The evaluation of eigenvalues and eigenvectors of real symmetric matrices by simultaneous iteration

M. Clint* and A. Jenningst

* Department of Computer Science, The Queen's University of Belfast
† Department of Civil Engineering, The Queen's University of Belfast

A method is described of obtaining all or a subset of the eigenvalues and corresponding eigenvectors of real symmetric matrices by iterating simultaneously with a number of trial vectors. Its relationship with a previous simultaneous iteration method is discussed and the results of some numerical tests are given. These methods are useful in cases where a few of the dominant eigenvalues and eigenvectors of large matrices are required.

(Received December 1968)

1. Introduction

The evaluation of the eigenvalues and eigenvectors of a symmetric matrix on a digital computer is normally performed by methods based on that of Householder (see Wilkinson, 1960) in which transformations are first carried out on the matrix to reduce it to tridiagonal form. This process is then followed by a procedure for obtaining the eigenvalues of the tridiagonal matrix and then one for constructing the required eigenvectors of the original matrix. In cases where only a few of the eigenvalues and eigenvectors of large matrices are required (particularly when these are sparse) or when good estimates of the eigenvectors are available from previous tests on similar matrices, iteration methods may be more efficient.

The well known iteration method is the power method attributed to Von Mises (see Bodewig, 1956) in which a trial vector is continually premultiplied by the matrix until the iterates become proportional to each other. This process yields the eigenvalue of largest modulus together with its eigenvector. An extension of this method (see Jennings, 1967) involves iterating simultaneously with two or more trial vectors and is a more general iterative approach which provides a powerful alternative to transformation methods for those problems which are well suited to iterative solution.

Within the basic concept of simultaneous iteration there are various ways of revising the set of trial vectors during each iteration. This paper derives a more fundamental approach related to which the original method is seen to be a special case.

2. The interaction matrix

Consider a symmetric matrix $\boldsymbol{A}$ of order n whose eigenvalues are λ_1, λ_2, etc., in order of magnitude, and whose corresponding normalised eigenvectors are $\boldsymbol{q}_1, \boldsymbol{q}_2$, etc., such that

$$\boldsymbol{A}\boldsymbol{q}_i = \lambda_i \boldsymbol{q}_i \tag{1}$$

and

$$\boldsymbol{q}_i^T \boldsymbol{q}_j = \delta_{ij} \tag{2}$$

where δ_{ij} is the Kronecker delta. Let the matrix $\boldsymbol{u} = [\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_m]$ represent a set of m trial vectors. Simultaneous premultiplication of all the trial vectors by the matrix produces a second set of vectors

$$\boldsymbol{v} = [\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_m]$$

given by

$$\boldsymbol{v} = \boldsymbol{A}\boldsymbol{u} \tag{3}$$

If the trial vectors happened to be eigenvectors of the matrix $\boldsymbol{A}$ then the product

$$\boldsymbol{B} = \boldsymbol{u}^T\boldsymbol{v} = \boldsymbol{u}^T\boldsymbol{A}\boldsymbol{u} \tag{4}$$

would be a diagonal matrix with m eigenvalues of $\boldsymbol{A}$ as elements. Although $\boldsymbol{B}$ will be symmetric it will not, in general, be a diagonal matrix. The size of the off-diagonal elements will depend on the degree of interaction between the various trial vectors when put through the premultiplication process of equation (3). The matrix $\boldsymbol{B}$ is of fundamental significance in simultaneous iteration methods and so the title 'interaction matrix' is proposed.

If the matrix of trial vectors is square ($m = n$) and orthogonal then a sufficient condition for the trial vectors to have converged to the eigenvectors of the matrix $\boldsymbol{A}$ would be that the interaction matrix is diagonal. However in the case where the trial vectors are not orthogonal or $m < n$, it is possible to have trial vectors which are not eigenvectors but which produce a diagonal interaction matrix. In this case the condition that the interaction matrix is diagonal is necessary but not sufficient for convergence.

