Jacobi-like Algorithms for Eigenvalue Decomposition of a Real Normal Matrix Using Real Arithmetic

B. B. Zhou and R. P. Brent

Computer Sciences Laboratory
The Australian National University
Canberra, ACT 0200, Australia

{*bing,rpb*}*@cslab.anu.edu.au*

Abstract

In this paper we introduce a method for designing efficient Jacobi-like algorithms for eigenvalue decomposition of a real normal matrix. The algorithms use only real arithmetic and achieve ultimate quadratic convergence. A theoretical analysis is conducted and some experimental results are presented.

1 Introduction

A real matrix A is said to be normal if it satisfies the equation

$$AA^T = A^T A$$

where A^T is the transpose of matrix A. This kind of matrix has the property that it can be reduced to a diagonal form through unitary similarity transformations

$$QA\bar{Q}^T = D$$

where Q is unitary and D is diagonal. In this paper we consider how to design a scalable Jacobi-like algorithm for efficient parallel computation of the eigenvalue decomposition of a real normal matrix over the real field. Real normal matrices are generalisations of real symmetric matrices. A real symmetric matrix is normal, but a real normal matrix is not necessarily symmetric. We shall focus our attention on the unsymmetric case although the method to be described applies to both cases.

The standard sequential method for eigenvalue decomposition of this kind of matrix is the QR algorithm. When massively parallel computation is considered, however, the parallel version of the QR algorithm for solving unsymmetric eigenvalue problems may not be very efficient because the algorithm is sequential in nature and not scalable.

One alternative to the QR method is a Jacobi method. Jacobi-based algorithms have recently attracted a lot of attention as they have a higher degree of potential parallelism. The Jacobi method, though originally designed for symmetric eigenvalue problems, can be extended to solve eigenvalue problems for unsymmetric normal matrices [3, 7]. However, one problem is that we have to use complex arithmetic even for real-valued normal matrices. Complex operations are expensive and should be avoided if possible. A quaternion-Jacobi method was recently introduced [4]. In this method a 4×4 symmetric matrix can be reduced to a 2×2 block diagonal form using one orthogonal similarity transformation. This method can also be extended to compute eigenvalues of a general normal matrix. However, the problems are that the original matrix has to be divided into a sum of a symmetric matrix and a skew-symmetric matrix, which is not an orthogonal operation, and that the algorithm cannot be used to solve the eigenvalue problem of near-normal matrices. Another parallel Jacobi-like algorithm, named the RTZ (which stands for Real Two-Zero) algorithm, was also proposed recently [5]. This method uses real arithmetic and orthogonal similarity transformations. It is claimed that quadratic convergence can be obtained when computing eigenvalues of a real near-normal matrix with real distinct eigenvalues. However, a serious problem with this method is that the process may not converge if the matrix has complex eigenvalues.

In this paper we describe a method for designing efficient Jacobi-like algorithms for eigenvalue decomposition of a real normal matrix. The designed algorithms use only real arithmetic and orthogonal similarity transformations. The theoretical analysis and experimental results show that ultimate quadratic convergence can be achieved for general real normal matrices with distinct eigenvalues.

Since the RTZ algorithm uses a similar idea to our method, an analysis of the RTZ algorithm is presented in Section 2. Our method is described in Section 3. In that section a theoretical analysis of convergence is also presented. Some experimental results are given in Section 4. Section 5 is the conclusion.

2 An Analysis of the RTZ Algorithm

We now show that when applying the RTZ algorithm, we can only obtain at best a linear convergence rate if the matrix has complex eigenvalues.

In the following a 4 × 4 matrix

$$A = \begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}, \qquad (1)$$

or its block form

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix},$$

is used to show how the convergence rate may be affected when applying the RTZ algorithm.

The basic idea for computing the eigenvalue decomposition of a 4 × 4 matrix using the RTZ algorithm is as follows: The first leading diagonal element a_{11} in A_{11} is chosen together with the first row in A_{12}, the first column in A_{21} and the whole of A_{22} to form a 3 × 3 matrix, that is,

$$A_1 = \begin{pmatrix} a_{11} & a_{13} & a_{14} \\ a_{31} & a_{33} & a_{34} \\ a_{41} & a_{43} & a_{44} \end{pmatrix}.$$

