

Tridiagonalizing Complex Symmetric Matrices in Waveguide Simulations

W.N. Gansterer¹, H. Schabauer¹, C. Pacher², and N. Finger²

¹ University of Vienna, Research Lab Computational Technologies and Applications
{wilfried.gansterer,hannes.schabauer}@univie.ac.at

² Austrian Research Centers GmbH - ARC, Smart Systems Division
{christoph.pacher,norman.finger}@arcs.ac.at

Abstract. We discuss a method for solving *complex symmetric (non-Hermitian)* eigenproblems $Ax = \lambda Bx$ arising in an application from optoelectronics, where reduced accuracy requirements provide an opportunity for trading accuracy for performance. In this case, the objective is to exploit the structural symmetry. Consequently, our focus is on a non-Hermitian tridiagonalization process. For solving the resulting complex symmetric tridiagonal problem, a variant of the Lanczos algorithm is used. Based on Fortran implementations of these algorithms, we provide extensive experimental evaluations. Runtimes and numerical accuracy are compared to the standard routine for non-Hermitian eigenproblems, LAPACK/zgeev. Although the performance results reveal that more work is needed in terms of increasing the fraction of Level 3 BLAS in our tridiagonalization routine, the numerical accuracy achieved with the non-Hermitian tridiagonalization process is very encouraging and indicates important research directions for this class of eigenproblems.

Keywords: Tridiagonalization, complex symmetric eigenvalue problems, waveguide simulation, optoelectronic.

1 Introduction

We discuss methods for efficiently tridiagonalizing a *complex symmetric (non-Hermitian)* matrix. The term *complex matrix* is used to denote a matrix which has at least one element with a nonzero imaginary part.

Tridiagonalization is an important preprocessing step in *reduction-based* (tridiagonalization-based) methods for computing eigenvalues and eigenvectors of dense real symmetric or complex Hermitian matrices. In the context considered here, the underlying complex symmetric eigenvalue problem (EVP) has similar structural, but different mathematical properties. Given matrices $\hat{A}, \hat{B} \in \mathbb{C}^{n \times n}$ with $\hat{A} = \hat{A}^\top$ (but $\hat{A} \neq \hat{A}^H$) and $\hat{B} = \hat{B}^\top$ (but $\hat{B} \neq \hat{B}^H$), the objective is to efficiently compute eigenvalues λ and eigenvectors y of the generalized EVP

$$\hat{A}y = \lambda \hat{B}y . \quad (1)$$

The main challenge is to find ways for utilizing the structural symmetry in the absence of the mathematical properties of Hermitian matrices.

Although problems of the type (1) do not occur as frequently in practice as real symmetric or complex Hermitian problems, there are some important applications where they arise (see, for example, [1,2,3]). The efforts summarized in this paper are motivated by the simulation of guided-wave multi-section devices in optoelectronics. As described in Section 2, techniques for numerically solving Maxwell's equations in this context lead to dense EVPs of the type (1).

Analogously to Hermitian problems, one possible approach for solving problem (1) starts with reducing it to *standard form* where $\hat{B} = I_n$. Complex symmetry allows for special techniques in this reduction step which are not discussed here. After that, a tridiagonalization process is performed on the standard EVP which results in a similar complex symmetric tridiagonal matrix T . After this tridiagonalization step, eigenvalues and eigenvectors of T are computed and the eigenvectors are backtransformed.

In the following, we focus on *symmetry-preserving* approaches for efficiently tridiagonalizing a complex symmetric matrix. This functionality constitutes a central step in the solution process outlined above and is one way of exploiting the available structure. A more detailed discussion of the other steps involved in solving (1) will be provided in a forthcoming paper.

Mathematically speaking, structural symmetry is not a very distinctive feature of complex matrices, since *every* matrix $A \in \mathbb{C}^{n \times n}$ is similar to a complex symmetric matrix [1]. In contrast to a real symmetric matrix a complex symmetric matrix A is not necessarily diagonalizable. Nevertheless, structural symmetry is of great interest for the development of space- and time-efficient algorithms. Obviously, half of the information in a complex symmetric matrix is redundant, and efficient algorithms should be able to take advantage of this fact in terms of memory requirements as well as in terms of computational effort. The utilization of this purely structural property in the absence of important mathematical properties of Hermitian matrices requires a trade-off in numerical stability. In order to perform a *symmetry preserving* similarity transformation, the transformation matrix $Q \in \mathbb{C}^{n \times n}$ needs to be *complex orthogonal* (but not unitary), that is, it has to satisfy $Q^T Q = I_n$.

