



Swapping 2×2 blocks in the Schur and generalized Schur form[☆]



Daan Camps^a, Nicola Mastronardi^b, Raf Vandebril^{a,*}, Paul Van Dooren^c

^a Department Computer Science, KU Leuven, Leuven, Belgium

^b Istituto per le Applicazioni del Calcolo "M. Picone", Consiglio Nazionale delle Ricerche, sede di Bari, Italy

^c Department of Mathematical Engineering, Catholic University of Louvain, Louvain-la-Neuve, Belgium

ARTICLE INFO

Article history:

Received 8 January 2019

Received in revised form 14 May 2019

Keywords:

Eigenvalues

Schur form

Reordering Schur form

Stability

ABSTRACT

In this paper we describe how to swap two 2×2 blocks in a real Schur form and a generalized real Schur form. We pay special attention to the numerical stability of the method. We also illustrate the stability of our approach by a series of numerical tests.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

The real Schur form of a standard eigenvalue problem $\lambda I_n - A$ is an orthogonal matrix decomposition that plays a fundamental role in numerical linear algebra. It is an upper block triangular form A_S that can be obtained under an orthogonal similarity transformation U :

$$A_S := U^T A U = \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1k} \\ 0 & A_{22} & \ddots & A_{2k} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & A_{kk} \end{bmatrix}, \quad (1)$$

where the diagonal blocks A_{ii} , $i = 1, \dots, k$ are of dimension 1×1 or 2×2 . This form is not only useful for computing the spectrum of A but also for computing an invariant subspace of the matrix A with a spectrum constrained to a particular region in the complex plane (see, e.g., [1–4]). For this, one needs to *reorder* the Schur form A_S by updating the orthogonal transformation U such that the diagonal blocks with their spectrum inside the considered region, appear first in the block triangular form. Such a reordering can always be performed by a series of swappings of two adjacent blocks in the real Schur form, as was shown in [1,5].

[☆] The work of D.C. and R.V. was supported by the Research Council KU Leuven, Belgium, project C14/16/056, Inverse-free Rational Krylov Methods: Theory and Applications; and the work of N.M. was partly supported by GNCS-INDAM, Italy and by CNR, Italy under the Short Term Mobility Program.

* Corresponding author.

E-mail addresses: daan.camps@cs.kuleuven.be (D. Camps), n.mastronardi@ba.iac.cnr.it (N. Mastronardi), raf.vandebril@cs.kuleuven.be (R. Vandebril), paul.vandooren@uclouvain.be (P. Van Dooren).

There is a corresponding matrix decomposition known as the generalized real Schur form of a so-called *regular pencil* $\lambda B - A$, implying that $\det(\lambda B - A)$ is not identically zero. The equivalent generalized Schur form is upper block triangular and can be obtained under orthogonal equivalence transformations:

$$\lambda B_S - A_S := Z^T(\lambda B - A)Q = \lambda \begin{bmatrix} B_{11} & B_{12} & \dots & B_{1k} \\ 0 & B_{22} & \ddots & B_{2k} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & B_{kk} \end{bmatrix} - \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1k} \\ 0 & A_{22} & \ddots & A_{2k} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & A_{kk} \end{bmatrix},$$

where the subpencils $(\lambda B_{ii} - A_{ii}), i = 1, \dots, k$ are of dimension 1×1 or 2×2 and where B_S is upper triangular. The concept corresponding to an invariant subspace here is that of a deflating subspace, and again one needs to reorder the blocks in the generalized Schur form by updating the orthogonal transformations Q and Z such that the diagonal blocks $(\lambda B_{ii} - A_{ii})$ with their spectrum inside the relevant region, appear first in the decomposition. The swapping problem for two consecutive blocks in the generalized Schur form was discussed in, e.g., [2,6,7].

The swapping of a scalar block with another scalar block or with a 2×2 block, was addressed adequately in the papers [1,2,6], for both the standard and the generalized Schur form. But the problem of swapping two 2×2 blocks still remains a delicate problem, which can fail when restricting oneself to a direct method [3]. In this paper we revisit the problem of swapping two 2×2 blocks in both the standard and the generalized Schur form. This problem gained a renewed interest because of its use in the new rational QZ method [8–10] which makes extensive use of the swapping of 2×2 blocks that may be poorly conditioned. It is therefore needed to ensure that these swappings can be performed in a numerically reliable manner.

