Stability of Structured Hamiltonian Eigensolvers

Tisseur, Françoise

2001

MIMS EPrint: **2007.224**

ISSN 1749-9097

SIAM J. Matrix Anal. Appl.
Vol. 23, No. 1, pp. 103–125

STABILITY OF STRUCTURED HAMILTONIAN EIGENSOLVERS*

FRANÇOISE TISSEUR†

Abstract. Various applications give rise to eigenvalue problems for which the matrices are Hamiltonian or skew-Hamiltonian and also symmetric or skew-symmetric. We define structured backward errors that are useful for testing the stability of numerical methods for the solution of these four classes of structured eigenproblems. We introduce the symplectic quasi-QR factorization and show that for three of the classes it enables the structured backward error to be efficiently computed. We also give a detailed rounding error analysis of some recently developed Jacobi-like algorithms of Faßbender, Mackey, and Mackey [*Linear Algebra Appl.*, to appear] for these eigenproblems. Based on the direct solution of 4×4, and in one case 8×8, structured subproblems these algorithms produce a complete basis of symplectic orthogonal eigenvectors for the two symmetric cases and a symplectic orthogonal basis for all the real invariant subspaces for the two skew-symmetric cases. We prove that, when the rotations are implemented using suitable formulae, the algorithms are strongly backward stable and we show that the QR algorithm does not have this desirable property.

Key words. Hamiltonian, skew-Hamiltonian, symmetric, skew-symmetric, symplectic, backward error, structure-preserving, rounding error, Jacobi algorithm, quaternion rotation

AMS subject classifications. 65F15, 65G05

PII. S0895479800368007

1. Introduction. This work concerns real structured Hamiltonian and skew-Hamiltonian eigenvalue problems where the matrices are either symmetric or skew-symmetric. We are interested in algorithms that are strongly backward stable for these problems. In general, a numerical algorithm is called *backward stable* if the computed solution is the true solution for slightly perturbed initial data. If, in addition, this perturbed initial problem has the same structure as the given problem, then the algorithm is said to be *strongly backward stable*.

There are three reasons for our interest in strongly backward stable algorithms. First, such algorithms preserve the algebraic structure of the problem and hence force the eigenvalues to lie in a certain region of the complex plane or to occur in particular kinds of pairings. Because of rounding errors, algorithms that do not respect the structure of the problem can cause eigenvalues to leave the required region [26]. Second, by taking advantage of the structure, storage and computation can be lowered. Finally, structure-preserving algorithms may compute eigenpairs that are more accurate than the ones provided by a general algorithm.

Structured Hamiltonian eigenvalue problems appear in many scientific and engineering applications. For instance, symmetric skew-Hamiltonian eigenproblems arise in quantum mechanical problems with time reversal symmetry [9], [23]. In response theory, the study of closed shell Hartree–Fock wave functions yields a linear response eigenvalue equation with a symmetric Hamiltonian [21]. Also, total least squares problems with symmetric constraints lead to the solution of a symmetric Hamiltonian problem [17].

*Received by the editors February 23, 2000; accepted for publication (in revised form) by V. Mehrmann November 24, 2000; published electronically May 3, 2001.

http://www.siam.org/journals/simax/23-1/36800.html

†Department of Mathematics, University of Manchester, Manchester, M13 9PL, England (ftisseur@ma.man.ac.uk, http://www.ma.man.ac.uk/~ftisseur). This work was supported by Engineering and Physical Sciences Research Council grant GR/L76532.

The motivation for this work comes from recently developed Jacobi algorithms for structured Hamiltonian eigenproblems [10]. These algorithms are structure-preserving, inherently parallelizable, and hence attractive for solving large-scale eigenvalue problems. Our first contribution is to define and show how to compute structured backward errors for structured Hamiltonian eigenproblems. These backward errors are useful for testing the stability of numerical algorithms. Our second contribution concerns the stability of these new Jacobi-like algorithms. We give a unified description of the algorithms for the four classes of structured Hamiltonian eigenproblems. This provides a framework for a detailed rounding error analysis and enables us to show that the algorithms are strongly backward stable when the rotations are implemented using suitable formulae.

The organization of the paper is as follows. In section 2 we recap the necessary background concerning structured Hamiltonians. In section 3 we derive computable structured backward errors for structured Hamiltonian eigenproblems. In section 4, we describe the structure-preserving QR-like algorithms proposed in [5] for structured Hamiltonian eigenproblems. We give a unified description of the new Jacobi-like algorithms and detail the Jacobi-like update for each of the four classes of structured Hamiltonian. In section 5 we give the rounding error analysis and in section 6 we use our computable backward errors to confirm empirically the strong stability of the Jacobi algorithms.

2. Preliminaries. A matrix $P \in \mathbb{R}^{2n\times 2n}$ is *symplectic* if $P^TJP = J$, where $J = \left[\begin{smallmatrix} 0 & I_n \\ -I_n & 0 \end{smallmatrix}\right]$ and I_n is the $n \times n$ identity matrix.

A matrix $H \in \mathbb{R}^{2n\times 2n}$ is *Hamiltonian* if $JH = (JH)^T$ is symmetric. Hamiltonian matrices have the form

$$H = \begin{bmatrix} E & F \\ G & -E^T \end{bmatrix},$$

where $E, F, G \in \mathbb{R}^{n\times n}$ and $F^T = F,\ G^T = G$. We denote the set of real Hamiltonian matrices by $\mathcal{H}_{2n}$.

A matrix $S \in \mathbb{R}^{2n\times 2n}$ is *skew-Hamiltonian* if $JS = -(JS)^T$ is skew-symmetric. Skew-Hamiltonian matrices have the form

$$S = \begin{bmatrix} E & F \\ G & E^T \end{bmatrix},$$

where $E, F, G \in \mathbb{R}^{n\times n}$ and $F^T = -F,\ G = -G^T$ are skew-symmetric. We denote the set of real skew-Hamiltonian matrices by $\mathcal{SH}_{2n}$.

Note that if $H \in \mathcal{H}_{2n}$, then $P^{-1}HP \in \mathcal{H}_{2n}$ and if $S \in \mathcal{SH}_{2n}$, then $P^{-1}SP \in \mathcal{SH}_{2n}$, where P is an arbitrary symplectic matrix. Thus symplectic similarities preserve Hamiltonian and skew-Hamiltonian structure. Also, symmetric and skew-symmetric structures are preserved by orthogonal similarity transformations. Therefore structure-preserving algorithms for symmetric or skew-symmetric Hamiltonian or skew-Hamiltonian eigenproblems have to use real symplectic orthogonal transformations, that is, matrices $U \in \mathbb{R}^{2n\times 2n}$ satisfying $U^TJU = J, \quad U^TU = I$. As in [10], we denote by $SpO(2n)$ the group of real symplectic orthogonal matrices. Any $U \in SpO(2n)$ can be written as $U = \left[\begin{smallmatrix} U_1 & -U_2 \\ U_2 & U_1 \end{smallmatrix}\right]$, where $U_1^TU_1 + U_2^TU_2 = I$ and $U_1^TU_2 = U_2^TU_1$.

In Tables 2.1 and 2.2, we summarize the structure of Hamiltonian and skew-Hamiltonian matrices that are either symmetric or skew-symmetric, their eigenvalue

Table 2.1
Properties of structured Hamiltonian matrices $H \in \mathcal{H}_{2n}$.

$JH = (JH)^T$	Structure	Eigenvalues	Canonical form
Symmetric $H = H^T$	$\begin{bmatrix} E & F \\ F & -E \end{bmatrix}$, $E = E^T$, $F = F^T$	real, pairs $\lambda, -\lambda$	$\begin{bmatrix} D & 0 \\ 0 & -D \end{bmatrix}$
Skew-symmetric $H = -H^T$	$\begin{bmatrix} E & F \\ -F & E \end{bmatrix}$, $E = -E^T$, $F = F^T$	pure imaginary, pairs $\lambda, \bar{\lambda}$	$\begin{bmatrix} 0 & -D \\ D & 0 \end{bmatrix}$

Table 2.2
Properties of structured skew-Hamiltonian matrices $S \in \mathcal{SH}_{2n}$.

$JS = -(JS)^T$	Structure	Eigenvalues	Canonical form
Symmetric $S = S^T$	$\begin{bmatrix} E & F \\ -F & E \end{bmatrix}$, $E = E^T$, $F = -F^T$	real, double	$\begin{bmatrix} D & 0 \\ 0 & D \end{bmatrix}$
Skew-symmetric $S = -S^T$	$\begin{bmatrix} E & F \\ F & -E \end{bmatrix}$, $E = -E^T$, $F = -F^T$	pure imaginary, double, pairs $\lambda, \bar{\lambda}$	$\begin{bmatrix} B & 0 \\ 0 & -B \end{bmatrix}$

properties, and their symplectic orthogonal canonical form. We use $D \in \mathbb{R}^{n\times n}$ to denote a diagonal matrix and $B \in \mathbb{R}^{n\times n}$ to denote a block-diagonal matrix that is the direct sum of 1×1 zero blocks and 2×2 blocks of the form $\left[\begin{smallmatrix} 0 & b \\ -b & 0 \end{smallmatrix}\right]$. These canonical forms are consequences of results in [19].

Next, we show that the eigenvectors of skew-symmetric Hamiltonian matrices can be chosen to have structure. This property is important when defining and deriving structured backward errors.

Lemma 2.1. *The eigenvectors of a skew-symmetric Hamiltonian matrix H can be chosen to have the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$ with $z \in \mathbb{C}^n$.*

Proof. Let $\left[\begin{smallmatrix} 0 & -D \\ D & 0 \end{smallmatrix}\right] = U^T H U$ be the canonical form of H with $U = \left[\begin{smallmatrix} U_1 & -U_2 \\ U_2 & U_1 \end{smallmatrix}\right]$ symplectic orthogonal. The matrix $X = \frac{1}{\sqrt{2}}\left[\begin{smallmatrix} I & I \\ -iI & iI \end{smallmatrix}\right]$ is unitary and diagonalizes the canonical form of H:

$$X^* \begin{bmatrix} 0 & -D \\ D & 0 \end{bmatrix} X = \begin{bmatrix} iD & 0 \\ 0 & -iD \end{bmatrix}.$$

Hence

$$UX = \frac{1}{\sqrt{2}} \begin{bmatrix} U_1 + iU_2 & U_1 - iU_2 \\ U_2 - iU_1 & U_2 + iU_1 \end{bmatrix}$$

is an eigenvector basis for H and this shows that the eigenvectors can be taken to have the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$ with $z \in \mathbb{C}^n$. ∎

Note that an eigenvector of a skew-symmetric Hamiltonian matrix does not necessarily have the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$. For instance, consider $H = \left[\begin{smallmatrix} 0 & -D \\ D & 0 \end{smallmatrix}\right]$ with $D = \mathrm{diag}(-d, d)$.