It is known that any real 3 × 3 matrix has at least one real eigenvalue. An eigenvector associated with a real eigenvalue of the above matrix can be obtained and used to generate a Householder matrix which is then applied to update the matrix A so that two off-diagonal elements a_{31} and a_{41} are annihilated. After that the second leading diagonal element a_{22} is chosen (together with A_{22}, the second row in A_{12} and the second column in A_{21}) and another 3 × 3 matrix

$$A_2 = \begin{pmatrix} a_{22} & a_{23} & a_{24} \\ a_{32} & a_{33} & a_{34} \\ a_{42} & a_{43} & a_{44} \end{pmatrix}$$

is formed. A Householder matrix is generated from an eigenvector (associated with a real eigenvalue) of this 3 × 3 matrix so that two other lower triangular off-diagonal elements a_{32} and a_{42} are eliminated through a similarity transformation. However, this destroys the zeros introduced previously. The two leading diagonal elements are thus chosen alternately and the process continues until all the elements in A_{21} are small enough to be considered as zero.

In the following discussions ϵ denotes a small positive number close to zero. If certain elements of a matrix are written as ϵ, we mean that the values of these elements are small and of the same order of ϵ, but they are not necessarily the same.

If the above RTZ procedure converges, all the elements in A_{21} and A_{12} become ϵ after a few iterations and each 3×3 matrix will thus have the same form as

$$B = \begin{pmatrix} b_{11} & \epsilon & \epsilon \\ \epsilon & b_{22} & b_{23} \\ \epsilon & b_{32} & b_{33} \end{pmatrix}. \qquad (2)$$

Assume that B is close to a normal block-diagonal matrix and written as

$$B = D + \epsilon F$$

where the elements in F satisfy $|f_{ij}| < 1$ and D has a form

$$D = \begin{pmatrix} d_{11} & 0 & 0 \\ 0 & d_{22} & d_{23} \\ 0 & d_{32} & d_{33} \end{pmatrix}.$$

If the eigenvalues of D are well separated, it is easy to prove according to the perturbation theory [10] that there exists an eigenvalue of B which satisfies the equation

$$|d_{11} - \lambda| = O(\epsilon) \qquad (3)$$

An eigenvector v associated with λ will also satisfy

$$\|v - v'\| = O(\epsilon) \qquad (4)$$

for v' an eigenvector associated with d_{11}, a real eigenvalue of D. Therefore, we have the following lemma.

Lemma 1 *Let a* 3 × 3 *matrix have a form as that in (2) and satisfy all the above mentioned conditions. Then there exists a real eigenvalue* λ*, such that*

$$|d_{11} - \lambda| = O(\epsilon).$$

An eigenvector associated with this eigenvalue has a form

$$v = (v_1 \;\; \epsilon \;\; \epsilon)^T. \qquad (5)$$

Proof. The first part of the lemma is directly obtained from (3). Let $v' = \left(v_1', v_2', v_3'\right)^T$ be an eigenvector associated with d_{11}. Then

$$(D - d_{11}I)v' = 0,$$

or

$$\begin{cases} d_1 v_1' + 0v_2' + 0v_3' = 0 \\ 0v_1' + d_2 v_2' + d_{23} v_3' = 0 \\ 0v_1' + d_{32} v_2' + d_3 v_3' = 0 \end{cases}$$

where $d_i = d_{ii} - d_{11}$. Since the eigenvalues of D are well separated, the determinant of the coefficient matrix from the second and the third equations will not be equal to zero. Thus v_2' and v_3' must be zero. We then have

$$v' = \left(v_1', 0, 0\right)^T.$$

From (4), therefore, v will have the form as that in (5). $\square$

Since the RTZ algorithm uses Householder transformations, it can be seen in the following that it is crucial for eigenvectors generated at each step to have the form as that in (5) in order to obtain ultimate quadratic convergence.

We now show how the lower triangular off-diagonal norm is affected when the generated Householder matrix is applied to update certain rows and columns of the matrix.

Assume that a Householder vector v is chosen based on vector $b = (b_1\ b_2\ b_3)^T$, that is, $v = (v_1\ b_2\ b_3)^T$ for $v_1 = b_1 + sign(b_1)||b||$ and $||b|| = \sqrt{b_1^2 + b_2^2 + b_3^2}$. The Householder matrix is then obtained as $H = I - 2vv^T/v^Tv$.

