

Solution of the Complete Symmetric Eigenproblem in a Virtual Memory Environment

Abstract: Algorithms for the complete symmetric eigenproblem are studied from the viewpoint of performance in a virtual memory environment. The preferred algorithm is based on tridiagonalization and the implicit QR algorithms. Implementation factors that affect execution time are discussed.

Introduction

Virtual memory and paging devices have considerably simplified the development of programs operating on large data arrays. Storage management problems that were left to the programmer's ingenuity are now automatically dealt with by the system.

However, some precautions must be taken in order to implement algorithms that are efficient in the context of virtual memory with paging. This question is of particular importance in the solution of matrix problems [1]. When designing a program involving matrix operations, the choice of the algorithms is guided not only by numerical considerations, such as stability, accuracy, and count of arithmetic operations, but also by other factors more intimately related to the virtual memory. A "good" algorithm should result in a program that minimizes the number and the extent of the virtual memory sweeps required in referring to the elements of the matrices involved.

In this paper, a few methods for the solution of the complete symmetric eigenproblem are examined in a virtual memory context, and the implementation of the preferred algorithm is studied. Additional considerations are also taken into account:

- that performance of the algorithm is also optimal in a real memory environment.
- that extra programming effort required by virtual memory considerations is minimal, and

- that the overall legibility of the resulting program is good in the sense that the code closely reflects the numerical algorithm.

Matrices are assumed to be stored row-wise as is the case in PL/I and BASIC. However, techniques similar to those described here apply as well to FORTRAN [1] (matrices stored column-wise). Experiments were run on an IBM System/3, Model 6 using the BASIC software paging facilities.

Choice of method

Two methods are considered for their numerical properties of stability, accuracy, and reliability [2]: Jacobi's method with threshold, and a combination of Householder's tridiagonalization with the implicit QR algorithm. These methods are briefly described below (detailed information is provided in the references).

• Jacobi's method with threshold

The matrix is diagonalized by means of similarity transformations [3, 4].

Given a positive threshold, each off-diagonal element of the matrix that is larger than the threshold in absolute value is annihilated in turn by a similarity transformation based on a plane rotation. For an off-diagonal subscripted element (i, j) , this implies the transformation of rows and columns i and j of the matrix. Once all off-diagonal

elements are smaller than the threshold, this threshold is decreased and the above procedure is repeated. The whole process is iterated until the threshold is negligible in terms of machine precision.

The eigenvectors are obtained by forming the product of the plane rotations involved in the diagonalization. The ultimate convergence of the process is quadratic.

For a matrix of order n , the minimum storage requirements are:

- One one-dimensional array of $n(n+1)/2$ locations for the symmetric matrix, and
- One $(n \times n)$ array, where the matrix of eigenvectors is constructed.

This widely known method is quite attractive for its elegance and its compactness. Yet, one can see the problems it raises in a virtual memory environment: there are extensive sweeps of the memory for each plane rotation, that is, the transformation of two rows and two columns of the matrix. When forming the matrix of the eigenvectors, extra sweeps are required from the array of the matrix to the array of the eigenvectors and back, and also within the array of the eigenvectors (transformation of two rows for each plane rotation). These sweeps are at the kernel of the iterative process.

• Householder's tridiagonalization and implicit QR

The original matrix is first reduced to a symmetric tridiagonal matrix by orthogonal similarity transformations based on Householder's matrices [5, 2]. The resulting tridiagonal matrix is then diagonalized by using the QR (or QL) algorithm with implicit shift [6, 7]. The ultimate convergence of the QR algorithm is cubic.

The matrix of eigenvectors is built by forming the product of the Householder matrices of the tridiagonalization stage and the plane rotations of the QR algorithm.

In a real memory environment, this method is the most efficient for the treatment of the complete eigenproblem.