In Sections 2 and 3 we revisit the swapping problem for the standard real Schur form, and for the generalized real Schur form and we present a new method to perform the swapping. In both sections we perform a detailed error analysis to bound the relative norm of the off diagonal errors, after the swapping operation. In case these errors are not sufficiently small, we show how to update the swapping by an iterative process to further reduce these errors. In Section 4 we compare the different possible approaches for swapping 2×2 blocks and we show statistics for the accuracy and need for iterative refinement.

2. The standard eigenvalue problem

The problem we address in this section is the swapping of two 2×2 blocks in a matrix A in real Schur form. We restrict ourselves to the case of a 4×4 matrix A

$$A := \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix} = \left[\begin{array}{cc|cc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{array} \right], \tag{2}$$

where the spectra of A_{11} and A_{22} are disjoint and not real

$$\Lambda(A_{11}) := \{\alpha_1 \pm j\beta_1\} \neq \Lambda(A_{22}) := \{\alpha_2 \pm j\beta_2\}, \quad \beta_i \neq 0, \quad i = 1, 2.$$

We are looking for a real orthogonal transformation Q such that the transformed matrix $\tilde{A} := Q^T A Q$ has the form

$$\tilde{A} := \begin{bmatrix} \tilde{A}_{11} & \tilde{A}_{12} \\ 0 & \tilde{A}_{22} \end{bmatrix} = \left[\begin{array}{cc|cc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{array} \right],$$

with $\Lambda(\tilde{A}_{11}) := \Lambda(A_{22}), \quad \Lambda(\tilde{A}_{22}) := \Lambda(A_{11})$.

The suggested construction is to find an orthonormal basis for the invariant subspace of A corresponding to the spectrum $\Lambda(A_{22})$, and to complete this to an orthonormal matrix Q . This can be obtained by the following method described in [3,5].

Lemma 1. *Let $A \in \mathbb{R}^{4 \times 4}$ be given as in (2), where the spectra of A_{11} and A_{22} are disjoint. Then the invariant subspace of A corresponding to the spectrum of A_{22} is spanned by the matrix*

$$\begin{bmatrix} -X \\ I_2 \end{bmatrix}, \quad \text{where } A_{11}X - XA_{22} = A_{12}. \tag{3}$$

Moreover, the orthogonal similarity transformation Q swaps the spectra of A_{11} and A_{22} if and only if

$$\begin{bmatrix} -X \\ I_2 \end{bmatrix} = Q \begin{bmatrix} R \\ 0 \end{bmatrix}, \tag{4}$$

where R is square and invertible.

The error analysis of Bai and Demmel [3] holds for the algorithmic implementation where X is first solved from the Sylvester equation (3), and then Q is constructed from the QR factorization (4).

Lemma 2 (Theorem 2 in [3]). Let \hat{X} be the computed solution of the Sylvester equation (3), let $E := -A_{12} - A_{11}\hat{X} + \hat{X}A_{22}$ be its residual and let \hat{Q} be the computed factor of the QR factorization

$$\begin{bmatrix} -\hat{X} \\ I_2 \end{bmatrix} = \hat{Q} \begin{bmatrix} \hat{R} \\ 0 \end{bmatrix}.$$

Then the computed similarity transformation satisfies

$$\begin{bmatrix} \tilde{A}_{11} & \tilde{A}_{12} \\ \Delta & \tilde{A}_{22} \end{bmatrix} := \hat{Q}^T \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix} \hat{Q}, \tag{5}$$

where

$$\|\Delta\|_2 \leq \|E\|_F / [1 + \sigma_2(X)^2].$$

It is pointed out in [3] that their error analysis does not imply that $\|\Delta\|_2 \leq \epsilon_M \|A\|_2$, where ϵ_M is the machine precision of the computer used. Nevertheless, the bound is often pessimistic and the observed errors often allow Δ to be safely dismissed.

In the rational QZ algorithm [8,10] $\mathcal{O}(n^2)$ swaps are executed, and failures occasionally happen due to inaccurate swaps of 2×2 blocks. In this paper, we try to address this problem when the above direct approach does not work properly.

For this we propose two adaptations of the so-called direct approach. The first one is the construction of the transformation Q from (4), the second one is an iterative refinement step to get Δ from (5) as small as possible.

