Electronic Transactions on Numerical Analysis.
Volume 4, pp. 37-45, March 1996.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

A NOTE ON THE ACCURACY OF SYMMETRIC EIGENREDUCTION ALGORITHMS *

K. VESELIĆ †

Abstract. We present some experimental results illustrating the fact that on highly ill–conditioned Hermitian matrices the relative accuracy of computed small eigenvalues by QR eigenreduction may drastically depend on the initial permutation of the rows and columns. Mostly there was an "accurate" permutation, but there does not seem to be an easy method to get at it. For banded matrices, like those from structural mechanics, the accurate pre–permutation, if it existed, was mostly non–banded. This is particularly true of tridiagonal matrices which shows that the tridiagonalization is not the only factor responsible for the inaccuracy of the eigenvalues.

Key words. LAPACK, QR method, Jacobi method, Hermitian matrices, eigenvalue computation.

AMS subject classification. 65F15.

1. Introduction.

Classical error analysis of common symmetric eigenreduction algorithms like QR[1] or Jacobi is based on two facts: (i) the use of orthogonal elementary transformations and (ii) the spectral norm estimate

$$|\delta\lambda| \leq \|\delta H\| \ , \tag{1.1}$$

for a perturbed eigenvalue $\lambda + \delta\lambda$ of a perturbed Hermitian matrix $H + \delta H$. Note that this estimate allows large relative errors in absolutely small eigenvalues. This is, of course, reflected in the output of the eigenreduction algorithms. For instance the positive definite matrix

$$\begin{pmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1.000001 \end{pmatrix} \tag{1.2}$$

has the small eigenvalue $\lambda = 3.333331 \cdot 10^{-7}$. The same eigenvalue, computed by single-precision ($macheps \approx 10^{-8}$) versions of the two algorithms, reads

$$\lambda_{QR} = 3.179 \cdot 10^{-7}, \ \lambda_{Jacobi} = 3.209 \cdot 10^{-7}, \tag{1.3}$$

which about reaches the estimate (1.1). However, according to the results in [5], a positive definite matrix H may determine its eigenvalues better than guaranteed by (1.1); we have

$$|\delta\lambda/\lambda| \leq cond(A) \max_{ij} |\delta H_{i,j}/H_{i,j}| \ , \tag{1.4}$$

where $A = DHD$ and D is any non-singular diagonal matrix. Similar, although less definite results hold for indefinite matrices ([10]). Moreover, in the positive definite case the Jacobi algorithm is relatively accurate i.e. it computes the eigenvalues with a (1.4)-like error. Exhaustive experiments ([5, 9, 8]) have shown that QR cannot

* Received May 10, 1995. Accepted for publication, January 8, 1996. Communicated by J. Demmel.
† Lehrgebiet Mathematische Physik, Fernuniversität Hagen, Postf. 940 58084 Hagen, Germany.
[1] We use the term QR for qr or ql based algorithms preceded by a tridiagonalisation step.

reach this type of the accuracy (cf. also a theoretical result in [3]). In the indefinite case Jacobi in its standard form—like the routine *jacobi* from [11]—does no better and another, more sophisticated version of Jacobi had to be designed ([9]). This method uses indefinite Cholesky decomposition as a preconditioner and continues with hyperbolic Jacobi iterations and is shown to have error behaviour of type (1.4) [8].

The aim of this note is to call attention to the phenomenon: *the accuracy of the QR algorithm may drastically depend on the permutation of rows and columns of the input matrix.* This was observed in extensive numerical experiments for the papers [9] and [8].[2] The same is valid for Jacobi (except, of course, in the definite case). We will illustrate and discuss this phenomenon on a set of examples. The examples will consist of *well-behaved* matrices, i.e., those which, in spite of their possible high condition, determine well their eigenvalues according to (1.4). Of course, our initial example (1.2) does not belong to the class of well-behaved matrices in the mentioned sense. Our experimental results using LAPACK solvers indicate the following facts:

