F = \mathcal{Q}\mathcal{M}_F\mathcal{Q}^T$, where

$$\mathcal{M}_F = \begin{bmatrix} 0 & 0 & Y^T \\ 0 & X_F & Z_F^T \\ Y & Z_F & W_F \end{bmatrix},$$

with $X_F = X + (U^{(2)})^T F U^{(2)}$, $Z_F = Z + S_m (U^{(1)})^T F U^{(2)}$, and $W_F = W + S_m (U^{(1)})^T F U^{(1)} S_m$. If F is low-rank, then the updates to W , X and Z can be cheaply computed.

When a sequence of saddle point matrices are solved, as in the quasi-Newton method, or in interior point methods, it may be necessary to update B and B^T as well as A . If the updates have special structure, the antitriangular factorization can be updated by low-rank approximations as in Griewank, Walther, and Korzec [8]. In the generic case, however, we require a low-rank update of the antitriangular factorizations of B and B^T , which can be obtained by extending the rank-1 update procedure described in Mastronardi and Van Dooren [17] or by using an updated QR factorization. Since the approaches are similar, we describe the QR approach here.

We consider the updated matrix

$$\mathcal{A}_{UP} = \mathcal{A} + uv^T + vu^T = \begin{bmatrix} A & (B + u_1 v_1^T)^T \\ (B + u_1 v_1^T) & 0 \end{bmatrix},$$

where $u = [0 \ u_1]^T$, $v = [v_1 \ 0]^T$, $u_1 \in \mathbb{R}^{m \times 1}$, and $v_1 \in \mathbb{R}^{n \times 1}$. If $B^T = U\hat{R}$, $\hat{R} = [R^T \ 0]^T$, is the QR decomposition of B^T , then a QR decomposition of $B^T + v_1 u_1^T$ is [6, section 12.5.1]

$$B^T + v_1 u_1^T = U_{UP} \hat{R}_{UP} = \begin{bmatrix} U_{UP}^{(1)} & U_{UP}^{(2)} \end{bmatrix} \begin{bmatrix} R_{UP} \\ 0 \end{bmatrix},$$

where $U_{UP} = UJ \in \mathbb{R}^{n \times n}$, $U_{UP}^{(1)} \in \mathbb{R}^{n \times m}$, and J is orthogonal. Since

$$Q_5^T \mathcal{Q}^T (\mathcal{A} + uv^T + vu^T) \mathcal{Q} Q_5^T = \begin{bmatrix} 0 & (\hat{R} + U^T v_1 u_1^T)^T \\ \hat{R} + U^T v_1 u_1^T & U^T A U \end{bmatrix},$$

then with

$$Q_6 = \begin{bmatrix} I_m & 0 \\ 0 & J \end{bmatrix}$$

we have that

$$Q_6^T Q_5^T \mathcal{Q}^T (\mathcal{A} + uv^T + vu^T) \mathcal{Q} Q_5^T Q_6 = \begin{bmatrix} 0 & R_{UP}^T & 0 \\ R_{UP} & (\hat{A}_{UP})_{11} & (\hat{A}_{UP})_{12} \\ 0 & (\hat{A}_{UP})_{21} & (\hat{A}_{UP})_{22} \end{bmatrix},$$

where $(\hat{A}_{UP})_{ij} = (U_{UP}^{(i)})^T A (U_{UP}^{(j)})$, $j = 1, 2$.

Finally, applying Q_5 as in Algorithm 2 gives the antitriangular form:

$$\mathcal{Q}_{UP}^T \mathcal{Q}^T (\mathcal{A} + uv^T + vu^T) \mathcal{Q} \mathcal{Q}_{UP} = \begin{bmatrix} 0 & 0 & Y_{UP}^T \\ 0 & X_{UP} & Z_{UP}^T \\ Y_{UP} & Z_{UP} & W_{UP} \end{bmatrix},$$

where $\mathcal{Q}_{UP} = Q_5^T Q_6 Q_5$, $W_{UP} = S_m(\hat{A}_{UP})_{11} S_m$, $X_{UP} = (\hat{A}_{UP})_{22}$, and $Z_{UP} = S_m(\hat{A}_{UP})_{12}$.

The computational cost associated with the update arises from the application of J , a composition of Givens rotations, to \mathcal{M} and $u_1 v_1^T$.