TABLE 3.1
t: Number of parameters defining H.

	Hamiltonian		Skew-Hamiltonian	
	$H = H^T$	$H = -H^T$	$H = H^T$	$H = -H^T$
t	$n^2 + n$	n^2	n^2	$n^2 - n$

Then $x^T = [i, 1, 1, i]$ is an eigenvector of H, corresponding to the eigenvalue $-id$, that is not of the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$.

3. Structured backward error. We begin by developing structured backward errors that can be used to test the strong stability of algorithms for our classes of Hamiltonian eigenproblems.

3.1. Definition. For notational convenience, the symbol H denotes from now on both Hamiltonian and skew-Hamiltonian matrices. Let $(\widetilde{x}, \widetilde{\lambda})$ be an approximate eigenpair for the structured Hamiltonian eigenvalue problem $Hx = \lambda x$, where $H \in \mathbb{R}^{2n\times 2n}$. A natural definition of the normwise backward error of an approximate eigenpair is

$$\eta(\widetilde{x}, \widetilde{\lambda}) = \min\Big\{\epsilon : (H + \Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x},\ \|\Delta H\| \le \epsilon \|H\|\Big\},$$

where we measure the perturbation in a relative sense and $\|\cdot\|$ denotes any vector norm and the corresponding subordinate matrix norm. Deif [8] derived the explicit expression for the 2-norm

$$\eta(\widetilde{x}, \widetilde{\lambda}) = \frac{\|r\|_2}{\|H\|_2 \|\widetilde{x}\|_2},$$

where $r = \widetilde{\lambda}\widetilde{x} - H\widetilde{x}$ is the residual. This shows that the normwise relative backward error is a scaled residual. The componentwise backward error is a more stringent measure of the backward error in which the components of the perturbation ΔH are measured individually:

$$\omega(\widetilde{x}, \widetilde{\lambda}) = \min\Big\{\epsilon : (H + \Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x},\ |\Delta H| \le \epsilon |H|\Big\}.$$

Here inequalities between matrices hold componentwise. Geurts [12] showed that

$$\omega(\widetilde{x}, \widetilde{\lambda}) = \max_{1\le i\le 2n} \frac{|r|_i}{(|H||\widetilde{x}|)_i}.$$

The componentwise backward error provides a more meaningful measure of the stability than the normwise version when the elements in H vary widely in magnitude. However, this measure is not entirely appropriate for our problems as it does not respect any structure (other than sparsity) in H. Bunch [2] and Van Dooren [25] have also discussed other situations when it is desirable to preserve structure in definitions of backward errors.

The four classes of structured Hamiltonian matrices we are dealing with are defined by $t \le n^2 + n$ real parameters that make up E and F (see Table 3.1). We write this dependence as $H = H[p]$ with $p \in \mathbb{R}^t$. Higham and Higham [13], [14] extend the notion of componentwise backward error to allow dependence of the perturbations on a set of parameters and they define structured componentwise backward errors. Following their idea and notation we define the structured relative normwise backward error by

$$\mu(\widetilde{x},\widetilde{\lambda}) = \min\Big\{\epsilon : (H+\Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x},\ H+\Delta H = H[p+\Delta p],\ \|\Delta H\|_F \le \epsilon\|H\|_F\Big\}, \tag{3.1}$$

where $H + \Delta H = H[p + \Delta p]$ implies that ΔH has the same structure as H. The structured relative componentwise backward error $\nu(\widetilde{x},\widetilde{\lambda})$ is defined as in (3.1) but with the constraint $\|\Delta H\|_F \le \epsilon\|H\|_F$ replaced by $|\Delta H| \le \epsilon|H|$.

In our case, the dependence of the data on the t parameters is linear. We naturally require $(\widetilde{x},\widetilde{\lambda})$ to have any properties forced upon the exact eigenpairs, otherwise the backward error will be infinite. In the next subsections, we give algorithms for computing these backward errors. We start by describing a general approach that was used in [13] in the context of structured linear systems and extend it to the case where the approximate solution lies in the complex plane.

3.2. A general approach for the computation of $\boldsymbol{\mu(\widetilde{x},\widetilde{\lambda})}$. Let $\widetilde{x} = \widetilde{u} + i\widetilde{v}$ and $\widetilde{\lambda} = \widetilde{\mu}+i\widetilde{\nu}$. By equating real and imaginary parts, the constraint $(H+\Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x}$ in (3.1) becomes

$$\begin{bmatrix} \Delta H & 0 \\ 0 & \Delta H \end{bmatrix} \begin{bmatrix} \widetilde{u} \\ \widetilde{v} \end{bmatrix} = \begin{bmatrix} (\widetilde{\mu}I - H)\widetilde{u} - \widetilde{\nu}\widetilde{v} \\ \widetilde{\nu}\widetilde{u} + (\widetilde{\mu}I - H)\widetilde{v} \end{bmatrix} = \begin{bmatrix} s_1 \\ s_2 \end{bmatrix}, \tag{3.2}$$

or equivalently $\Delta H\,[\,\widetilde{u}\quad \widetilde{v}\,] = [\,s_1 \quad s_2\,]$. Applying the vec operator (which stacks the columns of a matrix into one long vector), we obtain

$$\Big([\,\widetilde{u}\quad \widetilde{v}\,]^T \otimes I_{2n}\Big)\,\mathrm{vec}(\Delta H) = s, \quad s = \begin{bmatrix} s_1 \\ s_2 \end{bmatrix}, \tag{3.3}$$

where $\otimes$ denotes the Kronecker product. We refer to Lancaster and Tismenetsky [18, Chap. 12] for properties of the vec operator and the Kronecker product. By linearity we have

$$\mathrm{vec}(\Delta H) = B\Delta p \tag{3.4}$$

for $B \in \mathbb{R}^{4n^2\times t}$ of full rank and where Δp is the t-vector of parameters defining ΔH. There exists a diagonal matrix D_1 depending on the structure of H (symmetric/skew-symmetric Hamiltonian/skew-Hamiltonian) such that

$$\|\Delta H\|_F = \|D_1\Delta p\|_2. \tag{3.5}$$

Let $y = D_1\Delta p$ and $Y = ([\,\widetilde{u}\quad \widetilde{v}\,]^T \otimes I_{2n}) \in \mathbb{R}^{4n\times 4n^2}$. Using (3.4) we can rewrite (3.3) as $YBD_1^{-1}y = s$ with $YBD^{-1} \in \mathbb{R}^{4n\times t}$. Then, using (3.5),

$$\mu(\widetilde{x},\widetilde{\lambda}) = \min_{y\,\in\,\mathbb{R}^t} \Big\{\|y\|_2/\|H\|_F \ :\ YBD_1^{-1}y = s\Big\}. \tag{3.6}$$

This shows that the structured backward error is given in terms of the minimal 2-norm solution to an underdetermined system. If the underdetermined system is consistent, then the minimal 2-norm solution is given in terms of the pseudo-inverse by $y = (YBD_1^{-1})^+ s$. In this case

$$\mu(\widetilde{x}, \widetilde{\lambda}) = \|(YBD_1^{-1})^+ s\|_2 / \|H\|_F. \tag{3.7}$$

When H is a symmetric structured Hamiltonian, we can assume that $\widetilde{\lambda}$ and $\widetilde{x}$ are real. Therefore $\widetilde{\nu} = 0$ and $\widetilde{v} = 0$ and from (3.2) we have $[\, s_1 \quad s_2 \,] = (\widetilde{\mu} I - H)\,[\, \widetilde{u} \quad \widetilde{v} \,]$. Applying the vec operation gives

$$s = \Big([\, \widetilde{u} \quad \widetilde{v} \,]^T \otimes I_{2n} \Big) \operatorname{vec}(\widetilde{\mu} I - H) = Y \operatorname{vec}(\widetilde{\mu} I - H).$$

As $\widetilde{\mu} I - H$ is also a symmetric structured Hamiltonian, we have by linearity that $\operatorname{vec}(\widetilde{\mu} I - H) = B p_{\widetilde{\mu}}$, where $p_{\widetilde{\mu}}$ is the t-vector of parameters defining $\widetilde{\mu} I - H$. Then $s = YBD_1^{-1} D_1 p_{\widetilde{\mu}}$ lies in the range of YBD_1^{-1}. Therefore, the underdetermined system in (3.6) is consistent for symmetric Hamiltonians and for symmetric skew-Hamiltonians. For a skew-symmetric Hamiltonian, we can again prove consistency for pure imaginary approximate eigenvalues and approximate eigenvectors of the form in Lemma 2.1. We have not been able to prove that the underdetermined system is consistent for the skew-symmetric skew-Hamiltonian case.

As the dependence on the parameters is linear, in the definition of the structured relative componentwise backward error $\nu(\widetilde{x}, \widetilde{\lambda})$, we have the equivalence

$$|\Delta H| \le \epsilon |H| \iff |\Delta p| \le \epsilon |p|.$$

Let $D_2 = \operatorname{diag}(p_i)$ and $\Delta p = D_2 q$. Then the smallest ϵ satisfying $|\Delta p| \le \epsilon |p|$ is $\epsilon = \|q\|_\infty$. The minimal ∞-norm solution of $YBD_2\, q = s$ can be approximated by minimizing in the 2-norm. We have

$$\nu(\widetilde{x}, \widetilde{\lambda}) \le \|(YBD_2)^+ s\|_2 \le \sqrt{t+n}\; \nu(\widetilde{x}, \widetilde{\lambda}).$$

By looking at each problem individually, it is possible to reduce the size of the underdetermined system. Nevertheless, solution of the system by standard techniques still takes $O(n^3)$ operations. In the next section, we show that by using a symplectic quasi-QR factorization of the approximate eigenvector and residual (or some appropriate parts) we can derive expressions for $\mu(\widetilde{x}, \widetilde{\lambda})$ that are cheaper to compute for all the structured Hamiltonians of interest except for skew-symmetric skew-Hamiltonians. First, we define a symplectic quasi-QR factorization.