Related Work. Various non-reduction based methods for solving complex symmetric EVPs have been proposed, for example, based on the Jacobi method [4], on the Lanczos method [5], or on variants of the Jacobi-Davidson method [6]. For dense matrices, reduction-based methods tend to be more efficient. A modified conventional Householder-based reduction method has been described in [2]. The tridiagonalization of a dense complex symmetric matrix has also been investigated in [3]. In [2], the resulting tridiagonal complex symmetric problem is solved using a modified QR algorithm. Other related approaches for computing eigenvalues of a complex symmetric tridiagonal matrix were discussed in [7,8].

Synopsis. In Section 2 of this paper, the motivating application from optoelectronics is summarized. In Section 3, the tridiagonalization method investigated in this paper is discussed in detail, and Section 4 contains experimental results. Conclusions and future work are summarized in Section 5.

2 Guided-Wave Multisection Devices

The use of high-index contrast waveguides (WGs) in novel guided-wave devices for telecom- and sensing applications allows a very versatile tailoring of the flow of light. However, an efficient design requires the direct numerical solution of Maxwell's equations in inhomogeneous media. In many important cases such devices can be successfully modeled as follows: (i) in the x -direction (direction of propagation) the material parameters are piecewise constant, (ii) the material parameters and the optical fields do not depend on the y -coordinate, and (iii) in the z -direction the material parameters are allowed to vary arbitrarily. Usually, the z -dimension is of the order of up to several tens of wavelengths whereas the device extension into the x -direction is several hundreds of wavelengths.

A powerful numerical method for the solution of Maxwell's equations in such WG-based devices is the *eigenmode expansion technique* (which is often referred to as *mode-matching (MM) technique*) [9,10,11], where the electromagnetic field components in each subsection being homogeneous in the x -direction are represented in terms of a set of *local* eigenmodes. MM requires a small computational effort compared to other numerical techniques like two-dimensional finite-elements or FDTD which can be regarded as "brute-force" methods from the viewpoint of device physics. However, MM can only be as stable and efficient as the algorithms used to determine the required set of local WG modes. Due to the open boundary conditions (see Section 2.1) and materials with complex dielectric permittivities these local eigenmodes have typically complex eigenvalues which makes their correct classification very difficult: numerical instabilities can arise from an improper truncation of the mode spectrum. In a recently developed variant of the MM technique — the so-called *variational mode-matching (VMM)* [12] — this stability problem is avoided by applying a Galerkin approach to the local wave equations and taking into account the *whole* spectrum of the discretized operators.

2.1 The VMM-Approach

Within the 2D-assumption $\partial_y(\cdot) = 0$, Maxwell's equations for dielectric materials characterized by the dielectric permittivity $\varepsilon(x, z)$ take the form

$$\partial_x a \partial_x \phi + \partial_z a \partial_z \phi + k_0^2 b \phi = 0, \quad (2)$$

where $\phi = E_y$, $a = 1$, $b = \varepsilon$ for TE- and $\phi = H_y$, $a = \frac{1}{\varepsilon}$, $b = 1$ for TM-polarization, respectively; $k_0 = \frac{2\pi}{\lambda_0}$ (vacuum wavelength λ_0).

In the z -direction, the simulation domain is $0 \leq z \leq L$. To permit an accurate description of radiation fields, an artificial absorber (that mimics an open domain) has to be installed near $z = 0$ and $z = L$. For this purpose so-called *perfectly-matched layers (PMLs)* are used by employing the complex variable stretching approach [13], i. e., in the vicinity of the domain boundaries the coordinate z is extended into the complex plane: $z \mapsto \tilde{z} = z + \imath \int_0^z d\tau \sigma(\tau)$, where σ is the PML parameter determining the absorption strength. At $z = 0$ and $z = L$

Dirichlet- or Neumann boundary conditions (BCs) are set. However, they should not have a significant influence on the overall optical field since the physical BCs must be given by the PMLs. In the x -direction, the structure is divided into n_l local WGs, which expand over $x_{l-1} \leq x \leq x_l = x_{l-1} + d_l$ with $1 \leq l \leq n_l$.

