Electronic Transactions on Numerical Analysis.
Volume 11, pp. 85-93, 2000.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

CHOLESKY-LIKE FACTORIZATIONS OF SKEW-SYMMETRIC MATRICES*

PETER BENNER†, RALPH BYERS‡, HEIKE FASSBENDER§, VOLKER MEHRMANN¶, AND DAVID WATKINS ‖

Abstract. Every real skew-symmetric matrix B admits Cholesky-like factorizations $B = R^T JR$, where $J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$. This paper presents a backward-stable $\mathcal{O}(n^3)$ process for computing such a decomposition, in which R is a permuted triangular matrix. Decompositions of this type are a key ingredient of algorithms for solving eigenvalue problems with Hamiltonian structure.

Keywords: skew-symmetric matrices, matrix factorizations, Hamiltonian eigenproblems, complete pivoting.

AMS(MOS) subject classification: 15A23, 65F05.

1. Introduction.

Let $B \in \mathbb{R}^{2n,2n}$ be a skew-symmetric matrix. We will study decompositions of the form

$$B = R^T JR, \tag{1.1}$$

where

$$J = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}. \tag{1.2}$$

As will be shown in this paper, there always exists an R such that (1.1) holds. We present a stable $\mathcal{O}(n^3)$ algorithm that computes an R that has the form of a permuted triangular matrix.

Our motivation comes from eigenvalue problems with Hamiltonian structure. A matrix $H \in \mathbb{R}^{2n,2n}$ is said to be *Hamiltonian* if $(JH)^T = JH$ and *skew-Hamiltonian* if $(JH)^T = -JH$.

EXAMPLE 1. The study of corner singularities in anisotropic elastic materials [5, 6, 11, 9] leads to generalized eigenvalue problems of the form

$$\left(\begin{bmatrix} 0 & M \\ -K & 0 \end{bmatrix} - \lambda \begin{bmatrix} M & 0 \\ G & M \end{bmatrix} \right) \begin{bmatrix} v \\ w \end{bmatrix} = 0,$$

where $M = M^T \in \mathbb{R}^{n,n}$, $K = K^T \in \mathbb{R}^{n,n}$ and $G = -G^T \in \mathbb{R}^{n,n}$. All three matrices are large and sparse, as they are obtained from a finite element discretization. M is a positive definite mass matrix, and K is a negative definite matrix related to the stiffness matrix. In this pencil the first matrix is Hamiltonian and the second is skew-Hamiltonian. If we multiply the pencil by J on the left, we obtain the equivalent pencil

$$\left(\begin{bmatrix} -K & 0 \\ 0 & -M \end{bmatrix} - \lambda \begin{bmatrix} G & M \\ -M & 0 \end{bmatrix} \right) \begin{bmatrix} v \\ w \end{bmatrix} = 0. \tag{1.3}$$

*Received September 21, 2000. Accepted for publication October 20, 2000. Recommended by G. W. Stewart.

† Zentrum für Technomathematik, Universität Bremen, D-28334 Bremen, Germany.

‡ 405 Snow Hall, Department of Mathematics, University of Kansas, Lawrence, Kansas 66045, USA. This author was partially supported by National Science Foundation awards CCR-9732671, MRI-9977352, by the NSF EPSCoR/K*STAR program through the Center for Advanced Scientific Computing and through Sonderforschungsbereich 393, "Numerische Simulation auf massiv parallelen Rechnern".

§ Zentrum Mathematik, Technische Universität München, D-80290 München, Germany.

¶ Fakultät für Mathematik, Technische Universität Chemnitz, D-09107 Chemnitz, Germany. Partially supported through Sonderforschungsbereich 393, "Numerische Simulation auf massiv parallelen Rechnern"

‖ Department of Mathematics, Washington State University, Pullman, WA, 99164-3113, USA. Partially supported through Sonderforschungsbereich 393, "Numerische Simulation auf massiv parallelen Rechnern"

ETNA
Kent State University
etna@mcs.kent.edu

Now the first matrix is symmetric and the second is skew-symmetric.