Suppose that the values of b_2 and b_3 are of order ϵ, but b_1 is "large". We have

$$\begin{aligned}\frac{1}{||b||} &= \frac{1}{sign(b_1)b_1\sqrt{1+(b_2^2+b_3^2)/b_1^2}} \\ &\approx \frac{1}{sign(b_1)b_1\sqrt{1+\epsilon^2}} \\ &\approx \frac{1}{sign(b_1)b_1}(1-\epsilon^2).\end{aligned}$$

It is easy to verify that the Householder matrix has the structure

$$H \approx \begin{pmatrix} -1+\epsilon^2 & \epsilon & \epsilon \\ \epsilon & 1-\epsilon^2 & \epsilon^2 \\ \epsilon & \epsilon^2 & 1-\epsilon^2 \end{pmatrix}. \tag{6}$$

When using this matrix to update (left multiply) a column vector $(x_1\ 0\ 0)^T$, it is easy to see that, if x_1 is of order ϵ, the zero elements will become of order ϵ^2. If x_1 is "large", however, after the updating, the zero elements will become $O(\epsilon)$, rather than $O(\epsilon^2)$.

Suppose that all the elements in the off-diagonal blocks become small after a few sweeps and that the last two elements in the first column have been eliminated by the immediate previous iteration in which a_{11} and the corresponding 3×3 matrix are chosen. The updated matrix then has a form

$$A = \begin{pmatrix} a_{11} & a_{12} & \epsilon & \epsilon \\ a_{21} & a_{22} & \epsilon & \epsilon \\ 0 & \epsilon & a_{33} & a_{34} \\ 0 & \epsilon & a_{43} & a_{44} \end{pmatrix}.$$

Now we choose a_{22} and the corresponding 3×3 matrix as

$$\begin{pmatrix} a_{22} & \epsilon & \epsilon \\ \epsilon & a_{33} & a_{34} \\ \epsilon & a_{43} & a_{44} \end{pmatrix}$$

and apply the RTZ procedure again to annihilate the last two elements in the second column of matrix A. Assume that this 3×3 matrix has the same properties as the matrix described in Lemma 1. The eigenvector will be in the same form as that in (5). Therefore, the generated Householder matrix is of the form like (6).

The last three elements in the first column of matrix A will be affected in the updating procedure using this Householder matrix. If A_{11} has two real eigenvalues, we may reasonably assume that a_{21} and a_{12} are of order ϵ. (If not, an orthogonal similarity transformation can be applied to annihilate a_{21} without using complex arithmetic as A_{11} has two real eigenvalues.) When this Householder matrix is applied, the values of the zero elements in the first column will be of order ϵ^2. Quadratic convergence may then be achieved. If A_{11} (as well as A_{22}) has two complex eigenvalues, however, the value of a_{21} may no longer be small. The values of the zero elements in the first column will be increased to $O(\epsilon)$ as we discussed above. The convergence rate is thus only linear at best.

In the above we only considered a very simple problem, that is, the 4×4 case. When the matrix size is much larger, the slow local convergence can significantly affect the global convergence. In many cases the process may not even converge.

3 Our Method

To simplify our discussion, we assume that the matrix size is even. The basic idea of our method is described as follows: A real normal matrix is first divided into blocks. To avoid using complex arithmetic the size of each block is chosen to be 2×2 so that each pair of conjugate complex eigenvalues can be grouped into the same block. A sequence of orthogonal similarity transformations is then applied to annihilate the off-diagonal blocks in the lower triangle and this process continues until all the lower triangular off-diagonal blocks are considered as zero. The basic structure of the algorithm is depicted in Fig. 1.

In Fig. 1 NITN denotes the number of lower triangular off-diagonal blocks, or number of iterations in a sweep, which is equal to $n(n-2)/8$ for n the size of the problem. A counter NZCONT is used to check how many off-diagonal blocks in the lower triangle are zero. If this is equal to NITN, the process stops. The cyclic ordering is used for sequential computation. If we consider each block as a single element, The action of this algorithm is just the same as that of a nonblocked algorithm. Therefore, any existing efficient parallel orderings, for example, those in [1, 2, 8, 9, 11, 12], can be adopted to form an efficient parallel algorithm for solving the problem.

It should be noted that in the figure we do not show how to generate the orthogonal transformation matrices Q. This is a very important issue relating to performance. We shall show

```
NITN = n(n-2)/8
NZCONT = 0
REPEAT UNTIL NZCONT = NITN
  DO i = 1, n, 2
    DO j = j+2, n, 2
      B = ( A_ii  A_ij )
          ( A_ji  A_jj )
      IF (||A_ji|| not zero) THEN
          1. Find an orthogonal matrix Q
          such that B is reduced to a block
          triangular form through a similarity
          transformation.
          2. Update the corresponding rows
          and columns of A using Q.
      ELSE
          NZCONT = NZCONT + 1
      END IF
    END DO
  END DO
  IF (NZCONT ≠ NITN) NZCONT = 0
END REPEAT
```

Figure 1: The basic structure of our Jacobi-like algorithm.

that the quadratic convergence is achieved only through properly choosing the transformation matrices during the computation. In the following we discuss the convergence property of the algorithm.

