

Implicit double shift QR -algorithm for companion matrices

Marc Van Barel *Raf Vandebril*
Paul Van Dooren

Report TW 532, November 2008



Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Implicit double shift QR -algorithm for companion matrices

Marc Van Barel Raf Vandebril
Paul Van Dooren

Report TW 532, November 2008

Department of Computer Science, K.U.Leuven

Abstract

In this paper an implicit (double) shifted QR -method for computing the eigenvalues of companion and fellow matrices will be presented. Companion and fellow matrices are Hessenberg matrices, that can be decomposed into the sum of a unitary and a rank 1 matrix. The Hessenberg, the unitary as well as the rank 1 structures are preserved under a step of the QR -method. This makes these matrices suitable for the design of a fast QR -method.

Several techniques already exist for performing a QR -step. The implementation of these methods is highly dependent on the representation used. Unfortunately for most of the methods compression is needed since one is not able to maintain all three, unitary, Hessenberg and rank 1 structures.

In this manuscript an implicit algorithm will be designed for performing a step of the QR -method on the companion or fellow matrix based on a new representation consisting of Givens transformations. Moreover, no compression is needed as the specific representation of the involved matrices is maintained. Finally, also a double shift version of the implicit method is presented.

Keywords : Companion matrices, Fellow matrices, QR -algorithm, implicit method, multishift.

MSC : Primary : 65F15, Secondary : 11C20.

Implicit double shift QR -algorithm for companion matrices ^{*}

Marc Van Barel¹, Raf Vandebril¹, Paul Van Dooren²

¹ Katholieke Universiteit Leuven, Dept. of Computer Science, Celestijnenlaan 200A, B-3001 Leuven (Heverlee), Belgium

e-mail: {raf.vandebril, marc.vanbarel}@cs.kuleuven.ac.be

² Université Catholique de Louvain, Dept. of Mathematical Engineering, Avenue Georges Lemaitre 4, B-1348 Louvain-la-Neuve, Belgium

e-mail: paul.vandooren@uclouvain.be

December 3, 2008

Abstract In this paper an implicit (double) shifted QR -method for computing the eigenvalues of companion and fellow matrices will be presented. Companion and fellow matrices are Hessenberg matrices, that can be decomposed into the sum of a unitary and a rank 1 matrix. The Hessenberg, the unitary as well as the rank 1 structures are preserved under a step of the QR -method. This makes these matrices suitable for the design of a fast QR -method.

Several techniques already exist for performing a QR -step. The implementation of these methods is highly dependent on the representation used. Unfortunately for most of the methods compression is needed since one is not able to maintain all three, unitary, Hessenberg and rank 1 structures.

In this manuscript an implicit algorithm will be designed for performing a step of the QR -method on the companion or fellow matrix based on a new representation consisting of Givens transformations. Moreover, no compression is needed as the specific representation of the involved matrices is maintained. Finally, also a double shift version of the implicit method is presented.

Key words Companion matrices, Fellow matrices, QR -algorithm, implicit method, multishift

AMS-Classification 65F15, 11C20

1 Introduction

Computing roots of polynomials is an interesting problem with many challenges. Many methods act directly on the polynomials, e.g., bisection, Sturm sequences and so forth [1]. In actual implementations, however, the problem often is formulated in a matrix setting and the corresponding matrix problem is solved. When working with the standard monomial basis, computing roots of a polynomial coincides with computing eigenvalues of the associated matrix. The eigenvalues of a companion matrix equal the roots of the associated polynomial.

^{*} The research of the first two authors, was partially supported by the Research Council K.U.Leuven, project OT/05/40 (Large rank structured matrix computations), CoE EF/05/006 Optimization in Engineering (OPTEC), by the Fund for Scientific Research–Flanders (Belgium), G.0423.05 (RAM: Rational modeling: optimal conditioning and stable algorithms). All three authors were supported by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office, Belgian Network DYSCO (Dynamical Systems, Control, and Optimization). The first author has a grant as “Postdoctoraal Onderzoeker” from the Fund for Scientific Research–Flanders (Belgium).

Definition 1 Given a monic polynomial $p(z) = p_0 + p_1z + p_2z^2 + \dots + p_{n-1}z^{n-1} + z^n$ the associated companion matrix is of the following form:

$$C_p = \begin{bmatrix} 0 & & & -p_0 \\ 1 & \ddots & & -p_1 \\ & \ddots & & \vdots \\ & & 1 & -p_{n-1} \end{bmatrix}. \quad (1)$$

When changing the basis the idea remains the same, but the associated matrix changes, one obtains e.g., confederate, congenial, comrade matrices and so forth [2]. All these problems are worth investigating, but in this manuscript we will restrict ourselves to the class of companion and fellow matrices because it is the most widespread.

In the current LAPACK implementation the eigenvalues of the companion matrix are computed by the standard $O(n^3)$ Hessenberg eigenvalue solver. This is an unfortunate choice, since a companion matrix and its QR -iterates are highly structured as we will see later on.

Recently, many authors have been intrigued by hidden rank properties in dense matrices. Basic theoretical results on rank structured matrices [3,4], quickly evolved to system solvers for these matrices [5,6] and to eigenvalue methods based on these matrix structures [7,8,9].¹

A companion matrix, which is Hessenberg, can be written as the sum of a unitary and a rank 1 matrix. Since both the unitary and the rank 1 matrix are structured rank matrices, one wishes to exploit these structures to come to an efficient method. The brute-force Hessenberg eigenvalue solver uses $O(n^3)$ operations for computing the whole spectrum, whereas an efficient method exploiting all involved matrix structures would only use $O(n^2)$ operations. The problem has already been studied extensively. Let us make a summary of the existing methods crediting the authors who have contributed to the development of these methods.

Independently of each other Bindel et al. and Bini et al. [10,11] proved that the QR -iterates H_k have $\{1,3\}$ -quasiseparable structure when $H_0 = U + \mathbf{u}\mathbf{v}^H$ is a Hessenberg matrix that is unitary plus rank one. Hence, each Hessenberg matrix H_k can be represented using $O(n)$ parameters. Based on this fact, several algorithms were developed that performed QR -iteration steps on the Hessenberg matrices H_k . These algorithms differ in the way the Hessenberg matrix is represented, in the way the Hessenberg as well as the unitary plus rank one structure is maintained and in the explicit or implicit way of performing each QR -iteration step.

In the paper by Bini et al., the following relationship between C_p of (1) and its inverse C_p^{-1} is used (when p_0 is different from zero):

$$C_p = C_p^{-H} + UV^H,$$

with $U_k, V_k \in \mathbb{C}^{n \times 2}$. The QR -iteration step with shift is implemented in an explicit way. The experiments described in the paper show that the implementation of the algorithm has numerical difficulties (overflow/underflow, no convergence).

In the manuscript of Bindel et al. a larger class of matrices is considered, more precisely, the symmetric, skew symmetric or orthogonal matrices plus a low rank modification. The algorithm uses a compression technique to maintain the structure.

In Bini et al. [12] an alternative algorithm was developed to solve the Hessenberg, unitary plus rank one case $H_k = U_k + \mathbf{u}_k\mathbf{v}_k^H$. The authors showed that the matrix U_k can be written as a product of three sequences of 2×2 unitary transformations as represented in (7). The presented method is an explicit QR -algorithm, in which special compression is needed to maintain the unitary structure.

As in the paper of Bindel et al., a bulge chasing procedure is designed in [13]. In contrast to the previous papers this is done for a single as well as for a double shift. The representation that is used for the Hessenberg matrix H_k in each step is its QR -factorization, i.e., $H_k = Q_kR_k$. A chasing technique is used to chase the implicitly created bulge. Unfortunately this increases the rank of a certain involved factor. Compression is therefore needed to maintain a compact representation.

¹ The literature related to this subject is much more extended than these few references.

All the previous papers handled the case $H = U + \mathbf{u}\mathbf{v}^H$, i.e., when the Hessenberg matrix H is the sum of a unitary U and rank one matrix $\mathbf{u}\mathbf{v}^H$. Only the manuscript of Bindel et al. considered the more general class of symmetric, skew symmetric, or unitary plus rank one case. In the paper [14] by Bini et al. of 2005, the matrices $A = (a_{i,j})_{i,j}$ considered have the following form:

$$\begin{aligned} a_{ii} &= d_i + z_i \bar{w}_i \\ a_{ij} &= u_i t_{ij}^\times \bar{v}_j \quad \text{when } i > j \\ a_{ij} &= \bar{u}_j \overline{t_{ji}^\times} v_i + z_i \bar{w}_j - \bar{z}_j w_i \quad \text{when } i < j, \end{aligned} \quad (2)$$

with $t_{ij}^\times = t_{i-1} t_{i-2} \cdots t_{j+1}$ and given vectors $\mathbf{u}, \mathbf{v}, \mathbf{t}, \mathbf{z}, \mathbf{w}, \mathbf{d}^2$. This set is called the set of generalized companion matrices. It includes the arrowhead matrices, comrade matrices (symmetric tridiagonal plus rank one), diagonal plus rank one matrices, Note that each matrix of this set is $\{1, 3\}$ -quasiseparable. The authors prove that this set is closed under the application of each step of the QR -algorithm. Each step of the QR -algorithm is performed in an explicit way, i.e., the QR -factorization of $A_k - \mu_k I$ is computed and then the two factors multiplied in reverse order are added to $\mu_k I$.