4. Comparison with the nullspace method. The range space method for solving

$$(4.1) \quad \mathcal{A}x = \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix}$$

is applicable when A is invertible and is related to a block LDL^T decomposition [2, section 5] since

$$\mathcal{A} = \begin{bmatrix} I_n & 0 \\ BA^{-1} & I_m \end{bmatrix} \begin{bmatrix} A & 0 \\ 0 & -BA^{-1}B^T \end{bmatrix} \begin{bmatrix} I_n & A^{-1}B^T \\ 0 & I_m \end{bmatrix}.$$

The matrix factorization representation of the nullspace method is the antitriangular factorization, as we now show.

Given a basis for the nullspace¹ of B , such as $U^{(2)}$, and a particular solution \hat{u} of $Bu = g$, the nullspace method proceeds as follows [2, section 6], [21, section 15.2]:

1. Solve $(U^{(2)})^T A U^{(2)} v = (U^{(2)})^T (f - A\hat{u})$.
2. Set $u_* = U^{(2)} v + \hat{u}$.
3. Solve $BB^T p_* = B(f - Au_*)$.

Then (u_*, p_*) solves (4.1).

On the other hand, applying the antitriangularization (2.7) to (4.1) gives

$$(\mathcal{Q}^T \mathcal{A} \mathcal{Q})y = \mathcal{M}y = (\mathcal{Q}^T b), \quad y = \mathcal{Q}^T x,$$

or

$$\begin{bmatrix} 0 & 0 & Y^T \\ 0 & X & Z^T \\ Y & Z & W \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} g \\ (U^{(2)})^T f \\ S(U^{(1)})^T f \end{bmatrix}.$$

To recover u and p we must solve

$$(4.2a) \quad Y^T y_3 = g,$$

$$(4.2b) \quad X y_2 + Z^T y_3 = (U^{(2)})^T f,$$

$$(4.2c) \quad Y y_1 + Z y_2 + W y_3 = S(U^{(1)})^T f,$$

using the antitriangular substitution described in section 2. This is equivalent to applying the inverse of \mathcal{M} , which has upper block antitriangular structure, that is,

¹Note that any basis of the nullspace of B in general would be sufficient but we use $U^{(2)}$ here as this common choice corresponds to the antitriangular factorization.

$$(4.3) \quad \mathcal{M}^{-1} = \begin{bmatrix} Y^{-1}(ZX^{-1}Z^T - W)Y^{-T} & -Y^{-1}ZX^{-1} & Y^{-1} \\ -X^{-1}Z^TY^{-T} & X^{-1} & 0 \\ Y^{-T} & 0 & 0 \end{bmatrix} \quad \left. \begin{array}{l} \} m \\ \} n-m \\ \} m \end{array} \right\}.$$

Having obtained y , we recover u and p from

$$(4.4) \quad \begin{bmatrix} u \\ p \end{bmatrix} = \mathcal{Q}y = \begin{bmatrix} U^{(2)}y_2 + U^{(1)}S_my_3 \\ y_1 \end{bmatrix}.$$

We now show that solving (4.2a)–(4.2c) is equivalent to applying the nullspace method. From (4.2a), since $Y = S_mR$ and $B^T = U^{(1)}R$,

$$R^T(U^{(1)})^T U^{(1)}S_my_3 = B(U^{(1)}S_my_3) = g,$$

so that $\hat{u} = U^{(1)}S_my_3$ is a particular solution of $Bu = g$.

Since $X = \hat{A}_{22}$ and $Z = S\hat{A}_{12}$, where $\hat{A}_{ij} = (U^{(i)})^T AU^{(j)}$, $i, j = 1, 2$, as before, we have from (4.2b) that

$$(U^{(2)})^T AU^{(2)}y_2 = (U^{(2)})^T (f - A\hat{u}).$$

Substituting for \hat{u} and $W = S\hat{A}_{11}S$ in (4.2c) then gives that

$$\begin{aligned} Rp &= (U^{(1)})^T [f - A(U^{(2)}y_2 + \hat{u})] \\ R^T Rp &= (U^{(1)}R)^T [f - A(U^{(2)}y_2 + \hat{u})] \\ BB^T p &= B(f - Au_*), \end{aligned}$$

where $u_* = U^{(2)}y_2 + \hat{u}$.

