NUMERICAL LINEAR ALGEBRA WITH APPLICATIONS
Numer. Linear Algebra Appl. 2009; **16**:801–815
Published online 2 March 2009 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/nla.642

A twisted factorization method for symmetric SVD of a complex symmetric tridiagonal matrix

Wei Xu[1] and Sanzheng Qiao[2,*,†]

[1]*School of Software Engineering, Fudan University, Shanghai, People's Republic of China*
[2]*Department of Computing and Software, McMaster University, Hamilton, Ont., Canada L8S 4K1*

SUMMARY

This paper presents an $O(n^2)$ method based on the twisted factorization for computing the Takagi vectors of an n-by-n complex symmetric tridiagonal matrix with known singular values. Since the singular values can be obtained in $O(n^2)$ flops, the total cost of symmetric singular value decomposition or the Takagi factorization is $O(n^2)$ flops. An analysis shows the accuracy and orthogonality of Takagi vectors. Also, techniques for a practical implementation of our method are proposed. Our preliminary numerical experiments have verified our analysis and demonstrated that the twisted factorization method is much more efficient than the implicit QR method, divide-and-conquer method and Matlab singular value decomposition subroutine with comparable accuracy. Copyright © 2009 John Wiley & Sons, Ltd.

Received 10 April 2008; Revised 30 January 2009; Accepted 1 February 2009

KEY WORDS: twisted factorization; symmetric SVD; Takagi factorization; fast SVD; complex symmetric matrix

1. INTRODUCTION

For any complex symmetric matrix A, there exist a diagonal singular value matrix Σ and a unitary matrix U, such that

$$A = U\Sigma U^{\mathrm{T}}$$

This special form of the singular value decomposition (SVD) [1] of A is called symmetric SVD (SSVD) or Takagi factorization [2, 3]. One obvious advantage of this form is that it reflects the symmetry of A and thus saves the storage and computation by about half.

*Correspondence to: Sanzheng Qiao, Department of Computing and Software, McMaster University, Hamilton, Ont., Canada L8S 4K1.
†E-mail: qiao@mcmaster.ca

The SSVD of a complex symmetric matrix has applications in, for instance, the Grunsky inequalities [4], the computation of the near-best uniform polynomial or rational approximation of a high degree polynomial on a disk [5], and the complex independent component analysis problems [6]. However, Matlab and LAPACK [7] do not support complex symmetric structure and treat it as general complex. To compute the SVD of a complex symmetric matrix in LAPACK, the matrix is first reduced to a bidiagonal matrix, in which the symmetric structure is lost.

The algorithm proposed in this paper leads to an efficient method for computing the SSVD of a complex square Hankel matrix. A complex Hankel matrix is symmetric and can be reduced to a symmetric tridiagonal matrix in $O(n^2 \log n)$ [8]. Then, by applying the twisted factorization method, the symmetric SVD of the symmetric tridiagonal matrix can be efficiently obtained in $O(n^2)$. Thus, it leads to an $O(n^2 \log n)$ algorithm for computing the symmetric SVD of a complex square Hankel matrix [9]. Since a Toeplitz matrix can be transformed into a Hankel matrix by reversing its rows or columns, this method can be straightforwardly modified into a fast SVD algorithm for square Toeplitz matrices.

As in the case of a general matrix, computing the Takagi factorization of a complex symmetric matrix consists of two stages: tridiagonalization and diagonalization. In the first stage, a complex symmetric matrix A is reduced to a complex symmetric tridiagonal T using two-side Householder transformations or Lanczos method [8, 10]. In the second stage, the Takagi factorization

$$T = V \Sigma V^{\mathrm{T}}$$

of the complex symmetric tridiagonal T resulted from the first stage is computed. The methods for the second stage include the implicit QR method [3, 11] and the divide-and-conquer method [12]. The QR method requires $O(n^2)$ flops for computing all singular values, but additional $O(n^3)$ flops for all the Takagi vectors. The divide-and-conquer method [12] integrates the computation of the singular values and the computation of the Takagi vectors and in practice requires much less than $O(n^3)$ flops. Thus, the main cost of the second stage is the computation of the Takagi vector matrix V.

This paper presents a method for computing V given T and its computed singular values in Σ using $O(n^2)$ flops and $n^2 + O(n)$ storage space. In Section 2, we describe the twisted factorization, based on which, an efficient and stable method for computing the Takagi vectors is presented in Section 3. The accuracy and orthogonality of the computed vectors are analyzed in Section 4. Then, in Section 5, we address the issue of multiple and clustered singular values. Our numerical experiment results presented in Section 6 show that our method is efficient and accurate.

2. TWISTED FACTORIZATION

Let T be an n-by-n complex symmetric tridiagonal matrix, then $P = TT^{\mathrm{H}}$ is Hermitian pentadiagonal. In this section, we first present two decompositions of the shifted matrix:

$$P - \mu I = L D_L L^{\mathrm{H}} = U D_U U^{\mathrm{H}} \tag{1}$$

DOI: 10.1002/nla

where

$$L=\begin{bmatrix} 1 & & & & \\ l_1 & \ddots & & 0 & \\ m_1 & \ddots & \ddots & & \\ & \ddots & \ddots & \ddots & \\ 0 & & m_{n-2} & l_{n-1} & 1 \end{bmatrix}, \quad D_L=\operatorname{diag}(\alpha_1,\ldots,\alpha_n), \quad \alpha_i\in\mathbb{R}$$

and

$$U=\begin{bmatrix} 1 & u_1 & v_1 & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & v_{n-2} \\ & & & \ddots & u_{n-1} \\ & & & & 1 \end{bmatrix}, \quad D_U=\operatorname{diag}(\beta_1,\ldots,\beta_n), \quad \beta_i\in\mathbb{R}$$

Then, by combining the above two decompositions of the shifted $P-\mu I$, we construct the twisted factorization

$$P-\mu I=N_k D_k N_k^{\mathrm{H}}$$

It is called twisted factorization because N_k consists of part of the lower triangular L and part of the upper triangular U in (1). In the next section, based on the twisted factorization, we will show how to efficiently compute the eigenvectors of P, which are the left singular vectors of T but need not be the Takagi vectors. The transformation presented in [12] is then applied to convert the eigenvectors into the Takagi vectors.

To obtain the decomposition $P-\mu I=LD_LL^{\mathrm{H}}$, we compare the entries on the both sides of $P-\mu I=LD_LL^{\mathrm{H}}$. Specifically, the $(i+2,i)$-entry of $P-\mu I$ is $P_{i+2,i}$ and the corresponding entry of LD_LL^{H} is $m_i\alpha_i$. Thus, we have $m_i=P_{i+2,i}/\alpha_i$. The $(i+2,i+1)$-entries of the both sides are $P_{i+2,i+1}$ and $m_il_i^*\alpha_i+l_{i+1}\alpha_{i+1}$, where x^* denotes the complex conjugate of x. The $(i+2,i+2)$-entries are $P_{i+2,i+2}-\mu$ and $|m_i|^2\alpha_i+|l_{i+1}|^2\alpha_{i+1}+\alpha_{i+2}$. In summary, we have the following algorithm for the $\mathrm{LDL}^{\mathrm{H}}$ decomposition: $P-\mu I=LD_LL^{\mathrm{H}}$.