In this manuscript we will present a new manner of representing the companion matrix, which uses less parameters than the ones presented in the manuscripts above. Moreover, during the algorithm no compression is needed to maintain a low complexity algorithm. The representation is based on Givens transformations and easily admits an implicit QR -algorithm. The implicit version is based on structure restoring manipulations of the involved matrix, in order to maintain the same representation in each step of the QR -method. Once the single shift version is known, the double shift or multishift version is a straightforward extension.

The manuscript is organized as follows. In Section 2 some preliminary results are discussed. The relation between roots of polynomials and companion matrices is investigated and some comments on working with Givens transformations are given. In the following sections the different aspects to the development of the QR -methods are discussed. Section 3 discusses unitary plus rank one matrices, both the preservation of the structure and the representation are investigated. Section 4 and Section 5 discuss both the single and double shift techniques. Due to the special representation of the involved matrices the deflation is uncommon, Section 6 discusses how to perform deflation. Finally some details on the implementation and numerical experiments are presented.

Note: When finishing this manuscript, we became aware of another similar algorithm to compute the eigenvalues of a companion matrix. This method was explained by Boito at a workshop in Cortona [15].

2 Preliminary results

In this section we will firstly discuss the relation between polynomials and matrices. Secondly we will present some tools for working with graphical interpretations of Givens transformations.

2.1 Roots of polynomials

Suppose we are working with a polynomial $p(z)$ either in $\mathbb{R}_n[z]$ or $\mathbb{C}_n[z]$. The subscript n in $\mathbb{R}_n[z]$ and $\mathbb{C}_n[z]$ means that the considered polynomials are of degree less than or equal to n .

Without loss of generality we can assume p_n to be different from zero. Clearly the zeros of $p(z)$ equal the zeros of $p(z)/p_n$, hence we will assume in the remainder the polynomial $p(z)$ to be monic, i.e., with $p_n = 1$.

The companion matrix associated with a monic polynomial of degree n is given in Equation (1).

² This representation is called the quasiseparable representation.

Note 1 One can fairly easy see that the eigenvalues of the companion matrix coincide with the zeros of the associated polynomial $p(z)$, because $p(z) = \det(zI - C_p)$.


A theoretical, more broad context, not restricted to the standard monomial basis exists for proving the equality between eigenvalues of the associated matrix and roots of the polynomial. More information can be found for example in [2, 16].

2.2 Working with Givens transformations

Givens transformations are a powerful tool for working with structured matrices. E.g., the implicit QR -algorithm for Hessenberg matrices makes extensive use of Givens transformations [1] and especially for structured rank matrices [9] they are a valuable tool.

Also for structured rank matrices, Givens transformations play an important role. In this subsection we will discuss a graphical way of working with Givens transformations. Also interactions such as the “shift through” and “the fusion” of Givens transformations will be discussed.

The following graphical scheme represents the QR -factorization of a 6×6 Hessenberg matrix, based on Givens transformations. The scheme corresponds to $G_1 G_2 \dots G_5 R$. The matrix R is shown on the right and is clearly upper triangular. The Givens transformation G_5 is located above 1, G_4 in position 2 and so forth. One might consider this misleading, but one should think of the bottom line as a sort of timeline depicting which transformation needs to be applied first. Rewriting the formula we have $G_5^H \dots G_1^H H = R$. Hence G_1^H annihilates the first subdiagonal element, G_2^H the second and so forth.



In the remainder of the manuscript we will make extensive use of such schemes, with many more Givens transformations.


Givens transformations also interact with each other.

Lemma 1 Suppose two Givens transformations G_1 and G_2 are given:

$$G_1 = \begin{bmatrix} c_1 & -\bar{s}_1 \\ s_1 & \bar{c}_1 \end{bmatrix} \text{ and } G_2 = \begin{bmatrix} c_2 & -\bar{s}_2 \\ s_2 & \bar{c}_2 \end{bmatrix}.$$

Then we have that $G_1 G_2 = G_3$ is again a Givens transformation. We will call this the fusion of Givens transformations in the remainder of the text.

The proof is trivial. In our graphical schemes, we will depict this as follows:



The following lemma is very powerful and will give the possibility to interchange the order of Givens transformations and to obtain different patterns. Often Givens transformations of higher dimensions, say n , are considered. This means that the corresponding 2×2 Givens transformation is embedded in the identity matrix of dimension n , still changing only two rows when applied to the left.

Lemma 2 (Shift through lemma) Suppose three 3×3 Givens transformations \check{G}_1, \check{G}_2 and \check{G}_3 are given, such that the Givens transformations \check{G}_1 and \check{G}_3 act on the first two rows of a matrix, and \check{G}_2 acts on the second and third row (when applied on the left to a matrix).

Then there exist 3 Givens transformations \hat{G}_1, \hat{G}_2 and \hat{G}_3 such that

$$\check{G}_1 \check{G}_2 \check{G}_3 = \hat{G}_1 \hat{G}_2 \hat{G}_3,$$

where \hat{G}_1 and \hat{G}_3 work on the second and third row and \hat{G}_2 , works on the first two rows.

This result is well-known. The proof can be found in [17] and is simply based on the fact that one can factorize a 3×3 unitary matrix in different ways. Graphically we will depict this rearrangement as follows.

$$\begin{array}{c|ccc} \textcircled{1} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{2} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{3} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \hline & 3 & 2 & 1 \end{array} \quad \text{resulting in} \quad \begin{array}{c|ccc} \textcircled{1} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{2} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{3} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \hline & 3 & 2 & 1 \end{array}.$$

Of course there is also a variant in the other direction (from the right to the left scheme depicted by \curvearrowright).

Important to remark is that the fusion of Givens transformations can be considered as a special case of the shift through lemma. For our purposes we also need an additional operation, which is in fact also a special case of the shift through lemma, namely the scaled fusion.

Lemma 3 Suppose two Givens transformations³ G_1, G_2 and I_α , with $|\alpha| = 1$ are given,

$$G_1 = \begin{bmatrix} c_1 & -\bar{s}_1 \\ s_1 & \bar{c}_1 \end{bmatrix}, G_2 = \begin{bmatrix} c_2 & -\bar{s}_2 \\ s_2 & \bar{c}_2 \end{bmatrix} \text{ and } I_\alpha = \begin{bmatrix} 1 & 0 \\ 0 & \alpha \end{bmatrix}.$$

Then we have that $G_1 I_\alpha G_2 = G_3 I_\alpha = I_\alpha G_4$, with G_3 and G_4 again a Givens transformation. We will call this the scaled fusion of Givens transformations in the remainder of the text.

The proof involves straightforward computations, by embedding I_α in a 3×3 Givens transformation matrix, applying the shift through lemma and exploiting the fact that the matrix in which I_α is embedded is a diagonal matrix.

Graphically we will formulate this as:

$$\begin{array}{c|ccc} \textcircled{1} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{2} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \alpha & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \hline & 3 & 2 & 1 \end{array} \quad \text{resulting in} \quad \begin{array}{c|cc} \textcircled{1} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{2} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \alpha \text{ or } \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \hline & 2 & 1 \end{array} \quad \text{or} \quad \begin{array}{c|cc} \textcircled{1} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \begin{array}{c} \uparrow \\ \downarrow \end{array} \\ \textcircled{2} & \begin{array}{c} \uparrow \\ \downarrow \end{array} & \alpha \\ \hline & 2 & 1 \end{array}$$

Note 2 We note that the matrix I_α can be replaced by a matrix J_α of the following form:

$$J_\alpha = \begin{bmatrix} \alpha & 0 \\ 0 & 1 \end{bmatrix}.$$

leading to the following equations: $G_1 J_\alpha G_2 = G_3 J_\alpha = J_\alpha G_4$.

This scaled fusion is necessary since in the forthcoming QR -algorithm Givens transformations will tend to become diagonal. In the real case these Givens transformations equal the identity and hence there are no problems, but in the general complex setting they have a diagonal consisting of α and $\bar{\alpha}$ with $|\alpha| = 1$. Since these transformations will be surrounded by other Givens transformations some special operations are needed in order to manipulate them in an elegant way.

A final important operation is the shift through operation of length l .

Lemma 4 Suppose we have the following matrix product GWX , in which G denotes a Givens transformation acting on row 1 and 2. The matrices W and X are both unitary matrices consisting of a descending sequence of l Givens transformations. This means that both W and X consist of l successive Givens transformations. The i th Givens transformation G_i^W of W acts on row $i+1$ and $i+2$. The i th Givens transformation G_i^X of X acts on row i and $i+1$.


