# COMPUTING THE SVD OF A GENERAL MATRIX PRODUCT/QUOTIENT* 

GENE GOLUB ${ }^{\dagger}$, KNUT SØLNA ${ }^{\ddagger}$, AND PAUL VAN DOOREN§


#### Abstract

In this paper we derive a new algorithm for constructing a unitary decomposition of a sequence of matrices in product or quotient form. The unitary decomposition requires only unitary left and right transformations on the individual matrices and amounts to computing the generalized singular value decomposition of the sequence. The proposed algorithm is related to the classical Golub-Kahan procedure for computing the singular value decomposition (SVD) of a single matrix in that it constructs a bidiagonal form of the sequence as an intermediate result. When applied to two matrices this new method is an alternative way of computing the quotient and product SVD and is more economical than current methods.


Key words. mumerical methods, generalized singular values, products of matrices, quotients of matrices

## AMS subject classification. 65F15

## PII. S0895479897325578

Introduction. The two basic unitary decompositions of a matrix $A$ yielding some spectral information are the Schur form $A=U T U^{*}$-where $U$ is unitary and $T$ is upper triangular - and the singular value decomposition (SVD) $A=U \Sigma V^{*}$-where $U$ and $V$ are unitary and $\Sigma$ is diagonal (for the latter $A$ does not need to be square). It is interesting to note that both forms are usually computed by a $Q R$-like iteration [7]. The SVD algorithm of Golub-Kahan [6] is indeed an implicit $Q R$ algorithm applied to the Hermitian matrix $A^{*} A$. When looking at unitary decompositions involving two matrices, say, $A$ and $B$, a similar implicit algorithm was given in [10] and is known as the $Q Z$ algorithm. It computes $A=Q T_{a} Z^{*}$ and $B=Q T_{b} Z^{*}$, where $Q$ and $Z$ are unitary and $T_{a}$ and $T_{b}$ are upper triangular. This algorithm is in fact the $Q R$ algorithm again performed implicitly on the quotient $B^{-1} A$. The corresponding decomposition is therefore also known as the generalized Schur form.

When considering the generalized SVD of two matrices, appearing as a quotient $B^{-1} A$ or a product $B A$, the currently used algorithm is not of $Q R$ type but of a Jacobi type. The reason for this choice is that Jacobi methods easily extend to products and quotients. Unfortunately, the Jacobi algorithm typically has a (moderately) higher complexity than the $Q R$ algorithm. Yet, so far, nobody proposed an implicit $Q R$-like method for the SVD of a product or quotient of two matrices.

[^0]In this paper we show that, in fact, such an implicit algorithm is easy to derive and that it even extends straightforwardly to sequences of products/quotients of several matrices. Moreover, the complexity will be shown to be lower than for the corresponding Jacobi-like methods.

1. Implicit SVD. Consider the problem of computing the SVD of a matrix $A$ that is an expression of the following type:

$$
\begin{equation*}
A=A_{K}^{s_{K}} \cdots A_{2}^{s_{2}} \cdot A_{1}^{s_{1}} \tag{1}
\end{equation*}
$$

where $s_{i}= \pm 1$, i.e., a sequence of products or quotients of matrices. For simplicity we assume that the $A_{i}$ matrices are square $n \times n$ and invertible. It was pointed out in [3] that one can always perform a preliminary $Q R$-like reduction that extracts from a sequence of matrices with compatible dimensions another sequence of square invertible matrices with the same generalized singular values as the original sequence. We refer to [3] for the details of this reduction and will treat here only the case of square invertible matrices. While it is clear that one has to perform left and right transformations on $A$ to get $U^{*} A V=\Sigma$, these transformations will affect only $A_{K}$ and $A_{1}$. Beyond this, one can insert an expression $Q_{i}^{*} Q_{i}=I_{n}$ between every pair $A_{i+1}^{s_{i+1}} A_{i}^{s_{i}}$ in (1). If we also define $Q_{K} \doteq U$ and $Q_{0} \doteq V$, we arrive at the following expression:

$$
\begin{equation*}
U^{*} A V=\left(Q_{K}^{*} A_{K}^{s_{K}} Q_{K-1}\right) \cdots\left(Q_{2}^{*} A_{2}^{s_{2}} Q_{1}\right) \cdot\left(Q_{1}^{*} A_{1}^{s_{1}} Q_{0}\right) \tag{2}
\end{equation*}
$$

With the degrees of freedom present in these $K+1$ unitary transformations $Q_{i}$ at hand, one can now choose each expression $Q_{i}^{*} A_{i}^{s_{i}} Q_{i-1}$ to be upper triangular. Note that the expression $Q_{i}^{*} A_{i}^{s_{i}} Q_{i-1}=T_{i}^{s_{i}}$ with $T_{i}$ upper triangular can be rewritten as

$$
\begin{equation*}
Q_{i}^{*} A_{i} Q_{i-1}=T_{i} \quad \text { for } s_{i}=1, \quad Q_{j-1}^{*} A_{j} Q_{j}=T_{j} \quad \text { for } s_{j}=-1 \tag{3}
\end{equation*}
$$

From the construction of a normal $Q R$ decomposition, it is clear that while making the matrix $A$ upper triangular, this "freezes" only one matrix $Q_{i}$ per matrix $A_{i}$. The remaining unitary matrix leaves enough freedom to finally diagonalize the matrix $A$ as well. Since (2) computes the singular values of (1), it is clear that such a result can be obtained only by an iterative procedure. On the other hand, one intermediate form that is used in the Golub-Kahan SVD algorithm [6] is the bidiagonalization of $A$ and this can be obtained in a finite recurrence. We show in the next section that the matrices $Q_{i}$ in (2) can be constructed in a finite number of steps in order to obtain a bidiagonal $Q_{K}^{*} A Q_{0}$ in (2). In carrying out this task one should try to do as much as possible implicitly. Moreover, one would like the total complexity of the algorithm to be comparable to, or less than, the cost of $K$ singular value decompositions. This means that the complexity should be $\mathrm{O}\left(K n^{3}\right)$ for the whole process.
2. Implicit bidiagonalization. We now derive such an implicit reduction to bidiagonal form. Below $\mathcal{H}(i, j)$ denotes the group of Householder transformations having $(i, j)$ as the range of rows/columns they operate on. Similarly $\mathcal{G}(i, i+1)$ denotes the group of Givens transformations operating on rows/columns $i$ and $i+1$. We first consider the case where all $s_{i}=1$. We thus have only a product of matrices $A_{i}$ and in order to illustrate the procedure we show its evolution operating on a product of three matrices only, i.e., $A_{3} A_{2} A_{1}$. Below is a sequence of displays of the matrix product that illustrates the evolution of the bidiagonal reduction. Each display indicates the pattern of zeros (" 0 ") and nonzeros (" $x$ ") in the three matrices.