Under the necessary condition that ε does not depend on x Eq. (2) can be solved inside each local WG l with the separation ansatz

$$\phi^{(l)}(x, \tilde{z}) = \sum_{j=1}^{n_\varphi} \varphi_j(\tilde{z}) \sum_{\rho=1}^{n_\varphi} c_{j\rho}^{(l)} \left[\alpha_{\rho,+}^{(l)} e^{ik_0\nu_\rho^{(l)}(x-x_{l-1})} + \alpha_{\rho,-}^{(l)} e^{-ik_0\nu_\rho^{(l)}(x-x_l)} \right], \quad (3)$$

where ρ labels the local waveguide modes. The transverse shape functions $\varphi_j(\tilde{z})$ (the same set is used for all local WGs) must satisfy the outer boundary conditions. Apart from this constraint, the φ_j 's may be chosen rather freely allowing for adaptive refinement in z -regions where rapid variations of the field are expected. This ansatz reduces the 2D problem to a set of n_l 1D problems.

After inserting Eq. (3) into Eq. (2), Galerkin's method is applied to obtain a discretized version of Eq. (2) for each local WG l . Finally, the coefficients $\alpha_{\rho,\pm}^{(l)}$ are "mode-matched" by imposing the physical boundary conditions at all the x_l -interfaces [12].

2.2 The Complex Symmetric Eigenvalue Problem

For each local WG, the discretized version of Eq. (2) is a generalized EVP of the form

$$Ac_\rho = (\nu_\rho)^2 Bc_\rho, \quad (4)$$

where we have suppressed the index l for simplicity. Here, the ν_ρ 's are the modal refractive indices and the $c_{j\rho}$'s are the corresponding modal expansion coefficients occurring in Eq. (3). A is a sum of a mass- and a stiffness-matrix, $A_{mj} = \int d\tilde{z} \varphi_m(\tilde{z})b(\tilde{z})\varphi_j(\tilde{z}) - \frac{1}{k_0^2} \int d\tilde{z} (\partial_z \varphi_m(\tilde{z}))a(\tilde{z})(\partial_z \varphi_j(\tilde{z}))$, whereas B is a pure mass-matrix: $B_{mj} = \int d\tilde{z} \varphi_m(\tilde{z})a(\tilde{z})\varphi_j(\tilde{z})$.

The generalized EVP (4) has the following properties: (i) A and B are complex symmetric: the complex coordinate \tilde{z} originating from the PMLs (and the possibly complex material constants a and b) are responsible for the complex-valuedness; (ii) B is indefinite (due to the open boundary conditions represented by the PMLs and a possibly negative material constant a); (iii) the typical order of the matrices for 2D problems is 100–1000 (depending on the geometry and the required truncation order of the modal expansion—in 3D models the order can be much higher); (iv) the full spectrum of eigenpairs is required; (v) the required accuracy is of the order 10^{-8} for the eigenpairs corresponding to the lowest order WG modes (approx. 10% of the mode spectrum); a somewhat lower accuracy (approx. 10^{-6}) is acceptable for the remainder of the spectrum; (vi) depending on the WG geometry, some of the eigenvalues (especially those corresponding to the lowest order WG modes) may almost degenerate.

It is evident that an efficient eigenvalue solver which utilizes the symmetry of the EVP (4) as well as its special properties is a very important building block for efficient 2D and 3D optical mode solvers.

3 Methodology

The standard approach to solving a dense complex symmetric EVP (as available, for example, in LAPACK [14]) is to treat it as a nonsymmetric EVP: the complex symmetric matrix is reduced to upper Hessenberg form, from which eigenvalues and eigenvectors are computed using a QR iteration. The main motivation behind investigating tridiagonalization-based approaches as an alternative is the obvious potential for reducing storage and runtime requirements.

In order to preserve symmetry, *complex orthogonal* similarity transformations (COTs) Q are needed which satisfy $Q^T Q = I_n$. In general, $\|Q\|_2 \geq 1$ and thus the application of complex orthogonal matrices can increase numerical errors.

3.1 Splitting Methods

The real part R and the imaginary part S of a complex symmetric matrix $A = R + iS$ are real symmetric matrices. One basic idea, which has been introduced earlier [3], is to separate the tridiagonalization of R from the tridiagonalization of S as much as possible. More specifically, part of a column of R can be eliminated using a (real) orthogonal Householder transformation Q_R . After that, a (smaller) part of the corresponding column of S can be eliminated without causing any fill-in in R using another (real) orthogonal Householder transformation Q_I . Both of these operations are performed in real arithmetic, and both transformation matrices have norm one. Eventually, a single nonzero element below the subdiagonal in S remains to be eliminated. This operation has to be performed in complex arithmetic, using a 2×2 COT, whose norm cannot be bounded a priori. When the column elimination is first performed in R and then in S , we call the procedure *RI variant*. Analogously, it is possible to eliminate first in S and then in R . We call this procedure *IR variant*.

