EXPONENTIAL SPLITTINGS OF PRODUCTS OF MATRICES AND ACCURATELY COMPUTING SINGULAR VALUES OF LONG PRODUCTS

SUELY OLIVEIRA* AND DAVID E. STEWART†

Abstract. Accurately computing the singular values of long products of matrices is important for estimating Lyapunov exponents: $\lambda_i = \lim_{n\to\infty}(1/n)\log\sigma_i(A_n\cdots A_1)$. Algorithms for computing singular values of products in fact compute the singular values of a perturbed product $(A_n+E_n)\cdots(A_1+E_1)$. The question is how small are the relative errors of the singular values of the product with respect to these factorwise perturbations. In general, the relative errors in the singular values can be quite large. However, if the product has an exponential splitting, then the error in the singular values is $O(n^2 \max_i \kappa_2(A_i)\|E_i\|_F)$, uniformly in n. The *exponential splitting* property is not directly comparable with the notion of hyperbolicity in dynamical systems, but is similar in philosophy.

Key words. stability, SVD, products of matrices, Lyapunov exponents

AMS subject classifications. 65F15; secondary 34D08

1. Introduction.

The computation of singular values and eigenvalues of long products of matrices is of importance in areas such as dynamical systems theory, and control of periodic systems. To be useful, these computations must produce estimates of the singular values and eigenvalues that have high *relative* accuracy, not high accuracy relative to $\|A_nA_{n-1}\cdots A_2A_1\|$. That is, we wish to compute each singular value or eigenvalue with a small error, relative to the size of that singular value or eigenvalue, not just relative to the norm of the matrix. The eigenvalues and singular values of a long product $A_nA_{n-1}\cdots A_1$ tend to grow or decay exponentially in n. Unless the exponential growth rates in n for the various singular values are all the same (a rare occurrence), the singular values will diverge exponentially rapidly. Thus the product $A_nA_{n-1}\cdots A_1$ rapidly becomes ill-conditioned. Explicitly computing the product will usually only give accurate values for the largest eigenvalue or singular value. The exponentially smaller eigenvalues and singular values are typically swamped by the numerical noise produced by the explicit computation of the product.

In dynamical systems studies, it is often useful to estimate the Lyapunov exponents λ_i of a dynamical system $x_{t+1} = f(x_t)$. These can be defined as

$$\lambda_i = \lim_{n\to\infty}\frac{1}{n}\log\sigma_i(\nabla f(x_{n-1})\cdots\nabla f(x_1)\,\nabla f(x_0)).$$

While conditions have been developed under which the Lyapunov exponents are known to be stable with respect to perturbations in $\nabla f(x_i)$ [6, 15], here we investigate a related question: When can the singular values or eigenvalues of a long product be computed to high relative accuracy? The condition we present (exponential splittings) are similar to the conditions in [6, 15] for stability of the Lyapunov exponents. Furthermore, the condition of having an exponential splitting is numerically verifiable.

There are methods for computing eigenvalues and singular values which have good backward error properties, such as the periodic Schur algorithm of Bojanczyk, Golub and Van Dooren [3] for computing eigenvalues, the Jacobi-based algorithm of Bojanczyk, Ewerbring, Luk and Van Dooren [4], and the Rutishauser LRCH-based

*Computer Science Dept., University of Iowa, Iowa City, IA 52242 (oliveira@cs.uiowa.edu)
†Mathematics Dept., University of Iowa, Iowa City, IA 52242 (dstewart@math.uiowa.edu)

method of one of the authors [17]. Related algorithms include the graded QR algorithm of G.W. Stewart [18], and the QR algorithm of Abarbanel, Brown and Kennel [1]. Applied to the product $A_n A_{n-1} \cdots A_1$ of $k \times k$ matrices, the algorithms of [3, 4, 17] compute the exact decomposition of the *perturbed product*:

$$(A_n + E_n)(A_{n-1} + E_{n-1}) \cdots (A_1 + E_1)$$

where $\|E_i\|_F = O(f(k)\mathbf{u}\|A_i\|_F)$ where $f(k)$ is a modest function of k, and $\mathbf{u}$ is the unit roundoff or machine epsilon.

These algorithms can successfully handle products with hundreds to thousands of factor matrices, and still appear to produce very accurate results. This is in spite of the fact that such long products are usually extremely ill-conditioned with condition numbers typically 10^{100} to 10^{1000} and beyond. A partial explanation for the ability of the above algorithms to produce accurate estimates can be found in multiplicative perturbation theorems. The following two results can be found in D. Stewart [17, 16]. These results concern *outer multiplicative perturbations*, which are perturbations of the form $(I + G_+)A_n A_{n-1} \cdots A_2 A_1 (I + G_-)$ For singular values:

THEOREM 1.1. *If $A \in \mathbf{R}^{m\times k}$ and $E \in \mathbf{R}^{m\times m}$ with $\|E\|_2 < 1$, then $|\sigma_i((I + E)A)/\sigma_i(A) - 1| \le \|E\|_2$ for $i = 1, \ldots,$ rank(A).*

This is an easy corollary to the result mentioned in Horn and Johnson [13, §7.3, Problem 18, p. 423] that $\sigma_{i+j-1}(AB) \le \sigma_i(A)\sigma_j(B)$.

For eigenvalues:

THEOREM 1.2. *If A is diagonalizable with $X^{-1}AX = \Lambda =$ diag$(\lambda_1, \ldots, \lambda_k)$ and*

$$\kappa_1(X)\|E\|_1 < \min_{i,j: i\ne j} \frac{|\lambda_i - \lambda_j|}{|\lambda_i| + |\lambda_j|}, \tag{1.1}$$

then the eigenvalues $\widetilde{\lambda}_i$ of $(I + E)A$ are distinct and satisfy

$$\frac{|\widetilde{\lambda}_i - \lambda_i|}{|\lambda_i|} \le \kappa_1(X)\|E\|_1. \tag{1.2}$$

Note that the conditioning of Λ does not enter into the bounds; instead only the separation of the eigenvalues and the conditioning of the diagonalizing transformation enter into them.

A more general result on the perturbation of eigenvalues is proven in the appendix:

THEOREM 1.3. *Suppose that A is diagonalizable with $X^{-1}AX = \Lambda =$ diag$(\lambda_1, \ldots, \lambda_n)$. If $\beta = \kappa_1(X)\|E\|_1 < 1$, then for a suitable ordering $\widetilde{\lambda}_1, \widetilde{\lambda}_2, \ldots, \widetilde{\lambda}_n$ of the eigenvalues of $(I + E)A$, we have*

$$\frac{|\widetilde{\lambda}_i - \lambda_i|}{|\lambda_i|} \le (1 + \beta)\left(\frac{1+\beta}{1-\beta}\right)^{n-1} - 1. \tag{1.3}$$

Related results on multiplicative perturbations are given in Ipsen's survey article [14, §§5–7] and joint work of Eisenstat and Ipsen [8, 9, 10].