Lemma 2 *If a (real) normal matrix A is block triangular, it must be block diagonal, where the blocks on the main diagonal of A are either 1×1, or 2×2.*

The proof of the above lemma is simple and thus omitted (to save space). From this lemma we see that, if we can find an orthogonal matrix which reduces a real normal matrix to a block upper triangular form through a similarity transformation, the transformed matrix must be block diagonal. Therefore, it is possible just to try to reduce a real normal matrix to a block upper triangular form using orthogonal similarity transformations (over the real field). The matrix will eventually be reduced to block diagonal if the process converges.

We show in the following three lemmas that, if the orthogonal transformation matrices are chosen properly and if the process converges for a normal matrix with distinct eigenvalues, the convergence rate will ultimately be quadratic.

In the following discussion $||A||$ denotes the Frobenius norm of A.

Lemma 3 *If a normal matrix A is divided into a block matrix, then the main diagonal blocks have the following property:*

$$||A_{ii}A_{ii}^T - A_{ii}^T A_{ii}|| \le \sum_{k \ne i}(||A_{ki}||^2 + ||A_{ik}||^2) \qquad (7)$$

where A_{ij} is the block in the i^{th} row and j^{th} column of A.

Proof. Let $C = AA^T$ and $C' = A^T A$. We have $C_{ii} = \sum_k A_{ik}A_{ik}^T$ and $C'_{ii} = \sum_k A_{ki}^T A_{ki}$. Since matrix A is normal,

$$C_{ii} - C'_{ii} = \sum_k A_{ik}A_{ik}^T - \sum_k A_{ki}^T A_{ki} = 0,$$

or

$$A_{ii}A_{ii}^T - A_{ii}^T A_{ii} = \sum_{k \ne i}(A_{ki}^T A_{ki} - A_{ik}A_{ik}^T).$$

Thus

$$\begin{aligned} ||A_{ii}A_{ii}^T - A_{ii}^T A_{ii}|| &\le \sum_{k \ne i} ||A_{ki}^T A_{ki} - A_{ik}A_{ik}^T|| \\ &\le \sum_{k \ne i}(||A_{ki}^T A_{ki}|| + ||A_{ik}A_{ik}^T||) \\ &\le \sum_{k \ne i}(||A_{ki}||^2 + ||A_{ik}||^2). \end{aligned}$$

□

It can be seen from this lemma that each block on the main diagonal will be very close to a normal matrix if the norm of each off-diagonal block is small, that is,

$$||A_{ii}A_{ii}^T - A_{ii}^T A_{ii}|| = O(\epsilon^2)$$

if max$(||A_{ij}||) \le \epsilon$ for $i \ne j$.

The next lemma shows that, when the lower off-diagonal block of a 2×2 block matrix is annihilated through an orthogonal similarity transformation, the norm of its upper off-diagonal block will also be decreased if this block matrix is close to normal.

Lemma 4 *Assume that a 2×2 block matrix B, with the size of each block being 2×2, is close to a normal matrix and has the property*

$$||BB^T - B^T B|| = O(\eta) \qquad (8)$$

where η is a small positive number, and that B has four distinct nonzero eigenvalues λ_k which satisfy

$$min\,|\lambda_k| > c_1 > 0 \qquad (9)$$

and

$$\left|1 - \frac{\lambda_i}{\lambda_j}\right| > c_2 > 0 \qquad (10)$$

for $i \neq j$. If B is reduced to a block triangular form through an orthogonal similarity transformation, that is,

$$Q^T BQ = D \qquad (11)$$

where Q is a real orthogonal matrix and

$$D = \begin{pmatrix} D_{11} & D_{12} \\ 0 & D_{22} \end{pmatrix}, \qquad (12)$$

we then have

$$||D_{12}|| = O(\eta).$$

Proof. Since $B = QDQ^T$ and $B^T = QD^TQ^T$, we have

$$||BB^T - B^TB|| = ||Q(DD^T - D^TD)Q^T||.$$

Thus

$$||DD^T - D^TD|| = O(\eta). \qquad (13)$$

It is easy to see that the second block of the first column in $DD^T - D^TD$ is $D_{22}D_{12}^T - D_{12}^TD_{11}$. We know that $||D_{12}|| \neq 0$. Otherwise, B will be a normal matrix according to Lemma 2 and then the problem is trivial. Since $||D_{12}|| \neq 0$ and the eigenvalues are nonzero and distinct, $D_{22}D_{12}^T - D_{12}^TD_{11}$ must also be nonzero. It is easy to see, from (13), that its norm should be of order η, that is,