Algorithm 2.1
Given the Hermitian pentadiagonal $P=TT^{\mathrm{H}}$, this algorithm computes the $\mathrm{LDL}^{\mathrm{H}}$ decomposition of the shifted $P-\mu I=LD_LL^{\mathrm{H}}$.

$\alpha_1=P_{11}-\mu$; % $(1,1)$-entry
$l_1=P_{21}/\alpha_1$; % $(2,1)$-entry
$\alpha_2=P_{22}-\mu-|l_1|^2\alpha_1$; % $(2,2)$-entry
for $i=1:n-2$

DOI: 10.1002/nla

$m_i = P_{i+2,i}/\alpha_i$; % $(i+2,i)$-entry
$l_{i+1} = (P_{i+2,i+1} - m_i l_i^* \alpha_i)/\alpha_{i+1}$; % $(i+2,i+1)$-entry
$\alpha_{i+2} = P_{i+2,i+2} - \mu - |m_i|^2 \alpha_i - |l_{i+1}|^2 \alpha_{i+1}$;
% $(i+2,i+2)$-entry

end

The computational cost of the $\mathrm{LDL}^{\mathrm{H}}$ decomposition is $22n-37$ floating-point multiplications (real) and $11n-16$ floating-point additions (real), a total of $33n-53$ floating-point operations (multiplications or additions).

Similarly, we can compute the $\mathrm{UDU}^{\mathrm{H}}$ decomposition of $P-\mu I$ as follows.

Algorithm 2.2

Given the Hermitian pentadiagonal P, this algorithm computes the $\mathrm{UDU}^{\mathrm{H}}$ decomposition of the shifted $P-\mu I = U D_U U^{\mathrm{H}}$.

Let $\beta_{n+2} = \beta_{n+1} = v_n = v_{n-1} = u_n = 0$.

for $i = n-2:-1:1$
$\beta_{i+2} = P_{i+2,i+2} - \mu - |u_{i+2}|^2 \beta_{i+3} - |v_{i+2}|^2 \beta_{i+4}$; % $(i+2,i+2)$-entry
$u_{i+1} = (P_{i+1,i+2} - v_{i+1} u_{i+2}^* \beta_{i+3})/\beta_{i+2}$; % $(i+1,i+2)$-entry
$v_i = P_{i,i+2}/\beta_{i+2}$; % $(i,i+2)$-entry
end
$\beta_2 = P_{22} - \mu - |u_2|^2 \beta_3 - |v_2|^2 \beta_4$; % $(2,2)$-entry
$u_1 = (P_{12} - v_1 u_2^* \beta_3)/\beta_2$; % $(1,2)$-entry
$\beta_1 = P_{11} - \mu - |u_1|^2 \beta_2 - |v_1|^2 \beta_3$; % $(1,1)$-entry

The computational cost of the $\mathrm{UDU}^{\mathrm{H}}$ decomposition is $22n-16$ floating-point multiplications and $11n-4$ floating-point additions, a total of $33n-20$ floating-point operations. Now, given the $\mathrm{LDL}^{\mathrm{H}}$ and $\mathrm{UDU}^{\mathrm{H}}$ decompositions (1), we consider the twisted factorization of the shifted matrix

$$P - \mu I = N_k D_k N_k^{\mathrm{H}} \tag{2}$$

where $D_k = \mathrm{diag}(\alpha_1, \ldots, \alpha_{k-2}, \xi_k, \gamma_k, \beta_{k+1}, \ldots, \beta_n)$, $\xi_k, \gamma_k \in \mathbb{R}$ and

$$N_k = \begin{bmatrix}
1 & & & & & & & \\
l_1 & \ddots & & & 0 & & & \\
m_1 & \ddots & 1 & & & & & \\
 & \ddots & l_{k-2} & 1 & v_{k-1} & & & \\
 & & m_{k-2} & \eta_k & 1 & u_k & \ddots & \\
 & & & & & 1 & \ddots & v_{n-2} \\
 & & 0 & & & & \ddots & u_{n-1} \\
 & & & & & & & 1
\end{bmatrix} \tag{3}$$

DOI: 10.1002/nla

is a twisted combination of the partial lower triangular L and the partial upper triangular U. Note that for each index k, $1 \leqslant k \leqslant n$, there is a corresponding twisted factorization. However, given the $\mathrm{LDL}^{\mathrm{H}}$ and $\mathrm{UDU}^{\mathrm{H}}$ decompositions (1), there are only three unknowns $\xi_k, \gamma_k \in \mathbb{R}$ and $\eta_k \in \mathbb{C}$ in a twisted factorization. The following theorem shows a computation of the three unknowns and a relation between γ_k^{-1} and a Rayleigh quotient of $(P-\mu I)^{-1}$ when $P-\mu I$ is invertible.

Theorem 2.3
Given the $\mathrm{LDL}^{\mathrm{H}}$ and $\mathrm{UDU}^{\mathrm{H}}$ decompositions (1) of $P-\mu I$, defining $m_{-1}=m_0=\alpha_0=l_0=v_0=v_{n-1}=\beta_{n+1}=0$, then the three unknowns in the twisted factorization (2) are given by

$$\xi_k = \alpha_{k-1} - |v_{k-1}|^2 \beta_{k+1}, \quad k=1,\ldots,n$$

$$\eta_k = \begin{cases} 0, & k=1 \\ (l_{k-1}\alpha_{k-1} - u_k v_{k-1}^* \beta_{k+1})/\xi_k, & k=2,\ldots,n-1 \\ l_{n-1}, & k=n \end{cases}$$

$$\gamma_k = \begin{cases} \beta_k - |m_{k-2}|^2 \alpha_{k-2} - \xi_k |\eta_k|^2, & k=1,\ldots,n-1 \\ \alpha_n, & k=n \end{cases}$$

Also, when $P-\mu I$ is invertible,

$$\gamma_k^{-1} = \mathbf{e}_k^{\mathrm{T}} (P-\mu I)^{-1} \mathbf{e}_k \tag{4}$$

where $\mathbf{e}_k$ is the kth unit vector.