EXAMPLE 2. Linear quadratic optimal control problems for descriptor systems [8, 1] lead to eigenvalue problems of the form

$$\left(\begin{bmatrix} A & BB^T \\ C^TC & -A^T \end{bmatrix} - \lambda \begin{bmatrix} E & 0 \\ 0 & E^T \end{bmatrix}\right)\begin{bmatrix} v \\ w \end{bmatrix} = 0. \tag{1.4}$$

Again the first matrix is Hamiltonian and the second is skew-Hamiltonian, and again we can multiply by J to obtain an equivalent pencil

$$\left(\begin{bmatrix} C^TC & -A^T \\ -A & -BB^T \end{bmatrix} - \lambda \begin{bmatrix} 0 & E^T \\ -E & 0 \end{bmatrix}\right)\begin{bmatrix} v \\ w \end{bmatrix} = 0, \tag{1.5}$$

in which the first matrix is symmetric and the second is skew-symmetric.

Consider a generalized eigenvalue problem of the form $(A - \lambda B)v = 0$, where A is symmetric and B is skew-symmetric and nonsingular. We can often facilitate solution of this problem by transforming it to a standard eigenvalue problem $(H - \lambda I)z = 0$, in which H is a Hamiltonian matrix. Suppose we can factor B into a product $B = R^TJR$, as in (1.1). Then we can transform the pencil $A - \lambda B$ to $R^{-T}AR^{-1} - \lambda J$, in which the coefficient matrices are still symmetric and skew-symmetric, respectively. If we then premultiply by $J^{-1} = J^T$, we obtain $J^TR^{-T}AR^{-1} - \lambda I$, in which the coefficient matrix $H = J^TR^{-T}AR^{-1}$ is Hamiltonian. Equivalently, using $L = JA$ and $N = JB$ for some symmetric A and skew-symmetric B, we may consider a Hamiltonian/skew-Hamiltonian eigenproblem $(L-\lambda N)v = 0$ and use a factorization J^TR^TJR of the skew-Hamiltonian matrix N in order to transform this generalized eigenproblem to a standard Hamiltonian eigenproblem.

Transformations of this type were exploited in [1] and [9] to yield new structure-preserving algorithms for the computation of eigenvalues, eigenvectors, and deflating subspaces of Hamiltonian/skew-Hamiltonian pencils. The numerically stable factorization presented in this paper extends the algorithms proposed in [1] and [9] to all Hamiltonian/skew-Hamiltonian pencils. It is an important point that none of the methods proposed in [1, 9] requires explicit formation of the Hamiltonian matrix $H = J^TR^{-T}AR^{-1}$.

2. Skew-Symmetric Cholesky-like Factorizations.

In some cases a useful factorization of the form $B = R^TJR$ can be found by inspection.

EXAMPLE 3. Consider the matrix

$$B = \begin{bmatrix} 0 & E^T \\ -E & 0 \end{bmatrix}$$

from (1.5). By inspection

$$B = \begin{bmatrix} 0 & I \\ -E & 0 \end{bmatrix}\begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}\begin{bmatrix} 0 & E^T \\ -I & 0 \end{bmatrix} = R^TJR.$$

If E is nonsingular, then this is a useful factorization [1].

EXAMPLE 4. Consider the matrix

$$B = \begin{bmatrix} G & M \\ -M & 0 \end{bmatrix}$$

from (1.3). We have

$$B = \begin{bmatrix} I & -\frac{1}{2}G \\ 0 & M \end{bmatrix}\begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}\begin{bmatrix} I & 0 \\ \frac{1}{2}G & M \end{bmatrix} = R^TJR.$$

ETNA
Kent State University
etna@mcs.kent.edu

This is essentially the decomposition that was used in [9]. With this decomposition, the Hamiltonian matrix corresponding to the pencil $A - \lambda B$ of (1.3) is

$$\begin{aligned} H &= J^T R^{-T} A R^{-1} \\ &= \begin{bmatrix} 0 & -I \\ I & 0 \end{bmatrix} \begin{bmatrix} I & -\frac{1}{2}GM^{-1} \\ 0 & M^{-1} \end{bmatrix} \begin{bmatrix} -K & 0 \\ 0 & -M \end{bmatrix} \begin{bmatrix} I & 0 \\ -\frac{1}{2}M^{-1}G & M^{-1} \end{bmatrix} \\ &= \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix} \begin{bmatrix} I & -\frac{1}{2}G \\ 0 & I \end{bmatrix} \begin{bmatrix} K & 0 \\ 0 & M^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -\frac{1}{2}G & I \end{bmatrix}. \end{aligned}$$

It is relatively easy to see that any nonsingular skew-symmetric matrix has a decomposition of the form (1.1).

