Electronic Transactions on Numerical Analysis.
Volume 55, pp. 455–468, 2022.

ISSN 1068–9613.
DOI: 10.1553/etna_vol55s455

ETNA
Kent State University and
Johann Radon Institute (RICAM)

STRUCTURED SHIFTS FOR SKEW-SYMMETRIC MATRICES*

CHEN GREIF†

Abstract. We consider the use of a skew-symmetric block-diagonal matrix as a structured shift. Properties of Hamiltonian and skew-Hamiltonian matrices are used to show that the shift can be effectively used in the iterative solution of skew-symmetric linear systems or nonsymmetric linear systems with a dominant skew-symmetric part. Eigenvalue analysis and some numerical experiments confirm our observations.

Key words. skew-symmetric matrix, structured shift, Hamiltonian matrix, skew-Hamiltonian matrix, eigenvalue analysis, iterative solution of linear systems

AMS subject classifications. 65F08, 65F10, 65F50, 15A12, 15B57

1. Introduction.

The increasing importance of the numerical solution of problems related to Hamiltonian and port-Hamilitonian systems [5, 27] generates a potential interest in designing faster numerical solvers for skew-symmetric linear systems and nonsymmetric linear systems with a dominant skew-symmetric part. A real skew-symmetric matrix S satisfies

$$S^T = -S.$$

Such a matrix has zero diagonal elements, and its eigenvalues have been known for well over a century [26, 30] to be zero or purely imaginary, appearing in complex conjugate pairs. This immediately implies that skew-symmetric matrices are singular if their dimension is an odd number, and they may, of course, be singular also when their dimension is even. Direct solvers [6, 13, 15] require special pivoting strategies due to the presence of zeros on the diagonal. Iterative solvers [16, 18, 19] tend to be slow and challenging. Despite the Lanczos process being attractively simple for skew-symmetric matrices, the conjugate gradient (CG) method cannot be directly applied and in practice one has to resort to applying it to the normal equations, having implications on conditioning and convergence rate. Minimum residual methods face similar challenges [19, 22].

Hamiltonian and skew-Hamiltonian matrices are intimately connected to skew-symmetric matrices. Let us define them and describe a few of their properties; see [27] for a thorough review.

DEFINITION 1.1. *Consider the* $2n \times 2n$ *matrix*

$$\hat{J} = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix},$$

where I_n *denotes the* $n \times n$ *identity. Then a matrix* $A \in \mathbb{R}^{2n\times 2n}$ *is* <u>*Hamiltonian*</u> *if* $(A\hat{J})^T = A\hat{J}$. *A matrix* B *is* <u>*skew-Hamiltonian*</u> *if* $(B\hat{J})^T = -(B\hat{J})$.

The following three relations are equivalent [27, Theorem 2.1.1]:

1. A is Hamiltonian.
2. $A = \hat{J}H$, where H is symmetric.
3. $\hat{J}A$ is symmetric.

---

*Received August 18, 2021. Accepted March 22, 2022. Published online on May 12, 2022. Recommended by Peter Benner. The work of the author was supported by the Natural Sciences and Engineering Research Council of Canada.

†Department of Computer Science, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada (greif@cs.ubc.ca).

Among important properties of Hamiltonian matrices we mention *Hamiltonian eigensymmetry*: if λ is an eigenvalue of a real Hamiltonian matrix, then so are $-\lambda$, $\bar{\lambda}$, and $-\bar{\lambda}$; if λ is an eigenvalue of a real skew-Hamiltonian matrix, then so is $\bar{\lambda}$ and each eigenvalue has even algebraic multiplicity. See [10, Proposition 1.2] and [27, Proposition 2.3.1].

Let us also mention a special type of Hamiltonian matrices:

DEFINITION 1.2. *A matrix A is* Hamiltonian-positive *if it is Hamiltonian and its symmetric generator $H = \hat{J}^T A$ is positive definite.*

All eigenvalues of Hamiltonian-positive matrices are purely imaginary [1]. Properties and fast numerical computation of the eigenvalues of Hamiltonian and skew-Hamiltonian matrices have been extensively studied; see [9, 24, 25] and the references therein. In the context of iterative methods for eigenvalue problems, the symplectic Lanczos method [7, 8, 31] is a Krylov subspace method featuring short recurrences and based on the notion of J-orthogonality, i.e., orthogonality with respect to $\hat{J}$. The algorithm generates a so-called J-Hessenberg matrix on the projected subspace.

In this paper we consider a symmetric permutation of $\hat{J}$ and use it as a structured shift for solving skew-symmetric linear systems. The theory of Hamiltonian and skew-Hamiltonian matrices is useful for performing an eigenvalue analysis of the matrices involved. In Section 2 we introduce the idea of a skew-symmetric shift and in Section 3 we perform spectral analysis to illustrate its merits. In Section 4 we briefly discuss how skew-shifted linear systems may be iteratively solved. In Section 5 we analyze a variant of the Hermitian and skew-Hermitian splitting method using the proposed skew-symmetric shift. In Section 6 we offer brief concluding remarks.

2. Structured skew-symmetric shifts.

Let

$$J_2 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

and consider

$$J = I_n \otimes J_2. \tag{2.1}$$

The matrix J is a block-diagonal skew-symmetric $2n \times 2n$ matrix with respect to 2×2 blocks, with copies of J_2 along its main 2×2 block-diagonal. For example, for $n = 3$, the corresponding matrix J is 6×6 and is given by

$$J = \left[\begin{array}{cc|cc|cc} 0 & 1 & & & & \\ -1 & 0 & & & & \\ \hline & & 0 & 1 & & \\ & & -1 & 0 & & \\ \hline & & & & 0 & 1 \\ & & & & -1 & 0 \end{array}\right].$$

J is banded and there is no loss of generality here with respect to the notion of Hamiltonian or skew-Hamiltonian matrices, because $J = P\hat{J}P^T$, where P is a permutation matrix associated with the permutation vector

$$p = (1, 3, 5, \dots, 2n-1, 2, 4, 6, \dots, 2n).$$

We can then refer to permutations $\hat{H} = PHP^T$ and $\hat{S} = PSP^T$ in Definition 1.1 for Hamiltonian and skew-Hamiltonian matrices.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

It is easy to see that:

- J is orthogonal: $J^T = -J$ and $J^T J = JJ^T = I_{2n}$; from this it also follows that $J^2 = -I_{2n}$.
- The eigenvalues of J are $\pm i$, where
$$i = \sqrt{-1},$$
and each of them is of algebraic multiplicity n.