If we perform a singular value decomposition of X then we obtain such a factorization from

$$X = U_X \Sigma_X V_X^T \implies Q = \begin{bmatrix} U_X & 0 \\ 0 & V_X \end{bmatrix} \begin{bmatrix} C_X & S_X \\ -S_X & C_X \end{bmatrix}, \quad R = -S_X^{-1} V_X^T,$$

where $S_X := (I + \Sigma_X^2)^{-\frac{1}{2}}$ and $C_X := \Sigma_X S_X$ are diagonal matrices that can be computed elementwise from Σ_X . The recommended procedure for computing the diagonal elements c_i, s_i ($i = 1, 2$) of C_X and S_X is

$$\begin{cases} s_i := 1/\sqrt{1 + \sigma_i^2}, & c_i := s_i \cdot \sigma_i, & \text{if } \sigma_i \geq 1, \\ c_i := 1/\sqrt{1 + 1/\sigma_i^2}, & s_i := c_i/\sigma_i, & \text{if } \sigma_i < 1, \end{cases}$$

because this computes the smallest quantities to a higher relative accuracy. In the numerical examples given below, we show that this alternative construction for the transformation Q often gives better results than the use of a simple QR factorization.

But even for this approach, there is no guarantee that the similarity transformation (5) yields a negligible off-diagonal block Δ . We recommend then to perform an updating similarity to further reduce the norm of the block Δ in \tilde{A} . This can be done by applying a similarity transformation:

$$Q_{up}^T \tilde{A} Q_{up} = \begin{bmatrix} \hat{A}_{11} & \hat{A}_{12} \\ 0 & \hat{A}_{22} \end{bmatrix}, \quad Q_{up} := \begin{bmatrix} I & Y^T \\ -Y & I \end{bmatrix} \begin{bmatrix} R_Y & 0 \\ 0 & R_{Y^T} \end{bmatrix},$$

where R_Y and R_{Y^T} are normalizations to make Q_{up} orthonormal, and where Y is computed from the quadratic matrix equation

$$Y\tilde{A}_{12}Y + \tilde{A}_{22}Y - Y\tilde{A}_{11} - \Delta = 0.$$

Since the updating transformation Q_{up} is assumed to be close to the identity, $\|Y\|_2$ is very small and can then be well approximated by the solution of the linear equation

$$\tilde{A}_{22}Y - Y\tilde{A}_{11} = \Delta.$$

The solution Y of this linear system can again be computed using Kronecker products, and the implementation of the rotation is best performed via the singular value decomposition of Y since $Y := U_Y \Sigma_Y V_Y^T$ implies

$$Q_{up} = \begin{bmatrix} V_Y & 0 \\ 0 & U_Y \end{bmatrix} \begin{bmatrix} C_Y & S_Y \\ -S_Y & C_Y \end{bmatrix}, \quad R_Y = C_Y^{-1} V_Y^T, \quad R_{Y^T} = C_Y^{-1} U_Y^T,$$

where $C_Y := (I + \Sigma_Y^2)^{-\frac{1}{2}}$ and $S_Y := \Sigma_Y C_Y$ are diagonal matrices that can be computed elementwise from Σ_Y using the procedure described earlier. The updating transformation will not yield a new off-diagonal block that is exactly zero, but its norm can be expected to be of the order of $\|\Delta\|_2^2$, since this is the Riccati equation approach for computing invariant subspaces. This would imply that in most cases, one updating transformation Q_{up} is enough to be able to dismiss the new off-diagonal block. The numerical experiments given below, support this claim.

3. The generalized eigenvalue problem

In the corresponding problem for pencils of matrices we need to swap two 2×2 pencils in a pencil $\lambda B - A$ in real Schur form, with B and A like

$$\begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix} = \left[\begin{array}{cc|cc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ \hline 0 & 0 & \times & \times \\ 0 & 0 & 0 & \times \end{array} \right], \quad \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix} = \left[\begin{array}{cc|cc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{array} \right]. \tag{6}$$

Since we are looking for two pairs of disjoint complex conjugate eigenvalues, the matrix B is invertible and the spectrum of the above 4×4 pencil is the union of $\Lambda(B_{11}^{-1}A_{11}) := \{\alpha_1 \pm j\beta_1\}, \beta_1 \neq 0$ and $\Lambda(B_{22}^{-1}A_{22}) := \{\alpha_2 \pm j\beta_2\}, \beta_2 \neq 0$.