3.3. Symplectic quasi-QR factorization. We define the symplectic quasi-QR factorization of an $2n \times m$ matrix A by

$$A = QT, \quad T = \begin{bmatrix} T_1 \\ T_2 \end{bmatrix}, \tag{3.8}$$

where Q is real symplectic orthogonal, $T_1 \in \mathbb{R}^{n \times m}$ is upper trapezoidal, and $T_2 \in \mathbb{R}^{n \times m}$ is strictly upper trapezoidal. Such a symplectic quasi-QR factorization has also been discussed by Bunse-Gerstner [3, Cor. 4.5(ii)].Before giving an algorithm to compute this symplectic quasi-QR factorization, we need to describe two types of elementary orthogonal symplectic matrices that can be used to zero selected components of a vector.

A *symplectic Householder* matrix $H \in \mathbb{R}^{2n\times 2n}$ is a direct sum of $n\times n$ Householder matrices:

$$H(k,v) = \begin{bmatrix} P(k,v) & 0 \\ 0 & P(k,v) \end{bmatrix},$$

where

$$P(k,v) = \begin{cases} \operatorname{diag}\left(I_{k-1}, I_{n-k+1} - \dfrac{2}{v^T v} vv^T\right) & \text{if} \quad v \neq 0, \\ I_n & \text{otherwise,} \end{cases}$$

and v is determined such that for a given $x \in \mathbb{R}^n$, $P(k,v)x = y$ with $y(k+1\colon n) = 0$.

A *symplectic Givens rotation* $G(k,\theta) \in \mathbb{R}^{2n\times 2n}$ is a Givens rotation where the rotation is performed in the plane $(k, k+n)$, $1 \le k \le n$, that is, $G(k,\theta)$ has the form

$$G(k,\theta) = \begin{bmatrix} C & S \\ -S & C \end{bmatrix}, \quad \text{where} \quad \begin{array}{l} C = \operatorname{diag}(I_{k-1}, \cos\theta, I_{n-k}), \\ S = \operatorname{diag}(0_{k-1}, \sin\theta, 0_{n-k}), \end{array} \tag{3.9}$$

where θ is chosen such that for a given $x \in \mathbb{R}^{2n}$, $G(k,\theta)x = y$ with $y_{n+k} = 0$.

We use a combination of these orthogonal transformations to compute our symplectic quasi-QR factorization: symplectic Householder matrices are used to zero large portions of a vector and symplectic Givens are used to zero single entries.

ALGORITHM 3.1 (symplectic quasi-QR factorization). *Given a matrix* $A = \left[\begin{smallmatrix} A_1 \\ A_2 \end{smallmatrix}\right]$ *with* $A_1, A_2 \in \mathbb{R}^{n\times m}$, *this algorithm computes the symplectic quasi-QR factorization* (3.8).

$Q = I_{2n}$
For $k = 1\colon \min(n-1, m)$
    Determine $H_k^1 = H(k,v)$ with $x = A_2 e_k$; $A \leftarrow H_k^1 A$
    Determine $G_k = G(k,\theta)$ with $x = A e_k$; $A \leftarrow G_k A$
    Determine $H_k^2 = H(k,v)$ with $x = A_1 e_k$; $A \leftarrow H_k^2 A$
    $Q \leftarrow Q H_k^1 G_k^T H_k^2$
End
If $m \ge n$
    Determine $G_n = G(n,\theta)$ with $x = A e_n$; $A \leftarrow G_n A$, $Q \leftarrow Q G_n^T$
End

We illustrate the procedure for a generic 6×4 matrix:

$$\left[\begin{array}{cccc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{array}\right] \xrightarrow{H_1^1} \left[\begin{array}{cccc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \end{array}\right] \xrightarrow{G_1} \left[\begin{array}{cccc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \end{array}\right] \xrightarrow{H_1^2}$$

$$\left[\begin{array}{cccc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ \hline 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \end{array}\right] \xrightarrow{H_2^1} \left[\begin{array}{cccc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ \hline 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & 0 & \times & \times \end{array}\right] \xrightarrow{G_2} \left[\begin{array}{cccc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ \hline 0 & \times & \times & \times \\ 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{array}\right] \xrightarrow{H_2^2}$$

$$\left[\begin{array}{cccc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & 0 & \times & \times \\ \hline 0 & \times & \times & \times \\ 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{array}\right] \xrightarrow{G_3} \left[\begin{array}{cccc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & 0 & \times & \times \\ \hline 0 & \times & \times & \times \\ 0 & 0 & \times & \times \\ 0 & 0 & 0 & \times \end{array}\right].$$

3.4. Symmetric Hamiltonian eigenproblems. Let $r = (\widetilde{\lambda} I - H)\widetilde{x} = \Delta H \widetilde{x}$ be the residual vector and $QR = [\widetilde{x}\ r]$ be the symplectic quasi-QR factorization (3.8) with Q symplectic orthogonal and

$$R = \begin{bmatrix} e_{11} & e_{12} \\ 0 & e_{22} \\ \vdots & 0 \\ & \vdots \\ & e_{n+1,2} \\ \vdots & \vdots \\ 0 & 0 \end{bmatrix} \in \mathbb{R}^{2n \times 2}.$$

We have $Q^T \Delta H Q Q^T \widetilde{x} = Q^T r$, which is equivalent to

$$\Delta \widetilde{H} \begin{bmatrix} e_{11} \\ 0 \\ \vdots \\ \\ \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} e_{12} \\ e_{22} \\ 0 \\ \vdots \\ e_{n+1,2} \\ \vdots \\ 0 \end{bmatrix}, \tag{3.10}$$

where $\Delta \widetilde{H} = Q^T \Delta H Q$ is still a symmetric Hamiltonian matrix. Equation (3.10) defines the first column of $\Delta \widetilde{H}$. As $|e_{11}| = \|\widetilde{x}\| \neq 0$, we have $\Delta \tilde{h}_{11} = e_{12}/e_{11},\ \Delta \tilde{h}_{21} = e_{22}/e_{11},\ \Delta \tilde{h}_{n+1,1} = e_{n+1,2}/e_{11}$, and $\Delta \tilde{h}_{k,1} = 0$ for $k \neq 1, 2, n+1$. Let $\Delta \widetilde{E} = (\Delta \widetilde{E})^T$

and $\Delta\widetilde{F} = (\Delta\widetilde{F})^T$ be such that

$$\Delta\widetilde{E} = \frac{1}{e_{11}} \begin{bmatrix} e_{12} & e_{22} & 0 & \cdots & 0 \\ e_{22} & \times & \cdots & \cdots & \times \\ 0 & \vdots & & & \vdots \\ \vdots & \vdots & & & \vdots \\ 0 & \times & \cdots & \cdots & \times \end{bmatrix}, \quad \Delta\widetilde{F} = \frac{1}{e_{11}} \begin{bmatrix} e_{n+1,2} & 0 & 0 & \cdots & 0 \\ 0 & \times & \cdots & \cdots & \times \\ 0 & \vdots & & & \vdots \\ \vdots & \vdots & & & \vdots \\ 0 & \times & \cdots & \cdots & \times \end{bmatrix},$$

where the $\times$'s are arbitrary real coefficients. Then, any symmetric Hamiltonian of the form

$$\Delta H = Q \begin{bmatrix} \Delta\widetilde{E} & \Delta\widetilde{F} \\ \Delta\widetilde{F} & -\Delta\widetilde{E} \end{bmatrix} Q^T$$

satisfies $(H + \Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x}$. The Frobenius norm of ΔH is minimized by setting the $\times$'s to zero in the definition of $\Delta\widetilde{E}$ and $\Delta\widetilde{F}$. We obtain the following lemma.

LEMMA 3.2. *The backward error of an approximate eigenpair of a symmetric Hamiltonian eigenproblem is given by*

$$\mu(\widetilde{x}, \widetilde{\lambda}) = \frac{2}{|e_{11}|} \sqrt{\frac{e_{12}^2}{2} + e_{22}^2 + \frac{e_{n+1,2}^2}{2}} \Bigg/ \|H\|_F,$$

where $R = (e_{ij}) = Q^T[\widetilde{x}\ r]$ is the quasi-triangular factor in the symplectic quasi-QR factorization of $[\widetilde{x}\ r]$ with $r = (\widetilde{\lambda}I - H)\widetilde{x}$. We also have

$$\mu(\widetilde{x}, \widetilde{\lambda}) = \frac{\sqrt{2\|Q^T r\|^2 + 2(e_2^T Q^T r)^2}}{\|Q^T \widetilde{x}\|} \Bigg/ \|H\|_F,$$

where e_2 is the second unit vector.

3.5. Skew-symmetric Hamiltonian eigenproblems. For skew-symmetric Hamiltonian eigenproblems the technique developed in section 3.4 needs to be modified as in this case $r, \widetilde{x}$ are complex vectors and we want to define a real skew-symmetric Hamiltonian perturbation

$$\Delta H = \begin{bmatrix} \Delta E & \Delta F \\ -\Delta F & \Delta E \end{bmatrix}, \quad \Delta E = -\Delta E^T,\ \Delta F = \Delta F^T$$

so that $(H + \Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x}$.

In the definition of the structured backward error (3.1), we now assume that $\widetilde{\lambda}$ is pure imaginary and that $\widetilde{x}$ has the form $[\,\widetilde{z}^T \quad \pm i\widetilde{z}^T\,]^T$ (see Lemma 2.1). Taking the plus sign in $\widetilde{x}$, the equation $(H + \Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x}$ can be written as

$$\Delta E\widetilde{z} + i\Delta F\widetilde{z} = (\widetilde{\lambda}I - E)\widetilde{z} - iF\widetilde{z}, \tag{3.11}$$

$$-\Delta F\widetilde{z} + i\Delta E\widetilde{z} = F\widetilde{z} - i(\widetilde{\lambda}I - E)\widetilde{z}. \tag{3.12}$$

Multiplying (3.12) by $-i$ gives (3.11). Hence, we carry out the analysis with (3.11) only. Setting $\widetilde{\lambda} = i\widetilde{\mu},\ \widetilde{\mu} \in \mathbb{R},\ \widetilde{z} = \widetilde{u} + i\widetilde{v}$ in (3.11) and equating real and imaginary parts yields

$$\Delta E\widetilde{u} - \Delta F\widetilde{v} = -\widetilde{\mu}\widetilde{v} + F\widetilde{v} - E\widetilde{u},$$
$$\Delta E\widetilde{v} + \Delta F\widetilde{u} = \widetilde{\mu}\widetilde{u} - E\widetilde{v} - F\widetilde{u},$$

which is equivalent to $\Delta H w = s$ with $w = [\,\widetilde{u}^T \quad -\widetilde{v}^T\,]^T$ and $s = (\widetilde{\mu}J - H)w$. Using $x^T E x = 0$ and $F^T = F$, we show that w and s are orthogonal:

$$\begin{aligned} w^T s &= [\,\widetilde{u}^T \quad -\widetilde{v}^T\,]\,(\widetilde{\mu}J - H)\begin{bmatrix} \widetilde{u} \\ -\widetilde{v} \end{bmatrix} \\ &= -\widetilde{u}^T E\widetilde{u} - \widetilde{\mu}\widetilde{u}^T\widetilde{v} + \widetilde{u}^T F\widetilde{v} + \widetilde{\mu}\widetilde{v}^T\widetilde{u} - \widetilde{v}^T E\widetilde{v} - \widetilde{v}^T F\widetilde{u} = 0. \end{aligned}$$

For the other choice of sign with $\widetilde{x} = [\,\widetilde{z}^T \quad -i\widetilde{z}^T\,]^T$, the equation $(H+\Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x}$ is equivalent to $\Delta H w = s$ with $w = [\,\widetilde{u}^T \quad \widetilde{v}^T\,]^T$ and $s = -(\widetilde{\mu}J + H)w$. Here again, we can show that $w^T s = 0$.