First perform a Householder transformation $Q_{1}^{(1)} \in \mathcal{H}(1, n)$ on the rows of $A_{1}$ and the columns of $A_{2}$. Choose $Q_{1}^{(1)}$ to annihilate all but one element in the first column of $A_{1}$ :

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]
$$

Then perform a Householder transformation $Q_{2}^{(1)} \in \mathcal{H}(1, n)$ on the rows of $A_{2}$ and the columns of $A_{3}$. Choose $Q_{2}^{(1)}$ to annihilate all but one element in the first column of $A_{2}$ :

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]
$$

Then perform a Householder transformation $Q_{3}^{(1)} \in \mathcal{H}(1, n)$ on the rows of $A_{3}$. Choose $Q_{3}^{(1)}$ to annihilate all but one element in the first column of $A_{3}$ :

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]
$$

Note that this third transformation yields the same form also for the product of the three matrices:

At this stage we are interested in the first row of this product (indicated by boldface x's above). This row can be constructed as the product of the first row of $A_{3}$ with the matrices to the right of it, and this requires only $\mathrm{O}\left(K n^{2}\right)$ flops. Once this row is constructed we can find a Householder transformation $Q_{0}^{(1)} \in \mathcal{H}(2, n)$ operating on the last $(n-1)$ elements which annihilates all but two elements (the colon, ":", is used as it is in MATLAB):

$$
A_{3}(1,:) A_{2} A_{1} Q_{0}^{(1)}=\left[\begin{array}{lllll}
x & x & 0 & 0 & 0 \tag{4}
\end{array}\right] .
$$

This transformation is then applied to $A_{1}$ only and completes the first stage of the bidiagonalization since

$$
Q_{K}^{(1) *} A Q_{0}^{(1)}=\left[\begin{array}{ccccc}
x & x & 0 & 0 & 0 \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right]
$$

The second stage of the bidiagonalization is analogous to the first; it differs only in that the transformations operate only on rows/columns 2 to n. The Householder transformations $Q_{i}^{(2)} \in \mathcal{H}(2, n)$ for $1 \leq i \leq 3$ are chosen to eliminate elements 3 to $n$ in the second columns of $A_{i}$ in the manner described above. The transformation $Q_{0}^{(2)} \in \mathcal{H}(3, n)$ operates on the last $(n-2)$ elements of the second row of the product and annihilates all but two elements:

$$
A_{3}(2,:) A_{2} A_{1} Q_{0}^{(2)}=\left[\begin{array}{lllll}
0 & x & x & 0 & 0 \tag{5}
\end{array}\right]
$$

This transformation is applied to $A_{1}$ only, completing the second step of the bidiagonalization of $A$ :

$$
Q_{K}^{(2) *} Q_{K}^{(1) *} A Q_{0}^{(1)} Q_{0}^{(2)}=\left[\begin{array}{ccccc}
x & x & 0 & 0 & 0 \\
0 & x & x & 0 & 0 \\
0 & 0 & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & x & x & x
\end{array}\right]
$$

It is now clear from the context how to proceed further with this algorithm to obtain after $n-1$ stages:

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right]=\left[\begin{array}{lllll}
x & x & 0 & 0 & 0 \\
0 & x & x & 0 & 0 \\
0 & 0 & x & x & 0 \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right] .
$$

Note that we never construct the whole product $A=A_{3} A_{2} A_{1}$, but rather compute one of its rows when needed for constructing the transformations $Q_{0}^{(i)}$. The only matrices that are kept in memory and updated are the $A_{i}$ matrices and possibly $Q_{K}$ and $Q_{0}$ if we require the singular vectors of $A$ afterwards.

The complexity of this bidiagonalization step is easy to evaluate. Each matrix $A_{i}$ gets pre- and postmultiplied with essentially $n$ Householder transformations of decreasing range. For updating all $A_{i}$ we therefore need $10 K n^{3} / 3$ flops, and for updating $Q_{K}$ and $Q_{0}$ we need $4 n^{3}$ flops. For constructing the required row vectors of $A$ we need $(K-1) n^{3} / 3$ flops. Overall we thus need on the order of $11 K n^{3} / 3$ flops for the construction of the triangular $T_{i}$ and $4 n^{3}$ for the outer transformations $Q_{K}$ and $Q_{0}$. Essentially this is $11 n^{3} / 3$ flops per updated matrix.

If we now have some of the $s_{i}=-1$, we cannot use Householder transformations anymore on all matrices. Indeed, in order to construct the rows of $A$ when needed, the matrices $A_{i}$ for which $s_{i}=-1$ have to be triangularized first, say, with a $Q R$ factorization. The $Q R$ factorization is performed in an initial step and uses Householder transformations. From there on the same procedure as above is followed, but this implies using Givens rotations in certain steps of the bidiagonalization. For simplicity, we illustrate this on a matrix $A=A_{3} A_{2}^{-1} A_{1}$. We first apply a left transformation
$Q_{2}^{(0)}$ (using a sequence of Householder transformations) that triangularizes $A_{2}$ from the left and also apply this to the rows of $A_{1}$. The resulting triple then has the form

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x
\end{array}\right]
$$

We then apply a unitary transformation $Q_{1}^{(1)}$ to the rows of $A_{1}$ to eliminate elements 2 to $n$ in its first column using Givens transformations $G_{1} \in \mathcal{G}(n-1, n)$ until $G_{n-1} \in \mathcal{G}(1,2)$. Below we indicate in which order these zeros are created in the first column of $A_{1}$ by their index:

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x \\
x & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0_{4} & x & x & x & x \\
0_{3} & x & x & x & x \\
0_{2} & x & x & x & x \\
0_{1} & x & x & x & x
\end{array}\right]
$$

These transformations also have to be applied to the left of $A_{2}$, but the use of Givens rotations allows us to update the triangularized matrix $A_{2}$, while keeping it upper triangular: each time a Givens rotation applied to the left of $A_{2}$ destroys its triangular form, another Givens rotation is applied to the right of $A_{2}$ in order to restore its triangular form. (The same technique is used, for instance, in keeping the $B$ matrix upper triangular in the $Q Z$ algorithm applied to $B^{-1} A$.) Let $Q_{2}^{(1)}$ be the product of the Givens rotations applied to the right of $A_{2}$, then $Q_{2}^{(1)}$ also has to be applied to the right of $A_{3}$. Finally, for the column transformation $Q_{3}^{(1)}$ of $A_{3}$ eliminating elements 2 to $n$ of its first column, we can again use a Householder transformation $H_{5} \in \mathcal{H}(1, n)$. After this fifth transformation, the resulting triple has the form

$$
\left[\begin{array}{lllll}
x & x & x & x & x \\
0_{5} & x & x & x & x \\
0_{5} & x & x & x & x \\
0_{5} & x & x & x & x \\
0_{5} & x & x & x & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{array}\right]\left[\begin{array}{lllll}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x \\
0 & x & x & x & x
\end{array}\right],
$$

which clearly has a first column with only its leading element different from 0 . Its first row can easily be constructed, and we then apply a Householder transformation $Q_{0}^{(1)} \in \mathcal{H}(2, n)$ annihilating all but two elements as indicated in (4). This completes the first stage of the bidiagonalization of $A=A_{3} A_{2}^{-1} A_{1}$. Subsequent steps are similar but operate on matrices of decreasing dimensions.