Provided each of the factor matrices are well-conditioned, the factorwise perturbations $A_i + E_i$ can be represented as multiplicative perturbations $(I + F_i)A_i$ with $F_i = E_i A_i^{-1}$ and $\|F_i\|_F = \|E_i A_i^{-1}\|_F = O(f(k)\,\mathbf{u}\kappa_2(A_i))$. However, these do not necessarily translate into small *relative* perturbations of the singular values or eigenvalues, even if all of the factor matrices are well-conditioned. The reason is that

Theorems 1.1, 1.2, and 1.3 consider only *outer* multiplicative perturbations, while the perturbations of the inner factors $A_2, \ldots, A_{n-1}$ cannot be directly represented in terms of outer multiplicative perturbations. By contrast, if $n = 2$, then there are no inner perturbations, and relative errors can be bounded by quantities of order $\|E_1 A_1^{-1}\|_F + \|A_2^{-1} E_2\|_F = O(f(k)\mathbf{u}(\kappa_2(A_1) + \kappa_2(A_2)))$. This is closely related to a result of of Drmač [7, Thm. 3.1]:

THEOREM 1.4. *If σ_i are the singular values of $B^T C$ and $\sigma_i + \delta\sigma_i$ the singular values of $(B + \delta B)^T (C + \delta C)$, then*

$$\frac{|\delta\sigma_i|}{\sigma_i} \le \|B^\dagger \delta B\|_2 + \|C^\dagger \delta C\|_2 + \|B^\dagger \delta B\|_2 \|C^\dagger \delta C\|_2.$$

The purpose of this paper is to give tight bounds on the size of equivalent outer perturbations, given bounds on the inner perturbations. Such bounds show that current algorithms for computing eigenvalues or singular values for long products of matrices have small relative error for matrix products with exponential splittings.

It is possible to get very small bounds on equivalent outer perturbations for small inner perturbations, even if the product is extremely ill-conditioned. For example, let $B = \operatorname{diag}(1, 1/2)$. Then for any $0 \le p \le n$,

$$B^p \begin{bmatrix} 1 & \epsilon_1 \\ \epsilon_2 & 1 \end{bmatrix} B^{n-p} = \begin{bmatrix} 1 & 0 \\ 2^{-p}\epsilon_2 & 1 \end{bmatrix} B^k \begin{bmatrix} 1 & 2^{-(n-p)}\epsilon_1 \\ 0 & 1 - \epsilon_1\epsilon_2 \end{bmatrix}.$$

However, it is not always possible to transform small inner perturbations to small outer perturbations. Consider the following counter-example in D. Stewart [17]: Let $A = \operatorname{diag}(1/2, 1)$ and $B = \operatorname{diag}(1, 1/2)$. Choose n and $\epsilon_1, \epsilon_2 \ll 1$ so that $0 < 2^{-n} \ll \epsilon_1, \epsilon_2 \ll 1$. Let $\eta = 2^{-n}$. Then

$$A^n \begin{bmatrix} 1 & \epsilon_1 \\ \epsilon_2 & 1 \end{bmatrix} B^n = \begin{bmatrix} \eta & \epsilon_1\eta^2 \\ \epsilon_2 & \eta \end{bmatrix} = \eta \begin{bmatrix} 1 & \epsilon_1\eta \\ \epsilon_2/\eta & 1 \end{bmatrix}$$

which has a singular value of size $\approx \epsilon_2 \gg \eta$ (η is the only singular value of the unperturbed product $A^n B^n$).

The first example shows that B^n is a stable product, regardless of how large n is, while $A^n B^n$ is not. In the following section we will see that this is because B^n has an exponential splitting, while $A^n B^n$ does not.

2. Exponential splittings. Exponential splittings are closely related to the concept of hyperbolicity in dynamical systems. Similar properties have been discussed for demonstrating the stability of Lyapunov exponents with respect to perturbations to linear differential equations (see, for example, the discussions of exponential dichotomies on pp. 403–404, and pp. 407–408 of [6]). Exponential splittings lead to results showing that the singular values can be computed to high relative accuracy provided each factor is well conditioned.

The main result of this section is a proof that inner perturbations of long products with exponential splittings can be replaced by outer perturbations whose magnitudes are a modest multiple of the size of the inner perturbations. At the end of this section, this result is used to prove that the singular values of long products with exponential splittings can in fact be computed to high relative accuracy, in spite of the extreme ill-conditioning that occurs with long products. Eigenvalues of diagonalizable long products with exponential splittings can also be computed to high relative accuracy due to the stability results for outer perturbations above (Theorems 1.2 and 1.3).

DEFINITION 2.1. *An exponential splitting of an infinite product*

$$\cdots A_2 A_1 A_0 A_{-1} A_{-2} \cdots$$

of $k \times k$ matrices consists of a collection of factorizations

$$A_i = X_{i+1}^{-1} D_i X_i$$

where, for suitable constants $C, M > 0$, and $0 < \alpha < 1$,

1. *D_i is diagonal for all i,*
2. *$\kappa_2(X_i) = \|X_i\| \, \|X_i^{-1}\| \leq M$ for all i, and*
3. *for any $1 \leq r < s \leq k$, and $n = 1, 2, 3, \ldots$,*

$$\prod_{k=i}^{i+n-1} \frac{(D_k)_{ss}}{(D_k)_{rr}} \leq C\, \alpha^n. \tag{2.1}$$

Note that the product B^n of the previous section has a trivial exponential splitting: $D_i = B$ and $X_i = I$ for all i, giving $C = 1$, $\alpha = 1/2$ and $M = 1$. On the other hand, the product $A^n B^n$ does not admit an exponential splitting. This can be identified from the instability results of the previous section.

In this section it will be shown that if the infinite product $\cdots A_2 A_1 A_0 A_{-1} A_{-2} \cdots$ has an exponential splitting, then inner perturbations of any finite sub-product can be replaced by outer perturbations of similar size.

