

Structured Eigenvalue Problems

Heike Fassbender*¹ and Daniel Kressner**²

¹ Institut Computational Mathematics, TU Braunschweig, D-38023 Braunschweig

² Department of Mathematics, University of Zagreb, Croatia; Department of Computing Science, Umeå University, Sweden

Received 15 April 2005, revised 26 May 2005, accepted 1 June 2005

Key words Structured matrix, eigenvalue, invariant subspace, numerical methods, software.
MSC (2000) 65-F15

Most eigenvalue problems arising in practice are known to be structured. Structure is often introduced by discretization and linearization techniques but may also be a consequence of properties induced by the original problem. Preserving this structure can help preserve physically relevant symmetries in the eigenvalues of the matrix and may improve the accuracy and efficiency of an eigenvalue computation. The purpose of this brief survey is to highlight these facts for some common matrix structures. This includes a treatment of rather general concepts such as structured condition numbers and backward errors as well as an overview of algorithms and applications for several matrix classes including symmetric, skew-symmetric, persymmetric, block cyclic, Hamiltonian, symplectic and orthogonal matrices.

Copyright line will be provided by the publisher

1 Introduction

This survey is concerned with computing eigenvalues, eigenvectors and invariant subspaces of a structured square matrix A . In the scope of this paper, an $n \times n$ matrix A is considered to be *structured* if its n^2 entries depend on *less* than n^2 parameters.

It usually takes a long process of simplifications, linearizations and discretizations before one comes up with the problem of computing the eigenvalues or invariant subspaces of a matrix. These techniques typically lead to highly structured matrix representations, which, for example, may contain redundancy or inherit some physical properties from the original problem. As a simple example, let us consider a quadratic eigenvalue problem of the form

$$(\lambda^2 I_n + \lambda C + K)x = 0, \quad (1)$$

where $C \in \mathbb{R}^{n \times n}$ is skew-symmetric ($C = -C^T$), $K \in \mathbb{R}^{n \times n}$ is symmetric ($K = K^T$), and I_n denotes the $n \times n$ identity matrix. Eigenvalue problems of this type arise, e.g., from gyroscopic systems [96, 117] or Maxwell equations [108]; they have the physically relevant property that all eigenvalues appear in quadruples $\{\lambda, -\lambda, \bar{\lambda}, -\bar{\lambda}\}$, i.e., the spectrum is symmetric with respect to the real and imaginary axes. Linearization turns (1) into a matrix

* e-mail: h.fassbender@tu-bs.de

** Corresponding author: e-mail: kressner@cs.umu.se.




Fig. 1 Eigenvalues (‘.’) and approximate eigenvalues (‘*’) computed by a standard Arnoldi method (left picture) and by a structure-preserving Arnoldi method (right picture).

eigenvalue problem, e.g., the eigenvalues of (1) can be obtained from the eigenvalues of the matrix

$$A = \begin{bmatrix} -\frac{1}{2}C & \frac{1}{4}C^2 - K \\ I_n & -\frac{1}{2}C \end{bmatrix}. \quad (2)$$

This $2n \times 2n$ matrix is structured, its $4n^2$ entries depend only on the n^2 entries necessary to define C and K . The matrix A has the particular property that it is a *Hamiltonian* matrix, i.e., A is a two-by-two block matrix of the form

$$\begin{bmatrix} B & G \\ Q & -B^T \end{bmatrix}, \quad G = G^T, \quad Q = Q^T, \quad B, G, Q \in \mathbb{R}^{n \times n}.$$

Considering A to be Hamiltonian does not capture all the structure present in A but it captures an essential part: The spectrum of *any* Hamiltonian matrix is symmetric with respect to the real and imaginary axes. Hence, numerical methods that take this structure into account are capable to preserve the eigenvalue pairings of the original eigenvalue problem (1), despite the presence of roundoff and other approximation errors.

Figure 1, which displays the eigenvalues of a quadratic eigenvalue problem (1) stemming from a discretized Maxwell equation [108], illustrates this fact. The exact eigenvalues are represented by grey dots in each plot; one can clearly see the eigenvalue pairing $\{\lambda, -\lambda, \bar{\lambda}, -\bar{\lambda}\}$, for complex eigenvalues with nonzero real part and $\{\lambda, -\lambda\}$ for real and purely imaginary eigenvalues. Eigenvalue approximations (denoted by black crosses) have been computed with two different types of Arnoldi methods (a standard approach for computing approximations to some eigenvalues of large sparse matrices): in the plot on the left hand side the eigenvalue approximations were obtained from a few iterations of the standard Arnoldi method [51] while the ones in the plot on the right hand side were obtained from the same number of iterations of an Arnoldi method that takes Hamiltonian structures into account, see [96] and Section 3.7. It is particularly remarkable that the latter method produces not only pairs of eigenvalue approximations, but also purely imaginary approximations to purely imaginary eigenvalues.

Besides the preservation of such eigenvalue symmetries, there are several other benefits to be gained from using structure-preserving algorithms in place of general-purpose algorithms

for computing eigenvalues. These benefits include reduced computational time and improved eigenvalue/eigenvector accuracy.

This paper is organized as follows. In Section 2, we will present several general concepts related to (structured) eigenvalue problems. To be more specific, Section 2.1 provides a summary of common algorithms for solving eigenvalue problems while Section 2.2 discusses the efficiency gains to be expected from exploiting structure. Sections 2.3 and 2.4 are concerned with (structured) backward stability and condition numbers; both concepts are closely related to the accuracy of eigenvalues in finite-precision arithmetic. Section 3 treats several classes of matrices individually: symmetric, skew-symmetric, persymmetric, block cyclic, Hamiltonian, orthogonal, and others. This treatment includes a brief overview of existing algorithms and applications for each of these structures.

1.1 Preliminaries

The *eigenvalues* of a real or complex $n \times n$ matrix A are the roots of its characteristic polynomial $\det(A - \lambda I)$. Since the degree of the characteristic polynomial equals n , the dimension of A , it has n roots, so A has n eigenvalues. The eigenvalues may be real or complex, even if A is real (in that case, complex eigenvalues appear in pairs $\{\lambda, \bar{\lambda}\}$). The set of all eigenvalues will be denoted by $\lambda(A)$. A nonzero vector $x \in \mathbb{C}^n$ is called a (*right*) *eigenvector* of A if it satisfies $Ax = \lambda x$ for some $\lambda \in \lambda(A)$. A nonzero vector $y \in \mathbb{C}^n$ is called a *left eigenvector* if it satisfies $y^H A = \lambda y^H$, where y^H denotes the Hermitian transpose of y . Spaces spanned by eigenvectors remain invariant under multiplication by A , in the sense that $\text{span}\{Ax\} = \text{span}\{\lambda x\} \subseteq \text{span}\{x\}$. This concept generalizes to higher-dimensional spaces. A linear subspace $\mathcal{X} \subset \mathbb{C}^n$ with $Ax \in \mathcal{X}$ for all $x \in \mathcal{X}$ is called an *invariant subspace* of A .

If $n > 4$ there is no general closed formula for the roots of a polynomial in terms of its coefficients, and therefore one has to resort to a numerical technique in order to compute the roots by successive approximation. A further difficulty is that the roots may be very sensitive to small changes in the coefficients of the polynomial, and the effect of rounding errors in the construction of the characteristic polynomial is usually catastrophic (see Section 2.1 for an example). Numerically more reliable methods are obtained from the following observation. Eigenvalues and invariant subspaces can be easily obtained from the *Schur decomposition* of A : There is always a unitary matrix U ($U^H U = U U^H = I$) such that

$$T = U^H A U$$

is upper triangular, that is, $t_{ij} = 0$ for $i > j$. As $\lambda(A)$ equals $\lambda(T)$, the eigenvalues of A can be read off from the diagonal entries of T . Moreover, if we partition $U = [X, X_\perp]$, where $X \in \mathbb{C}^{n \times k}$ for some $1 \leq k < n$, then we may rewrite the relation $AU = TU$ as

$$A[X, X_\perp] = [X, X_\perp] \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, \quad T_{11} \in \mathbb{C}^{k \times k}, \quad T_{22} \in \mathbb{C}^{(n-k) \times (n-k)}. \quad (3)$$

This shows that $AX = XT_{11}$, i.e., the columns of X form an orthonormal basis for the invariant subspace \mathcal{X} belonging to $\lambda(T_{11})$. Bases for invariant subspaces belonging to other eigenvalues can be obtained by changing the order of the eigenvalues on the diagonal of T [51].

In case the matrix A is real, one would like to restrict the equivalence transformation which reveals the eigenvalues and corresponding invariant subspaces to a real transformation. As A

may have complex eigenvalues, a reduction to the above introduced Schur form is no longer possible. We must lower our expectations and be content with the calculation of an alternative decomposition known as the *real Schur decomposition*: There always exists an orthogonal matrix Q ($Q^T Q = Q Q^T = I$) such that $Q^T A Q = T_R$ is upper quasi-triangular, that is, T_R is block upper triangular with either 1×1 or 2×2 diagonal blocks. Each 2×2 diagonal block corresponds to a pair of complex conjugate eigenvalues.