Notice that all transformations $Q_{i}$ and $Q_{i-1}$ applied to a matrix $A_{i}$ with index $s_{i}=-1$ have to be of Givens type, which require more flops than Householder transformations for the same number of annihilated elements. So the more negative exponents we have, the more expensive the overall algorithm becomes. Without loss of generality, we can assume that at most half of the indices are equal to -1 , since otherwise we can compute the SVD of $A^{-1}$ rather than that of $A$ (all matrices $A_{i}$ were assumed to be invertible). The situation with the highest computational complexity is thus when every other index $s_{i}$ is negative, since then all transformations
but one have to be of Givens type. Let us analyze the case where $K$ is even and $s_{2 i}=-1, s_{2 i-1}=1, i=1, \ldots, \frac{K}{2}$. The preliminary $Q R$ reduction of the matrices $A_{2 i}$ requires $\frac{4}{3} n^{3} \frac{K}{2}$ flops and $2 n^{3} \frac{K}{2}$ flops for also updating the matrices $A_{2 i-1}$ with these transformations. From there on, each matrix undergoes $\frac{n(n-1)}{2}$ Givens rotations on the left and on the right. For the triangular matrices $A_{2 i}$ this requires a total of $3 n^{2} \frac{K}{2}$ flops, whereas for the (originally dense) matrices $A_{2 i-1}$ this requires a total of $5 n^{3} \frac{K}{2}$ flops. For constructing the required row vectors of $A$ we need $(K-1) n^{3} / 3$ flops as before. Finally, updating the matrix $Q_{0}$ via Givens transformations and $Q_{K}$ via Householder transformations requires $2 n^{2}$ and $3 n^{3}$ flops, respectively. The worst-case complexity of the general case is thus $6 n^{3} K$ flops for obtaining the triangular matrices $T_{i}$ and $5 n^{3}$ flops for the outer transformations $Q_{0}$ and $Q_{K}$. This is about $60 \%$ more than for the product case.
3. Error analysis. In the previous section we showed how to obtain an estimate of a bidiagonal decomposition of the matrix product/quotient. We now turn to the problem of obtaining accurate estimates of the singular values. This warrants a discussion of the errors committed in the bidiagonalization step.

The use of Householder and Givens transformations for all operations in the bidiagonalization step guarantees that the obtained matrices $T_{i}$ in fact correspond to slightly perturbed data as follows:

$$
\text { (6) } T_{i}=Q_{i}^{*}\left(A_{i}+\delta A_{i}\right) Q_{i-1}, s_{i}=1, \quad T_{j}=Q_{j-1}^{*}\left(A_{j}+\delta A_{j}\right) Q_{j}, s_{j}=-1
$$

where

$$
\begin{equation*}
\left\|\delta A_{i}\right\| \leq \epsilon c_{n}\left\|A_{i}\right\|, \quad\left\|Q_{i}^{*} Q_{i}-I_{n}\right\| \leq \epsilon d_{n} \tag{7}
\end{equation*}
$$

with $\epsilon$ the machine precision and $c_{n}, d_{n}$ moderate constants depending on the problem size $n$. This is obvious since each element transformed to zero can indeed be put equal to zero without affecting the $\epsilon$ bound (see [11], [7]).

The situation is different for the elements of $A$ since they are not stored explicitly in the computer. How does one proceed further to compute the generalized singular values of $A$ ? Once the triangular matrices $T_{i}$ are obtained, it is easy and cheap to reconstruct the bidiagonal:

$$
T_{K}^{s_{k}} \cdots T_{2}^{s_{2}} \cdot T_{1}^{s_{1}}=B=\left[\begin{array}{ccccc}
q_{1} & e_{2} & o_{1,3} & \cdots & o_{1, n}  \tag{8}\\
& q_{2} & e_{3} & \ddots & \vdots \\
& & \ddots & \ddots & o_{n-2, n} \\
& & & \ddots & e_{n} \\
& & & & q_{n}
\end{array}\right]
$$

and then compute the singular values of the bidiagonal in a standard way. The diagonal elements $q_{i}$ are indeed just a product of the corresponding diagonal elements of the $T_{j}$ matrices, possibly inverted:

$$
q_{i}=t_{K_{i, i}}^{s_{K}} \cdots t_{2_{i, i}}^{s_{2}} \cdot t_{1_{i, i}}^{s_{1}}
$$

and the off-diagonal elements $e_{i}$ can be computed from the corresponding $2 \times 2$ diagonal blocks (with index $i-1$ and $i$ ) of the $T_{j}$ matrices. It is clear that the $q_{i}$ can be computed in a backward stable way since all errors can be superimposed on the
diagonal elements $t_{j_{i, i}}$ of the matrices $T_{j}$. For the errors incurred when computing the $e_{i}$ one needs a more detailed analysis. We show below that the backward errors can be superimposed on the off-diagonal elements $t_{j_{i-1, i}}$ of $T_{j}$ without creating any conflicts with previously constructed backward errors, and we derive bounds for these backward errors. From the vector recurrence

$$
\left[\begin{array}{l}
e  \tag{9}\\
q
\end{array}\right]:=\left[\begin{array}{ll}
t_{j_{i-1, i-1}} & t_{j_{i-1, i}} \\
0 & t_{j_{i, i}}
\end{array}\right]^{s_{j}} \cdot\left[\begin{array}{l}
e \\
q
\end{array}\right]
$$

we easily derive the following algorithm used for computing $q_{i}$ and $e_{i}$ for $i=1, \ldots, n$. $q:=1 ; e:=0$;
for $j=1: K$

$$
\begin{aligned}
& \text { if } s_{j}=1 \text {, then } e:=e * t_{j_{i-1, i-1}}+q * t_{j_{i-1, i}} ; q:=q * t_{j_{i, i}} \text {; } \\
& \text { else } q:=q / t_{j_{i, i}} ; e:=\left(e-q * t_{j_{i-1, i}}\right) / t_{j_{i-1, i-1}} ;
\end{aligned}
$$

end
$q_{i}:=q ; e_{i}:=e ;$
Note that for $i=1$ the same recurrence holds without the expressions involving $e$. From these recurrences it is clear that the calculation of $q_{i}$ involves one flop per step $j$ and hence a total of $K$ rounding errors which can be superimposed on the diagonal elements $t_{j_{i, i}}$ :

$$
\begin{align*}
q_{i} & =\operatorname{comp}\left(t_{K_{i_{i}}}^{s_{K}} \cdots t_{1_{i, i}}^{s_{1}}\right) \\
& =\bar{t}_{K_{i, i}}^{s_{K}} \cdots \bar{t}_{1_{i, i}}^{s_{1}} \tag{10}
\end{align*} \quad \text { with } \quad \bar{t}_{j_{i, i}}=t_{j_{i, i}}\left(1+\epsilon_{i, j}\right), \quad\left|\epsilon_{i, j}\right|<\epsilon
$$

with comp denoting a floating point operator. For the calculation of $e_{i}$ there are 3 flops per step $j$ and hence a total of $3 K$ roundings which have to be superimposed on the $t_{j_{i-1, i}}$ elements. Fortunately, $e_{j}$ is a sum of $K$ terms which contain each a different element $t_{j_{i-1, i}}$ as a factor. We illustrate this for $K=4$ and $s_{j}=1$, highlighting the relevant elements:

$$
\begin{align*}
e_{i}=\operatorname{comp} & \left(\mathbf{t}_{4_{i-1, i}} \cdot t_{3_{i, i}} \cdot t_{2_{i, i}} \cdot t_{1_{i, i}}\right. \\
& +t_{4_{i-1, i-1}} \cdot \mathbf{t}_{3_{i-1, i}} \cdot t_{2_{i, i}} \cdot t_{1_{i, i}}  \tag{11}\\
& +t_{4_{i-1, i-1}} \cdot t_{3_{i-1, i-1}} \cdot \mathbf{t}_{2_{i-1, i}} \cdot t_{1_{i, i}} \\
& \left.+t_{4_{i-1, i-1}} \cdot t_{3_{i-1, i-1}} \cdot t_{2_{i-1, i-1}} \cdot \mathbf{t}_{1_{i-1, i}}\right)
\end{align*}
$$