Any trial vector may be expressed as a linear combination of the eigenvectors, hence

$$\boldsymbol{u} = \boldsymbol{qC} \tag{5}$$

where $\boldsymbol{C}$ is an $n \times m$ matrix of coefficients and $\boldsymbol{q} = [\boldsymbol{q}_1, \boldsymbol{q}_2, \ldots, \boldsymbol{q}_n]$. Substituting for $\boldsymbol{u}$ in equation (4) gives

$$\boldsymbol{B} = \boldsymbol{C}^T\boldsymbol{q}^T\boldsymbol{AqC} = \boldsymbol{C}^T\boldsymbol{\Lambda}\,\boldsymbol{C} \tag{6}$$

where $\boldsymbol{\Lambda}$ is a diagonal matrix of the eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$.

A method for improving the accuracy of eigenvector estimates where all eigenvectors are present is given by Patton (1966). This may be described as a simultaneous iteration method and operates with unorthogonal trial vectors. The method to be adopted in this paper follows the method of Jennings (1967) in that only orthogonal sets of trial vectors are used satisfying the relationship

$$\boldsymbol{u}^T\boldsymbol{u} = \boldsymbol{I} \tag{7}$$

and hence

$$\boldsymbol{C}^T\boldsymbol{C} = \boldsymbol{I} \tag{8}$$

If $m = n$ then from equations (6) and (8) it is seen that $\boldsymbol{C}^T$ is the matrix of eigenvectors of $\boldsymbol{B}$ and the diagonal elements of $\boldsymbol{\Lambda}$ are the eigenvalues of the original matrix $\boldsymbol{A}$. Therefore if $\boldsymbol{u}$ is an orthogonal set of vectors which approximate closely the eigenvectors of $\boldsymbol{A}$, the matrix $\boldsymbol{B}$ is a unitary transformation of $\boldsymbol{A}$ which is nearly diagonal and is thus easy to reduce by the Jacobi method. When $m < n$ there are not enough algebraic equations to define the coefficients of $\boldsymbol{C}$ from equations (6) and (8) and hence a less direct approach has to be adopted.

3. The significance of the eigen-reduction of the interaction matrix

Numbering the eigenvalues of $\boldsymbol{A}$ in descending order of modulus and writing

$$\boldsymbol{u} = \boldsymbol{q}_a\boldsymbol{C}_a + \boldsymbol{q}_b\boldsymbol{C}_b \tag{9}$$

where $\boldsymbol{q}_a$ is a matrix of the first m eigenvectors of $\boldsymbol{A}$, $\boldsymbol{q}_b$ is a matrix of the eigenvectors $m + 1$ to n and $\boldsymbol{C}_a$ and $\boldsymbol{C}_b$ are coefficient matrices, $\boldsymbol{C}_a$ being square, equation (6) becomes

$$\boldsymbol{B} = \boldsymbol{C}_a^T\boldsymbol{\Lambda}_a\boldsymbol{C}_a + \boldsymbol{C}_b^T\boldsymbol{\Lambda}_b\boldsymbol{C}_b \tag{10}$$

where $\boldsymbol{\Lambda}_a$ and $\boldsymbol{\Lambda}_b$ are diagonal matrices with elements $\lambda_1, \ldots, \lambda_m$ and $\lambda_{m+1}, \ldots, \lambda_n$ respectively.

Equation (8) becomes

$$\boldsymbol{C}_a^T\boldsymbol{C}_a + \boldsymbol{C}_b^T\boldsymbol{C}_b = \boldsymbol{I} \tag{11}$$

If the coefficients $\boldsymbol{C}_b$ are assumed to be small compared with the coefficients $\boldsymbol{C}_a$ then the eigenvalues of $\boldsymbol{B}$ are seen to be approximately the m largest eigenvalues of $\boldsymbol{A}$ and the matrix of eigenvectors of $\boldsymbol{B}$ (called $\boldsymbol{P}$) is approximately $\boldsymbol{C}_a^T$.

4. Iteration procedure

The set of vectors $\boldsymbol{v}$ formed from equation (3) can be expressed in the form

$$\boldsymbol{v} = \boldsymbol{q}_a\boldsymbol{\Lambda}_a\boldsymbol{C}_a + \boldsymbol{q}_b\boldsymbol{\Lambda}_b\boldsymbol{C}_b. \tag{12}$$

By ignoring the contributions involving the coefficients $\boldsymbol{C}_b$ in equations (11) and (12), a linear manipulation of this vector set giving a closer approximation to the first m eigenvectors is given by

$$\boldsymbol{w} = \boldsymbol{vP} \simeq \boldsymbol{vC}_a^T \simeq \boldsymbol{q}_a\boldsymbol{\Lambda}_a. \tag{13}$$