Proof
The $(k-1,k-1)$-entry of $N_k D_k N_k^{\mathrm{H}}$ is $|m_{k-3}|^2\alpha_{k-3} + |l_{k-2}|^2\alpha_{k-2} + \xi_k + |v_{k-1}|^2\beta_{k+1}$ and the corresponding entry in $L D_L L^{\mathrm{H}}$ is $|m_{k-3}|^2\alpha_{k-3} + |l_{k-2}|^2\alpha_{k-2} + \alpha_{k-1}$. Equating the two entries, we have $\xi_k = \alpha_{k-1} - |v_{k-1}|^2\beta_{k+1}$. Equating the $(k,k-1)$-entries in $N_k D_k N_k^{\mathrm{H}}$ and $L D_L L^{\mathrm{H}}$ leads to $\xi_k \eta_k = l_{k-1}\alpha_{k-1} - u_k v_{k-1}^* \beta_{k+1}$. While comparing the (k,k)-entries of $N_k D_k N_k^{\mathrm{H}}$ and $U D_U U^{\mathrm{H}}$ gives $|m_{k-2}|^2\alpha_{k-2} + \xi_k|\eta_k|^2 + \gamma_k = \beta_k$. The trivial cases when $k=1,2,3,n$ can be verified similarly.

Since the kth column of N_k is the kth unit vector, $N_k \mathbf{e}_k = \mathbf{e}_k$. It then follows that $N_k^{-1}\mathbf{e}_k = \mathbf{e}_k$ and the Rayleigh quotient

$$\mathbf{e}_k^{\mathrm{T}} (P-\mu I)^{-1} \mathbf{e}_k = \mathbf{e}_k^{\mathrm{T}} (N_k^{\mathrm{H}})^{-1} D_k^{-1} N_k^{-1} \mathbf{e}_k = \mathbf{e}_k^{\mathrm{T}} D_k^{-1} \mathbf{e}_k = \gamma_k^{-1}$$

when $P-\mu I$ is nonsingular. □

The cost of computing ξ_k, η_k and γ_k is 11 floating-point multiplications and 7 floating-point additions, a total of 18 floating-point operations, without counting the operations for $u_k v_{k-1}^* \beta_{k+1}$, whose complex conjugate is computed in Algorithm 2.2.

3. COMPUTING TAKAGI VECTORS

Using the twisted factorization described in the previous section, we now present an efficient method for computing the eigenvectors of P, which are the left singular vectors of T.

DOI: 10.1002/nla

Suppose that $P = Q\Lambda Q^{\mathrm{H}}$ is an eigendecomposition of P, where $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$ is the eigenvalue matrix and $Q = [\mathbf{q}_1, \ldots, \mathbf{q}_n] = [q_{i,j}]$ is the unitary eigenvector matrix. As we know, if we have a good approximation μ of an eigenvalue λ_i, one of effective ways of computing the corresponding eigenvector $\mathbf{q}_i$ is the shifted inverse power method. For each index k of the twisted factorization (2), the solution $\mathbf{z}_k$ for

$$(P - \mu I)\mathbf{z}_k = \gamma_k \mathbf{e}_k \tag{5}$$

is the result of one inverse power iteration with shift μ and initial vector $\gamma_k \mathbf{e}_k$. How accurate is the approximation $\mathbf{z}_k$ of the eigenvector $\mathbf{q}_i$? After normalizing $\mathbf{z}_k$ in (5), we have

$$\left\| P\frac{\mathbf{z}_k}{\|\mathbf{z}_k\|_2} - \lambda_i \frac{\mathbf{z}_k}{\|\mathbf{z}_k\|_2} \right\|_2 = \left\| (\mu - \lambda_i)\frac{\mathbf{z}_k}{\|\mathbf{z}_k\|_2} + \frac{\gamma_k}{\|\mathbf{z}_k\|_2}\mathbf{e}_k \right\|_2 \leqslant |\lambda_i - \mu| + \frac{|\gamma_k|}{\|\mathbf{z}_k\|_2}$$

This shows that the error in $\mathbf{z}_k$ as the eigenvector corresponding to λ_i is bounded by the error in μ as an approximation of λ_i plus $|\gamma_k|/\|\mathbf{z}_k\|_2$. The following theorem [13] gives an upper bound for the term $|\gamma_k|/\|\mathbf{z}_k\|_2$.

Theorem 3.1 (*Dhillon* [*13*])
Suppose that $P - \mu I$ is invertible and

$$(P - \mu I)\mathbf{z}_k = \mathbf{e}_k \gamma_k \quad \text{for } k = 1, \ldots, n$$

Then, if the kth entry $q_{k,i}$ of $\mathbf{q}_i$ is nonzero,

$$\begin{aligned} \frac{|\gamma_k|}{\|\mathbf{z}_k\|_2} &= \frac{|\lambda_i - \mu|}{|q_{k,i}|}(1 + (|q_{k,i}|^{-2} - 1)\mathscr{A})^{-1/2} \\ &\leqslant \frac{|\lambda_i - \mu|}{|q_{k,i}|} \\ &\leqslant \sqrt{n}|\lambda_i - \mu| \end{aligned}$$

for at least one k. Here $\mathscr{A}$ is a weighted arithmetic mean of $[|\lambda_i - \mu|/|\lambda_j - \mu|]^2$, $j \neq i$, and $0 < \mathscr{A} < (|\lambda_i - \mu|/\mathrm{gap}(\mu))^2$, where $\mathrm{gap}(\mu) = \min_{j \neq i} |\lambda_j - \mu|$.

The above theorem implies that if μ is a good approximation of λ_i, then there exists at least one index k so that $\mathbf{z}_k$ is a good approximation of the eigenvector $\mathbf{q}_i$ associated with the eigenvalue λ_i of P. The accuracy can be as good as $|\lambda_i - \mu|$ by a factor of at most $\sqrt{n} + 1$. Since the term $|\gamma_k|/\|\mathbf{z}_k\|_2$ is dependent of the index k, to get good accuracy, we find the smallest $|\gamma_k|$ by computing the twisted factorization (2) for $k = 1, \ldots, n$. Note that $\|\mathbf{z}_k\|_2 \geqslant 1$ since, from (5) and (4), the kth entry of $\mathbf{z}_k$:

$$\mathbf{e}_k^{\mathrm{T}}\mathbf{z}_k = \gamma_k \mathbf{e}_k^{\mathrm{T}}(P - \mu I)^{-1}\mathbf{e}_k = 1$$

implying that $|\gamma_k|/\|\mathbf{z}_k\|_2 \leqslant |\gamma_k|$. Once the index k for the smallest $|\gamma_k|$ is found, we solve for $\mathbf{z}_k$ in (5).

In the following, we show that $\mathbf{z}_k$ can be solved efficiently and stably in $O(n)$ operations by exploiting the structure of N_k in the twisted factorization (2).