The $3 K$ rounding errors can thus easily be superimposed on these different elements $t_{j_{i-1, i}}, j=1, \ldots, K$. But since we have already superimposed errors on the alldiagonal elements $t_{j_{i, i}}$ we have to add these perturbations here as well. For $s_{j}=1$ we thus have

$$
\begin{align*}
e_{i}= & \bar{t}_{4_{i-1, i}} \cdot \bar{t}_{3_{i, i}} \cdot \bar{t}_{2_{i, i}} \cdot \bar{t}_{1_{i, i}} \\
& +\bar{t}_{4_{i-1, i-1}} \cdot \bar{t}_{3_{i-1, i}} \cdot \bar{t}_{2_{i, i}} \cdot \bar{t}_{1_{i, i}}  \tag{12}\\
& +\bar{t}_{4_{i-1, i-1}} \cdot \bar{t}_{3_{i-1, i-1}} \cdot \bar{t}_{t_{i-1, i}} \cdot \bar{t}_{1_{i, i}} \\
& +\bar{t}_{4_{i-1, i-1}} \cdot \bar{t}_{3_{i-1, i-1}} \cdot \bar{t}_{2_{i-1, i-1}} \cdot \bar{t}_{1_{i-1, i}}
\end{align*}
$$

where $(K-1)$ additional roundings are induced for each factor. Therefore, we have $\bar{t}_{j_{i-1, i}}=t_{j_{i-1, i}}\left(1+\eta_{i, j}\right), \quad\left|\eta_{i, j}\right|<(4 K-1) \epsilon /(1-(4 K-1) \epsilon)$. When some of the $s_{j}=-1$ the above expression is similar: the $t_{j_{i, i}}$ then appear as inverses, some + signs change to - signs, and an additional factor $1 /\left(t_{j_{i-1, i-1}} t_{j_{i, i}}\right)$ appears in the $j$ th term if $s_{j}=-1$. So in the worst case $(K+1)$ additional roundings are introduced for each factor and the obtained bound is then $\left|\eta_{i, j}\right|<(4 K+1) \epsilon /(1-(4 K+1) \epsilon)$. In the worst case the errors yield a backward perturbation $\left\|\delta T_{j}\right\|$ which is thus bounded
by $5 K \epsilon\left\|T_{j}\right\|$ and hence much smaller than the errors $\delta A_{j}$ incurred in the triangularization process. The perturbation effect of computing the elements $q_{i}$ and $e_{i}$ is thus negligible compared to that of the triangularization. We thus showed that the computed bidiagonal corresponds exactly to the bidiagonal of the product of slightly perturbed triangular matrices $T_{j}$, that in turn satisfy the bounds (6)-(7). Unfortunately, nothing of the kind can be guaranteed for the elements $o_{i, j}$ in (8), which are supposed to be zero in exact arithmetic. Notice that the element $e_{i+1}$ is obtained as the norm of the vector on which a Householder transformation is applied:

$$
\left|e_{i+1}\right|=\left\|T_{K}^{s_{K}}(i, i: n) T_{K-1}^{s_{K-1}}(i: n, i: n) \cdots T_{1}^{s_{1}}(i: n, i+1: n)\right\|,
$$

where we used the MATLAB notation for subarrays: $T_{1}^{s_{1}}(i: n, i+1: n)$ is thus the submatrix of $T_{1}^{s_{1}}$ with row indices $i$ to $n$ and column indices $i+1$ to $n$. If all $s_{i}=1$ we can obtain by straightforward perturbation results of matrix vector products, a bound of the type

$$
\left|o_{i, j}\right| \leq \epsilon c_{n}\left\|T_{K}(i, i: n)\right\| \cdot\left\|T_{K-1}(i: n, i: n)\right\| \cdots\left\|T_{1}(i: n, i: n)\right\| .
$$

If not all $s_{i}=1$ we need to also use perturbation results of solutions of systems of equations, since we need to evaluate the last $n-i$ components of the vector $e_{i}^{*} T_{K}^{s_{K}}(i$ : $n, i: n) T_{K-1}^{s_{K-1}}(i: n, i: n) \cdots T_{1}^{s_{1}}(i: n, i: n)$ and this requires a solution of a triangular system of equations each time a power $s_{j}=-1$ is encountered. In this case the bound would become

$$
\left|o_{i, j}\right| \leq \epsilon c_{n}\left\|T_{K}^{s_{K}}(i, i: n)\right\| \cdot\left\|T_{K-1}^{s_{K-1}}(i: n, i: n)\right\| \cdots\left\|T_{1}^{s_{1}}(i: n, i: n)\right\| \kappa,
$$

where $\kappa$ is the product of all condition numbers of the inverted triangular systems (and hence much larger than 1). These are much weaker bounds than asking the offdiagonal elements of $A$ to be $\epsilon$ smaller than the ones on the bidiagonal. This would be the case, e.g., if instead we had

$$
\left|o_{i, j}\right| \leq \epsilon c_{n}\left\|T_{K}^{s_{K}}(i, i: n) T_{K-1}^{s_{K-1}}(i: n, i: n) \cdots T_{1}^{s_{1}}(i: n, i+1: n)\right\|=\epsilon c_{n}\left|e_{i+1}\right|
$$

Such a bound would guarantee high relative accuracy in the singular values computed from the bidiagonal only [4]. Hence, this is the kind of result one would hope for. These two bounds can in fact be very different when significant cancellations occur between the individual matrices, e.g., if

$$
\|A\| \ll\left\|A_{K}^{s_{K}}\right\| \cdots\left\|A_{2}^{s_{2}}\right\| \cdot\left\|A_{1}^{s_{1}}\right\| .
$$

