

An Algorithm for the Complete Solution of Quadratic Eigenvalue Problems

SVEN HAMMARLING, Numerical Algorithms Group Ltd. and The University of Manchester
CHRISTOPHER J. MUNRO, Rutherford Appleton Laboratory
FRANÇOISE TISSEUR, The University of Manchester

We develop a new algorithm for the computation of all the eigenvalues and optionally the right and left eigenvectors of dense quadratic matrix polynomials. It incorporates scaling of the problem parameters prior to the computation of eigenvalues, a choice of linearization with favorable conditioning and backward stability properties, and a preprocessing step that reveals and deflates the zero and infinite eigenvalues contributed by singular leading and trailing matrix coefficients. The algorithm is backward-stable for quadratics that are not too heavily damped. Numerical experiments show that our MATLAB implementation of the algorithm, `quadeig`, outperforms the MATLAB function `polyeig` in terms of both stability and efficiency.

Categories and Subject Descriptors: F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems—Computations on matrices; G.1 [Mathematics of Computing]: Numerical Analysis; G.1.3. [Numerical Analysis]: Numerical Linear Algebra—Eigenvalues and eigenvectors (direct and iterative methods)

General Terms: Algorithms, Performance, Reliability

Additional Key Words and Phrases: Quadratic eigenvalue problem, deflation, linearization, companion form, backward error, condition number, scaling, eigenvector

ACM Reference Format:

Hammarling, S., Munro, C. J., and Tisseur, F. 2013. An algorithm for the complete solution of quadratic eigenvalue problems. *ACM Trans. Math. Softw.* 39, 3, Article 18 (April 2013), 19 pages.
DOI: <http://dx.doi.org/10.1145/2450153.2450156>

1. INTRODUCTION

Eigensolvers for quadratic eigenvalue problems (QEPs) $Q(\lambda)x = 0$, $y^*Q(\lambda) = 0$, where

$$Q(\lambda) = \lambda^2 A_2 + \lambda A_1 + A_0, \quad A_i \in \mathbb{C}^{n \times n}$$

are often absent from the numerical linear algebra component of software libraries. Users are therefore left to choose a $2n \times 2n$ linearization $L(\lambda) = A - \lambda B$ of $Q(\lambda)$,

The work of C. J. Munro was supported by a scholarship from the School of Mathematics at the University of Manchester. The work of F. Tisseur was supported by Engineering and Physical Sciences Research Council Grant EP/1005293 and a Fellowship from the Leverhulme Trust.

Authors' addresses: S. Hammarling, Numerical Algorithms Group Ltd., Oxford, UK and School of Mathematics, The University of Manchester, Manchester, M13 9PL, UK; email: sven@nag.co.uk; C. J. Munro, Scientific Computing Department, Rutherford Appleton Laboratory, Harwell Oxford, Oxfordshire, OX11 0QX, UK; email: chris.j.munro@roscopmail.com; F. Tisseur, The University of Manchester, Manchester, M13 9PL, UK; email: ftisseur@ma.man.ac.uk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 0098-3500/2013/04-ART18 \$15.00
DOI: <http://dx.doi.org/10.1145/2450153.2450156>

for example,

$$\mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} A_0 & 0 \\ 0 & I \end{bmatrix} - \lambda \begin{bmatrix} -A_1 & -A_2 \\ I & 0 \end{bmatrix}, \quad (1)$$

call an eigensolver for generalized eigenproblems, and finally recover the eigenvectors of $Q(\lambda)$ from those of the linearized problem $L(\lambda)z = 0$, $w^*L(\lambda) = 0$. However, in doing so it is important to understand the influence of the linearization process on the accuracy and stability of the computed solution. Indeed, solving the QEP by applying a backward stable algorithm (e.g., the QZ algorithm [Moler and Stewart 1973]) to a linearization can be backward unstable for the QEP [Tisseur 2000]. Also, unless the block structure of the linearization is respected (and it is not by standard techniques), the conditioning of the solutions of the larger linear problem can be worse than that for the original quadratic, since the class of admissible perturbations is larger. For example, eigenvalues that are well-conditioned for $Q(\lambda)$ may be ill-conditioned for $L(\lambda)$ [Higham et al. 2006; Higham et al. 2008]. For these reasons, a numerical algorithm for solving QEPs needs to be carefully designed.

QEPs arise in a wide variety of applications in science and engineering such as in the dynamic analysis of mechanical systems, where the eigenvalues represent vibrational frequencies. Many practical examples of QEPs can be found in the NLEVP collection [Betcke et al. 2013] and the survey article by Tisseur and Meerbergen [2001]. In applications, the leading coefficient A_2 and/or the trailing coefficient A_0 can be singular. Regular quadratics (i.e., $\det Q(\lambda) \neq 0$) with singular A_0 and/or A_2 have zero and/or infinite eigenvalues. Theoretically, the QZ algorithm handles infinite eigenvalues well [Watkins 2000]. However, experiments by Kågström and Kressner [2006] show that if infinite eigenvalues are not extracted before starting the QZ steps, then they may never be detected due to the effect of rounding errors in floating point arithmetic. In one quadratic eigenvalue problem occurring in the vibration analysis of rail tracks under excitation arising from high-speed trains [Hilliges et al. 2004; Mackey et al. 2006a, p.18], the deflation of zero and infinite eigenvalues had a significant impact on the quality of the remaining computed finite eigenvalues.

In this work we present an eigensolver for the complete solution of dense QEPs that incorporates key recent developments on the numerical solution of polynomial eigenvalue problems, namely a scaling of the problem parameters prior to the computation [Fan et al. 2004; Gaubert and Sharify 2009], a choice of linearization with favorable conditioning and backward stability properties [Higham et al. 2006, 2007, 2008; Tisseur 2000], and a preprocessing step that reveals and deflates the zero and infinite eigenvalues contributed by singular leading and trailing matrix coefficients. The preprocessing step may also detect nonregularity, but that is not guaranteed. Our algorithm takes advantage of the block structure of the chosen linearization. We have implemented it as a MATLAB function called `quadeig`, which can make use of functions from the NAG Toolbox for MATLAB [NAG]. Our eigensolver can in principle be extended to matrix polynomials of degree higher than two but such an extension is not straightforward and requires further analysis, which is left to future work.

The remaining sections cover the following topics. In Section 2 we discuss the influence of backward error and condition number on the choice of linearization and give the motivation for our particular choice of linearization. In Section 3 we discuss scaling of the QEP, concentrating particularly on eigenvalue parameter scaling, and make some remarks on heavily damped problems. Section 4 looks at how to preprocess the QEP in order to deflate zero and infinite eigenvalues, and how to take advantage of the preprocessing in order to block-triangularize our choice of linearization. Section 5 discusses the recovery of the left and right eigenvectors of the QEP from those of the linearization. Section 6 presents numerical experiments, particularly using problems

from the NLEVP collection [Betcke et al. 2013], which demonstrate how our MATLAB code `quadeig` outperforms `polyeig` on QEPs.

2. CHOICE OF LINEARIZATIONS

Formally, $\mathcal{L}(\lambda) = \mathcal{A} - \lambda\mathcal{B}$ is a linearization of $\mathcal{Q}(\lambda)$ if there exist unimodular matrix polynomials $E(\lambda)$ and $F(\lambda)$ (that is, $\det E(\lambda)$ and $\det F(\lambda)$ are nonzero constants) such that

$$E(\lambda)\mathcal{L}(\lambda)F(\lambda) = \begin{bmatrix} \mathcal{Q}(\lambda) & 0 \\ 0 & I_n \end{bmatrix}.$$

Hence $\det(\mathcal{A} - \lambda\mathcal{B})$ agrees with $\det\mathcal{Q}(\lambda)$ up to a nonzero constant multiplier, so \mathcal{L} and \mathcal{Q} have the same eigenvalues. Research on linearizations of matrix polynomials has been very active in recent years, with achievements including generalization of the definition of linearization [Lancaster and Psarrakos 2005; Lancaster 2008], derivation of new (structured) linearizations [Amiraslani et al. 2009; Antoniou and Vologiannidis 2004, 2006; Higham et al. 2006; Mackey et al. 2006b, 2006c] and analysis of the influence of the linearization process on the accuracy and stability of computed solutions [Higham et al. 2006, 2007, 2008].

For a given quadratic \mathcal{Q} , there are infinitely many linearizations (the pencil (1) is just one example). These linearizations can have widely varying eigenvalue condition numbers [Higham et al. 2006], and approximate eigenpairs of $\mathcal{Q}(\lambda)$ computed via linearization can have widely varying backward errors [Higham et al. 2007]. In the following sections we define the terms backward error and condition number more precisely and focus on a particular linearization that our algorithm will employ.

2.1. Backward Error and Condition Number

In order to define backward errors and condition numbers valid for all λ , including ∞ , we rewrite $\mathcal{Q}(\lambda)$ and $\mathcal{L}(\lambda)$ in the homogeneous form

$$\mathcal{Q}(\alpha, \beta) = \alpha^2 A_2 + \alpha\beta A_1 + \beta^2 A_0, \quad \mathcal{L}(\alpha, \beta) = \beta A - \alpha B,$$

and identify λ with any pair $(\alpha, \beta) \neq (0, 0)$ for which $\lambda = \alpha/\beta$.