We can now carry on the analysis as in section 3.4. Let $QR = [w\ s]$ be the symplectic quasi-QR factorization of $[w\ s]$. As $w^T s = 0$, we have that $e_{12} = 0$. We obtain ΔH by solving the underdetermined system

$$\Delta\widetilde{H}\begin{bmatrix} e_{11} \\ 0 \\ \vdots \\ \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ e_{22} \\ 0 \\ \vdots \\ e_{n+1,2} \\ \vdots \\ 0 \end{bmatrix}, \quad \Delta H = Q\Delta\widetilde{H}Q^T. \tag{3.13}$$

Lemma 3.3. *The backward error of an approximate eigenpair $(\widetilde{x},\widetilde{\lambda})$ of a skew-symmetric Hamiltonian eigenproblem with $\widetilde{\lambda}$ pure imaginary and $\widetilde{x}$ of the form $\widetilde{x} = [\,\widetilde{z}^T \quad \pm i\widetilde{z}^T\,]^T$ is given by*

$$\mu(\widetilde{x},\widetilde{\lambda}) = \left(\frac{2}{|e_{11}|}\sqrt{e_{22}^2 + \frac{e_{n+1,2}^2}{2}}\right)\Big/\|H\|_F,$$

where $R = Q^T[w\ s]$ is the quasi-triangular factor in the symplectic quasi-QR factorization of $[w\ s]$ with $\widetilde{z} = \widetilde{u} + i\widetilde{v}$ and

$$w = \begin{cases} [\,\widetilde{u}^T \quad -\widetilde{v}^T\,]^T & \text{if } \widetilde{x} = [\,\widetilde{z}^T \quad i\widetilde{z}^T\,]^T, \\ [\,\widetilde{u}^T \quad \widetilde{v}^T\,]^T & \text{otherwise,} \end{cases} \qquad s = \begin{cases} (\widetilde{\mu}J - H)w & \text{if } \widetilde{x} = [\,\widetilde{z}^T \quad i\widetilde{z}^T\,]^T, \\ -(\widetilde{\mu}J + H)w & \text{otherwise.} \end{cases}$$

We also have

$$\mu(\widetilde{x},\widetilde{\lambda}) = \left(\frac{\sqrt{2\|Q^T s\|_2^2 + 2|e_2^T Q^T s|^2}}{\|Q^T\widetilde{x}\|_2}\right)\Big/\|H\|_F,$$

where e_2 is the second unit vector.

3.6. Symmetric skew-Hamiltonian eigenproblems. The analysis for symmetric skew-Hamiltonian eigenproblems is similar to that in section 3.4. The only difference comes from noting that

$$\begin{aligned} (J\widetilde{x})^T r &= [\,\widetilde{x}_2^T \quad -\widetilde{x}_1^T\,]\begin{bmatrix} \widetilde{\lambda}I - E & -F \\ F & \widetilde{\lambda}I - E \end{bmatrix}\begin{bmatrix} \widetilde{x}_1 \\ \widetilde{x}_2 \end{bmatrix} \\ &= \widetilde{x}_2^T(\widetilde{\lambda}I - E)\widetilde{x}_1 - \widetilde{x}_2^T F\widetilde{x}_2 - \widetilde{x}_1^T F\widetilde{x}_1 - \widetilde{x}_1^T(\widetilde{\lambda}I - E)\widetilde{x}_2 = 0 \end{aligned}$$

using $v^T F v = 0$ and $\widetilde{x}_1^T(\widetilde{\lambda} I - E)\widetilde{x}_2 = \widetilde{x}_2^T(\widetilde{\lambda} I - E)\widetilde{x}_1$. Instead of computing a symplectic quasi-QR factorization of $[\widetilde{x}\ r]$, we compute a symplectic quasi-QR factorization of $[J\widetilde{x}\ r]$ in order to introduce one more zero in the triangular factor R. We summarize the result in the next lemma.

LEMMA 3.4. *The backward error of an approximate eigenpair of a symmetric skew-Hamiltonian eigenproblem is given by*

$$\mu(\widetilde{x}, \widetilde{\lambda}) = \frac{2}{|e_{11}|}\sqrt{e_{22}^2 + \frac{e_{n+1,2}^2}{2}} \Big/ \|H\|_F, \tag{3.14}$$

where $R = Q^T[J\widetilde{x}\ r]$ is the quasi-triangular factor in the symplectic quasi-QR factorization of $[J\widetilde{x}\ r]$ with $r = (\widetilde{\lambda} I - H)\widetilde{x}$. We also have

$$\mu(\widetilde{x}, \widetilde{\lambda}) = \left(\frac{\sqrt{2\|Q^T r\|_2^2 + 2|e_2^T Q^T r|^2}}{\|Q^T J\widetilde{x}\|_2}\right) \Big/ \|H\|_F.$$

3.7. Comments. Lemmas 3.2–3.4 provide an explicit formula for the backward error that can be computed in $O(n^2)$ operations.

For skew-symmetric skew-Hamiltonian matrices H, the eigenvectors are complex with no particular structure. The constraint $(H + \Delta H)\widetilde{x} = \widetilde{\lambda}\widetilde{x}$ in (3.1) can be written in the form $\Delta H[\Re(\widetilde{x}), \Im(\widetilde{x})] = [\Re(r), \Im(r)]$, where $r = (H - \widetilde{\lambda} I)\widetilde{x}$ is the residual. We were unable to explicitly construct matrices ΔH satisfying this constraint via a symplectic QR factorization of $[\Re(\widetilde{x}), \Im(\widetilde{x}), \Re(r), \Im(r)]$. Thus, in this case, we have to use the approach described in section 3.2 to compute $\mu(\widetilde{x}, \widetilde{\lambda})$, which has the drawback that it requires $O(n^3)$ operations.

4. Algorithms for Hamiltonian eigenproblems.

A simple but inefficient approach to solve structured Hamiltonian eigenproblems is to use the (symmetric or unsymmetric as appropriate) QR algorithm on the $2n \times 2n$ structured Hamiltonian matrix. This approach is computationally expensive and uses $4n^2$ storage locations. Moreover, the QR algorithm does not use symplectic orthogonal transformations and is therefore not structure-preserving.

Benner, Merhmann, and Xu's method [1] for computing the eigenvalues and invariant subspaces of a real Hamiltonian matrix uses the relationship between the eigenvalues and invariant subspaces of H and an extended $4n \times 4n$ Hamiltonian matrix. Their algorithm is structure-preserving for the extended Hamiltonian matrix but is not structure-preserving for H. Therefore, it is not strongly backward stable in the sense of this paper.

4.1. QR-like algorithms.

Bunse-Gerstner, Byers, and Mehrmann [5] provide a chart of numerical methods for structured eigenvalue problems, most of them based on QR-like algorithms. In this section, we describe their recommended algorithms for our structured Hamiltonian eigenproblems. In the limited case where $\text{rank}(F) = 1$, Byer's Hamiltonian QR algorithm [6] based on symplectic orthogonal transformations yields a strongly backward stable algorithm.

For symmetric Hamiltonian eigenproblems, the quaternion QR algorithm [4] is suggested. The quaternion QR algorithm is an extension of the Francis QR algorithm for complex or real matrices to quaternion matrices. This algorithm uses exclusively quaternion unitary similarity transformations so that it is backward stable. Compared with the standard QR algorithm for symmetric matrices, this algorithm cuts the storage and work requirements approximately in half. However, its implementation requires quaternion arithmetic and it is not clear whether it is strongly backward stable.

A skew-symmetric Hamiltonian H is first reduced via symplectic orthogonal transformations to block antidiagonal form $\left[\begin{smallmatrix} 0 & T \\ -T & 0 \end{smallmatrix}\right]$, where the blocks are symmetric tridiagonal. The complete solution is obtained via the symmetric QR algorithm applied to T. The whole algorithm is strongly backward stable as it uses only real symplectic orthogonal transformations that are known to be backward stable.

For symmetric skew-Hamiltonian problems, the use of the "X-trick" is suggested:

$$X^T H X = \begin{bmatrix} E - iF & 0 \\ 0 & E + iF \end{bmatrix} \quad \text{with} \quad X = \frac{1}{\sqrt{2}} \begin{bmatrix} I & I \\ -iI & iI \end{bmatrix}. \tag{4.1}$$

The eigenvalues of H are computed from the eigenvalue of the Hermitian matrices $E - iF$ or $E + iF$, using the Hermitian QR algorithm for instance. One drawback of this approach is that it uses complex arithmetic and does not provide a real symplectic orthogonal eigenvector basis. Hence the algorithm does not preserve the "realness" of the original matrix.

Finally, for the skew-symmetric skew-Hamiltonian case, H is reduced to block-diagonal form via a finite sequence of symplectic orthogonal transformations. The blocks are themselves tridiagonal and skew-symmetric. Then Paardekooper's Jacobi algorithm [22] or the algorithm in [11] for skew-symmetric tridiagonal matrices can be used to obtain the complete solution. The whole algorithm is strongly backward stable.