PROPOSITION 2.1. *Let $B \in \mathbb{R}^{m,m}$ be skew-symmetric. If B is nonsingular, then there exists $R \in \mathbb{R}^{m,m}$ such that B can be factored as $B = R^T J R$.*

Proof. This is an easy consequence of the spectral decomposition $B = Q^T X Q$ (Murnaghan canonical form), see e.g. [10], in which Q is orthogonal and X is block diagonal. Each block on the main diagonal is 2×2 and has the form

$$\begin{bmatrix} 0 & \delta_i \\ -\delta_i & 0 \end{bmatrix},$$

with $\delta_i > 0$, where $\pm\imath\delta_i$ is a pair of complex conjugate (purely imaginary) eigenvalues of B. If we let $D = \mathrm{diag}\{\sqrt{\delta_1}, \sqrt{\delta_1}, \ldots, \sqrt{\delta_n}, \sqrt{\delta_n}\}$, then $X = D\hat{J}D = D^T\hat{J}D$, where $\hat{J}$ is the block diagonal matrix with 2×2 blocks

$$\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

running down the main diagonal. Clearly, $\hat{J}$ is permutationally similar to J. Indeed,

$$\hat{J} = PJP^T,$$

where P is the perfect-shuffle permutation

$$P = \begin{bmatrix} e_1, e_3, \ldots, e_{2n-1}, e_2, e_4, \ldots, e_{2n} \end{bmatrix}, \tag{2.1}$$

with e_j the jth unit vector in $\mathbb{R}^{2n}$. Combining these factorizations, we obtain $B = R^T J R$, where $R = P^T D Q$. ☐

Not only have we shown the existence of a decomposition $B = R^T J R$, we have sketched a practical (and stable) algorithm for computing one. The only nontrivial step is the computation of the Murnaghan form. Efficient methods for computing this form are given in [10, 12]. The decomposition is also usable in practice, since R^{-1} is easily accessible; each of its factors P^T, D, and Q, has an easily computable inverse. However, this factorization has substantial drawbacks. The iterative algorithms needed to compute the Murnaghan form use substantially more arithmetic work than the direct factorization algorithm presented below. Furthermore, and perhaps more importantly, the Murnaghan decomposition does not preserve sparseness. Q is essentially a matrix of eigenvectors of B. Even if B is extremely sparse (as, e.g., for problems as in Example 1), Q will be full.

2.1. Triangular Factorizations. For sparse problems we require a different factorization. Chances for preserving sparseness are better if R is obtained through a process similar to Gaussian elimination, which leads to a triangular R. This technique has proven successful in computing sparse LU and Cholesky factorizations; see, e.g, [2, 3]. We begin with a theorem

ETNA
Kent State University
etna@mcs.kent.edu

that tells what can be done without pivoting. Of greatest interest to us is the even-dimensional case. However, we include the odd-dimensional case for completeness.

For notational convenience we will produce decompositions of the form $B = R^T \hat{J} R$, using the shuffled matrix $\hat{J}$ instead of J. We then obtain the desired results by deshuffling.

We will use the following notation. When we want to emphasize the dimension of $\hat{J}$, we write $\hat{J}_{2n}$. Thus $\hat{J}_{2n} = P J_{2n} P^T$, where P is as in (2.1) and

$$J_{2n} = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}.$$

In addition we define

$$\hat{J}_{2n+1} = \begin{bmatrix} \hat{J}_{2n} & 0 \\ 0 & 0 \end{bmatrix}.$$

Finally $B[< i >]$ will denote the ith leading principal submatrix of the matrix B.

THEOREM 2.2.

i) Let $B \in \mathbb{R}^{2n,2n}$ be a skew-symmetric matrix such that $\det B[< 2j >] \neq 0$ for $j = 1, \ldots, n$. Then B has a unique factorization

$$B = R^T \hat{J}_{2n} R, \tag{2.2}$$

where R is upper triangular with $r_{2j-1,2j} = 0$, $r_{2j-1,2j-1} > 0$ and $r_{2j,2j} = \pm r_{2j-1,2j-1}$ for $j = 1, \ldots, n$. Thus R has 2×2 blocks of the form

$$\begin{bmatrix} r & 0 \\ 0 & \pm r \end{bmatrix}$$

running down the main diagonal.

ii) Let $B \in \mathbb{R}^{(2n+1),(2n+1)}$ be a skew-symmetric matrix such that $\det B[< 2j >] \neq 0$ for $j = 1, \ldots, n$. Then B has a unique factorization

$$B = R^T \hat{J}_{2n+1} R, \tag{2.3}$$

with R as in part i), along with the additional condition $r_{2n+1,2n+1} = 0$.