2 General Concepts

In this section, we briefly summarize some general concepts related to general and structured eigenvalue computations.

2.1 Algorithms

Algorithms for solving structured eigenvalue problems are often extensions or special cases of existing algorithms for solving the general, unstructured eigenvalue problem. One can roughly divide these algorithms into direct methods and iterative methods.

Direct methods aim at computing *all* eigenvalues and (optionally) eigenvectors/invariant subspaces of a matrix A . The most widely used direct method for general nonsymmetric matrices is the *QR algorithm* [51]. This algorithm, which is behind the MATLAB [93] command `eig` , computes a sequence of matrix decompositions converging to a complete Schur form of A by implicitly applying a QR decomposition in each iteration. *Jacobi-like methods* [45, 112] apply a sweep of elementary transformations, such as Givens rotations, in each iteration. They are well-suited for parallel computation but their slow speed of convergence makes them often inferior to the QR algorithm unless there is some special structure in A , e.g., A is symmetric or close to being upper triangular. Moreover, Jacobi's method can sometimes compute tiny eigenvalues much more accurately than other methods [41].

Iterative methods aim at computing a *selected set* of eigenvalues and (optionally) the eigenvectors/invariant subspaces belonging to these eigenvalues. The most widely known iterative method is probably the *power method* [51]. Starting from some random vector v , one constructs a sequence v, Av, A^2v, A^3v, \dots by repeated matrix-vector multiplication. This sequence converges to an eigenvector belonging to the dominant eigenvalue (i.e., the eigenvalue of largest magnitude), provided that there is only one such eigenvalue. Besides being incapable of obtaining several eigenvalues or eigenvectors, the power method may also suffer from very slow convergence in the presence of several nearly dominant eigenvalues. To find other eigenvalues and eigenvectors, the power method can be applied to $(A - \sigma I)^{-1}$ for some shift σ , an algorithm called *inverse iteration*. The largest eigenvalue of $(A - \sigma I)^{-1}$ is $1/(\lambda_i - \sigma)$ where λ_i is the closest eigenvalue of A to σ , so that one can choose which eigenvalues to find by choosing σ .

The *Arnoldi method* [5, 51] achieves faster convergence by considering not only the last iterate of the power method but the whole space spanned by all previous iterates, the so called *Krylov subspace* $\mathcal{K}_k(A, v) := \text{span}\{v, Av, \dots, A^{k-1}v\}$. Using Krylov subspaces also gives the flexibility to approximate several eigenvalues and the associated invariant subspaces. The increased memory requirements of the Arnoldi method can be limited by employing restarting strategies, see, e.g., [110]. This leads to the implicitly restarted Arnoldi algorithm [85] (IRA), which forms the basis of MATLAB's `eigs` .




Fig. 2 Computed roots ('x') of the polynomial $(\lambda - 1)(\lambda - 2) \cdots (\lambda - 25)$, after its coefficients have been stored in double-precision.

A different class of iterative methods is obtained by applying *Newton's method* to a system of nonlinear equations satisfied by an eigenvalue/eigenvector pair (λ, x) . One such system is given by

$$f(\lambda, x) := \begin{bmatrix} Ax - \lambda x \\ w^H x - 1 \end{bmatrix} = 0,$$

where w is a suitably chosen vector. It turns out that Newton's method applied to this system is equivalent to inverse iteration, see also [95] for more details. The Jacobi-Davidson method (see the review in the present volume of this journal or [5, 109]) is a closely related Newton-like method, but so far little is known about adapting this method to structured eigenvalue problems. There are other functions f that may be used, but not all of them lead to practically useful methods for computing eigenvalues. For example, using $p(\lambda) = \det(\lambda I - A)$ by explicitly constructing the coefficients of the polynomial $p(\lambda)$ is generally *not* advisable. Often, already storing these coefficients in finite-precision arithmetic leads to an unacceptably high error in the eigenvalues; leave alone any other roundoff error implied by the typically wildly varying magnitudes. Figure 2, which goes back to an example by Wilkinson [134], illustrates this effect. However, it should be mentioned that iterative methods based on $p(\lambda)$ are suitable for some classes of matrices such as tridiagonal matrices, but great care in representing $p(\lambda)$ must be taken, see [13] and the references therein.

2.2 Efficiency

Direct methods generally require $\mathcal{O}(n^3)$ computational time to compute the eigenvalues of a general $n \times n$ matrix. The fact that a structured matrix depends on less than n^2 parameters gives rise to the hope that an algorithm taking advantage of the structure may require less effort than a general-purpose algorithm. For example, the QR algorithm applied to a symmetric matrix requires roughly 10% of the floating point operations (flops) required by the same algorithm applied to a general matrix [51]. For other structures, such as skew-Hamiltonian and Hamiltonian matrices, this figure can be less dramatic [9]. Moreover, in view of recent progress made in improving the performance of general-purpose algorithms [21, 22], it may

require considerable implementation efforts to turn this reduction of flops into an actual reduction of computational time.

The convergence of an iterative method strongly depends on the properties of the matrix A and the subset of eigenvalues to be computed, which makes the needed computational effort rather difficult to predict. For example, methods based on Krylov subspaces require in each iteration a matrix-vector multiplication, which can take up to $2n^2$ flops. This figure may reduce significantly for a structured matrix A , e.g., to roughly twice the number of nonzero entries for sparse matrices. Some structures, such as skew-Hamiltonian and persymmetric matrices [96, 125], induce some additional structure in the Krylov subspace making it possible to reduce the computational effort even further.

2.3 Backward error and backward stability

Any numerical method for computing eigenvalues is affected by round-off errors due to the limitations of finite-precision arithmetic. Already representing the entries of a matrix A in double-precision generally causes a relative error of about 10^{-16} . Unless some further information is available, the best we can therefore expect from our numerical method is that it computes the *exact* eigenvalues of a slightly *perturbed* matrix $A + E$, where $\|E\|_F$ (the Frobenius norm of E) is not much greater than $10^{-16} \times \|A\|_F$. A numerical method satisfying this property is called (numerically) *backward stable*. Methods known to be backward stable include the QR algorithm, most Jacobi-like methods, the power method, the Arnoldi method, IRA, and many of the better Newton-like methods.

A simple way to check backward stability is to compute the residual $r = A\hat{x} - \hat{\lambda}\hat{x}$ of a computed eigenvalue/eigenvector pair $(\hat{\lambda}, \hat{x})$ with $\|\hat{x}\|_2 = 1$. Then $(\hat{\lambda}, \hat{x})$ is an exact eigenpair of the perturbed matrix $A - r\hat{x}^H$ and the backward error is consequently given by $\|r\hat{x}^H\|_F = \|r\|_2$. If only $\hat{\lambda}$ is available then a suitable backward error is given by $\sigma_{\min}(A - \hat{\lambda}I)$, which coincides with the 2-norm of the smallest perturbation E such that $\hat{\lambda}$ is an eigenvalue of $A + E$. (Here, σ_{\min} denotes the smallest singular value of a matrix.)

The notion of *structured backward error* not only requires $(\hat{\lambda}, \hat{x})$ to be an exact eigenpair of $A + E$ but also requires $A + E$ to stay within the considered class of structured matrices. For example, if A is symmetric then a suitable symmetric E is given by $E = -(r\hat{x}^T + \hat{x}r^T - (r^T\hat{x})\hat{x}\hat{x}^T)$. Since $\|r\|_2 \leq \|E\|_2 \leq \|r\|_2/\sqrt{2}$, this implies that the structured backward error with respect to symmetric matrices is rather close to the standard backward error. This is not for all structures the case, see [116] for more details. Depending on the point of view, the problem of finding the structured backward error can also be regarded as a structured inverse eigenvalue problem [36] or computing the smallest structured singular value [66, 73].

A numerical method is called *strongly backward stable* if the computed eigenpairs satisfy small structured backward errors [24]. Strong backward stability implies backward stability, but the converse is generally not true.

2.4 Condition numbers and pseudospectra

Knowing that a computed eigenvalue satisfies a small (structured) backward error does not necessarily admit conclusions about the accuracy of this eigenvalue. This also depends on the




Fig. 3 Unstructured pseudospectra of the matrix A in (6) for different perturbation levels ε .

eigenvalue condition number $\kappa(\lambda)$, which is formally defined as follows:

$$\kappa(\lambda) := \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sup\{|\hat{\lambda} - \lambda| : E \in \mathbb{C}^{n \times n}, \|E\|_2 \leq \varepsilon\}, \quad (4)$$

where $\hat{\lambda}$ is the eigenvalue of the perturbed matrix $A + E$ closest to λ . In other words, $\kappa(\lambda)$ measures the worst-case first order influence of a perturbation E on the accuracy of λ . If $\kappa(\lambda)$ is large, then λ is said to be ill-conditioned. Eigenvalues with small condition number are said to be well-conditioned.