DOI: 10.1002/nla

As we know that when μ is a good approximation of an eigenvalue λ_i of P, the system (5) is ill-conditioned. However, using the twisted factorization $P-\mu I=N_k D_k N_k^{\mathrm{H}}$ and noting that the kth column of the twisted factor N_k is $\mathbf{e}_k$, we can reformulate the system (5) into an equivalent but simpler one:

$$N_k^{\mathrm{H}}\mathbf{z}_k=\gamma_k D_k^{-1}N_k^{-1}\mathbf{e}_k=\mathbf{e}_k$$

Solving $N_k^{\mathrm{H}}\mathbf{z}_k=\mathbf{e}_k$ is not only more efficient than solving (5) but also more stable because the ill-conditioning of (5) caused by the small $|\gamma_k|$ in D_k is avoided. From the structure of N_k^{H} (3), the entries z_j of the solution $\mathbf{z}_k$ are given by

$$\begin{aligned}
z_k &= 1\\
z_{k-1} &= -\eta_k^*, \quad k>1\\
z_j &= -l_j^* z_{j+1}-m_j^* z_{j+2}, \quad j=k-2,k-3,\ldots,1\\
z_{k+1} &= v_{k-1}^*\eta_k^*-u_k^*, \quad v_0=0\\
z_j &= -v_{j-2}^* z_{j-2}-u_{j-1}^* z_{j-1}, \quad j=k+2,k+3,\ldots,n
\end{aligned} \tag{6}$$

The computation of $\mathbf{z}_k$ requires $8n$ floating-point multiplications and $6n$ floating-point additions, a total of $14n$ operations.

The following algorithm summarizes the procedure of computing the eigenvector given a computed eigenvalue $\hat{\lambda}_i$ of P. The last step in the following algorithm requires $5n$ floating-point operations. Adding the costs of all steps, we obtain the total cost of $85n-55$ for computing an eigenvector.

Algorithm 3.2 (*Computing eigenvector*)
Given the Hermitian pentadiagonal matrix $P=TT^{\mathrm{H}}$ and $\mu=\hat{\lambda}_i$, a computed eigenvalue of P, this algorithm computes an approximation $\hat{\mathbf{q}}_i$ of the eigenvector $\mathbf{q}_i$ corresponding to λ_i.

1. Compute the LDL$^{\mathrm{H}}$ and UDU$^{\mathrm{H}}$ decompositions (1) of $P-\mu I$ using Algorithms 2.1 and 2.2.
2. Applying Theorem 2.3, for $i=1,\ldots,n$, compute the twisted factorizations $P-\mu I=N_i D_i N_i^{\mathrm{H}}$ and find k such that $|\gamma_k|=\min_i|\gamma_i|$.
3. Solve for $\mathbf{z}_k$ in $N_k^{\mathrm{H}}\mathbf{z}_k=\mathbf{e}_k$ using (6).
4. Set $\hat{\mathbf{q}}_i=\mathbf{z}_k/\|\mathbf{z}_k\|_2$.

Remark
A refinement technique can be integrated into Algorithm 3.2. We first compute an eigenvector $\tilde{\mathbf{q}}_i$ corresponding to μ via Algorithm 3.2. Then, μ is refined by $\tilde{\mathbf{q}}_i^H P\tilde{\mathbf{q}}_i$. Finally, we compute $\hat{\mathbf{q}}_i$ through Algorithm 3.2 again with the refined μ. In fact, our refinement is equivalent to reiterate the inverse power method, for each eigenvector corresponding to a simple eigenvalue. Obviously, the cost of computation is almost doubled. However, our experiments show that the orthogonality of eigenvectors is improved through the refinement.

The computed eigenvectors $\hat{\mathbf{q}}_i$ are the left singular vectors of T, but they need not be the Takagi vectors [12]. The transformations in [12] can be used to transform the eigenvectors of P to Takagi

DOI: 10.1002/nla

vectors of T. The cost for each Takagi vector transformation is $O(n)$, thus the total cost of the transformations to the Takagi vectors of T is at most $O(n^2)$.

4. ACCURACY AND ORTHOGONALITY

Now that we have an efficient algorithm for computing the eigenvectors of P, in this section we investigate the accuracy and orthogonality of the computed eigenvectors. The accuracy of a computed eigenvector $\hat{\mathbf{q}}_i$ is measured by the sine of the angle between the computed $\hat{\mathbf{q}}_i$ and the exact eigenvector $\mathbf{q}_i$ corresponding to the eigenvalue λ_i of P. The orthogonality of two computed eigenvectors $\hat{\mathbf{q}}_i$ and $\hat{\mathbf{q}}_j$ is measured by the cosine of the angle between them.

Replacing $\mathbf{z}_k$ in (5) with $\|\mathbf{z}_k\|_2\hat{\mathbf{q}}_i$, we get

$$(P-\mu I)\hat{\mathbf{q}}_i = \frac{\gamma_k}{\|\mathbf{z}_k\|_2}\mathbf{e}_k$$

Since $\mathbf{e}_k^{\mathrm{T}}\hat{\mathbf{q}}_i = 1/\|\mathbf{z}_k\|_2$, the above equation can be written as

$$(P+E)\hat{\mathbf{q}}_i = \mu\hat{\mathbf{q}}_i$$

where $E=-\gamma_k\mathbf{e}_k\mathbf{e}_k^{\mathrm{T}}$, which indicates that the computed eigenvector $\hat{\mathbf{q}}_i$ is the exact eigenvector corresponding to the eigenvalue μ of $P+E$. In other words, $(\mu, \hat{\mathbf{q}}_i)$ is an eigenpair of the perturbed $P+E$, whereas $(\lambda_i, \mathbf{q}_i)$ is an eigenpair of P. This formulation allows us to apply the following Second $\sin\theta$ Theorem of eigenspaces by Davis and Kahan [14] to study the accuracy of the computed $\hat{\mathbf{q}}_i$.

Theorem 4.1 (*Second* $\sin\theta$ *Theorem* [*14*])
Suppose $A, \widehat{A}\in\mathbb{C}^{n\times n}$ are Hermitian and $X=[X_1 X_2]$ and $\widehat{X}=[\widehat{X}_1\ \widehat{X}_2]$ are unitary, where $X_1, \widehat{X}_1 \in \mathbb{C}^{n\times l}$ $(1\leqslant l\leqslant n-1)$, such that

$$X^{\mathrm{H}}AX=\begin{bmatrix} A_{11} & 0 \\ 0 & A_{22}\end{bmatrix} \quad \text{and} \quad \widehat{X}^{\mathrm{H}}\widehat{A}\widehat{X}=\begin{bmatrix} \widehat{A}_{11} & 0 \\ 0 & \widehat{A}_{22}\end{bmatrix}$$

where $A_{11}, \widehat{A}_{11}\in C^{l\times l}$. Let $\mathscr{X}=\mathrm{Range}(X_1)$, $\widehat{\mathscr{X}}=\mathrm{Range}(\widehat{X}_1)$,

$$R=A\widehat{X}_1-\widehat{X}_1\widehat{A}_{11}$$

and

$$\delta=\min_{i,j}\{|\hat{\lambda}_i-\lambda_j|:\hat{\lambda}_i\in\lambda(\widehat{A}_{11}),\ \lambda_j\in\lambda(A_{22})\}>0$$

Then, we have

$$\|\sin\theta(\mathscr{X},\widehat{\mathscr{X}})\|_F\leqslant\frac{\|R\|_F}{\delta}$$

where $\theta(\mathscr{X},\widehat{\mathscr{X}})$ denotes the angle between the two subspaces $\mathscr{X}$ and $\widehat{\mathscr{X}}$.