1. For most matrices there is an initial permutation improving the QR accuracy in the sense of (1.4). This means that it was rather difficult to find matrices on which no permutation carried improvement in accuracy. Here, of course, the evidence is poor, since testing all permutations restricts the dimension to at most 7-8. LAPACK version 1.0 had the tendency to prefer the column-norm decreasing permutation, version 2.0 the reverse, but in neither case the effect was strong enough to recommend this pre–permutation in general. Even more erratic was the behaviour on banded matrices like those from structural mechanics, where an optimal permutation was never banded.
2. In particular, on initially tridiagonal matrices a most accurate permutation is mostly non-tridiagonal; this implies that the tridiagonalization part of the algorithm is not exclusively responsible for the loss of accuracy (cf. [3]).
3. The QR accuracy (for the same initial permutation) may vary with the hardware used (which usually implies a different compiler). In particular, the 'V' option (with eigenvectors) may differ drastically from the 'N' option (no eigenvectors).
4. For the symmetric Jacobi method the column-decreasing order almost always visibly improves the accuracy–unless, of course, the matrix is positive definite and Jacobi is accurate for any permutation [5]. This effect is not strong enough to recommend symmetric Jacobi method for indefinite matrices.
5. Of all used methods the modified Jacobi [9] stands out because of its usually high accuracy, which recommends it as a method of choice when accurate eigenreducing is desired.
6. Examples of matrices on which even our modified Jacobi lacks the expected accuracy are very rare [9],[8]. This inaccuracy is known to be due to the insufficiency of the complete pivoting strategy of indefinite Cholesky preconditioner [2] and on known examples of this kind an *ad hoc* change of the pivoting strategy will produce the expected accuracy. In any case, in all our experiments no matrix was found on which a QR solver with an 'omniscient pre–permutation' would beat our modified Jacobi, again with an 'omniscient' pivoting strategy.

Our QR algorithm is *ssyev, dsyev* from LAPACK, Version 1.0. The Jacobi al-

[2] The first who called our attention to this phenomenon was C. Moler (personal communication, 1990).

gorithms are *sjac*–a standard Jacobi algorithm–as well as *ssyevj* and *dsyevj*–implicit Jacobi algorithms (available from netlib)–based on [9] and [8] which produce a greater relative accuracy.[3] Fortran programs are implemented on an IBM RISC-6000 and this is the standard situation. "Critical cases" have been re-run on a 486-PC with the same fortran codes, but, of course, other hardware and compiler. The latter results are reported only if they sensibly deviate from the standard ones.

We also repeated some experiments using LAPACK 2.0 routines. The results are qualitatively the same but may greatly differ on individual matrices. These are presented at the end of the paper. Again, it is worthwhile to mention the significant, sometimes drastic difference change in accuracy between LAPACK QR solvers with and without eigenvector option.

We accepted common initial digits of the eigenvalues, obtained by *dsyev* and *dsyevj*, as "correct". A matrix is accepted as well-behaved if *ssyevj* computed its eigenvalues correctly. All tested matrices had, of course, a condition number of at least 10^6 so that the classical, norm error analysis of small eigenvalues is of little use in single precision computations. We skipped reporting standard Jacobi (*sjac*) results whenever they are well-known to be accurate, e.g., in positive definite cases.

2. Random matrices. We tested several hundreds of matrices of the dimension $n \leq 50$ and the form

$$H = DAD, \tag{2.1}$$

where D is a random diagonal matrix. The matrix A was chosen in two ways: (i) with random elements between -1 and 1 and (ii) as UD_0U^T, where D_0 is a given diagonal matrix and U a random orthogonal matrix. In case (i) pre–permuting the matrix as above was always a full success. For instance, take a typical such matrix

$$\begin{pmatrix} -9.990E+07 & -8.740E+03 & -4.440E+07 \\ -8.740E+03 & -1.890E-01 & 9.160E+02 \\ -4.440E+07 & 9.160E+02 & 3.130E+07 \end{pmatrix}, \tag{2.2}$$

which, decreasingly permuted reads

$$\begin{pmatrix} -9.990E+07 & -4.440E+07 & -8.740E+03 \\ -4.440E+07 & 3.130E+07 & 9.160E+02 \\ -8.740E+03 & 9.160E+02 & -1.890E-01 \end{pmatrix}. \tag{2.3}$$