The above properties apply to $\hat{J}$ as well. The matrix J was used, for example, in [7, 14] in the context of the SR decomposition and symplectic Lanczos for Hamiltonian matrices.

A common way of handling the numerical difficulties related to solving linear algebra problems with a given skew-symmetric matrix, S, is to shift it if possible, namely to add a scaled identity to it. The standard shift is $S + \alpha I$, where α is a real scalar and I is the identity. The eigenvalues of S are shifted by α and the eigenvectors are preserved. The spectral effect of the shift is thus simple and fully understood. In the case of singular skew-symmetric matrices, the standard shift allows for eliminating zero eigenvalues and hence the singularity. Similarity transformations are easy to perform. If $Q^T SQ = T$ with Q orthogonal and T tridiagonal and skew-symmetric, then $Q^T(S + \alpha I)Q = T + \alpha I$. Iterative linear solvers have been based on shifts and have been proven to be effective; see [21, 22, 23, 29, 33].

The standard shift does not preserve skew-symmetry. Therefore, a solver that relies on the skew-symmetry of the original matrix S needs to be adjusted if S is shifted. It is thus useful to extend our set of computational tools by considering structure-preserving alternatives.

This leads us to the main idea of this work.

DEFINITION 2.1. *Given a real scalar α and a skew-symmetric matrix, S, consider*

$$S(\alpha) = S + \alpha J, \tag{2.2}$$

where J is defined in (2.1). We call this operation a skew-symmetric shift *and refer to the matrix $S(\alpha)$ as* skew-shifted *with respect to $S \equiv S(0)$.*

The operation defined in (2.2) preserves skew-symmetry. We refer to it as a shift even though it does not behave the same way as the standard shift by a scaled identity matrix. The skew-symmetric shift does not generally preserve the eigenvectors and does not shift the eigenvalues additively by α, but it does, nonetheless, have a shifting effect on the spectrum, which we analyze next.

3. The eigenvalues of skew-shifted matrices.

Let us consider the effect of the skew-symmetric shift on the eigenvalues of S. We will assume without loss of generality that $\alpha > 0$. (With some trivial adjustments, all forthcoming analytical observations can be adapted to $\alpha \leq 0$.)

We start with a simple example, which allows us to make some basic analytical observations. Suppose S is a $2n \times 2n$ skew-symmetric matrix that has the same nonzero structure as J:

$$S = D \otimes J_2,$$

where D is a diagonal $n \times n$ matrix. Without loss of generality, let us assume that D has positive diagonal entries, $d_j,\ j = 1, \dots, n$. The eigenvalues of S are given by

$$\lambda_j(S) = \pm i d_j, \qquad j = 1, \dots, n.$$

By known properties of eigenvalues of Kronecker products, the eigenvalues of the skew-shifted matrix are

$$\lambda_j(S(\alpha)) = \pm i(d_j + \alpha), \qquad j = 1, \dots, n.$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

The magnitudes of the shifted purely imaginary eigenvalues are $|i(d_j + \alpha)| = d_j + \alpha,\ j = 1, \ldots, n,$ and the effect on the eigenvalues for this simple case is different from the effect caused by a standard shift by the identity matrix. For the matrix S,

$$\lambda_j(S + \alpha I) = \pm i d_j + \alpha, \qquad j = 1, \ldots, n.$$

The shifted eigenvalues, when a scaled identity is used, move in one direction only, along the horizontal (real) axis, and their magnitudes are $\sqrt{d_j^2 + \alpha^2},\ j = 1, \ldots, n$. Since $\sqrt{d_j^2 + \alpha^2} \leq d_j + \alpha$, the skew-symmetric shift moves away the eigenvalues further from 0 compared to the standard shift.

Figure 3.1 illustrates the difference in the effect of the skew-symmetric shift vs. the standard shift for this simple example.

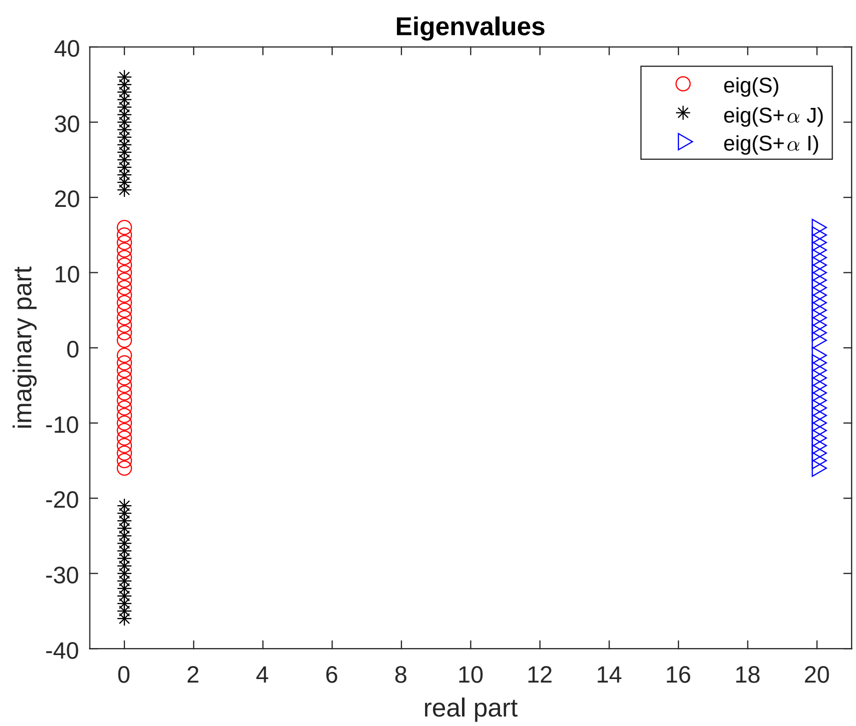

FIG. 3.1. *An illustration of the effect of a skew-symmetric shift for a simple block-diagonal case. The matrix is given by $S = D \otimes J_2$, where D is a diagonal 16×16 matrix with uniformly-spaced positive diagonal entries from 1 to 16. Shown are the eigenvalues of S, $S + \alpha J$ and $S + \alpha I$, with $\alpha = 20$.*

For general skew-symmetric matrices that do not have the structure of J we typically cannot explicitly compute the eigenvalues of the skew-shifted matrix and Figure 3.1 is no longer representative of the spectral distribution. But we can make a few observations about the effect on the spectrum for α sufficiently large.