Thus, solving a system with the antitriangular factorization is equivalent to applying the nullspace method with the QR nullspace basis and with $\hat{u} = U^{(1)}S_my_3$. From (4.4) we then have that $u = u^* = U^{(2)}y_2 + \hat{u}$ and $p = y_1$.

Note that no antitriangular solves are required in the nullspace method, even though we are solving a linear system with a block antitriangular matrix. This is because the permutation matrix S_m that transforms the upper triangular matrix R to antitriangular form occurs as $S_m^2 = I$ in (4.2a) and can be eliminated from (4.2c).

We have seen that the antitriangular factorization allows us to view the nullspace method as a factorization rather than as a procedure, similar to other direct solvers such as Gaussian elimination, which can be written as the product of structured matrices. This idea could of course be generalized to other factorizations with different representations of the nullspace.

5. Eigenvalue bounds. Of interest when solving saddle point systems are the eigenvalues of A and the (negative) Schur complement $BA^{-1}B^T$ when it exists, i.e., when A is invertible. Since the Schur complement involves the inverse of the $n \times n$ matrix A , its eigenvalues can be particularly difficult to approximate. Here we give bounds for the eigenvalues of both matrices that depend only on the eigenvalues of X and W and the singular values of Y .

Since

$$\widehat{M} = \begin{bmatrix} X & Z^T \\ Z & W \end{bmatrix} = \begin{bmatrix} (U^{(2)})^T \\ S_m(U^{(1)})^T \end{bmatrix} A [U^{(2)}U^{(1)}S_m]$$

for Algorithm 2,² the eigenvalues of A are identical to those of \widehat{M} . This means that Cauchy's interlacing theorem can be used to bound the eigenvalues of A .

LEMMA 1. Let $\lambda_1(A) \leq \lambda_2(A) \leq \dots \leq \lambda_n(A)$ be the eigenvalues of A , $0 < \lambda_1(X) \leq \lambda_2(X) \leq \dots \leq \lambda_{n-m}(X)$ be the eigenvalues of X , and $\lambda_1(W) \leq \lambda_2(W) \leq \dots \leq \lambda_m(W)$ be the eigenvalues of W . Then,

$$\begin{aligned}\lambda_k(A) &\leq \lambda_k(X) \leq \lambda_{k+m}(A), \quad k = 1, \dots, n-m, \\ \lambda_k(A) &\leq \lambda_k(W) \leq \lambda_{k+n-m}(A), \quad k = 1, \dots, m.\end{aligned}$$

Proof. The results follow from the similarity of \widehat{M} and A and by applying the interlacing theorem [10, Theorem 4.3.15] to \widehat{M} using X or W . \square

Also of interest when A is positive definite are the eigenvalues of $BA^{-1}B^T$. To bound these we first prove the following lemma.

LEMMA 2. Assume that A is positive definite. Let $0 < \lambda_1(\widetilde{W}) \leq \lambda_2(\widetilde{W}) \leq \dots \leq \lambda_m(\widetilde{W})$ be the eigenvalues of $\widetilde{W} = S_m(U^{(1)})^T A^{-1} U^{(1)} S_m$ and $0 < \lambda_1(W^{-1}) \leq \lambda_2(W^{-1}) \leq \dots \leq \lambda_m(W^{-1})$ be the eigenvalues of W^{-1} . Then,

$$\lambda_k(W^{-1}) \leq \lambda_k(\widetilde{W}), \quad k = 1, \dots, m.$$

Proof. First note that for any $x \neq 0$, the Cauchy-Schwarz inequality gives that $(x^T x)^2 = (x^T A^{\frac{1}{2}} A^{-\frac{1}{2}} x)^2 \leq (x^T A x)(x^T A^{-1} x)$ or

$$\frac{x^T x}{x^T A x} \leq \frac{x^T A^{-1} x}{x^T x}.$$

Using this and the orthogonality of S_m and $U^{(1)}$, the Courant-Fischer theorem [10, p. 180] gives

$$\begin{aligned}\lambda_k(W^{-1}) &= \min_{\dim(S)=k} \max_{\substack{x \in S \\ x \neq 0}} \frac{x^T S_m ((U^{(1)})^T A U^{(1)})^{-1} S_m x}{x^T S_m^T S_m x} \\ &= \min_{\dim(S)=k} \max_{\substack{y \in S \\ y \neq 0}} \frac{y^T (U^{(1)})^T U^{(1)} y}{y^T (U^{(1)})^T A U^{(1)} y} \\ &\leq \min_{\dim(S)=k} \max_{\substack{y \in S \\ y \neq 0}} \frac{y^T (U^{(1)})^T A^{-1} U^{(1)} y}{y^T (U^{(1)})^T U^{(1)} y} \\ &= \min_{\dim(S)=k} \max_{\substack{z \in S \\ z \neq 0}} \frac{z^T S_m (U^{(1)})^T A^{-1} U^{(1)} S_m z}{z^T z} = \lambda_k(\widetilde{W}). \quad \square\end{aligned}$$

Lemma 2 can be used to bound the eigenvalues of $BA^{-1}B^T$ as follows.