Here the experiments yielded the following maximal relative errors in all eigenvalues:

$$\begin{array}{rcccc} & dsyev & ssyevj & ssyev & sjac \\ unpermuted & 1.2E-08 & 2.9E-07 & 2.4E+01 & 5.5E-07 \\ permuted & 2.0E-15 & 2.9E-07 & 3.8E-07 & 5.3E-07. \end{array} \tag{2.4}$$

Here, and in the following the eigenvalues computed by *dsyevj* are taken as correct. The improvement produce by a permutation on values from *ssyev* is dramatic. A similar situation was observed in case (ii). However, in that case, for a cleverly chosen D_0 a decreasing permutation produced a less significant improvement. For instance, consider the matrix

[3] The beginning letter "s" or "d" in the algorithm name mean as usual single or double precision version, respectively.

ETNA
Kent State University
etna@mcs.kent.edu

$$
\begin{pmatrix}
3.646E+04 & -6.317E+02 & 2.389E+02 & -1.124E+04 \\
-6.317E+02 & 1.791E+04 & 1.097E+01 & -5.161E+02 \\
2.389E+02 & 1.097E+01 & 5.501E+00 & 1.952E+02 \\
-1.124E+04 & -5.161E+02 & 1.952E+02 & 2.195E+04
\end{pmatrix}. \tag{2.5}
$$

Its decreasingly ordered permutation is

$$
\begin{pmatrix}
3.646E+04 & -1.124E+04 & -6.317E+02 & 2.389E+02 \\
-1.124E+04 & 2.195E+04 & -5.161E+02 & 1.952E+02 \\
-6.317E+02 & -5.161E+02 & 1.791E+04 & 1.097E+01 \\
2.389E+02 & 1.952E+02 & 1.097E+01 & 5.501E+00
\end{pmatrix}. \tag{2.6}
$$

Here the relative errors were

$$
\begin{array}{rcccc}
 & dsyev & ssyevj & ssyev & sjac \\
unpermuted & 8E-10 & 4E-05 & 6E-02 & 5E-05 \\
permuted & 5E-13 & 4E-05 & 1E-01 & 1E-04
\end{array} \tag{2.7}
$$

which, effectively, means that the permutation produced no improvement. A search over all possible permutations, however, found

$$
\begin{pmatrix}
5.501E+00 & 1.952E+02 & 2.389E+02 & 1.097E+01 \\
1.952E+02 & 2.195E+04 & -1.124E+04 & -5.161E+02 \\
2.389E+02 & -1.124E+04 & 3.646E+04 & -6.317E+02 \\
1.097E+01 & -5.161E+02 & -6.317E+02 & 1.791E+04
\end{pmatrix} \tag{2.8}
$$

with the maximal errors

$$
\begin{array}{cccc}
dsyev & ssyevj & ssyev & sjac \\
1E-13 & 4E-05 & 9E-05 & 3E-06,
\end{array} \tag{2.9}
$$

which are satisfactory.

3. Some special matrices. We begin with a tridiagonal matrix whose well-behavedness is well-known ([1, 4]):

$$
H = \begin{pmatrix}
0 & w & 0 & 0 & 0 & 0 \\
w & 0 & 1 & 0 & 0 & 0 \\
0 & 1 & 0 & w & 0 & 0 \\
0 & 0 & w & 0 & 1 & 0 \\
0 & 0 & 0 & 1 & 0 & w \\
0 & 0 & 0 & 0 & w & 0
\end{pmatrix} \tag{3.1}
$$

with $w = 1E-03$. Here the eigenvalues occur in plus-minus pairs. The correct value of small eigenvalue pair is

$$\pm 9.999990E-10.$$

The routine *ssyev* computed the values $1.2295756E-09$ and $-8.1312388E-10$ and *sjac* computed the values $2.9E-08$ and $-3.5E-11$. When the matrix was permuted to a columns decreasing order, *ssyev* improved a little: $9.92E-10$, $-1.01E-09$. The corresponding *sjac* eigenvalues were accurate: $\pm 9.9999886E-10$. However on a 486-PC *sjac* gave only $1.00003E-09$, $-9.9997E-10$. The best results for *ssyev* result were obtained with the permutation