We are looking for two real orthogonal transformations Q and Z such that the transformed pencil $\lambda \tilde{B} - \tilde{A} := Z^T(\lambda B - A)Q$ has \tilde{B} and \tilde{A} of the form

$$\begin{bmatrix} \tilde{B}_{11} & \tilde{B}_{12} \\ 0 & \tilde{B}_{22} \end{bmatrix} = \left[\begin{array}{cc|cc} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ \hline 0 & 0 & \times & \times \\ 0 & 0 & 0 & \times \end{array} \right], \quad \begin{bmatrix} \tilde{A}_{11} & \tilde{A}_{12} \\ 0 & \tilde{A}_{22} \end{bmatrix} = \left[\begin{array}{cc|cc} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \hline 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{array} \right],$$

with $\Lambda(\tilde{B}_{11}^{-1}\tilde{A}_{11}) := \Lambda(B_{22}^{-1}A_{22}), \Lambda(\tilde{B}_{22}^{-1}\tilde{A}_{22}) := \Lambda(B_{11}^{-1}A_{11})$. The suggested construction is to find a pair of orthonormal bases for the right and left deflating subspaces of $\lambda B - A$ corresponding to the spectrum of $B_{22}^{-1}A_{22}$ and $B_{11}^{-1}A_{11}$, respectively, and to complete these bases to orthonormal transformations Q and Z , respectively. This can be obtained by the following direct method, described in [7].

Lemma 3. *Let the pencil $\lambda B - A$ be given as in (6), where the spectra of $B_{11}^{-1}A_{11}$ and $B_{22}^{-1}A_{22}$ are disjoint. Then the right and left deflating subspaces of $\lambda B - A$ corresponding to the spectrum of $B_{22}^{-1}A_{22}$ and $B_{11}^{-1}A_{11}$, respectively, are spanned by the columns and rows of the matrices*

$$\begin{bmatrix} -X \\ I_2 \end{bmatrix}, \quad [I_2 \quad Y], \quad \text{where} \quad \begin{cases} A_{11}X - YA_{22} = A_{12}, \\ B_{11}X - YB_{22} = B_{12}. \end{cases} \tag{7}$$

Moreover, the orthogonal equivalence transformations Q and Z swap the spectra of the diagonal blocks in $Z^T(\lambda B - A)Q$ if and only if

$$\begin{bmatrix} -X \\ I_2 \end{bmatrix} = Q \begin{bmatrix} R_X \\ 0 \end{bmatrix}, \quad \text{and} \quad [I_2 \quad Y] = [0 \quad R_Y] Z^T, \tag{8}$$

where R_X and R_Y are square and invertible.

The error analysis of Kågström’s algorithm [7] holds again for a particular algorithmic implementation where X and Y are first solved from the generalized Sylvester equations (7), and then Q and Z are constructed from the QR factorizations (8).

Lemma 4 (Theorem 3.1 in [7]). *Let \hat{X} and \hat{Y} be the computed solutions of the generalized Sylvester equation (7), let $E := -A_{12} - A_{11}\hat{X} + \hat{Y}A_{22}$ and $F := -B_{12} - B_{11}\hat{X} + \hat{Y}B_{22}$ be their residuals and let \hat{Q} and \hat{Z} be the computed factors of the QR factorizations*

$$\begin{bmatrix} -\hat{X} \\ I_2 \end{bmatrix} = \hat{Q} \begin{bmatrix} \hat{R}_X \\ 0 \end{bmatrix}, \quad \begin{bmatrix} I_2 \\ \hat{Y}^T \end{bmatrix} = \hat{Z} \begin{bmatrix} 0 \\ \hat{R}_Y^T \end{bmatrix}.$$

Then the computed equivalence transformation satisfies

$$\begin{bmatrix} \lambda \tilde{B}_{11} - \tilde{A}_{11} & \lambda \tilde{B}_{11} - \tilde{A}_{12} \\ \lambda \Delta_B - \Delta_A & \lambda \tilde{B}_{22} - \tilde{A}_{22} \end{bmatrix} := \hat{Z}^T \begin{bmatrix} \lambda B_{11} - A_{11} & \lambda B_{12} - A_{12} \\ 0 & \lambda B_{22} - A_{22} \end{bmatrix} \hat{Q},$$

where

$$\|\Delta_A\|_2 \leq \|E\|_F / \sqrt{(1 + \sigma_2(X)^2)(1 + \sigma_2(Y)^2)},$$

$$\|\Delta_B\|_2 \leq \|F\|_F / \sqrt{(1 + \sigma_2(X)^2)(1 + \sigma_2(Y)^2)}.$$

It is again pointed out in [7] that this error analysis does not imply that the off-diagonal block $\lambda \Delta_B - \Delta_A$ can be dismissed. We therefore recommend to apply the same techniques here as for the standard eigenvalue problem: an alternative manner of computing the transformation and an iterative refinement step.