THEOREM 2.2. *Suppose that $\cdots A_2 A_1 A_0 A_{-1} A_{-2} \cdots$ has an exponential splitting with constants C, M and α as described above. There is a number $\widetilde{C}$ depending only on C and α such that if $\|F_i\|_F \leq \gamma/M \leq 1/(8M\widetilde{C}n^2)$ for all i, then there are matrices $G_\pm$ such that $\|G_\pm\|_F \leq \widetilde{C}n^2\gamma$ and*

$$A_t(I+F_t)A_{t-1} \cdots A_{t-n+1}(I+F_{t-n+1})A_{t-n} = (I+G_+)A_t A_{t-1} \cdots A_{t-n+1} A_{t-n}(I+G_-). \blacksquare$$

Before we prove this theorem, we need some lemmas. The first step of the proof of Theorem 2.2 reduces the problem to dealing with perturbed products of the D_i's. Once this is done, it should be noted that the lower triangular perturbations should be moved left, and the upper triangular perturbations should be moved right, in the product; so we need to decompose inner multiplicative perturbations into products of lower and upper triangular parts.

LEMMA 2.3. *If $I + E = (I + E_L)(I + E_D)(I + E_U)$ where E_L is strictly lower triangular, E_D diagonal, E_U strictly upper triangular, and $\|E\|_2 < 1$, then*

$$\|E_L\|_F, \|E_D\|_F, \|E_U\|_F \leq \frac{\|E\|_F}{1 - \|E\|_2}.$$

Proof. The proof is based on the results of Barrlund [2] and Sun [19] as cited in Higham [12, Thm. 9.14]: If $A = LU$ is the LU factorization of A and $A + \Delta A = (L + \Delta L)(U + \Delta U)$ is the LU factorization of $A + \Delta A$, then if $G = L^{-1}\, \Delta A\, U^{-1}$ with $\|G\|_2 < 1$, then

$$\max\left\{ \frac{\|\Delta L\|_F}{\|L\|_2}, \frac{\|\Delta U\|_F}{\|U\|_2} \right\} \leq \frac{\|G\|_F}{1 - \|G\|_2}.$$

Applying this to $A = I$ (so that $L = U = I$) and $\Delta A = E$, we find that for the LU factorization $I + E = (I + E_L)(I + E_{DU})$, where E_{DU} is upper triangular, but not necessarily strictly upper triangular, $\|E_L\|_F, \|E_{DU}\|_F \le \|E\|_F/(1 - \|E\|_2)$ and the desired inequality holds for E_L.

Since E_D is just the diagonal part of E_{DU}, we have $\|E_D\|_F \le \|E_{DU}\|_F$, and the desired inequality holds for E_D.

By applying the result of Barrlund and Sun to $I + E^T = (I + E_U^T)(I + E_{DL}^T)$, the desired inequality also holds for E_U. □

The second lemma we need gives simple bounds for products of matrices.

LEMMA 2.4. *If* $I + C = (I + A)(I + B)$*, then* $1 + \|C\| \le (1 + \|A\|)(1 + \|B\|)$ *for any matrix norm* $\|\cdot\|$.

Proof. By direct calculation, $C = A + B + AB$; taking norms and the obvious inequalities gives the result. □

Lemma 2.4 gives multiplicative bounds; we will later need to turn this into additive bounds using Lemma 2.5.

LEMMA 2.5. *If* $(1 + a) \le \prod_{i=1}^n (1 + b_i/(1 - b_i))$*,* $b_i \ge 0$ *for all* i*, and* $\sum_i b_i < 1$*, then*

$$a \le \frac{\sum_{i=1}^n b_i}{1 - \sum_{i=1}^n b_i}.$$

Proof. Suppose that $\alpha, \beta \ge 0$ and that $\alpha + \beta < 1$. Then

$$\begin{aligned}\left(1 + \frac{\alpha}{1-\alpha}\right)\left(1 + \frac{\beta}{1-\beta}\right) &= \frac{1}{1-\alpha}\,\frac{1}{1-\beta} \\ &= \frac{1}{1 - \alpha - \beta + \alpha\beta} \\ &\le 1 + \frac{\alpha + \beta}{1 - (\alpha + \beta)}.\end{aligned}$$

Applying this inequality inductively yields the result. □

Proof. (Theorem 2.2) We consider a perturbed product of p matrices:

$$A_t(I + F_t)A_{t-1} \cdots A_{t-p+1}(I + F_{t-p+1})A_{t-p}.$$

We will represent this perturbed product by an outer perturbation

$$(I + G_{t,p,+})A_t A_{t-1} \cdots A_{t-p+1} A_{t-p}(I + G_{t,p,-})$$

and obtain bounds $\|G_{t,p,+}\|_F, \|G_{t,p,-}\|_F \le h(p,\gamma)/(1 - h(p,\gamma))$ provided $h(p,\gamma) \le \theta$, where $\gamma = \max_{i=t-p+1,\ldots,t} \kappa_2(A_i)\|F_i\|_F$, using the exponential splitting of Definition 2.1. We recursively define

$$h(n,\gamma) = \frac{1 + C\alpha^{\lfloor n/2 \rfloor}/(1-\theta)}{1-\theta}(3\,h(\lceil n/2 \rceil, \gamma) + \gamma) \tag{2.2}$$

for all $n \ge N$, provided the resulting value is no larger than θ, where θ and N are constants to be chosen later. The constants C and α are obtained from the exponential splitting property.

Note that $A_i = X_{i+1}^{-1} D_i X_i$ for all i. Without loss of generality, scale the X_i's so that $\|X_i\| = \|X_i^{-1}\| = \sqrt{\kappa_2(X_i)}$. This scaling does not change the singular values except by a scale factor, which does not change the relative perturbations. Thus we can assume that $\|X_i\|, \|X_i^{-1}\| \le \sqrt{M}$. For such a choice of X_i,

$$\begin{aligned}&A_t(I+F_t)A_{t-1}\cdots A_{t-p+1}(I+F_{t-p+1})A_{t-p} = \\ &\quad X_{t+1}^{-1}D_t(I+X_tF_tX_{t-1}^{-1})D_{t-1}\cdots D_{t-p+1}(I+X_{t-p+1}F_{t-p+1}X_{t-p}^{-1})D_{t-p}X_{t-p}.\end{aligned}$$

Now, write $G_i = X_i F_i X_{i-1}^{-1}$. By the boundedness of $\kappa_2(X_i)$,

$$\|G_i\|_F \le \|X_i\|\,\|F_i\|_F\,\|X_{i-1}^{-1}\| = \sqrt{\kappa_2(X_i)\kappa_2(X_{i-1})}\|F_i\|_F \le M\|F_i\|_F.$$

Thus we have the perturbed product

$$X_{t+1}^{-1}D_t(I+G_t)D_{t-1}\cdots D_{t-p+1}(I+G_{t-p+1})D_{t-p}X_{t-p}.$$

We now consider the problem of obtaining equivalent outer multiplicative perturbations for segments of this perturbed product

$$\begin{aligned}&D_i(I+G_i)D_{i-1}\cdots D_{i-q+1}(I+G_{i-q+1})D_{i-q} = \\ &\qquad (I+G_+)D_iD_{i-1}\cdots D_{i-q+1}D_{i-q}(I+G_-).\end{aligned}$$

If $q = 1$, there is nothing to do; there are no “inner perturbations” and we can take $h(1,\gamma) = \gamma$.