THEOREM 3.1. *Suppose without loss of generality that the eigenvalues of S are given by*

$$\lambda_j(S) = \pm i \mu_j, \qquad j = 1, \ldots, n,$$

where $\mu_j \geq 0$, listed in descending order:

$$\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n \geq 0.$$

Then, if $\alpha > \mu_1$, the skew-shifted matrix $S(\alpha) = S + \alpha J$ is nonsingular and its eigenvalues are grouped within two distinct intervals on the imaginary axis, denoted by $i\mathcal{I}$ and corresponding to rotating the interval

$$\mathcal{I} = \Big[-\mu_1 - \alpha, \mu_1 - \alpha\Big] \bigcup \Big[-\mu_1 + \alpha, \mu_1 + \alpha\Big] \subset \mathbb{R} \tag{3.1}$$

by 90 degrees onto the imaginary axis.

Proof. Weyl's eigenvalue inequality theorem [32] leads to several useful eigenvalue inequalities, and one such result [20, Corollary 4.3.15] states that if C and D are two Hermitian matrices, say of size $n \times n$, whose eigenvalues are ordered in decreasing order, $\lambda_1(C) \geq \lambda_2(C) \geq \cdots \geq \lambda_n(C)$ and $\lambda_1(D) \geq \lambda_2(D) \geq \cdots \geq \lambda_n(D)$, then the eigenvalues of $C + D$, $\lambda_j(C + D), j = 1, \ldots, n$, satisfy

$$\lambda_j(C) + \lambda_n(D) \leq \lambda_j(C + D) \leq \lambda_j(C) + \lambda_1(D).$$

Since J and S are skew-symmetric, the matrices iJ and iS are Hermitian (indefinite) and the eigenvalues of iS are given by $-\mu_1 \leq -\mu_2 \leq \cdots \leq -\mu_n \leq 0 \leq \mu_n \leq \cdots \leq \mu_2 \leq \mu_1$, whereas the eigenvalues of iJ are 1 and -1, each of algebraic multiplicity n. Setting $C = i\alpha J$ and $D = iS$ and assuming $\alpha > \mu_1$, we have $iS(\alpha) = C + D$ and obtain the following $2n$ conditions on its eigenvalues:

$$-\mu_1 - \alpha \leq \lambda_j^-(iS(\alpha)) \leq \mu_1 - \alpha < 0, \qquad j = 1, \ldots, n,$$

and

$$0 < -\mu_1 + \alpha \leq \lambda_j^+(iS(\alpha)) \leq \mu_1 + \alpha, \qquad j = 1, \ldots, n.$$

Thus, the eigenvalues of $S(\alpha)$ are within the intervals defined by $i\mathcal{I}$ with $\mathcal{I}$ defined in (3.1), as stated in the theorem. Since $\alpha > \mu_1$, the two intervals defining $\mathcal{I}$ are decoupled and are located on both sides of the origin, and $S(\alpha)$ is nonsingular as claimed.

We remark that it is possible to obtain the bound on the maximum modulus directly by using the fact that S and $S + \alpha J$ are unitarily diagonalizable. The spectral radius of S satisfies $\rho(S) = \max_j |\lambda_j(S)| = \mu_1 = \|S\|_2$, and

$$\begin{aligned}\max_j |\lambda_j(S + \alpha J)| &= \rho(S + \alpha J) = \|S + \alpha J\|_2 \\ &\leq \|S\|_2 + \alpha = \rho(S) + \alpha = \mu_1 + \alpha,\end{aligned}$$

giving the absolute value of the endpoints of $i\mathcal{I}$, which are symmetric about the origin. □

COROLLARY 3.2. *Under the conditions of Theorem 3.1, an upper bound on the spectral condition number of $S(\alpha) = S + \alpha J$ is given by*

$$\kappa_2(S + \alpha J) \leq \frac{\alpha + \mu_1}{\alpha - \mu_1}. \tag{3.2}$$

Consequently, if $\kappa(S) > 1$, a sufficient condition for $S + \alpha J$ to be better conditioned than S is

$$\alpha \geq \mu_1 \frac{\mu_1 + \mu_n}{\mu_1 - \mu_n} = \mu_1 \frac{\kappa(S) + 1}{\kappa(S) - 1} > \mu_1. \tag{3.3}$$

Proof. Denote the eigenvalues of $S + \alpha J$ by $\lambda_j(S + \alpha J), j = 1, \dots, n$. By skew-symmetry, the condition number of $S + \alpha J$ is given by

$$\kappa(S + \alpha J) = \frac{\max_j |\lambda_j(S + \alpha J)|}{\min_j |\lambda_j(S + \alpha J)|}.$$

Using the symmetry of $\mathcal{I}$ about the origin and taking its endpoints, it can be algebraically confirmed that if (3.3) is satisfied, then by (3.2)

$$\kappa(S + \alpha J) \leq \frac{\alpha + \mu_1}{\alpha - \mu_1} \leq \frac{\mu_1}{\mu_n} = \kappa(S),$$

which gives a sufficient condition for $S + \alpha J$ to be better conditioned than S. ∎

Let us make a few remarks on the results stated in Corollary 3.2.