The matrix product GWX can then be rewritten as

$$GWX = \hat{W} \hat{X} \hat{G},$$

where \hat{G} is now a Givens transformation acting on row $l+1$ and $l+2$. The unitary matrices \hat{W} and \hat{X} are again descending sequences of l Givens transformations.

³ In fact these transformations are rotations. Throughout the manuscript we assume to be working with rotations. More information on Givens rotations can be found in [18].

Before proving the statement we will depict this graphically. Initially on the left we have the following scheme, in which G is denoted by a \times . The unitary matrices W and X are both sequences of 4 Givens transformations, W ranges from position 8 up to 5 and X ranges from position 4 up to 1.




On the right we see the resulting scheme which is equivalent to the first one. The Givens transformation \hat{G} is again denoted by a \times . The new sequence \hat{W} ranges from 9 upto 6 and \hat{X} ranges from 5 up to 2.


In the text we will refer to this as a shift through operation of length l . One can also perform a shift through operation from right to left. This will be referred to as a shift through operation of length l to the left.

Proof The proof consists of nothing else then a successive application of the shift through operation. Without loss of generality we assume $l = 4$. Hence we will prove the transition from left to right in Scheme (3).


Reordering permits the application of the shift through operation. This leads to the right scheme. We have marked another Givens transformation now with \times .



Rewriting gives us the possibility to apply again the shift through operation, resulting after rewriting in the right scheme.



Clearly the undesired Givens transformation moves down, creating two new sequences, starting in positions 9 and 8 and removing the sequences in positions 4 up to 1. Continuing this procedure, and applying two more times the shift through operation gives us Scheme (4).



This corresponds to the desired situation. Mathematically we obtain

$$\hat{W}\hat{X}\hat{G},$$

where \hat{W} denotes the sequence of Givens transformations ranging from position 9 up to position 6 and \hat{X} denotes the sequence ranging from position 5 up to position 2. The resulting Givens transformation in position 1 is \hat{G} .

Note 3 One can easily see that the shift through operation is a special case of Lemma 4, where the two sequences W and X are of length 1.

3 Unitary plus rank one matrices

In this manuscript we will develop QR -methods for computing eigenvalues of Hessenberg, unitary plus rank one matrices. This is an important class of matrices as the companion and the fellow matrix are both of this form. Hence, the problem of computing zeros of polynomials can be translated to the eigenproblem of a Hessenberg, unitary plus rank one matrix. Let us divide this section into small parts discussing some ingredients for developing an implicit QR -method.

3.1 Structure under a QR -step

Let us denote the Hessenberg matrix we are working with as follows

$$H = U + \mathbf{u}\mathbf{v}^H, \quad (5)$$

with H Hessenberg, U unitary and \mathbf{u} and \mathbf{v} two vectors.

Suppose we have a shift μ and we would like to perform a step of the QR -iteration on the matrix H . We have the following formulas

$$\begin{aligned} H - \mu I &= QR \\ \hat{H} = RQ + \mu I &= Q^H H Q. \end{aligned}$$

Applying now the similarity transformation on the terms of Equation (5) we obtain the following relations:

$$\begin{aligned} \hat{H} &= Q^H H Q = Q^H U Q + Q^H \mathbf{u}\mathbf{v}^H Q \\ &= \hat{U} + \hat{\mathbf{u}}\hat{\mathbf{v}}^H, \end{aligned}$$

with \hat{H} Hessenberg, \hat{U} unitary and $Q^H \mathbf{u}$ and $Q^H \mathbf{v}$ two vectors.

Hence the Hessenberg as well as the unitary plus rank 1 structure is preserved under a step of the QR -iteration. This is essential for developing an efficient implicit QR -method exploiting the matrix structure.

3.2 A representation for the unitary matrix

Since the unitary matrix in the splitting of the Hessenberg matrix is in general a dense matrix, we want to represent this matrix as efficiently as possible. Even though the matrix is full, it has a structured rank part.


Consider

$$U = H - \mathbf{u}\mathbf{v}^H,$$

since the matrix H is Hessenberg and has zeros below the subdiagonal, the matrix U needs to be of rank 1 below the subdiagonal. The matrix U is therefore unitary and the elements \boxtimes make up the structured rank part (rank 1) in this 6×6 example:

$$U = \begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \boxtimes & \times & \times & \times & \times & \times \\ \boxtimes & \boxtimes & \times & \times & \times & \times \\ \boxtimes & \boxtimes & \boxtimes & \times & \times & \times \\ \boxtimes & \boxtimes & \boxtimes & \boxtimes & \times & \times \end{bmatrix}.$$


We will construct the representation of this matrix in a graphical form. In the first scheme we already annihilated some elements in the last row of the matrix U , by a single Givens transformation acting on row five and six. Three elements are annihilated by a single Givens transformation, since they are coming from a low rank part in the matrix U .




Let us briefly explain the scheme above as they will be used throughout the entire manuscript. The right part of the scheme depicts a matrix, in which the elements \boxtimes satisfy the low rank constraints. The bracket above 1 indicates a single Givens transformation acting on rows **5** and **6** of the matrix. Mathematically, the scheme depicts

$$U = G_1 U_1, \quad (6)$$


where G_1 corresponds to the Givens transformation in position 1, and U_1 is the matrix on the right having extra zeros in the last row. In fact the scheme here presents a sort of factorization of the original unitary matrix U . The elements to be annihilated by the second Givens transformation are marked in the scheme with \otimes . Applying a second transformation gives us $U = G_1 G_2 U_2$. In the following scheme we see the matrix U_2 on the right, Givens transformation G_1 above 2 and G_2 above 1. Recall that the bottom line depicts a time line, indicating in which order the Givens transformation have to be applied on the matrix U_2 (see Subsection 2.2).




After the Givens transformation acting on row three and four is performed we have completely removed the low rank part. The matrix remaining on the right side is now a generalized Hessenberg matrix having two subdiagonals.



Peeling of the second subdiagonal, removing successively all elements marked with \otimes , will give us an extra descending sequence of Givens transformations. These Givens transformations can be found in positions 1 up to 4.



In the final step the remaining subdiagonal from the matrix on the right is removed. This will be done by a descending sequence of 5 Givens transformations. We obtain the following scheme.



In fact we have computed now a QR -factorization of the unitary matrix consisting of three sequences of Givens transformations: $U = QR$. Since the matrix R is upper triangular and also unitary, the matrix will be diagonal. Moreover, the final sequence of Givens transformations can be chosen such that the diagonal matrix R has all diagonal elements except the last one equal to 1, hence $R = I_\alpha$, with $|\alpha| = 1$. Graphically this is depicted as follows (we do not depict the ones on the diagonal of I_α):

$$\begin{array}{c|cccccccc}
 \textcircled{1} & & & & & & & & \\
 \textcircled{2} & & & & & & & & \\
 \textcircled{3} & & & & & & & & \\
 \textcircled{4} & & & & & & & & \\
 \textcircled{5} & & & & & & & & \\
 \textcircled{6} & & & & & & & & \\
 \hline
 & 9 & 8 & 7 & 6 & 5 & 4 & 3 & 2 & 1 & \alpha
 \end{array} \quad (8)$$

Denoting in Scheme (7) the sequence from 12 to 10 by V , the sequence from 9 to 6 by W and the sequence from 5 to 1 by X , we have factored the unitary matrix U as the product of three unitary matrices $U = VWXI_\alpha$. We use three different symbols for the unitary matrices, because the use of too many sub and superscripts would make the forthcoming implicit method difficult to read.

This will be the representation of the unitary matrix in the sum $H = U + \mathbf{u}\mathbf{v}^H$. Due to the relation between the low rank parts in the matrices $\mathbf{u}\mathbf{v}^H$ and U , we can decompose the vector \mathbf{u} as follows.

$$\begin{array}{c|cccc}
 \textcircled{1} & & & & \times \\
 \textcircled{2} & & & & \times \\
 \textcircled{3} & & & & \times \\
 \textcircled{4} & & & & 0 \\
 \textcircled{5} & & & & 0 \\
 \textcircled{6} & & & & 0 \\
 \hline
 & 9 & 8 & 7 & \\
 \end{array} \quad (9)$$

Important to remark is that the Givens transformations in position 9, 8 and 7 in Scheme (9) are exactly the same as the ones in the corresponding position of Scheme (8).

This means that we get the following equalities:

$$\begin{aligned}
 H &= U + \mathbf{u}\mathbf{v}^H \\
 &= VWXI_\alpha + \mathbf{u}\mathbf{v}^H \\
 &= V(WXI_\alpha + V^H\mathbf{u}\mathbf{v}^H) \\
 &= V(WXI_\alpha + \hat{\mathbf{u}}\mathbf{v}^H),
 \end{aligned}$$

where the vector $\hat{\mathbf{u}}$ has only the first three elements different from zero (this is the vector on the right in Scheme (9)).

The representation designed in this section is the one that will be used for developing an implicit QR -method for unitary plus low rank matrices.

In the next section we will discuss the implementation of the implicit single shifted QR -method based on this representation.