Consider the perturbed product of length $n = p + q$:

$$D_i(I+G_i)D_{i-1}\cdots D_{i-p-q+1}(I+G_{i-p-q+1})D_{i-p-q}.$$

By the complete induction hypothesis,

$$\begin{aligned}D_i(I+G_i)D_{i-1}\cdots D_{i-q+1}(I+G_{i-q+1})D_{i-q} &= \\ (I+\widehat{G}_+)D_iD_{i-1}\cdots D_{i-q+1}D_{i-q}(I+\widehat{G}_-)&,\\ D_{i-q-1}(I+G_{i-q-1})D_{i-q-2}\cdots D_{i-p-q+1}(I+G_{i-p-q+1})D_{i-p-q} &= \\ (I+\widetilde{G}_+)D_{i-q-1}D_{i-q-2}\cdots D_{i-p-q+1}D_{i-p-q}(I+\widetilde{G}_-)&\end{aligned}$$

where $\|\widehat{G}_\pm\|_F \le h(q,\gamma)/(1-h(q,\gamma))$ and $\|\widetilde{G}_\pm\|_F \le h(p,\gamma)/(1-h(p,\gamma))$. Then

$$\begin{aligned}&D_i(I+G_i)D_{i-1}\cdots D_{i-p-q+1}(I+G_{i-p-q+1})D_{i-p-q} = \\ &\qquad (I+\widehat{G}_+)D_iD_{i-1}\cdots D_{i-q+1}D_{i-q}(I+\widehat{G}_-)(I+G_{i-q})(I+\widetilde{G}_+) \\ &\qquad\quad D_{i-q-1}D_{i-q-2}\cdots D_{i-p-q+1}D_{i-p-q}(I+\widetilde{G}_-).\end{aligned}$$

Write $I+\overline{G} = (I+\widehat{G}_-)(I+G_{i-q})(I+\widetilde{G}_+)$, so that

$$\begin{aligned}1+\|\overline{G}\|_F &\le (1+\|\widehat{G}_-\|_F)(1+\|G_{i-q}\|_F)(1+\|\widetilde{G}_+\|_F) \\ &\le (1+\frac{h(q,\gamma)}{1-h(q,\gamma)})(1+\gamma)(1+\frac{h(p,\gamma)}{1-h(p,\gamma)}) \\ &\le 1+\frac{h(q,\gamma)+\gamma+h(p,\gamma)}{1-(h(q,\gamma)+\gamma+h(p,\gamma))}.\end{aligned}$$

Now we decompose $I+\overline{G} = (I+\overline{G}_L)(I+\overline{G}_D)(I+\overline{G}_U)$ where $\overline{G}_L$ is strictly lower triangular, $\overline{G}_D$ is diagonal, and $\overline{G}_U$ is strictly upper triangular. This is possible as long as $\|\overline{G}\|_2 < 1$. By Lemma 2.3,

$$\|\overline{G}_L\|_F, \|\overline{G}_D\|_F, \|\overline{G}_U\|_F \le \|\overline{G}\|_F/(1-\|\overline{G}\|_2) \le \|\overline{G}\|_F/(1-\|\overline{G}\|_F).$$

Provided $\|\overline{G}\|_F \leq \theta < 1$, $\|\overline{G}_L\|_F, \|\overline{G}_D\|_F, \|\overline{G}_U\|_F \leq \|\overline{G}\|_F/(1-\theta)$.

Thus

$$\|\overline{G}_L\|_F, \|\overline{G}_D\|_F, \|\overline{G}_U\|_F \leq \frac{(h(q,\gamma)+\gamma+h(p,\gamma)/(1-\theta)}{1-(h(q,\gamma)+\gamma+h(p,\gamma))/(1-\theta)}.$$

By the exponential splitting property,

$$D_i D_{i-1}\cdots D_{i-q+1}D_{i-q}(I+\overline{G}_L) = (I+\widehat{\overline{G}}_L)D_i D_{i-1}\cdots D_{i-q+1}D_{i-q}$$

where $(\widehat{\overline{G}}_L)_{rs} = (\overline{G}_L)_{rs}\prod_{k=i-q}^{i}((D_k)_{rr}/(D_k)_{ss})$, and so $|(\widehat{\overline{G}}_L)_{rs}| \leq C\alpha^q|(\overline{G}_L)_{rs}|$ for $r < s$. Since $\overline{G}_L$ is strictly lower triangular, $|\widehat{\overline{G}}_L| \leq C\alpha^q|\overline{G}_L|$ by (2.1), and $\|\widehat{\overline{G}_L}\|_F \leq C\alpha^q\|\overline{G}_L\|_F \leq C\alpha^q\|\overline{G}\|_F/(1-\|\overline{G}\|_F)$ by (2.1). Also, $(I+\overline{G}_D)D_iD_{i-1}\cdots D_{i-q+1}D_{i-q} = D_iD_{i-1}\cdots D_{i-q+1}D_{i-q}(I+\overline{G}_D)$ since diagonal matrices commute.

Similarly,

$$\begin{aligned}(I+\overline{G}_U)D_{i-q-1}D_{i-q-2}\cdots D_{i-p-q+1}D_{i-p-q} &=\\ &D_{i-q-1}D_{i-q-2}\cdots D_{i-p-q+1}D_{i-p-q}(I+\widehat{\overline{G}}_U)\end{aligned}$$

where $\|\widehat{\overline{G}}_U\|_F \leq C\alpha^p\|\overline{G}_U\|_F \leq C\alpha^p\|\overline{G}\|_F/(1-\|\overline{G}\|_F)$, and

$$\begin{aligned}(I+\overline{G}_D)D_{i-q-1}D_{i-q-2}\cdots D_{i-p-q+1}D_{i-p-q} &=\\ &D_{i-q-1}D_{i-q-2}\cdots D_{i-p-q+1}D_{i-p-q}(I+\overline{G}_D).\end{aligned}$$

Thus

$$\begin{aligned}D_i(I+G_i)D_{i-1}&\cdots D_{i-p-q+1}(I+G_{i-p-q+1})D_{i-p-q}\\ &= (I+\widehat{G}_+)(I+\widehat{\overline{G}}_L)(I+\overline{G}_D)D_iD_{i-1}\cdots D_{i-p-q+1}D_{i-p-q}\\ &\qquad (I+\overline{G}_D)(I+\widehat{\overline{G}}_U)(I+\widetilde{G}_-)\\ &= (I+G_+)D_iD_{i-1}\cdots D_{i-p-q+1}D_{i-p-q}(I+G_-).\end{aligned}$$

Thus

$$(I+G_+) = (I+\widehat{G}_+)(I+\widehat{\overline{G}}_L)(I+\overline{G}_D)^{1/2}$$

where $\|\widehat{G}_+\|_F \leq h(q,\gamma)/(1-h(q,\gamma))$, $\|\widehat{\overline{G}}_L\|_F \leq C\alpha^q\|\overline{G}\|_F/(1-\|\overline{G}\|_F) \leq (C\alpha^q/(1-\theta))\|\overline{G}\|_F$, and $\|\overline{G}_D\|_F \leq \|\overline{G}\|_F/(1-\|\overline{G}\|_F) \leq \|\overline{G}\|_F/(1-\theta)$.