For the constructions of Q and Z from the matrices X and Y we thus recommend to first compute the singular value decompositions

$$X = U_X \Sigma_X V_X^T, \quad Y^T = U_Y \Sigma_Y V_Y^T$$

from which we can construct

$$Q = \begin{bmatrix} U_X & 0 \\ 0 & V_X \end{bmatrix} \begin{bmatrix} C_X & S_X \\ -S_X & C_X \end{bmatrix}, \quad R_X = -S_X^{-1}V_X^T,$$

where $S_X := (I + \Sigma_X^2)^{-\frac{1}{2}}$ and $C_X := \Sigma_X S_X$, and

$$Z = \begin{bmatrix} V_Y & 0 \\ 0 & U_Y \end{bmatrix} \begin{bmatrix} C_Y & -S_Y \\ S_Y & C_Y \end{bmatrix}, \quad R_Y^T = C_Y^{-1}V_Y^T,$$

where $C_Y := (I + \Sigma_Y^2)^{-\frac{1}{2}}$ and $S_Y := \Sigma_Y C_Y$.

The updating of the equivalence transformation to further reduce the norm of the block $\lambda\Delta_B - \Delta_A$ in $\lambda\tilde{B} - \tilde{A}$ can be done by applying an equivalence transformation

$$\begin{bmatrix} \lambda\hat{B}_{11} - \hat{A}_{11} & \lambda\hat{B}_{12} - \hat{A}_{12} \\ 0 & \lambda\hat{B}_{22} - \hat{A}_{22} \end{bmatrix} := Z_{up}^T \begin{bmatrix} \lambda\tilde{B}_{11} - \tilde{A}_{11} & \lambda\tilde{B}_{11} - \tilde{A}_{12} \\ \lambda\Delta_B - \Delta_A & \lambda\tilde{B}_{22} - \tilde{A}_{22} \end{bmatrix} Q_{up},$$

$$Q_{up} := \begin{bmatrix} I & X^T \\ -X & I \end{bmatrix} \begin{bmatrix} R_X & 0 \\ 0 & R_{X^T} \end{bmatrix}, \quad Z_{up} := \begin{bmatrix} I & Y^T \\ -Y & I \end{bmatrix} \begin{bmatrix} R_Y & 0 \\ 0 & R_{Y^T} \end{bmatrix},$$

where R_X, R_{X^T}, R_Y and R_{Y^T} are normalization factors to make Q_{up} and Z_{up} orthonormal. The (X, Y) are computed from the system of quadratic matrix equations

$$\tilde{\Delta}_A - \tilde{A}_{22}X + Y\tilde{A}_{11} - Y\tilde{A}_{12}X = 0, \quad \tilde{\Delta}_B - \tilde{B}_{22}X + Y\tilde{B}_{11} - Y\tilde{B}_{12}X = 0.$$

These can be approximated by the system of linear equations

$$\tilde{\Delta}_A = \tilde{A}_{22}X - Y\tilde{A}_{11}, \quad \tilde{\Delta}_B = \tilde{B}_{22}X - Y\tilde{B}_{11},$$

since $\|X\|_2$ and $\|Y\|_2$ are very small. The solution (X, Y) of this linear system can be computed using Kronecker products.