$$||D_{22}D_{12}^T - D_{12}^TD_{11}|| = O(\eta). \qquad (14)$$

Since B has four nonzero eigenvalues, both D_{11} and D_{22} have full rank. The following inequality holds:

$$||D_{12}^T - D_{22}^{-1}D_{12}^TD_{11}|| \leq ||D_{22}^{-1}|| ||D_{22}D_{12}^T - D_{12}^TD_{11}||.$$

According to one of our assumptions, that is, the smallest eigenvalue of matrix B is greater than c_1, $||D_{22}^{-1}||$ will be less than $1/c_1$. We thus have

$$||D_{12}^T - D_{22}^{-1}D_{12}^TD_{11}|| = O(\eta).$$

Let

$$D_{11} = Q_1R_1Q_1^H$$

and

$$D_{22}^{-1} = Q_2R_2Q_2^H$$

be the eigenvalue decompositions of D_{11} and D_{22}^{-1} for Q_1 and Q_2 orthonormal and R_1 and R_2 upper triangular, and define

$$Q_2^HD_{12}^TQ_1 = E.$$

we then have

$$\begin{aligned} ||E - R_2ER_1|| &= ||Q_2(E - R_2ER_1)Q_1^H|| \\ &= ||D_{12}^T - D_{22}^{-1}D_{12}^TD_{11}|| \\ &= O(\eta). \end{aligned}$$

Let

$$E = \begin{pmatrix} e_{11} & e_{12} \\ e_{21} & e_{22} \end{pmatrix},$$

$$R_1 = \begin{pmatrix} r_{11}^{(1)} & r_{12}^{(1)} \\ 0 & r_{22}^{(1)} \end{pmatrix}$$

and

$$R_2 = \begin{pmatrix} r_{11}^{(2)} & r_{12}^{(2)} \\ 0 & r_{22}^{(2)} \end{pmatrix}.$$

Expanding $G = E - R_2ER_1$, we have the following four elements:

$$\begin{aligned} g_{11} &= (1 - r_{11}^{(2)}r_{11}^{(1)})e_{11} - r_{12}^{(2)}r_{11}^{(1)}e_{21}, \\ g_{12} &= (1 - r_{11}^{(2)}r_{22}^{(1)})e_{12} - r_{11}^{(2)}r_{12}^{(1)}e_{11} - \\ &\quad r_{12}^{(2)}r_{12}^{(1)}e_{21} - r_{12}^{(2)}r_{22}^{(1)}e_{22}, \\ g_{21} &= (1 - r_{22}^{(2)}r_{11}^{(1)})e_{21} \\ g_{22} &= (1 - r_{22}^{(2)}r_{22}^{(1)})e_{22} - r_{22}^{(2)}r_{12}^{(1)}e_{21}. \end{aligned}$$

Since $r_{ii}^{(1)}$ and $1/r_{ii}^{(2)}$ are eigenvalues of B and all the above elements should be of order η, it is then easy to verify by using (9) and (10) that all the elements in E must be of order η. Therefore, we obtain

$$\begin{aligned} ||D_{12}|| &= ||Q_2EQ_1^H|| \\ &= ||E|| \\ &= O(\eta). \end{aligned}$$

□

Assume that a normal matrix is divided into blocks of size 2 × 2 and that the norm of each off-diagonal blocks is of order ϵ. We further assume that each A_{ij} in (7) is a 2 × 2 block submatrix. Thus the norm of each off-diagonal submatrix A_{ij} for $i \neq j$ must also have an order of ϵ. From Lemma 3 we may obtain

$$||A_{ii}A_{ii}^T - A_{ii}^TA_{ii}|| = O(\epsilon^2).$$

After an orthogonal similarity transformation which reduces A_{ii} to a block upper triangular matrix, the norm of the upper off-diagonal block in A_{ii} should have the same order as $||A_{ii}A_{ii}^T - A_{ii}^TA_{ii}||$ according to the above lemma. Therefore, it will be of order ϵ^2.