The advantages of splitting methods seem obvious: Most of the computation can be done in real arithmetic, only one third of the transformations are potentially unstable, and this danger can easily be monitored because of the low dimensions of the COTs (only order two).

Complex Orthogonal Transformations. The transformation matrix

$$G := \frac{1}{\sqrt{z^2 + s^2}} \begin{pmatrix} z & s \\ -s & z \end{pmatrix}, \quad z = z_1 + iz_2 \in \mathbb{C}, \quad z_1, z_2, s \in \mathbb{R}, \quad (5)$$

defines a COT since $G^T G = I_2$. Consequently, GAG^T is a similarity transformation of A . In the RI variant, a COT G_{RI} has to be determined such that

$$G_{RI} \begin{pmatrix} a + ib \\ ic \end{pmatrix} = \begin{pmatrix} d + ie \\ 0 \end{pmatrix},$$

where $a, b, c, d, e \in \mathbb{R}$ and $c \neq 0$. Choosing the parameters $z = s(\frac{b}{c} - i\frac{a}{c})$, $s \neq 0$ arbitrary, the COT is given as

$$G_{RI} = \frac{1}{\sqrt{b^2 - a^2 + c^2 - i(2ab)}} \begin{pmatrix} b - ia & c \\ -c & b - ia \end{pmatrix}. \quad (6)$$

In the IR variant, a COT G_{IR} has to be determined such that

$$G_{\text{IR}} \begin{pmatrix} a + ib \\ c \end{pmatrix} = \begin{pmatrix} d + ie \\ 0 \end{pmatrix}.$$

With $z = s \left(\frac{a}{c} + i \frac{b}{c} \right)$, $s \neq 0$ arbitrary, the COT is given as

$$G_{\text{IR}} = \frac{1}{\sqrt{a^2 - b^2 + c^2 + i(2ab)}} \begin{pmatrix} a + ib & c \\ -c & a + ib \end{pmatrix}. \tag{7}$$

3.2 Numerical Aspects

In a splitting method, the complex orthogonal transformations (5) are the only non-unitary transformations, all other transformations used have unit norm. If $\|G\|_2 \gg 1$ the accuracy of the tridiagonalization process could be influenced negatively. G is a normal matrix, and thus its spectral norm is given by its largest eigenvalue in modulus:

$$\|G\|_2 = \left(\frac{1 + \gamma}{1 - \gamma} \right)^{1/4} \quad \text{with} \quad \gamma = \frac{2|z_2 s|}{z_1^2 + z_2^2 + s^2}. \tag{8}$$

If γ approaches one, the accuracy of the tridiagonalization process may deteriorate. For G_{RI} and G_{IR} , respectively, γ in (8) becomes

$$\gamma_{\text{RI}} = \frac{2|ac|}{a^2 + b^2 + c^2}, \quad \gamma_{\text{IR}} = \frac{2|bc|}{a^2 + b^2 + c^2}.$$

We observe that the freedom in choosing the parameter s does not help in controlling the norm of the COT, since γ_{RI} and γ_{IR} are *independent* of s . During the tridiagonalization process, monitoring the norms of the COTs makes it possible to detect potentially large errors. Various strategies have been suggested to avoid large norms, such as the recovery transformations proposed in [3].

Adaptive Elimination Order. The order of processing R and S can be determined independently in each iteration of the tridiagonalization process. For both variants, the norm of each COT can be precomputed with only marginal overhead. Based on this information, the COT with the smaller norm can be selected and the corresponding variant carried out. Obviously, this heuristic choice is only a *local* minimization and there is no guarantee that it minimizes the accumulated norm of all COTs in the tridiagonalization process. Comparison to and combination with recovery transformations are topics of ongoing work.