In the above theorem, setting

$$A=P,\quad \widehat{A}=P+E,\quad X_1=\mathbf{q}_i \quad \text{and} \quad \widehat{X}_1=\hat{\mathbf{q}}_i$$

DOI: 10.1002/nla

we have $l=1$, $A_{11}=\lambda_i$, $\widehat{A}_{11}=\mu$,

$$\delta=\min_{j\neq i}|\lambda_j-\mu|=\text{gap}(\mu) \quad \text{and} \quad R=P\hat{\mathbf{q}}_i-\mu\hat{\mathbf{q}}_i=\frac{\gamma_k}{\|\mathbf{z}_k\|_2}\mathbf{e}_k$$

Then, from Theorem 4.1, using the inequality $|\gamma_k|/\|\mathbf{z}_k\|_2\leqslant\sqrt{n}|\lambda_i-\mu|$ in Theorem 3.1, we obtain the following theorem on an upper bound for the sine of the angle between $\mathbf{q}_i$ and $\hat{\mathbf{q}}_i$.

Theorem 4.2
Let $(\lambda_i,\ \mathbf{q}_i)$ be an eigenpair of P, $\mu\approx\lambda_i$, and $\hat{\mathbf{q}}_i$ is computed by Algorithm 3.2, then

$$|\sin\theta(\hat{\mathbf{q}}_i,\mathbf{q}_i)|\leqslant\frac{\sqrt{n}|\lambda_i-\mu|}{\text{gap}(\mu)}$$

where $\text{gap}(\mu)=\min_{j\neq i}|\lambda_j-\mu|$.

This theorem shows that if μ is a good approximation of λ_i, that is, $|\lambda_i-\mu|$ is small, then the computed eigenvector $\hat{\mathbf{q}}_i$ is a good approximation of the eigenvector $\mathbf{q}_i$ corresponding to the eigenvalue λ_i provided that λ_i is not clustered with the other eigenvalues.

Now we present a theorem on the orthogonality of the eigenvectors computed by Algorithm 3.2. The orthogonality between two computed eigenvectors $\hat{\mathbf{q}}_i$ and $\hat{\mathbf{q}}_j$ is measured by the cosine of the angle between them. The following theorem gives an upper bound for the cosine.

Theorem 4.3
Let $(\lambda_i,\mathbf{q}_i)$ and $(\lambda_j,\mathbf{q}_j)$, $i\neq j$, be two eigenpairs of P and $\hat{\mathbf{q}}_i$ and $\hat{\mathbf{q}}_j$ be the eigenvectors computed by Algorithm 3.2 using the computed eigenvalues $\hat{\lambda}_i\approx\lambda_i$ and $\hat{\lambda}_j\approx\lambda_j$, then

$$\cos\theta(\hat{\mathbf{q}}_i,\hat{\mathbf{q}}_j)\leqslant\frac{\sqrt{n}|\hat{\lambda}_i-\lambda_i|}{\text{gap}(\hat{\lambda}_i)}+\frac{\sqrt{n}|\hat{\lambda}_j-\lambda_j|}{\text{gap}(\hat{\lambda}_j)} \tag{7}$$

Proof
Assuming $1\leqslant i\leqslant k<j\leqslant n$ and denoting $\mathscr{V}_1$ and $\mathscr{V}_2$ as the subspaces spanned by $\mathbf{q}_1,\ldots,\mathbf{q}_k$ and $\mathbf{q}_{k+1},\ldots,\mathbf{q}_n$ respectively, we have

$$\begin{aligned}|\cos\theta(\hat{\mathbf{q}}_i,\hat{\mathbf{q}}_j)| &\leqslant\left|\cos\left(\frac{\pi}{2}-\theta(\hat{\mathbf{q}}_i,\mathscr{V}_1)-\theta(\hat{\mathbf{q}}_j,\mathscr{V}_2)\right)\right|\\ &\leqslant|\sin\theta(\hat{\mathbf{q}}_i,\mathscr{V}_1)|+|\sin\theta(\hat{\mathbf{q}}_j,\mathscr{V}_2)|\end{aligned}$$

since $\mathscr{V}_1$ and $\mathscr{V}_2$ are orthogonal to each other. Also, since $\mathbf{q}_i\in\mathscr{V}_1$ and $\mathbf{q}_j\in\mathscr{V}_2$, we have

$$\theta(\hat{\mathbf{q}}_i,\mathscr{V}_1)\leqslant\theta(\hat{\mathbf{q}}_i,\mathbf{q}_i) \quad \text{and} \quad \theta(\hat{\mathbf{q}}_j,\mathscr{V}_2)\leqslant\theta(\hat{\mathbf{q}}_j,\mathbf{q}_j)$$

The inequality (7) then follows from Theorem 4.2. □

This theorem shows that if $\hat{\lambda}_i$ and $\hat{\lambda}_j$ are good approximations of λ_i and λ_j, respectively, and both λ_i and λ_j are not clustered with the other eigenvalues, then the eigenvectors $\hat{\mathbf{q}}_i$ and $\hat{\mathbf{q}}_j$ computed by Algorithm 3.2 have good orthogonality.

DOI: 10.1002/nla

5. MULTIPLE AND CLUSTERED SINGULAR VALUES

In the previous discussion, we have assumed that the singular values of T, which are the square roots of the eigenvalues of $P=TT^{\mathrm{H}}$, are simple. In this section, we consider the case of multiple singular values of T.

We assume that T is irreducible, that is there are no zero entries on its subdiagonal. Otherwise, T is block diagonal and can be deflated into smaller irreducible matrices. If T is irreducible, then $P=TT^{\mathrm{H}}$ is Hermitian pentadiagonal with nonzero entries on its second subdiagonal. It is well known that if a Hermitian tridiagonal matrix of order n is irreducible, then its rank is at least $n-1$. This can be extended to Hermitian pentadiagonal matrices. If a Hermitian pentadiagonal matrix of order n has no zero entries on its second subdiagonal, then it is of rank at least $n-2$. This implies that the multiplicity of any eigenvalue of an irreducible P is at most two, since $\mathrm{rank}(P-\sigma^2 I)\geqslant n-2$, for any $\sigma\in\mathbb{R}$.