The normwise backward error of an approximate (right) eigenpair (x, α, β) of $\mathcal{Q}(\alpha, \beta)$ is defined by

$$\eta_{\mathcal{Q}}(x, \alpha, \beta) = \min\{\epsilon : (\mathcal{Q}(\alpha, \beta) + \Delta\mathcal{Q}(\alpha, \beta))x = 0, \|\Delta A_i\|_2 \leq \epsilon\|A_i\|_2, i = 0:2\}, \quad (2)$$

where $\Delta\mathcal{Q}(\alpha, \beta) = \alpha^2\Delta A_2 + \alpha\beta\Delta A_1 + \beta^2\Delta A_0$. The backward error $\eta_{\mathcal{L}}(z, \alpha, \beta)$ for an approximate eigenpair (z, α, β) of the pencil $\mathcal{L}(\alpha, \beta) = \beta A - \alpha B$ is defined in a similar way. Explicit expressions are given by Higham et al. [2007]:

$$\eta_{\mathcal{Q}}(x, \alpha, \beta) = \frac{\|\mathcal{Q}(\alpha, \beta)x\|_2}{\left(\sum_{i=0}^2 |\alpha|^i |\beta|^{2-i} \|A_i\|_2\right) \|x\|_2}, \quad \eta_{\mathcal{L}}(z, \alpha, \beta) = \frac{\|\mathcal{L}(\alpha, \beta)z\|_2}{(|\beta|\|A\|_2 + |\alpha|\|B\|_2) \|z\|_2}. \quad (3)$$

The definitions and explicit expressions for the backward error $\eta_{\mathcal{Q}}(y^*, \alpha, \beta)$ and $\eta_{\mathcal{L}}(w^*, \alpha, \beta)$ of left approximate eigenpairs (y^*, α, β) and (w^*, α, β) of \mathcal{Q} and \mathcal{L} are analogous to those for right eigenpairs.

A normwise condition number $\kappa_{\mathcal{Q}}(\alpha, \beta)$ of a simple eigenvalue (α, β) of \mathcal{Q} can be defined as

$$\kappa_{\mathcal{Q}}(\alpha, \beta) = \max_{\|\Delta A\| \leq 1} \frac{\|K(\alpha, \beta)\Delta A\|_2}{\|[\alpha, \beta]\|_2},$$

where $K(\alpha, \beta) : (\mathbb{C}^{n \times n})^3 \rightarrow T_{(\alpha, \beta)}\mathbb{P}_1$ is the differential of the map from (A_0, A_1, A_2) to (α, β) in projective space, and $T_{(\alpha, \beta)}\mathbb{P}_1$ denotes the tangent space at (α, β) to \mathbb{P}_1 , the

projective space of lines through the origin in \mathbb{C}^2 . Here $\Delta A = (\Delta A_0, \Delta A_1, \Delta A_2)$ and $\|\Delta A\| = \|[\omega_0^{-1}\Delta A_0, \omega_1^{-1}\Delta A_1, \omega_2^{-1}\Delta A_2]\|_F$ with $\omega_i = \|A_i\|_2$. An extension of a result of Dedieu and Tisseur [2003, Thm. 4.2] that treats the unweighted Frobenius norm yields the explicit formula [Higham et al. 2006]:

$$\kappa_{\mathcal{Q}}(\alpha, \beta) = \left(\sum_{i=0}^2 |\alpha|^{2i} |\beta|^{2(2-i)} \|A_i\|_2^2 \right)^{1/2} \frac{\|y\|_2 \|x\|_2}{|y^*(\bar{\beta}D_\alpha \mathcal{Q} - \bar{\alpha}D_\beta \mathcal{Q})|_{(\alpha, \beta)x}}, \quad (4)$$

where $D_\alpha \equiv \frac{\partial}{\partial \alpha}$ and $D_\beta \equiv \frac{\partial}{\partial \beta}$, and x, y are right and left eigenvectors of \mathcal{Q} associated with (α, β) . The eigenvalue condition number $\kappa_{\mathcal{L}}(\alpha, \beta)$ for the pencil $\mathcal{L}(\alpha, \beta) = \beta\mathcal{A} - \alpha\mathcal{B}$ is defined in a similar way, and an explicit formula is given by

$$\kappa_{\mathcal{L}}(\alpha, \beta) = \sqrt{|\beta|^2 \|\mathcal{A}\|_2^2 + |\alpha|^2 \|\mathcal{B}\|_2^2} \frac{\|w\|_2 \|z\|_2}{|w^*(\bar{\beta}D_\alpha \mathcal{L} - \bar{\alpha}D_\beta \mathcal{L})|_{(\alpha, \beta)z}}, \quad (5)$$

where z, w are right and left eigenvectors of \mathcal{L} associated with (α, β) . Note that the denominators of the expressions (4) and (5) are nonzero for simple eigenvalues. Also, these expressions are independent of the choice of representative of (α, β) and of the scaling of the eigenvectors. Let (α, β) and $(\tilde{\alpha}, \tilde{\beta})$ be the original and perturbed simple eigenvalues, normalized such that $\|(\alpha, \beta)\|_2 = 1$ and $[\alpha, \beta][\tilde{\alpha}, \tilde{\beta}]^* = 1$. Then the angle between the original and perturbed eigenvalues satisfies [Higham et al. 2007] the following:

$$|\theta((\alpha, \beta), (\tilde{\alpha}, \tilde{\beta}))| \leq \kappa_{\mathcal{Q}}(\alpha, \beta) \|\Delta A\| + o(\|\Delta A\|). \quad (6)$$

Backward error and conditioning are complementary concepts. The product of the condition number (4) with one of the backward errors in (3) provides an approximate upper bound on the angle between the original and computed eigenvalues. The eigenvalue condition numbers (4) and backward errors (3) are computed optionally by our algorithm.

2.2. Companion Linearizations

Companion linearizations are the most commonly used linearizations in practice. Several forms exist. Perhaps the best known are the first and second companion linearizations of \mathcal{Q} , given by

$$C_1(\lambda) = \begin{bmatrix} A_1 & A_0 \\ -I & 0 \end{bmatrix} - \lambda \begin{bmatrix} -A_2 & 0 \\ 0 & -I \end{bmatrix}, \quad C_2(\lambda) = \begin{bmatrix} A_1 & -I \\ A_0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} -A_2 & 0 \\ 0 & -I \end{bmatrix}. \quad (7)$$

Note that $C_2(\lambda)$ is the block transpose of $C_1(\lambda)$. Other companion forms such as

$$C_3(\lambda) = \begin{bmatrix} A_0 & 0 \\ 0 & I \end{bmatrix} - \lambda \begin{bmatrix} -A_1 & -A_2 \\ I & 0 \end{bmatrix}, \quad C_4(\lambda) = \begin{bmatrix} A_0 & 0 \\ 0 & I \end{bmatrix} - \lambda \begin{bmatrix} -A_1 & I \\ -A_2 & 0 \end{bmatrix},$$

are obtained, for example, by taking the reversal of the first or second companion form of the reversal of \mathcal{Q} ,

$$\text{rev}(\mathcal{Q}(\lambda)) := \lambda^2 A_0 + \lambda A_1 + A_2, \quad (8)$$

or simply by swapping the block rows or block columns of the above linearizations.

Companion linearizations have a number of desirable properties.

- (a) They are always linearizations even if $\mathcal{Q}(\lambda)$ is nonregular. Moreover, they are strong linearizations, that is, they preserve the partial multiplicities of infinite eigenvalues [Lancaster and Psarrakos 2005].
- (b) The left and right eigenvectors of $\mathcal{Q}(\lambda)$ are easily recovered from those of the companion form [Grammont et al. 2011; Higham et al. 2007] and (9) for C_2 .

- (c) If the quadratic is well scaled (i.e., $\|A_i\|_2 \approx 1$, $i = 0:2$), companion linearizations have good conditioning and backward stability properties (see Section 3).

Among companion linearizations $C_i(\lambda) = A_i - \lambda B_i$ we are looking for one for which:

- (d) the B_i matrix is in block upper triangular form, thereby reducing the computational cost of the Hessenberg-triangular reduction step of the QZ algorithm;
(e) the linearization can easily be transformed to a block upper triangular form revealing zero and infinite eigenvalues, if any.

The first and second companion linearizations in (7) satisfy desideratum (d), and we will show in Section 4 that, in the presence of singular leading and trailing matrix coefficients, desideratum (e) can easily be achieved for the second companion linearization. Hence our eigensolver will use $C_2(\lambda)$.

Concerning property (b), the second companion form $C_2(\lambda)$ in (7) has right eigenvectors z and left eigenvectors w of the form

$$z = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{cases} \begin{bmatrix} \alpha x \\ -\beta A_0 x \end{bmatrix} & \text{if } \alpha \neq 0, \\ \begin{bmatrix} \beta x \\ \alpha A_2 x + \beta A_1 x \end{bmatrix} & \text{if } \beta \neq 0, \end{cases} \quad w = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \bar{\alpha} y \\ \bar{\beta} y \end{bmatrix}, \quad (9)$$

where x, y are right and left eigenvectors of $Q(\lambda)$ with eigenvalue $\lambda = \alpha/\beta$. The formulas in (9) show that x can be recovered from the first n entries of z or by solving $A_0 x = z_2$ when A_0 is nonsingular, whereas y can be recovered from either the n first entries or the last n entries of w .