In general the predicted vectors $\boldsymbol{w}$ will not be an orthogonal set and before they can be used as new starting vectors they have to be subjected to an orthogonalising process. The method described in Jennings (1967) can be adopted in which each subdominant vector is orthogonalised with respect to all vectors dominant to it. The process can be summarised as a linear transformation

$$\bar{\boldsymbol{u}} = \boldsymbol{wS} \tag{14}$$

where $\boldsymbol{S}$ is obtained by a reduction of the matrix

$$\boldsymbol{D} = \boldsymbol{w}^T\boldsymbol{w}. \tag{15}$$

If the vectors are in descending order according to the modulus of their predicted λ values, then $\boldsymbol{S}$ will be an upper triangular matrix equal to $\boldsymbol{L}^{-T}$ ($\boldsymbol{L}$ being the Choleski lower triangular matrix). In the general case the largest remaining diagonal element is taken as pivot at any stage in the reduction. This will confuse the simple triangular form of $\boldsymbol{S}$ but the computational process will be very similar.

Thus an iterative cycle has been formed in which the size of the discrepancy between successive trial vectors

$$\boldsymbol{e}_i = (\text{sign } \lambda_i)\bar{\boldsymbol{u}}_i - \boldsymbol{u}_i \tag{16}$$

may be used to determine when sufficient accuracy has been obtained to halt the process.

The full iteration procedure can be summarised in the following way:

1. Multiply $\boldsymbol{v} = \boldsymbol{Au}$;
2. Form the interaction matrix $\boldsymbol{B} = \boldsymbol{u}^T\boldsymbol{v}$;
3. Derive the estimates of λ_a and form $\boldsymbol{P}$, the matrix of eigenvectors of $\boldsymbol{B}$;
4. Form $\boldsymbol{w} = \boldsymbol{vP}$;
5. Form $\boldsymbol{D} = \boldsymbol{w}^T\boldsymbol{w}$;
6. Derive $\boldsymbol{S}$ from $\boldsymbol{D}$;
7. Form $\bar{\boldsymbol{u}} = \boldsymbol{vS}$;
8. Test for convergence, returning to step 1 if further iteration is required.

The premultiplication $\boldsymbol{v} = \boldsymbol{Au}$ has the effect of diminishing the components of the lower eigenvectors and hence at each iteration the coefficients $\boldsymbol{C}_b$ will be reduced. Because the coefficients $\boldsymbol{C}_b$ only occur in squared form in equations (10) and (11) it should be expected that the eigenvalues will converge more rapidly than the corresponding eigenvectors. The only operation involving the matrix $\boldsymbol{A}$ is a multiplication and therefore simultaneous iteration methods may be particularly suitable for obtaining eigenvalues and eigenvectors of sparse matrices.

5. Significance of linear and non-linear prediction

In cases where the interaction matrix has a very strong diagonal its matrix of eigenvectors will also have strong elements, and the eigenvectors can be so ordered that these strong elements all lie on the leading diagonal. If a typical element of the eigenvector matrix is P_{ij} then from the eigenvalue properties of $\boldsymbol{B}$

$$\sum_{k=1}^{n} b_{ik}P_{kj} = \lambda_j P_{ij}. \tag{17}$$

Putting $i = j$ and ignoring minor contributions in the summation

$$\lambda_i \simeq b_{ii}. \tag{18}$$

Similarly for $i \neq j$

$$b_{ii}P_{ij} + b_{ij}P_{jj} \simeq \lambda_j P_{ij}, \tag{19}$$

which gives

$$P_{ij} \simeq -b_{ij}/(b_{ii} - b_{jj}). \tag{20}$$

Equation (18) is the formula used for estimating the eigenvalues in Jennings (1967) and equation (20) is identical to the linear prediction. The non-linear prediction employed in this reference is thus seen to be a simple approximation to the eigenvectors of the interaction matrix which approaches the linear prediction and therefore the true eigenvectors as the matrix $\boldsymbol{B}$ approaches diagonal form. It can therefore be expected to be efficient in the later rounds of iteration but the question arises as to whether a more accurate method of predicting the eigenvectors of $\boldsymbol{B}$ would have a beneficial effect in the initial stages of the iteration.