1. If $\kappa(S) = 1$, S is perfectly conditioned and skew-shifting it is of no practical interest. Therefore, we assume $\kappa > 1$ and α in (3.3) is well defined. On the other hand, when S is very ill-conditioned, namely $\kappa(S) \gg 1$, (3.3) means $\alpha \gtrapprox \mu_1$.
2. We note that (3.2) is an analytical upper bound for $\kappa(S + \alpha J)$, whereas $\kappa(S)$ is known exactly in terms of the eigenvalues of S. In practice, we have observed that the range of values of α for which $S + \alpha J$ is better conditioned than S is larger than the range predicted by the analysis, which is just a sufficient but not a necessary condition.
3. The bound on the condition number goes down monotonically towards 1 as α increases, and may suggest taking α as large as possible. A similar situation holds for the standard shift. However, aiming to minimize the condition number of $S + \alpha J$ is not the only consideration, as the discussion in Section 5 shows. Therefore, we cannot always aim for selecting an arbitrarily large value of α.

4. Solving skew-shifted systems. Consider a CG-type solver for solving the linear system

$$(S + \alpha J)x = b. \tag{4.1}$$

Due to the skew-symmetry, double steps must be considered, and this amounts to applying CG to the normal equations [19]. In [18] we presented a version of CGNE that uses Golub-Kahan bidiagonalization [17]. Algorithm 1, which contains trivial notational modifications of [18, Algorithm 4], lays out the iterative process.

Algorithm 1 Golub-Kahan bidiagonalization-based CGNE for $(S + \alpha J)x = b$.

Given b, compute $u_1\beta_1 := b$, $v_1\gamma_1 := -(S + \alpha J)u_1$, $\tau_1 := \beta_1/\gamma_1$, $x_1 := v_1\tau_1$
for $j = 1$ step 1 until convergence **do**
  $u_{j+1}\beta_{j+1} := (S + \alpha J)v_j - u_j\gamma_j$
  $v_{j+1}\gamma_{j+1} := -(S + \alpha J)u_{j+1} - v_j\beta_{j+1}$
  $\tau_{j+1} := -\tau_j\beta_{j+1}/\gamma_{j+1}$
  $x_{j+1} := x_j + v_{j+1}\tau_{j+1}$
end for

EXAMPLE 4.1. Consider the convection-diffusion equation in three dimensions with constant convective coefficients,

$$-\Delta u + (\sigma, \tau, \mu)\nabla u = f,$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

on the unit cube, subject to homogeneous Dirichlet boundary conditions. We discretize it using a uniform mesh with mesh size $h = 1/(n+1)$ and applying centered finite differences. This yields an $n^3 \times n^3$ linear system. Denoting the mesh Reynolds numbers by $\beta = \frac{\sigma h}{2}, \gamma = \frac{\tau h}{2}$, and $\delta = \frac{\mu h}{2}$, the discretization is numerically stable if $\beta, \gamma, \delta < 1$. Those three values determine the skew-symmetric part of the discrete operator associated with the linear system. We discard the symmetric part and use the skew-symmetric part as the matrix for the linear system; we denote it by S.

We take size 4096×4096 and apply Algorithm 1 to solve linear system (4.1). The infinity norm of S is 3.6, and it is an effective upper bound on the maximal eigenvalue of S, which is approximately 3.54. We pick three values of α. For $\alpha = 1$ there is no convergence. The middle value, $\alpha = 4$, is a bit larger than the maximal modulus of S and its infinity norm and it guarantees that $S + \alpha J$ is nonsingular by Theorem 3.1, as well as better conditioned than S, by Corollary 3.2. Convergence in this case is slow but steady. For $\alpha = 10$ convergence is fast and the residual norm is reduced to approximately 10^{-7} within eight iterations, i.e., 16 matrix-vector products.

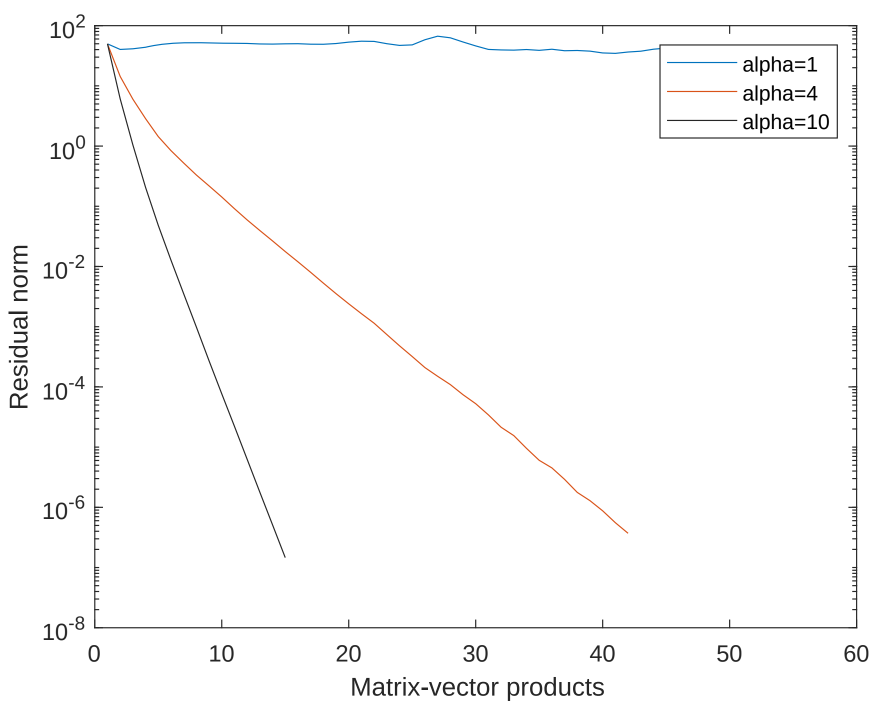

FIG. 4.1. *Convergence behavior of CGNE (Algorithm 1) for the* 4096×4096 *skew-symmetric linear system described in Example 4.1. We use* $\beta = 0.5$, $\gamma = 0.6, \delta = 0.7$. *The right-hand side, b, was set such that the solution of the linear system is the vector of 1s. The initial guess was zero, and the initial residual norm was* $\|b\| = 49.83$.