Since the off-diagonal norm of A_{ii} is reduced from $O(\epsilon)$ to $O(\epsilon^2)$ and the norms of other A_{ij}s in (7) are not affected during the updating procedure, we thus obtain a steady decrease in off-diagonal norm during the computation. In the following we show that, if the orthogonal transformation matrices are chosen properly, ultimate quadratic convergence can be achieved.

Explicitly write matrices B and Q described in Lemma 4 as 2×2 matrices, that is,

$$B = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix}$$

and

$$Q = \begin{pmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{pmatrix}.$$

and assume that $||B_{12}|| = O(\epsilon)$ and $||B_{21}|| = O(\epsilon)$. Then we have the following lemma.

Lemma 5 *Assume that the eigenvalues of B_{11}, B_{22}, D_{11} and D_{22} are γ_{i1}, γ_{i2}, λ_{i1} and λ_{i2} for $i = 1, 2$, respectively. If $\lambda_{ij} > c_1 > 0$ for $i, j = 1, 2$ and*

$$\left|1 - \frac{\gamma_{jk}}{\lambda_{il}}\right| > c_2 > 0$$

for $l \neq k$, the norms of both Q_{12} and Q_{21} in the generated orthogonal matrix will then be of order ϵ, that is,

$$||Q_{12}|| = O(\epsilon)$$

and

$$||Q_{21}|| = O(\epsilon).$$

Proof. From (11) we have

$$\begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} \begin{pmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{pmatrix} =$$

$$\begin{pmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{pmatrix} \begin{pmatrix} D_{11} & D_{12} \\ 0 & D_{22} \end{pmatrix}. \quad (15)$$

Thus

$$B_{21}Q_{11} + B_{22}Q_{21} = Q_{21}D_{11},$$

or

$$B_{21}Q_{11} = Q_{21}D_{11} - B_{22}Q_{21},$$

or

$$B_{21}Q_{11}D_{11}^{-1} = Q_{21} - B_{22}Q_{21}D_{11}^{-1} \quad (16)$$

and

$$B_{11}Q_{12} + B_{12}Q_{22} = Q_{11}D_{12} + Q_{12}D_{22},$$

or

$$(B_{12}Q_{22} - Q_{11}D_{12})D_{22}^{-1} = Q_{12} - B_{11}Q_{12}D_{22}^{-1}. \quad (17)$$

Since $||B_{21}|| = O(\epsilon)$, $||B_{12}|| = O(\epsilon)$ and $||D_{12}|| = O(\epsilon^2)$ from Lemma 4 with $\eta = \epsilon^2$, the norms of the left-side of the equations in (16) and (17) must be of order ϵ. We thus have

$$||Q_{21} - B_{22}Q_{21}D_{11}^{-1}|| = O(\epsilon)$$

and

$$||Q_{12} - B_{11}Q_{12}D_{22}^{-1}||. = O(\epsilon)$$

With the same technique used in Lemma 4, we may easily obtain $||Q_{21}|| = O(\epsilon)$ and $||Q_{12}|| = O(\epsilon)$. □

The condition set in Lemma 5 implies that implicit permutation on rows and columns between blocks is not allowed when applying orthogonal similarity transformations to B. If this condition is satisfied, both $||Q_{12}||$ and $||Q_{21}||$ will be of order ϵ. Using this orthogonal transformation matrix to update a vector $(v_1\ v_2\ 0\ 0)^T$, the zero elements in the vector will be of order ϵ^2 if both v_1 and v_2 are of order ϵ.

We now summarize the results obtained from the above three lemmas. Under the conditions that the process converges and the matrix to be decomposed has distinct eigenvalues, we have proved that, when annihilating the lower off-diagonal block of submatrix B in Fig 1, the norm of the upper-diagonal block is reduced from $O(\epsilon)$ to $O(\epsilon^2)$. Thus the off-diagonal norm of the matrix has a steady decreasing during the computation. We have also proved that the zero blocks in the lower triangle will be of order ϵ^2 through updating stages if the orthogonal transformation matrices are chosen properly. If a reasonable and systematic ordering is also applied, the off-diagonal norm of the matrix will be decreased from $O(\epsilon)$ to $O(\epsilon^2)$ after each sweep of computation. Thus ultimate quadratic convergence is obtained.

It should be noted that in proving the above lemmas we assumed that B in Fig. 1 is generated from the adjacent blocks, that is, we set $j = i + 1$. If this condition is not satisfied, we may use permutation matrices. Since permutation matrices are orthogonal, the results will not be affected.