The definition (4) readily implies $|\hat{\lambda} - \lambda| \lesssim \kappa(\lambda)\|E\|_2$. Thus, a backward stable method attains high accuracy for fairly well-conditioned eigenvalues. If λ is a simple eigenvalue, i.e., λ is a simple root of $\det(\lambda I - A)$, then $\kappa(\lambda) = 1/|y^H x|$ with x and y being right and left eigenvectors belonging to λ normalized such that $\|x\|_2 = \|y\|_2 = 1$ [51]. For any normal matrix N (that is $N^H N = N N^H$) there exists a unitary matrix Q such that $Q^H N Q = \text{diag}(\lambda_1, \dots, \lambda_n)$ is diagonal. Hence, $N q_j = \lambda_j q_j$ and $q_j^H N = q_j^H \lambda_j$. Its every right eigenvector is also a left eigenvector belonging to the same eigenvalue and $\kappa(\lambda) = 1$ for all its eigenvalues. Hence, eigenvalues of normal matrices are well-conditioned.

Pseudospectra provide more detailed information on the behavior of eigenvalues under perturbations [118]. For a given perturbation level $\varepsilon > 0$, the pseudospectrum $\Lambda_\varepsilon(A)$ of A is the union of the eigenvalues of all perturbed matrices $A + E$ with $\|E\|_2 \leq \varepsilon$:

$$\Lambda_\varepsilon(A) := \bigcup \{\Lambda(A + E) : E \in \mathbb{C}^{n \times n}, \|E\|_2 \leq \varepsilon\}. \quad (5)$$

Figure 3 displays pseudospectra for the matrix

$$A = \begin{bmatrix} 2 & 1 & 2 \\ -1 & 2 & 2 \\ 0 & 0 & 0 \end{bmatrix}. \quad (6)$$

Provided that all eigenvalues are simple, the pseudospectrum approaches, as ε tends to zero, discs around the eigenvalues and the radii of these discs divided by ε coincide with the corresponding eigenvalue condition numbers, see, e.g., [73].

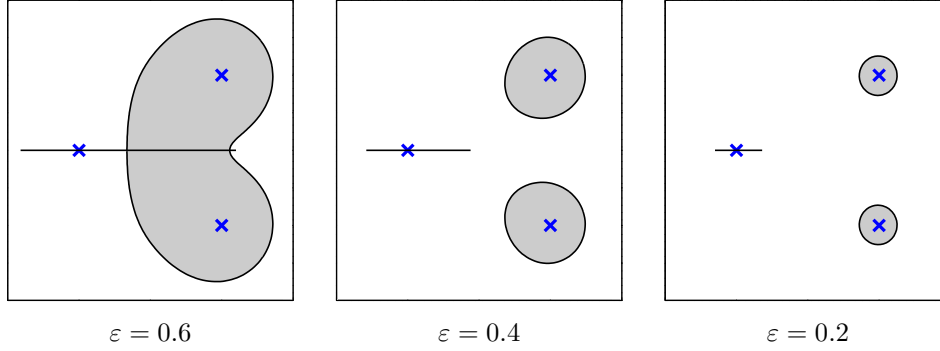


Fig. 4 Real pseudospectra of the matrix A in (6) for different perturbation levels ε .

This is no longer true if we only admit structured perturbations in (5), i.e., matrices E that belong to a certain matrix class struct . The corresponding *structured pseudospectrum* $\Lambda_\varepsilon^{\text{struct}}$ may approach completely different geometrical shapes [66, 73]. For example, if we allow for real perturbations only, then $\text{struct} \equiv \mathbb{R}^{n \times n}$. $\Lambda_\varepsilon^{\mathbb{R}^{n \times n}}$ is known to approach ellipses. This is demonstrated in Figure 4 (a line can be regarded as a degenerated ellipse). The larger half diameters of these ellipses divided by ε yield the *structured eigenvalue condition numbers*, formally defined as

$$\kappa^{\text{struct}}(\lambda) := \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sup\{|\hat{\lambda} - \lambda| : E \in \text{struct}, \|E\|_2 \leq \varepsilon\}. \quad (7)$$

It turns out that $\kappa^{\mathbb{R}^{n \times n}}(\lambda) \geq \kappa(\lambda)/\sqrt{2}$, see [30]. Since $\kappa^{\text{struct}}(\lambda) \leq \kappa(\lambda)$ holds for any struct , this implies that there is no significant difference between $\kappa^{\mathbb{R}^{n \times n}}(\lambda)$ and $\kappa(\lambda)$. Surprisingly, the same statement holds for a variety of other structures [18, 55, 74, 107], see also Section 3. Note that (7) implies $|\hat{\lambda} - \lambda| \lesssim \kappa^{\text{struct}}(\lambda)\|E\|_2$ for all $E \in \text{struct}$. Explicit expressions for $\kappa^{\text{struct}}(\lambda)$ have been derived in [64], provided that struct is a linear matrix space.

Defining an appropriate condition number $\kappa(\mathcal{X})$ for an invariant subspace \mathcal{X} of A is more complicated due to the fact that we have to work with subspaces and rely on a meaningful notion of distances between subspaces. Such a notion is provided by the canonical angles [111, 113] and the resulting $\kappa(\mathcal{X})$ can be interpreted as follows. Let the columns of X form an orthonormal basis for \mathcal{X} . Then there is a matrix \hat{X} such that the columns of \hat{X} form a basis for an invariant subspace $\hat{\mathcal{X}}$ of the slightly perturbed matrix $A + E$ and

$$\|\hat{X} - X\|_F \lesssim \kappa(\mathcal{X})\|E\|_F$$

for all $E \in \mathbb{C}^{n \times n}$. A (possibly smaller) $\kappa^{\text{struct}}(\mathcal{X})$ is similarly defined and yields the same bound for all E restricted to struct . Structured condition numbers for eigenvectors and invariant subspaces have been studied in [31, 64, 68, 77, 79, 114]. The (structured) perturbation analysis of quadratic matrix equations is a closely related area, which is comprehensively treated in [76, 115].

3 More on Some Structures

In this section, we aim to provide a brief discussion on selected aspects of some common matrix structures.

3.1 Sparsity and related structures

Sparsity is probably the most ubiquitous structure in (numerical) linear algebra. While a sparse matrix cannot be expected to have any particular eigenvalue or eigenvector properties, it always admits the efficient calculation of its eigenvalues and eigenvectors [5], see also Section 2.2.

In [101], it was shown that if struct denotes all matrices having an assigned sparsity pattern then the corresponding structured eigenvalue condition number equals $\|(yx^H)|_{\text{struct}}\|_F/|y^Hx|$, where $(yx^H)|_{\text{struct}}$ means the restriction of the rank-one matrix yx^H to the sparsity structure of struct . For example if the perturbation is upper triangular then $(yx^H)|_{\text{struct}}$ is the upper triangular part of yx^H .

No methods are known to be strongly backward stable for an arbitrary sparsity structure struct . Hence, it is generally difficult to achieve accuracy gains promised by a small value of $\|(yx^H)|_{\text{struct}}\|_F$.

3.2 Symmetric matrices

Probably the two most fundamental properties of a symmetric matrix $A \in \mathbb{R}^{n \times n}$, $A = A^T$, are that every eigenvalue is real and every right eigenvector is also a left eigenvector belonging to the same eigenvalue. Both facts immediately follow from the observation that a Schur decomposition of A always takes the form $Q^T A Q = \text{diag}(\lambda_1, \dots, \lambda_n)$.

It is simple to show that the structured eigenvalue and invariant subspace condition numbers are equal to the corresponding unstructured condition numbers, i.e., $\kappa^{\text{symm}}(\lambda) = \kappa(\lambda) = 1$ and

$$\kappa^{\text{symm}}(\mathcal{X}) = \kappa(\mathcal{X}) = \frac{1}{\min\{|\mu - \lambda| : \lambda \in \Lambda_1, \mu \in \Lambda_2\}},$$

where \mathcal{X} is a simple invariant subspace belonging to an eigenvalue subset $\Lambda_1 \subset \lambda(A)$, and $\Lambda_2 = \lambda(A) \setminus \Lambda_1$. Moreover, many numerical methods, such as QR, Arnoldi and Jacobi-Davidson, automatically preserve symmetric matrices and are strongly backward stable.

These facts should not lead to the wrong conclusion that the preservation of symmetric matrices is not important. Algorithms tailored to symmetric matrices (e.g., divide and conquer or Lanczos methods) take much less computational effort and sometimes achieve high relative accuracy in the eigenvalues [41] and – having the right representation of A at hand – even in the eigenvectors [43, 44]. To illustrate these accuracy gains, let us consider a 20×20 matrix $A = DHD$ with

$$H = \begin{bmatrix} 1 & 0.1 & \cdots & 0.1 \\ 0.1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0.1 \\ 0.1 & \cdots & 0.1 & 1 \end{bmatrix}, \quad D = \text{diag}(10^{20}, 10^{19}, \dots, 10^1). \quad (8)$$




Fig. 5 Relative errors $|\hat{\lambda} - \lambda|/\lambda$ for the eigenvalues of the matrix $A = DHD$ as defined in (8) computed by the QR algorithm and by the Jacobi algorithm.