For a matrix of order n , the storage requirements are:

- One two-dimensional array of n^2 locations for the matrix of eigenvectors, which initially contains the symmetric matrix (see Ref. 5, procedure tred2), and
- Two one-dimensional arrays of n locations for the diagonal and subdiagonal elements of the symmetric tridiagonal matrix.

In addition to being more economical of storage, this method presents the substantial advantage that only one large array is involved in the iterative process (QR) for the computation of the eigenvectors. This means that the number and the extent of the virtual memory sweeps is more limited than in the Jacobi algorithm. One more factor in favor of the QR algorithm is better convergence properties.

• Remarks on inverse iteration for the eigenvectors

A third possibility was considered because no operation on large arrays was imbedded in an iterative process: a combination of tridiagonalization and QR for the eigenvalues, and Varah's inverse iteration for the eigenvectors [8].

Although this approach leads to a program sometimes as efficient as one based on the QR algorithm, it presents serious drawbacks: The resulting program is about twice as voluminous as the program based on tridiagonalization and QR, and the storage requirements are higher than for the Jacobi method. Furthermore, the orthogonality of the eigenvectors is not guaranteed, the eigenvectors corresponding to clustered eigenvalues may be poorly determined.

The preferred algorithm is thus based on tridiagonalization and the QR algorithm for the computation of the eigenvalues and eigenvectors.

Implementation of the tridiagonalization—QR method

• Tridiagonalization

Let A be a symmetric matrix. Its reduction to tridiagonal form can be expressed by the following:

$$\begin{aligned} A^{(0)} &= A, \\ A^{(k+1)} &= P^{(k)} A^{(k)} P^{(k)}, \quad k = 0, \dots, n-3, \end{aligned} \quad (1)$$

where $A^{(k)}$ is a symmetric matrix that is tridiagonal in its last k rows and where $P^{(k)}$ is the Householder matrix that reduces the $(n-k)$ th row of $A^{(k)}$ to tridiagonal form. $A^{(n-2)}$ is the resulting tridiagonal matrix similar to A .

$P^{(k)}$ is orthogonal and of the form

$$P^{(k)} = I + \alpha^{(k)} v^{(k)} v^{(k)T},$$

where $\alpha^{(k)}$ is a scalar and $v^{(k)}$ a vector chosen to annihilate the pertinent elements of row $(n-k)$ of $A^{(k)}$

Equation (1) is equivalent to

$$A^{(k+1)} = A^{(k)} + w h^T + h w^T, \quad (2)$$

where

$$w = \alpha^{(k)} v^{(k)} \text{ and}$$

$$h = A^{(k)} v^{(k)} + [(v^{(k)T} A^{(k)} v^{(k)})/2] w. \quad (3)$$

(Superscripts applicable to w and h have been omitted for clarity.)

Since the successive $A^{(k)}$'s are symmetric, only their lower triangular part is formed, including the diagonal.

Equation (2), rewritten as

$$a_{i,j}^{(k+1)} = a_{i,j}^{(k)} + w_i h_j + w_j h_i,$$

shows that, once vectors w and h are computed, the elements of the array containing $A^{(k)}$ can be modified to produce $A^{(k+1)}$ in the order that best suits the structure of the virtual memory, that is, row-wise. The main computational part of Eq. (3) is the calculation of $A^{(k)} v^{(k)}$.

$v^{(k)}$ is defined by

$$\begin{aligned} v_i^{(k)} &= a_{n-k,i}^{(k)}, & 1 \leq i < n-k-1, \\ v_i^{(k)} &= a_{n-k,i}^{(k)} - p^{(k)}, & i = n-k-1, \\ v_i^{(k)} &= 0 \text{ otherwise,} \end{aligned}$$

with

$$p^{(k)} = -\text{sign}(a_{n-k,n-k-1}^{(k)}) \left[\sum_{i=1}^{n-k-1} (a_{n-k,i}^{(k)})^2 \right]^{1/2}.$$

Since the last $(k+1)$ components of $v^{(k)}$ are zero, the multiplication $A^{(k)} v^{(k)}$ involves the leading principal submatrix of $A^{(k)}$ of order $(n-k-1)$ only. To efficiently perform the multiplication, one would like to address each pertinent element of the lower triangular part of $A^{(k)}$ once and only once, the matrix being swept row-wise.