Proof. We begin by proving part i) by induction on n. If $n = 1$ we have $B = \left[\begin{smallmatrix} 0 & v \\ -v & 0 \end{smallmatrix}\right]$, where $v \neq 0$ by the assumption $\det B[< 2 >] \neq 0$. From the equation

$$\begin{bmatrix} r & 0 \\ 0 & s \end{bmatrix} \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} r & 0 \\ 0 & s \end{bmatrix} = \begin{bmatrix} 0 & rs \\ -rs & 0 \end{bmatrix}$$

we see that we need to choose r and s so that $rs = v$. Since we want $r > 0$ and $s = \pm r$, we must choose r and s as follows. If $v > 0$, take $r = +\sqrt{v}$ and $s = r$; if $v < 0$, take $r = +\sqrt{-v}$ and $s = -r$. This establishes existence and uniqueness of the factorization in the case $n = 1$.

For the induction step partition B as

$$B = \begin{bmatrix} B_{11} & B_{12} \\ -B_{12}^T & B_{22} \end{bmatrix},$$

where $B_{11} \in \mathbb{R}^{2,2}$. As we have just demonstrated, B_{11} has a unique decomposition $B_{11} = R_{11}^T \hat{J}_2 R_{11}$, where $R_{11} = \text{diag}\{r, \pm r\}$. Let $R_{12} = \hat{J}_2^T R_{11}^{-T} B_{12}$ and $S = B_{22} - R_{12}^T \hat{J}_2 R_{12}$. Then

$$B = \begin{bmatrix} B_{11} & B_{12} \\ -B_{12}^T & B_{22} \end{bmatrix} = \begin{bmatrix} R_{11}^T & 0 \\ R_{12}^T & I \end{bmatrix} \begin{bmatrix} \hat{J}_2 & \\ & S \end{bmatrix} \begin{bmatrix} R_{11} & R_{12} \\ 0 & I \end{bmatrix}. \tag{2.4}$$

ETNA
Kent State University
etna@mcs.kent.edu

The Schur complement $S \in \mathbb{R}^{2n-2,2n-2}$ is skew-symmetric. It also satisfies $\det S[<2j>] \neq 0$ for $j = 1, \ldots, n-1$, inheriting the property from B. This is an easy consequence of (2.4). Therefore, by the induction hypothesis, there is a unique upper triangular R_{22} of the specified form, such that $S = R_{22}^T \hat{J}_{2n-2} R_{22}$. Substituting this expression for S into (2.4), we obtain

$$B = \begin{bmatrix} B_{11} & B_{12} \\ -B_{12}^T & B_{22} \end{bmatrix} = \begin{bmatrix} R_{11}^T & 0 \\ R_{12}^T & R_{22}^T \end{bmatrix} \begin{bmatrix} \hat{J}_2 & \\ & \hat{J}_{2n-2} \end{bmatrix} \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix}, \tag{2.5}$$

which is the desired decomposition.

We easily check uniqueness. Equation (2.5) implies $B_{12} = R_{11}^T \hat{J}_2 R_{12}$, which determines R_{12} uniquely, since $R_{11}^T \hat{J}_2$ is nonsingular. Equation (2.5) also implies $B_{22} = R_{12}^T \hat{J}_2 R_{12} + R_{22}^T \hat{J}_{2n-2} R_{22}$, which forces $R_{22}^T \hat{J}_{2n-2} R_{22} = S$. As part of the induction hypothesis, R_{22} is unique.