4 The single shift method

There exist several variants for performing an implicit QR -method on unitary plus low rank matrices. Unfortunately most of these algorithms are not capable of preserving both the unitary and the low rank structure. Numerical roundoff creates loss of the low rank structure and hence a form of rank compression is needed. This compression can create undesired results such as loss of accuracy. Often also an increase of parameters for representing the unitary plus low rank matrix is needed to deal with the loss of structure.

In this section we will discuss how to perform an implicit QR -step on the companion matrix, exploiting thereby the developed representation.

4.1 Implicit QR-methods

Assume we have $A - \mu I = QR$, with A , e.g., of tridiagonal, Hessenberg or semiseparable form. Perform now a unitary similarity transformation $\hat{A} = \hat{Q}^H A \hat{Q}$ with \hat{Q} having $\hat{Q}\mathbf{e}_1 = Q\mathbf{e}_1$ such that \hat{A} has the same structural constraints as assumed for A . Then one can easily prove that the matrix \hat{A} is essentially the same as the matrix coming from an explicit QR-step with shift μ performed on the matrix A .

An implicit QR-method consists of performing the transformation $\hat{Q}^H A \hat{Q}$ in such a manner that the full matrix Q is not needed beforehand, but is determined on the fly.

The global flow of an implicit method is the following.

- Compute an orthogonal transformation \hat{Q}_1 , such that $\hat{Q}_1^H (A - \mu I) \mathbf{e}_1 = \beta \mathbf{e}_1$.
- Perform a similarity transformation with \hat{Q}_1 on the matrix A : $\hat{Q}_1^H A \hat{Q}_1$.
- Perform now a new similarity transformation with the matrix \hat{Q}_2 , with $\hat{Q}_2 \mathbf{e}_1 = \mathbf{e}_1$, such that $\hat{Q}_2^H \hat{Q}_1^H A \hat{Q}_1 \hat{Q}_2$ has the same structural constraints as the original matrix A .

The first two items are often named the initialization step, whereas the third step is often called the chasing step, since in the Hessenberg and tridiagonal case, bulges are chased away to restore the Hessenberg or tridiagonal structure. This is also the subdivision used for the next two subsections.

4.2 Initialization

The implicit QR-step performed on a Hessenberg matrix is determined by the first Givens transformation G_1 (corresponds to \hat{Q}_1 from the previous subsection), such that

$$G_1^H (H - \mu I) \mathbf{e}_1 = \pm \| (H - \mu I) \mathbf{e}_1 \| \mathbf{e}_1.$$

We will now perform the similarity transformation $G_1^H H G_1$, exploiting the factorization of the matrix H

$$H = V (W X I_\alpha + \hat{\mathbf{u}} \mathbf{v}^H). \quad (10)$$

The idea is to obtain again a Hessenberg matrix, after the complete QR-step is finished. Moreover, we want the resulting Hessenberg matrix to satisfy the same representation as shown in (10). Of course after applying the first similarity transformation based on G_1 , the structure will not be exactly as desired. Hence extra Givens transformations will be constructed to restore the structure. This will be the chasing. Throughout the entire procedure we want the representation to remain as close as possible to the original representation of H .

Let us perform the first similarity transformation on the matrix $H = H_0 = V_0 (W_0 X_0 I_\alpha + \hat{\mathbf{u}}_0 \mathbf{v}_0^H)$. We obtain

$$\begin{aligned} H_1 &= G_1^H V_0 (W_0 X_0 I_\alpha + \hat{\mathbf{u}}_0 \mathbf{v}_0^H) G_1 \\ &= G_1^H V_0 (W_0 X_0 I_\alpha G_1 + \hat{\mathbf{u}}_0 \mathbf{v}_0^H G_1) \\ &= G_1^H V_0 (W_0 X_0 I_\alpha G_1 + \hat{\mathbf{u}}_0 \mathbf{v}_1^H), \end{aligned}$$


with $\mathbf{v}_1^H = \mathbf{v}_0^H G_1$. In the remainder of the derivations, all intermediate variables will be depicted with a tilde or a hat and all final values for representing H_1 will be $V_1, W_1, X_1, \mathbf{u}_1$ and \mathbf{v}_1 .

Since the Givens transformation G_1^H acts on the first two rows, and V_0 acts on the rows 3 up to 6, they commute and we obtain the following:

$$\begin{aligned} H_1 &= V_0 (G_1^H W_0 X_0 I_\alpha G_1 + G_1^H \hat{\mathbf{u}}_0 \mathbf{v}_1^H) \\ &= V_0 (G_1^H W_0 X_0 I_\alpha G_1 + \tilde{\mathbf{u}}_1 \mathbf{v}_1^H), \end{aligned}$$


with $\tilde{\mathbf{u}}_1 = G_1^H \hat{\mathbf{u}}_0$, having only the first three elements different from zero.

It seems that the matrix H_1 is already in the good form, except for the unitary matrix $G_1^H W_0 X_0 I_\alpha G_1$. Let us see how to change this matrix. The Givens transformation G_1^H is shown in position 10, position 9 up to 6 represents W_0 , position 5 up to 1 represents X_0 and position 0 is reserved for G_1 and I_α . The undesired Givens transformations are marked with a \times .



We will start by removing the Givens transformation in position 10. Neglecting the transformations in position 1 and 0, we can apply a shift through operation of length 4.

The left scheme depicts the result. One can easily fuse the Givens transformation in position 2 with the Givens transformation in position 1. Hence we have removed already 1 undesired transformation. This results in the right scheme.




This corresponds to the following relations:

$$H_1 = V_0 (\tilde{W}_1 \tilde{X}_1 I_\alpha G_1 + \tilde{\mathbf{u}}_1 \mathbf{v}_1^H),$$

the unitary matrices \tilde{W}_1 and \tilde{X}_1 are two descending sequences of transformations.

The only remaining undesired Givens transformation is G_1 in the middle of the formula. The idea is now to drag this Givens transformation completely through the other matrices such that it appears before the matrix V_0 and moreover acts on row 2 and row 3. It has to act on row 2 and row 3, so that choosing the next Givens transformation G_2 can be chosen to annihilate this transformation and such that it does not interfere with row 1 anymore.

Let us continue, by trying to remove also the second undesired Givens transformation, the one in position 0. One can fairly easily apply two times the shift through operation to obtain the following scheme. In fact one applies a shift through operation of length 2 to the left.



Now there is an unwanted transformation in position 10. Denote this transformation with \tilde{G}_1 . In formulas we obtain now:

$$H_1 = V_0 (\tilde{G}_1 \hat{W}_1 X_1 I_\alpha + \tilde{\mathbf{u}}_1 \mathbf{v}_1^H),$$

in which the sequences of Givens transformations \hat{W}_1 and X_1 have changed again.

This gives us the following relations:

$$\begin{aligned} H_1 &= V_0 \tilde{G}_1 (\hat{W}_1 X_1 I_\alpha + \tilde{G}_1^H \tilde{\mathbf{u}}_1 \mathbf{v}_1^H), \\ &= V_0 \tilde{G}_1 (\hat{W}_1 X_1 I_\alpha + \hat{\mathbf{u}}_1 \mathbf{v}_1^H), \end{aligned}$$

Two terms are of importance here: $V_0 \tilde{G}_1$ and $\tilde{G}_1^H \hat{\mathbf{u}}_1$. Clearly the vector $\hat{\mathbf{u}}_1$, will fill in one of its zeros. The vector will have now four nonzero elements.

The matrix product $V_0\tilde{G}_1$ does not essentially change the Givens structure of the matrix V_0 . A single fusion of two Givens transformations removes the undesired transformation \tilde{G}_1 , without destroying the structure of the matrix V_0 . In the scheme below, the undesired Givens transformation can be found in position 0, whereas the Givens transformations in positions 3 up to 1 make up the matrix V_0 .

$$\begin{array}{c|c}
 \textcircled{1} \\
 \textcircled{2} \\
 \textcircled{3} \\
 \textcircled{4} \\
 \textcircled{5} \\
 \textcircled{6} \\
 \hline
 3 \ 2 \ 1 \ 0
 \end{array}
 \rightarrow
 \begin{array}{c|c}
 \textcircled{1} \\
 \textcircled{2} \\
 \textcircled{3} \\
 \textcircled{4} \\
 \textcircled{5} \\
 \textcircled{6} \\
 \hline
 3 \ 2 \ 1
 \end{array}
 \quad (13)$$

The resulting matrix \tilde{V}_1 has the same Givens pattern as the original matrix V_0 .

We obtain the following relations:

$$H_1 = \tilde{V}_1 (\hat{W}_1 X_1 I_\alpha + \hat{\mathbf{u}}_1 \mathbf{v}_1^H). \quad (14)$$

But this is not sufficient yet. We want to obtain a similar factorization as of the original matrix H . Hence, the four nonzero elements in the vector $\hat{\mathbf{u}}_1$ need to be transformed into three nonzero elements.