3. SCALINGS

Ideally, we would like the linearization C_2 that we use to be as well-conditioned as the original quadratic Q and for it to lead, after recovering an approximate eigenpair of Q from one of C_2 , to a backward error of the same order of magnitude as that for C_2 . Following the methodology described in Grammont et al. [2011] and Higham et al. [2007, 2006], we find that, with w and z denoting approximate left and right eigenvectors of $C_2(\lambda)$,

$$\frac{1}{\sqrt{2}} \frac{1}{\|p(\alpha, \beta)\|_2} \leq \frac{\kappa_{C_2}(\alpha, \beta)}{\kappa_Q(\alpha, \beta)} \leq 2^{3/2} \frac{\max(1, \max_{i=0:2} \|A_i\|_2)^2}{\|p(\alpha, \beta)\|_2}, \quad (10)$$

$$\frac{1}{\sqrt{2}} \leq \frac{\eta_Q(z_1, \alpha, \beta)}{\eta_{C_2}(z, \alpha, \beta)} \leq 2^{3/2} \frac{\max(1, \max_{i=0:2} \|A_i\|_2)}{\|p(\alpha, \beta)\|_1} \frac{\|z\|_2}{\|z_1\|_2}, \quad (11)$$

$$\frac{1}{\sqrt{2}} \leq \frac{\eta_Q(w_k^*, \alpha, \beta)}{\eta_{C_2}(w^*, \alpha, \beta)} \leq 2^{5/2} \frac{\max(1, \max_{i=0:2} \|A_i\|_2)^2}{\|p(\alpha, \beta)\|_1} \frac{\|w\|_2}{\|w_k\|_2}, \quad k = 1, 2, \quad (12)$$

where

$$p(\alpha, \beta) = [|\alpha|^2 \|A_2\|_2 \quad |\alpha|\|\beta\| \|A_1\|_2 \quad |\beta|^2 \|A_0\|_2]^T$$

and the eigenvalue (α, β) has been normalized so that $|\alpha|^2 + |\beta|^2 = 1$. The bounds (10)–(12) reveal that if $\|A_i\|_2 \approx 1$, $i = 0:2$, then

$$\kappa_Q(\alpha, \beta) \approx \kappa_{C_2}(\alpha, \beta), \quad \eta_Q(x, \alpha, \beta) \approx \eta_{C_2}(z, \alpha, \beta), \quad \eta_Q(y^*, \alpha, \beta) \approx \eta_{C_2}(w^*, \alpha, \beta)$$

for all eigenvalues (α, β) . Indeed, under that condition, $\|p(\alpha, \beta)\|_1 \approx 1$ and for an exact right eigenvector z the ratio $\|z\|_2/\|z_1\|_2$ is bounded by about $2\sqrt{2}$; this can be seen from the first equation in (9) if $|\alpha| \geq |\beta|$ and the second if $|\alpha| \leq |\beta|$. Also, for an exact left

eigenvector w , (9) shows that $\|w\|_2/\|w_k\|_2 \in [1, \sqrt{2}]$ by taking $k = 1$ if $|\alpha| \geq |\beta|$ and $k = 2$ if $|\alpha| \leq |\beta|$. We assume that these bounds also hold for the approximate eigenvectors.

When the coefficient matrices A_i , $i = 0:2$ vary largely in norm, numerical experiments in Higham et al. [2007] and Higham et al. [2006] show that $\kappa_Q \ll \kappa_{C_2}$ and $\eta_Q \gg \eta_{C_2}$ can happen, affecting the quality of computed solutions (see for example the beam problem in Higham et al. [2008]). We now discuss scaling strategies to overcome this issue.

3.1. Eigenvalue Parameter Scaling

An eigenvalue parameter scaling converts $Q(\lambda) = \lambda^2 A_2 + \lambda A_1 + A_0$ to $\tilde{Q}(\mu) = \mu^2 \tilde{A}_2 + \mu \tilde{A}_1 + \tilde{A}_0$, and is dependent on two nonzero scalar parameters γ and δ , where

$$\lambda = \gamma\mu, \quad Q(\lambda)\delta = \mu^2(\gamma^2\delta A_2) + \mu(\gamma\delta A_1) + \delta A_0 \equiv \tilde{Q}(\mu). \quad (13)$$

It has the important property that

$$\eta_Q(x, \alpha, \beta) = \eta_{\tilde{Q}}(x, \tilde{\alpha}, \tilde{\beta}),$$

where $\mu = \tilde{\alpha}/\tilde{\beta}$, so this scaling has no effect on the backward error for the quadratic, but it does affect the backward error for the linearization. In general, $\kappa_Q(\alpha, \beta) \neq \kappa_Q(\tilde{\alpha}, \tilde{\beta})$, but the eigenvalue parameter scaling leaves unchanged the λ eigenvalue condition number

$$\begin{aligned} \kappa_Q(\lambda) &= \limsup_{\epsilon \rightarrow 0} \left\{ \frac{|\Delta\lambda|}{\epsilon|\lambda|} : (Q(\lambda + \Delta\lambda) + \Delta Q(\lambda + \Delta\lambda))(x + \Delta x) = 0, \right. \\ &\quad \left. \|\Delta A_i\|_2 \leq \epsilon\omega_i, \quad i = 0:2 \right\} \\ &= \frac{(|\lambda|^2\|A_2\|_2 + |\lambda|\|A_1\|_2 + \|A_0\|_2)\|x\|_2\|y\|_2}{|\lambda|\|y^*(2\lambda A_1 + A_2)x|}, \end{aligned}$$

(see Tisseur [2000] and Higham et al. [2006]).

3.1.1. Fan, Lin, and Van Dooren Scaling. Fan et al. [2004] showed that when A_0 and A_2 are nonzero,

$$\gamma = \sqrt{\|A_0\|_2/\|A_2\|_2} =: \gamma_{FLV}, \quad \delta = 2/(\|A_0\|_2 + \|A_1\|_2\gamma) =: \delta_{FLV} \quad (14)$$

solves the problem of minimizing the maximum distance of the coefficient matrix norms from 1:

$$\min_{\gamma, \delta} \max\{\|\tilde{A}_0\|_2 - 1, \|\tilde{A}_1\|_2 - 1, \|\tilde{A}_2\|_2 - 1\}.$$

With this choice of parameters, it is shown in Higham et al. [2007] that

$$\max_{i=0:2} \|\tilde{A}_i\|_2 \leq 2, \quad \frac{1}{2} \leq \|p(\tilde{\alpha}, \tilde{\beta})\|_2^{-1} \leq \frac{\sqrt{3}}{2} \min \left\{ 1 + \tau_Q, \frac{1}{|\tilde{\alpha}\tilde{\beta}|} \right\}, \quad (15)$$

where

$$\tau_Q = \frac{\|A_1\|_2}{(\|A_2\|_2\|A_0\|_2)^{1/2}}. \quad (16)$$

—When $\tau_Q \lesssim 1$, which in the terminology of damped mechanical systems means that the problem is not too heavily damped, then on using (15) the bounds (10)–(12) imply that $\kappa_{C_2} \approx \kappa_{\tilde{Q}}$ for all eigenvalues, and $\eta_{\tilde{Q}} \approx \eta_{C_2}$ for both left and right eigenpairs. Hence if the eigenpairs of $C_2(\lambda)$ are computed with a small backward error (which is the case if we use the QZ algorithm), then we can recover eigenpairs for $Q(\lambda)$ with a small backward error.

—When $\tau_Q \gtrsim 1$, $\|p(\tilde{\alpha}, \tilde{\beta})\|_2^{-1}$ will still be of order 1 if $|\tilde{\alpha}||\tilde{\beta}| = |\tilde{\alpha}|\sqrt{1-|\tilde{\alpha}|^2} = O(1)$, which is the case unless $|\mu| = |\tilde{\alpha}|/|\tilde{\beta}| = |\tilde{\alpha}|/\sqrt{1-|\tilde{\alpha}|^2}$ is small or large.

3.1.2. Tropical Scaling. Gaubert and Sharify [2009] propose an eigenvalue parameter scaling based on the tropical roots of the max-times scalar quadratic polynomial (also called tropical polynomial)

$$q_{\text{trop}}(x) = \max(\|A_2\|_2 x^2, \|A_1\|_2 x, \|A_0\|_2), \quad x \in [0, \infty).$$

This polynomial has a double tropical root¹

$$\gamma_{\text{trop}}^+ = \gamma_{\text{trop}}^- = \sqrt{\|A_0\|_2 / \|A_2\|_2} = \gamma_{FLV}$$

when $\tau_Q \leq 1$ and two distinct tropical roots

$$\gamma_{\text{trop}}^+ = \frac{\|A_1\|_2}{\|A_2\|_2}, \quad \gamma_{\text{trop}}^- = \frac{\|A_0\|_2}{\|A_1\|_2}, \quad (\gamma_{\text{trop}}^+ > \gamma_{\text{trop}}^-)$$

when $\tau_Q > 1$.

Gaubert and Sharify prove that when the tropical roots are well-separated and A_2, A_1 are well-conditioned, then the n largest eigenvalues in modulus are of the order of γ_{trop}^+ . Similarly, if $\gamma_{\text{trop}}^+ \gg \gamma_{\text{trop}}^-$ and A_1, A_0 are well-conditioned, then the n smallest eigenvalues in modulus are of the order of γ_{trop}^- . They show experimentally that if $Q(\lambda)$ is scaled as in (13) with

$$\gamma = \gamma_{\text{trop}}^+, \quad \delta = (q_{\text{trop}}(\gamma_{\text{trop}}^+))^{-1}, \quad (17)$$

then the large eigenvalues in magnitude of $\tilde{Q}(\mu)$ are computed with a small backward error by the QZ algorithm (via first companion linearization). Similarly, the choice of parameters

$$\gamma = \gamma_{\text{trop}}^-, \quad \delta = (q_{\text{trop}}(\gamma_{\text{trop}}^-))^{-1} \quad (18)$$

experimentally yields small backward errors for eigenvalues of small magnitude. This behavior is confirmed theoretically by our bounds (10)–(12). Indeed, for the choices (17) and (18) of parameters, $\max(1, \max_{i=0:2} \|A_i\|_2)^2 = 1$ and

$$\|p(\tilde{\alpha}, \tilde{\beta})\|_1 = O(1) \text{ if } \begin{cases} \gamma = \gamma_{\text{trop}}^+ \text{ and } |\mu| \geq 1, & \text{or equivalently, } |\lambda| \geq \gamma_{\text{trop}}^+, \\ \gamma = \gamma_{\text{trop}}^- \text{ and } |\mu| \leq 1, & \text{or equivalently, } |\lambda| \leq \gamma_{\text{trop}}^-. \end{cases}$$

Hence, when $\tau_Q > 1$, tropical scaling guarantees optimal backward errors and conditioning for the eigenvalues of the scaled quadratic outside the unit circle when (17) is used and for those in the unit circle when (18) is used.