Now consider the odd-dimensional case ii). Suppose $B \in \mathbb{R}^{m,m}$, where $m = 2n+1$. The case $m = 1$ is trivial. If $m > 1$, we deduce the result from the even case $m-1$. Partition B as

$$B = \begin{bmatrix} \tilde{B} & v \\ -v^T & 0 \end{bmatrix},$$

where $\tilde{B} \in \mathbb{R}^{2n,2n}$. Since we have already established the even case, $\tilde{B}$ has a unique decomposition $\tilde{B} = \tilde{R}^T \hat{J}_{2n} \tilde{R}$ of the specified form. Let $w = \hat{J}_{2n}^T \tilde{R}^{-T} v$. Then

$$B = \begin{bmatrix} \tilde{B} & v \\ -v^T & 0 \end{bmatrix} = \begin{bmatrix} \tilde{R}^T & 0 \\ w^T & 0 \end{bmatrix} \begin{bmatrix} \hat{J}_{2n} & \\ & 0 \end{bmatrix} \begin{bmatrix} \tilde{R} & w \\ 0 & 0 \end{bmatrix},$$

which is the desired decomposition. One easily checks that this equation determines $\tilde{R}$ and w uniquely. □

In the odd-dimensional case, the condition $r_{2n+1,2n+1} = 0$ was specified arbitrarily in order to obtain uniqueness of the factorization. In fact, $r_{2n+1,2n+1}$ can be given any value.

2.2. Factorization with Complete Pivoting. If we allow pivoting, we can drop the assumptions on the nonsingularity of principal submatrices. To preserve skew-symmetry in a skew-symmetric B we only consider symmetric permutations of B, that is $B \leftarrow PBP^T$, where P is a permutation matrix. Row permutations or column permutations alone destroy skew-symmetry.

THEOREM 2.3. *Let $B \in \mathbb{R}^{m,m}$ be skew-symmetric and let* $\operatorname{rank}(B) = 2s$. *Then there exists a permutation matrix Q such that*

$$B = Q^T R^T \hat{J} R Q, \tag{2.6}$$

where R is upper triangular, $r_{2j-1,2j} = 0$, $r_{2j-1,2j-1} = r_{2j,2j} > 0$ for $j = 1, \ldots, s$, $|r_{jk}| \leq r_{jj}$ for $k > j$, $j = 1, \ldots, 2s$, and $r_{jk} = 0$ for $j > 2s$. In general Q is not uniquely determined. Once Q has been fixed, R is uniquely determined.

Proof. The proof is by induction on s. If $s = 0$, the correct decomposition is $B = 0^T \hat{J} 0$. Now suppose $s > 0$, and assume the result holds for matrices of rank $2(s-1)$. We use a complete pivoting strategy. Search B for its largest (positive) entry. Suppose this is in position (i, j), noting that $i \neq j$. If $i \neq 1$, interchange rows 1 and i. The largest entry is now in position $(1, j)$. Now interchange columns 1 and i. If $j = 1$, the largest entry is now in position $(1, i)$; let $k = i$. If $j \neq 1$, then since $i \neq j$, the largest entry remains in position $(1, j)$; let $k = j$. Now if $k \neq 2$, interchange columns 2 and k and also rows 2 and k. The result is a skew-symmetric matrix $\tilde{B} = Q_1^T B Q_1$, whose largest entry is in the $(1,2)$ position. Q_1 is a permutation matrix.

Partition $\tilde{B}$ as

$$\tilde{B} = \begin{bmatrix} \tilde{B}_{11} & \tilde{B}_{12} \\ -\tilde{B}_{12}^T & \tilde{B}_{22} \end{bmatrix},$$

where $\tilde{B}_{11} \in \mathbb{R}^{2,2}$. $\tilde{B}_{11} = \left[\begin{smallmatrix} 0 & v \\ -v & 0 \end{smallmatrix}\right]$, where v is the largest (positive) entry of $\tilde{B}$. Let $r = \sqrt{v}$ and $R_{11} = \left[\begin{smallmatrix} r & 0 \\ 0 & r \end{smallmatrix}\right]$. Then $\tilde{B}_{11} = R_{11}^T \hat{J}_2 R_{11}$. Let $\tilde{R}_{12} = \hat{J}_2^T R_{11}^{-T} \tilde{B}_{12}$ and $S = \tilde{B}_{22} - \tilde{R}_{12}^T \hat{J}_2 \tilde{R}_{12}$. Then

$$\tilde{B} = \begin{bmatrix} \tilde{B}_{11} & \tilde{B}_{12} \\ -\tilde{B}_{12}^T & \tilde{B}_{22} \end{bmatrix} = \begin{bmatrix} R_{11}^T & 0 \\ \tilde{R}_{12}^T & I \end{bmatrix} \begin{bmatrix} \hat{J}_2 & \\ & S \end{bmatrix} \begin{bmatrix} R_{11} & \tilde{R}_{12} \\ 0 & I \end{bmatrix}. \tag{2.7}$$

From the definition of $\tilde{R}_{12}$, we see that each of its entries has the form $\tilde{b}/r$, where $\tilde{b}$ is an entry of $\tilde{B}_{12}$. Since $|\tilde{b}|/r \leq v/r = r$, we conclude that all entries in the first two rows are bounded in modulus by r.