4 Experimental Evaluation

In our experiments, we used the following routines: `zsyssta` reduces a generalized EVP (\hat{A}, \hat{B}) to a standard EVP (A) , and `zsyevv` solves the standard complex

symmetric EVP. The latter consists of a newly implemented RI tridiagonalization (`zsytridi`), `compev` [15] for computing eigenvalues and `inverm` [15] for computing corresponding eigenvectors of the complex symmetric tridiagonal. `zsyevg` tests the accuracy of the tridiagonalization process by first calling `zsytridi`, followed by a call of `LAPACK/zgeev` on the resulting tridiagonal matrix.

The codes were run on a Sun Fire v40z with 4 dual-core Opteron 875 CPUs (2.2 GHz) and 24 GB main memory. Suse Linux Enterprise Server 10, the GNU Fortran 95 compiler, LAPACK version 3.1.1, GOTO BLAS 1.20, and the AMD Core Math Library (ACML 4.0.1) were used. We experimented with random test matrices with elements in $[0, 2]$ as well as with a real application case.

4.1 Numerical Accuracy

Denoting with (λ_i, x_i) the eigenpairs computed by `LAPACK/zgeev`, and with $(\tilde{\lambda}_i, \tilde{x}_i)$ the eigenpairs computed by `zsyevv`, an eigenvalue error \mathfrak{E} and a residual error \mathfrak{R} have been computed according to

$$\mathfrak{E} := \max_i \frac{|\tilde{\lambda}_i - \lambda_i|}{|\lambda_i|}, \quad \mathfrak{R} = \max_i \frac{\|(A - \tilde{\lambda}_i I_n)\tilde{x}_i\|_2}{\|A\|_2}, \quad i \in \{1, \dots, n\}.$$

Fig. 1 illustrates that the loss of accuracy in the tridiagonalization process itself is surprisingly low! Although the total values of \mathfrak{E} and \mathfrak{R} of `zsyevv` increase up to 10^{-6} , most of this error is due to the Lanczos variant used for solving the tridiagonal problem. The error introduced by the RI tridiagonalization is only about two orders of magnitudes higher than the one of `LAPACK/zgeev`.

1D Waveguide Problem. The waveguide structure is a Si/SiO_x twin waveguide operated in TM-polarization at a wavelength $\lambda_0 = 1.55 \mu\text{m}$. The dielectric constants are $\epsilon_{\text{Si}} = 12.96$ and $\epsilon_{\text{SiO}_x} = 2.25$. The core thickness and -separation are $0.5 \mu\text{m}$ and $0.25 \mu\text{m}$, respectively. The z -extension of the model domain, terminated by electrically perfectly conducting walls, is $10 \mu\text{m}$. The PML-layer thickness is $1 \mu\text{m}$ with the PML-parameter $\sigma = 1$. As shape functions, localized linear hat functions and polynomial bubble functions with a degree up to 24 were used.

For reducing the generalized problem (4) to standard form, we computed a generalized (complex) symmetric Cholesky factor F of \hat{B} . With $\|\hat{B} - FF^T\|_2 = 1.8 \cdot 10^{-16}$, the accuracy of this factorization is satisfactory for our test case. The eigenpairs (λ_i, x_i) of the resulting standard problem computed using GNU OCTAVE were compared with the eigenpairs $(\tilde{\lambda}_i, \tilde{x}_i)$ computed by our routine `zsyevv`. Backtransformation of the eigenvectors leads to a weighted residual error

$$\max_{i=1, \dots, n} \frac{\|(\hat{A} - \tilde{\lambda}_i \hat{B})\tilde{y}_i\|_2}{\|\hat{A}\|_2 \|\hat{B}\|_2} = 3.8 \cdot 10^{-14},$$

which is a very satisfactory accuracy (for this test case, $\|\hat{A}\|_2 = 928, \|\hat{B}\|_2 = 2$).