3.2. Further Remarks Concerning Heavily Damped Problems

The analysis in Section 3.1 implies that small backward errors cannot be guaranteed for all eigenvalues when $\tau_Q \gtrsim 1$. Hence we may want to transform the problem to an equivalent one for which $\tau_Q \approx 1$. To that effect, we can

(i) use a homogeneous rotation

$$\begin{bmatrix} c & s \\ -s & c \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} =: \begin{bmatrix} \tilde{\alpha} \\ \tilde{\beta} \end{bmatrix}, \quad c, s \in \mathbb{R}, \quad c^2 + s^2 = 1, \quad (19)$$

¹At a tropical root, the maximum is attained by at least two monomials.

and define a rotated quadratic $\tilde{Q}(\tilde{\alpha}, \tilde{\beta})$ via

$$Q(\alpha, \beta) = \sum_{j=0}^2 (c\tilde{\alpha} - s\tilde{\beta})^j (s\tilde{\alpha} + c\tilde{\beta})^{2-j} A_j = \sum_{j=0}^2 \tilde{\alpha}^j \tilde{\beta}^{2-j} \tilde{A}_j =: \tilde{Q}(\tilde{\alpha}, \tilde{\beta}) \quad (20)$$

with

$$\tilde{A}_2 = Q(c, s), \quad \tilde{A}_1 = -2csA_2 + (c^2 - s^2)A_1 + 2csA_0, \quad \tilde{A}_0 = Q(-s, c).$$

It is not difficult to find c, s such that $\|\tilde{A}_1\|_2 \lesssim (\|\tilde{A}_2\|_2 \|\tilde{A}_0\|_2)^{1/2}$ (for instance, $c = s = 1/\sqrt{2}$ usually achieves this inequality);

- (ii) use a diagonal scaling or balancing of the type discussed in Lemonnier and Van Dooren [2006] and Betcke [2008].

However, in both cases, examples exist for which after scaling back, $\eta_Q \gg nu$, so we do not pursue this idea.

3.3. Summary

Based on the analysis in Section 3.1 and remarks in Section 3.2, the default in our implementation is to apply the Fan, Lin, and Van Dooren scaling (14) to the quadratic $Q(\lambda)$ prior to building the second companion linearization $C_2(\lambda)$ when $\tau_Q \lesssim 1$, and not to scale otherwise. The tropical scaling (17), (18) is also implemented, but is left as an option. We refer to Experiment 3 in Section 6 for numerical illustrations of these eigenvalue parameter scalings.

4. DEFLATION OF 0 AND ∞ EIGENVALUES

The eigenvalues of a regular $n \times n$ quadratic $Q(\lambda)$ are the zeros of the characteristic polynomial $\det Q(\lambda) = \det A_2 \lambda^{2n} + \text{lower-order terms}$, so when A_2 is nonsingular, $Q(\lambda)$ has $2n$ finite eigenvalues. When A_2 is singular, $Q(\lambda)$ has d finite eigenvalues to which we add $2n - d$ infinite eigenvalues, where d is the degree of $\det Q(\lambda)$. Note that λ is an eigenvalue of Q if and only if $1/\lambda$ is an eigenvalue of the reversal of Q in (8), where 0 and ∞ are regarded as reciprocals. If $r_0 = \text{rank}(A_0) < n$, then Q has at least $n - r_0$ zero eigenvalues and if $r_2 = \text{rank}(A_2) < n$, Q has at least $n - r_2$ infinite eigenvalues. As an example, the quadratic

$$Q(\lambda) = \lambda^2 \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} + \lambda A_1 + \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

with A_1 such that $\det Q(\lambda) \not\equiv 0$ has at least one infinite eigenvalue and at least one zero eigenvalue. If $A_1 = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$, then the remaining eigenvalues are ∞ and -1 .

Let us denote by

$$\mathcal{N}_r(A) = \{x \in \mathbb{C}^n : Ax = 0\}, \quad \mathcal{N}_l(A) = \{y \in \mathbb{C}^n : y^*A = 0\}$$

the right and left nullspaces, respectively, of $A \in \mathbb{C}^{n \times n}$. Note that the right and left eigenvectors of Q , associated with the 0 and ∞ eigenvalues, generate the right and left nullspace of A_0 and A_2 , respectively.

Our algorithm checks the ranks of A_0 and A_2 and when one or both of them are singular, it deflates the corresponding zero and infinite eigenvalues.

4.1. Rank and Nullspace Determination

A QR factorization with column pivoting can be used to determine the rank of an $n \times n$ matrix A . This factorization has the form

$$Q^*AP = \begin{matrix} & k & n-k \\ & \begin{bmatrix} R_{11} & R_{12} \\ 0 & 0 \end{bmatrix} \end{matrix}, \quad (21)$$

where Q is unitary, P is a permutation matrix, and R_{11} is upper triangular and non-singular Golub and Van Loan [1996, p. 248]. Then $\text{rank}(A) = \text{rank}(R_{11}) = k$. For sufficiently small $\|E\|_2$, it is shown in Higham [1990, Thm. 5.2] that $A + E$ has the QR factorization with column-pivoting:

$$\bar{Q}^*(A + E)P = \begin{matrix} & k & n-k \\ & \begin{bmatrix} \bar{R}_{11} & \bar{R}_{12} \\ 0 & \bar{R}_{22} \end{bmatrix} \end{matrix}, \quad (22)$$

with

$$\frac{\|\bar{R}_{22}\|_2}{\|A\|_2} \leq \frac{\|E\|_2}{\|A\|_2} (1 + \|R_{11}^{-1}R_{12}\|_2) + O\left(\frac{\|E\|_2}{\|A\|_2}\right)^2, \quad (23)$$

and where $\|R_{11}^{-1}R_{12}\|_2 \leq ((n-k)(4^k - 1)/3)^{1/2}$. Although the latter bound is nearly attainable, numerical experience shows that $\|R_{11}^{-1}R_{12}\|_2$ is usually small. Hence if $A + E$ is close to a rank k matrix, then \bar{R}_{22} will usually be small. Our algorithm sets \bar{R}_{22} to zero if $\|\bar{R}_{22}\|_2 \leq \text{tol}$, where the tolerance tol can be specified by the user. By default, $\text{tol} = nu\|A\|_2$, where u is the unit round-off. With the default tolerance we can overestimate the rank, but this does not affect the stability of our algorithm. Indeed, we only use QR factorizations with column-pivoting to deflate zero and infinite eigenvalues. If the rank is overestimated, then we deflate fewer eigenvalues than we could have done. The QZ algorithm then has to solve a generalized eigenproblem of larger dimension than really necessary.

Although there are other more sophisticated methods for estimating the rank [Chandrasekaran and Ipsen 1994], QR with column-pivoting performs well in our algorithm in practice (see Section 6), and has the advantage of being available in LAPACK and as a MATLAB built-in function.

Note that the last $n-k$ columns of Q in (21) span the left null space of A . A basis for the right nullspace of A is obtained by postmultiplying (21) by a sequence of Householder transformations H_1, \dots, H_k that reduce R_{12} to zero. This leads to a complete orthogonal decomposition of A ,

$$Q^*AZ = \begin{matrix} & k & n-k \\ & \begin{bmatrix} T_{11} & 0 \\ 0 & 0 \end{bmatrix} \end{matrix}, \quad (24)$$

where $Z = PH_1 \cdots H_k$ (see Golub and Van Loan [1996, p. 250]). Then the last $n-k$ columns of Z span the right nullspace of A . The decomposition (24) will be needed in the following sections.

4.2. Block Triangularization of $C_2(\lambda)$

Throughout this section we assume that $r_0 := \text{rank}(A_0) \leq \text{rank}(A_2) =: r_2$ (if $r_0 > r_2$, we work with $\text{rev}(Q(\lambda))$ instead of $Q(\lambda)$). Let

$$Q_i^*A_iP_i = \begin{matrix} & r_i & n-r_i \\ & \begin{bmatrix} R_{11}^{(i)} & R_{12}^{(i)} \\ 0 & 0 \end{bmatrix} \end{matrix} = \begin{bmatrix} R^{(i)} \\ 0 \end{bmatrix}, \quad i = 0, 2, \quad (25)$$

be QR factorizations with column-pivoting of A_0 and A_2 . With the help of these factorizations and another complete orthogonal decomposition when both A_0 and A_2 are singular (i.e., $r_0, r_2 < n$), we show how to transform the second companion form

$$C_2(\lambda) = \begin{bmatrix} A_1 & -I \\ A_0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} -A_2 & 0 \\ 0 & -I \end{bmatrix}$$

into block upper triangular form

$$QC_2(\lambda)V = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ 0 & A_{22} & A_{23} \\ 0 & 0 & 0_{n-r_0} \end{bmatrix} - \lambda \begin{bmatrix} B_{11} & B_{12} & B_{13} \\ 0 & 0_{n-r_2} & B_{23} \\ 0 & 0 & B_{33} \end{bmatrix}, \quad (26)$$

where the $2n \times 2n$ matrices are partitioned conformably. Note that if A_{22} (or B_{33}) is singular, then $\det Q(\lambda) = \det C_2(\lambda) \equiv 0$, and hence $Q(\lambda)$ is nonregular. When A_{22} and B_{33} are nonsingular, (26) reveals $n - r_0$ zero eigenvalues and $n - r_2$ infinite eigenvalues. The remaining eigenvalues are those of the $(r_0 + r_2) \times (r_0 + r_2)$ pencil $A_{11} - \lambda B_{11}$.