The Schur complement S is skew-symmetric and has rank $2(s-1)$. Therefore, by the induction hypothesis, $S = \tilde{Q}_2^T R_{22}^T \hat{J}_{m-2} R_{22} \tilde{Q}_2$, where $\tilde{Q}_2$ is a permutation matrix, and R_{22} is an upper-triangular matrix satisfying the hypotheses of the theorem. Substituting this expression for S in (2.7), we obtain

$$\begin{aligned} \tilde{B} &= \begin{bmatrix} R_{11}^T & 0 \\ \tilde{R}_{12}^T & \tilde{Q}_2^T R_{22}^T \end{bmatrix} \begin{bmatrix} \hat{J}_2 & \\ & \hat{J}_{m-2} \end{bmatrix} \begin{bmatrix} R_{11} & \tilde{R}_{12} \\ 0 & R_{22}\tilde{Q}_2 \end{bmatrix} \\ &= \begin{bmatrix} I_2 & \\ & \tilde{Q}_2^T \end{bmatrix} \begin{bmatrix} R_{11}^T & 0 \\ \tilde{Q}_2 \tilde{R}_{12}^T & R_{22}^T \end{bmatrix} \hat{J}_m \begin{bmatrix} R_{11} & \tilde{R}_{12}\tilde{Q}_2^T \\ 0 & R_{22} \end{bmatrix} \begin{bmatrix} I_2 & \\ & \tilde{Q}_2 \end{bmatrix}. \end{aligned}$$

Letting $R_{12} = \tilde{R}_{12}\tilde{Q}_2^T$,

$$R = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix}, \quad Q_2 = \begin{bmatrix} I_2 & \\ & \tilde{Q}_2 \end{bmatrix},$$

and $Q = Q_2 Q_1$, we have $\tilde{B} = Q_1 B Q_1^T = Q_2^T R^T \hat{J} R Q_2$, i.e. $B = Q^T R^T \hat{J} R Q$. $\square$

The proof of Theorem 2.3 is constructive; it yields an algorithm, which can be sketched

ETNA
Kent State University
etna@mcs.kent.edu

as follows. The algorithm writes R over B.

```
for j = 1, ..., ⌊m/2⌋
  ⎡ find (ii, jj) such that B[ii, jj] = max B[2j − 1 : m, 2j − 1 : m]
  ⎢ if B[ii, jj] = 0, then
  ⎢   ⎡ rank = 2j − 2
  ⎢   ⎣ return
  ⎢ if jj = 2j − 1, then kk ← ii, else kk ← jj
  ⎢ if ii ≠ 2j − 1, then interchange rows, columns ii and 2j − 1
  ⎢ if kk ≠ 2j, then interchange rows, columns kk and 2j
  ⎢ r ← +√(B[2j − 1, 2j])
  ⎢ B[2j − 1, 2j − 1] ← r;
  ⎢ B[2j, 2j] ← r;
  ⎢ B[2j − 1, 2j] ← 0
  ⎢ B[2j − 1 : 2j, 2j + 1 : m] ← −r^{-1} Ĵ_2 B[2j − 1 : 2j, 2j + 1 : m]
  ⎢ B̂ ← B[2j − 1 : 2j, 2j + 1 : m]^T Ĵ_2 B[2j − 1 : 2j, 2j + 1 : m]
  ⎣ B[2j + 1 : m, 2j + 1 : m] ← B[2j + 1 : m, 2j + 1 : m] − B̂
end for
rank = 2⌊m/2⌋
return
```

To this brief description we add a few details. By symmetry we need only store and operate on the upper half of B. In particular, only about half of the operations indicated in the construction and application of the update $\hat{B}$ (bottom two lines in the loop) need to be performed. The permutation matrix Q is determined by keeping a record of the column interchanges. Each column interchange should be applied to all rows of the array, not just the current Schur complement $B[2j-1:m, 2j-1:m]$. This way the update $R_{12} = \tilde{R}_{12}\tilde{Q}_2^T$, indicated in the proof of Theorem 2.3, is performed automatically.