5. Hermitian/skew-Hermitian splitting iterative schemes. The Hermitian and skew-Hermitian splitting (HSS) iteration scheme [4] for solving a nonsymmetric linear system

ETNA
Kent State University and
Johann Radon Institute (RICAM)

$Ax = b$ is given by

$$\begin{cases} (\alpha I + H)x_{k+\frac{1}{2}} &= (\alpha I - S)x_k + b, \\ (\alpha I + S)x_{k+1} &= (\alpha I - H)x_{k+\frac{1}{2}} + b, \end{cases} \tag{5.1}$$

where H and S are the symmetric and skew-symmetric parts, respectively, of A:

$$H = \frac{(A + A^T)}{2}, \qquad S = \frac{(A - A^T)}{2}.$$

In [4], where the scheme was first introduced, the authors assumed that H is positive definite and derived an optimal value of α that minimizes an upper bound on the spectral radius of the iteration matrix. Krylov subspace methods can accelerate the convergence of the scheme while relaxing the boundedness requirement on the spectral radius of the iteration matrix. Spectral analysis of this scheme and extensions have been given in several papers; see, for example, [2, 3, 4, 11, 12].

5.1. A skew-shifted HSS iterative scheme. We now consider a scheme based on replacing the standard shift in the original HSS scheme (5.1) by the skew-symmetric shift. Consider

$$\begin{cases} (\alpha J + H)x_{k+\frac{1}{2}} &= (\alpha J - S)x_k + b, \\ (\alpha J + S)x_{k+1} &= (\alpha J - H)x_{k+\frac{1}{2}} + b. \end{cases} \tag{5.2}$$

In the second half-iteration $S + \alpha J$ is skew-symmetric. Therefore, we can directly apply a standard short-recurrence Krylov subspace method for skew-symmetric matrices without having to resort to special schemes such as [21, 22, 33].

Let us establish the nonsingularity of the two matrices in (5.2) that are to be inverted.

PROPOSITION 5.1. *For α sufficiently large, the matrices $\alpha J + H$ and $\alpha J + S$ are nonsingular, and therefore scheme (5.2) is well defined.*

Proof. The matrix $\alpha J + H$ is nonsingular if and only if $J^T(\alpha J + H) = \alpha I - JH$ is nonsingular, and for the latter, if H is symmetric positive definite then JH is Hamiltonian-positive (Definition 1.2) and all its eigenvalues are purely imaginary. Therefore $\alpha I - JH$ is nonsingular for $\alpha \neq 0$, which implies that $\alpha J + H$ is nonsingular.

As for $\alpha J + S$, by Theorem 3.1 and Corollary 3.2, for α sufficiently large $\alpha J + S$ is nonsingular and better conditioned than S. ∎

Scheme (5.2) is mathematically equivalent to applying HSS to the transformed linear system $-JAx = -Jb$ by considering the split $JA = (-JH) + (-JS)$:

$$\begin{cases} (\alpha I - JH)x_{k+\frac{1}{2}} &= (\alpha I + JS)x_k - Jb, \\ (\alpha I - JS)x_{k+1} &= (\alpha I + JH)x_{k+\frac{1}{2}} - Jb. \end{cases}$$

Let us denote the iteration matrix for (5.2) as

$$T_{JJ} = (\alpha J + S)^{-1}(\alpha J - H)(\alpha J + H)^{-1}(\alpha J - S).$$

Equivalently, in terms of standard shifts of Hamiltonian/skew-Hamiltonian matrices, we have

$$T_{JJ} = (\alpha I - JS)^{-1}(\alpha I + JH)(\alpha I - JH)^{-1}(\alpha I + JS).$$

It is possible to avoid the loss of symmetry in solving for $H + \alpha J$ in (5.2) by considering an alternative scheme that skew-shifts S by αJ and shifts H by the standard shift:

$$\begin{cases} (\alpha I + H)x_{k+\frac{1}{2}} &= (\alpha I - S)x_k + b, \\ (\alpha J + S)x_{k+1} &= (\alpha J - H)x_{k+\frac{1}{2}} + b. \end{cases} \tag{5.3}$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

The parameter α in the first half of the iteration may be replaced by a different parameter in the second half of the iteration. This scheme preserves symmetry for iterations involving inversion of a shifted version of H and skew-symmetry for iterations involving a skew-shifted version of S. Here, too, we can obtain a mathematically equivalent scheme involving shifted Hamiltonian and skew-Hamiltonian matrices:

$$\begin{cases} (\alpha I + H)x_{k+\frac{1}{2}} & = (\alpha I - S)x_k + b, \\ (\alpha I - JS)x_{k+1} & = (\alpha I + JH)x_{k+\frac{1}{2}} - Jb. \end{cases}$$

The iteration matrix for (5.3) is

$$T_{IJ} = (\alpha J + S)^{-1}(\alpha J - H)(\alpha I + H)^{-1}(\alpha I - S),$$

and here, too, it is possible to express T_{IJ} as a product of shifted Hamiltonian and skew-Hamiltonian matrices.

5.2. Spectral analysis. Let us assess the convergence of (5.2). We use here the technique applied in [4]. Applying a similarity transformation

$$\widehat{T}_{JJ} = (\alpha J + S)\, T_{JJ}\, (\alpha J + S)^{-1},$$

T_{JJ} and $\widehat{T}_{IJ}$ have have same eigenvalues. We write the latter as a product of two matrices:

$$\widehat{T}_{JJ} = \underbrace{(\alpha J - H)(\alpha J + H)^{-1}}_{\widehat{T}_H}\underbrace{(\alpha J - S)(\alpha J + S)^{-1}}_{\widehat{T}_S}.$$

The spectral radius of T_{JJ} is bounded as follows:

$$\rho(T_{JJ}) = \rho(\widehat{T}_{JJ}) \leq \|\widehat{T}_H\|_2 \|\widehat{T}_S\|_2,$$

and we now examine each of the matrices $\widehat{T}_S$ and $\widehat{T}_H$.

LEMMA 5.2. *If H is symmetric positive definite, the eigenvalues of*

$$\widehat{T}_H = (\alpha J - H)(\alpha J + H)^{-1}$$

are all equal to 1 in modulus.

Proof. Let $\widehat{T}_H x = \lambda x$. The eigenvalue problem can be written as a generalized eigenvalue problem of the form

$$(\alpha J - H)y = \lambda(\alpha J + H)y.$$

Alternatively, $\widehat{T}_H$ can be written as a Cayley transform of $-JH$:

$$\widehat{T}_H = (\alpha I + JH)(\alpha I - JH)^{-1}.$$

Either way, we get that the eigenvalues of $\widehat{T}_H$ are related to those of JH by

$$JHy = \left(\alpha \frac{\lambda - 1}{\lambda + 1}\right) y.$$

Since H is positive definite, JH is Hamiltonian-positive and all its eigenvalues are purely imaginary [1] (see Definition 1.2). Suppose its eigenvalues are given by $\pm i\mu_j, j = 1, \ldots, n$.