We consider three cases.

- (i) $r_0 = r_2 = n$. In this case there are no zeroes and no infinite eigenvalues. However, we make use of the factorization of A_2 in (25) to reduce the leading coefficient $\begin{bmatrix} -A_2 & 0 \\ 0 & -I \end{bmatrix}$ of the linearization to upper triangular form, which is a necessary step in the QZ algorithm. This is achieved with

$$Q = \begin{bmatrix} Q_2^* & 0 \\ 0 & I_n \end{bmatrix}, \quad V = \begin{bmatrix} P_2 & 0 \\ 0 & I_n \end{bmatrix},$$

so that

$$QC_2(\lambda)V = \begin{bmatrix} Q_2^* A_1 P_2 & -Q_2^* \\ A_0 P_2 & 0 \end{bmatrix} - \lambda \begin{bmatrix} -R^{(2)} & 0 \\ 0 & -I \end{bmatrix} = A_{11} - \lambda B_{11}.$$

- (ii) $r_0 < r_2 = n$. In this case there are at least $n - r_0$ zero eigenvalues, which we deflate with

$$Q = \begin{bmatrix} Q_2^* & 0 \\ 0 & Q_0^* \end{bmatrix}, \quad V = \begin{bmatrix} P_2 & 0 \\ 0 & Q_0 \end{bmatrix},$$

so that

$$QC_2(\lambda)V = \begin{matrix} & n & r_0 & n-r_0 \\ & X_{11} & X_{12} & X_{13} \\ n & X_{21} & 0 & 0 \\ r_0 & X_{21} & 0 & 0 \\ n-r_0 & 0 & 0 & 0 \end{matrix} - \lambda \begin{bmatrix} -R^{(2)} & 0 & 0 \\ 0 & -I_{r_0} & 0 \\ 0 & 0 & -I_{n-r_0} \end{bmatrix}, \quad (27)$$

where $X_{11} = Q_2^* A_1 P_2$, $[X_{12} \ X_{13}] = -Q_2^* Q_0$ and $X_{21} = R^{(0)} P_0^* P_2$. The pencil (27) is in the form (26) with $A_{11} = \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & 0 \end{bmatrix}$ and $B_{11} = \begin{bmatrix} -R^{(2)} & 0 \\ 0 & -I_{r_0} \end{bmatrix}$ of dimension $(n + r_0) \times (n + r_0)$. Note that, as in case (i), B_{11} is upper triangular.

- (iii) $r_0 \leq r_2 < n$. There are at least $n - r_0$ zero eigenvalues and at least $n - r_2$ infinite eigenvalues that we deflate as follows. With

$$\tilde{Q} = \begin{bmatrix} Q_2^* & 0 \\ 0 & Q_0^* \end{bmatrix}, \quad \tilde{V} = \begin{bmatrix} I_n & 0 \\ 0 & Q_0 \end{bmatrix}$$

we obtain

$$\tilde{Q}C_2(\lambda)\tilde{V} = \begin{matrix} & r_2 & n-r_2 & r_0 & n-r_0 \\ & X_{11} & X_{12} & X_{13} & X_{14} \\ n-r_2 & X_{21} & X_{22} & X_{23} & X_{24} \\ r_0 & X_{31} & X_{32} & 0 & 0 \\ n-r_0 & 0 & 0 & 0 & 0 \end{matrix} - \lambda \begin{bmatrix} Y_{11} & Y_{12} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -I_{r_0} & 0 \\ 0 & 0 & 0 & -I_{n-r_0} \end{bmatrix}, \quad (28)$$

where $\begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} = Q_2^* A_1$, $\begin{bmatrix} X_{13} & X_{14} \\ X_{23} & X_{24} \end{bmatrix} = -Q_2^* Q_0$, $[X_{31} \ X_{32}] = R^{(0)} P_0^*$, and $[Y_{11} \ Y_{12}] = -R^{(2)} P_2^*$. Let

$${}_{n-r_2} \begin{bmatrix} r_2 & n-r_2 & r_0 \\ X_{21} & X_{22} & X_{23} \end{bmatrix} = Q_3 \begin{bmatrix} n-r_2 & r_0+r_2 \\ R_3 & 0 \end{bmatrix} Z_3$$

be a complete orthogonal decomposition and let

$$Q = \begin{bmatrix} I_{r_2} & 0 & 0 & 0 \\ 0 & 0 & I_{r_0} & 0 \\ 0 & Q_3^* & 0 & 0 \\ 0 & 0 & 0 & I_{n-r_0} \end{bmatrix} \tilde{Q}, \quad V = \tilde{V} \begin{bmatrix} n+r_0 & n-r_0 \\ Z_3^* & 0 \\ 0 & I_{n-r_0} \end{bmatrix} \begin{bmatrix} 0 & I_{n-r_2} & 0 \\ I_{r_2+r_0} & 0 & 0 \\ 0 & 0 & I_{n-r_0} \end{bmatrix}.$$

Then easy calculations show that $QC_2(\lambda)V$ has the form (26) with $A_{22} = R_3$.

When A_0 or A_2 have low rank, the block triangularization (26) of $C_2(\lambda)$ yields a substantial saving in work for the eigenvalues/eigenvectors computation, as it reduces the size of the pencil to which the QZ algorithm is applied. In the worst case, that is when A_0 and A_2 are both nonsingular, we only make use of one of the two QR factorizations with column-pivoting. The execution times provided in Section 6 indicate that this extra cost is negligible compared with the overall cost.

5. LEFT AND RIGHT EIGENVECTORS

From (26) we have that

$$\Lambda(Q(\lambda)) = \Lambda(C_2(\lambda)) = \Lambda(A_{11} - \lambda B_{11}) \cup \underbrace{\{\infty, \dots, \infty\}}_{n-r_2 \text{ times}} \cup \underbrace{\{0, \dots, 0\}}_{n-r_0 \text{ times}},$$

where $\Lambda(A)$ denotes the spectrum of A . Our implementation uses the QZ algorithm to compute the Schur decomposition of $A_{11} - \lambda B_{11}$, and hence its eigenvalues.

The computation of the eigenvectors of $Q(\lambda)$ depends on whether the eigenvectors are right or left eigenvectors, whether they correspond to deflated eigenvalues or not, and whether the Fan, Lin, and Van Dooren scaling is used or not, as we now explain.

5.1. Right Eigenvectors

When A_0 or A_2 or both are singular, the vectors spanning their right nullspaces $\mathcal{N}_r(A_0)$ and $\mathcal{N}_r(A_2)$ are right eigenvectors associated with the 0 and ∞ eigenvalues of $Q(\lambda)$. These nullspaces can be obtained from (25) by zeroing $R_{12}^{(i)}$, $i = 0, 2$ with the help of r_j Householder reflectors to yield a complete orthogonal decomposition as in (24), that is,

$$Q_j^* A_j Z_j = \begin{bmatrix} r_j & n-r_j \\ T_{11}^j & 0 \\ 0 & 0 \end{bmatrix}, \quad j = 0, 2.$$

The last $n - r_0$ columns of Z_0 are eigenvectors of Q with eigenvalue 0 and the last $n - r_2$ columns of Z_2 are eigenvectors of Q with eigenvalue ∞ .

For the right eigenvectors associated with the nondeflated eigenvalues note that if $\tilde{z} \in \mathbb{C}^{r_0+r_2}$ is a right eigenvector of the pencil $A_{11} - \lambda B_{11}$, which is easily obtained once the Schur decomposition of $A_{11} - \lambda B_{11}$ is available, then

$$z = \begin{bmatrix} n \\ n \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = V \begin{bmatrix} \tilde{z} \\ 0 \end{bmatrix}$$

is a right eigenvector of $C_2(\lambda)$, where V is as in Section 4.2. We also know that z must have the form displayed in (9). To recover an approximate eigenvector x of Q from z , we proceed as follows.

- If the scaling (14) is used and $\tau_Q \lesssim 1$ with τ_Q as in (16), then $x = z_1$, since the analysis in Section 3.1.1 guarantees a small backward error for the eigenpair (z_1, α, β) .
- If A_0 is singular, then $x = z_1$.
- Otherwise $\alpha \neq 0$, and if $\beta = 0$, then $x = z_1$. If $\beta \neq 0$, we use the QR factorization of A_0 in (25) to solve the linear system $z_2 = -\beta A_0 x_2$ for x_2 . Then $x = z_1$, if $\eta_Q(z_1, \alpha, \beta) \leq \eta_Q(x_2, \alpha, \beta)$, otherwise, $x = x_2$.

5.2. Left Eigenvectors

When A_0 is singular, the last $n - r_0$ columns of Q_0 in (25) are eigenvectors of Q associated with the $n - r_0$ deflated zero eigenvalues and when A_2 is singular, the last $n - r_2$ columns of Q_2 in (25) are left eigenvectors of Q associated with the deflated $n - r_2$ infinite eigenvalues.

For the left eigenvectors of Q corresponding to the nondeflated eigenvalues, we first compute the left eigenvectors of $C_2(\lambda)$ corresponding to these eigenvalues. For this, we use the Schur decomposition of $A_{11} - \lambda B_{11}$ to obtain the Schur decomposition of $C_2(\lambda)$, from which the left eigenvectors are easily computed. In exact arithmetic, if w is a left eigenvector of C_2 , then it has the form displayed in (9), but in floating point arithmetic, w_1 and w_2 are generally not parallel,

$$w = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \bar{\alpha} y_1 \\ \bar{\beta} y_2 \end{bmatrix}.$$

Then, to recover an approximate left eigenvector y of Q from w we proceed as follows.