4 Experimental Results

Through our discussion we do not give any particular method for obtaining orthogonal transformation matrices. Any method which may generate orthogonal transformation matrices satisfying the condition set in Lemma 5 can be used. One scheme is as follows: When the eigenvalues of either A_{ii} or A_{jj} in submatrix B in Fig. 1 are real, we can apply the RTZ procedure so that a local quadratic convergence is achieved and the generated orthonormal matrix will have the form described in Lemma 5. When both eigenvalues of A_{ii} and A_{jj} are complex, we can simply use the QR procedure to reduce B to a block triangular form. With a combination of these two procedures we can obtain an efficient Jacobi-like algorithm.

The algorithm described above has been implemented on a sequential machine. The stopping criteria used in our experiment is the same as that in EISPACK [6], that is, an off-diagonal element a_{ij} is considered as zero if $|a_{ij}| \leq (|a_{ii}| + |a_{jj}|) * \epsilon_{mach}$ for ϵ_{mach} the machine precision. A block is considered as zero if all the elements in it are considered as zero and the computation stops if all the lower triangular off-diagonal blocks are considered as zero.

The test matrices used in the experiment are generated by computing QDQ^T, where Q is an orthogonal matrix and D is a block diagonal matrix. Each block in D is of size 2×2, that is,

$$D_{ii} = \begin{pmatrix} d_1 & d_2 \\ d_3 & d_4 \end{pmatrix}$$

where the four elements are positive random numbers smaller than one. When $d_2 = -d_3$ and $d_1 = d_4$, we have two complex eigenvalues $d_1 \pm id_2$. Otherwise, we set $d_2 = d_3 = 0$ for two real eigenvalues.

In our experiment we choose four different matrices. The first matrix has distinct real eigenvalues, the second one has half of its eigenvalues real and the other half complex. and the third contains distinct complex eigenvalues. The fourth matrix is similar to the second matrix except it has three real eigenvalues of multiplicity four.

Some experimental results are presented in Tables 1 and 2. Table 1 gives the lower block triangular norms after each sweep for computing eigenvalues of the matrices of size 40×40. It can be seen that ultimate quadratic convergence can be obtained for decomposing a matrix with distinct eigenvalues, especially for matrices with only real eigenvalues where the convergence is better than quadratic. We can also see from both tables that is will take a few more sweeps to converge for a matrix with complex eigenvalues. The more complex eigenvalues, the slower the speed though the quadratic convergence property is maintained in the extreme case when all eigenvalues are complex.

For a matrix with repeated eigenvalues, however, quadratic convergence can only be observed at a few middle sweeps. This result holds for problems with different sizes. This phenomenon needs to be studied further.

Sweep	Lower Block Triangular Norm	
	Matrix 1	Matrix 2
0	2.002949651	2.491182060
1	0.9298566189	1.658914630
2	0.3230377708	0.9562437960
3	8.148755450D-02	0.3041956127
4	2.272463954D-02	5.732517299D-02
5	5.984053421D-04	2.915087621D-03
6	1.583025639D-08	1.729400600D-06
7	4.396127327D-16	1.214607626D-13
8	4.396127327D-16	2.996108412D-16
9		2.996108412D-16

Sweep	Lower Block Triangular Norm	
	Matrix 3	Matrix 4
0	2.785430894	2.162572292
1	1.975427113	1.622702464
2	1.424859846	1.137339662
3	0.9068788996	0.5313679428
4	0.5020392998	0.1254582065
5	0.1631977468	8.449010583D-03
6	2.753458690D-02	2.825586595D-05
7	2.454913944D-04	1.090020755D-08
8	2.197732988D-08	4.401366214D-10
9	2.682697256D-16	3.441166540D-14
10	2.408358962D-16	2.180213396D-16
11	2.408358962D-16	2.180213396D-16

Table 1: Sweeps and lower blocks triangular norms for 40×40 matrices.

5 Conclusions

In this paper we first gave an analysis of the RTZ algorithm. We showed that at most a linear convergence rate can be obtained when this algorithm is applied to decompose a normal matrix which has complex eigenvalues. However, our analysis indicates that quadratic convergence may be achieved if the given matrix has only real eigenvalues. Thus the algorithm is still useful for eigenvalue decomposition of a normal or near normal matrix if the matrix has only real eigenvalues.