6. Jacobi eigen-reduction of interaction matrix

In the case where the number of total vectors is much smaller than the order of the matrix $\boldsymbol{A}$ then the evaluation of the eigenvalues and eigenvectors of $\boldsymbol{B}$ will be numerically much shorter than the eigen-reduction of $\boldsymbol{A}$ and it is therefore feasible to use any numerical procedure for it. However as the results are to be fed back into an iterative process, high accuracy is not required. The Jacobi method (see Ralston and Wilf, 1960) commends itself from the point of view of simplicity and the fact that low accuracy results may be obtained fairly quickly particularly in the latter stages of iteration when the interaction matrix will be almost diagonal.

7. Repeated premultiplication

The process of forming the interaction matrix, performing an eigen-reduction of it and using this to predict the eigenvectors could be described as performing an interaction analysis. An interaction analysis consists of steps 2, 3 and 4 in the iteration cycle of paragraph 4. Steps 5, 6 and 7 could be described as an orthogonalising sequence. The iteration procedure can therefore be briefly summarised by the following four processes taken in sequence:

(*a*) Premultiplication;
(*b*) Interaction Analysis;
(*c*) Orthogonalising Sequence;
(*d*) Test for Convergence.

In cases where m is less than n but not so much less that the interaction analysis and orthogonalising procedure take insignificant amounts of computing time, then it may be desirable to perform an interaction analysis less frequently than once every premultiplication. This may be done by adding one or more premultiplication operations between the interaction analysis and the orthogonalising sequence. If this procedure is adopted care has to be taken to see that the accuracy required for all the eigenvectors can be achieved.

8. Accuracy of the eigenvectors

In order to investigate the maximum possible accuracy that may be achieved using simultaneous iteration methods the effect of rounding errors will be considered on the assumption that convergence is otherwise complete. The rounding errors in a vector $\boldsymbol{u}_i$ may be separated into components in each of the eigenvectors as follows:

$$\boldsymbol{u}_i = \boldsymbol{q}_i + \Sigma e_{ij}\boldsymbol{q}_j. \tag{21}$$

The coefficients e_{ij} will be of order $\epsilon = 10^{-s}$ where s is the number of significant figures in the computation. On premultiplication these error components will be magnified in proportion to their associated eigenvalues giving

$$\boldsymbol{v}_i = \lambda_i\boldsymbol{q}_i + \Sigma\,(e_{ij}\lambda_j + {}^{e}{}_{ij}\lambda_i)\boldsymbol{q}_j, \tag{22}$$

where $\bar{e}_{ij}$ are rounding errors introduced in the premultiplication which will be assumed to be of order ϵ. The largest errors in a vector $\boldsymbol{v}_i$ will therefore have a relative magnitude $\lambda_1\epsilon/\lambda_i$.

On forming the interaction matrix a general off-diagonal element at position (i, j) will have error terms of order $\epsilon\lambda_i$, $\epsilon\lambda_j$ and $\epsilon^2\lambda_k$ where k can take any value other than i and j. As the off-diagonal elements will tend to be small the linearising assumption of paragraph 5 may be used to show that the vector $\boldsymbol{w}_i$ will be given by

$$\boldsymbol{w}_i = \lambda_i\boldsymbol{q}_i + \Sigma\,\bar{\bar{e}}_{ij}\boldsymbol{q}_j, \tag{23}$$

where $\bar{\bar{e}}_{ij}$ consists of terms of order $\epsilon\lambda_i^2/(\lambda_i - \lambda_j)$, $\epsilon\lambda_i\lambda_j/(\lambda_i - \lambda_j)$ and $\epsilon^2\lambda_k\lambda_j/(\lambda_i - \lambda_j)$. There will be a magnification of rounding errors in $\boldsymbol{w}_i$ if any of the coefficients $\bar{\bar{e}}_{ij}$ exceeds $\epsilon\lambda_i$. Three circumstances in which these coefficients can significantly exceed $\epsilon\lambda_i$ are as follows:

(*a*) $\lambda_j \gg \lambda_i$: In this case the errors introduced in $\boldsymbol{w}_i$ are not important because there is no corresponding error in the vector $\boldsymbol{w}_j$; hence, as $\boldsymbol{w}_i$ is orthogonalised with respect to $\boldsymbol{w}_j$ in the orthogonalising process, this error will not be carried forward to the next trial vector $\boldsymbol{u}_i$.
(*b*) $\lambda_i - \lambda_j \ll \lambda_i$;
(*c*) $\lambda_i < \epsilon\lambda_k$.