- If the scaling (14) is used and $\tau_Q \lesssim 1$, then $y = w_1$ if $|\lambda| = |\alpha/\beta| \geq 1$ and $y = w_2$ if $|\lambda| < 1$, since with this choice for y , the analysis in Section 3 guarantees a small backward error for the left eigenpair (y^*, α, β) .
- Otherwise, we take as approximate left eigenvector of Q with eigenvalue $\lambda = \alpha/\beta$, the vector $w_i \neq 0$ with smallest backward error $\eta_Q(w_i^*, \alpha, \beta)$.

5.3. Further Remarks

As explained in the previous two sections, there are situations for which we have two choices for the eigenvectors. For example, suppose that (x_1, α, β) and (x_2, α, β) are two approximate right eigenpairs of a single eigenpair (x, α, β) of Q . Then we can try to find $a = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}$ such that $\eta_Q(a_1 x_1 + a_2 x_2, \alpha, \beta)$ is minimized. In other words, we need to solve the following problem:

$$\min_{a \in \mathbb{C}^2} \frac{\|Q(\alpha, \beta)Xa\|_2}{\|Xa\|_2},$$

where $X = [x_1 \ x_2]$. For that we can take the GSVD of the pair of $n \times 2$ matrices $(Q(\alpha, \beta)X, X)$,

$$Q(\alpha, \beta)X = UCY, \quad X = VSY,$$

where U, V are unitary, Y is nonsingular, and $C = \text{diag}(c_1, c_2)$, $S = \text{diag}(s_1, s_2)$ with $c_1, c_2, s_1, s_2 \geq 0$. Thus, if we let $a = Y^{-1}b$,

$$\min_{a \in \mathbb{C}^2} \frac{\|Q(\alpha, \beta)Xa\|_2^2}{\|Xa\|_2^2} = \min_{b \in \mathbb{C}^2} \frac{\|Cb\|_2^2}{\|Sb\|_2^2} = \min_{b \in \mathbb{C}^2} \frac{b^* C^* C b}{b^* S^* S b},$$

which is the smallest eigenvalue of $C^* C - \lambda S^* S$. So the minimum is achieved at $b = e_i$, where $|c_i/s_i|$ is minimal. Hence $a = Y^{-1}e_i$. In practice, we found that this approach does not decrease the backward error much. The matrix Y is often very close to being singular, producing a new eigenvector with a larger backward error than that for x_1 and x_2 . Therefore, we did not implement this approach.

ALGORITHM 1: Quadratic Eigenvalue Solver—`quadeig`

Input: $n \times n$ matrices A_2, A_1, A_0 defining the quadratic $Q(\lambda)$.

Optional: parameter `pscale` such that

$$\text{pscale} = \begin{cases} 0 : \text{Fan, Lin and Van Dooren scaling if } \tau_Q < 10 \text{ and} \\ \quad \text{no scaling otherwise (default).} \\ 1 : \text{no eigenvalue parameter scaling} \\ 2 : \text{Fan, Lin and Van Dooren scaling} \\ 3 : \text{tropical scaling with smallest tropical root } \gamma_{trop}^- \\ 4 : \text{tropical scaling with largest tropical root } \gamma_{trop}^+ \end{cases}$$

and a tolerance `tol` for rank decision (default: $\text{tol} = nu \max_i(\|A_i\|)$).

Output: Eigenvalues,
right/left eigenvectors (optional), see Section 5,
eigenvalue condition numbers (optional),
backward errors of computed right/left eigenpairs (optional).

- 1 Apply eigenvalue parameter scaling `pscale` as described in Section 3.
- 2 Compute the block triangularization (26) of $C_2(\lambda)$ using `tol` (see Section 4).
- 3 Compute eigenvalues of $Q(\lambda)$ using Schur decomposition of $A_{11} - \lambda B_{11}$.
- 4 Compute right/left eigenvectors of $Q(\lambda)$ as described in Section 5 (optional).
- 5 Compute eigenvalue condition numbers (optional). Compute backward errors of computed right/left eigenpairs (optional).

5.4. Algorithm

The main steps of our quadratic eigensolver are presented in Algorithm 1. They consist of two preprocessing steps: an eigenvalue parameter scaling and a deflation of the zero and infinite eigenvalues contributed by singular leading and trailing matrix coefficients, as described in Sections 3 and 4. These preprocessing steps are followed by the computation of the eigenvalues with the QZ algorithm. The eigenvector computation described in Section 5 is optional, as are the computation of the eigenvalue condition numbers and backward errors of approximate eigenpairs.

6. NUMERICAL EXPERIMENTS

We now describe a collection of numerical experiments designed to give insight into `quadeig`, our MATLAB implementation of Algorithm 1, its performance in floating-point, and the implementation issues. Our computations were done in MATLAB 7.13.0 (R2011b) under Windows XP (SP3) with a Pentium E6850, for which $u = 2^{-53} \approx 1.1 \times 10^{-16}$. When available, `quadeig` makes use of some functions from the NAG Toolbox for MATLAB, namely `f08bh`² to reduce the upper trapezoidal matrix $[R_{11} \ R_{12}]$ in (21) to upper triangular form $[R \ 0]$, and `f08yk`³ to compute some of the left generalized eigenvectors of $C_2(\lambda)$ as well as their complex analogs. The use or not of these LAPACK routines through the NAG Toolbox does not affect the numerical stability of our implementation of Algorithm 1, just its efficiency.

In all computations involving $\|A_i\|_2$, $i = 0, 2$, (e.g., Fan, Lin, and Van Dooren scaling, tropical scaling, condition numbers, backward errors), the 2-norm is replaced by the Frobenius norm $\|\cdot\|_F$.

We compare the performance of `quadeig` to that of the MATLAB function `polyeig`, which, for quadratics, calls the QZ algorithm on the linearization (1) `polyeig` returns

²`f08bh` corresponds to the LAPACK routine `DTZRZF`.

³`f08yk` corresponds to the LAPACK routine `DTGEVC`.

Table I. Largest Backward Errors of Eigenpairs (computed by `polyeig` and `quadeig` on quadratic eigenvalue problems from the NLEVP collection)

Problem	n	τ_Q	polyeig		quadeig	
			$\eta_Q^{\max}(x, \alpha, \beta)$	$\eta_Q^{\max}(x, \alpha, \beta)$	$\eta_Q^{\max}(y^*, \alpha, \beta)$	
acoustic_wave_1d	10	2.1e-1	2.3e-15	6.5e-16	5.5e-16	
acoustic_wave_2d	30	2.1e-1	5.6e-16	6.2e-16	6.4e-16	
bicycle	2	4.2e-1	1.5e-15	6.1e-17	5.2e-17	
bilby	5	4.2e+0	2.6e-16	6.0e-16	3.5e-16	
cd_player	60	9.3e+3	1.7e-10	7.4e-16	1.8e-15	
closed_loop	2	2.5e+0	1.8e-16	8.4e-16	1.5e-16	
damped_beam	200	2.1e-4	3.3e-9	9.9e-16	8.7e-16	
dirac	80	1.6e+0	4.1e-15	1.2e-15	1.6e-15	
gen_hyper2	15	1.4e+0	6.1e-16	5.5e-16	4.9e-16	
gen_tantipal2	16	1.5e+0	4.8e-16	4.7e-16	4.1e-16	
gen_tpal2	16	1.9e+0	6.9e-16	6.1e-16	6.9e-16	
intersection	10	4.9e-1	3.8e-17	4.7e-17	8.5e-17	
hospital	24	6.6e-2	3.0e-13	6.2e-16	6.2e-16	
metal_strip	9	7.8e+0	4.0e-14	6.4e-16	4.0e-16	
mobile_manipulator	5	8.3e-2	1.2e-18	6.2e-17	6.4e-17	
omnicam1	9	1.1e+0	1.7e-15	9.4e-17	3.0e-17	
omnicam2	15	1.9e+0	4.8e-17	6.6e-17	2.3e-16	
pdde_stability	225	4.4e+1	5.0e-14	1.5e-14	1.3e-14	
power_plant	8	6.7e-1	1.3e-8	3.8e-16	4.9e-17	
qep1	3	2.5e+0	2.2e-16	7.3e-17	6.2e-17	
qep2	3	1.3e+0	9.6e-17	8.7e-17	8.7e-17	
qep3	3	7.5e-1	9.4e-17	1.2e-16	5.1e-17	
qep5	3	1.2e+0	4.4e-16	2.8e-16	2.0e-16	
railtrack	1005	1.8e+1	2.0e-8	2.4e-15	9.6e-15	
relative_pose_6pt	10	5.8e+0	1.1e-15	5.2e-16	2.9e-16	
shaft	400	1.1e-6	8.5e-8	1.0e-15	9.6e-16	
sign1	81	2.0e+0	2.3e-16	9.4e-16	9.6e-16	
sign2	81	1.8e+0	5.8e-16	1.6e-15	1.0e-15	
sleeper	10	3.5e+0	3.7e-15	3.5e-16	2.8e-16	
speaker_box	107	2.1e-5	1.5e-11	2.2e-16	3.9e-16	
spring	5	8.1e+0	5.6e-16	5.6e-16	4.9e-16	
spring_dashpot	10	6.4e-3	1.0e-15	1.3e-16	1.2e-16	
wing	3	1.8e-1	4.5e-15	3.6e-16	4.1e-16	
wiresaw1	10	1.5e-2	9.2e-15	5.6e-16	5.6e-16	
wiresaw2	10	7.3e-2	4.8e-14	9.8e-16	9.6e-16	

eigenvalues, right eigenvectors, and eigenvalue condition numbers, but does not return left eigenvectors and backward errors of approximate eigenpairs.

Experiment 1. We ran `quadeig` with the default options (except for the calls to the NAG Toolbox, which were switched off) on all the QEPs from the NLEVP collection [Betcke et al. 2013] of size $n < 1250$ and tested its numerical stability. For each problem, Table I displays the largest backward error for the right eigenpairs returned by `polyeig` and the largest backward errors for the right and left eigenpairs returned by `quadeig`. We make the following observations.