From Figure 5 it can be concluded that the Jacobi algorithm computes all eigenvalues of A with high relative accuracy while the QR algorithm fails to achieve this goal.

Besides the well-known algorithms for symmetric eigenproblems such as the QR algorithm, the Lanczos algorithm (a Krylov subspace method tailored for symmetric matrices) and the Jacobi method, which are discussed in most basic books on numerical methods (see, e.g., [51, 131, 40]), there are a number of more recent developments. For example, a divide-and-conquer algorithm was first introduced in 1981 [39]. Its quite subtle implementation was only discovered ten years later [58]. The algorithm is one of the fastest now available for computing all eigenvalues and all eigenvectors. Bisection can be used to find just a subset of the eigenvalues of a symmetric tridiagonal matrix, e.g., those in the interval $[a, b]$. If only $k \ll n$ eigenvalues are required, bisection can be much faster than QR. Inverse iteration can then be used to independently compute the corresponding eigenvectors. The numerical difficulties associated with that approach have recently been resolved [42, 43, 44]. Of course, there is a symmetric version of the Jacobi-Davidson algorithm (see the review in the present volume of this journal or [5, 109]).

An overview of all developments related to symmetric eigenvalue problems is far beyond the scope of this survey; we refer to [37, 40, 103] for introductions to this flourishing branch of eigenvalue computation.

3.3 Skew-symmetric matrices

Any eigenvalue of a skew-symmetric matrix $A \in \mathbb{R}^{n \times n}$, $A = -A^T$, is purely imaginary. If n is odd then there is at least one zero eigenvalue. As for symmetric matrices, any right eigenvector is also a left eigenvector belonging to the same eigenvalue (this is true for any normal matrix). The real Schur form of A takes the form

$$Q^T A Q = \begin{bmatrix} 0 & \alpha_1 \\ -\alpha_1 & 0 \end{bmatrix} \oplus \cdots \oplus \begin{bmatrix} 0 & \alpha_k \\ -\alpha_k & 0 \end{bmatrix} \oplus 0 \oplus \cdots \oplus 0.$$

for some real scalars $\alpha_1, \dots, \alpha_k \neq 0$.

While the structured condition number for a nonzero eigenvalue always satisfies $\kappa^{\text{skew}}(\lambda) = \kappa(\lambda) = 1$, we have for a simple zero eigenvalue $\kappa^{\text{skew}}(0) = 0$ but $\kappa(0) = 1$ [107]. Again, there is nothing to be gained for an invariant subspace \mathcal{X} ; it is simple to show $\kappa^{\text{skew}}(\mathcal{X}) = \kappa(\mathcal{X})$.

Skew-symmetric matrices have received much less attention than symmetric matrices, probably due to the fact that skew-symmetry plays a less important role in applications. Again, the QR and the Arnoldi algorithm automatically preserve skew-symmetric matrices. Variants of the QR algorithm tailored to skew-symmetric matrices have been discussed in [72, 128]. Jacobi-like algorithms for skew-symmetric matrices and other Lie algebras have been developed and analyzed in [59, 60, 61, 75, 102, 106, 133].

3.4 Persymmetric matrices

A real $2n \times 2n$ matrix A is called *persymmetric* if it is symmetric with respect to the anti-diagonal. E.g., A takes the following form for $n = 2$:

$$A = \left[\begin{array}{cc|cc} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{23} \\ \hline a_{31} & a_{12} & a_{22} & a_{12} \\ a_{41} & a_{31} & a_{21} & a_{11} \end{array} \right].$$

If we *additionally* assume A to be *symmetric*, then we can write

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{12}^T & F_n A_{11}^T F_n \end{bmatrix},$$

where F_n denotes the $n \times n$ flip matrix, i.e., F_n has ones on the anti-diagonal and zeros everywhere else. Note that A is also a *centrosymmetric* matrix since $A = F_{2n} A F_{2n}$. A practically relevant example of such a matrix is the Gramian of a set of frequency exponentials $\{e^{\pm i\lambda_k t}\}$, which plays a role in the control of mechanical and electric vibrations [105]. Employing the orthogonal matrix $U = \frac{1}{\sqrt{2}} \begin{bmatrix} I & F_n \\ -F_n & I \end{bmatrix}$ we have

$$U^T A U = \begin{bmatrix} A_{11} - A_{12} F_n & 0 \\ 0 & F_n A_{11} F_n + F_n A_{12} \end{bmatrix}, \tag{9}$$

where we used the symmetry of A_{11} and the persymmetry of A_{12} . This is a popular trick when dealing with centrosymmetric matrices [132]. (A similar but technically slightly more complicated trick can be used if A has odd dimension.) Thus,

$$\lambda(A) = \lambda(A_{11} - A_{12} F_n) \cup \lambda(A_{11} + A_{12} F_n).$$

Perhaps more importantly, if these two eigenvalue sets are disjoint then any eigenvector belonging to $\lambda(A_{11} - A_{12} F_n)$ takes the form

$$x = U \begin{bmatrix} \tilde{x} \\ 0 \end{bmatrix} = \begin{bmatrix} \tilde{x} \\ -F_n \tilde{x} \end{bmatrix}$$

for some $\tilde{x} \in \mathbb{R}^n$. This property of x is sometimes called *center-skew-symmetry*. Analogously, any eigenvector belonging to $\lambda(A_{11} + A_{12} F_n)$ is *center-symmetric*. While the structured and unstructured eigenvalue condition numbers for symmetric persymmetric matrices

are the same [74, 107], there can be a significant difference in the invariant subspace condition numbers. With respect to structured perturbations, the separation between $\lambda(A_{11} - A_{12}F_n)$ and $\lambda(A_{11} + A_{12}F_n)$, which can be arbitrarily small, does not play any role for invariant subspaces belonging to eigenvalues from one of the two eigenvalue sets [31]. It can thus be important to retain the symmetry structure of the eigenvectors in finite-precision arithmetic. Krylov subspace methods achieve this goal automatically if the starting vector is chosen to be center-symmetric (or center-skew-symmetric). This property has been used in [125] to construct a structure-preserving Lanczos process. Efficient algorithms for performing matrix-vector products with centrosymmetric matrices have been investigated in [48, 97, 100].

We now briefly consider a *persymmetric* and *skew-symmetric* matrix A ,

$$A = \begin{bmatrix} A_{11} & A_{12} \\ -A_{12}^T & F_n A_{11}^T F_n \end{bmatrix}.$$

Again, the orthogonal matrix U can be used to reduce A :

$$U^T A U = \begin{bmatrix} 0 & A_{11} F_n + A_{12} \\ F_n A_{11} - A_{12}^T & 0 \end{bmatrix}.$$

Hence, the eigenvalues of A are the positive and negative square roots of the eigenvalues of the matrix product $(A_{11} F_n + A_{12})(F_n A_{11} - A_{12}^T)$, see also [105].

Structure-preserving Jacobi algorithms for symmetric persymmetric and skew-symmetric persymmetric matrices have been recently developed in [90].

3.5 Product eigenvalue problems and block cyclic matrices

The task of computing the eigenvalues of a matrix product arises from various applications in systems and control theory, see, e.g., [7, 23, 89, 123]. Moreover, it has been recently shown that many structured eigenvalue problems can be viewed as particular instances of the product eigenvalue problem [130].

For simplicity, let us consider computing the eigenvalues of an $n \times n$ product of three matrices only: $\Pi = ABC$. At first sight, such an eigenvalue problem does not match the definition of a structured eigenvalue problem stated in the beginning of this survey; Π is an $n \times n$ matrix but depends on the $3n^2$ entries defining its factors A, B, C . However, any product eigenvalue problem can be equivalently seen as an embedded block cyclic structured eigenvalue problem:

$$\mathcal{A} = \begin{bmatrix} 0 & 0 & C \\ B & 0 & 0 \\ 0 & A & 0 \end{bmatrix}.$$

The fact that \mathcal{A}^3 is a block diagonal matrix with diagonal entries CAB, BCA, ABC implies that λ is an eigenvalue of \mathcal{A} if and only if λ^3 is an eigenvalue of ABC (which has the same eigenvalues as CAB and BCA). Also, any invariant subspace of ABC can be related to an invariant subspace of \mathcal{A} [87]. This connection admits the application of structured perturbation results to obtain factor-wise perturbation results for the product eigenvalue problem [11, 35, 81, 86].

The essential key to develop numerically stable algorithms for solving a product eigenvalue problem is to avoid the explicit computation of the matrix product. The periodic QR algorithm [17, 63, 119] is a strongly backward stable method for computing eigenvalues and invariant subspaces of \mathcal{A} , which in turn implies factor-wise backward stability for computing eigenvalues and invariant subspaces of Π . In [78], it has been shown numerically equivalent to the standard QR algorithm applied to a permuted version of \mathcal{A} . Methods for solving product eigenvalue problems with large and (possibly) sparse factors can be found in [80, 88]. Using these methods instead of applying a standard method to Π often results in higher accuracy for the computed eigenvalues, especially for those of small magnitude. This well-appreciated advantage has driven the development of many reliable algorithms for solving instances of the product eigenvalue problem, such as the Golub-Reinsch algorithm for the singular value decomposition [50] or the QZ algorithm for the generalized eigenvalue problem [99].