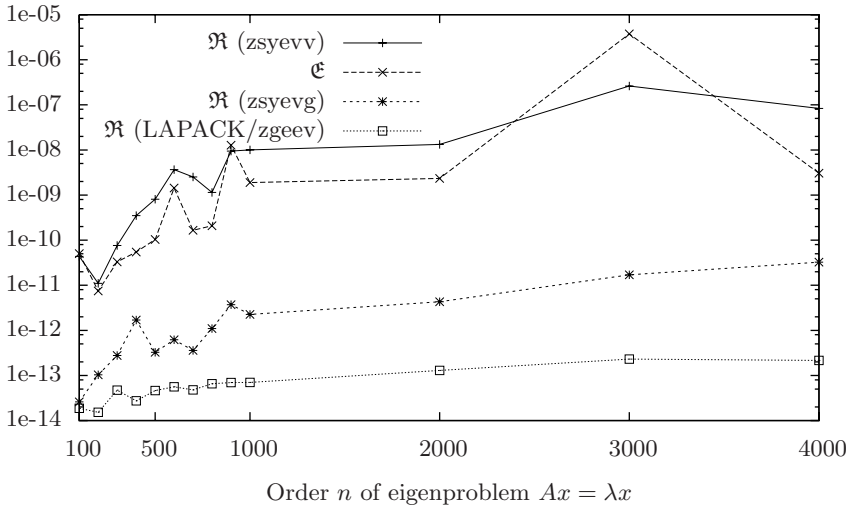


Fig. 1. Accuracy of zsyevv, LAPACK/zgeev, and zsyevg operating on random matrices

4.2 Runtime Performance

We compared our routine zsyevv to LAPACK/zgeev using two different implementations of the BLAS. Fig. 2 shows that the current version of zsyevv is faster than zgeev *only* if the ACML BLAS is used. With the overall faster GOTO BLAS, zgeev outperforms our implementation. At first sight, this result is disappointing. Despite the exploitation of the structure, the new routine is slower than the more general routine for nonsymmetric problems for the best BLAS available.

A more detailed analysis helps to pinpoint the reason. Table 1 shows the percentages of the total runtimes which each of the two routines spent in their different parts for the two different BLAS versions. For our routine zsyevv, the tridiagonalization part zsytridi clearly dominates the computation time for all problem sizes and for both BLAS versions. This shows that our current code zsytridi is unable to take advantage of the faster GOTO BLAS. Three different parts of LAPACK/zgeev have been timed separately: zgehrrd reduces the complex matrix A to upper Hessenberg form, zhseqr computes the eigenvalues of the Hessenberg matrix, and ztrevc computes corresponding eigenvectors. The runtime for all other code parts of LAPACK/zgeev is summed under “rest”. The picture here is quite different. For the ACML BLAS, the operations on the Hessenberg matrix clearly dominate for all problem sizes, whereas for the faster GOTO BLAS, the percentages for the three dominating parts become very similar for large problem sizes.

Summarizing, we observe that our current code cannot utilize a faster BLAS. This is not surprising, since so far it is dominated by Level 2 BLAS operations and more effort is needed to increase the fraction of Level 3 BLAS operations.

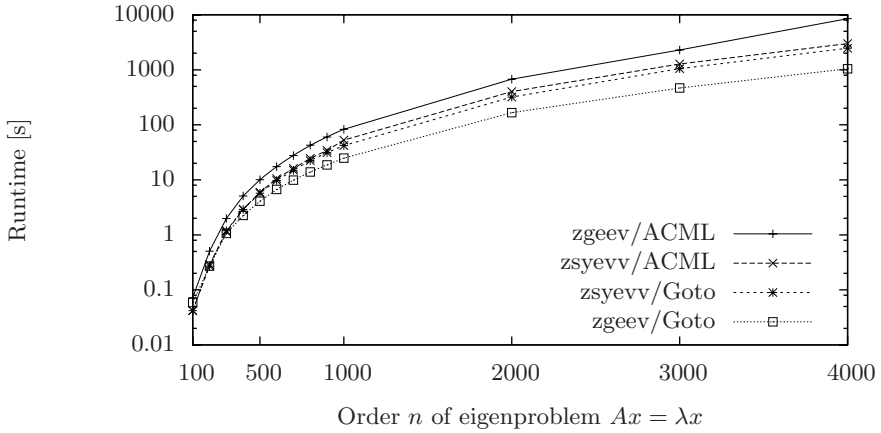


Fig. 2. Runtimes of zsyevv and LAPACK/zgeev operating on random matrices

Table 1. Percentages of runtimes spent in parts of zsyevv and LAPACK/zgeev

		zsyevv			LAPACK/zgeev			
BLAS	n	zsytridi	compev	inverm	zgehrd	zhseqr	ztrevc	rest
ACML	500	87.2	6.5	6.3	8.1	82.6	6.2	3.1
	2000	93.9	1.5	4.6	7.2	84.2	6.2	2.4
	4000	94.5	0.8	4.7	4.4	90.3	3.8	1.5
GOTO	500	87.3	6.5	6.2	15.3	66.9	12.0	5.9
	2000	92.7	1.9	5.4	22.7	50.6	18.9	7.8
	4000	93.7	1.0	5.3	28.6	37.8	23.9	9.7

5 Conclusions and Future Work

Motivated by application problems arising in optoelectronics, a tridiagonalization process for complex symmetric matrices based on complex orthogonal transformations has been investigated. Compared to the standard LAPACK routine for nonsymmetric eigenproblems, the loss of numerical accuracy caused by the potentially unstable tridiagonalization process is surprisingly low in practice. However, partly in contrast to results published earlier [16], the performance benefits achieved are not yet satisfactory, especially for highly optimized BLAS.