The flop count, for a dense B, is about $\frac{1}{3}(m^3-(m-2s)^3)$, i.e. $\frac{1}{3}m^3$ in the high-rank case and $2sm^2$ in the low-rank case. In addition the pivot searches make about $\frac{1}{12}m^3$ comparisons in all.

Less expensive pivoting strategies are possible. For example, the pivot search could be confined to the top two rows, unless no suitable pivot is found there. One can also consider pivoting strategies that take sparseness into account, compromising on stability in order to obtain a sparser factor R. Incomplete factorizations can also be obtained from the given algorithm analogously to incomplete LU or Cholesky decompositions by prescribing a level of allowed fill-in or using a drop tolerance.

The desired factorization of B is an easy corollary of Theorem 2.3.

COROLLARY 2.4. *Let $B \in \mathbb{R}^{2n,2n}$ be skew-symmetric. Then*

$$B = \mathcal{R}^T J \mathcal{R},$$

where $\mathcal{R} = P^T R Q$, P is the perfect-shuffle permutation (2.1), R is an upper-triangular matrix satisfying the hypotheses of Theorem 2.3, and Q is a permutation matrix.

Analogously, we obtain a result for skew-Hamiltonian matrices using the isomorphism $N \to JN$ which maps skew-Hamiltonian matrices to skew-symmetric matrices and vice versa.

COROLLARY 2.5. *Let $N \in \mathbb{R}^{2n,2n}$ be skew-Hamiltonian. Then*

$$N = J^T \mathcal{R}^T J \mathcal{R},$$

where $\mathcal{R}$ is the factor of JN as in Corollary 2.4.

ETNA
Kent State University
etna@mcs.kent.edu

3. Stability of the Factorization. Standard techniques of backward error analysis can be applied to obtain the following result (cf. [4, Theorem 9.3]).

THEOREM 3.1. *Let B be skew-symmetric with rank $2s$. Suppose we compute the decomposition $B = Q^T R^T \hat{J} R Q$ using floating point arithmetic with unit roundoff u. Then the computed R satisfies $B + E = Q^T R^T \hat{J} R Q$, where*

$$|E| \leq 2suQ^T\,|R^T|\,|\hat{J}|\,|R|\,Q + \mathcal{O}(u^2).$$

This result is valid for the unpivoted algorithm, but it therefore holds for all pivoting strategies, since pivoting is equivalent to applying the unpivoted algorithm to a matrix for which the interchanges have been performed in advance. So long as only modest element growth occurs in the computation of R, the algorithm is backward stable.

The objective of the complete pivoting strategy is to discourage element growth. There is both theoretical and empirical evidence that it succeeds. We begin with a theoretical bound. Let $B^{(0)}$ be the skew-symmetric matrix B at the beginning of the above algorithm, and let $B^{(r)}$ be the Schur complement matrix $B[2r+1:m, 2r+1:m]$ at the end of the rth pass through the outer loop. Define the *element growth factor* γ by

$$\gamma = \max_r \frac{\max_{ij} |B_{ij}^{(r)}|}{\max_{ij} |B_{ij}^{(0)}|}.$$

It is easy to show that $\max_{ij} |r_{ij}| \leq \sqrt{\gamma}$. Wilkinson's classical analysis of the growth factor for Gaussian elimination with complete pivoting (GCP) [13] carries over with only trivial modifications to the above algorithm. It can be shown that for a nonsingular $2n \times 2n$ skew-symmetric matrix B,

$$\gamma \leq \sqrt{(2n)4^1 6^{1/2} 8^{1/3} \cdots (2n)^{1/(n-1)}}. \tag{3.1}$$

This bound has the same order of magnitude as Wilkinson's classical bound on the element growth factor for GCP [13].

Mountains of numerical evidence accumulated over the years have shown that the GCP bound is quite pessimistic. For example, for matrices of dimension $n = 100$ and $n = 200$ the bound guarantees that the growth factor is no greater than 3571 and 28298, respectively. However, numerical experience has shown that the actual growth factor is usually much less than n. Indeed, it is difficult to find matrices for which the growth factor is much greater than n. See, for example, the discussion in [4, pp. 180–181]. On the basis of this evidence, GCP has been pronounced stable.