So

$$\begin{aligned}1+\|G_+\|_F &\leq (1+\|\widehat{G}_+\|_F)(1+\|\widehat{\overline{G}}_L\|_F)(1+\|\overline{G}_D\|_F)\\ &\leq (1+\frac{h(q,\gamma)}{1-h(q,\gamma)})(1+\frac{C\alpha^q\|\overline{G}\|_F/(1-\theta)}{1-C\alpha^q\|\overline{G}\|_F/(1-\theta)})(1+\frac{\|\overline{G}\|_F}{1-\|\overline{G}\|_F})\\ &\leq 1+\frac{h(q,\gamma)+(1+C\alpha^q/(1-\theta))\|\overline{G}\|_F}{1-(h(q,\gamma)+(1+C\alpha^q/(1-\theta))\|\overline{G}\|_F)}\end{aligned}$$

Since

$$\|\overline{G}\|_F \leq \frac{h(q,\gamma)+\gamma+h(p,\gamma)}{1-(h(q,\gamma)+\gamma+h(p,\gamma))} \leq \theta$$

we have $h(q,\gamma)+\gamma+h(p,\gamma) \leq \theta$, and so

$$\|G_+\|_F \leq \frac{h(p+q,\gamma)}{1-h(p+q,\gamma)}$$

provided

$$h(p+q,\gamma) \geq h(q,\gamma) + \frac{1+C\alpha^q/(1-\theta)}{1-\theta}(h(q,\gamma)+\gamma+h(p,\gamma)). \tag{2.3}$$

Applying the same arguments to the right-hand outer perturbation gives

$$\|G_-\|_F \leq \frac{h(p+q,\gamma)}{1-h(p+q,\gamma)}$$

provided

$$h(p+q,\gamma) \geq h(p,\gamma) + \frac{1+C\alpha^p/(1-\theta)}{1-\theta}(h(q,\gamma)+\gamma+h(p,\gamma)). \tag{2.4}$$

If we choose $p = \lfloor n/2 \rfloor$ and $q = \lceil n/2 \rceil$, then the inequalities in (2.3, 2.4) become equalities by (2.2). Thus the induction step holds for n.

Also note that we can write $h(n,\gamma) = h_n\,\gamma$ where

$$h_n = \frac{1+C\alpha^{\lfloor n/2 \rfloor}/(1-\theta)}{1-\theta}(3\,h_{\lceil n/2 \rceil}+1), \qquad h_1 = 1,$$

and the desired inequalities hold for $k = 1, 2, \ldots, n$ provided $\gamma \leq \theta/h_n$.

The main task now is to determine the rate of growth of h_n. Fix $0 < \theta \leq 1/8$. Then we can find an N such that $(1+C\alpha^{\lfloor N/2 \rfloor}/(1-\theta))/(1-\theta) \leq \frac{8}{7}(1+\frac{8}{7}C\alpha^{\lfloor N/2 \rfloor}) < \frac{4}{3}$. For $n \geq N$,

$$h_n \leq 4\,h_{\lceil n/2 \rceil} + \frac{4}{3}.$$

By the Master Theorem [5, p. 62], there is a constant $\widetilde{C}$ so that $h_n \leq \widetilde{C}\,n^2$ for all $n \geq 1$.

The desired bound therefore holds for all n where

$$\gamma = \max_i \|G_i\|_F \leq \max_i \kappa_2(A_i)\|F_i\|_F \leq \theta/(\widetilde{C}n^2) = 1/(8\widetilde{C}\,n^2).$$

□

Note that the bounds are not sharp; in particular, the exponent of n in the bound $\|G_\pm\|_F < \widetilde{C}n^2\gamma$ can be reduced by further analysis to be arbitrarily close to $\log_2(3) \approx 1.58$ by making N large and $\theta > 0$ small.

The relative stability of the singular values of long products $A_pA_{p-1}\cdots A_1$ can be easily deduced from the above result:

Theorem 2.6. *Suppose the infinite product $\cdots A_3A_2A_1$ has an exponential splitting with constants C, α and M, and $\widetilde{C}$ is the constant described in Theorem 2.2. Also, if $\Phi_p = A_pA_{p-1}\cdots A_1$ and $\widehat{\Phi}_p = (A_p+E_p)(A_{p-1}+E_{p-1})\cdots(A_1+E_1)$ with $\kappa_2(A_i)\|E_i\|_F \leq \gamma/M$ for all i, where $\gamma \leq 1/(8\widetilde{C}\,p^2)$, then*

$$\left|\frac{\sigma_i(\widehat{\Phi}_p)}{\sigma_i(\Phi_p)} - 1\right| \leq \widetilde{C}p^2\gamma(2+\widetilde{C}p^2\gamma).$$

Proof. By Theorem 2.2, there are matrices $G_\pm$ where

$$\widehat{\Phi}_p = (I+G_+)\Phi_p(I+G_-)$$

and $\|G_\pm\|_F \le \widetilde{C}p^2\gamma$. Now

$$1-\|G_-\|_F \le \frac{\sigma_i((I+G_+)\Phi_p(I+G_-))}{\sigma_i((I+G_+)\Phi_p)} \le 1+\|G_-\|_F,$$
$$1-\|G_+\|_F \le \frac{\sigma_i((I+G_+)\Phi_p)}{\sigma_i(\Phi_p)} \le 1+\|G_+\|_F,$$

using the inequalities $\|G_\pm\|_2 \le \|G_\pm\|_F$.