3.6 Orthogonal matrices

All the eigenvalues of a real orthogonal matrix Q lie on the unit circle. Moreover, as orthogonality implies normality, any right eigenvector is also a left eigenvector belonging to the same eigenvalue. Orthogonal eigenvalue problems have a number of applications in digital signal processing, see [1, 2] for an overview.

The set of orthogonal matrices $\text{orth} = \{A : A^T A = I\}$ forms a smooth manifold, and the tangent space of orth at A is given by $\{AH : H \in \text{skew}\}$. By the results in [74] this implies

$$\begin{aligned} \kappa^{\text{orth}}(\lambda) &= \sup\{|x^H AHx| : H \in \text{skew}, \|AH\|_2 = 1\} \\ &= \sup\{|x^H Hx| : H \in \text{skew}, \|H\|_2 = 1\}. \end{aligned}$$

Hence, if $\lambda = \pm 1$ then x can be chosen to be real which implies $x^T Hx = 0$ for all x and consequently $\kappa^{\text{orth}}(\lambda) = 0$, provided that λ is simple. Similarly, if λ is complex, we have $\kappa^{\text{orth}}(\lambda) = \kappa(\lambda) = 1$ and hence $c_2^{\text{orth}}(\lambda) = 1$. A more general perturbation analysis of orthogonal and unitary eigenvalue problems, based on the Cayley transform, can be found in [16].

Once again, the QR algorithm automatically preserves orthogonal matrices. To make it work (efficiently), it is important to take the fact that the underlying matrix is orthogonal into account. A careful choice of shifts can lead to cubic convergence or even ensure global convergence [46, 126, 127]. Even better, an orthogonal (or unitary) Hessenberg matrix can be represented by $\mathcal{O}(n)$ so called *Schur parameters* [52, 27]. This makes it possible to implement the QR algorithm very efficiently [53]; for an extension to unitary and orthogonal matrix pairs, see [25]. In [54], a divide and conquer approach for unitary Hessenberg matrices was proposed, based on the methodology of Cuppen's divide and conquer approach for symmetric tridiagonal matrices [38]. Krylov subspace methods for orthogonal and unitary matrices have been developed and analyzed in [15, 26, 62, 70, 71].

3.7 Hamiltonian matrices

As already mentioned in the introduction, one of the most remarkable properties of a Hamiltonian matrix

$$H = \begin{bmatrix} A & G \\ Q & -A^T \end{bmatrix}, \quad G = G^T, \quad Q = Q^T,$$

is that its eigenvalues always occur in pairs $\{\lambda, -\lambda\}$, if $\lambda \in \mathbb{R} \cup i\mathbb{R}$, or in quadruples $\{\lambda, -\lambda, \bar{\lambda}, -\bar{\lambda}\}$, if $\lambda \in \mathbb{C} \setminus (\mathbb{R} \cup i\mathbb{R})$. Hamiltonian matrices arise in applications related to linear control theory for continuous-time systems [6] and quadratic eigenvalue problems [96, 117], to name only a few. Deciding whether a certain Hamiltonian matrix has purely imaginary eigenvalues is the most critical step in algorithms for computing the stability radius of a matrix or the H_∞ norm of a linear time-invariant system, see, e.g., [20, 29].

The described eigenvalue pairings often reflect important properties of the underlying application and should thus be preserved in finite-precision arithmetic. QR-like algorithms that achieve this goal have been developed in [10, 28, 120] while Krylov subspace methods tailored to Hamiltonian matrices can be found in [8, 9, 49, 96, 129]. An efficient strongly backward stable method for computing invariant subspaces of H has recently been proposed in [34].

Concerning structured perturbation results for Hamiltonian matrices, see [31, 74, 77, 116] and the references therein. It turns out that there is no or little difference between $\kappa(\lambda)$ and $\kappa^{\text{Hamil}}(\lambda)$, the (structured) eigenvalue condition numbers. Also, $\kappa(\mathcal{X})$ and $\kappa^{\text{Hamil}}(\mathcal{X})$ are equal for the important case that \mathcal{X} is the invariant subspace belonging to all eigenvalues in the open left half plane.

There is much more to be said about Hamiltonian eigenvalue problems, see [9] for a recent survey.

3.8 Symplectic matrices

A matrix M is called *symplectic* iff

$$MJM^T = J,$$

or equivalently, $M^TJM = J$, where

$$J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}.$$

The symplectic matrices form a group under multiplication. The eigenvalues of symplectic matrices occur in reciprocal pairs: If λ is an eigenvalue of M with right eigenvector x , then λ^{-1} is an eigenvalue of M with left eigenvector $(Jx)^T$. The computation of eigenvalues and eigenvectors of such matrices is an important task in applications like the discrete linear-quadratic regulator problem, discrete Kalman filtering, the solution of discrete-time algebraic Riccati equations and certain large, sparse quadratic eigenvalue problems. See, e.g., [82, 84, 94, 96] for applications and further references. Symplectic matrices also occur when solving linear Hamiltonian difference systems [14].

Note that a Cayley transform turns a symplectic matrix into a Hamiltonian one and vice versa. This explains the close resemblance of the spectra of Hamiltonian and symplectic matrices. Moreover, every Hamiltonian matrix H satisfies $HJ = (HJ)^T$. Unfortunately, the Hamiltonian and the symplectic eigenproblems are (despite common believe) quite different. The symplectic eigenproblem is much more difficult than the Hamiltonian one. The relation between the two eigenproblems is best described by comparing it with the relation between symmetric and orthogonal eigenproblems or the Hermitian and unitary eigenproblems. In all these cases, the underlying algebraic structures are an algebra and a group acting on this algebra. For the algebra (Hamiltonian, symmetric, Hermitian matrices), the structure is explicit,

i.e., can be read off the matrix by viewing it. In contrast, the structure of a matrix contained in a group (symplectic, orthogonal, unitary matrices) is given only implicitly. It is very difficult to make this structure explicit. If the “group” eigenproblem is to be solved using a method that exploits the given structure, than this is relatively easy for orthogonal or unitary matrices as one works with the standard scalar product. Additional difficulties for the symplectic problem arise from the fact that one has to work with an indefinite inner product.

The described eigenvalue pairings often reflect important properties of the underlying application and should thus be preserved in finite-precision arithmetic. QR-like algorithms that achieve this goal have been developed as well as Krylov subspace methods tailored to symplectic matrices. An efficient strongly backward stable method for computing invariant subspaces of M is not known so far. More on the algorithms and theoretical results for symplectic matrices are comprehensively summarized in [47].

Concerning structured perturbation results for symplectic matrices, see [74, 116] and the references therein. If λ is a simple eigenvalue for M , then so is $1/\lambda$. It turns out that there is no difference between $\kappa(\lambda)$ and $\kappa(1/\lambda)$, the unstructured eigenvalue condition numbers, but the structured ones differ.

3.9 Other structures

Here, we shortly comment on other structured eigenvalue problems frequently encountered in the literature. By no means should this list or the provided references be regarded as complete.

Hankel and Toeplitz matrices: Toeplitz matrices have constant entries on each diagonal parallel to the main diagonal; they belong to the larger class of persymmetric matrices. Toeplitz structure occur naturally in a variety of applications; tridiagonal Toeplitz matrices are commonly the result of discretizing differential equation problems. $F_n T$ is a Hankel matrix for each Toeplitz matrix T and F_n as the Section 3.4. Hankel matrices arise naturally in problems involving power moments. The Hankel and Toeplitz structure is rich in special properties. Besides admitting the fast computation of matrix-vector products [121], Hankel and Toeplitz matrices have a number of interesting eigenvalue properties, see [19] and the references therein.

Multivariate eigenvalue problems: These are multiparameter eigenvalue problems of the form $W_i(\lambda)x_i = 0$, for $i = 1, \dots, k$, $x_i \in \mathbb{C}^{n_i} \setminus \{0\}$, $\lambda = (\lambda_1, \dots, \lambda_k) \in \mathbb{C}^k$ and $W_i(\lambda) = V_{i0} - \sum_{j=0}^k \lambda_j V_{ij}$ where $V_{ij} \in \mathbb{C}^{n_i \times n_i}$. A k -tuple λ that satisfies the equation $W_i(\lambda)x_i = 0$, for $i = 1, \dots, k$, is called an eigenvalue and the tensor product $x = x_1 \otimes \dots \otimes x_k$ is the corresponding right eigenvector. Such eigenvalue problems arise in a variety of applications [4], particularly in mathematical physics when the method of separation of variables is used to solve boundary value problems [124]. The result of the separation is a multiparameter system of ordinary differential equations. The multiparameter eigenproblems can be considered as particularly structured generalized eigenvalue problems [3]; see [67, 68, 104] for algorithms and perturbation results based on this connection.

Nonnegative matrices: A nonnegative matrix A is a matrix whose entries are all nonnegative, $a_{ij} \geq 0$ for all i, j . These matrices provide a vast area of research because of

strong links to Markov chains, graph clustering and other practically important applications [12]. See [83] and the references therein for structure-exploiting algorithms with applications to computing Google's PageRank. More on structured perturbation results can be found, e.g., in [33, 69, 98].