4.2. Jacobi-like algorithms. Byers [7] adapted the nonsymmetric Jacobi algorithm [24] to the special structure of Hamiltonian matrices. The Hamiltonian Jacobi algorithm based on symplectic Givens rotations and symplectic double Jacobi rotations of the form $J \otimes I_{2n}$, where J is a 2×2 Jacobi rotation, preserves the Hamiltonian structure. This Jacobi algorithm, when it converges, builds a Hamiltonian Schur decomposition [7, Thm. 1]. For symmetric H, this Jacobi algorithm converges to the canonical form $\left[\begin{smallmatrix} D & 0 \\ 0 & -D \end{smallmatrix}\right]$ and is strongly backward stable. For skew-symmetric Hamiltonian H, this Jacobi algorithm does not converge as the symplectic orthogonal canonical form for H is not Hamiltonian triangular.

Recently, Faßbender, Mackey, and Mackey [10] developed Jacobi algorithms for structured Hamiltonian eigenproblems that preserve the structure and produce a complete basis of symplectic orthogonal eigenvectors for the two symmetric cases and a symplectic orthogonal basis for all the real invariant subspaces for the two skew-symmetric cases. These Jacobi algorithms are based on the direct solution of 4×4, and in one case 8×8, subproblems using appropriate transformations. The algorithms work entirely in real arithmetic. Note that "realness" of the initial matrix can be viewed as additional structure that these Jacobi algorithms preserve. We give a unified description of these Jacobi-like algorithms for the four classes of structured Hamiltonian eigenproblems under consideration.

Let $H \in \mathbb{R}^{2n \times 2n}$ be a structured Hamiltonian matrix (see Table 2.1 and 2.2). These Jacobi methods attempt to reduce the quantity (off-diagonal norm)

$$\text{off}(H) = \sqrt{\sum_{i=1}^{2n} \sum_{j \in \mathcal{S}} |h_{ij}|^2},$$

where $\mathcal{S}$ is a set of indices depending on the structure of the problem using a sequence of symplectic orthogonal transformations $H \leftarrow SHS^T$ with $S \in \mathbb{R}^{2n\times 2n}$. The aim is that H converges to its canonical form. In the following, we note $A_{i,j,i+n,j+n}$ the restriction to the $(i, j, i+n, j+n)$ plane of A.

ALGORITHM 4.1. *Given a structured Hamiltonian matrix* $H \in \mathbb{R}^{2n\times 2n}$ *and a tolerance* tol > 0*, this algorithm overwrites* H *with its approximate canonical form* PHP^T*, where* P *is symplectic orthogonal and* $\text{off}(PHP^T) \le \text{tol}\,\|H\|_F$.

$P = I_{2n}$
$\epsilon = \text{tol}\,\|H\|_F$
while $\text{off}(H) > \epsilon$
 Choose (i, j)
 Compute a symplectic orthogonal S
  such that $(SHS^T)_{i,j,i+n,j+n}$ is in canonical form.
 $H \leftarrow SHS^T$ preserving structure
 $P \leftarrow SP$ preserving structure
end

Note that the pair (i, j) uniquely determines a 4×4 principal submatrix

$$H_{i,j,n+i,n+j} = \begin{bmatrix} h_{ii} & h_{ij} & h_{i,i+n} & h_{i,j+n} \\ h_{ji} & h_{jj} & h_{j,i+n} & h_{j,j+n} \\ h_{i+n,i} & h_{i+n,j} & h_{i+n,i+n} & h_{i+n,j+n} \\ h_{j+n,i} & h_{j+n,j} & h_{j+n,i+n} & h_{j+n,j+n} \end{bmatrix} \tag{4.2}$$

that also inherits the Hamiltonian or skew-Hamiltonian structure together with the symmetry or skew-symmetry property. There are many ways of choosing the indices (i, j) but this choice does not affect the rest of the analysis. We refer to $n(n-1)/2$ updates as a sweep. Each sweep must be complete, that is, every part of the matrix must be reached. We see immediately that any complete sweep of the $(1, 1)$ block of H consisting of 2×2 principal submatrices generates a corresponding complete sweep of H.

For each 4×4 target submatrix, a symplectic orthogonal matrix that directly computes the corresponding canonical form is constructed and embedded into the $2n \times 2n$ identity matrix in the same way that the 4×4 target has been extracted.

For skew-symmetric skew-Hamiltonians, the 4×4 based Jacobi algorithm does not converge. The aim of these Jacobi algorithms is to move the weight to the diagonal of either the diagonal blocks or off-diagonal blocks. That cannot be done for a skew-symmetric skew-Hamiltonian because these diagonals are zero. There is no safe place where the norm of the target submatrix can be kept. However, if an 8×8 skew-symmetric skew-Hamiltonian problem is solved instead, the 2×2 diagonal blocks of H become a safe place for the norm of target submatrices and the resulting 8×8 based Jacobi algorithm is expected to converge. The complete sweep is defined by partitioning E in 2×2 blocks, leaving 2×1 and 1×2 blocks along the rightmost and lower edges when n is odd. Hence, in this case we must also be able to directly solve 6×6 subproblems.

Immediately, we see that the difficult part in deriving these algorithms is to define the appropriate symplectic orthogonal transformation S that computes the canonical form of the restriction to the $(i, j, i+n, j+n)$ plane of H. Faßbender, Mackey, and Mackey [10] show that by using a quaternion representation of the 4×4 symplectic orthogonal group, as well as 4×4 Hamiltonian and skew-Hamiltonian matrices in the tensor square of the quaternion algebra, we can define and construct 4×4 symplectic orthogonal matrices R that do the job. These transformations are based on rotations of $\mathbb{P} \cong \mathbb{R}^3$, the subspace of pure quaternions.

We need to give all the required transformations in a form suitable for rounding error analysis and also to facilitate the description of the structure preserving Jacobi algorithms. We start by defining two types of quaternion rotations. This enables us to encode the formulas in [10] into one. Let $e_s \neq e_1$ be a standard basis vector of $\mathbb{R}^4$ and $p \in \mathbb{R}^4$ such that $p \neq 0$, $e_1^T p = 0$ (p is a pure quaternion), and $p/\|p\|_2 \neq e_s$. Let

$$x^T = \|p\|_2 e_1^T + e_s^T \begin{bmatrix} 0 & -p_2 & -p_3 & -p_4 \\ p_2 & 0 & p_4 & -p_3 \\ p_3 & -p_4 & 0 & p_2 \\ p_4 & p_3 & -p_2 & 0 \end{bmatrix}. \tag{4.3}$$

We define the left quaternion rotation by

$$Q_L(p,s) = \frac{1}{\|x\|_2} \begin{bmatrix} x_1 & -x_2 & -x_3 & -x_4 \\ x_2 & x_1 & -x_4 & x_3 \\ x_3 & x_4 & x_1 & -x_2 \\ x_4 & -x_3 & x_2 & x_1 \end{bmatrix}. \tag{4.4}$$

Q_L is symplectic orthogonal and not difficult to compute. We have $x_1 = \|p\|_2 + p_s$ and the other components of x are just permutations of the coordinates of p.

We define the right quaternion rotation by

$$Q_R(p,s) = \frac{1}{\|x\|_2} \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ -x_2 & x_1 & -x_4 & x_3 \\ -x_3 & x_4 & x_1 & -x_2 \\ -x_4 & -x_3 & x_2 & x_1 \end{bmatrix}. \tag{4.5}$$

The matrix Q_R is orthogonal. It is symplectic when $s \neq 3$ and $x_2 = x_4 = 0$.

Let $p = [\, p_1 \quad p_2 \quad p_3 \quad p_4 \,]^T \in \mathbb{R}^4$ be nonzero. Following [10], we define the 4×4 symplectic orthogonal Givens rotation associated with p by

$$G(p) = \frac{1}{\|p\|_2} \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ -p_2 & p_1 & p_4 & -p_3 \\ -p_3 & -p_4 & p_1 & p_2 \\ -p_4 & p_3 & -p_2 & p_1 \end{bmatrix}. \tag{4.6}$$

We now have all the tools needed to define the symplectic orthogonal transformations that directly compute the canonical form for each of the 4×4 structured Hamiltonian eigenproblems of interest. We refer to [10] for more details about how these transformations have been derived.

4.2.1. Symmetric Hamiltonian. Let $H \in \mathbb{R}^{4\times 4}$ be a symmetric Hamiltonian matrix. The canonical form of H is obtained in two steps: first H is reduced to 2×2 block diagonal form and then the complete diagonalization is obtained by using a double Jacobi rotation.

For the first step we consider the singular value decomposition of the 3×3 matrix

$$A = \begin{bmatrix} \frac{1}{2}(h_{11} + h_{22}) & 0 & \frac{1}{2}(h_{13} + h_{24}) \\ h_{14} & 0 & -h_{12} \\ \frac{1}{2}(h_{24} - h_{13}) & 0 & \frac{1}{2}(h_{11} - h_{22}) \end{bmatrix} = U \Sigma V^T.$$

Let u_1 and v_1 be the left and right singular vectors corresponding to the largest singular value σ_1 and let $u = \left[\begin{smallmatrix} 0 \\ u_1 \end{smallmatrix}\right]$, $v = \left[\begin{smallmatrix} 0 \\ v_1 \end{smallmatrix}\right]$. We have $A^T u_1 = \sigma_1 v_1$ and $e_2^T A^T u_1 = 0$

so that $e_2^T v_1 = 0$. Hence, for $s = 2$ and $p = v$, the vector x in (4.3) is such that $x_2 = x_4 = 0$, which implies that the right quaternion rotation $Q_R(v, 2)$ is symplectic orthogonal. As shown in [10], the product $Q = Q_L(u, 2)Q_R(v, 2)$ block diagonalizes H, that is, $QHQ^T = \operatorname{diag}(\widetilde{E}, -\widetilde{E})$. Complete diagonalization is obtained by using a double Jacobi rotation $J(\theta) \otimes I_2$, where θ is chosen such that $J(\theta) = \left[\begin{smallmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{smallmatrix}\right]$ diagonalizes $\widetilde{E}$.

In summary, the symplectic orthogonal transformation S used in Algorithm 4.1 is equal to the identity matrix except in the $(i, j, n+i, n+j)$ plane, where the $(i, j, n+i, n+j)$-restriction matrix is given by

$$S_{i,j,n+i,n+j} = \big(J(\theta) \otimes I_2\big) Q_L(u, 2) Q_R(v, 2).$$

4.2.2. Skew-symmetric Hamiltonian. Let $H \in \mathbb{R}^{4\times 4}$ be a skew-symmetric Hamiltonian matrix and let $p \in \mathbb{R}^4$ be defined from the elements of H by

$$p = [\,0, \quad h_{21}, \quad \tfrac{1}{2}(h_{31} - h_{42}), \quad h_{41}\,]^T.$$

It is easy to verify that for $S = Q_L(p, 3)$,

$$SHS^T = \begin{bmatrix} 0 & 0 & -\|p\|_2 + b & 0 \\ 0 & 0 & 0 & \|p\|_2 + b \\ \|p\|_2 - b & 0 & 0 & 0 \\ 0 & -\|p\|_2 - b & 0 & 0 \end{bmatrix},$$

where $b = \frac{1}{2}(h_{13} + h_{24})$.