This can be achieved as shown in the sequence of BASIC statements below, where N is the order of the matrix, A , the array of $A^{(k)}$, V , the array of $v^{(k)}$, X , the array of $A^{(k)} v^{(k)}$, and M represents $(n-k-1)$.

```

1  FOR I = 1 TO M
2  S = 0
3  T = V(I)
4  FOR J = 1 TO I - 1
5  U = A(I, J)
6  X(J) = X(J) + U*T
7  S = S + U*V(J)
8  NEXT J
9  X(I) = S + A(I, I)*T
10 NEXT I

```

In the implementation of the tridiagonalization process, it is essential for good performance that a separate one-dimensional array be allocated for vector $v^{(k)}$, although its elements can be retrieved from the $(n-k)$ th row of the array containing $A^{(k)}$. The nonzero parts of w and h are stored in one-dimensional arrays, where the diagonal and subdiagonal elements of the resulting tridiagonal matrix are stored as they are obtained. Once the matrix has been reduced to tridiagonal form, the product of the Householder transformations is formed as a first step in the construction of the matrix of eigenvectors. This product is computed in the array used for the $A^{(k)}$ s. In order to allow this product to be expressed in terms of row operations, the product $P^{(0)} P^{(1)} \cdots P^{(n-3)}$ (premultiplying matrices) is formed, which eventually leads to the matrix of left eigenvectors (transpose of the matrix of eigenvectors).

When the off-tridiagonal part of the $(n-k-1)$ row of $A^{(k)}$ is negligible in terms of machine precision, it is advantageous to skip a transformation and to write $A^{(k+1)} \equiv A^{(k)}$. The following simple test can be used for this purpose. The transformation is skipped when

$$\text{fl} \left\{ \text{fl} \left[\sum_{i=1}^{n-k-2} (a_{n-k,i}^{(k)})^2 \right] + \text{fl} \{ (a_{n-k,n-k-1}^{(k)})^2 \} \right\} \equiv \text{fl} \{ (a_{n-k,n-k-1}^{(k)})^2 \},$$

where $\text{fl} \{ \}$ denotes values obtained by machine floating-point computation.

• QR algorithm for a symmetric tridiagonal matrix

The following is the QR form of the implicit QL algorithm of [6, 7]. The shift at the k th iteration is $\sigma^{(k)}$, $\{e_i^{(k)}; i = 2, \dots, n\}$ are the subdiagonal elements of the tridiagonal matrix, and $\{d_i^{(k)}; i = 1, \dots, n\}$ are the diagonal elements. At the k th iteration, the shift is chosen to be the eigenvalue of the bottom 2×2 principal submatrix that is closer to $d_n^{(k)}$. When $e_n^{(k)}$ can be considered as negligible, $d_n^{(k)}$ is taken as an eigenvalue and the order n of the matrix is reduced by one (by neglecting the last row and column). This process is repeated on the successive reduced matrices until the whole matrix is diagonalized. The formulas reflecting one iteration are

$$c_1 = s_1 = 1, q_1 = d_1^{(k)} - \sigma^{(k)}, u_1 = 0.$$

$$h_i = c_{i-1} e_i^{(k)},$$

$$p_i = s_{i-1} e_i^{(k)},$$

$$e_{i=1}^{(k+1)} = (p_i^2 + q_{i-1}^2)^{1/2},$$

$$s_i = p_i / e_{i-1}^{(k+1)},$$

$$c_i = q_i / e_{i-1}^{(k+1)},$$

$$g_i = d_{i-1}^{(k)} - u_{i-1},$$

$$t_i = (d_i^{(k)} - g_i) s_i + 2 c_i h_i,$$

$$u_i = s_i t_i,$$

$$d_{i-1}^{(k+1)} = g_i + u_i,$$

$$q_i = c_i t_i - h_i,$$

$$d_n^{(k+1)} = d_n^{(k)} - u_n, e_n^{(k+1)} = q_n;$$

where $i = 2, \dots, n$.