The last two cases are both circumstances in which the eigenvectors would be expected to have a degree of indeterminacy matching the loss in the predicted accuracy. When there is a pair of equal eigenvalues convergence cannot be obtained for the eigenvectors associated with the equal pair, and slight loss in accuracy of the other eigenvectors has been recorded in test cases.

However if instead of one premultiplication, h premultiplications are carried out in each iteration loop then the rounding error associated with the largest eigenvector in the estimate of $\boldsymbol{q}_i$ will magnify to order $\lambda_i^h\epsilon/\lambda_1^h$. Although this particular error component should be eliminated on orthogonalisation it could dominate the vector if h is large and thereby impair the accuracy of the predicted eigenvectors.

Errors will also arise due to coefficients of the matrix $\boldsymbol{C}_b$. These coefficients will diminish at each iteration but the slowest to diminish will be those associated with eigenvalue $m + 1$. At each premultiplication the component of the $(m + 1)$th eigenvector in the trial vector converging on the jth eigenvector will be reduced by the

factor λ_{m+1}/λ_j. The interaction analysis will not be effective in eliminating this error component as it is only concerned with interaction between the first m trial vectors. Hence neglecting the disturbing effect of this analysis and of the orthogonalising procedure on the convergence process, the accuracy of the jth vector after k iterations will be of the order of $(\lambda_{m+1}/\lambda_j)^k$. This rule gives an approximate idea of the total number of premultiplications required to obtain a given accuracy although in practice the number of iterations will be affected to a certain extent by the method and accuracy of the predictions for the eigenvectors of the interaction matrix. This rule can be used to show that in most cases it will be advantageous to carry more trial vectors than necessary in order to speed the convergence. For instance, if the eigenvalues of a matrix are

$$10, 9, 6, 4{\cdot}5, 3, 2, 1, 0{\cdot}5, \ldots$$

and if four eigenvectors are required to four-figure accuracy, then, by iterating with only four trial vectors, about 23 premultiplications would be required, whereas iterating with five or six trial vectors the number of premultiplications is likely to be reduced to about 12 and 7 respectively. Even allowing for the larger size of the matrices this is likely to introduce a saving in computation time. Other advantages are a greater probability that very close eigenvalues will not have a serious effect on the convergence, an approximate knowledge of the next subdominant eigenvalues and eigenvectors and more assurance that the calculated eigenvectors will be within the stipulated tolerance (when convergence is slow the change in a vector at each iteration will be less than the total error in the vector at this stage).

9. Numerical results

Tests were carried out on the matrix given in **Table 1** using four methods which will be given the code names A1, A3, B1 and B3. These four methods are described as follows:

A1: The method of reference (see Jennings, 1967).

A3: Similar to A1 but with three premultiplications in each iteration loop.

B1: The method described in paragraph 6 where the eigen-reduction of the interaction matrix was performed by the Jacobi method. The tolerance for the Jacobi method was such that the off-diagonal elements were reduced to less than 10^{-6}.

B3: Similar to B1 but with three premultiplications in each iteration loop.

Table 2 shows the number of iterations required to achieve six-figure accuracy in the eigenvectors using different numbers of trial vectors for each of the four methods. The corresponding eigenvalues always achieved the maximum machine accuracy of 11 figures. In each case the $\boldsymbol{u}$ matrix was taken as a segment of the unit matrix for the first round of iteration. In brackets is shown the number of premultiplications where this is different from the number of iterations. **Table 3** shows the recorded computing times in secs on an ICT 1905 computer. These times are given for comparison between the four methods but should not be used as a general guide to the speed because they are elapsed times and

Table 2

Number of iterations for four alternative methods (Where number of premultiplications is different it is shown in brackets)

METHOD	A1	A3	B1	B3
8 trial vectors 4 roots and vectors required	14	8(24)	12	5(15)
3 trial vectors 3 roots and vectors required	11	6(18)	10	4(12)