$$\begin{pmatrix} 0 & 0 & 1 & 0 & 0 & w \\ 0 & 0 & 0 & w & 0 & 0 \\ 1 & 0 & 0 & 0 & w & 0 \\ 0 & w & 0 & 0 & 1 & 0 \\ 0 & 0 & w & 1 & 0 & 0 \\ w & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \tag{3.2}$$

where the computed small eigenvalues were $9.993E-10$, $-1.001E-09$. Note howerver that this accuracy was obtained on 14 other permutations! For matrix (3.2), however, *sjac* performed poorly obtaining values $2.3E-11$ and $-4.3E-08$. This example shows that inaccuracies of the QR method do not lie only in the tridiagonalization step. Quite often a permutation which needed to be subsequently tridiagonalized produced better accuracy than the initial tridiagonal matrix did!

We now take some examples from structural mechanics. Consider the eigenvalue problem

$$Kx = \lambda Mx, \tag{3.3}$$

where M, the mass matrix, is diagonal and positive definite whereas the K, the stiffness matrix, is positive definite and sparse. Here in fact, our experiments are performed on the derived matrix $M^{-1/2}KM^{-1/2}$.

Longitudinal vibrations of a beam are described by

$$K = (n+1)^2 \begin{pmatrix} 2 & -1 & & & \\ -1 & 2 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & -1 \\ & & & -1 & 2 \end{pmatrix}. \tag{3.4}$$

First take $n = 5$ and $M = diag(1, 1, 10^6, 1, 1)$ (one heavy mass point). This yields the matrix

$$\begin{pmatrix} 72 & -36 & 0 & 0 & 0 \\ -36 & 72 & -0.036 & 0 & 0 \\ 0 & -0.036 & 0.000072 & -0.036 & 0 \\ 0 & 0 & -0.036 & 72 & -36 \\ 0 & 0 & 0 & -36 & 72 \end{pmatrix}$$

with the eigenvalues

$$\begin{array}{l} 1.08000012E+02 \\ 1.08000000E+02 \\ 3.60000360E+01 \\ 3.60000000E+01 \\ 2.39999800E-05. \end{array}$$

Here the smallest eigenvalue is of interest. The routine *ssyev* produced the value

ETNA
Kent State University
etna@mcs.kent.edu

$1.96E-05$ ($2.312E-05$ on 486-PC). For the column-decreasing permutation,

$$\begin{pmatrix} 72 & 0 & -36 & 0 & -0.036 \\ 0 & 72 & 0 & -36 & -0.036 \\ -36 & 0 & 72 & 0 & 0 \\ 0 & -36 & 0 & 72 & 0 \\ -0.036 & -0.036 & 0 & 0 & 0.000072 \end{pmatrix}, \tag{3.5}$$

ssyev gave $1.2E-05$. However, there was a permutation,

$$\begin{pmatrix} 0.000072 & 0 & -0.036 & -0.036 & 0 \\ 0 & 72 & -36 & 0 & 0 \\ -0.036 & -36 & 72 & 0 & 0 \\ -0.036 & 0 & 0 & 72 & -36 \\ 0 & 0 & 0 & -36 & 72 \end{pmatrix}, \tag{3.6}$$

where *ssyev* gave $2.39999899E-05$, and there were 23 more permutations with this accuracy!

Another example from structural mechanics is that of a transversally vibrating beam where K from (3.4) is replaced by its square. Take now $n=6$, and let $M = diag(1, 1, 10^6, 10^6, 1, 1)$. This yields the matrix

$$\begin{pmatrix} 12005 & -9604 & 2.401 & 0 & 0 & 0 \\ -9604 & 14406 & -9.604 & 2.401 & 0 & 0 \\ 2.401 & -9.604 & 0.014406 & -0.009604 & 2.401 & 0 \\ 0 & 2.401 & -0.009604 & 0.014406 & -9.604 & 2.401 \\ 0 & 0 & 2.401 & -9.604 & 14406 & -9604 \\ 0 & 0 & 0 & 2.401 & -9604 & 12005 \end{pmatrix}. \tag{3.7}$$

The correct eigenvalues are

$$\begin{array}{c} 2.2884245E+04 \\ 2.2884243E+04 \\ 3.5267703E+03 \\ 3.5267620E+03 \\ 8.4034726E-03 \\ 1.7149986E-04. \end{array} \tag{3.8}$$