One could observe that the bidiagonalization procedure is in fact a Lanczos procedure [6]. Therefore, there is a tendency to find first the dominant directions of the expression $A_{K}^{s_{K}} \cdots A_{1}^{s_{1}}$ and hence also those directions where there is less cancellation between the different factors. We will see in the examples below that such a phenomenon indeed occurs which is a plausible explanation for the good accuracy obtained. One way to test the performance of the algorithm in cases with very small singular values is to generate powers of a symmetric matrix $A=S^{K}$. The singular values will be the powers of the absolute values of the eigenvalues of $S$ :

$$
\sigma_{i}(A)=\left|\lambda_{i}(S)\right|^{K}
$$

and hence will have a large dynamic range. The same should be true for the bidiagonal of $A$ and the size of the $o_{i, j}$ will then become critical for the accuracy of the singular
values when computed from the bidiagonal elements $q_{i}, e_{i}$. This and several other examples are discussed in the next section. There we observe a very high relative accuracy even for the smallest singular values. The only explanation we can give for this is that as the bidiagonalization proceeds, it progressively finds the largest singular values first and creates submatrices that are of smaller norm. These then do not really have cancellation between them, but instead the decreasing size of the bidiagonal elements is the result of decreasing elements in each transformed matrix $A_{i}$. In other words, a grading is created in each of the transformed matrices. We believe this could be explained by the fact that the bidiagonalization is a Lanczos procedure and that such grading is often observed there when the matrix has a large dynamic range of eigenvalues. In practice it is of course always possible to evaluate bounds for the elements $\left|o_{i, j}\right|$ and thereby obtain estimates of the accuracy of the computed singular values.

The consequence of the above is that the singular values of such sequences can be computed (or better, "estimated") at high relative accuracy from the bidiagonal only! Notice that the bidiagonalization requires 4 to $6 K n^{3}$ flops but that the subsequent SVD of the bidiagonal is essentially free since it is $\mathrm{O}\left(n^{2}\right)$.
4. Singular vectors and iterative refinement. If one wants the singular vectors as well as the singular values at a guaranteed accuracy, one can start from the bidiagonal $B$ as follows. First compute the bidiagonal,

$$
B=Q_{K}^{*} A Q_{0}=T_{K}^{s_{K}} \cdots T_{2}^{s_{2}} \cdot T_{1}^{s_{1}}
$$

and then the SVD of $B$,

$$
B=U \Sigma V^{*}
$$

where we choose the diagonal elements of $\Sigma$ to be ordered in decreasing order. We then proceed by propagating the transformation $U$ (or $V$ ) and updating each $T_{i}$ so that they remain upper triangular. Since the neglected elements $o_{i, j}$ were small, the new form

$$
\hat{Q}_{K}^{*} A \hat{Q}_{0}=\hat{T}_{K}^{s_{K}} \cdots \hat{T}_{2}^{s_{2}} \cdot \hat{T}_{1}^{s_{1}}
$$

will be upper triangular, and nearly diagonal. This is the ideal situation to apply one sweep of Kogbetliantz's algorithm. Since this algorithm is quadratically convergent when the diagonal is ordered [2], one sweep should be enough to obtain $\epsilon$-small offdiagonal elements.

The complexity of this procedure is as follows. If we use only Givens transformations, we can keep all matrices upper triangular by a subsequent Givens correction. Such a pair takes $6 n$ flops per matrix and we need to propagate $n^{2} / 2$ of those. That means $3 n^{3}$ per matrix. The cost of one Kogbetliantz sweep is exactly the same since we propagate the same amount of Givens rotations. We therefore arrive at the following total count for our algorithm:

4 to $6 \mathrm{Kn}^{3}$ for triangularizing $A_{i} \rightarrow T_{i}$,
4 to $5 n^{3}$ for constructing $Q_{K}$ and $Q_{0}$,
$8 n^{3}$ for computing $U$ and $V$,
$3 K n^{3}$ for updating $T_{i} \rightarrow \hat{T}_{i}$,
$3 K n^{3}$ for one last Kogbetliantz sweep.
The total amount of flops after the bidiagonalization is thus comparable to applying 2 Kogbetliantz sweeps, whereas the Jacobi-like methods typically require 5 to 10 sweeps!

Moreover, our method allows us to select a few singular values and only compute the corresponding singular vectors. The matrices $Q_{K}$ and $Q_{0}$ can be stored in factored form and inverse iteration performed on $B$ to find its selected singular vector pairs and then transformed back to pairs of $A$ using $Q_{K}$ and $Q_{0}$.
5. Numerical examples. Computing the SVD of a general product/quotient of matrices is, as suggested above, a delicate numerical problem. In this section we analyze the accuracy of the $Q R$-like algorithm described in this paper using several examples of varying degree of difficulty. The examples are chosen in order to illustrate the following points already discussed in the paper.
(a) Implicit methods are more reliable than explicit methods. This is of course well known, but we illustrate it with some striking examples.
(b) The bidiagonal computed by the $Q R$-like method yields singular values computed to high relative accuracy even when their dynamical range is very large.
(c) The bidiagonal has a typical "graded" structure when the singular values have a wide dynamical range and its "off-bidiagonal" elements are negligible with respect to the bidiagonal elements in the same row. This is due to its connection to the Lanczos procedure as discussed earlier.
(d) The connection with the Lanczos procedure also allows us to terminate the bidiagonalization early and yet has a good estimate of the dominant singular values.
Points (a)-(d) illustrate the good (relative) accuracy that can be obtained from this procedure even without using iterative refinement based on Kogbetliantz's algorithm. The following points now compare the $Q R$-like and Kogbetliantz approaches.
(e) The bidiagonalization and Kogbetliantz methods have comparable accuracy in "difficult" examples with strong cancellation in the product.
(f) The typical number of Kogbetliantz steps ( 6 to 10) needed for convergence yields a much slower method than mere bidiagonalization. Moreover, the results are comparable, even when the Kogbetliantz iteration is continued further.
(g) The accuracy obtained from the bidiagonal only is already better on average than that of Kogbetliantz.
Finally, we illustrate that good (relative) accuracy is obtained also for a matrix quotient.
(h) The accuracy obtained by bidiagonalization of a matrix quotient is high, even when compared to the accuracy obtained if the inverted factors are explicitly known.
These points illustrate the power of this $Q R$-like method. Note that in all examples we use only the basic part of the algorithm without the iterative refinement step. All calculations were carried out in MATLAB on a Silicon Graphics Indigo workstation with IEEE floating point standard. For computing the singular values of the computed bidiagonal we use the method due to Fernando and Parlett [5].
(a) Implicit versus explicit. Let us consider the following products:

$$
A_{1}[n, m]=T_{n}^{m}
$$

where $T_{n}$ is a $n \times n$ symmetric Toeplitz matrix whose first column is $[2,-1,0,0, \ldots, 0]$ (singular values and singular vectors of such matrices are known [8]). Since it contains only integers we can form these powers of $T_{n}$ without any rounding errors, which is important for our comparison. The accuracy obtained by computing the SVD of this explicitly formed product is displayed in Figure 1. The interpretation of the figure is