Then

$$\alpha\frac{\lambda_j-1}{\lambda_j+1}=\pm i\mu_j, \qquad j=1,\dots,n,$$

from which it follows that the eigenvalues of $\widehat{T}_H$ are given by

$$\lambda_j=\frac{\alpha\pm i\mu_j}{\alpha\mp i\mu_j}, \qquad j=1,\dots,n.$$

Therefore $|\lambda_j|=1$ for all $j=1,\dots,n$. ∎

LEMMA 5.3. *The eigenvalues of*

$$\widehat{T}_S=(\alpha J-S)(\alpha J+S)^{-1}$$

satisfy

$$\lambda_j=\frac{\alpha+\nu_j}{\alpha-\nu_j},$$

where ν_j are the eigenvalues of the skew-Hamiltonian matrix JS.

Proof. Repeating the steps of the proof of Lemma 5.2, the eigenvalues of $\widehat{T}_S$ satisfy the generalized eigenvalue problem

$$(\alpha J-S)y=\lambda(\alpha J+S)y,$$

from which it follows that

$$JSy=\left(\alpha\frac{\lambda-1}{\lambda+1}\right)y.$$

The required result follows. Similarly to Lemma 5.2, it can also be obtained by writing down the Cayley transform of the skew-Hamiltonian matrix $-JS$:

$$\widehat{T}_S=(\alpha I+JS)(\alpha I-JS)^{-1},$$

and proceeding to compute the eigenvalues. ∎

Given that the spectral radii of $\widehat{T}_S$ and $\widehat{T}_H$ may not be closely below their respective spectral norms, convergence of the scheme cannot be analytically guaranteed. Nonetheless, from Lemmas 5.2 and 5.3 it follows that the availability of a suitable analytical choice of α may primarily depend on the spectral distribution of JS. Since the spectral radius of $\widehat{T}_H$ is 1, a reasonable strategy is to aim to minimize the spectral radius of $\widehat{T}_S$. We show below a scenario where this may be possible to accomplish.

DEFINITION 5.4. *Let S be a nonsingular $2n\times 2n$ skew-symmetric matrix. Denote its main 2×2 diagonal blocks by*

$$T_\ell=\begin{bmatrix}0 & s_{\ell,\ell+1}\\ -s_{\ell,\ell+1} & 0\end{bmatrix}, \qquad \ell=1,3,\dots,n.$$

We say that S is skew-diagonally dominant *if*

$$\|T_\ell\|=|s_{\ell,\ell+1}|\geq\sum_{j\neq\ell,\ell+1}\left(|s_{\ell,j}|+|s_{\ell+1,j}|\right), \qquad \forall\,\ell=1,3,\dots,n.$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

Definition 5.4 basically extends the standard notion of diagonally-dominant matrices to be defined with respect to 2×2 blocks, which is suitable for skew-symmetric matrices, due to having zero diagonal elements. The skew-symmetric shift makes a given skew-symmetric matrix skew-diagonally dominant for a sufficiently large value of α.

Going back to Lemma 5.3, it may be possible to find a suitable value of α if the skew-symmetric part of A is skew-diagonally dominant. For example, in the simple case where $S \approx \varphi J$ with $\varphi > 0$, the lemma indicates that taking $\alpha \approx \varphi$ may provide a good choice for the shift. (In practice we may want to take $\alpha \gtrapprox \varphi$ to improve the conditioning and spectral distribution of the skew-shifted matrix.)

While Lemma 5.2 applies to H symmetric positive definite (for which JH is Hamiltonian-positive), scheme (5.2) may still work for H mildly indefinite. The Hamiltonian-positivity is lost in this case, but if the spectral radius stays close to 1, it may still be possible to obtain convergence of the scheme.

The scheme (5.3) is potentially more practical than (5.2) because the symmetry with respect to H-related solves is preserved. Eigenvalue problems

$$(\alpha J - H)x = \lambda(\alpha I + H)x\,; \qquad (\alpha I - S)y = \mu(\alpha J + S)y$$

need to be solved in order to evaluate the spectral radius of the iteration matrix. These eigenvalue problems may be rewritten as

$$(\lambda + 1)(J - \lambda I)^{-1}Hx = \alpha x; \qquad (\mu + 1)(I - \mu J)^{-1}Sy = \alpha y,$$

where the matrices $J - \lambda I$ and $I - \mu J$ are block diagonal with respect to 2×2 blocks.

5.3. Acceleration using Krylov subspace iterations. From the analysis in Section 5.2 it is evident that it is difficult to find a choice of the parameter α that guarantees boundedness of the spectral radius of the iteration matrix, and hence convergence. We thus consider relaxing this requirement by using the proposed scheme as a preconditioner for a Krylov subspace solver.

Similarly to the algebraic derivations in [12, Section 3], we can write schemes (5.2) and (5.3) as

$$x_{k+1} = Tx_k + c,$$

where $T = I - M^{-1}A$ and $c = M^{-1}b$. For (5.2), the iteration matrix is $T = T_{JJ}$ and we use the notation $M = M_{JJ}$. Similarly, we denote by $M = M_{IJ}$ the matrix corresponding to $T = T_{IJ}$, related to (5.3). Adapting the algebraic steps of [12, Section 3] to the current scheme, it follows that

$$M_{JJ} = -\frac{1}{2\alpha}(\alpha I - JH)(\alpha J + S)$$

for (5.2), and

$$M_{IJ} = \frac{1}{2\alpha}(\alpha I + H)(I - J)(\alpha J + S)$$

for (5.3).