4. Numerical results

Our MATLAB implementation of the swapping methods is made publicly available on <https://github.com/campsd/swap22>. We first test the swapping of 2×2 blocks for the standard eigenvalue problem. To this end, we randomly generated 4×4 block triangular matrices with two pairs of complex conjugate eigenvalues. The diagonal blocks are generated as:

$$A_{11} = \begin{bmatrix} a & bk \\ -b/k & a \end{bmatrix}, \quad A_{22} = \begin{bmatrix} a + r_1g & (b + r_2g)k \\ -(b + r_2g)/k & a + r_1g \end{bmatrix},$$

with a, b, r_1 and r_2 random values drawn from the standard normal distribution. We abuse notation and denote by λ_1 the complex pair of eigenvalues $a \pm jb$ and by λ_2 the pair $(a + r_1g) \pm j(b + r_2g)$; with $|\lambda_1 - \lambda_2|$ we denote the maximum of the difference of the corresponding eigenvalues. We use the parameters g and k to change respectively the gap $|\lambda_1 - \lambda_2|$ and the condition number of the eigenvector. Both $|\lambda_1 - \lambda_2|$ and k are varied from 10^{-12} up to 10^{12} in 30 logarithmically spaced increments. For every combination, 20 random matrices have been generated for which the eigenvalues are swapped. Fig. 1 shows the percentage of times the swap needed iterative refinement when Q was either computed from the SVD of X or from the QR factorization of $\begin{bmatrix} -X \\ I \end{bmatrix}$. Fig. 1 indicates that the first method outperforms the second. Fig. 2 shows the relative backward error for both cases. This figure indicates backward stability in both cases, but again the first method outperforms the second. Finally, Figs. 3 and 4 show the relative forward errors on the eigenvalues after the swap for respectively the SVD and QR update.

The right part of Fig. 4 shows that there are two regions in the parameter domain where the relative forward error on $\hat{\lambda}_2$, computed with the QR update, is of the order of 10^{-12} which is significantly higher than the error with the SVD update which is of order 10^{-15} .

As the generalized eigenvalue problem poses additional challenges in comparison to the standard eigenvalue problem, we also perform a numerical experiment to illustrate the reordering problem of 2×2 blocks in generalized Schur form. We used the following strategy to generate 4×4 block triangular pencils with prescribed gap and eigenvector condition number: the matrix B is first randomly generated as a block triangular matrix having blocks of dimension 2×2 with its entries drawn from the standard normal distribution, afterwards the diagonal blocks A_{11} and A_{22} are computed such that:

$$B_{11}^{-1}A_{11} = \begin{bmatrix} a & bk \\ -b/k & a \end{bmatrix}, \quad B_{22}^{-1}A_{22} = \begin{bmatrix} a + r_1g & (b + r_2g)k \\ -(b + r_2g)/k & a + r_1g \end{bmatrix},$$

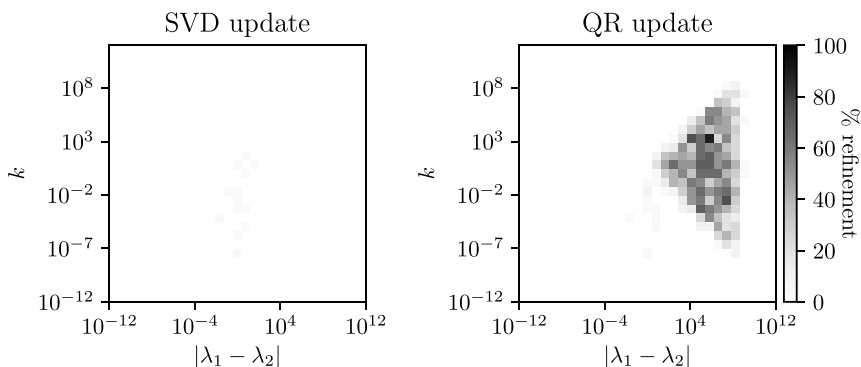


Fig. 1. Percentage of times the swap required iterative refinement.

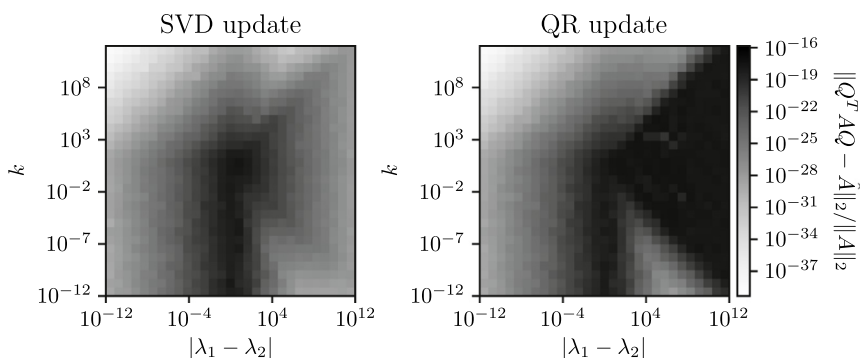


Fig. 2. Relative backward error of the swap.