- `polyeig` is unstable on several examples that are highlighted in bold in the table.
- As predicted by the analysis in Section 3, `quadeig`, unlike `polyeig`, returns right and left eigenpairs with backward errors close to the machine precision for quadratics $Q(\lambda)$ that are not too heavily damped (i.e., $\tau_Q \lesssim 1$).

Table II. Execution Time (in seconds for the computation of all eigenvalues Λ and all right eigenpairs (Λ, X) of $n \times n$ quadratics in the NLEVP collection, where $r_i = \text{rank}(A_i), i = 0, 2$)

Problem	n	r_0	r_2	polyeig		quadeig	
				Λ	(Λ, X)	Λ	(Λ, X)
acoustic_wave_2D	870			118.0	202.8	114.1	190.8
damped_beam	1000			95.6	155.7	98.3	166.9
spring	1000			181.7	272.1	93.8	170.2
shaft	400	400	199	1.7	5.4	1.4	2.7
railtrack	1005	67	67	17.3	70.0	4.8	6.3
railtrack2	1410	705	705	128.8	306.2	79.2	113.5

- The improvements in the backward errors for the damped_beam, hospital, power_plant, shaft, and speaker_box problems are due to the Fan, Lin, and Van Dooren scaling. Note that for the latter problem, the stiffness matrix A_0 has numerical rank $106 < n = 107$. The smallest eigenvalue is computed as 0.16 by polyeig and as 0 by quadeig's deflation process.
- For the railtrack problem, $\text{rank}(A_0) = \text{rank}(A_2) = 67 \ll 1005$ and the improvement in the backward errors of quadeig over polyeig is due to our deflation strategy.
- The cd_player problem is heavily damped with $\tau_Q = 9.3 \times 10^3 \gg 1$ and, as a result, quadeig uses no scaling. It computes eigenpairs with small backward errors. Note that these results are not predicted by the upper bounds on the backward errors in (11)–(12), which are large for some of the eigenpairs.

Experiment 2. Table II displays the execution time of polyeig and quadeig for the computation of the eigenvalues Λ alone, and the eigenvalues/eigenvectors (Λ, X) for some large QEPs in the NLEVP collection, some of which have singular leading and trailing coefficient matrices so that deflation takes place in quadeig. For a fair comparison, we made use of the NAG Toolbox (see first paragraph of Section 6) for the execution times reported in Table II. We make the following comments.

- When no deflation occurs, and with the exception of the spring problem, the execution time for the computation of all the eigenvalues is about the same for quadeig and polyeig. This shows that our preprocessing step, which involves two QR factorizations with column-pivoting to check for possible deflation, does not affect the overall execution time.
- polyeig is twice slower than quadeig on the spring problem for the computation of the eigenvalues. The output of the following MATLAB commands shows that the QZ algorithm is slower on the linearization polyeig uses than on the linearization $C_2(\lambda)$ used by quadeig. The reason for this is unclear.

```
n = 1000;
coeffs = nlevp('spring',n);
A0 = coeffs{1}; A1 = coeffs{2}; A2 = coeffs{3};
disp('Linearization C2')
A = [A1 -eye(n); A0 zeros(n)]; B = [-A2 zeros(n); zeros(n) -eye(n)];
tic, eig(A,B); toc
disp('Linearization (1.1) as in polyeig')
A = [A0 zeros(n); zeros(n) eye(n)]; B = [-A1 -A2; eye(n) zeros(n)];
tic, eig(A,B); toc

Linearization C2
Elapsed time is 92.997492 seconds.
Linearization (1.1) as in polyeig
Elapsed time is 182.523154 seconds.
```

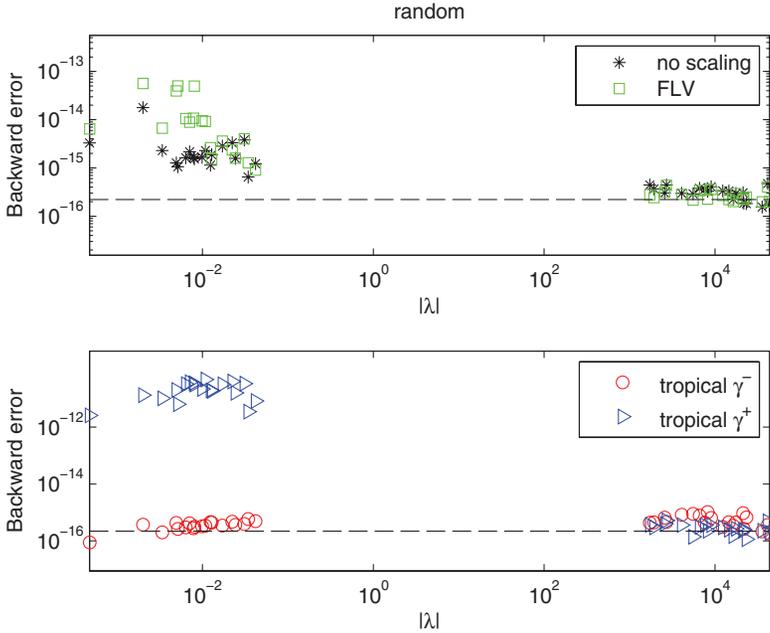


Fig. 1. Backward errors of right eigenpairs computed by `quadeig` and different choices for the eigenvalue parameter scaling `pscale`. For the random problem, $\tau_Q \approx 10^3$ and $\gamma_{trop}^- \approx 10^{-2}$, $\gamma_{trop}^+ \approx 10^4$.

—For the computation of all eigenpairs, Table II also shows that `quadeig` can be faster than `polyeig`, in particular when the leading and trailing coefficient matrices have low rank. For such problems, the deflation of the zero and infinite eigenvalues speeds up the execution time by reducing the size of the problem to which the QZ algorithm is applied. We note that, in contrast with `quadeig`, the eigenpairs computed by `polyeig` are inaccurate for the damped_beam problem (see Experiment 1 and Higham et al. [2008]).

Experiment 3. To give insight into the influence of the parameter `pscale` of `quadeig`, we generated the following heavily damped problems.

- (1) The random problem is a quadratic generated in MATLAB with the commands


```
n = 30; randn('state', 31);
A0 = 1e1*randn(n); A1 = 1.e3*randn(n); A2 = 1e-1*randn(n);
```
- (2) The modified hospital problem is the hospital problem from the NLEVP collection for which the damping matrix is multiplied by 10^3 .
- (3) The modified acoustic problem is the acoustic_wave_2D problem from the NLEVP collection for which the damping matrix is multiplied by 10^2 .

For each problem, we ran `quadeig` with `pscale` = 1 (no eigenvalue parameter scaling), `pscale` = 2 (Fan, Lin, and Van Dooren scaling, FLV), `pscale` = 3 (tropical scaling with smallest root γ_{trop}^-) and `pscale` = 4 (tropical scaling with largest root γ_{trop}^+). The backward errors for the right eigenpairs are shown on Figures 1 to 3.

The top plot on each figure shows that our eigensolver `quadeig` with no scaling can be unstable for heavily damped problems. Note that no scaling is the default option when $\tau_Q > 10$. These plots also show that the Fan, Lin and Van Dooren scaling, `quadeig` does not always produce eigenpairs with small backward errors when $\tau_Q \gg 1$.

The bottom plot of each figure illustrates the tropical scaling. These plots confirm the analysis of Section 3.1.2, which shows that `quadeig` with `pscale` = 3 (`pscale` = 4)

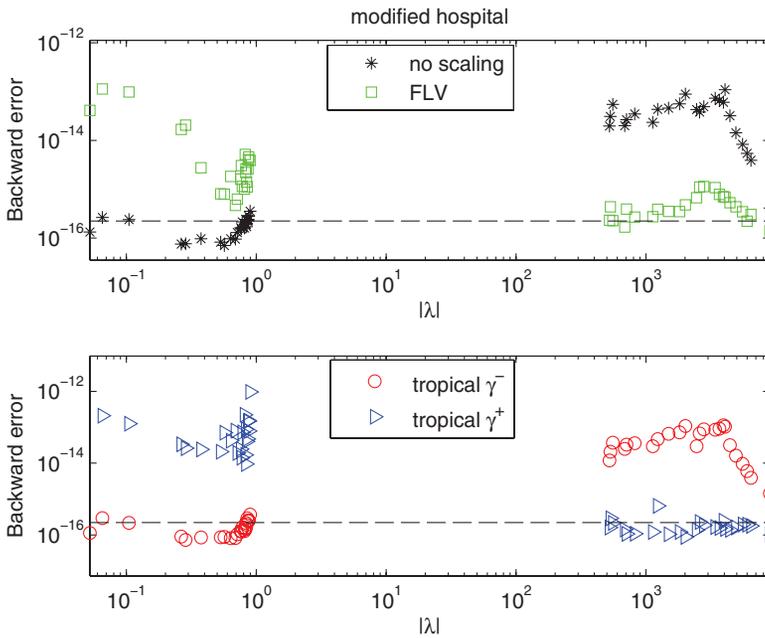


Fig. 2. Backward errors of right eigenpairs computed by `quadeig` and different choices for the eigenvalue parameter scaling `pscale`. For the modified hospital problem, $\tau_Q \approx 66$ and $\gamma_{trop}^- \approx 0.8$, $\gamma_{trop}^+ \approx 3.7 \times 10^3$.