We can thus take M_{IJ} or M_{JJ} as preconditioners for a suitable Krylov subspace solver, for example GMRES. In practice there is no need to keep the constants $\pm\frac{1}{2\alpha}$ in the preconditioners because they do not make a difference in the iteration; this is done also in [12, Section 4]. The nonsingularity of the two potential preconditioners follows from Proposition 5.1.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

The matrix M_{IJ} is potentially more appealing as a preconditioner for GMRES than M_{JJ}, because it involves inversions that can be carried out easily. The matrix $I - J$ is block diagonal with respect to 2×2 blocks and it is trivial to solve for directly. The shifted matrix $\alpha I + H$ is symmetric positive definite, and CG can be used here as an inner iteration. The skew-shifted matrix $\alpha J + S$ is well conditioned for α sufficiently large, and we have shown in Section 4 that CGNE can be effectively applied to solve such linear systems. Based on the theory of inner-outer iterations [28], it is possible to apply inner iterations of CG and CGNE to a crude tolerance throughout the GMRES outer iteration.

EXAMPLE 5.5. We return to the convection-diffusion equation in three dimensions with constant convective coefficients, specified in Example 4.1, but allow a shift that makes the symmetric part of the coefficient matrix indefinite. In other words, we allow having a Helmholtz-type operator rather than a Laplacian for the symmetric part of the discrete differential operator. The problem corresponds to a discretization of the PDE

$$-(\Delta + k^2)u + (\sigma, \tau, \mu)\nabla u = f$$

on the unit cube, subject to homogeneous Dirichlet boundary conditions. We use the same discretization as described in Example 4.1, and make two choices of mesh Reynolds numbers. On the left-hand side of Figure 5.1 we show the result of using the same mesh Reynolds numbers used in Example 4.1, and on the right-hand side we show an example where the matrix is skew-diagonally dominant, as per Definition 5.4. We select k so that the symmetric part of the coefficient matrix (scaled by h^2, where h is the mesh size) is $H = L - \xi I$, where L is the discrete Laplacian (scaled by h^2) and I is the identity matrix, and test for $\xi = 0$ (for which H is the discrete Laplacian) as well $\xi = 1$ and $\xi = 2$. We use M_{IJ} (scaled by -2α). Details on the convective coefficients used are given in the caption of Figure 5.1.

We apply a loose inner stopping criterion, to accelerate convergence. Our experiments show that the iterative scheme is robust even when an inner tolerance as low as 0.01 and a maximum of five or six matrix-vector products for each of the CG and the CGNE schemes are imposed. In the last few outer iterations, close to convergence, the tolerance for the inner iterations was tightened. As evident from the figures, the overall computational cost of the scheme is modest.

In the skew-diagonally dominant case, the preconditioned Krylov scheme is rather robust with respect to higher values of ξ and we show convergence also for $\xi = 2$. When the matrix is not skew-diagonally, dominant convergence deteriorates more rapidly when H becomes indefinite, as can be observed on the left-hand plot, but the scheme still converges.

We note that the stationary scheme (5.3) converges for this example when the matrix is strongly skew-diagonally dominant and H is positive definite or mildly indefinite, but did not converge otherwise, for example in the setup of the left-hand figure. Altogether it is less robust as a standalone solver than the preconditioned Krylov solver.

6. Concluding remarks. We have introduced structured shifts for skew-symmetric matrices, which are shown to be useful in solving linear systems with a strong skew-symmetric component. Hamiltonian and skew-Hamiltonian matrices provide a valuable tool for performing an eigenvalue analysis. An HSS-like solver using skew-symmetric shifts has been introduced. Obtaining analytical bounds on convergence is challenging, but it is possible to obtain some results connecting the bounds to the spectra of skew-symmetric, Hamiltonian, and skew-Hamiltonian matrices, and the scheme seems robust when used as a preconditioner for Krylov subspace solvers.

Acknowledgments. I am grateful to Volker Mehrmann for his careful reading and insightful comments on an earlier version of this paper and to the referees for their very helpful reviews.

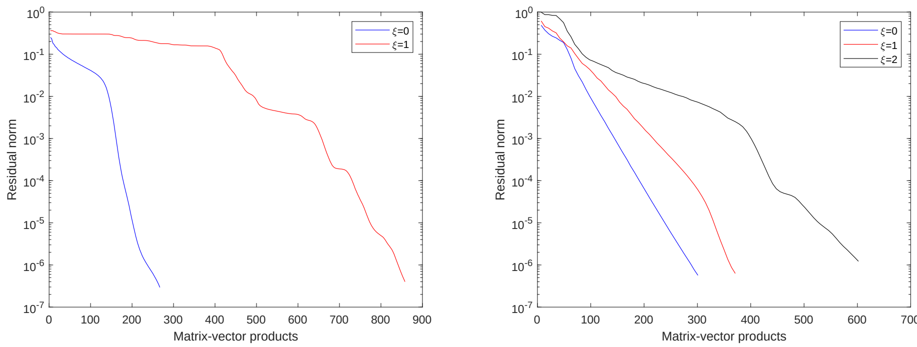

FIG. 5.1. *Convergence behavior of GMRES, preconditioned with M_{IJ} applied to variations of the problem described in Example 5.5. On the left-hand side we use $\beta = 0.5$, $\gamma = 0.6$, and $\delta = 0.7$. The values of ξ are given in the legends of the figures. On the right-hand side $\beta = 5$, $\gamma = 0.6$, and $\delta = 0.7$. We select for all experiments $\alpha = 10$. The right-hand side was set so that the solution is the vector of all 1s. The initial guess was the zero vector.*

REFERENCES

[1] P. AMODIO, *A symplectic Lanczos-type algorithm to compute the eigenvalues of positive definite Hamiltonian matrices*, in Computational science—ICCS 2003, Part II, P. M. A. Sloot, D. Abramson, A. V. Bogdanov, J. J. Dongarra, A. Y. Zomaya, and Y. E. Gorbachev, eds., vol. 2658 of Lecture Notes in Comput. Sci., Springer, Berlin, 2003, pp. 139–148.

[2] Z.-Z. BAI, *Quasi-HSS iteration methods for non-Hermitian positive definite linear systems of strong skew-Hermitian parts*, Numer. Linear Algebra Appl., 25 (2018), Art. e2116, 19 pages.

[3] Z.-Z. BAI, G. H. GOLUB, L.-Z. LU, AND J.-F. YIN, *Block triangular and skew-Hermitian splitting methods for positive-definite linear systems*, SIAM J. Sci. Comput., 26 (2005), pp. 844–863.