In Sections 2 and 3, we have shown that if μ is a good approximation of a simple eigenvalue λ_i of P, then there exists an index k such that γ_k in the twisted factorization $P-\mu I=N_k D_k N_k^{\mathrm{H}}$ is small and the corresponding eigenvector can be obtained by solving for $\mathbf{z}_k$ in $N_k^{\mathrm{H}}\mathbf{z}_k=\mathbf{e}_k$. In this section, we will show that, in the case of a multiple eigenvalue $\lambda_i=\lambda_{i+1}$, there exist two indices k_1 and k_2, $k_1\neq k_2$, so that γ_{k_1} and γ_{k_2} in the twisted factorizations indexed by k_1 and k_2 are small and the associated eigenvectors can be obtained by solving $N_{k_1}^{\mathrm{H}}\mathbf{z}_{k_1}=\mathbf{e}_{k_1}$ and $N_{k_2}^{\mathrm{H}}\mathbf{z}_{k_2}=\mathbf{e}_{k_2}$.

Theorem 5.1
Let $P-\mu I=Q\Lambda Q^{\mathrm{H}}$ be an eigendecomposition, where $\Lambda=\mathrm{diag}(\lambda_1,\ldots,\lambda_n)$ and $Q=[q_{i,j}]$. Suppose that $\lambda_i=\lambda_{i+1}$ is a multiple eigenvalue, then there exist k_1 and k_2, $k_1\neq k_2$, such that

$$\max(|q_{k_1,i}|,|q_{k_1,i+1}|)\geqslant 1/\sqrt{2n} \quad \text{and} \quad \max(|q_{k_2,i}|,|q_{k_2,i+1}|)\geqslant 1/\sqrt{2n}$$

that is, at least one of $|q_{k_1,i}|$ and $|q_{k_1,i+1}|$ is not smaller than $1/\sqrt{2n}$ and the same for $|q_{k_2,i}|$ and $|q_{k_2,i+1}|$.

Proof
Since Q is unitary, $\sum_{k=1}^{n}|q_{k,i}|^2=1$, which implies that there exists k_1: $|q_{k_1,i}|^2\geqslant 1/n$. Making use of $\sum_{j=1}^{n}|q_{k_1,j}|^2=1$, we immediately obtain

$$\min(|q_{k_1,i}|^2,|q_{k_1,i+1}|^2)\leqslant\tfrac{1}{2}$$

Without loss of generality, we assume $|q_{k_1,i+1}|^2=\min(|q_{k_1,i}|^2,|q_{k_1,i+1}|^2)\leqslant\frac{1}{2}$. It then follows that $\sum_{k\neq k_1}|q_{k,i+1}|^2\geqslant\frac{1}{2}$, showing that there exists k_2 ($k_2\neq k_1$): $|q_{k_2,i+1}|^2>1/(2n)$. Therefore, we have $|q_{k_1,i}|\geqslant 1/\sqrt{n}>1/\sqrt{2n}$ and $|q_{k_2,i+1}|>1/\sqrt{2n}$, $k_1\neq k_2$. □

Similar to Theorem 3.1, we show an upper bound for $|\gamma_k|/\|\mathbf{z}_k\|_2$, $k=k_1,k_2$, for the case of multiple eigenvalues $\lambda_i=\lambda_{i+1}$.

Theorem 5.2
Suppose that $P-\mu I$ is invertible and

$$(P-\mu I)\mathbf{z}_k=\gamma_k\mathbf{e}_k \quad \text{for } k=1,2,\ldots,n$$

DOI: 10.1002/nla

Let $\lambda_i = \lambda_{i+1}$ be a multiple eigenvalue of P. If $\max(|q_{k,i}|, |q_{k,i+1}|) \geqslant 1/\sqrt{2n}$, then

$$\frac{|\gamma_k|}{\|\mathbf{z}_k\|_2} \leqslant \sqrt{2n}|\lambda_i - \mu|$$

Proof
Since $P - \mu I$ is invertible, $\mathbf{z}_k = \gamma_k (P - \mu I)^{-1}\mathbf{e}_k$, implying

$$\begin{aligned}\|\mathbf{z}_k\|_2^2 &= |\gamma_k|^2 \mathbf{e}_k^{\mathrm{T}} Q(\Lambda^* - \mu I)^{-1}(\Lambda - \mu I)^{-1} Q^{\mathrm{H}}\mathbf{e}_k \\ &= |\gamma_k|^2 \sum_{j=1}^{n} \frac{|q_{k,j}|^2}{|\lambda_j - \mu|^2}\end{aligned}$$

Without loss of generality, we assume that $|q_{k,i}| = \max(|q_{k,i}|,\ |q_{k,i+1}|) \geqslant 1/\sqrt{2n}$. It then follows that

$$\begin{aligned}\frac{\|\mathbf{z}_k\|_2^2}{|\gamma_k|^2} &= \frac{|q_{k,i}|^2}{|\lambda_i - \mu|^2} + \frac{|q_{k,i+1}|^2}{|\lambda_i - \mu|^2} + \sum_{j \neq i, i+1} \frac{|q_{k,j}|^2}{|\lambda_j - \mu|^2} \\ &= \frac{|q_{k,i}|^2}{|\lambda_i - \mu|^2}\left(1 + \frac{|q_{k,i+1}|^2}{|q_{k,i}|^2} + \sum_{j \neq i, i+1} \frac{|q_{k,j}|^2}{|q_{k,i}|^2}\left|\frac{\lambda_i - \mu}{\lambda_j - \mu}\right|^2\right) \\ &= \frac{|q_{k,i}|^2}{|\lambda_i - \mu|^2}\left(1 + \frac{|q_{k,i+1}|^2}{|q_{k,i}|^2} + \frac{1 - |q_{k,i}|^2 - |q_{k,i+1}|^2}{|q_{k,i}|^2}\mathscr{A}_2\right) \qquad (8)\end{aligned}$$

where

$$\mathscr{A}_2 = \sum_{j \neq i, i+1} \omega_j \left|\frac{\lambda_i - \mu}{\lambda_j - \mu}\right|^2, \quad \sum_{j \neq i, i+1} \omega_j = 1 \text{ and } \omega_j \geqslant 0$$

is a weighted arithmetic mean of $\{|\lambda_i - \mu|^2/|\lambda_j - \mu|^2, j \neq i, i+1\}$. Thus,

$$0 < \mathscr{A}_2 < \left(\frac{|\lambda_i - \mu|}{\mathrm{gap}_2(\mu)}\right)^2$$

where $\mathrm{gap}_2(\mu) = \min_{j \neq i, i+1} |\lambda_j - \mu|$. It follows from (8) that

$$\begin{aligned}\frac{|\gamma_k|}{\|\mathbf{z}_k\|_2} &= \frac{|\lambda_i - \mu|}{|q_{k,i}|}\left(1 + \frac{|q_{k,i+1}|^2}{|q_{k,i}|^2} + \frac{1 - |q_{k,i}|^2 - |q_{k,i+1}|^2}{|q_{k,i}|^2}\mathscr{A}_2\right)^{-1/2} \\ &\leqslant \frac{|\lambda_i - \mu|}{|q_{k,i}|} \\ &\leqslant \sqrt{2n}|\lambda_i - \mu|\end{aligned}$$