To do so, we need to rewrite Equation (14). Denote the Givens transformation in Scheme (12) in position 9, with G^l , this means $G^l \hat{W}_1^l = \hat{W}_1$, we will drag this Givens transformation out of the brackets on the left. We denote all changed variables with l to clearly indicate that one extra Givens has moved to the left. Until we will move the Givens back inside the brackets we will indicate this on the affected elements.

$$\begin{aligned}
 H_1 &= \tilde{V}_1 G^l \left(\hat{W}_1^l X_1 I_\alpha + G^{lH} \hat{\mathbf{u}}_1 \mathbf{v}_1^H \right) \\
 &= \tilde{V}_1^l \left(\hat{W}_1^l X_1 I_\alpha + \hat{\mathbf{u}}_1^l \mathbf{v}_1^H \right)
 \end{aligned}$$

in which $G^{lH} \hat{\mathbf{u}}_1 = \hat{\mathbf{u}}_1^l$ still has four elements different from zero, the matrix \tilde{V}_1^l is now a sequence having one more Givens transformation than \tilde{V}_1 .

Graphically $\tilde{V}_1 \hat{W}_1 = \tilde{V}_1^l \hat{W}_1^l$ is depicted as follows.


$$\begin{array}{c|c|c}
 \textcircled{1} \\
 \textcircled{2} \\
 \textcircled{3} \\
 \textcircled{4} \\
 \textcircled{5} \\
 \textcircled{6} \\
 \hline
 7 \ 6 \ 5 \ 4 \ 3 \ 2 \ 1
 \end{array}
 =
 \begin{array}{c|c|c}
 \textcircled{1} \\
 \textcircled{2} \\
 \textcircled{3} \\
 \textcircled{4} \\
 \textcircled{5} \\
 \textcircled{6} \\
 \hline
 7 \ 6 \ 5 \ 4 \ 3 \ 2 \ 1
 \end{array}
 \quad (15)$$

The vertical line depicts the splitting of the product of the Givens transformations. Clearly the Givens transformation in position 4 on the left scheme has moved to the left part in the right scheme. In the left scheme, the matrix \tilde{V}_1 consists of the Givens in position 7 to 5, \hat{W}_1 contains the Givens from 4 to 1, whereas in the right scheme the new \tilde{V}_1^l consists of the Givens transformations in position 7 up to 4 and \hat{W}_1^l contains the Givens transformations in position 3 up to 1.

Construct now a Givens transformation acting on row 3 and 4 of the vector $\hat{\mathbf{u}}_1^l$ and annihilating the element in position 4. Let us denote this Givens transformation by \hat{G}_1 such that $\hat{G}_1^H \hat{\mathbf{u}}_1^l = \mathbf{u}_1^l$ with \mathbf{u}_1^l having only the first three elements different from zero. Let us plug this in into the formulas.


$$\begin{aligned}
 H_1 &= \tilde{V}_1^l \left(\hat{W}_1^l X_1 I_\alpha + \hat{G}_1 \hat{G}_1^H \hat{\mathbf{u}}_1^l \mathbf{v}_1^H \right) \\
 &= \tilde{V}_1^l \left(\hat{G}_1 \hat{G}_1^H \hat{W}_1^l X_1 I_\alpha + \hat{G}_1 \mathbf{u}_1^l \mathbf{v}_1^H \right) \\
 &= \tilde{V}_1^l \hat{G}_1 \left(\left(\hat{G}_1^H \hat{W}_1^l \right) X_1 I_\alpha + \mathbf{u}_1^l \mathbf{v}_1^H \right)
 \end{aligned}$$

Let us take a closer look now at $\hat{G}_1^H \hat{W}_1^l$. Graphically this is depicted in the scheme below. The Givens transformations from 1 up to 3 make up the matrix \hat{W}_1^l and the Givens \hat{G}_1^H is shown in position 4. Applying one fusion, removes the Givens transformation in position 4.



This results in a sequence of Givens transformations denoted as $\hat{G}_1^H \hat{W}_1^l = W_1^l$.

The matrix product $\tilde{V}_1^l \hat{G}_1^H$ is of the following form. Applying the shift through lemma gives us the following scheme.



This scheme can be written as $G_2 V_1^l$, in which V_1^l incorporates the Givens from position 1 to 4 and G_2 is shown in the fifth position. This Givens transformation G_2 will determine the next chasing step. First we need however to bring the Givens transformation in position 1 again inside the brackets.

Plugging all of this into the equations above gives us:

$$H_1 = G_2 V_1^l \left(W_1^l X_1 I_\alpha + \mathbf{u}_1^l \mathbf{v}_1^H \right).$$

Rewriting now the formula above by bringing the Givens transformation \tilde{G}_1^l in position 1 of the matrix V_1^l inside the brackets does not change the formulas significantly. We obtain:

$$H_1 = G_2 V_1 \left(W_1 X_1 I_\alpha + \mathbf{u}_1 \mathbf{v}_1^H \right), \text{ with } V_1 = V_1^l \tilde{G}_1^H, \quad W_1 = \tilde{G}_1 W_1^l, \quad \mathbf{u}_1 = \tilde{G}_1 \mathbf{u}_1^l.$$

The equation depicted here is almost of the desired structure, the matrix H_1 is almost in Hessenberg form, except for a bulge in position (2, 1). The representation is also almost in the correct form, only the Givens transformation G_2 is undesired. We will remove this transformation (removing thereby also the bulge in H_1), by applying another similarity transformation with G_2 .

4.3 The chasing

We have now performed the initial step, which was the most difficult step. We will continue now the algorithm, in an implicit way. We will only depict how to remove the Givens transformation G_2 , the procedure can be continued in a similar fashion.

Since we want our matrix H to become again of unitary plus low rank form obeying the designed representation, we want to remove the disturbance G_2 . Performing the unitary similarity transformation with this Givens transformation removes the transformation in the left, but creates an extra transformation on the right. We obtain the following:

$$\begin{aligned} H_2 &= G_2^H H_1 G_2 \\ &= G_2^H G_2 V_1 \left(W_1 X_1 I_\alpha + \mathbf{u}_1 \mathbf{v}_1^H \right) G_2 \\ &= V_1 \left(W_1 X_1 I_\alpha G_2 + \mathbf{u}_1 \mathbf{v}_1^H G_2 \right) \\ &= V_1 \left(W_1 X_1 I_\alpha G_2 + \mathbf{u}_1 \mathbf{v}_2^H \right), \end{aligned}$$

where $\mathbf{v}_2^H = \mathbf{v}_1^H G_2$.

Similarly as in the initial step we can drag G_2 through W_1 and X_1 , this means applying a shift through operation of length 2. We obtain $W_1X_1G_2 = \tilde{G}_2W_2X_2$. The reader can create the schemes to see on which rows the transformation is performed. This gives us

$$\begin{aligned} H_2 &= V_1 (\tilde{G}_2W_2X_2I_\alpha + \mathbf{u}_1\mathbf{v}_2^H), \\ &= V_1\tilde{G}_2 (W_2X_2I_\alpha + \tilde{G}_2^H\mathbf{u}_1\mathbf{v}_2^H). \end{aligned}$$

Since the Givens transformation \tilde{G}_2^H acts on row 4 and row 5, $\tilde{G}_2^H\mathbf{u}_1 = \mathbf{u}_1$ (\mathbf{u}_1 has only the first three elements different from zero). Applying a final shift through operation for $V_1\tilde{G}_2$ we obtain $G_3V_2 = V_1\tilde{G}_2$ giving us (with $\mathbf{u}_2 = \mathbf{u}_1$):

$$H_2 = G_3V_2 (W_2X_2I_\alpha + \mathbf{u}_2\mathbf{v}_2^H).$$

We have performed now a step of the chasing method since the Givens transformation G_3 has shifted down one position w.r.t. the Givens transformation G_2 .

The chasing procedure can be continued in a straightforward manner. Unfortunately this approach does not allow us to determine the last Givens transformation G_{n-1} . Let us show what goes wrong and how to solve this problem in the next subsection.

4.4 The last Givens transformation

Suppose we have performed step $n-3$ and we have the following situation:

$$H_{n-3} = G_{n-2}V_{n-3} (W_{n-3}X_{n-3}I_\alpha + \mathbf{u}_{n-3}\mathbf{v}_{n-3}^H).$$

Performing the similarity transformation determined by G_{n-2} results in

$$\begin{aligned} H_{n-2} &= V_{n-3} (W_{n-3}X_{n-3}I_\alpha G_{n-2} + \mathbf{u}_{n-3}\mathbf{v}_{n-3}^H G_{n-2}) \\ &= V_{n-3} (W_{n-3}X_{n-3}I_\alpha G_{n-2} + \mathbf{u}_{n-2}\mathbf{v}_{n-2}^H). \end{aligned}$$

The Givens transformation G_{n-2} works on rows $n-2$ and $n-1$, hence we obtain the left graphical schemes for the first term of the equation $W_{n-3}X_{n-3}I_\alpha G_{n-2}$. Applying once the shift through operation moves the Givens transformation in position 0 to position 6, depicted on the right.