Combining these two gives

$$(1-\|G_-\|_F)(1-\|G_+\|_F) \le \frac{\sigma_i(\widehat{\Phi}_p)}{\sigma_i(\Phi_p)} \le (1+\|G_-\|_F)(1+\|G_+\|_F).$$

Substituting the bounds for $\|G_\pm\|_F$ and subtracting one gives the desired result. □

3. Numerical results. In order to illustrate the theory, some numerical calculations were performed on a number of different products and their perturbations. The algorithm used to compute the singular values of the products is a previously unpublished one. It first performs a product QR factorization as described in [17]. Then $R = R_pR_{p-1}\cdots R_2R_1$ is computed in the form $D\widehat{R}$ where D is diagonal and $\widehat{R}$ is upper triangular and $\widehat{r}_{ii} = \pm 1$ for all i. Thus the entries of $\widehat{R}$ are well-scaled. The *logarithms* of the entries in D are stored A Jacobi method is used to compute the singular value decomposition $D\widehat{R} = UD\widehat{\Sigma}V^T$. Since D is usually extremely ill-conditioned, the formulas used to compute the Jacobi rotations need to be properly designed to avoid excessive round-off error and also to avoid over- and under-flow. The last factor R_1 is replaced by R_1V^T. Then the product QR factorization of the RV^T is computed and the process repeats. This continues until the strictly upper triangular part of $\widehat{R}$ has norm less then the tolerance, which was chosen to be $10^{-14}p$. This method can be shown to have small backward error provided the number of outer iterations is small.

The test results were obtained by perturbing products of the form

1. $A\cdot A\cdots A = A^n$,
2. $A\cdot C\cdot A\cdot C\cdots A\cdot C = (AC)^{n/2}$, and
3. $A\cdot A\cdots A\cdot B\cdot B\cdots B = A^{n/2}B^{n/2}$,

where $A = \mathrm{diag}(2,1)$, $B = \mathrm{diag}(1,2)$, and $C = \mathrm{diag}(1,1.8)$. Note that the third class of products are not stable for large n. The perturbations were pseudo-random matrices generated by MATLAB^{TM}'s `rand` command multiplied by a factor α ranging from 10^{-14} to 10^{-6}. This command produces matrices with entries in the range $[0,1)$ which are intended to be independent uniformly distributed random values. Note that the mean of the entries in the perturbing matrices is $\frac{1}{2}\alpha$; using mean zero perturbing matrices tends to produce a slower asymptotic growth in the perturbations of the singular values due to an averaging effect.

The values of α used are 10^{-14}, 10^{-12}, 10^{-10}, 10^{-8} and 10^{-6}. Each family of products was tested for n = 30, 100, 300, 3000, and 10000. For each product and each value of α, ten runs were done and means of the changes in the logarithms of the singular values were computed. For reference, the logarithms of the singular values of the unperturbed products were computed. Since the matrices A, B and C used are diagonal matrices, the only errors for the unperturbed products are incurred in computing and adding the logarithms of the diagonal entries of A, B, and C.

For the products A^n and $(AC)^{n/2}$, the ratios between the errors in $\log(\sigma_i)$ and the perturbation level were essentially constant, showing a linear relationship between

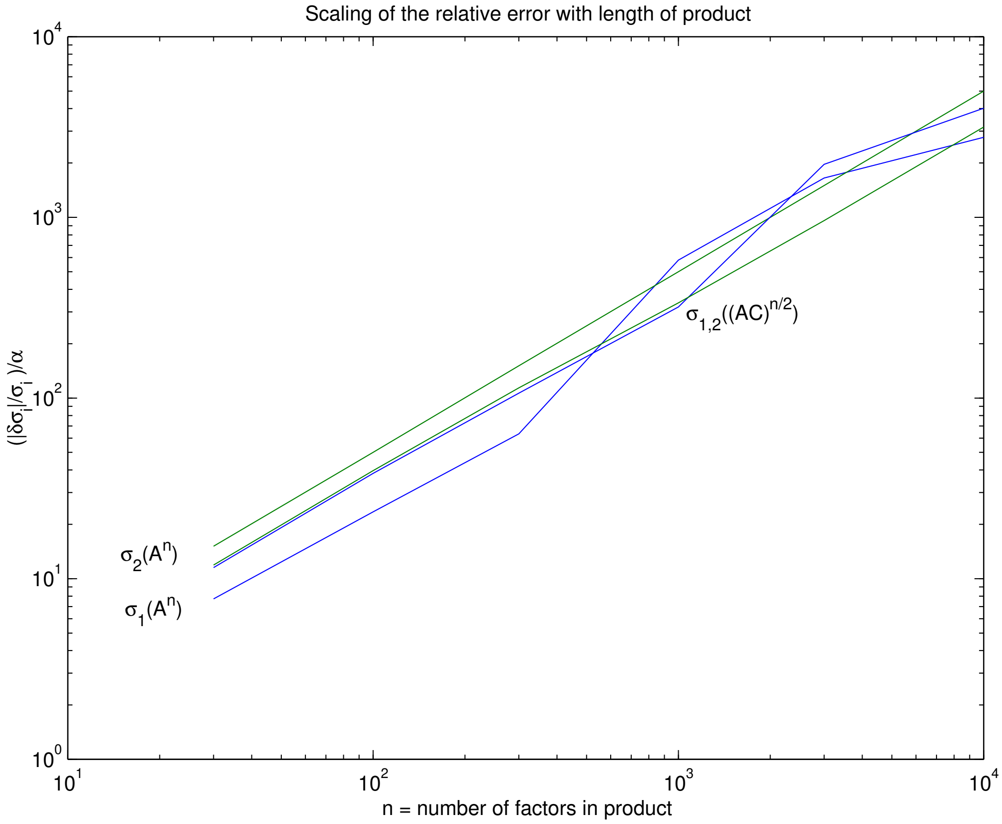

FIG. 3.1. *Scaling of errors with length of product n*

the size of the perturbations and the errors in $\log(\sigma_i)$. The errors in $\log(\sigma_i)$ divided by α are plotted against the length of the product in Figure 3.1. Since $|\log(\sigma_i + \delta\sigma_i) - \log(\sigma_i)| \ll 1$, this difference is very nearly $|\delta\sigma_i|/\sigma_i$. Note that the errors scale proportionally to n, rather than $n^{1.58}$ or n^2. This suggests that the exponent in the error bound can probably be reduced further.

However, for behavior of the errors for the test product $A^{n/2}B^{n/2}$, which does not have an exponential splitting, is rather different for large n. The errors in the logarithms of the singular values are plotted against n in Figure 3.2. Note that the errors in $\log(\sigma_i)$ for $n = 30$ are about $2^{n/2} \approx 3.2 \times 10^4$ times the perturbation level. For larger values of n, the errors in $\log(\sigma_i)$ become quite large indicating that the computed singular values become meaningless for these perturbed products.