This QR transformation consists of performing on the tridiagonal matrix $(n-1)$ similarity transformations based on plane rotations. The plane rotations apply to pairs of consecutive rows and columns of the matrix, taken in their natural order: $(1, 2), (2, 3), \dots$. The i th plane rotation is defined by the parameters c_{i+1} and s_{i+1} .

The matrix of left eigenvectors is obtained by premultiplying the matrix of accumulated transformations

[illegible]

The formulas describing the QR iteration show that practically no performance improvement is to be expected from a special formulation of the algorithm for virtual memory.

It should be noted that no algorithmic modification of the method was necessary for an adaptation to a virtual memory environment. The critical points are:

- These implementation details do not affect the performance of the method in a real memory environment.

For the eigenvalues only, the advantage of the tridiagonalization-QR is substantially increased, since no operation on a large array is imbedded in the iterative process. Furthermore, there is no need to form the product of the transformations involved in the tridiagonalization stage.

The matrix-vector multiplication then becomes

Note on FORTRAN implementation

The matrix of *right* eigenvectors is constructed by forming the product $\mathbf{P}^{(n-3)} \mathbf{P}^{(n-2)} \cdots \mathbf{P}^{(0)}$ and then by postmultiplying this product by the transposes of the above QR plane rotations. All these operations are performed column-wise.

In general, the time savings obtained by using the combination of tridiagonalization-QR rather than Jacobi's method are very substantial. These savings are highly dependent on the size of the arrays, the size of the virtual memory page and other characteristics of the system used.

With arrays allowing for a maximum matrix order of 25, the combination tridiagonalization-QR was about five times as fast as Jacobi's method for random matrices of order 20.

1. C. B. Moler, "Matrix Computations with Fortran and Paging," *Communications ACM* 15, No. 4 (April, 1972).
2. J. H. Wilkinson, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, 1965.
3. J. Greenstadt, "The Determination of the Characteristic Roots of a Matrix by the Jacobi Method," *Mathematical Methods for Digital Computers*, Vol. I, Editors Ralston and Wilf, John Wiley & Sons, New York, 1960.
4. H. Rutishauser, "The Jacobi Method for Real Symmetric Matrices," *Handbook for Automatic Computation, Vol. II, Linear Algebra*, Editors Wilkinson and Reinsch, Springer, Berlin, 1971.

5. R. S. Martin, C. Reinsch, and J. H. Wilkinson, "Householder Tridiagonalization of a Symmetric Matrix," *Handbook for Automatic Computation, Vol. II, Linear Algebra*, Editors Wilkinson and Reinsch, Springer, Berlin, 1971.
6. A. Dubrulle, "A Short Note on the Implicit QL Algorithm for Symmetric Tridiagonal Matrices," *Numerische Mathematik*, **15**, p. 450 (1970).
7. A. Dubrulle, R. S. Martin, and J. H. Wilkinson, "The Implicit QL Algorithm," *Handbook for Automatic Computation, Vol. II, Linear Algebra*, Editors Wilkinson and Reinsch, Springer, Berlin, 1971.
8. G. Peters, and J. H. Wilkinson, "The Calculation of Specified Eigenvectors by Inverse Iteration," *Handbook for Automatic Computation, Vol. II, Linear Algebra*, Editors Wilkinson and Reinsch, Springer, Berlin, 1971.

Received June 2, 1972

The author is located at the IBM Data Processing Division Development Center, Palo Alto, California. 94301