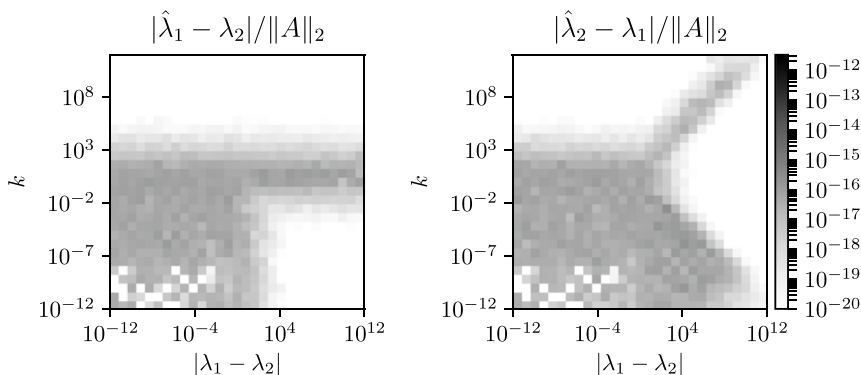


Fig. 3. Relative forward error on the eigenvalues after the swap computed with the SVD update.

with a, b, r_1 and r_2 random values drawn from the standard normal distribution. In this case, we independently varied g and k from 10^{-6} up to 10^6 in 20 logarithmically spaced increments. The results are averaged over 20 runs and summarized in Figs. 5 and 6 which show respectively the average number of iterative refinement iterations that were required to obtain a dismissible off-diagonal block and the forward error on the eigenvalues after the swapping transformation as computed with the SVD.

The results differ from the standard eigenvalue problem. Fig. 5 reveals that in this experiment the SVD update required on average more refinement iterations than the QR update. Fig. 6 shows that if the eigenvectors are well-conditioned, the relative forward error on the eigenvalues is very small. If the condition number increases, the relative forward error can grow to the order of 10^{-10} . The forward error obtained with the QR update behaves qualitatively similar to Fig. 5. For

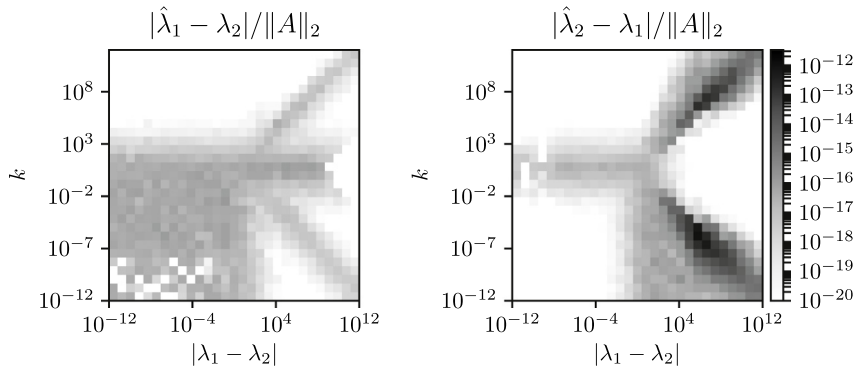


Fig. 4. Relative forward error on the eigenvalues after the swap computed with the QR update.

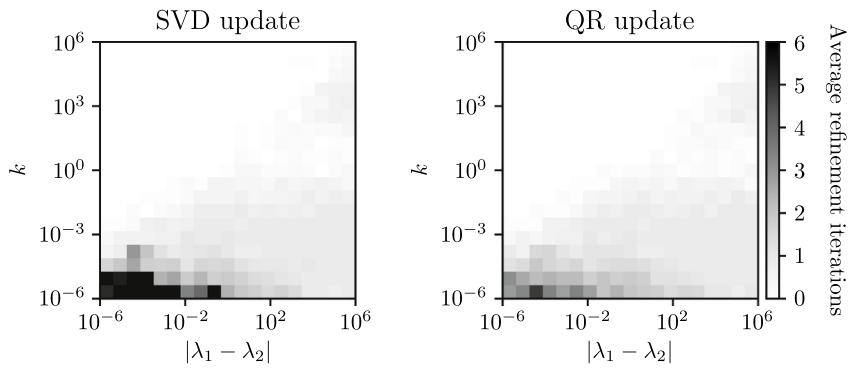


Fig. 5. Average number of iterative refinement iterations.