FIG. 1. The relative accuracy obtained by computing the SVD for the explicitly formed product of Toeplitz matrices. In (a) and (b) solid, dashed, and dotted lines correspond to $A_{1}[n, m]$ for $n=10$ and $m \in\{8,16,32\}$; in (c) and (d), $n \in\{10,20,40\}$ and $m=8$. The lines interpolate the relative accuracies for the different singular values.
as follows. Figures 1 (a) and $1(\mathrm{~b})$ correspond to $n=10$ and $m \in\{8,16,32\}$, whereas Figures 1(c) and 1(d) correspond to $n \in\{10,20,40\}$ and $m=8$. The associated results are indicated by the solid, dashed, and dotted lines, respectively. Notice that each line corresponds to a different range of singular values as expected for matrices $T_{n}^{m}$ with different values of $m$ and $n$. In Figures 1(a) and 1(c) we plot the relative accuracy of the singular values as a function of the actual magnitude of the singular value. The lines interpolate the observed relative accuracies, that is, the values $\left|\sigma_{i}-\hat{\sigma}_{i}\right| / \sigma_{i}$. In Figures 1(b) and 1(d) we plot the maximum absolute error in the left singular vector elements, that is, $\max _{j}\left[\left|u_{j i}-\hat{u}_{j i}\right|\right]$, with $u_{j i}$ being the elements of the $i$ th singular vector, also as a function of the magnitude of the corresponding singular value. From the figure it is clear that the relative accuracy of the computed decomposition is quickly lost as we form powers of the matrix. Moreover, the situation is aggravated as we increase the dimension, and hence the condition number, of the matrix. The explanation lies of course in the fact that roundoff errors in a matrix $A$ are typically proportional to $\epsilon\|A\|$. For the product $A_{1}[n, m]$ this tends to have a catastrophic effect on the accuracy of the smallest singular values since they are smaller than $\epsilon\left\|A_{1}\right\|$.

Let us now use the $Q R$-like SVD algorithm. The result is shown in Figure 2 and illustrates that we have obtained a relative accuracy which is essentially independent of the magnitude of the associated singular value. This shows the need for implicit methods.
(b) Relative accuracy of implicit methods. A nonsymmetric example along the same vein is given in Figure 3. The interpretation of the figure is as for Figure 2, but now we consider the product


Fig. 2. Relative accuracy of the SVD estimate obtained by the $Q R$-like algorithm for the product of Toeplitz matrices. In (a) and (b) solid, dashed, and dotted lines correspond to $A_{1}[n, m]$ for $n=10$ and $m \in\{8,16,32\}$; in (c) and (d), $n \in\{10,20,40\}$ and $m=8$. Note that also the smallest singular values are computed with high relative accuracy.


Fig. 3. Relative accuracy of the SVD estimate obtained by the $Q R$-like algorithm for the product of nonsymmetric matrices. In (a) and (b) solid, dashed, and dotted lines correspond to $A_{2}[n, m]$ for $n=10$ and $m \in\{8,16,32\}$; in (c) and (d), $n \in\{10,20,40\}$ and $m=8$.


Fig. 4. Grading in the bidiagonal decomposition computed by the $Q R$-like algorithm. Figure (a) corresponds to $A_{1}[20,16]$ and (b) to $A_{2}[20,16]$. The *'s show the magnitude of the diagonal coefficients, the o's the magnitude of the upper bidiagonal coefficients.

$$
A_{2}[n, m]=\left(D_{n} D_{n}^{*}\right)^{m}
$$

Thus, there are $2 m$ matrices in the matrix product. The matrix $D_{n}$ is obtained by explicitly forming

$$
\begin{equation*}
D_{n} \equiv U_{n} \Sigma_{n} V_{n}^{*} \tag{13}
\end{equation*}
$$

where the matrices $U_{n}$ and $V_{n}$ are randomly chosen orthogonal matrices. They are defined by the singular vectors of a matrix with independent mean zero unit variance Gaussian entries. Furthermore, $\Sigma_{n}$ is the leading part of a $10 k \times 10 k$ diagonal matrix with diagonal equal to the Kronecker product:

$$
[10,9.9,9,8,7,6,5,4,3,2] \otimes\left[10^{-1}, 10^{-2}, \ldots, 10^{-k}\right]
$$

The motivation for choosing the product in this way is that we obtain an example in which the matrices involved are nonsymmetric and for which we "know" the actual singular values and can examine the obtained relative accuracy. The result is much as above. Using the implicit procedure for computing the singular values returns singular values whose relative accuracies are rather insensitive to the actual magnitude of the corresponding singular value.
(c) Graded bidiagonal. That the merits of the algorithm can be understood in terms of the Lanczos connection is confirmed by the next example explained in Figure 4. Here we have plotted the magnitude of the coefficients of the computed bidiagonal for the products $A_{1}[20,16]$ and $A_{2}[20,16]$, respectively, in Figures 4(a) and 4(b). The *'s show the absolute values of the diagonal coefficients and the $o$ 's the absolute values of the upper bidiagonal coefficients in the computed bidiagonal. We see that in both cases a grading has indeed been obtained. The algorithm picks out the dominant directions first, leading to a grading in the computed decomposition. The "effective condition number" of remaining subproblems are therefore successively reduced, and the associated singular values can apparently be obtained with high relative accuracy.

The high accuracy obtained above suggests that the "off-bidiagonal" elements in the transformed product are indeed small relative to the bidiagonal. This is confirmed by the next figure. In Figure 5 we plot, indicated by $*$, the norm of the off-bidiagonal elements normalized by the norm of the bidiagonal elements. That is, after the transformation to upper triangular form we explicitly form the product of the matrices in the product and compute for each row $j,\left\|o_{j,(j+2): n}\right\| /\left\|o_{j, j:(j+1)}\right\|$, with $o_{i, j}$ being the


Fig. 5. Relative accuracies of the bidiagonal elements $(*)$ and of the computed singular values (o). Figure (a) corresponds to $A_{1}[20,16]$ and (b) to $A_{2}[20,16]$.


FIG. 6. Accuracy of estimate of dominant singular value obtained from leading part of computed bidiagonal. The accuracy is plotted as a function of the dimension of the leading submatrix. Figure (a) corresponds to $A_{1}[80,16]$ and (b) to $A_{2}[80,16]$.
elements in the computed product. In exact arithmetic this quantity should be zero. The $o$ 's in the figure are the relative accuracies in the computed singular values. Note that the grading and relative smallness of the off-bidiagonal elements make it possible to compute even the smallest singular values with high relative accuracy.
(d) Dominant singular value. A consequence of the Lanczos connection is furthermore that we can obtain good estimates for the dominant singular values of the product without computing the full bidiagonal. This is illustrated in Figure 6. Here we plot the estimate of the dominant singular value we obtain by computing the corresponding singular value for the leading parts of the computed bidiagonal, $\hat{B}(1: i, 1: i)$. We plot the relative accuracy of this estimate as a function of $i$ in Figures 6(a) and $6(\mathrm{~b})$, corresponding to the products $A_{1}[80,16]$ and $A_{2}[80,16]$, respectively. The plots show that we need compute only a part of the bidiagonal in order to obtain a good estimate of the dominant singular value.
(e) Examples with strong cancellation. Here we consider examples with a significant cancellation in the product. That is, a subsequence of the matrices in the product is associated with a large dynamic range relative to that of the product whose associated singular values might be only mildly, or not at all, graded. The following example illustrates this:

$$
\begin{aligned}
& A_{3}[10, m]=D_{10}^{m} D_{10}^{-m} \\
& A_{4}[10, m]=\left(\begin{array}{ll}
D_{10} & D_{10}^{-1}
\end{array}\right)^{m}
\end{aligned}
$$