We then proposed a method for designing efficient Jacobi-like algorithms for eigenvalue decomposition of a real normal matrix. Both theoretical analysis and experimental results show that ultimate quadratic convergence can be achieved even if the given matrix has complex eigenvalues. It is expected that ultimate quadratic, or near-quadratic convergence is achievable when the algorithms are applied

Matrix size	40	80	120	160	200
Matrix 1	8	9	10	11	11
Matrix 2	9	11	11	12	13
Matrix 3	11	12	13	13	14
Matrix 4	11	12	13	13	13

Table 2: Sweeps taken for matrices of various sizes.

to compute the eigenvalue decomposition of a near-normal matrix.

The algorithm we used in the experiment combines the RTZ algorithm and the QR algorithm. When the testing matrices have distinct eigenvalues, we can obtain quadratic convergence, which is consistent with our theoretical analysis. However, only near-quadratic convergence was observed in the experiment when a matrix has repeated eigenvalues. An interesting problem is thus if we can obtain quadratic convergence when using Jacobi-like algorithm for computing (over the real field) the eigenvalue decomposition of a normal matrix with repeated eigenvalues.

It should be noted that the algorithm used in our experiment is not the only candidate for an efficient Jacobi-like algorithm. Through our discussion we did not restrict ourselves to use any particular method for generating orthogonal transformations. Any method which generates the required transformation matrices can be used. For example, the QR procedure may be used not only in the case that all the eigenvalues of the 2×2 block matrix B in Fig. 1 are complex, but also in cases that some or all of the eigenvalues are real. However, a naive implementation using the QR procedure may lead to very slow convergence. Special care has to be taken in order to achieve ultimate quadratic convergence. This issue is discussed in [13].

References

[1] R. P. Brent, "Parallel algorithms for digital signal processing", *Proceedings of NATO Advanced Study Institute on Numerical Linear Algebra, Digital Signal Processing and Parallel Algorithms*, Leuven, Belgium, August, 1988.

[2] R. P. Brent and F. T. Luk, "The solution of singular-value and symmetric eigenvalue problems on multiprocessor arrays", *SIAM J. Sci. and Statist. Comput.*, 6, 1985, pp. 69-84.

[3] P. J. Eberlein, "On the Schur decomposition of a matrix for parallel computation", *IEEE Trans. on Computers*, Vol. C-36, No. 2, 1987, pp. 167-174.

[4] N. Mackey, "Hamilton and Jacobi meet again: quaternions and the eigenvalue problem", *SIAM J. Matrix Anal. Appl.*, Vol. 16, No. 2, pp. 421-435.

[5] M. Mantharam and P. J. Eberlein, "The real two-zero algorithm: a parallel algorithm to reduce a real matrix to a real Schur form", *IEEE Trans. on Parallel and Distributed Systems*, Vol. 6, No. 1, 1995, pp. 48-62.

[6] B. T. Smith, J. M. Boyle, J. J. Dongarra, B. S. Garbow, Y. Ikebe, V. C. Klema and C. B. Moler, *Matrix Eigensystem Routines - EISPACK Guide*, Vol. 6 of Lecture Notes in Computer Science, Springer-verlag, Berlin, 2 ed., 1976.

[7] G. W. Stewart, "A Jacobi-like algorithm for computing the Schur decomposition of a nonhermitian matrix", *SIAM J. Sci. and Statist. Comput.*, Vol. 6, No. 4, 1985, pp. 853-864.

[8] F. T. Luk and H. Park, "On parallel Jacobi orderings", *SIAM J. Sci. and Statist. Comput.*, 10, 1989, pp. 18-26.

[9] R. Schreiber, "Solving eigenvalue and singular value problems on an undersized systolic array", *SIAM. J. Sci. Stat. Comput.*, 7, 1986, pp. 441-451.

[10] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, England, 1965.

[11] B. B. Zhou and R. P. Brent, "A parallel ordering algorithm for efficient one-sided Jacobi SVD computations", to appear in *Proc. of Sixth IASTED-ISMM International Conference on Parallel and Distributed Computing and Systems*, Washington, DC, October 1994.

[12] B. B. Zhou and R. P. Brent, "On the parallel implementation of the one-sided Jacobi algorithm for singular value decompositions", to appear in *Proc. of 3rd Euromicro Workshop on Parallel and Distributed Processing*, San Remo, Italy, January 1995.

[13] B. B. Zhou and R. P. Brent, "Block Jacobi-like algorithms for eigenvalue decomposition of a real normal matrix", in preparation.