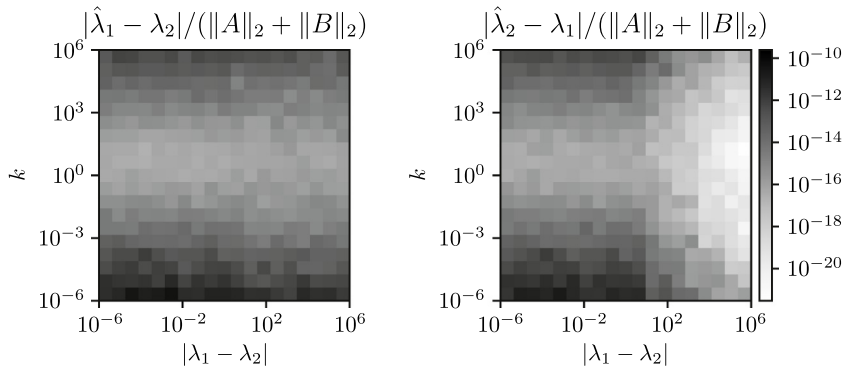


Fig. 6. Relative forward error on the eigenvalues after the swap computed with the SVD update.

this reason, we have chosen to omit the figure for the sake of brevity. We remark that in all problems, the computed swapping transformations were backward stable.

5. Concluding remarks

In this paper we have revisited the problem of swapping two 2×2 blocks on the diagonal of a standard or generalized Schur form. We have introduced two simple modifications of the direct swapping techniques introduced and analyzed in [3,5,7]. The modifications are based on the use of the generalized singular value decomposition of an orthogonal matrix, and the iterative refinement of invariant and deflating subspaces, based on Riccati equations and their linear approximations. These modifications reveal substantial improvements.

References

- [1] G. Stewart, HQR3 and EXCHNG. Fortran subroutines for calculating and ordering the eigenvalues of a real upper hessenberg matrix, *ACM Trans. Math. Software* 2 (3) (1976) 275–280.
- [2] P. Van Dooren, A generalized eigenvalue approach for solving Riccati equations, *SIAM J. Sci. Stat. Comput.* 2 (1981) 121–135.
- [3] Z. Bai, J. Demmel, On swapping diagonal blocks in real Schur form, *Linear Algebra Appl.* 186 (1993) 73–95.
- [4] A. Bojanczyk, P. Van Dooren, Reordering diagonal blocks in the real Schur form, in: M. Moonen, G. Golub, B.D. Moor (Eds.), *Linear Algebra for Large Scale and Real-Time Applications*, in: NATO ASI Series (Series E: Applied Sciences), Springer, Dordrecht, 1993, pp. 351–352.
- [5] B. Ng, B. Parlett, Development of an Accurate Algorithm for EXP(Bt), Part I, Programs To Swap Diagonal Blocks, Univ. of California, Berkeley, 1988.
- [6] P. Van Dooren, ALGORITHM 590: DSUBSP and EXCHQZ: Fortran routines for computing deflating subspaces with specified spectrum, *ACM Trans. Math. Software* 8 (1983) 376–382.
- [7] B. Kågström, A direct method for reordering eigenvalues in the generalized Schur form of a regular matrix pair (A, B) , in: M. Moonen, G. Golub, B.D. Moor (Eds.), *Linear Algebra for Large Scale and Real-Time Applications*, in: NATO ASI Series (Series E: Applied Sciences), Springer, Dordrecht, 1993, pp. 195–218.
- [8] D. Camps, K. Meerbergen, R. Vandebril, A rational QZ method, 2018, accepted for publication in *Simax*, [arXiv:1802.04094](https://arxiv.org/abs/1802.04094), math.NA.
- [9] D. Camps, K. Meerbergen, R. Vandebril, A multishift, multipole rational QZ method with aggressive early deflation, 2019, submitted for publication, [arXiv:1902.10954](https://arxiv.org/abs/1902.10954), math.NA.
- [10] M. Berljafa, S. Güttel, Generalized rational Krylov decompositions with an application to rational approximation, *SIAM J. Matrix Anal. Appl.* 36 (2) (2015) 894–916.