COROLLARY 3. Let $0 < \lambda_1(W^{-1}) \leq \lambda_2(W^{-1}) \leq \dots \leq \lambda_m(W^{-1})$ be the eigenvalues of W^{-1} and $0 < \lambda_1(BA^{-1}B^T) \leq \lambda_2(BA^{-1}B^T) \leq \dots \leq \lambda_m(BA^{-1}B^T)$ be the eigenvalues of $BA^{-1}B^T$. Then,

$$\theta_k \lambda_k(W^{-1}) \leq \lambda_k(BA^{-1}B^T), \quad k = 1, \dots, m,$$

where $\sigma_m(Y)^2 \leq \theta_k \leq \sigma_1(Y)^2$.

²An analogous transform holds for Algorithm 1.

Proof. From the QR decomposition (2.4) of B^T we have that $BA^{-1}B^T = Y^T\widetilde{W}Y$ and so $\lambda_k(BA^{-1}B^T) = \lambda_k(Y^T\widetilde{W}Y)$. By Ostrowski's theorem [10, Theorem 4.5.9], it follows that $\lambda_k(Y^T\widetilde{W}Y) = \theta_k\lambda_k(\widetilde{W})$. Combining this with the inequality in Lemma 2 gives the result. \square

Thus, the antitriangular factorization gives lower bounds on all the eigenvalues of the Schur complement $BA^{-1}B^T$. In particular, it bounds from below the smallest eigenvalue, which can be useful when bounding the eigenvalues of \mathcal{A} or when approximating inf-sup constants [23]. Note that since Y is antitriangular its singular values are relatively easy to compute.

6. The antitriangular factorization and preconditioning. When the saddle point system (4.1) is too large to be solved by a direct method, an iterative method such as a Krylov subspace method is usually applied. Unfortunately, however, these iterative methods typically converge slowly when applied to saddle point problems unless preconditioners are used. Many preconditioners for saddle point matrices have been proposed [2, section 10], [3], but we focus here on block preconditioners and show how they can be factored by the antitriangular factorization in section 2. We first discuss block diagonal and block triangular preconditioners and then describe constraint preconditioners, showing that in this latter case the same orthogonal transformation converts \mathcal{A} and \mathcal{P} to antitriangular form. We assume throughout that \mathcal{A} in (4.1) is factorized in antitriangular form (2.7), i.e., that $\mathcal{A} = Q\mathcal{M}Q^T$.

We briefly mention the block diagonal matrix

$$\mathcal{P}_D = \begin{bmatrix} T & 0 \\ 0 & V \end{bmatrix},$$

where $T \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{m \times m}$ are symmetric. Often T is chosen to approximate A and V to approximate the Schur complement $BA^{-1}B^T$. Indeed, if $T = A$ and $V = BA^{-1}B^T$, then $\mathcal{P}_D^{-1}\mathcal{A}$ has three eigenvalues, 1 and $(1 \pm \sqrt{5})/2$ [13, 18].

Applying Q in a similarity transform gives

$$Q^T\mathcal{P}_DQ = \begin{bmatrix} V & 0 & 0 \\ 0 & \widehat{T}_{22} & \widehat{T}_{12}^TS \\ 0 & S\widehat{T}_{12} & S\widehat{T}_{11}S \end{bmatrix},$$

where $\widehat{T}_{ij} = (U^{(i)})^T T U^{(j)}$, $i, j = 1, 2$. Thus, the transformed preconditioner is also block diagonal, with an $m \times m$ block followed by an $n \times n$ block. Note that since \mathcal{P}_D is positive definite, $Q^T\mathcal{P}_DQ$ cannot have significant block antidiagonal structure.