Unfortunately we cannot shift the Givens transformation through anymore, a single fusion can be applied and the undesired transformation vanishes. We obtain the following equation:

$$H_{n-2} = V_{n-2} (W_{n-2}X_{n-2}I_\alpha + \mathbf{u}_{n-2}\mathbf{v}_{n-2}^H).$$

The representation is of the desired form, but we know that the matrix H_{n-2} is not yet of the correct form, this matrix is of Hessenberg form with an extra bulge in position $(n, n-1)$. This final Givens transformation can hence not be determined implicitly anymore.

We determine this Givens transformation explicitly by computing the matrix vector product $H_{n-2}\mathbf{e}_{n-2}$, determine now G_{n-1} such that $\mathbf{e}_n^H G_{n-1}^H H_{n-2} \mathbf{e}_{n-2} = 0$.

The final similarity transformation results in:

$$H_{n-1} = G_{n-1}^H V_{n-2} (W_{n-2}X_{n-2}I_\alpha G_{n-1} + \mathbf{u}_{n-2}\mathbf{v}_{n-2}^H G_{n-1}).$$

Applying a fusion for the product $G_{n-1}^H V_{n-2}$ and a special fusion for the product $X_{n-2}I_\alpha G_{n-1}$ results in

$$H_{n-1} = V_{n-1} (W_{n-1}X_{n-1}I_\alpha + \mathbf{u}_{n-1}\mathbf{v}_{n-1}^H),$$


which is a new Hessenberg matrix, i.e., a sum of a unitary and a rank 1 matrix, represented using the desired representation. This completes the single shifted QR -step.

cannot monitor convergence in position $H(n, n - 1)$ in this fashion. This element has to be computed separately. Also the element $H(2, 1)$ cannot be monitored in this fashion, we will come back to this in Subsection 6.5.

Instead of subdividing the problem and trying to design two separate representations, one for the upper left and one for the lower right part, we will simply continue working with this representation. We know that the line indicates the splitting in the Hessenberg matrix, hence we want to work on the part above and below the line separately. This will be depicted in forthcoming subsections.

Continuing performing steps of the QR -algorithm on this representation, will result in a quasi upper triangular matrix R , which is represented as the sum of a unitary and a rank one matrix. This quasi upper triangular matrix R has the eigenvalues of the original matrix on its block diagonal.


The Givens pattern of the final unitary matrix (omitting the row scaling factors) has the following form. We see that during the reducing procedure we will transform the Givens transformations from the inner sequence into unitary diagonal matrices.



Let us illustrate how to continue working on both parts of the Hessenberg matrix. The part above and the part below the dashed line. The shift and the initial Givens transformation can be computed by explicitly computing some columns of the Hessenberg matrix. This will cost at most $O(n)$ operations.


6.2 Upper part

Let us see how the algorithm changes in case of the following scheme (only the unitary factor is represented). The original Hessenberg matrix has a zero in position $H(6, 5)$ and we want to perform a QR -step on the top part, i.e., the first 5 rows. Since the standard QR procedure applied on a reducible Hessenberg matrix breaks down when reaching the zero subdiagonal element, we would also expect a certain break during the chasing procedure in this case.




The first initial disturbing Givens transformation is applied in an identical way as before the deflation. Nothing new here. The next Givens similarity transformation acting on rows and columns 2 and 3 can be performed as in the standard case, just like the transformation on rows and columns 3 and 4. The next similarity transformation acting on row and columns 4 and 5 is a special one. In the standard case, this Givens transformation does not create another bulge. Hence, the matrix is again of Hessenberg form and the QR procedure breaks down. One cannot compute the following Givens transformation acting on row 5 and 6 anymore.

Let us see how this translates to our case. Suppose we have the following situation, just before applying the Givens similarity transformation. Since only the unitary matrix is essentially involved in this chasing procedure we will not depict any operations on the rank one structure.




Performing the Givens similarity transformation will remove the undesired Givens transformation in position 14 and create another one in the second position. In the left scheme this situation is depicted. A shift through operation is depicted and the result is shown in the right scheme.




In the right scheme we see that we can apply another shift through operation, due to the fact that the Givens transformation in position 7 crosses the dashed line. In Section 4.4 this Givens transformation fused with another one and vanished. Hence, in this case, we do not need to determine the last Givens transformation by explicitly computing the fourth column. The final step is therefore essentially different from the final step in Section 4.4.

So when dragging the remaining undesired Givens transformation completely to the left we obtain the following scheme.



When performing the next transformation, the algorithm should stop since the element $H(6,5)$ in the Hessenberg matrix is zero. This means that at the end of this similarity transformation, all undesired Givens transformations should be removed. Let us see how this translates towards the representation.




After the similarity transformation we obtain the left scheme. A fusion is depicted in the right scheme.

As a result we obtain again a similar represented Hessenberg matrix. This process can easily be repeated until another Givens transformation in the second sequence is close enough to diagonal form. The right scheme shows the application of a scaled fusion (scaled due to the parameter δ_6) and hence all undesired transformations are gone.


6.3 Bottom part

Applying a QR -step on the bottom part only differs from Section 4 in the way the first Givens transformation is performed. The final Givens transformation again has to be determined by computing the penultimate column.

Assume we have the situation as depicted in Scheme (21). Since we would like to perform a QR -step on the bottom part, the initial similarity transformation will act immediately on rows and columns 4 and 5. Suppose the transformation is determined and the similarity transformation is performed. This results in the following scheme (the Givens transformations are applied directly below the dashed line).



The Givens transformation on the left marked with \times is shifted through such that it appears in the middle. The next scheme clearly shows that only one fusion needs to be applied to remove the undesired Givens transformation from the left.



As a result we obtain sort of the standard situation. The Givens transformation on the right is not troublesome, dragging it through completely to the left as done before, it will come out acting on rows 5 and 6. Hence it can be used for determining the next Givens transformation.

Remark 1 Important to notice is the fact that the vector \mathbf{u} from the rank 1 part is not influenced in the above procedure anymore. Only the vector \mathbf{v} has to be updated when performing a Givens transformation on the right. This means that performing the initial Givens transformation is a much cheaper procedure than in the upper part.

6.4 More general cases

Of course when running the algorithm cases like the following will occur:

$$H = \left[\begin{array}{c|c|c} H_1 & \times & \times \\ \hline 0 & H_2 & \times \\ \hline 0 & 0 & H_3 \end{array} \right]. \tag{22}$$

To apply a QR -step on the matrix H_2 , a combination of techniques from the previous two sections needs to be made. Performing the initial step is described in Subsection 6.3, whereas performing the last Givens is done like in Subsection 6.2. QR -steps performed on these middle blocks are cheaper than the ones on H_1 and H_3 .

6.5 Special cases

Two cases were not treated before, namely the possibilities of $H(n, n - 1)$ and $H(2, 1)$ being zero. The appearance of a zero element in both these positions is not reflected in the representation. Hence one has to compute these elements explicitly and monitor their behavior. They can be

The Givens transformation marked with \times can easily be brought inside the brackets, without creating extra undesired elements in the vector \mathbf{u} . Applying a final fusion removes this Givens transformation. This brings us back in the standard case, and one can continue with the standard chasing procedures.

7 Numerical experiments

7.1 Balancing strategy

Because of the possible variation in magnitude of coefficients of the polynomial, normally a pre-processing step is applied to the matrix as, e.g., in [13]. In fact this is a balancing strategy, meaning that the variation in magnitude of the elements of the matrix is reduced. Generally DCD^{-1} is computed with D a diagonal matrix. To obtain again a unitary plus low rank matrix, the diagonal elements have to be chosen as follows $d_i = \beta^i$ (for some β). This results in

$$DCD^{-1} = \beta \begin{bmatrix} 0 & 0 & \dots & 0 & \frac{p_0}{\beta^n} \\ 1 & 0 & \dots & 0 & \frac{p_1}{\beta^{n-1}} \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & \ddots & \ddots & 0 & \frac{p_{n-2}}{\beta^2} \\ 0 & 0 & \dots & 1 & \frac{p_{n-1}}{\beta} \end{bmatrix} = \beta \hat{C}.$$

For an appropriate β , the eigenvalues of matrix C are then obtained by computing the eigenvalues of matrix \hat{C} and then multiply them by the factor β . Empirically, they found the following criterion for choosing β to be useful. Choose β such that

$$\text{Range}\{|\hat{c}_1|, |\hat{c}_2|, \dots, |\hat{c}_n|, 1\} := \frac{\max\{|\hat{c}_1|, |\hat{c}_2|, \dots, |\hat{c}_n|, 1\}}{\min\{|\hat{c}_1|, \dots, |\hat{c}_n|, 1\}}$$

is as small as possible. With $\hat{c}_i = a_i/\beta^i$.