[4] Z.-Z. BAI, G. H. GOLUB, AND M. K. NG, *Hermitian and skew-Hermitian splitting methods for non-Hermitian positive definite linear systems*, SIAM J. Matrix Anal. Appl., 24 (2003), pp. 603–626.

[5] C. BEATTIE, V. MEHRMANN, H. XU, AND H. ZWART, *Linear port-Hamiltonian descriptor systems*, Math. Control Signals Systems, 30 (2018), Art. 17, 27 pages.

[6] P. BENNER, R. BYERS, H. FASSBENDER, V. MEHRMANN, AND D. WATKINS, *Cholesky-like factorizations of skew-symmetric matrices*, Electron. Trans. Numer. Anal., 11 (2000), pp. 85–93.
http://etna.ricam.oeaw.ac.at/vol.11.2000/pp85-93.dir/pp85-93.pdf

[7] P. BENNER AND H. FASSBENDER, *An implicitly restarted symplectic Lanczos method for the Hamiltonian eigenvalue problem*, Linear Algebra Appl., 263 (1997), pp. 75–111.

[8] P. BENNER, H. FASSBENDER, AND M. STOLL, *A Hamiltonian Krylov-Schur-type method based on the symplectic Lanczos process*, Linear Algebra Appl., 435 (2011), pp. 578–600.

[9] P. BENNER, D. KRESSNER, AND V. MEHRMANN, *Skew-Hamiltonian and Hamiltonian eigenvalue problems: theory, algorithms and applications*, in Proceedings of the Conference on Applied Mathematics and Scientific Computing, Z. Drmač, M. Marušić, and Z. Tutek, eds., Springer, Dordrecht, 2005, pp. 3–39.

[10] P. BENNER, V. MEHRMANN, AND H. XU, *A note on the numerical solution of complex Hamiltonian and skew-Hamiltonian eigenvalue problems*, Electron. Trans. Numer. Anal., 8 (1999), pp. 115–126.
http://etna.ricam.oeaw.ac.at/vol.8.1999/pp115-126.dir/pp115-126.pdf

[11] M. BENZI, M. J. GANDER, AND G. H. GOLUB, *Optimization of the Hermitian and skew-Hermitian splitting iteration for saddle-point problems*, BIT, 43 (2003), pp. 881–900.

[12] M. BENZI AND G. H. GOLUB, *A preconditioner for generalized saddle point problems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 20–41.

[13] J. R. BUNCH, *A note on the stable decompostion of skew-symmetric matrices*, Math. Comp., 38 (1982), pp. 475–479.

[14] A. BUNSE-GERSTNER AND V. MEHRMANN, *A symplectic QR like algorithm for the solution of the real algebraic Riccati equation*, IEEE Trans. Automat. Control, 31 (1986), pp. 1104–1113.

[15] I. S. DUFF, *The design and use of a sparse direct solver for skew symmetric matrices*, J. Comput. Appl. Math., 226 (2009), pp. 50–54.

[16] S. C. EISENSTAT, *Equivalence of Krylov subspace methods for skew-symmetric linear systems*, Tech. Rep., Department of Computer Science, Yale University.
Published at http://arxiv.org/abs/1512.00311, 2015.

[17] G. GOLUB AND W. KAHAN, *Calculating the singular values and pseudo-inverse of a matrix*, J. Soc. Indust.

Appl. Math. Ser. B Numer. Anal., 2 (1965), pp. 205–224.

[18] C. Greif, C. C. Paige, D. Titley-Peloquin, and J. M. Varah, *Numerical equivalences among Krylov subspace algorithms for skew-symmetric matrices*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 1071–1087.

[19] C. Greif and J. M. Varah, *Iterative solution of skew-symmetric linear systems*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 584–601.

[20] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd. ed., Cambridge University Press, Cambridge, 2013.

[21] R. Idema and C. Vuik, *A minimal residual method for shifted skew-symmetric systems*, Tech. Rep. 07-09, Delft University of Technology, Delft, 2007.

[22] E. Jiang, *Algorithm for solving shifted skew-symmetric linear system*, Front. Math. China, 2 (2007), pp. 227–242.

[23] M. Manguoğlu and V. Mehrmann, *A two-level iterative scheme for general sparse linear systems based on approximate skew-symmetrizers*, Electron. Trans. Numer. Anal., 54 (2021), pp. 370–391. http://etna.ricam.oeaw.ac.at/vol.54.2021/pp370-391.dir/pp370-391.pdf

[24] V. Mehrmann and P. Van Dooren, *Structured backward errors for eigenvalues of linear port-Hamiltonian descriptor systems*, SIAM J. Matrix Anal. Appl., 42 (2021), pp. 1–16.

[25] V. Mehrmann and D. Watkins, *Structure-preserving methods for computing eigenpairs of large sparse skew-Hamiltonian/Hamiltonian pencils*, SIAM J. Sci. Comput., 22 (2000), pp. 1905–1925.

[26] W. H. Metzler, *On certain properties of symmetric, skew symmetric, and orthogonal matrices*, Amer. J. Math., 15 (1893), pp. 274–282.

[27] K. R. Meyer and D. C. Offin, *Introduction to Hamiltonian Dynamical Systems and the N-Body Problem*, 3rd. ed., Springer, Cham, 2017.

[28] V. Simoncini and D. B. Szyld, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477.

[29] D. B. Szyld and O. B. Widlund, *Variational analysis of some conjugate gradient methods*, East-West J. Numer. Math., 1 (1993), pp. 51–74.

[30] H. Taber, *On certain properties of symmetric, skew symmetric, and orthogonal matrices*, Proc. Lond. Math. Soc., 22 (1890/91), pp. 449–469.

[31] D. S. Watkins, *On Hamiltonian and symplectic Lanczos processes*, Linear Algebra Appl., 385 (2004), pp. 23–45.

[32] H. Weyl, *Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung)*, Math. Ann., 71 (1912), pp. 441–479.

[33] O. B. Widlund, *A Lanczos method for a class of nonsymmetric systems of linear equations*, SIAM J. Numer. Anal., 15 (1978), pp. 801–812.