Table 1

12 × 12 Matrix with approximate eigenvalues

21·54	27·62	35·67	41·42	48·36	55·32	4·43	11·27	18·42	25·43	32·38	39·32
	73·68	103·40	132·90	164·00	194·80	10·82	36·50	67·23	98·25	129·00	159·70
		192·40	265·90	342·30	419·10	17·30	65·22	136·10	212·80	289·40	365·90
			428·30	576·00	722·20	23·71	94·30	210·90	355·10	504·30	653·10
				843·60	1,090·30	30·05	122·90	285·10	501·90	747·70	977·30
					1,510·60	36·42	151·60	359·30	649·30	996·90	1,373·00
						10·81	17·21	23·53	29·88	36·25	42·61
	Symmetric						49·90	78·73	107·40	136·10	164·90
								157·70	232·10	306·20	380·40
									383·80	530·60	677·50
										788·10	1,036·90
											1,424·70

Approximate eigenvalues:

5,333	236	187	38·1	32·9	14·8	12·7	8·6	6·8	6·5	4·1	2·8

not processor times for development programmes, the efficiency of which could be significantly improved. Comparing methods A1 and B1 it is seen that the use of the Jacobi eigen-reduction decreases slightly the number of iterations required, but may not lead to a reduction in computing time because of the extra time involved in the more exact eigen-reduction.

Table 3

Computing time for four alternative methods

METHOD	A1	A3	B1	B3
8 trial vectors 4 roots and vectors required	61	52	67	39
3 trial vectors 3 roots and vectors required	18	14	24	10

When using three multiplications per iteration the total number of iterations was always reduced at the expense of having a greater number of premultiplications. The increase in computing speed obtained showed that for these examples the computing time was more dependent on the number of iterations than on the number of premultiplications, and that the more efficient method was the one in which the more accurate eigen-reduction was employed. For large matrices, in which the ratio m/n is small, it would be expected that the number of premultiplications would become the dominant factor in determining the computing time, making B1 the most suitable method.

The possibility of reducing the stipulated tolerance for the Jacobi routine was investigated but it was found that too low a Jacobi tolerance could prevent convergence of the main iteration loop. This is a danger inherent in nesting one iteration loop inside another. However an alternative procedure was tested in which a fixed number of rotations was performed at each entry to the Jacobi routine thus eliminating the iterative nature of the inner loop. It was found that by adopting this procedure for methods B1 and B3 an approximate 5% gain in computing time could be obtained. For the given example the optimum number of rotations was found to be about m^2. As the number of rotations in each eigen-reduction was reduced the number of iterations tended to increase so as to be comparable with the appropriate values for methods A1 and A3. On the other hand, with a large number of rotations the eigen-reduction of the interaction matrix is effectively accurate and so the number of iterations will be the same as that obtained if any direct method were employed to perform the eigen-reduction. Use of the Jacobi routine in this way may thus be regarded as a bridge between the approximate eigen-reduction of the interaction matrix embodied in the non-linear prediction of Jennings (1967) and accurate methods of performing the eigen-reduction.

10. Conclusion

A more general concept of the simultaneous iteration method of obtaining the eigenvalues and eigenvectors of symmetric matrices has been derived than in the development previously given in Jennings (1967). The significance of the interaction matrix has been discussed, in particular its eigensolution. Some numerical tests have been performed using a Jacobi reduction to obtain eigenvector estimates of the interaction matrix with both one and three premultiplications for every interaction analysis. The results of these tests show that some increase in computing efficiency may be obtained over the method of Jennings (1967) at the expense of a slightly more complicated technique.

References

BODEWIG, E. (1956). *Matrix Calculus*, North Holland Publishing Co., Amsterdam.

JENNINGS, A. (1967). A Direct Iteration Method for Obtaining the Latent Roots and Vectors of a Symmetric Matrix, *Proc. Camb. Phil. Soc.*, Vol. 63, pp. 755–765.

PATTON, P. C. (1966). A New Numerical Method for the Simultaneous Computation of the Eigenvectors and Eigenvalues of Arbitrary Dynamical Matrices, Dr. Ing. Thesis, Stuggart.

RALSTON, A., and WILF, H. S. (1960). *Mathematical Methods for Digital Computers*, John Wiley and Sons Inc.

WILKINSON, J. H. (1960). Householder's Method for the Solution of the Algebraic Eigenproblem, *The Computer Journal*, Vol. 2, pp. 23–27.