The effort summarized here motivates various further research activities. Methodologically, the performance results indicate the need for blocked approaches. This suggests that *non-splitting* methods, where A is not split into real and imaginary part, can be an attractive alternative. For the optoelectronics problem, the matrices in (4) can be made *banded* in some situations by choosing appropriate shape functions. This motivates the investigation of efficient algorithms for generalized banded complex symmetric eigenvalue problems.

References

1. Horn, R.A., Johnson, C.R.: Matrix Analysis. Cambridge University Press, Cambridge (1985)
2. Ohnami, K., Mikami, Y.: Resonance scattering in a two-dimensional non-integrable system. *J. Phys. A* 25, 4903–4912 (1992)
3. Bar-On, I., Ryaboy, V.: Fast diagonalization of large and dense complex symmetric matrices, with applications to quantum reaction dynamics. *SIAM J. Sci. Comput.* 18, 1412–1435 (1997)
4. Leung, A.Y.T., Liu, Y.F.: A generalized complex symmetric eigensolver. *Comput. and Structures* 43, 1183–1186 (1992)
5. Cullum, J.K., Willoughby, R.A.: A practical procedure for computing eigenvalues of large sparse nonsymmetric matrices. In: Cullum, J.K., Willoughby, R.A. (eds.) *Proceedings of the IBM Europe Institute Workshop on Large Scale Eigenvalue Problems*, pp. 193–223. North-Holland, Amsterdam (1986)
6. Arbenz, P., Hochstenbach, M.E.: A Jacobi–Davidson method for solving complex symmetric eigenvalue problems. *SIAM J. Sci. Comput.* 25(5), 1655–1673 (2004)
7. Luk, F., Qiao, S.: Using complex-orthogonal transformations to diagonalize a complex symmetric matrix. In: Luk, F.T. (ed.) *Advanced Signal Processing: Algorithms, Architectures, and Implementations VII, Proc. SPIE.*, vol. 162, pp. 418–425 (1997)
8. Cullum, J.K., Willoughby, R.A.: A *QL* procedure for computing the eigenvalues of complex symmetric tridiagonal matrices. *SIAM J. Matrix Anal. Appl.* 17, 83–109 (1996)
9. Sudbo, A.S.: Film mode matching: A versatile numerical method for vector mode field calculations in dielectric waveguides. *Pure and Appl. Optics* 2, 211–233 (1993)
10. Franza, O.P., Chew, W.C.: Recursive mode matching method for multiple waveguide junction modeling. *IEEE Trans. Microwave Theory Tech.* 44, 87–92 (1996)
11. Bienstman, P., Baets, R.: Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers. *Optical and Quantum Electronics* 33, 327–341 (2001)
12. Finger, N., Pacher, C., Boxleitner, W.: Simulation of Guided-Wave Photonic Devices with Variational Mode-Matching, April 2007. *American Institute of Physics Conference Series*, vol. 893, pp. 1493–1494 (2007)
13. Teixeira, F.L., Chew, W.C.: General closed-form PML constitutive tensors to match arbitrary bianisotropic and dispersive linear media. *IEEE Microwave Guided Wave Lett.* 8, 223–225 (1998)
14. Anderson, E., Bai, Z., Bischof, C.H., Blackford, S., Demmel, J.W., Dongarra, J.J., Du, C.J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D.C.: *Lapack Users' Guide*, 3rd edn. SIAM Press, Philadelphia (1999)
15. Cullum, J.K., Willoughby, R.A.: *Lanczos Algorithms for Large Symmetric Eigenvalue Computations*, vol. 1. Theory, vol. 2. Programs, Birkhäuser, Boston, MA (1985)
16. Bar-On, I., Paprzycki, M.: High performance solution of the complex symmetric eigenproblem. *Numerical Algorithms* 18, 195–208 (1998)