Here *ssyev* computed values $9.3E-03$ and $3.7E-04$ (values $8.6E-03$, and $-9.8E-04$ on the 486-PC) as the smallest eigenvalues. With the column-norm decreasing ordering,

$$\begin{pmatrix} 14406 & 0 & -9604 & 0 & 2.401 & -9.604 \\ 0 & 14406 & 0 & -9604 & -9.604 & 2.401 \\ -9604 & 0 & 12005 & 0 & 0 & 2.401 \\ 0 & -9604 & 0 & 12005 & 2.401 & 0 \\ 2.401 & -9.604 & 0 & 2.401 & 0.014406 & -0.009604 \\ -9.604 & 2.401 & 2.401 & 0 & -0.009604 & 0.014406 \end{pmatrix}, \tag{3.9}$$

ssyev computed values $1E-02$ and $7E-04$ which was not a success. A permutation which produced optimal results is

$$\begin{pmatrix} 0.014406 & 0 & 2.401 & 2.401 & -0.009604 & -9.604 \\ 0 & 12005 & -9604 & 0 & 2.401 & 0 \\ 2.401 & -9604 & 14406 & 0 & -9.604 & 0 \\ 2.401 & 0 & 0 & 12005 & 0 & -9604 \\ -0.009604 & 2.401 & -9.604 & 0 & 0.014406 & 2.401 \\ -9.604 & 0 & 0 & -9604 & 2.401 & 14406 \end{pmatrix}. \tag{3.10}$$

Here *ssyev* computed values $8.402E-03$ and $1.702E-04$, but the PC did much worse obtaining values $7.75E-03$ and $-6.03E-04$.

Some earlier experiments of ours seem also to indicate that accuracy gets even worse, if a "band preserving" tridiagonalisation is used.

Note also a recent result in [6] where a simple criterion is given to check whether a rotation is dangerous –at least in the positive definite case. It was emphasized also , both in [9] and in [6], that in cases where a rotation is dangerous that it is better to factor H (this need not always be just the Cholesky factor). Subsequently one should perform a singular value decomposition on a factor This has the advantage that it deals with a matrix whose condition number is the square root of the original one. Note that the Cholesky decomposition of a well-behaved positive definite matrix is always accurate [5]. But, even if the matrix is not well behaved, a careful evaluation of its factor may save the accuracy. In [6] the Hilbert matrix is given as an example of this. Our matrix (1.2) is not less interesting in this respect. Actually, this matrix can be written as

$$H = LL^T$$

with

$$L^T = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 0.001 \end{pmatrix}.$$

The matrix L^T determines its singular values well, being bidiagonal; see [4]. So, methods for computing accurate factors of matrices with given structure may be a promising subject for future research.

We finally produce a well-behaved matrix on which even our modified Jacobi is inaccurate. Take

$$H = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & \varepsilon \end{pmatrix} \quad , \quad \varepsilon = 1E-07 \tag{3.11}$$

with the small, well defined eigenvalue $5.000000E-08$. Here QR was correct on the permutations $1, 2, 3$ and $1, 3, 2$. The best *sjac* value $4.85E-08$ was obtained on the permutation $3, 1, 2$. Here even our reference algorithm *ssyevj* was not much better. This is due to the fact that in this case the indefinite symmetric decomposition [2] (which is the first part of *ssyevj* and which here begins with a simple Gaussian elimination step) is inaccurate [9, 10, 8].[4] According to [10, 8] the well-behavedness of the matrix H in (3.11) is due to the factorization

$$H = GJG^T$$

[4] It was probably Rutishauser [7] who first produced this matrix as an example where the common complete pivoting in Gaussian elimination is sensibly less accurate than some other strategies.

ETNA
Kent State University
etna@mcs.kent.edu

with

$$G = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & \sqrt{\varepsilon} \end{pmatrix}, \ J = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

This factorization is obtained, when the indefinite symmetric decomposition algorithm is set to make an initial 2×2 elimination step, instead of the 1×1 one step which is prescribed by the strategy given in [2]. With this change *ssyevj* becomes accurate [8]. This example shows, that even in the indefinite case, good factorizations may be essential for accuracy.