4.2.3. Symmetric skew-Hamiltonian. Let $H \in \mathbb{R}^{4\times 4}$ be a symmetric skew-Hamiltonian matrix and let $p \in \mathbb{R}^4$ be defined from the elements of H by

$$p = [\,0, \quad -h_{14}, \quad \tfrac{1}{2}(h_{11} - h_{22}), \quad h_{12}\,]^T.$$

Then, $S = Q_L(p, 3)$ diagonalizes H and

$$SHS^T = \begin{bmatrix} b + \|p\|_2 & 0 & 0 & 0 \\ 0 & b - \|p\|_2 & 0 & 0 \\ 0 & 0 & b + \|p\|_2 & 0 \\ 0 & 0 & 0 & b - \|p\|_2 \end{bmatrix},$$

where $b = \frac{1}{2}(h_{11} + h_{22})$.

4.2.4. Skew-symmetric skew-Hamiltonian. For the convergence of the Jacobi algorithm to be possible we need to solve an 8×8 subproblem. The matrix $H \in \mathbb{R}^{8\times 8}$ is block diagonalized with three 4×4 symplectic Givens rotations of the form (4.6) and one symplectic Givens rotation of the form (3.9). Let G be the product of these rotations. We have

$$GHG^T = \begin{bmatrix} \widetilde{E} & 0 \\ 0 & -\widetilde{E} \end{bmatrix}, \tag{4.7}$$

where $\widetilde{E} \in \mathbb{R}^{4\times 4}$ is tridiagonal and skew-symmetric. The complete 2×2 block-diagonalization is obtained by directly transforming $\widetilde{E}$ into its real Schur form as follows. In [20], Mackey showed that the transformation $Q = Q_L(q_1, 2)Q_R(q_2, 2)$ with

$$q_1 = [\,0, -\tfrac{1}{2}(\tilde{e}_{12} + \tilde{e}_{34}), 0, -\tfrac{1}{2}\tilde{e}_{23}\,]^T, \qquad q_2 = [\,0, \tfrac{1}{2}(\tilde{e}_{12} - \tilde{e}_{34}), 0, -\tfrac{1}{2}\tilde{e}_{23}\,]^T$$

directly computes the real Schur form of $\widetilde{E}$, that is,

$$Q\widetilde{E}Q^T = \begin{bmatrix} B_1 & 0 \\ 0 & B_2 \end{bmatrix} \quad \text{with} \quad B_1 = \begin{bmatrix} 0 & s_2 - s_1 \\ s_1 - s_2 & 0 \end{bmatrix}, \quad B_2 = \begin{bmatrix} 0 & -s_1 - s_2 \\ s_1 + s_2 & 0 \end{bmatrix},$$

where $s_1 = \|q_1\|_2$ and $s_2 = \|q_2\|_2$. Then $S = (Q \otimes I_2)G$ is the symplectic orthogonal transformation that computes the real Schur form of the 8×8 skew-symmetric skew-Hamiltonian H:

$$SHS^T = \begin{bmatrix} B_1 & 0 & 0 & 0 \\ 0 & B_2 & 0 & 0 \\ 0 & 0 & -B_1 & 0 \\ 0 & 0 & 0 & -B_2 \end{bmatrix}.$$

When n is odd, we have to solve a 6×6 subproblem for each complete sweep of the Jacobi algorithm. As for the 8×8 case, the 6×6 skew-symmetric skew-Hamiltonian H is first reduced to the form (4.7), where $\widetilde{E} \in \mathbb{R}^{3\times 3}$ is tridiagonal and skew-symmetric. This is done by using just one 4×4 symplectic Givens rotation followed by one 2×2 symplectic Givens rotation. Let

$$\widetilde{E}_{aug} = \begin{bmatrix} 0 & 0 \\ 0 & \widetilde{E} \end{bmatrix} \in \mathbb{R}^{4\times 4}$$

and $q = [\,0 \quad -\frac{1}{2}\tilde{e}_{23}, 0, -\frac{1}{2}\tilde{e}_{12}\,]^T$. Then $Q = Q_L(q,4)Q_R(q,4)$ computes directly the real Schur form of $\widetilde{E}_{aug}$. Moreover, we have $e_1^T Q e_1 = 1$, so that $Q = \mathrm{diag}(1, \widetilde{Q})$ and $\widetilde{Q}\widetilde{E}\widetilde{Q}^T = B$, where

$$B = \begin{bmatrix} 0 & -b & 0 \\ b & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{with} \quad b = 2\|q\|_2.$$

5. Error analysis of the Jacobi algorithms. In floating point arithmetic, Algorithm 4.1 computes an approximate canonical form $\widehat{T}$ such that

$$\widehat{T} =: P(H + \Delta H)P^T,$$

where P is symplectic orthogonal, and an approximate basis of symplectic orthogonal eigenvectors $\widehat{P}$. We want to derive bounds for $\|\Delta H\|, \|\widehat{P}\widehat{P}^T - I\|$, and $\|\widehat{P}J\widehat{P}^T - J\|$.

5.1. Preliminaries. We use the standard model for floating point arithmetic [16]

$$fl(x \,\mathrm{op}\, y) = (x \,\mathrm{op}\, y)(1+\delta), \quad |\delta| \le u, \quad \mathrm{op} = +, -, *, /, \tag{5.1}$$

where u is the unit roundoff. We assume that (5.1) holds also for the square roots operation. To keep track of the higher terms in u we make use of the following result [16, Lem. 3.1].

LEMMA 5.1. *If $|\delta_i| \le u$ and $\rho_i = \pm 1$ for $i = 1{:}n$, and $nu < 1$, then*

$$\prod_{i=1}^{n} (1+\delta_i)^{\rho_i} = 1 + \theta_n, \qquad \text{where} \qquad |\theta_n| \le \frac{nu}{1-nu} =: \gamma_n.$$

We define

$$\widetilde{\gamma}_k = \frac{pku}{1 - pku},$$

where p denotes a small integer constant whose value is unimportant. In the following, computed quantities will be denoted by hats.

First, we consider the construction of a 4×4 Givens rotation and left and right quaternion rotations.

LEMMA 5.2. *Let a 4×4 Givens rotation $G = G(p)$ be constructed according to (4.6) with $p \in \mathbb{R}^4$. Then the computed $\widehat{G}$ satisfies $|\widehat{G} - G| \leq \gamma_5 |G|$.*

Proof. This result is a straightforward extension of Lemma 18.6 in [16] concerning 2×2 Givens rotations. ∎

The rounding error properties of right and left quaternion rotations require more attention. When $p_s < 0$, the computation of $\|p\|_2 + p_s$ and therefore the computation of $Q_L(p,s)$ or $Q_R(p,s)$ is affected by cancellation. This problem can be overcome by using another formula as shown in the next lemma.

LEMMA 5.3. *Let 4×4 left and right quaternion rotations $Q_L = Q_L(p,s)$ and $Q_R = Q_R(p,s)$ be constructed according to*

$$Q_L(p,s) = \frac{1}{\sqrt{2\|p\|_2 \alpha}} \begin{bmatrix} \alpha & -x_2 & -x_3 & -x_4 \\ x_2 & \alpha & -x_4 & x_3 \\ x_3 & x_4 & \alpha & -x_2 \\ x_4 & -x_3 & x_2 & \alpha \end{bmatrix}, \tag{5.2}$$

$$Q_R(p,s) = \frac{1}{\sqrt{2\|p\|_2 \alpha}} \begin{bmatrix} \alpha & x_2 & x_3 & x_4 \\ -x_2 & \alpha & -x_4 & x_3 \\ -x_3 & x_4 & \alpha & -x_2 \\ -x_4 & -x_3 & x_2 & \alpha \end{bmatrix}, \tag{5.3}$$

where

$$[\,x_2 \quad x_3 \quad x_4\,] = \begin{cases} [\,0 \quad p_4 \quad -p_3\,] & \text{if } s = 2, \\ [\,-p_4 \quad 0 \quad p_2\,] & \text{if } s = 3, \\ [\,p_3 \quad -p_2 \quad 0\,] & \text{if } s = 4 \end{cases}$$

and

$$\alpha = \begin{cases} \|p\|_2 + p_s \ \ \text{if } p_s \geq 0, \\ \displaystyle\sum_{i=2, i \neq s}^{4} \left(p_i^2\right) / (\|p\|_2 - p_s) \quad \text{otherwise}, \end{cases} \tag{5.4}$$

with $p \in \mathbb{R}^4$ given. Then the computed $\widehat{Q_L}$ and $\widehat{Q_R}$ satisfy

$$|\widehat{Q_L} - Q_L| \leq \widetilde{\gamma}_1 |Q_L|, \quad |\widehat{Q_R} - Q_R| \leq \widetilde{\gamma}_1 |Q_R|.$$

Proof. It is straightforward to verify that the expressions for $Q_L(p,s)$ and $Q_R(p,s)$ in (5.2) and (5.3) agree with the definitions in (4.4) and (4.5).

We have $fl(\|p\|_2) = \|p\|_2(1 + \theta_4)$ with $|\theta_4| \leq \gamma_4$. If $p_s \geq 0$,

$$fl(\alpha) = (\|p\|_2(1 + \theta_4) + p_s)(1 + \delta) = \|p\|_2(1 + \theta_5) + p_s(1 + \delta).$$

As $p_s \geq 0$, there exists θ_5' such that $fl(\alpha) = (\|p\|_2 + p_s)(1+\theta_5')$ with $|\theta_5'| \leq \gamma_5$. If $p_s < 0$, using the same argument we have

$$fl(\|p\|_2 - p_s) = (\|p\|_2 - p_s)(1+\theta_5), \quad |\theta_5| \leq \gamma_5.$$

We also have

$$fl\left(\sum_{i=2,i\neq s}^{4} p_i^2\right) = \left(\sum_{i=2,i\neq s}^{4} p_i^2\right)(1+\theta_3).$$

Using [16, Lem. 3.3] we have

$$fl(\alpha) = \frac{(\sum_{i=2,i\neq s}^{4} p_i^2)(1+\theta_3)}{(\|p\|_2 - p_s)(1+\theta_5)}(1+\delta) = \alpha(1+\theta_9), \quad |\theta_9| \leq \gamma_9,$$

and

$$fl(\sqrt{2\|p\|_2\alpha}) = \sqrt{2\|p\|_2\alpha}(1+\theta_{16}), \quad |\theta_{16}| \leq \gamma_{16}.$$

Hence, we certainly have

$$fl((Q_L)_{ij}) \leq (Q_L)_{ij}(1+\theta_{26}), \quad |\theta_{26}| \leq \gamma_{26} \leq \widetilde{\gamma}_1. \qquad \square$$

In the following we use the term *elementary symplectic orthogonal matrix* to describe any double Givens rotation, 4×4 Givens rotation, or left or right quaternion rotation that is embedded as a principal submatrix of the identity matrix $I \in \mathbb{R}^{2n\times 2n}$.