FIG. 7. The figure compares the accuracy of the computed $S V D$ using, respectively, the $Q R$-like (*) and the Kogbetliantz (o) algorithms. The considered matrix products exhibit strong cancellation. Figures (a) and (b) correspond to $A_{3}[n, m]$ and (c) and (d) to $A_{4}[n, m]$. In (a) and (c), $n=10$ and $m \in\{2,4,8,16,32,64\}$; in (b) and (d), $n \in\{10,14,18,22,26,30\}$ and $m=8$. The $Q R$ like algorithm provides accurate singular value estimates at a lower computational cost than the Kogbetliantz algorithm.
with $D_{10}$ defined as in (13). Note that in this example we compute $D_{10}^{-1}$ explicitly and use the product form of the algorithm. In Figure 7 subplots (a) and (b) correspond to $A_{3}$ and (c) and (d) to $A_{4}$. For the various product sizes we plot the maximum relative error over the computed singular values. We do so as a function of the "product condition number," defined as the product of the condition numbers of the matrices involved, in this case $\kappa_{D}^{2 m}$. In Figures 7(a) and 7(c) we let $n=10$ and $m \in$ $\{2,4,8,16,32,64\}$, whereas in Figures $7(\mathrm{~b})$ and $7(\mathrm{~d})$ we let $n \in\{10,14,18,22,26,30\}$ and $m=8$. Note that for both of the above matrix products the product condition number is much larger than its actual condition number. Figures $7(\mathrm{a})$ and $7(\mathrm{~b})$ show that for the matrix products which are associated with a significant cancellation there is a loss in relative accuracy. The o's in the plot correspond to computing the decomposition by the Kogbetliantz algorithm, fixing the number of sweeps to 12 to avoid issues of convergence tests. Note that the accuracy obtained thereby is not much better than that obtained by the bidiagonalization part of the $Q R$-like algorithm, that is, without iterative refinement.
(f) Convergence and complexity. We next turn to the special but important case when $m=2$ and compare the performance of the algorithm with that of the Kogbetliantz algorithm. In Figures 8(a) and 8(b) we consider the product $A_{2}[40,1]$. The dashed lines correspond to the relative accuracy obtained by $2,4,6$, and 10 sweeps of the Kogbetliantz algorithm. The relatively slow convergence of some singular values corresponds to those being closely spaced. Note that even 10 sweeps of Kogbetliantz's algorithm do not return an approximation with accuracy beyond that obtained by the


Fig. 8. Convergence of Kogbetliantz algorithm when computing the SVD of the pair of matrices defined by $A_{2}[40,1]$. The accuracies after $2,4,6$, and 10 sweeps are shown. The bottom solid line shows the accuracy obtained with the $Q R$-like algorithm without iterative refinement.



Fig. 9. Comparison of accuracy of the computed $S V D$ obtained, respectively, by the $Q R$-like and the Kogbetliantz algorithms for the square of a collection of random matrices. For each matrix realization the cross is the maximum relative error with the Kogbetliantz algorithm over the maximum relative error for the $Q R$-like algorithm.
$Q R$-like SVD algorithm without iterative refinement as shown by the solid line.
(g) Comparison of accuracy. The two plots in Figure 9 are obtained as follows. We consider the products defined by

$$
A_{5}[n, 2]=N_{n} N_{n}^{*}
$$

with $N_{n}$ being an $n \times n$ random matrix whose coefficients are normally distributed and with $n \in\{20,40,60,80\}$. The SVD was first computed (via MATLAB) and we then reconstructed $N_{n}$ from this SVD. Hence it is reasonable to assume that the singular values of $N_{n}$ and $A_{5}[n, 2]$ are known "exactly." Let $\hat{\sigma}_{i}$ and $\tilde{\sigma}_{i}$ represent the estimates of the singular values associated with, respectively, the $Q R$-like and the Kogbetliantz algorithms. In the latter case we used 10 sweeps, whereas in the former we did not include iterative refinement. Similarly let $\hat{u}_{i j}$ and $\tilde{u}_{i j}$ represent the coefficients in the left singular vectors. We then plot the ratio $\max _{i}\left[\left|\sigma_{i}-\tilde{\sigma}_{i}\right| / \sigma_{i}\right] / \max _{i}\left[\mid \sigma_{i}-\right.$ $\left.\hat{\sigma}_{i} \mid / \sigma_{i}\right]$ in Figure 9(a) and the ratio $\max _{i j}\left[\left|u_{i j}-\tilde{u}_{i j}\right|\right] / \max _{i j}\left[\left|u_{i j}-\hat{u}_{i j}\right|\right]$ in Figure $9(\mathrm{~b})$. The +'s correspond to different realizations of the matrix $N_{n}$. We see that the $Q R$-like algorithm typically yields a more accurate approximation despite its lower computational cost.
(h) Example with matrix quotients. In this last example we compute the SVD of matrix quotients involving inverted matrices. As described above, we then have


Fig. 10. The figure compares the accuracy of the computed $S V D$ when the $Q R$-like algorithm based on, respectively, the quotient representation $(+)$ and the product representation (*) of the matrix is being used. Figures (a) and (b) correspond to $A_{6}[n, m]$ and (c) and (d) to $A_{7}[n, m]$. In (a) and (c), $n=10$ and $m \in\{2,4,8,16,32\}$ and in (b) and (d), $n \in\{10,14,18,22,26\}$ and $m=8$. Note that the accuracy obtained from the quotient representation is similar to the accuracy obtained from the product representation.
to carry out an initial step in the bidiagonalization where the $Q R$ factorizations of the inverted matrices are computed. We construct the matrix quotients such that we explicitly can compute the inverses and compare the accuracy of the product version of the algorithm, based on these explicitly computed inverses, with the quotient version.

First, define the matrix quotient $A_{6}$ :

$$
A_{6}[n, m]=A_{1}^{-1} A_{2}^{-1} \cdots A_{m}^{-1} A_{m+1} A_{m+2} \cdots A_{2 m}
$$

where

$$
\begin{aligned}
A_{i} & =Q_{n}^{(i)} \Sigma_{n}^{-1} Q_{n}^{(i-1) *} \\
A_{m+i} & =Q_{n}^{(m+i-1)} \Sigma_{n} Q_{n}^{(m+i) *}
\end{aligned}
$$

for $1 \leq i \leq m$ and with $\Sigma_{n}$ defined as in (13). Moreover, the $Q_{n}^{(i)}$ are independent random orthogonal $n \times n$ matrices. As above these are defined by the singular vectors of a matrix with independent mean zero unit variance Gaussian entries.