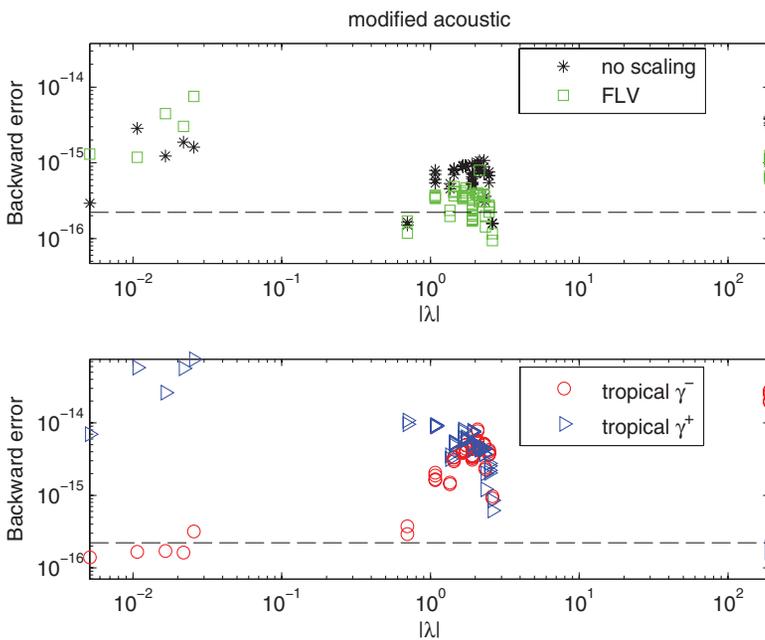


Fig. 3. Backward errors of right eigenpairs computed by `quadeig` and different choices for the eigenvalue parameter scaling `pscale`. For the modified acoustic problem, $\tau_Q \approx 20$ and $\gamma_{trop}^- \approx 0.1$, $\gamma_{trop}^+ \approx 40$.

leads to small backward errors for eigenvalues λ such that $|\lambda| \lesssim \gamma_{trop}^-$ ($|\lambda| \gtrsim \gamma_{trop}^+$). The plots also show that for problems with well-conditioned coefficient matrices A_i (this is the case for the random and modified hospital problems) and for which $\gamma_{trop}^- \ll \gamma_{trop}^+$, the eigenvalues of $Q(\lambda)$ split into two groups: those whose modulus are close to γ_{trop}^- and those whose modulus are close to γ_{trop}^+ (see Gaubert and Sharify [2009, Thm. 2]). For the modified acoustic problem, A_1 is singular and the above result does not hold, as illustrated on Figure 3.

Based on these experiments and the analysis in Section 3.1.2, we recommend the use of tropical scaling with γ_{trop}^- to users interested in the smaller eigenvalues in magnitude, and the use of tropical scaling with γ_{trop}^+ to users interested in the larger eigenvalues in magnitude.

7. CONCLUSIONS

We have described a new algorithm for the computation of all the eigenvalues and optionally the right and left eigenvectors of dense quadratic matrix polynomials. Our algorithm incorporates the scaling of Fan, Lin, and Van Dooren for problems that are not too heavily damped; a choice of linearization with favorable conditioning and backward stability properties; and a preprocessing step that reveals and deflates the zero and infinite eigenvalues contributed by singular leading and trailing matrix coefficients.

The algorithm is backward-stable for quadratics that are not too heavily damped. We note that QEPs in applications are often not heavily damped. For heavily damped problems, our algorithm offers the option of using an eigenvalue parameter scaling based on tropical roots, for which we can show that the smallest tropical root leads to small backward errors for eigenvalues small in absolute value, and the largest tropical root leads to small backward errors for the largest eigenvalues in absolute value.

Numerical experiments show that our MATLAB implementation of the algorithm, `quadeig`, outperforms the MATLAB function `polyeig` in terms of both stability and efficiency. The MATLAB function `quadeig` is available from the Web page <http://www.maths.manchester.ac.uk/~ftisseur/misc/> and a FORTRAN implementation is under development.

ACKNOWLEDGMENTS

The authors thank the referees for their valuable suggestions.

REFERENCES

- AMIRASLANI, A., CORLESS, R. M., AND LANCASTER, P. 2009. Linearization of matrix polynomials expressed in polynomial bases. *IMA J. Numer. Anal.* 29, 141–157.
- ANTONIOU, E. N. AND VOLOGIANNIDIS, S. 2004. A new family of companion forms of polynomial matrices *Electron. J. Linear Algebra* 11, 78–87.
- ANTONIOU, E. N. AND VOLOGIANNIDIS, S. 2006. Linearizations of polynomial matrices with symmetries and their applications. *Electron. J. Linear Algebra* 15, 107–114.
- BETCKE, T. 2008. Optimal scaling of generalized and polynomial eigenvalue problems. *SIAM J. Matrix Anal. Appl.* 30, 4, 1320–1338.
- BETCKE, T., HIGHAM, N. J., MEHRMANN, V., SCHRÖDER, C., AND TISSEUR, F. 2013. NLEVP: A collection of nonlinear eigenvalue problems. *ACM Trans. Math. Softw.* 39, 2, Article 7.
- CHANDRASEKARAN, S. AND IPSEN, I. C. F. 1994. On rank-revealing factorisations. *SIAM J. Matrix Anal. Appl.* 15, 2, 592–622.
- DEDIEU, J.-P. AND TISSEUR, F. 2003. Perturbation theory for polynomial eigenvalue problems in homogeneous form. *Linear Algebra Appl.* 358, 1–3, 71–94.
- FAN, H.-Y., LIN, W.-W., AND VAN DOOREN, P. 2004. Normwise scaling of second order polynomial matrices. *SIAM J. Matrix Anal. Appl.* 26, 1, 252–256.

- GAUBERT, S. AND SHARIFY, M. 2009. Tropical scaling of polynomial matrices. In *Lecture Notes in Control and Information Sciences*, vol. 389, Springer, Berlin, 291–303.
- GOLUB, G. H. AND VAN LOAN, C. F. 1996. *Matrix Computations* 3rd Ed. Johns Hopkins University Press, Baltimore, MD.
- GRAMMONT, L., HIGHAM, N. J., AND TISSEUR, F. 2011. A framework for analyzing nonlinear eigenproblems and parameterized linear systems. *Linear Algebra Appl.* 435, 3, 623–640.
- HIGHAM, N. J. 1990. Analysis of the Cholesky decomposition of a semi-definite matrix. In *Reliable Numerical Computation*, M. G. Cox and S. J. Hammarling, Eds., Oxford University Press, 161–185.
- HIGHAM, N. J., LI, R.-C., AND TISSEUR, F. 2007. Backward error of polynomial eigenproblems solved by linearization. *SIAM J. Matrix Anal. Appl.* 29, 4, 1218–1241.
- HIGHAM, N. J., MACKEY, D. S., MACKEY, N., AND TISSEUR, F. 2006. Symmetric linearizations for matrix polynomials. *SIAM J. Matrix Anal. Appl.* 29, 1, 143–159.
- HIGHAM, N. J., MACKEY, D. S., AND TISSEUR, F. 2006. The conditioning of linearizations of matrix polynomials. *SIAM J. Matrix Anal. Appl.* 28, 4, 1005–1028.
- HIGHAM, N. J., MACKEY, D. S., TISSEUR, F., AND GARVEY, S. D. 2008. Scaling, sensitivity and stability in the numerical solution of quadratic eigenvalue problems. *Int. J. Numer. Methods Eng.* 73, 3, 344–360.
- HILLIGES, A., MEHL, C., AND MEHRMANN, V. 2004. On the solution of palindromic eigenvalue problems. In *Proceedings of the European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS'04)*. P. Neittaanmaki, et al. Eds. <http://www.mit.jyu.fi/eccomas2004/proceedings/proceed.html>.
- KÅGSTRÖM, B. AND KRESSNER, D. 2006. Multishift variants of the QZ algorithm with aggressive early deflation. *SIAM J. Matrix Anal. Appl.* 29, 199–227.
- LANCASTER, P. 2008. Linearization of regular matrix polynomials. *Electron. J. Linear Algebra* 17, 21–27.
- LANCASTER, P. AND PSARRAKOS, P. 2005. A note on weak and strong linearizations of regular matrix polynomials. Numerical analysis rep. 470, Manchester Centre for Computational Mathematics, Manchester, UK.
- LEMONNIER, D. AND VAN DOOREN, P. M. 2006. Balancing regular matrix pencils. *SIAM J. Matrix Anal. Appl.* 28, 1, 253–263.
- MACKEY, D. S., MACKEY, N., MEHL, C., AND MEHRMANN, V. 2006a. Structured polynomial eigenvalue problems: Good vibrations from good linearizations. MIMS EPrint 2006.38, The University of Manchester, UK.
- MACKEY, D. S., MACKEY, N., MEHL, C., AND MEHRMANN, V. 2006b. Structured polynomial eigenvalue problems: Good vibrations from good linearizations. *SIAM J. Matrix Anal. Appl.* 28, 4, 1029–1051.
- MACKEY, D. S., MACKEY, N., MEHL, C., AND MEHRMANN, V. 2006c. Vector spaces of linearizations for matrix polynomials. *SIAM J. Matrix Anal. Appl.* 28, 4, 971–1004.
- MOLER, C. B. AND STEWART, G. W. 1973. An algorithm for generalized matrix eigenvalue problems. *SIAM J. Numer. Anal.* 10, 2, 241–256.
- NAG Toolbox for MATLAB. NAG Ltd., Oxford. <http://www.nag.co.uk/>.
- TISSEUR, F. 2000. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra Appl.* 309, 339–361.
- TISSEUR, F. AND MEERBERGEN, K. 2001. The quadratic eigenvalue problem. *SIAM Rev.* 43, 2, 235–286.
- WATKINS, D. S. 2000. Performance of the QZ algorithm in the presence of infinite eigenvalues. *SIAM J. Matrix Anal. Appl.* 22, 364–375.

Received December 2011; revised May 2012; accepted September 2012