Similarly, the block lower triangular preconditioner

$$\mathcal{P}_T = \begin{bmatrix} T & 0 \\ B & V \end{bmatrix}$$

gives

$$Q^T\mathcal{P}_TQ = \begin{bmatrix} V & 0 & Y^T \\ 0 & \widehat{T}_{22} & \widehat{T}_{12}^TS \\ 0 & S\widehat{T}_{12} & S\widehat{T}_{11}S \end{bmatrix}.$$

The corresponding upper triangular preconditioner has an analogous form.

Constraint preconditioners [12, 14, 19, 20]

$$(6.1) \quad \mathcal{P}_C = \begin{bmatrix} T & B^T \\ B & 0 \end{bmatrix},$$

on the other hand, preserve the constraints of \mathcal{A} exactly but replace A by a symmetric approximation T . Precisely because the constraints are preserved,

$$(6.2) \quad \mathcal{M}_C = \mathcal{Q}^T \mathcal{P}_C \mathcal{Q} = \begin{bmatrix} 0 & 0 & Y^T \\ 0 & \hat{T}_{22} & \hat{T}_{12}S \\ Y & S\hat{T}_{12} & S\hat{T}_{11}S \end{bmatrix},$$

where $\hat{T}_{ij} = U_i^T T U_j$, $i, j = 1, 2$, is an antitriangular matrix when \hat{T}_{22} is positive definite.

It is known that $\mathcal{P}_C^{-1} \mathcal{A}$ has at least $2m$ unit eigenvalues, with the remainder being the eigenvalues λ of $(U^{(2)})^T A U^{(2)} v = \lambda (U^{(2)})^T T U^{(2)} v$ [12, 14]. (Note that we could use any basis for the nullspace of B in place of $U^{(2)}$.) Since A is positive definite on the nullspace of B , any nonunit eigenvalues are real, although negative eigenvalues will occur when $(U^{(2)})^T T U^{(2)}$ is not positive definite. These facts are also easily discerned from the antitriangular forms. Since $\mathcal{P}_C^{-1} \mathcal{A} = \mathcal{Q} \mathcal{M}_C^{-1} \mathcal{M} \mathcal{Q}^T$, which is similar to $\mathcal{M}_C^{-1} \mathcal{M}$, explicitly obtaining \mathcal{M}_C^{-1} as in (4.3) and multiplying by \mathcal{M} gives a block upper triangular matrix from which the eigenvalues are immediately obvious. Indeed, Keller, Gould, and Wathen [12] use a flipped version of (6.2) to investigate the eigenvalues of $\mathcal{P}_C^{-1} \mathcal{A}$.

The matrix $\mathcal{Q}^T \mathcal{P}_C \mathcal{Q}$ may be applied using the procedure outlined in section 4, and it makes clear the equivalence between constraint preconditioners and the nullspace method that has previously been observed [7, 14, 22]. Conversely, any matrix \mathcal{N} in antitriangular form with $Y = SR$ defines a constraint preconditioner

$$\mathcal{P}_C = \mathcal{Q} \mathcal{N} \mathcal{Q}^T = \begin{bmatrix} T & B \\ B & 0 \end{bmatrix}$$

for \mathcal{A} , with

$$T = U^{(2)} X (U^{(2)})^T + U^{(1)} S Z (U^{(2)})^T + U^{(2)} Z^T S (U^{(1)})^T + U^{(1)} S W S (U^{(1)})^T.$$

We note that alternative factorizations of constraint preconditioners, some of which rely on a basis for the nullspace of B , have also been proposed. In particular, the Schilders' factorization [4, 5, 24] reorders the variables so that $B = [B_1 \ B_2]$ with $B_1 \in \mathbb{R}^{m \times m}$ nonsingular. In this case

$$N = \begin{bmatrix} -B_1^{-1} B_2 \\ I \end{bmatrix}$$

is a basis for the nullspace and is important to the factorization.

For a true, inertia-revealing antitriangular factorization T must be positive definite on the nullspace of B , so that \hat{T}_{22} is symmetric positive definite. If we are only interested in preconditioning (and not in the inertia of \mathcal{P}_C) we only require that \hat{T}_{22} is invertible.

In summary, we see that applying the same orthogonal similarity transform that makes \mathcal{A} antitriangular to \mathcal{P}_D , \mathcal{P}_T , and \mathcal{P}_C results in preconditioners with specific structures. The antitriangular form of \mathcal{P}_C makes the equivalence between constraint preconditioners and the nullspace method clear, reveals the eigenvalues of $\mathcal{P}_C^{-1} \mathcal{A}$, and may provide other insights into the properties of constraint preconditioners.