Second, define the matrix quotient $A_{7}$ :

$$
A_{7}[n, m]=A_{1}^{-1} A_{2} A_{3}^{-1} \cdots A_{2 m-2} A_{2 m-1}^{-1} A_{2 m}
$$

with

$$
\begin{aligned}
A_{2 i-1} & =Q_{n}^{(2 i-1)} \Sigma_{n}^{-1} Q_{n}^{(2 i-2) *} \\
A_{2 i} & =Q_{n}^{(2 i-1)} \Sigma_{n} Q_{n}^{(2 i) *}
\end{aligned}
$$



FIG. 11. The figure compares the accuracy of the computed $S V D$ when the $Q R$-like algorithm based on, respectively, the quotient representation ( + ) and the product representation (*) of the matrix is being used. The considered matrix quotients exhibit strong cancellation. Figures (a) and (b) correspond to $A_{3}[n, m]$ and (c) and (d) to $A_{4}[n, m]$. In (a) and (c), $n=10$ and $m \in$ $\{2,4,8,16,32,64\}$ and in (b) and (d), $n \in\{10,14,18,22,26,30\}$ and $m=8$.
for $1 \leq i \leq m$ and with $\Sigma_{n}$ and $Q_{n}^{(i)}$ defined as above. Note that these quotients are associated with a large dynamic range.

The resulting relative accuracy obtained when varying the matrix dimension and the number of matrices in the quotient is shown in Figure 10. Figures 10(a) and $10(\mathrm{~b})$ correspond to the quotient $A_{6}$ and Figures $10(\mathrm{c})$ and $10(\mathrm{~d})$ correspond to the quotient $A_{7}$. In Figures 10 (a) and $10(\mathrm{c}), n=10$ and $m \in\{2,4,8,16,32\}$, whereas in Figures $10(\mathrm{~b})$ and $10(\mathrm{~d}), n \in\{10,14,18,22,26\}$ and $m=8$. The + 's show the accuracy obtained with the $Q R$-like algorithm based on the quotient and without iterative refinement. The $*$ 's show the accuracy obtained with the product form of the $Q R$-like algorithm without iterative refinement; note that in this case the inverses are explicitly computed. We see that the relative accuracy obtained when we do not assume knowledge of the inverses is comparable to, or even somewhat better than, the accuracy obtained if these are known!

Finally, reconsider the matrices of example (e) that exhibit strong cancellation. We compute as above the decomposition based on both the product form and the quotient form. The result is shown in Figure 11. Figures 11(a) and 11(b) correspond to the quotient $A_{3}$ and Figures $11(\mathrm{c})$ and 11(d) correspond to the quotient $A_{4}$. In Figures 11(a) and 11(c), $n=10$ and $m \in\{2,4,8,16,32,64\}$, whereas in Figures 11(b) and $11(\mathrm{~d}), n \in\{10,14,18,22,26,30\}$ and $m=8$. The figure shows that computation based on the quotient form gives a relative accuracy that is in general somewhat better than the accuracy based on the product form, at the cost of a slightly higher flop count.
6. Concluding remarks. The algorithm presented in this paper nicely complements the unitary decompositions for sequences of matrices defined for the generalized $Q R[3]$ and Schur decompositions [1]. These decompositions find applications in sequences of matrices defined from discretizations of ordinary differential equations occurring, for instance, in 2-point boundary value problems [9] or control problems [1]. We expect that they will lead to powerful tools for analyzing as well as solving problems in these application areas.

We want to stress here that in all examples it turned out to be sufficient to compute the bidiagonal $B$ of the expression $A^{s_{K}} \cdots A^{s_{1}}$ and then the singular values of $B$, without any further iterative refinement. This is rather surprising. The bounds obtained on the accuracy of the bidiagonal are much worse than what was observed in the examples. This point and the connection to the Lanczos process need further analysis. That we get accurate approximations for the leading order bidiagonals might be useful when solving ill-posed or inverse problems.

The main advantage of the new method lies exactly in the fact that this bidiagonal is so accurate. If no iterative refinement is needed, then the method requires 5 to 10 times less flops than Kogbetliantz! If iterative refinement is needed, then the method should still be superior since the work then amounts essentially to the work of two Kogbetliantz steps.

Finally, we point out that there is recent work on computing singular values to high relative accuracy via the Kogbetliantz algorithm. This work is based on extracting particular scalings from the factors. So far this has been applied to problems involving three factors only. Extensions to several matrices and whether these methods perhaps could be combined with the bidiagonal approach in an advantageous way are still open problems. Those methods and the ideas developed in this paper are, we believe, related. In both methods grading in the factors, obtained either explicitly or implicitly, is important.

## REFERENCES

[1] A. Bojanczyk, G. Golub, and P. Van Dooren, The periodic Schur form. Algorithms and applications, in Proceedings of the SPIE Conference, San Diego, CA, 1992, pp. 31-42.
[2] J. P. Charlier and P. Van Dooren, On Kogbetliantz's SVD algorithm in the presence of clusters, Linear Algebra Appl., 95 (1987), pp. 135-160.
[3] B. De Moor and P. Van Dooren, Generalizations of the singular value and $Q R$ decomposition, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 993-1014.
[4] J. Demmel, M. Gu, S. Eisenstat, I. Slapnicar, and K. Veselic, Computing the singular value decomposition with high relative accuracy, Linear Algebra Appl., submitted.
[5] K. V. Fernando and B. N. Parlett, Accurate singular values and differential qd algorithms, Numer. Math., 67 (1994), pp. 191-229.
[6] G. GOlub and W. Kahan, Calculating the singular values and pseudo-inverse of a matrix, SIAM J. Numer. Anal., 2 (1965), pp. 205-224.
[7] G. Golub and C. Van Loan, Matrix Computations, 3rd ed., The Johns Hopkins University Press, Baltimore, MD, 1996.
[8] R. Gregory and D. Karney, A Collection of Matrices for Testing Computational Algorithms, Wiley-Interscience, New York, 1969.
[9] R. M. Mattheij and S. J. Wright, Parallel stabilized compactification for ODEs with parameters and multipoint conditions, Appl. Numer. Math., 13 (1993), pp. 305-333.
[10] C. B. Moler and G. W. Stewart, An algorithm for the generalized matrix eigenvalue problem, SIAM J. Numer. Anal., 10 (1973), pp. 241-256.
[11] J. H. Wilkinson, The Algebraic Eigenvalue Problem, Clarendon Press, Oxford, 1965.


[^0]:    *Received by the editors August 20, 1997; accepted for publication (in revised form) by L. Eldén November 15, 1999; published electronically May 31, 2000. This paper contains research results of the Belgian Programme on Interuniversity Poles of Attraction, initiated by the Belgian State, Prime Minister's Office for Science, Technology and Culture. The scientific responsibility rests with its authors.
    http://www.siam.org/journals/simax/22-1/32557.html
    ${ }^{\dagger}$ Computer Science Department, Stanford University, Stanford, CA 94305-9025 (golub@ sccm.stanford.edu). This author was partially supported by the National Science Foundation under grants DMS-9105192 and DMS-9403899.
    ${ }^{\ddagger}$ SC-CM, Stanford University, Stanford, CA 94305-9025 (solna@sccm.stanford.edu). This author was partially supported by The Research Council of Norway.
    ${ }^{\text {§ }}$ Cesame, Université Catholique de Louvain, Louvain-la-Neuve B1348, Belgium (vdooren@ anma.ucl.ac.be). This author was partially supported by the National Science Foundation under grant CCR-96-19596.