Appendix: A multiplicative perturbation theorem for eigenvalues. In this appendix the following theorem is proven:

THEOREM 3.1. *Suppose that A is diagonalizable with $X^{-1}AX = \Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$. Then for a suitable ordering $\widetilde{\lambda}_1, \widetilde{\lambda}_2, \ldots, \widetilde{\lambda}_n$, of the eigenvalues of $(I+E)A$, and if $\beta = \kappa_1(X)\|E\|_1 < 1$, we have*

$$\frac{|\widetilde{\lambda}_i - \lambda_i|}{|\lambda_i|} \le (1+\beta)\left(\frac{1+\beta}{1-\beta}\right)^{n-1} - 1.$$

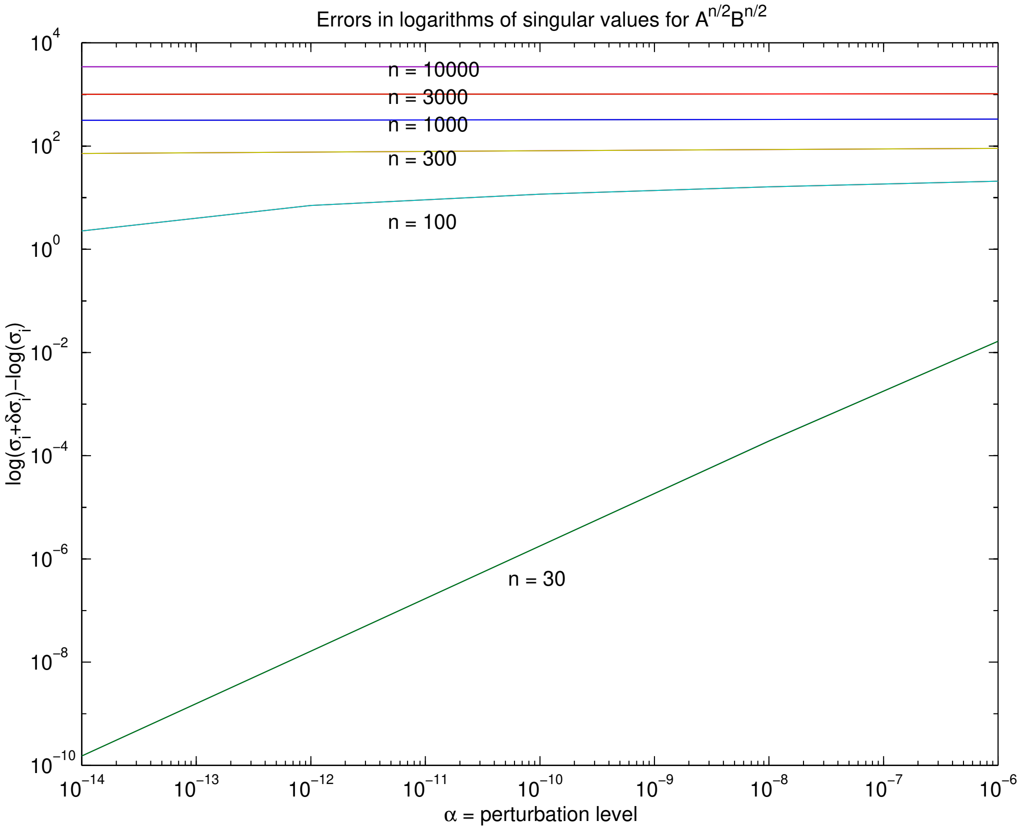

FIG. 3.2. *Errors in* $\log(\sigma_i)$ *for a product without exponential splitting*

Proof. The proof is based on Gershgorin's theorem (see [13, §6.1, p. 344] or [11, §7.2, p. 341]). First note that

$$X^{-1}(I+E)AX = (I+X^{-1}EX)X^{-1}AX = (I+X^{-1}EX)\Lambda.$$

Let $F = X^{-1}EX$. Then $\|F\|_1 \leq \kappa_1(X)\|E\|_1 = \beta$. Now we find the Gershgorin disks for $(I+F)\Lambda$ using column sums. The jth Gershgorin disk is

$$\begin{aligned} D_j &= \{\, z \in \mathbf{C} \mid |z-(1+f_{jj})\lambda_j| \leq |\lambda_j| \textstyle\sum_{i\neq j} |f_{ij}| \,\} \\ &\subseteq \{\, z \in \mathbf{C} \mid |z-\lambda_j| \leq |\lambda_j|\, \|F\|_1 \,\} \\ &\subseteq \{\, z \in \mathbf{C} \mid |z-\lambda_j| \leq |\lambda_j|\, \beta \,\} = \widehat{D}_j. \end{aligned}$$

Since there are exactly k eigenvalues of $(I+E)A$ in each connected component of $\bigcup_j D_j$ containing exactly k disks, then any connected component of $\bigcup_j \widehat{D}_j$ with exactly k disks $\widehat{D}_j$ contains exactly k eigenvalues.

Suppose that $\widehat{D}_i \cap \widehat{D}_j \neq \emptyset$. Let $z \in \widehat{D}_i \cap \widehat{D}_j$. Then $|z-\lambda_i| \leq |\lambda_i|\,\beta$ and $|z-\lambda_j| \leq |\lambda_j|\,\beta$. So $|\lambda_i-\lambda_j| \leq |\lambda_i - z| + |z-\lambda_j| \leq (|\lambda_i|+|\lambda_j|)\beta$. Also, $|\lambda_j| \leq |\lambda_i| + |\lambda_j-\lambda_i| \leq |\lambda_i| + (|\lambda_i|+|\lambda_j|)\beta$. Rearranging gives $|\lambda_j| \leq [(1+\beta)/(1-\beta)]|\lambda_i|$, and also $|\lambda_i-\lambda_j| \leq (2\beta)/(1-\beta)|\lambda_i|$.

Choose an ordering of the eigenvalues $\widetilde{\lambda}_1, \ldots, \widetilde{\lambda}_n$ so that $\widetilde{\lambda}_i$ is in the same connected component of $\bigcup_j \widehat{D}_j$ as λ_i, for each i. Such an ordering is possible by Gershgorin's theorem.