7.2 The implementation

The algorithm using a double shift as explained in Section 5 was implemented in Matlab⁵. To avoid the fact that the Givens transformations for the last two columns have to be computed explicitly, we do one iteration step with the two shifts equal to zero, applied to the polynomial $z^2 p(z)$ which has clearly two additional zeros for $z = 0$. This iteration step brings these two zeros as eigenvalues of the bottom 2×2 block of the (extended) companion matrix. So, after this iteration step, one can deflate these two artificial roots and no explicit computation of the Givens transformations based on the last two columns has to be computed again. We will refer to our implementation as **VanBarel**. This software can be requested from the authors. We wanted to compare this implementation with all other methods for which we were able to obtain an implementation. We had access to the following implementations: **Delvaux** corresponds to a Matlab-implementation of the method proposed in [19], **Bini** denotes the Matlab-implementation for the method from [12], and **eig no bal** and **eig bal** correspond to the eig function of Matlab without and with balancing, respectively (with balancing is meant an extra balancing inside the function **eig**).

7.3 The experiments

Two different errors will be considered in the experiments. The maximum eigenvalue error (Max Eig Err) is defined as the distance between the sets $\lambda(C)$ and $\tilde{\lambda}(C)$:

$$\text{dist}(\lambda(C), \tilde{\lambda}(C)) = \max\left\{ \max_{\tilde{\lambda} \in \tilde{\lambda}(C)} \|\tilde{\lambda} - \lambda(C)\|, \max_{\lambda \in \lambda(C)} \|\lambda - \tilde{\lambda}(C)\| \right\},$$

⁵ Matlab is a registered trademark of The MathWorks, Inc.

where $\|\lambda - \tilde{\lambda}(C)\| = \min_{\tilde{\lambda} \in \tilde{\lambda}(C)} |\lambda - \tilde{\lambda}|$, $\tilde{\lambda}(C)$ are the eigenvalues computed with our algorithm and $\lambda(C)$ the eigenvalues computed by the Matlab function `eig` (assume that these eigenvalues are exact). The maximal componentwise relative error in the coefficients of the polynomials (Max Coeff Err) is defined as :

$$\max_i \left(\frac{|p_i - \tilde{p}_i|}{|p_i|} \right),$$

where \tilde{p}_i, p_i are the coefficients of the polynomial obtained by the computed and exact eigenvalues, respectively. The coefficients \tilde{p}_i are computed in multiple precision.

Experiment 1: In this experiment, we investigate the robustness of the different methods with respect to the choice of the scaling factor β described in Subsection 7.1.

To do so, we consider the following monic polynomials each of degree 20 where almost all of these polynomials have a large variation in the magnitude of their coefficients[20, 13]:

1. Wilkinson polynomial with roots: k with $k = 1, \dots, 20$.
2. Monic polynomial with roots: $[-2.1 : 0.2 : 1.7]$.
3. Monic polynomial with roots: 2^k with $k = -10, \dots, 9$.
4. Scaled Wilkinson polynomial with roots: $\frac{k}{n}$ with $k = 1, \dots, 20$.
5. Reversed Wilkinson polynomial with roots: $\frac{1}{k}$ with $k = 1, \dots, 20$.
6. Polynomial with well separated roots: $\frac{1}{2^k}$ with $k = 1, \dots, 20$.
7. Polynomial $p(z) = (20!) \sum_{k=0}^{20} \frac{z^k}{k!}$.
8. Polynomial $p(z) = z^{20} + z^{19} + \dots + z + 1$.

For the polynomials 1) – 6) the roots are known, for 7) – 8) this is not the case. Because of the wide range of the coefficients of the polynomial, i.e. for the Wilkinson polynomial $|A_i| \in [1, 10^{20}]$, balancing, which is discussed in Subsection 7.1, is of high importance. Therefore different balancing factors β are applied to the matrix C . Figure 1-2, contains the result of varying the balancing factor $\beta = 2^j, j = -10, \dots, 10$ for the eight polynomials of degree 20. The result of the maximal componentwise error on the coefficients is shown in logarithmic scale for the five different methods: VanBarel, Delvaux, Bini, and `eig no bal` and `eig bal`.

The figures show that, when the methods Delvaux, Bini, and `eig no bal` give a (rather) accurate solution, which is not always the case, it is only for some specific scaling factors β . For the method VanBarel the choice of the scaling parameter β is not so crucial to obtain an accurate solution. As long as the scaling factor is not too large, the method works. This shows that this method is very robust with respect to the other $O(n^2)$ methods that we are aware of. The function `eig bal` of Matlab always works well. Note that this is an $O(n^3)$ method as well as `eig` with no balancing.

In Table 1, the maximal componentwise error on the coefficients (Max Coeff Err) for the optimal choice of the parameter β is shown for the methods Bini, Delvaux and VanBarel. This table shows that the method VanBarel gives the most accurate results when scaling is necessary.

Experiment 2: To check that the computational complexity is of order $O(n^2)$, we constructed random polynomials of degree n , by computing the coefficients $A_l \in \mathbb{C}, l = 0, \dots, n$ by using randomly distributed complex numbers. The matrix sizes were chosen as $n = 25 \cdot 2^{i-1}$, with $i = 1, \dots, 6$. We assume that no balancing ($\beta = 1$) is necessary. For each of these sizes n , 5 samples were considered.

Figure 3(a) shows the ratio T_{i+1}/T_i averaged over the 5 samples for the method Delvaux, Bini and VanBarel. As you can see this is of the order $O(n^2)$. Figure 3(b) shows the average number of QR iterations per cycle. For the method Delvaux the average is between 2 and 3. For the method Bini it is a little bit more, between 5 and 6. For VanBarel, it is between 1 and 2.

In Table 2, the maximal componentwise error on the coefficients (Max Coeff Err) and the maximal eigenvalue error (Max Eig Err) are shown for the methods Delvaux, Bini and VanBarel.

	Bini		Delvaux		Van Barel	
	Max Coeff Err	β	Max Coeff Err	β	Max Coeff Err	β
1	10^{-10}	2^{-9}	10^{-11}	2^3	10^{-14}	2^{-9}
2	10^{-10}	1	10^{-11}	1	10^{-12}	2^{-1}
3	10^9	2^{-1}	10^{-6}	2^{-2}	10^{-11}	2^{-9}
4	10^{-9}	2^{-2}	10^{-11}	2^{-1}	10^{-14}	2^{-6}
5	10^{-11}	2^{-3}	10^{-11}	2^{-3}	10^{-13}	2^{-8}
6	10^9	2^{-20}	10^{-6}	2^{-12}	10^{-11}	2^{-19}
7	10^{-11}	2^3	10^{-13}	2^3	10^{-14}	2^2
8	10^{-14}	1	10^{-14}	1	10^{-13}	1

Table 1. Experiment 1: Maximum Coefficient error for the best balancing for the polynomials 1-8.

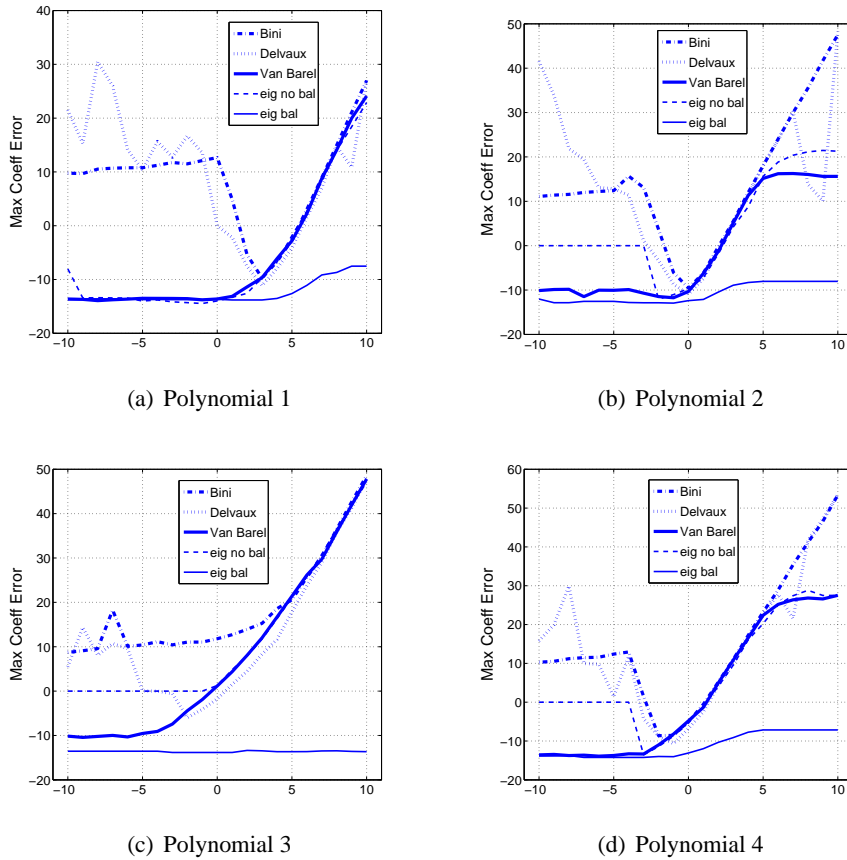


Figure 1. Experiment 1: Looking for the optimal balancing factor $\beta = 2^j$ for polynomials 1 – 4.