Polynomial eigenvalue problems: In a matrix polynomial eigenvalue problem $\lambda \in \mathbb{C}$ and $x \in \mathbb{C}^n$ is sought such that $p(\lambda)x = (M_0 + \lambda M_1 + \dots + \lambda^k M_k)x = 0$ with coefficient matrices $M_j \in \mathbb{C}^{n \times n}$, $j = 0, \dots, k$. A standard approach in order to solve such eigenvalue problems is the use of linearization. Any linearization of a matrix polynomial gives rise to a structured eigenvalue problems [92]. See [95] for a survey of theoretical and algorithmic work on how to exploit the structure of such linearizations and further structure induced by the coefficients of the matrix polynomial.

Palindromic eigenvalue problems: A matrix polynomial $p(\lambda) = M_0 + \lambda M_1 + \dots + \lambda^k M_k$ with coefficient matrices $M_j \in \mathbb{C}^{n \times n}$, $j = 0, \dots, k$ is called palindromic if $p(\lambda) = \lambda^k (p(\lambda^{-1}))^T = \lambda^k M_0^T + \lambda^{k-1} M_1^T + \dots + \lambda M_{k-1}^T + M_k^T$. This class of structured generalized eigenvalue problems has recently been investigated [65, 91].

Hierarchical matrices: The sign function iteration preserves such matrices and can be used to compute spectral projectors and invariant subspaces very efficiently [56, 57].

Semi-separable matrices: Developing efficient and structure-preserving algorithms for semi-separable and related matrices has recently become an active field of research, see, e.g., [32, 122]. Little is known on structured perturbation results.

4 Conclusions and Outlook

In this paper, we have summarized some of the existing knowledge on structured eigenvalue problems. It turns out that quite often the structure of a problem is reflected in the eigenvalues, e.g., eigenvalues appearing in pairs or quadruples. Using standard eigensolvers the special structures of these problems are neglected, often leading to unstructured rounding errors which destroy the eigenvalue pairings. Structure-preserving algorithms prohibit this effect. Moreover, such methods can reduce computational time and improve upon eigenvalue accuracy.

A number of matrix eigenvalue problems arise from the discretization and linearization of nonlinear infinite-dimensional eigenvalue problems. Sometimes an unfortunate choice of the discretization and/or linearization hides relevant structures of the problem. In such cases, it is worthwhile to reconsider the original problem trying to capture these structures.

Recent improvements of the QR algorithm (e.g., aggressive early deflation [22]) may be extended to structured algorithms, but little work has been done in this direction so far. A commonly underappreciated aspect is the development of publicly available software for structured eigenvalue problems.

Acknowledgements The work of the second author was supported by the DFG Research Center MATHEON "Mathematics for key technologies" in Berlin and by an Emmy Noether fellowship from the DFG. The authors would like to thank Michael Karow for providing the MATLAB functions that were used to draw the pseudospectra displayed in Figures 3 and 4.

References

- [1] G. S. Ammar, W. B. Gragg, and L. Reichel. On the eigenproblem for orthogonal matrices. In *Proc. IEEE Conference on Decision and Control*, pages 1963–1966, 1985.
- [2] G. S. Ammar, W. B. Gragg, and L. Reichel. Direct and inverse unitary eigenproblems in signal processing: An overview. In M. S. Moonen, G. H. Golub, and B. L. R. De Moor, editors, *Linear algebra for large scale and real-time applications (Leuven, 1992)*, volume 232 of *NATO Adv. Sci. Inst. Ser. E Appl. Sci.*, pages 341–343, Dordrecht, 1993. Kluwer Acad. Publ.
- [3] F. V. Atkinson. *Multiparameter eigenvalue problems*. Academic Press, New York, 1972.
- [4] F.V. Atkinson. Multiparameter spectral theory. *Bull. Amer. Math. Soc.*, 74:1–27, 1968.
- [5] Z. Bai, J. W. Demmel, J. J. Dongarra, A. Ruhe, and H. van der Vorst, editors. *Templates for the Solution of Algebraic Eigenvalue Problems*. Software, Environments, and Tools. SIAM, Philadelphia, PA, 2000.
- [6] P. Benner. Computational methods for linear-quadratic optimization. *Supplemento ai Rendiconti del Circolo Matematico di Palermo, Serie II*, No. 58:21–56, 1999.
- [7] P. Benner, R. Byers, V. Mehrmann, and H. Xu. Numerical computation of deflating subspaces of skew-Hamiltonian/Hamiltonian pencils. *SIAM J. Matrix Anal. Appl.*, 24(1), 2002.
- [8] P. Benner and H. Faßbender. An implicitly restarted symplectic Lanczos method for the Hamiltonian eigenvalue problem. *Linear Algebra Appl.*, 263:75–111, 1997.
- [9] P. Benner, D. Kressner, and V. Mehrmann. Skew-Hamiltonian and Hamiltonian eigenvalue problems: Theory, algorithms and applications. In Z. Drmač, M. Marušić, and Z. Tutek, editors, *Proceedings of the Conference on Applied Mathematics and Scientific Computing, Brijuni (Croatia), June 23-27, 2003*, pages 3–39. Springer-Verlag, 2005.
- [10] P. Benner, V. Mehrmann, and H. Xu. A numerically stable, structure preserving method for computing the eigenvalues of real Hamiltonian or symplectic pencils. *Numer. Math.*, 78(3):329–358, 1998.
- [11] P. Benner, V. Mehrmann, and H. Xu. Perturbation analysis for the eigenvalue problem of a formal product of matrices. *BIT*, 42(1):1–43, 2002.
- [12] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*, volume 9 of *Classics in Applied Mathematics*. SIAM, Philadelphia, PA, 1994. Revised reprint of the 1979 original.
- [13] D. A. Bini, L. Gemignani, and F. Tisseur. The Ehrlich-Aberth method for the nonsymmetric tridiagonal eigenvalue problem. NA Report 428, Manchester Centre for Computational Mathematics, 2003. To appear in *SIAM J. Matrix Anal. Appl.*
- [14] M. Bohner. Linear Hamiltonian difference systems: Disconjugacy and Jacobi–type conditions. *J. Math. Anal. Appl.*, 199:804–826, 1996.
- [15] B. Bohnhorst. *Beiträge zur numerischen Behandlung des unitären Eigenwertproblems*. PhD thesis, Fakultät für Mathematik, Universität Bielefeld, Bielefeld, Germany, 1993.
- [16] B. Bohnhorst, A. Bunse-Gerstner, and H. Faßbender. On the perturbation theory for unitary eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 21(3):809–824, 2000.
- [17] A. Bojanczyk, G. H. Golub, and P. Van Dooren. The periodic Schur decomposition; algorithm and applications. In *Proc. SPIE Conference*, volume 1770, pages 31–42, 1992.
- [18] A. Böttcher and S. Grudsky. Structured condition numbers of large Toeplitz matrices are rarely better than usual condition numbers. *Numer. Linear Algebra Appl.*, 12:95–102, 2005.
- [19] A. Böttcher and K. Rost. Topics in the numerical linear algebra of Toeplitz and Hankel matrices. *GAMM Mitteilungen*, 27:174–188, 2004.
- [20] S. Boyd, V. Balakrishnan, and P. Kabamba. A bisection method for computing the \mathcal{H}_∞ norm of a transfer matrix and related problems. *Math. Control, Signals, Sys.*, 2:207–219, 1989.
- [21] K. Braman, R. Byers, and R. Mathias. The multishift *QR* algorithm. I. Maintaining well-focused shifts and level 3 performance. *SIAM J. Matrix Anal. Appl.*, 23(4):929–947, 2002.
- [22] K. Braman, R. Byers, and R. Mathias. The multishift *QR* algorithm. II. Aggressive early deflation. *SIAM J. Matrix Anal. Appl.*, 23(4):948–973, 2002.