4. Recent LAPACK version (added in proof). We briefly present the experimental results on the same types of matrices described above using the recent LAPACK version 2.0 compiled with Microsoft FORTRAN Visual Workbench v. 1.1 with its default options on a 486-PC. Qualitatively these results are similar to those described above. In most cases the version without eigenvectors (*syev* with the flag 'N') was more accurate than the one with them (*syev* with the flag 'V'). It was difficult to find a matrix of small dimension on wich both options would be inaccurate in all permutations. Although *syev* uses QR or QL version depending of the direction of grading, we found matrices on which the initial permutation, say 1,2,3,4 produced very different results from those on the reverse one 4,3,2,1. Here the column-norm increasing order has a slight advantage, but not significant enough to be generally recommended.

We illustrate this with two examples. The first example is

$$\begin{pmatrix} -1.0000000000000000E+9 & -4.2948904163016274E-1 & 2.0647076965986236 \\ -4.2948904163016274E-1 & -5.5792431861865712E-2 & 8.1059562122547863 \\ 2.0647076965986236E+0 & 8.1059562122547863E+0 & 1.9719676162003452 \end{pmatrix}, \tag{4.1}$$

which has a condition number about 10^8. Here both 'N' and 'V' option were fully accurate on permutations 3,2,1 and 2,3,1. On 1,3,2 we had three correct figures, all others were even less accurate.

The second example is

$$\begin{pmatrix} 3E+0 & 1E+0 & 1E+0 & -1E+0 \\ 1E+0 & 1E+8 & 1E+0 & 1E+0 \\ 1E+0 & 1E+0 & 3E+8 & 1E+0 \\ -1E+0 & 1E+0 & 1E+0 & 1E+0 \end{pmatrix}, \tag{4.2}$$

which has (correctly rounded) eigenvalues

$$3.0000000E+8 \quad 1.0000000E+8 \quad 3.4142136E+0 \quad 5.8578641E-1,$$

and its condition number is about 10^8. Here no permutation and no option was able to produce a single accurate digit for the smallest eigenvalue. For the second smallest eigenvalue some permutations computed (with the 'N' flag) one accurate digit.[5]

Ackowledgement. The author is indebted to I. Slapničar, Z. Drmač, E. Pietzsch and W. Schumann for their comments and/or their help with the software. He is also indebted to the referee for his very useful suggestions.

[5] We tried this matrix with some more single-precision QR codes at hand. None produced better results.

REFERENCES

[1] J. Barlow and J. Demmel, *Computing Accurate Eigensystems of Scaled Diagonally Dominant Matrices*, SIAM J. Num. Anal., 27(1988), pp. 762-791.
[2] J. R. Bunch, B. N. Parlett, *Direct methods for solving symmetric indefinite solutions of linear equations*, SIAM J. Numer. Anal., 8 (1971), pp. 639-655.
[3] J. Demmel, *The Inherent Inaccuracy of Implicit Tridiagonal QR*, IMA Preprint Series 963, University of Minnesota, Minneapolis, 1992.
[4] J. Demmel and W. Kahan,*Accurate Singular Values of Bidiagonal Matrices*, SIAM J. Sci. Stat. Comp., 11 (1990), pp. 873-912.
[5] J. Demmel and K. Veselić,*Jacobi's method is more accurate than QR*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 1204-1245.
[6] R. Mathias,Accurate Eigensystem Computations by Jacobi Methods IMA preprint, 1993.
[7] H. Rutishauser, *Vorlesungen Über Numerische Mathematik*, Birkhäuser Stuttgart 1976.
[8] I. Slapničar,*Accurate symmetric eigenreduction by a Jacobi method*, Fernuniversität Hagen, Ph. D. Thesis 1992.
[9] K. Veselić, *A Jacobi eigenreduction algorithm for definite matrix pairs*, Numer. Math. 64 (1993), pp. 241-269.
[10] K. Veselić and I. Slapničar, Floating-point perturbations of Hermitian matrices, Linear Algebra Appl., 195 (1993), pp. 81-116.
[11] J. H. Wilkinson and C. Reinsch, *Handbook for automatic computation: Vol. II–Linear Algebra*, Springer-Verlag, New York, 1971.