7. Conclusions. We have considerably simplified the antitriangular factorization for symmetric indefinite matrices of Mastronardi and Van Dooren in the specific and common case of saddle point matrices. This leads to the observation that this factorization is equivalent to the well-known nullspace method. We have shown that the factorization is strongly stable and that low-rank updates to A and B can be efficiently incorporated into an existing antitriangular factorization. The blocks X , Z , and Y can be used to obtain bounds on the eigenvalues of A and the Schur complement. Additionally, we have considered the form of this antitriangular factorization for popular constraint preconditioning and block diagonal and block triangular preconditioning, showing how specific structures are preserved.

Acknowledgments. The authors would like to thank Gil Strang for inspiring this work and for valuable suggestions. We are also grateful to David Titley-Peloquin and the editor and referees for their helpful comments and advice that greatly improved the manuscript.

REFERENCES

- [1] S. BELLAVIA, V. DE SIMONE, D. DI SERAFINO, AND B. MORINI, *Updating Constraint Preconditioners for KKT Systems in Quadratic Programming via Low-Rank Corrections*, arXiv: 1312.0047, 2013.
- [2] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.
- [3] M. BENZI AND A. J. WATHEN, *Some preconditioning techniques for saddle point problems*, in Model Order Reduction: Theory, Research Aspects and Applications, W. H. A. Schilders, H. A. van der Vorst, and J. Rommes, eds., Math. Ind. 13, Springer-Verlag, Berlin, 2008, pp. 195–211.
- [4] 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.
- [5] H. S. DOLLAR AND A. J. WATHEN, *Approximate factorization constraint preconditioners for saddle-point problems*, SIAM J. Sci. Comput., 27 (2006), pp. 1555–1572.
- [6] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
- [7] 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.
- [8] A. GRIEWANK, A. WALTHER, AND M. KORZEC, *Maintaining factorized KKT systems subject to rank-one updates of Hessians and Jacobians*, Optim. Methods Softw., 22 (2007), pp. 279–295.
- [9] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, 2002.
- [10] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1990.
- [11] V. E. HOWLE AND S. A. VAVASIS, *An iterative method for solving complex-symmetric systems arising in electrical power modeling*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 1150–1178.
- [12] 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.
- [13] YU. A. KUZNETSOV, *Efficient iterative solvers for elliptic finite element problems on nonmatching grids*, Russian J. Numer. Anal. Math. Modelling, 10 (1995), pp. 187–211.
- [14] 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.
- [15] H. MAHAWAR AND V. SARIN, *Parallel iterative methods for dense linear systems in inductance extraction*, Parallel Computing, 29 (2003), pp. 1219–1235.
- [16] N. MASTRONARDI AND P. VAN DOOREN, *An algorithm for solving the indefinite least squares problem with equality constraints*, BIT, 54 (2014), pp. 201–218.

- [17] N. MASTRONARDI AND P. VAN DOOREN, *The antitriangular factorization of symmetric matrices*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 173–196.
- [18] 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.
- [19] B. A. MURTAGH AND M. A. SAUNDERS, *Large-scale linearly constrained optimization*, Math. Program., 14 (1978), pp. 41–72.
- [20] B. A. MURTAGH AND M. A. SAUNDERS, *A projected Lagrangian algorithm and its implementation for sparse nonlinear constraints*, Math. Program. Stud., 16 (1982), pp. 84–117.
- [21] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer, New York, 1999.
- [22] I. PERUGIA, V. SIMONCINI, AND M. ARIOLI, *Linear algebra methods in a mixed approximation of magnetostatic problems*, SIAM J. Sci. Comput., 21 (1999), pp. 1085–1101.
- [23] J. PESTANA AND A. J. WATHEN, *Natural Preconditioners for Saddle Point Problems*, Tech. report 1754, Mathematical Institute, University of Oxford, 2013.
- [24] W. H. A. SCHILDERS, *A Preconditioning Technique for Indefinite Systems*, Tech. report RANA00-18, Technische Universiteit Eindhoven, 2000, <http://repository.tue.nl/544475>.
- [25] J.-G. SUN, *Structured backward errors for KKT systems*, Linear Algebra Appl., 288 (1999), pp. 75–88.