DOI: 10.1002/nla

since

$$\frac{1-|q_{k,i}|^2-|q_{k,i+1}|^2}{|q_{k,i}|^2}\mathscr{A}_2\geqslant 0 \quad \text{and} \quad \frac{|q_{k,i+1}|^2}{|q_{k,i}|^2}\geqslant 0$$

which completes the proof. □

The above two theorems show that if we have a good approximation μ of a multiple eigenvalue $\lambda_i=\lambda_{i+1}$ of P, then there exist two indices k_1 and k_2, $k_1\neq k_2$, so that $|\gamma_{k_1}|/\|\mathbf{z}_{k_1}\|_2$ and $|\gamma_{k_2}|/\|\mathbf{z}_{k_2}\|_2$ produced by Theorem 5.2 are as small as $|\lambda_i-\mu|$ up to a factor of at most $\sqrt{2n}$. It then follows from Theorem 4.2 that the computed eigenvectors $\hat{\mathbf{q}}_i=\mathbf{z}_{k_1}/\|\mathbf{z}_{k_1}\|_2$ and $\hat{\mathbf{q}}_{i+1}=\mathbf{z}_{k_2}/\|\mathbf{z}_{k_2}\|_2$ satisfy

$$|\sin\theta(\hat{\mathbf{q}}_i,\mathbf{q}_i)|, |\sin\theta(\hat{\mathbf{q}}_{i+1},\mathbf{q}_{i+1})|\leqslant\frac{\sqrt{2n}|\lambda_i-\mu|}{\mathrm{gap}_2(\mu)}$$

where $\mathrm{gap}_2(\mu)=\min_{j\neq i,i+1}|\lambda_j-\mu|$.

In summary, we present the following algorithm for multiple singular values in $146n-149$ floating-point operations.

Algorithm 5.3

Given the Hermitian pentadiagonal matrix $P=TT^{\mathrm{H}}$ and $\mu=\hat{\lambda}_i=\hat{\lambda}_{i+1}$, a computed multiple eigenvalue of P, this algorithm computes approximations of the eigenvectors corresponding to the multiple eigenvalue $\lambda_i=\lambda_{i+1}$.

1. Compute the $\mathrm{LDL}^{\mathrm{H}}$ and $\mathrm{UDU}^{\mathrm{H}}$ decompositions (1) of $P-\mu I$ using Algorithms 2.1 and 2.2.
2. Applying Theorem 2.3, for $i=1,\ldots,n$, compute the twisted factorizations $P-\mu I=N_iD_iN_i^{\mathrm{H}}$ and find k_1 and k_2, $k_1\neq k_2$, such that $|\gamma_{k_1}|$ and $|\gamma_{k_2}|$ are the smallest among $|\gamma_i|$.
3. Solve for $\mathbf{z}_{k_1}$ in $N_{k_1}^{\mathrm{H}}\mathbf{z}_{k_1}=\mathbf{e}_{k_1}$ and $\mathbf{z}_{k_2}$ in $N_{k_2}^{\mathrm{H}}\mathbf{z}_{k_2}=\mathbf{e}_{k_2}$ using (6).
4. Set $\hat{\mathbf{q}}_i=\mathbf{z}_{k_1}/\|\mathbf{z}_{k_1}\|_2$ and $\hat{\mathbf{q}}_{i+1}=\mathbf{z}_{k_2}/\|\mathbf{z}_{k_2}\|_2$.

We have found that if the two indices k_1 and k_2 found in step 2 of Algorithm 5.3 are close, for example $k_2=k_1+1$, the eigenvectors $\mathbf{z}_{k_1}$ and $\mathbf{z}_{k_2}$ can be almost parallel. If k_1 and k_2 are far apart, then $\mathbf{z}_{k_1}$ and $\mathbf{z}_{k_2}$ are linearly independent, which is desirable. Thus, we propose the following strategy of choosing k_1 and k_2. First, we find k_1 such that $|\gamma_{k_1}|=\min_i|\gamma_i|$. Then, if there is only one isolated second smallest $|\gamma_{k_2}|$, then we choose k_2 as the second index. If there is a cluster of several equally small $|\gamma_i|$ next to $|\gamma_{k_1}|$, then among them we choose an index k_2 which is far apart from k_1.

In practice, however, eigenvalues can be clustered instead of exactly equal. It is known that the computed eigenvectors corresponding to clustered eigenvalues can lose orthogonality. In addition to the deflation and the strategy of choosing the indices k_1 and k_2 described above, we propose the following technique to deal with the issue of computing the Takagi vectors corresponding to clustered singular values. If an irreducible symmetric tridiagonal matrix has more than two tightly clustered singular values, we then group them in pairs. The left singular vectors corresponding to each tightly clustered singular value pair are computed using the strategy of choosing the indices k_1 and k_2 to improve the linear independency between them. The left singular vectors corresponding to all clustered singular values are orthogonalized by the modified Gram–Schmidt method when they are converted into the Takagi vectors. Our empirical results presented in Section 6 show that

DOI: 10.1002/nla

this combination of deflation, index selection and orthogonalization works well in dealing with clustered singular values.

6. NUMERICAL EXPERIMENTS

We implemented our twisted factorization Algorithms 3.2 and 5.3 incorporated with the refinement technique in Matlab and tested their accuracy and efficiency. Our experiments were carried out on a server with two 2.4 GHz Xeon CPUs, 1 GB RAM and an 80 GB disk. Random complex symmetric tridiagonal matrices with predetermined singular values were generated as follows. First, a vector $\mathbf{s}$ of n singular values was initialized. Then, a random unitary matrix U was generated by the QR decomposition of a random complex matrix. Finally, a complex symmetric tridiagonal matrix T was obtained by tridiagonalizing the product $U\Sigma U^{\mathrm{T}}$, where $\Sigma=\mathrm{diag}(\mathbf{s})$ using two-side Householder transformations. Thus, the complex symmetric tridiagonal T had predetermined singular values in $\mathbf{s}$.

We tested our algorithm on the following five matrices with various singular value clusters. Some singular values are even identical to working precision. We applied the techniques for dealing with clustered singular values described in the previous section. When an entry on the subdiagonal of a tridiagonal matrix was smaller than the square root of the machine precision, it was set to zero and the tridiagonal matrix was deflated. The singular values were computed using the implicit QR method [11]. The error in the computed singular vector $\hat{\mathbf{s}}$ was measured by

$$\Delta_v=\|\mathbf{s}-\hat{\mathbf{s}}\|_2$$

The error in the computed Takagi factorization was measured by

$$\Delta_t=\|\hat{V}\hat{\Sigma}\hat{V}^{\mathrm{T}}-T\|_2 \quad \text{where } \hat{\Sigma}=\mathrm{diag}(\hat{\mathbf{s}})$$

and $\hat{V}$ was computed by Algorithms 3.2 and 5.3. The orthogonality of the computed Takagi vectors was measured by

$$\Delta_o=\|\hat{V}\hat{V}^{\mathrm{H}}-I\|_2$$

Example 1 (*Nested clusters, Dhillon and Parlett* [*15*])
A 13×13 real symmetric tridiagonal matrix with spectrum, ε, 1, 1 ± 10^{-15}, 1 ± 10^{-12}, 1 ± 10^{-9}, 1 ± 10^{-6}, 1 ± 10^{-3} and 2, where ε is the machine precision.