We have proved that any computed elementary symplectic orthogonal matrix $\widehat{P} = fl(P)$ used by the Jacobi algorithm satisfies a bound of the form

$$|\widehat{P} - P| \leq \widetilde{\gamma}_1 |P|. \tag{5.5}$$

LEMMA 5.4. *Let $x \in \mathbb{R}^{2n\times 2n}$ and consider the computation of $y = \widehat{P}x$, where $\widehat{P}$ is a computed elementary symplectic orthogonal matrix satisfying* (5.5). *The computed $\widehat{y}$ satisfies*

$$\widehat{y} = P(x + \Delta x), \quad \|\Delta x\|_2 \leq \widetilde{\gamma}_1 \|x\|_2,$$

where P is the exact elementary symplectic orthogonal matrix.

Proof. The vector $\widehat{y}$ differs from x only in elements $i, j, i+n$, and $j+n$. We have

$$\widehat{y}_i = e_i^T P x + \Delta y_i, \quad |\Delta y_i| \leq \widetilde{\gamma}_1 |e_i^T P||x|.$$

We obtain similar results for $\widehat{y}_j$, $\widehat{y}_{n+i}$, and $\widehat{y}_{n+j}$. Hence,

$$\widehat{y} = Px + \Delta y, \quad |\Delta y| \leq \widetilde{\gamma}_1 |P|\,|x|.$$

As $\|\,|P|\,\|_2 \leq 2$, we have $\|\Delta y\|_2 \leq 2\widetilde{\gamma}_1\|x\|_2 = \widetilde{\gamma}_1'\|x\|_2$. Finally, we define $\Delta x = P^T \Delta y$ and note that $\|\Delta x\|_2 = \|\Delta y\|_2$. $\square$

Now, we consider the pre- and postmultiplication of a matrix H by an approximate elementary symplectic orthogonal matrix $\widehat{P}$.

LEMMA 5.5. *Let $H \in \mathbb{R}^{2n\times 2n}$ and $P \in \mathbb{R}^{2n\times 2n}$ be any elementary symplectic orthogonal matrix such that $fl(P)$ satisfies* (5.5). *Then,*

$$fl(PH) = P(H + \Delta H), \ \|\Delta H\|_F \leq \widetilde{\gamma}_1 \|H\|_F,$$
$$fl(PHP^T) = P(H + \Delta H)P^T, \ \|\Delta H\|_F \leq \widetilde{\gamma}_1 \|H\|_F.$$

Proof. Let h_i be the ith column of H. By Lemma 5.4 we have

$$fl(Ph_i) = P(h_i + \Delta h_i), \quad \|\Delta h_i\|_2 \le \widetilde{\gamma}_1 \|h_i\|_2.$$

The same result holds for h_j, h_{n+i}, and h_{n+j} and the other columns of H are unchanged. Hence, $fl(PH) = P(H + \Delta H)$, where $\|\Delta H\|_F \le \widetilde{\gamma}_1 \|H\|_F$. Similarly, $fl(\widehat{B}P^T) = (\widehat{B} + \Delta\widehat{B})P^T$ with $\|\Delta\widehat{B}\|_F \le \widetilde{\gamma}_1 \|\widehat{B}\|_F$. Then, with $\widehat{B} = fl(PH)$ we have

$$fl(PHP^T) = (PH + P\Delta H + \Delta\widehat{B})P^T = P(H + \Delta H + P^T \Delta\widehat{B})P^T,$$

with $\|\Delta\widehat{B}\|_F \le \widetilde{\gamma}_1(1 + \widetilde{\gamma}_1)\|H\|_F$ and therefore $\|\Delta H + P^T \Delta\widehat{B}\|_F \le \widetilde{\gamma}_1 \|H\|_F$. □

As a consequence of Lemma 5.5, if H_{k+1} is the matrix obtained after one Jacobi update with S_k (which is the product up to six elementary symplectic orthogonal matrices), we have

$$\widehat{H}_{k+1} = S_k(\widehat{H}_k + \Delta\widehat{H}_k)S_k^T, \quad \|\Delta\widehat{H}_k\|_F \le \widetilde{\gamma}_1 \|\widehat{H}_k\|_F, \tag{5.6}$$

where S_k is the exact transformation for $\widehat{H}_k$.

Up to now, we made no assumption on H. If H is a structured Hamiltonian matrix, the $(i, j, n+i, n+j)$-restriction of RHR^T is in canonical form. For instance, if H is a skew-symmetric Hamiltonian matrix, in a computer implementation the diagonal elements of H are not computed but are set to zero. Also, h_{ij}, $h_{i,j+n}$ and by skew-symmetry h_{ji}, $h_{j+n,i}$ are set to zero. But by forcing these elements to be zero, we are making the error smaller so the bounds still hold.

Because of the structure of the problem, both storage and the flop count can be reduced by a factor of four. Any structured Hamiltonian matrix needs less than n^2+n storage locations. If only the t parameters defining H are computed, the structure in the error is preserved and ΔH has the same structure as H. It is easy to see that the bounds in Lemma 5.6 are still valid with the property that ΔH has the same structure as H.

THEOREM 5.6. *Algorithm* 4.1 *for structured Hamiltonians H compute a canonical form $\widehat{T}$ such that*

$$\widehat{T} = P(H + \Delta H)P^T, \quad P^T P = I, \quad P^T J P = J,$$

where ΔH has the same structure as H and $\|\Delta H\|_F \le \widetilde{\gamma}_k \|H\|_F$, where k is the number of symplectic orthogonal transformations S_i applied for each Jacobi update.

The computed basis of symplectic orthogonal eigenvectors $\widehat{P} = fl(S_k \dots S_2 S_1)$ satisfies

$$\|\widehat{P}^T \widehat{P} - I\|_F \le \widetilde{\gamma}_k \quad \text{and} \quad \|\widehat{P}^T J \widehat{P} - J\|_F \le \widetilde{\gamma}_k. \tag{5.7}$$

Proof. From (5.6), one Jacobi update of H satisfies

$$\widehat{H}_1 = fl(S_1 H S_1^T) = S_1(H + \Delta H_1)S_1^T, \quad \|\Delta H_1\|_F \le \widetilde{\gamma}_1 \|H\|_F.$$

For the second update we have

$$\begin{aligned}\widehat{H}_2 = fl(S_2 \widehat{H}_1 S_2^T) &= S_2(\widehat{H}_1 + \Delta\widehat{H}_1)S_2^T, \quad \|\Delta\widehat{H}_1\|_F \le \widetilde{\gamma}_1 \|\widehat{H}_1\|_F \\ &= S_2 S_1(H + \Delta H_1 + S_1^T \Delta\widehat{H}_1 S_1)S_1^T S_2^T \\ &= S_2 S_1(H + \Delta H_2)S_1^T S_2^T,\end{aligned}$$

where $\|\Delta H_2\|_F \le \|\Delta H_1\|_F + \|\Delta \widehat{H}_1\|_F \le \widetilde{\gamma}_1(1+(1+\widetilde{\gamma}_1))\|H\|_F \le \widetilde{\gamma}_2\|H\|_F$. Continuing in this fashion, we find that, after k updates,

$$\widehat{H}_k = S_k \dots S_1(H + \Delta H_k)S_1^T \dots S_k^T \quad \text{with} \quad \|\Delta H_k\|_F \le \widetilde{\gamma}_k\|H\|_F.$$

In a similar way, using the first part of Lemma 5.5 we have

$$\begin{aligned} \widehat{P}_1 &= fl(S_1 I) = S_1 + \Delta P_1, \quad \|\Delta P_1\|_F \le \widetilde{\gamma}_1, \\ \widehat{P}_2 &= fl(S_2\widehat{P}_1) = S_2(S_1 + \Delta P_1) + \Delta\widehat{P}_2, \quad \|\Delta \widehat{P}_2\|_F \le \widetilde{\gamma}_1\|\widehat{P}_1\| \\ &= S_2S_1 + \Delta P_2, \quad \|\Delta P_2\|_F \le \widetilde{\gamma}_1 + \widetilde{\gamma}_1(1+\widetilde{\gamma}_1) \le \widetilde{\gamma}_2. \end{aligned}$$

After k updates, $\widehat{P}_k = fl(S_k\widehat{P}_{k-1}) = S_k \dots S_1 + \Delta P_k, \quad \|\Delta P_k\|_F \le \widetilde{\gamma}_k$, and (5.7) follows readily. ◻

Theorem 5.6 shows that the computed eigenvalues are the exact eigenvalues of a nearby structured Hamiltonian matrix and that the computed basis of eigenvectors is orthogonal and symplectic up to machine precision. This proves the strong backward stability of the Jacobi algorithms.

6. Numerical experiments. To illustrate our results we present some numerical examples. All computations were carried out in MATLAB, which has unit roundoff $u = 2^{-53} \approx 2.2 \times 10^{-16}$.

For symmetric Hamiltonians, symmetric skew-Hamiltonians, and skew-symmetric Hamiltonians with approximate eigenvector $\widehat{x}$ of the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$, computing $\mu(\widehat{x}, \widehat{\lambda})$ in (3.1) involves a symplectic quasi-QR factorization of a $2n \times 2$ matrix, which can be done in order n^2 flops, a cost negligible compared with the $O(n^3)$ cost of the whole eigendecomposition.

For skew-symmetric Hamiltonians with approximate eigenvector $\widehat{x}$ not of the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$, and for skew-symmetric skew-Hamiltonians, the computation of $\mu(\widehat{x}, \widehat{\lambda})$ requires $O(n^3)$ flops as we have to find the minimal 2-norm solution of a large underdetermined system in (3.6). Thus, in this case, $\mu(\widehat{x}, \widehat{\lambda})$ is not a quantity we would compute routinely in the course of solving a problem.