- [23] R. Bru, R. Cantó, and B. Ricarte. Modelling nitrogen dynamics in citrus trees. *Mathematical and Computer Modelling*, 38:975–987, 2003.
- [24] J. R. Bunch. The weak and strong stability of algorithms in numerical linear algebra. *Linear Algebra Appl.*, 88/89:49–66, 1987.
- [25] A. Bunse-Gerstner and L. Elsner. Schur parameter pencils for the solution of the unitary eigenproblem. *Linear Algebra Appl.*, 154/156:741–778, 1991.
- [26] A. Bunse-Gerstner and H. Faßbender. Error bounds in the isometric Arnoldi process. *J. Comput. Appl. Math.*, 86(1):53–72, 1997.
- [27] A. Bunse-Gerstner and C. He. On a Sturm sequence of polynomials for unitary Hessenberg matrices. *SIAM J. Matrix Anal. Appl.*, 16(4):1043–1055, 1995.
- [28] R. Byers. A Hamiltonian QR algorithm. *SIAM J. Sci. Statist. Comput.*, 7(1):212–229, 1986.
- [29] R. Byers. A bisection method for measuring the distance of a stable to unstable matrices. *SIAM J. Sci. Statist. Comput.*, 9:875–881, 1988.
- [30] R. Byers and D. Kressner. On the condition of a complex eigenvalue under real perturbations. *BIT*, 44(2):209–215, 2004.
- [31] R. Byers and D. Kressner. Structured condition numbers for invariant subspaces, 2005. In preparation.
- [32] S. Chandrasekaran and M. Gu. A divide-and-conquer algorithm for the eigendecomposition of symmetric block-diagonal plus semiseparable matrices. *Numer. Math.*, 96(4):723–731, 2004.
- [33] G. E. Cho and C. D. Meyer. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra Appl.*, 335:137–150, 2001.
- [34] D. Chu, X. Liu, and V. Mehrmann. A numerically backwards stable method for computing the Hamiltonian Schur form. Preprint 24-2004, Institut für Mathematik, TU Berlin, 2004.
- [35] E. K.-W. Chu and W.-W. Lin. Perturbation of eigenvalues for periodic matrix pairs via the Bauer-Fike theorems. *Linear Algebra Appl.*, 378:183–202, 2004.
- [36] M. T. Chu and G. H. Golub. Structured inverse eigenvalue problems. *Acta Numer.*, 11:1–71, 2002.
- [37] J. K. Cullum and R. A. Willoughby. *Lanczos algorithms for large symmetric eigenvalue computations. Vol. 1*, volume 41 of *Classics in Applied Mathematics*. SIAM, Philadelphia, PA, 2002. Reprint of the 1985 original.
- [38] J. J. M. Cuppen. A divide and conquer method for the symmetric tridiagonal eigenproblem. *Numer. Math.*, 36(2):177–195, 1980/81.
- [39] J.J.M. Cuppen. A divide and conquer method for the symmetric tridiagonal eigenproblem. *Numer. Math.*, 36:177–195, 1981.
- [40] J. W. Demmel. *Applied Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.
- [41] J. W. Demmel and K. Veselić. Jacobi’s method is more accurate than QR. *SIAM J. Matrix Anal. Appl.*, 13(4):1204–1245, 1992.
- [42] I. S. Dhillon. Current inverse iteration software can fail. *BIT*, 38(4):685–704, 1998.
- [43] I. S. Dhillon and B. N. Parlett. Orthogonal eigenvectors and relative gaps. *SIAM J. Matrix Anal. Appl.*, 25(3):858–899, 2003.
- [44] I. S. Dhillon and B. N. Parlett. Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices. *Linear Algebra Appl.*, 387:1–28, 2004.
- [45] P. J. Eberlein. A Jacobi-like method for the automatic computation of eigenvalues and eigenvectors of an arbitrary matrix. *J. Soc. Indust. Appl. Math.*, 10:74–88, 1962.
- [46] P. J. Eberlein and C. P. Huang. Global convergence of the QR algorithm for unitary matrices with some results for normal matrices. *SIAM J. Numer. Anal.*, 12:97–104, 1975.
- [47] H. Faßbender. *Symplectic methods for the symplectic eigenproblem*. Kluwer Academic/Plenum Publishers, New York, 2000.
- [48] H. Faßbender and K. D. Ikramov. Computing matrix-vector products with centrosymmetric and centrohermitian matrices. *Linear Algebra Appl.*, 364:235–241, 2003.

- [49] W. R. Ferng, W.-W. Lin, and C.-S. Wang. The shift-inverted J -Lanczos algorithm for the numerical solutions of large sparse algebraic Riccati equations. *Comput. Math. Appl.*, 33(10):23–40, 1997.
- [50] G. H. Golub and C. Reinsch. Singular value decomposition and least squares solution. *Numer. Math.*, 14:403–420, 1970.
- [51] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.
- [52] W. B. Gragg. Positive definite Toeplitz matrices, the Hessenberg process for isometric operators, and the Gauss quadrature on the unit circle. In *Numerical methods of linear algebra (Russian)*, pages 16–32. Moskov. Gos. Univ., Moscow, 1982.
- [53] W. B. Gragg. The QR algorithm for unitary Hessenberg matrices. *J. Comput. Appl. Math.*, 16:1–8, 1986.
- [54] W. B. Gragg and L. Reichel. A divide and conquer method for unitary and orthogonal eigenproblems. *Numer. Math.*, 57(8):695–718, 1990.
- [55] S. Graillat. Structured condition number and backward error for eigenvalue problems. Research Report RR2005-01, LP2A, University of Perpignan, France, 2005.
- [56] L. Grasedyck. Existence of a low rank or \mathcal{H} -matrix approximant to the solution of a Sylvester equation. *Numer. Linear Algebra Appl.*, 11(4):371–389, 2004.
- [57] L. Grasedyck, W. Hackbusch, and B. N. Khoromskij. Solution of large scale algebraic matrix Riccati equations by use of hierarchical matrices. *Computing*, 70(2):121–165, 2003.
- [58] M. Gu and S.C. Eisenstat. A divide and conquer method for the symmetric tridiagonal eigenproblem. *SIAM J. Matrix Anal. Appl.*, 16:172–191, 1995.
- [59] D. Hacon. Jacobi’s method for skew-symmetric matrices. *SIAM J. Matrix Anal. Appl.*, 14(3):619–628, 1993.
- [60] V. Hari. On the quadratic convergence of the Paardekooper method. I. *Glas. Mat. Ser. III*, 17(37)(1):183–195, 1982.
- [61] V. Hari and N. H. Rhee. On the quadratic convergence of the Paardekooper method. II. *Glas. Mat. Ser. III*, 27(47)(2):369–391, 1992.
- [62] S. Helsen, A. B. J. Kuijlaars, and M. Van Barel. Convergence of the isometric Arnoldi process. *SIAM J. Matrix Anal. Appl.*, 26(3):782–809, 2005.
- [63] J. J. Hench and A. J. Laub. Numerical solution of the discrete-time periodic Riccati equation. *IEEE Trans. Automat. Control*, 39(6):1197–1210, 1994.
- [64] D. J. Higham and N. J. Higham. Structured backward error and condition of generalized eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 20(2):493–512, 1999.
- [65] A. Hilliges, C. Mehl, and V. Mehrmann. On the solution of palindromic eigenvalue problems. In *Proceedings of ECCOMAS, Jyväskylä, Finland*, 2004.
- [66] D. Hinrichsen and A. J. Pritchard. *Mathematical Systems Theory I*, volume 48 of *Texts in Applied Mathematics*. Springer-Verlag, 2005.
- [67] M. E. Hochstenbach and B. Plestenjak. A Jacobi-Davidson type method for a right definite two-parameter eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 24(2):392–410, 2002.
- [68] M. E. Hochstenbach and B. Plestenjak. Backward error, condition numbers, and pseudospectra for the multiparameter eigenvalue problem. *Linear Algebra Appl.*, 375:63–81, 2003.
- [69] I. Ipsen. Sensitivity of PageRank, 2005. In preparation.
- [70] C. F. Jagels and L. Reichel. The isometric Arnoldi process and an application to iterative solution of large linear systems. In *Iterative methods in linear algebra (Brussels, 1991)*, pages 361–369. North-Holland, Amsterdam, 1992.
- [71] C. F. Jagels and L. Reichel. A fast minimal residual algorithm for shifted unitary matrices. *Numer. Linear Algebra Appl.*, 1(6):555–570, 1994.
- [72] B. Kågström. The QR algorithm to find the eigensystem of a skew-symmetric matrix. Report UMINF-14.71, Institute of Information Processing, University of Umeå, Sweden, 1971.
- [73] M. Karow. *Geometry of Spectral Value Sets*. PhD thesis, Universität Bremen, Fachbereich 3 (Mathematik & Informatik), Bremen, Germany, 2003.