Example 2 (*Wilkinson matrix, Dhillon and Parlett* [*15*])
The 101×101 Wilkinson matrix that has various eigenvalue clusters of pairs. The rightmost cluster is the tightest, with λ_{100} and λ_{101} identical to working accuracy.

Example 3 (*Uniform distribution* ($\sqrt{\varepsilon}$ *apart*))
A 400×400 random complex symmetric tridiagonal matrix with singular values:

$$\sigma_1=\varepsilon,\ \sigma_j=1+(j-1)\sqrt{\varepsilon},\ j=2,\ldots,399 \quad \text{and} \quad \sigma_{400}=2$$

where ε is the machine precision.

DOI: 10.1002/nla

Example 4 (*Uniform distribution* (ε *to 1*))
A 400×400 random complex symmetric tridiagonal matrix with singular values evenly spaced between ε and 1:

$$\sigma_j=\varepsilon+(j-1)\tau,\quad j=1,2,\ldots,400$$

where $\tau=(1-\varepsilon)/399$.

Example 5 (*Clustered at 1*)
Example 5A 400×400 random complex symmetric tridiagonal matrix with singular values:

$$\sigma_i=1+v\varepsilon\quad\text{for } i=1,\ldots,388 \text{ and } \sigma_{400}=\varepsilon$$

where v is normally distributed between -1 and 1.

Table I shows the errors in the computed SSVDs of the above five examples.

For performance, we tested our algorithm on random complex symmetric tridiagonal matrices of five different sizes. For each size, we generated five complex symmetric tridiagonal matrices. Their singular values were random numbers uniformly distributed between 0 and 1. We ran our twisted factorization (Twist) method, divide-and-conquer (DAC) method [12] and the Matlab's `svd` function. In the divide-and-conquer method, when the size of a submatrix T is less than or equal to 10, its Takagi factorization is computed directly by the implicit QR method [11]. For fair comparison, we used Matlab's `svd` function, instead of the implicit QR method used in the previous experiment, to compute the singular values for our twisted factorization (Twist) method. Table II shows the average running time and the average factorization error Δ_t of the five matrices of same size. The results in Table II demonstrate that our method is significantly more efficient

Table I. Errors in the computed SSVDs of the five testing matrices with clustered singular values.

Example	Δ_o	Δ_t	Δ_v
Nested	9.8586E−11	9.8586E−11	6.6400E−11
Wilkinson	1.0109E−10	9.0317E−10	9.8164E−13
$\sqrt{\varepsilon}$ apart	6.5221E−15	8.2247E−13	1.5264E−13
ε to 1	1.7985E−12	8.1968E−13	8.0259E−14
Clustered at 1	4.6373E−16	1.5076E−14	5.4858E−14

Table II. The performance and accuracy comparison of the Twisted factorization (Twist) method, divide-and-conquer (DAC) method and SVD subroutine in Matlab (SVD).

	Running time (s)			Δ_t		
Matrix size	Twist	DAC	SVD	Twist	DAC	SVD
100	0.41	0.65	0.06	5.3425E−13	9.5324E−14	5.6786E−15
200	1.29	2.62	0.90	6.2342E−12	4.1002E−13	1.5234E−14
400	4.67	9.32	6.73	1.0123E−11	1.1023E−11	3.2111E−14
800	18.08	44.12	79.99	3.2123E−11	3.9821E−11	3.7456E−14
1600	70.07	253.62	1210.20	5.2398E−11	5.3252E−11	5.3421E−14

DOI: 10.1002/nla

than both the divide-and-conquer method and the Matlab's `svd` function even for matrices of moderately large size.

REFERENCES

1. Golub GH, Van Loan CF. *Matrix Computations* (3rd edn). The Johns Hopkins University Press: Baltimore, MD, 1996.
2. Horn RA, Johnson CR. *Matrix Analysis*. Cambridge University Press: Cambridge, 1996.
3. Bunse-Gerstner A, Gragg WB. Singular value decomposition of complex symmetric matrices. *Journal of Computational and Applied Mathematics* 1988; **21**:41–54.
4. Pommerenke C. *Univalent Functions*. Vandenhoeck & Ruprecht: Göttingen, 1975.
5. Trefethen LN. Near-circularity of the curve in complex Chebyshev approximation. *Journal of Approximation Theory* 1981; **31**:344–367.
6. Comon P. Independent component analysis, a new concept? *Signal Processing* 1994; **36**:287–314.
7. Anderson E, Bai Z, Bischof C, Demmel J, Dongarra J, Du Croz J, Greenbaum A, Hammarling S, McKenney A, Ostrouchov S, Sorensen D. *LAPACK Users' Guide* (3rd edn). SIAM: Philadelphia, PA, 1999.
8. Qiao S. Orthogonalization techniques for the Lanczos tridiagonalization of complex symmetric matrices. In *Advanced Signal Processing Algorithms*, *Architectures and Implementations XIV*, Luk FT (ed.). *Proceedings of SPIE* 2004; **5559**:423–434.
9. Xu W, Qiao S. A fast SVD algorithm for square Hankel matrices. *Linear Algebra and its Applications* 2008; **428**:550–563.
10. Qiao S, Liu G, Xu W. Block Lanczos tridiagonalization of complex symmetric matrices. In *Advanced Signal Processing Algorithms*, *Architectures and Implementations XV*, Luk FT (ed.). *Proceedings of SPIE* 2005; **5910**: 285–295.
11. Luk FT, Qiao S. A fast singular value algorithm for Hankel matrices. In *Fast Algorithms for Structured Matrices*: *Theory and Applications*, Olshevsky V (ed.). Contemporary Mathematics, vol. 323. American Mathematical Society: Providence, RI, 2003; 169–177.
12. Xu W, Qiao S. A divide-and-conquer method for the Takagi factorization. *SIAM Journal on Matrix Analysis and Applications* 2008; **30**:142–153.
13. Dhillon JS. A new $O(n^2)$ algorithm for the symmetric tridiagonal eigenvalue/eigenvector problem. *Ph.D. Thesis*, UCB, Berkeley, CA, 1997.
14. Davis C, Kahan WM. The rotation of eigenvectors by a perturbation III. *SIAM Journal on Numerical Analysis* 1970; **7**:1–46.
15. Dhillon JS, Parlett BP. Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices. *Linear Algebra and its Applications* 2004; **387**:1–28.

DOI: 10.1002/nla