Note that in our implementation of the Jacobi-like algorithm for skew-symmetric Hamiltonians we choose the approximate eigenvectors to be the columns of $P\left[\begin{smallmatrix} I & I \\ -iI & -iI \end{smallmatrix}\right]$, where P is the accumulation of the symplectic orthogonal transformations used by the algorithm to build the canonical form. In this case, the approximate eigenvectors $\widehat{x}$ are guaranteed to be of the form $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$.

To test the strong stability of numerical algorithms for solving structured Hamiltonian eigenproblems, we applied the direct search maximization routine `mdsmax` of the MATLAB Test Matrix Toolbox [15] to the function

$$f(E, F) = \max_{1 \le i \le 2n} \mu(\widehat{x}_i, \widehat{\lambda}_i),$$

where $(\widehat{\lambda}_i, \widehat{x}_i)$ are the computed eigenpairs. In this way we carried out a search for problems on which the algorithms performs unstably.

As expected from the theory, we could not generate examples for which the structured backward error for the Jacobi-like algorithms is large: $\mu(\widehat{x}, \widehat{\lambda}) < nu\|H\|_F$ in all our tests.

The symmetric QR algorithm does not use symplectic orthogonal transformations and is therefore not structure-preserving. To our surprise, we could not generate examples of symmetric Hamiltonian and symmetric skew-Hamiltonian matrices for which

TABLE 6.1
Backward error of the eigenpair for $\lambda = 0.741i$ *of the* 4×4 *skew-symmetric Hamiltonian defined by* (6.1).

	$\eta_{\max}(\widehat{x},\widehat{\lambda})$	$\omega_{\max}(\widehat{x},\widehat{\lambda})$	$\mu_{\max}(\widehat{x},\widehat{\lambda})$
QR algorithm	2×10^{-16}	4×10^{-16}	9×10^{-2}
Jacobi-like algorithm	5×10^{-17}	1×10^{-16}	9×10^{-17}

TABLE 6.2
Backward errors of the approximation of the eigenvalue 0 *for a* 30×30 *random skew-symmetric skew-Hamiltonian matrix.*

	$\lvert\widehat{\lambda}\rvert$	$\eta(\widehat{x},\widehat{\lambda})$	$\omega(\widehat{x},\widehat{\lambda})$	$\mu(\widehat{x},\widehat{\lambda})$
QR algorithm	3×10^{-11}	1×10^{-16}	6×10^{-16}	7×10^{-7}
Jacobi-like algorithm	0	6×10^{-17}	4×10^{-16}	1×10^{-15}

any of the eigenpairs computed by the symmetric QR algorithm has a large backward error. However, the QR algorithm does not compute a symplectic orthogonal basis of eigenvectors and also, it is easy to generate examples for which the $\pm\lambda$ structure for symmetric Hamiltonians and eigenvalue multiplicity 2 structure for symmetric skew-Hamiltonians is not preserved. If we generalize the definition of the structured backward error of a single eigenpair to a set of k eigenpairs, the symmetric QR algorithm is likely to produce sets of eigenpairs with an infinite structured backward error. The QR-like algorithm for symmetric skew-Hamiltonians is likely to provide eigenvectors that are complex instead of real, yielding an infinite structured backward error in (3.14).

The good backward stability of individual eigenpairs computed by the QR algorithm does not hold for the skew-symmetric Hamiltonian case. For instance, we considered the skew-symmetric Hamiltonian eigenproblem

$$(6.1)\quad H = \begin{bmatrix} E & F \\ -F & E \end{bmatrix}, \text{ with } E = \begin{bmatrix} 0 & 0.75 \\ -0.75 & 0 \end{bmatrix}, \; F = \begin{bmatrix} -0.1875 & 0.0938 \\ 0.0938 & -0.125 \end{bmatrix},$$

whose eigenvalues are distinct: $\Lambda(H) = \{0.803i, -0.803i, 0.741i, -0.741i\}$. In Table 6.1, we give the normwise, componentwise, and structured normwise backward error of the eigenpair for $\lambda = 0.741i$ computed by the unsymmetric QR algorithm and the skew-symmetric Jacobi algorithm. The QR algorithm does not use symplectic orthogonal transformations and the computed eigenvectors do not have the structure $\left[\begin{smallmatrix} z \\ \pm iz \end{smallmatrix}\right]$. Therefore, for the computation of $\mu_{\max}(\widehat{x},\widehat{\lambda})$, we use the general formula (3.7).

In the skew-symmetric skew-Hamiltonian case, when n is odd, 0 is an eigenvalue of multiplicity two and is not always well approximated with the unsymmetric QR algorithm. We generated a random 15×15 E and F. We give in Table 6.2 the backward errors associated with the approximation of the eigenvalue 0 for both the QR algorithm and Jacobi algorithm.

7. Conclusion. The first contribution of this work is to extend existing definitions of backward errors in a way appropriate to structured Hamiltonian eigenproblems. We provided computable formulae that are inexpensive to evaluate except for skew-symmetric skew-Hamiltonians. Our numerical experiments showed that for symmetric structured Hamiltonian eigenproblems, the symmetric QR algorithm computes eigenpairs with a small structured backward error but the algebraic properties of the problem are not preserved.

Our second contribution is a detailed rounding error analysis of the new Jacobi algorithms of Faßbender, Mackey, and Mackey [10] for structured Hamiltonian eigenproblems. These algorithms are structure-preserving, inherently parallelizable, and hence attractive for solving large-scale eigenvalue problems. We proved their strong stability when the left and right quaternion rotations are implemented using our formulae (5.2), (5.3). Jacobi algorithms are easy to implement and offer a good alternative to QR algorithms, namely, the unsymmetric QR algorithm, which we showed to be not strongly backward stable for skew-symmetric Hamiltonian and skew-Hamiltonian eigenproblems, and the algorithm for symmetric skew-Hamiltonians based on applying the QR algorithm to (4.1), which does not respect the "realness" of the problem.

Acknowledgments. I thank Nil Mackey for pointing out the open question concerning the strong stability of the Jacobi algorithms for structured Hamiltonian eigenproblems and for her suggestion in fixing the cancellation problem when computing the quaternion rotations. I also thank Steve Mackey for his helpful comments on an earlier manuscript.

REFERENCES

[1] P. Benner, V. Mehrmann, and H. Xu, *A new method for computing the stable invariant subspace of a real Hamiltonian matrix*, J. Comput. Appl. Math., 86 (1997), pp. 17–43.

[2] J. R. Bunch, *The weak and strong stability of algorithms in numerical linear algebra*, Linear Algebra Appl., 88/89 (1987), pp. 49–66.

[3] A. Bunse-Gerstner, *Matrix factorizations for symplectic QR-like methods*, Linear Algebra Appl., 83 (1986), pp. 49–77.

[4] A. Bunse-Gerstner, R. Byers, and V. Mehrmann, *A quaternion QR algorithm*, Numer. Math., 55 (1989), pp. 83–95.

[5] A. Bunse-Gerstner, R. Byers, and V. Mehrmann, *A chart of numerical methods for structured eigenvalue problems*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 419–453.

[6] R. Byers, *A Hamiltonian QR algorithm*, SIAM J. Sci. Statist. Comput., 7 (1986), pp. 212–229.

[7] R. Byers, *A Hamiltonian-Jacobi algorithm*, IEEE Trans. Automat. Control, 35 (1990), pp. 566–570.

[8] A. Deif, *A relative backward perturbation theorem for the eigenvalue problem*, Numer. Math., 56 (1989), pp. 625–626.

[9] J. Dongarra, J. R. Gabriel, D. D. Kölling, and J. H. Wilkinson, *The eigenvalue problem for Hermitian matrices with time reversal symmetry*, Linear Algebra Appl., 60 (1984), pp. 27–42.

[10] H. Faßbender, D. S. Mackey, and N. Mackey, *Hamilton and Jacobi come full circle: Jacobi algorithms for structured Hamiltonian eigenproblems*, Linear Algebra Appl., to appear.

[11] K. V. Fernando, *Accurately counting singular values of bidiagonal matrices and eigenvalues of skew-symmetric tridiagonal matrices*, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 373–399.

[12] A. J. Geurts, *A contribution to the theory of condition*, Numer. Math., 39 (1982), pp. 85–96.

[13] D. J. Higham and N. J. Higham, *Backward error and condition of structured linear systems*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 162–175.

[14] D. J. Higham and N. J. Higham, *Structured backward error and condition of generalized eigenvalue problems*, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 493–512.

[15] N. J. Higham, *The Test Matrix Toolbox for* Matlab (*version* 3.0), Numerical Analysis Report 276, Manchester Centre for Computational Mathematics, Manchester, England, 1995.

[16] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, 1996.

[17] P. Lancaster and L. Rodman, *Algebraic Riccati Equations*, Oxford University Press, New York, 1995.

[18] P. Lancaster and M. Tismenetsky, *The Theory of Matrices*, 2nd ed., Academic Press, London, 1985.

[19] W.-W. Lin, V. Mehrmann, and H. Xu, *Canonical forms for Hamiltonian and symplectic matrices and pencils*, Linear Algebra Appl., 302/303 (1999), pp. 469–533.

[20] N. Mackey, *Hamilton and Jacobi meet again: Quaternions and the eigenvalue problem*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 421–435.

[21] J. Olsen, H. Jørgen, A. Jensen, and P. Jørgensen, *Solution of the large matrix equations which occur in response theory*, J. Comput. Phys., 74 (1988), pp. 265–282.

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

[23] N. Rösch, *Time-reversal symmetry, Kramers' degeneracy and the algebraic eigenvalue problem*, Chemical Physics, 80 (1983), pp. 1–5.

[24] G. W. Stewart, *A Jacobi-like algorithm for computing the Schur decomposition of a non-Hermitian matrix*, SIAM J. Sci. Statist. Comput., 6 (1985), pp. 853–864.

[25] P. M. Van Dooren, *Structured linear algebra problems in digital signal processing*, in Numerical Linear Algebra, Digital Signal Processing and Parallel Algorithms, G. H. Golub and P. M. Van Dooren, eds., vol. F70 of NATO ASI Series, Springer-Verlag, Berlin, 1991, pp. 361–384.

[26] C. F. Van Loan, *A symplectic method for approximating all the eigenvalues of a Hamiltonian matrix*, Linear Algebra Appl., 61 (1984), pp. 233–251.