Since the complete pivoting process in the above algorithm is very similar to GCP, we expected that the bound (3.1) would also prove pessimistic. This expectation has been realized. We implemented the above algorithm in MATLAB [7] and found that when applied to random matrices, the growth factor is modest. We then used `fminsearch` from MATLAB's optimization toolbox to search for skew-symmetric matrices whose growth factor is large. For each of three choices of m, we called `fminsearch` fifty times. Each call to `fminsearch` started with an $m \times m$ skew-symmetric matrix with nontrivial entries drawn from the normal distribution with mean zero and variance one. For $m = 2n = 10$, the bound (3.1) is 18.7, but the largest element growth factor that `fminsearch` found was 3. For $m = 16$ the bound (3.1) is 46.0, but the largest element growth factor that we observed was 3.36. For $m = 20$ the bound (3.1) is 72.8, but the largest element growth factor that we observed was 4.37. Based on these findings, as well as the evidence that has been accumulated for GCP, we assert that the above algorithm is stable.

ETNA
Kent State University
etna@mcs.kent.edu

4. Conclusions. Every real skew-symmetric matrix B admits Cholesky-like factorizations $B = R^T JR$, where $J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$. Factorizations of this type are a key ingredient of algorithms for solving eigenvalue problems with Hamiltonian structure. Factorizations in which R is a permuted triangular matrix can be computed by an $\mathcal{O}(n^3)$ process similar to Gaussian elimination. If complete pivoting is used, the process is numerically backward stable.

REFERENCES

[1] P. BENNER, R. BYERS, V. MEHRMANN, AND H. XU, *Numerical computation of deflating subspaces of embedded Hamiltonian pencils*, Tech. Rep. 99-15, Sonderforschungsbereich 393, Numerische Simulation auf massiv parallelen Rechnern, Fakultät für Mathematik, TU Chemnitz, 09107 Chemnitz, Germany, 1999. Available online from `http://www.tu-chemnitz.de/sfb393/sfb99pr.html`.

[2] I. S. DUFF, A. M. ERISMAN, AND J. K. REID, *Direct methods for sparse matrices*, The Clarendon Press, Oxford University Press, New York, second ed., 1989. Oxford Science Publications.

[3] A. GEORGE AND J. W. H. LIU, *Computer solution of large sparse positive definite systems*, Prentice-Hall Inc., Englewood Cliffs, N.J., 1981. Prentice-Hall Series in Computational Mathematics.

[4] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM Publications, Philadelphia, PA, 1996.

[5] V. KOZLOV, V. MAZ'YA, AND J. ROSSMANN, *Spectral properties of operator pencils generated by elliptic boundary value problems for the Lamé system*, Rostocker Math. Kolloq., 51 (1997), pp. 5–24.

[6] D. LEGUILLON, *Computation of 3d-singularities in elasticity*, in Boundary value problems and integral equations in nonsmooth domains, M. Costabel et al., ed., vol. 167 of Lecture Notes in Pure and Appl. Math., New York, 1995, Marcel Dekker, pp. 161–170. Proceedings of the conference, held at the CIRM, Luminy, France, May 3-7, 1993.

[7] *MATLAB, Version 5*. The MathWorks, inc., 24 Prime Park Way, Natick, MA 01760-1500, USA, 1996.

[8] V. MEHRMANN, *The Autonomous Linear Quadratic Control Problem, Theory and Numerical Solution*, no. 163 in Lecture Notes in Control and Inform. Sci., Springer-Verlag, Heidelberg, July 1991.

[9] V. MEHRMANN AND D. WATKINS, *Structure-preserving methods for computing eigenpairs of large sparse skew-Hamiltoninan/Hamiltonian pencils*, Tech. Rep. SFB393/00-02, Fakultät für Mathematik, TU Chemnitz, 09107 Chemnitz, FRG, 2000. Available online from `http://www.tu-chemnitz.de/sfb393/sfb00pr.html`.

[10] M. PAARDEKOOPER, *An eigenvalue algorithm for skew-symmetric matrices*, Numer. Math., 17 (1971), pp. 189–202.

[11] H. SCHMITZ, K. VOLK, AND W. L. WENDLAND, *On three-dimensional singularities of elastic fields near vertices*, Numer. Methods Partial Differential Equations, 9 (1993), pp. 323–337.

[12] R. WARD AND L. GRAY, *Eigensystem computation for skew-symmetric and a class of symmetric matrices*, ACM Trans. Math. Software, 4 (1978), pp. 278–285.

[13] J. WILKINSON, *Error analysis of direct methods of matrix inversion*, J. Assoc. Comput. Mach., 8 (1961), pp. 281–330.