For each i there is a chain of pairs of overlapping disks connecting λ_i with $\widetilde{\lambda}_i$. Consider a chain of pairs of overlapping disks: $\widehat{D}_{i_1} \cap \widehat{D}_{i_2} \neq \emptyset$, $\widehat{D}_{i_2} \cap \widehat{D}_{i_3} \neq \emptyset$, $\ldots$, $\widehat{D}_{i_{k-1}} \cap \widehat{D}_{i_k} \neq \emptyset$ where $\widetilde{\lambda}_{i_1} \in \widehat{D}_{i_k}$. Then by the above remarks about overlapping pairs of disks,

$$
\begin{aligned}
|\lambda_{i_{j+1}} - \lambda_{i_j}| &\leq \frac{2\beta}{1-\beta}|\lambda_{i_j}|, \qquad (a)\\
|\lambda_{i_{j+1}}| &\leq \frac{1+\beta}{1-\beta}|\lambda_{i_j}|, \qquad (b)
\end{aligned}
\tag{3.1}
$$

for $j = 1, 2, \ldots, k-1$. Then (3.1b) immediately gives

$$
|\lambda_{i_j}| \leq \left(\frac{1+\beta}{1-\beta}\right)^{j-1} |\lambda_{i_1}|.
$$

Now using (3.1a),

$$
\begin{aligned}
|\lambda_{i_k} - \lambda_{i_1}| &\leq \sum_{j=1}^{k-1} |\lambda_{i_{j+1}} - \lambda_{i_j}| \leq \sum_{j=1}^{k-1} |\lambda_{i_j}| \frac{2\beta}{1-\beta}\\
&\leq \sum_{j=1}^{k-1} |\lambda_{i_1}| \left(\frac{1+\beta}{1-\beta}\right)^{j-1} \frac{2\beta}{1-\beta}\\
&\leq |\lambda_{i_1}| \left[\left(\frac{1+\beta}{1-\beta}\right)^{k-1} - 1\right].
\end{aligned}
$$

Then

$$
\begin{aligned}
\frac{|\widetilde{\lambda}_{i_1} - \lambda_{i_1}|}{|\lambda_{i_1}|} &\leq \frac{|\widetilde{\lambda}_{i_1} - \lambda_{i_k}|}{|\lambda_{i_1}|} + \frac{|\lambda_{i_1} - \lambda_{i_k}|}{|\lambda_{i_1}|}\\
&\leq \frac{|\lambda_{i_k}|}{|\lambda_{i_1}|} \frac{|\widetilde{\lambda}_{i_1} - \lambda_{i_k}|}{|\lambda_{i_k}|} + \frac{|\lambda_{i_1} - \lambda_{i_k}|}{|\lambda_{i_1}|}\\
&\leq \left(\frac{1+\beta}{1-\beta}\right)^{k-1} \beta + \left(\frac{1+\beta}{1-\beta}\right)^{k-1} - 1\\
&= (1+\beta)\left(\frac{1+\beta}{1-\beta}\right)^{k-1} - 1.
\end{aligned}
$$

Thus with a suitable ordering of $\widetilde{\lambda}_1, \ldots, \widetilde{\lambda}_n$,

$$
\frac{|\widetilde{\lambda}_i - \lambda_i|}{|\lambda_i|} \leq (1+\beta)\left(\frac{1+\beta}{1-\beta}\right)^{k-1} - 1,
$$

and since there cannot be more than n disks in a chain of overlapping pairs, the desired result follows. ◻

Note that for small β, $|\widetilde{\lambda}_i - \lambda_i|/|\lambda_i| \leq (2n-1)\beta + O(n\beta)^2$. In practice, there are usually very few Gershgorin disks that overlap; if there are no more than k disks in a connected component of $\bigcup_j \widehat{D}_j$, then the above bound can be reduced to

$$
\frac{|\widetilde{\lambda}_i - \lambda_i|}{|\lambda_i|} \leq (1+\beta)\left(\frac{1+\beta}{1-\beta}\right)^{k-1} - 1 = (2k-1)\beta + O(k\beta)^2.
$$

REFERENCES

[1] H. D. I. Abarbanel, R. Brown, and M. B. Kennel, *Local Lyapunov exponents computed from observed data*, Chaos: An Int. J. of Nonlinear Science, 2 (1992), pp. 343–365.

[2] A. Barrlund, *Perturbation bounds for the LDL^H and LU decompositions*, BIT, 31 (1991), pp. 358–363.

[3] A. Bojanczyk, G. H. Golub, and P. V. Dooren, *The periodic Schur decomposition: Algorithms and applications*, Tech. Report NA-92-07, Stanford University, August 1992.

[4] A. W. Bojanczyk, M. Ewerbring, F. T. Luk, and P. V. Dooren, *An accurate product SVD algorithm*, in SVD and Signal Processing II, R. J. Vaccaro, ed., Amsterdam, New York, 1991, Elsevier Science Publishers, pp. 113–131.

[5] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, MIT Press and McGraw-Hill, New York, St. Louis, San Francisco and Montreal, 1990.

[6] L. Dieci, R. D. Russell, and E. S. Van Vleck, *On the computation of Lyapunov exponents for continuous dynamical systems*, SIAM J. Numer. Anal., 34 (1997), pp. 402–423.

[7] Z. Drmač, *Accurate computation of the product-induced singular value decomposition with applications*, SIAM J. Numer. Anal., 35 (1998), pp. 1969–1994.

[8] S. C. Eisenstat and I. C. F. Ipsen, *Relative perturbation techniques for singular value problems*, SIAM J. Matrix Anal. Appl., 32 (1995), pp. 1972–1988.

[9] ———, *Relative perturbation results for eigenvalues and eigenvectors of diagonalisable matrices*, BIT, 38 (1998), pp. 502–509.

[10] ———, *Three absolute perturbation bounds for matrix eigenvalues imply relative bounds*, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 149–158 (electronic).

[11] G. Golub and C. V. Loan, *Matrix Computations*, Johns Hopkins Press, 2nd ed., 1989. 1st edition published 1983 by North Oxford Academic.

[12] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, PA, 1996.

[13] R. A. Horne and C. A. Johnson, *Matrix Analysis*, Cambridge Uni. Press, Cambridge, 1985.

[14] I. C. F. Ipsen, *Relative perturbation results for matrix eigenvalues and singular values*, Acta Numerica, 7 (1998), pp. 151–201.

[15] R. J. Sacker and G. R. Sell, *A spectral theory for linear differential systems*, J. Differential Equations, 27 (1978), pp. 320–358.

[16] D. E. Stewart, *Product algorithms for eigensystems*, Tech. Report ACTR-31-03-94, Australian National University, Centre for Mathematics and its Applications, 1994. Avail. at `http://wwwmaths.anu.edu.au/research.groups/advcomp/acreports`.

[17] ———, *A new algorithm for the SVD of a long product of matrices*, Electron. Trans. Numer. Anal., 5 (1997), pp. 29–47.

[18] G. W. Stewart, *On graded QR decompositions of products of matrices*, Electronic Trans. Numer. Anal., 3 (1995), pp. 39–49.

[19] J. Sun, *Componentwise perturbation bounds for some matrix decompositions*, BIT, 32 (1992), pp. 702–714.