- [74] M. Karow, D. Kressner, and F. Tisseur. Structured eigenvalue condition numbers. Numerical Analysis Report No. 467, Manchester Centre for Computational Mathematics, Manchester, England, 2005.
- [75] M. Kleinstauber, U. Helmke, and K. Hüper. Jacobi's algorithm on compact Lie algebras. *SIAM J. Matrix Anal. Appl.*, 26(1):42–69, 2004.
- [76] M. Konstantinov, D.-W. Gu, V. Mehrmann, and P. Petkov. *Perturbation theory for matrix equations*, volume 9 of *Studies in Computational Mathematics*. North-Holland Publishing Co., Amsterdam, 2003.
- [77] M. Konstantinov, V. Mehrmann, and P. Petkov. Perturbation analysis of Hamiltonian Schur and block-Schur forms. *SIAM J. Matrix Anal. Appl.*, 23(2):387–424, 2001.
- [78] D. Kressner. The periodic QR algorithm is a disguised QR algorithm, 2003. To appear in *Linear Algebra Appl.*
- [79] D. Kressner. Perturbation bounds for isotropic invariant subspaces of skew-Hamiltonian matrices, 2003. To appear in *SIAM J. Matrix Anal. Appl.*
- [80] D. Kressner. A Krylov-Schur algorithm for matrix products, 2004. Submitted.
- [81] D. Kressner. *Numerical Methods and Software for General and Structured Eigenvalue Problems*. PhD thesis, TU Berlin, Institut für Mathematik, Berlin, Germany, 2004. Revised version to appear in *Lecture Notes in Computational Science and Engineering*, Springer, Heidelberg.
- [82] P. Lancaster and L. Rodman. *The Algebraic Riccati Equation*. Oxford University Press, Oxford, 1995.
- [83] A. N. Langville and C. D. Meyer. Deeper inside pagerank. *Internet Math.*, 1(3):335–380, 2004.
- [84] A. J. Laub. Invariant subspace methods for the numerical solution of Riccati equations. In S. Bittanti, A. J. Laub, and J. C. Willems, editors, *The Riccati Equation*, pages 163–196. Springer-Verlag, Berlin, 1991.
- [85] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK users' guide*. SIAM, Philadelphia, PA, 1998. Solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods.
- [86] W.-W. Lin and J.-G. Sun. Perturbation analysis for the eigenproblem of periodic matrix pairs. *Linear Algebra Appl.*, 337:157–187, 2001.
- [87] W.-W. Lin, P. Van Dooren, and Q.-F. Xu. Periodic invariant subspaces in control. In *Proc. of IFAC Workshop on Periodic Control Systems, Como, Italy*, 2001.
- [88] K. Lust. *Numerical Bifurcation Analysis of Periodic Solutions of Partial Differential Equations*. PhD thesis, Department of Computer Science, KU Leuven, Belgium, 1997.
- [89] K. Lust. Improved numerical Floquet multipliers. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 11(9):2389–2410, 2001.
- [90] D. S. Mackey, N. Mackey, and D. Dunlavy. Structure preserving algorithms for perplectic eigenproblems. *Electron. J. Linear Algebra*, 13:10–39, 2005.
- [91] D. S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann. Palindromic polynomial eigenvalue problems: Good vibrations from good linearizations, 2005. Submitted.
- [92] D. S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann. Vector spaces of linearizations for matrix polynomials, 2005. Submitted.
- [93] The MathWorks, Inc., Cochituate Place, 24 Prime Park Way, Natick, Mass, 01760. *MATLAB Version 6.5*, 2002.
- [94] V. Mehrmann. *The Autonomous Linear Quadratic Control Problem, Theory and Numerical Solution*. Number 163 in *Lecture Notes in Control and Information Sciences*. Springer-Verlag, Heidelberg, 1991.
- [95] V. Mehrmann and H. Voss. Nonlinear eigenvalue problems: A challenge for modern eigenvalue methods. *GAMM Mitteilungen*, 27:121–152, 2004.
- [96] V. Mehrmann and D. S. Watkins. Structure-preserving methods for computing eigenpairs of large sparse skew-Hamiltonian/Hamiltonian pencils. *SIAM J. Sci. Comput.*, 22(6):1905–1925, 2000.
- [97] A. Melman. Symmetric centrosymmetric matrix-vector multiplication. *Linear Algebra Appl.*, 320(1-3):193–198, 2000.

- [98] C. D. Meyer and G. W. Stewart. Derivatives and perturbations of eigenvectors. *SIAM J. Numer. Anal.*, 25(3):679–691, 1988.
- [99] C. B. Moler and G. W. Stewart. An algorithm for generalized matrix eigenvalue problems. *SIAM J. Numer. Anal.*, 10:241–256, 1973.
- [100] Z. J. Mou. Fast algorithms for computing symmetric/Hermitian matrix-vector products. *Electronic Letters*, 27:1272–1274, 1991.
- [101] S. Noschese and L. Pasquini. Eigenvalue condition numbers: zero-structured versus traditional. Preprint 28, Mathematics Department, University of Rome La Sapienza, Italy, 2004.
- [102] M. H. C. Paardekooper. An eigenvalue algorithm for skew-symmetric matrices. *Numer. Math.*, 17:189–202, 1971.
- [103] B. N. Parlett. *The Symmetric Eigenvalue Problem*, volume 20 of *Classics in Applied Mathematics*. SIAM, Philadelphia, PA, 1998. Corrected reprint of the 1980 original.
- [104] B. Plestenjak. A continuation method for a right definite two-parameter eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 21(4):1163–1184, 2000.
- [105] R. M. Reid. Some eigenvalue properties of persymmetric matrices. *SIAM Rev.*, 39(2):313–316, 1997.
- [106] N. H. Rhee and V. Hari. On the cubic convergence of the Paardekooper method. *BIT*, 35(1):116–132, 1995.
- [107] S. M. Rump. Eigenvalues, pseudospectrum and structured perturbations, 2005. Submitted.
- [108] F. Schmidt, T. Friese, L. Zschiedrich, and P. Deuffhard. Adaptive multigrid methods for the vectorial Maxwell eigenvalue problem for optical waveguide design. In W. Jäger and H.-J. Krebs, editors, *Mathematics. Key Technology for the Future*, pages 279–292, 2003.
- [109] G. L. G. Sleijpen and H. A. Van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425, 1996.
- [110] D. C. Sorensen. Implicit application of polynomial filters in a k -step Arnoldi method. *SIAM J. Matrix Anal. Appl.*, 13:357–385, 1992.
- [111] G. W. Stewart. Error and perturbation bounds for subspaces associated with certain eigenvalue problems. *SIAM Rev.*, 15:727–764, 1973.
- [112] G. W. Stewart. A Jacobi-like algorithm for computing the Schur decomposition of a non-Hermitian matrix. *SIAM J. Sci. Statist. Comput.*, 6(4):853–864, 1985.
- [113] G. W. Stewart and J.-G. Sun. *Matrix Perturbation Theory*. Academic Press, New York, 1990.
- [114] J.-G. Sun. Stability and accuracy: Perturbation analysis of algebraic eigenproblems. Technical report UMINF 98-07, Department of Computing Science, University of Umeå, Umeå, Sweden, 1998. Revised 2002.
- [115] J.-G. Sun. Perturbation analysis of algebraic Riccati equations. Technical report UMINF 02-03, Department of Computing Science, University of Umeå, Umeå, Sweden, 2002.
- [116] F. Tisseur. A chart of backward errors for singly and doubly structured eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 24(3):877–897, 2003.
- [117] F. Tisseur and K. Meerbergen. The quadratic eigenvalue problem. *SIAM Rev.*, 43(2):235–286, 2001.
- [118] L. N. Trefethen and M. Embree. *Spectra and Pseudospectra: The Behavior of Non-Normal Matrices and Operators*. Princeton University Press, 2005. To appear.
- [119] C. F. Van Loan. A general matrix eigenvalue algorithm. *SIAM J. Numer. Anal.*, 12(6):819–834, 1975.
- [120] C. F. Van Loan. A symplectic method for approximating all the eigenvalues of a Hamiltonian matrix. *Linear Algebra Appl.*, 61:233–251, 1984.
- [121] C. F. Van Loan. *Computational Frameworks For the Fast Fourier Transform*. SIAM, Philadelphia, PA, 1992.
- [122] R. Vandebril. *Semiseparable matrices and the symmetric eigenvalue problem*. PhD thesis, KU Leuven, Departement Computerwetenschappen, 2004.
- [123] A. Varga and P. Van Dooren. Computational methods for periodic systems - an overview. In *Proc. of IFAC Workshop on Periodic Control Systems, Como, Italy*, pages 171–176, 2001.

- [124] H. Volkmer. *Multiparameter Problems and Expansion Theorems*, volume 1356 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1988.
- [125] H. Voss. A symmetry exploiting Lanczos method for symmetric Toeplitz matrices. *Numer. Algorithms*, 25(1-4):377–385, 2000.
- [126] T.-L. Wang and W. B. Gragg. Convergence of the shifted QR algorithm for unitary Hessenberg matrices. *Math. Comp.*, 71(240):1473–1496, 2002.
- [127] T.-L. Wang and W. B. Gragg. Convergence of the unitary QR algorithm with a unimodular Wilkinson shift. *Math. Comp.*, 72(241):375–385, 2003.
- [128] R. C. Ward and L. J. Gray. Eigensystem computation for skew-symmetric matrices and a class of symmetric matrices. *ACM Trans. Math. Software*, 4(3):278–285, 1978.
- [129] D. S. Watkins. On Hamiltonian and symplectic Lanczos processes. *Linear Algebra Appl.*, 385:23–45, 2004.
- [130] D. S. Watkins. Product eigenvalue problems. *SIAM Rev.*, 47:3–40, 2005.
- [131] D.S. Watkins. *Fundamentals of Matrix Computation*. Wiley, Chichester, UK, 1991.
- [132] J. R. Weaver. Centrosymmetric (cross-symmetric) matrices, their basic properties, eigenvalues, and eigenvectors. *Amer. Math. Monthly*, 92(10):711–717, 1985.
- [133] N. J. Wildberger. Diagonalization in compact Lie algebras and a new proof of a theorem of Kostant. *Proc. Amer. Math. Soc.*, 119(2):649–655, 1993.
- [134] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Clarendon Press, Oxford, 1965.