For the errors it can be seen that if n increases the errors increase slightly. Looking at the maximal eigenvalue error the results obtained with method *Delvaux* are slightly better.

8 Conclusion

In this paper an algorithm to compute the eigenvalues of a monic polynomial is described. This is done by computing the eigenvalues of the corresponding companion matrix. We showed how the QR algorithm can be adjusted such that the algorithm exploits unitary plus rank one structure of the matrices involved and how the matrices can be represented by $O(n)$ parameters using sequences of Givens transformations. The numerical performance of the algorithm was demonstrated by means of numerical experiments showing the good accuracy and efficiency of the algo-

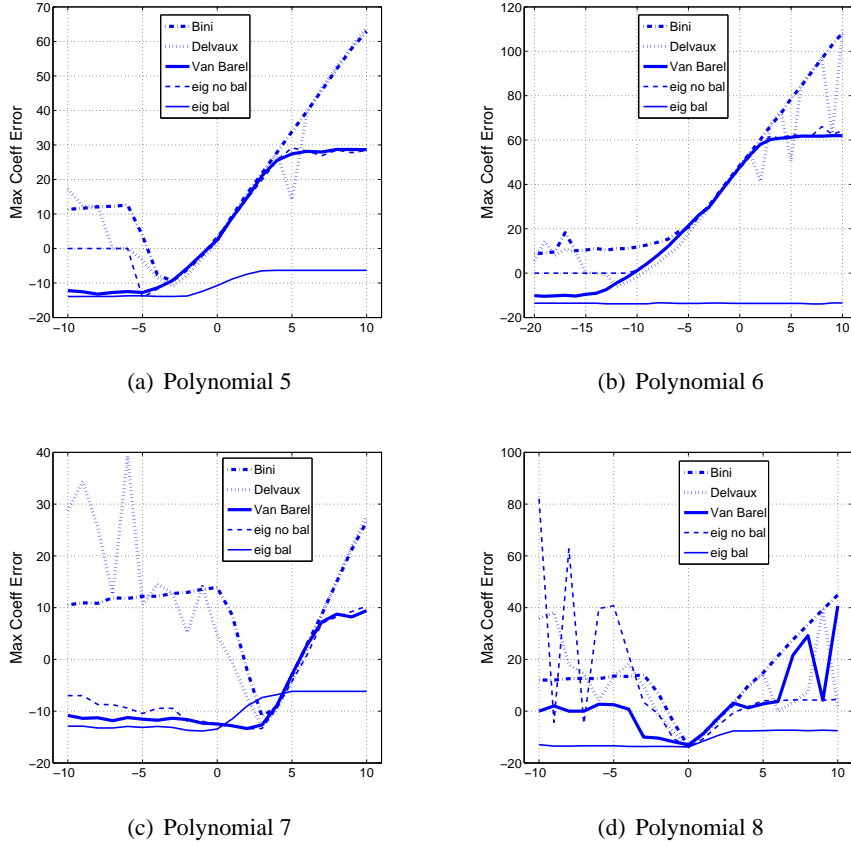


Figure 2. Experiment 1: Looking for the optimal balancing factor $\beta = 2^j$ for polynomials 5 – 8.

n	Delvaux		Bini		VanBarel	
	Max Coeff Err	Max Eig Err	Max Coeff Err	Max Eig Err	Max Coeff Err	Max Eig Err
25	2.141×10^{-14}	4.276×10^{-15}	3.287×10^{-14}	1.226×10^{-14}	1.758×10^{-13}	2.210×10^{-14}
50	4.4378×10^{-14}	4.650×10^{-15}	1.494×10^{-13}	2.150×10^{-14}	1.342×10^{-12}	3.834×10^{-14}
100	1.734×10^{-13}	1.106×10^{-14}	6.401×10^{-13}	7.639×10^{-14}	9.349×10^{-12}	1.292×10^{-13}
200	5.689×10^{-13}	3.774×10^{-14}	3.570×10^{-12}	3.551×10^{-13}	8.111×10^{-11}	4.093×10^{-13}
400	3.636×10^{-12}	1.468×10^{-13}	2.016×10^{-11}	9.245×10^{-12}	7.671×10^{-10}	1.534×10^{-12}
800	2.970×10^{-11}	4.262×10^{-13}	1.134×10^{-10}	4.367×10^{-11}	7.075×10^{-9}	8.901×10^{-12}
1600	2.022×10^{-10}	3.329×10^{-12}	5.605×10^{-10}	3.565×10^{-10}	6.507×10^{-8}	2.961×10^{-11}

Table 2. Experiment 2: numerical results for methods Delvaux, Bini and VanBarel.

rithm. Especially the robustness with respect to the scaling parameter β is as far as we know not observed in any other $O(n^2)$ method.

In the near future, we plan to make an implementation in C or C++ to study the efficiency of the method compared to standard eigenvalue solvers, e.g., from LAPACK. We also want to prove the numerical stability of our algorithm.

Acknowledgments

We would like to thank Katrijn Frederix for helping us to generate the figures and tables in the section on the numerical experiments.




Figure 3. Experiment 2: a) time ratio, b) average of the number of iterations per cycle.

References

1. G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, Maryland, USA, third edition, 1996.
2. S. Barnett. *Polynomials and Linear Control Systems*. Monographs and Textbooks in Pure and Applied Mathematics. Marcel Dekker, Inc., New York, USA, 1983.
3. R. Vandebril, M. Van Barel, and N. Mastronardi. A note on the representation and definition of semiseparable matrices. *Numerical Linear Algebra with Applications*, 12(8):839–858, October 2005.
4. M. Fiedler and T. L. Markham. Completing a matrix when certain entries of its inverse are specified. *Linear Algebra and its Applications*, 74:225–237, 1986.
5. Y. Eidelman and I. C. Gohberg. On a new class of structured matrices. *Integral Equations and Operator Theory*, 34:293–324, 1999.
6. S. Delvaux and M. Van Barel. A QR -based solver for rank structured matrices. *SIAM Journal on Matrix Analysis and Applications*, 30(2):464–490, 2008.
7. R. Vandebril, M. Van Barel, and N. Mastronardi. An implicit QR -algorithm for symmetric semiseparable matrices. *Numerical Linear Algebra with Applications*, 12(7):625–658, 2005.
8. Y. Eidelman, I. C. Gohberg, and V. Olshevsky. The QR iteration method for Hermitian quasiseparable matrices of an arbitrary order. *Linear Algebra and its Applications*, 404:305–324, July 2005.
9. S. Delvaux and M. Van Barel. The explicit QR-algorithm for rank structured matrices. Technical Report TW459, Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200A, 3000 Leuven (Heverlee), Belgium, May 2006.
10. D. A. Bini, F. Daddi, and L. Gemignani. On the shifted QR iteration applied to companion matrices. *Electronic Transactions on Numerical Analysis*, 18:137–152, 2004.
11. D. Bindel, S. Chandrasekaran, J. W. Demmel, D. Garmire, and M. Gu. A fast and stable nonsymmetric eigensolver for certain structured matrices. Technical report, Department of Computer Science, University of California, Berkeley, California, USA, May 2005.
12. D. A. Bini, Y. Eidelman, L. Gemignani, and I. C. Gohberg. Fast QR eigenvalue algorithms for Hessenberg matrices which are rank-one perturbations of unitary matrices. *SIAM Journal on Matrix Analysis and Applications*, 29(2):566–585, 2007.
13. S. Chandrasekaran, M. Gu, J. Xia, and J. Zhu. A fast QR algorithm for companion matrices. *Operator Theory: Advances and Applications*, 179:111–143, 2007.
14. D. A. Bini, L. Gemignani, and V. Y. Pan. Fast and stable QR eigenvalue algorithms for generalized companion matrices and secular equations. *Numerische Mathematik*, 100(3):373–408, 2005.
15. D. A. Bini, P. Boito, Y. Eidelman, L. Gemignani, and I. C. Gohberg. A fast QR eigenvalue method for a class of structured matrices, 2008.
16. R. Vandebril, M. Van Barel, and N. Mastronardi. *Matrix Computations and Semiseparable Matrices, Volume II: Eigenvalue and Singular Value Methods*. Johns Hopkins University Press, 2008.
17. R. Vandebril, M. Van Barel, and N. Mastronardi. A parallel QR-factorization/solver of structured rank matrices. Technical Report TW474, Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200A, 3000 Leuven (Heverlee), Belgium, October 2006.
18. D. Bindel, J. W. Demmel, W. Kahan, and O. A. Marques. On computing Givens rotations reliably and efficiently. *ACM Transactions on Mathematical Software*, 28(2):206–238, June 2002.
19. S. Delvaux, K. Frederix, and M. Van Barel. An algorithm for computing the eigenvalues of block companion matrices. Technical report, Department of Computer Science, Katholieke Universiteit Leuven, 2008. In preparation.
20. A. Edelman and H. Murakami. Polynomial roots from companion matrix eigenvalues. *Mathematics of Computation*, 64(210):763–776, April 1995.