SIAM REVIEW
Vol. 18, No. 4, October 1976

ILL-CONDITIONED EIGENSYSTEMS AND THE COMPUTATION OF THE JORDAN CANONICAL FORM*

G. H. GOLUB† AND J. H. WILKINSON‡

Abstract. The solution of the complete eigenvalue problem for a nonnormal matrix A presents severe practical difficulties when A is defective or close to a defective matrix. Moreover, in the presence of rounding errors, one cannot even determine whether or not a matrix is defective. Several of the more stable methods for computing the Jordan canonical form are discussed, together with the alternative approach of computing well-defined bases (usually orthogonal) of the relevant invariant subspaces.

1. Introduction. From the standpoint of classical algebra, the algebraic eigenvalue problem has been completely solved. The problem is the subject of classical *similarity* theory, and the fundamental result is embodied in the Jordan canonical form (J.c.f.). Most mathematicians encounter similarity theory in an abstract setting, but since we are concerned here with practical algorithms, we first review the basic result purely in matrix terms.

The J.c.f. is described with reference to matrices known as *elementary Jordan* blocks. A Jordan block of order r associated with an eigenvalue λ_i will be denoted by $J_k(\lambda_i)$, and its general form is adequately illustrated by the definition

$$J_4(\lambda_i) = \begin{bmatrix} \lambda_i & 1 & 0 & 0 \\ 0 & \lambda_i & 1 & 0 \\ 0 & 0 & \lambda_i & 1 \\ 0 & 0 & 0 & \lambda_i \end{bmatrix} \tag{1.1}$$

The basic theorem is that given any $n \times n$ matrix with complex elements, there exists a nonsingular matrix X such that

$$X^{-1}AX = J, \qquad AX = XJ, \tag{1.2}$$

where J, the J.c.f. of A, is block diagonal, each diagonal matrix being an elementary Jordan block. Apart from the ordering of the blocks along the diagonal of J (which can be arbitrary), the J.c.f. is unique, although X is far from unique. It will be convenient to order the blocks in some standard way. Unless reference is made to the contrary, we assume that the $|\lambda_i|$ are in order of nonincreasing magnitude and that the blocks associated with a specific λ_i are ordered to be of nondecreasing size. Thus if the matrix A of order 12 has only 2 distinct eigenvalues λ_1 and λ_2 with $|\lambda_1| \geqq |\lambda_2|$, and λ_1 is associated with 2 blocks of order 2 and one of order 3 while

* Received by the editors January 17, 1975, and in revised form July 8, 1975. This invited paper was prepared in part and published under Contract DA HC 19-69-C-0022 with the U.S. Army Reserve Office.

† Computer Science Department, Stanford University, Stanford, California 94305. This work was supported in part by the National Science Foundation under Grant GJ35135X and by the Atomic Energy Commission under Grant AT(04-3)-326PA #30.

‡ Division of Numerical Analysis and Computing, Department of Industry, National Physical Laboratory, Teddington, Middlesex, England.

λ_2 is associated with one block of order 2 and one of order 3, its J.c.f. will be presented in the form

$$\begin{bmatrix} J_2(\lambda_1) & & & & \\ & J_2(\lambda_1) & & & \\ & & J_3(\lambda_1) & & \\ & & & J_2(\lambda_2) & \\ & & & & J_3(\lambda_2) \end{bmatrix} \tag{1.3}$$

Here λ_1 is an eigenvalue of multiplicity $2 + 2 + 3 = 7$ and λ_2 of multiplicity $2 + 3 = 5$. The example illustrates that there may be more than one block of a given dimension associated with a specific λ_i.

Let us consider the significance of the existence of a block $J_r(\lambda_i)$ in J, where $J_r(\lambda_i)$ starts in rows and columns s and ends in rows and columns t, and

$$r = t - s + 1. \tag{1.4}$$

Equating columns s to t on both sides of equation (1.2), we have

$$\begin{aligned} Ax_s &= \lambda_i x_s, & (A - \lambda_i I)x_s &= 0, \\ Ax_{s+1} &= \lambda_i x_{s+1} + x_s, & (A - \lambda_i I)x_{s+1} &= x_s, \\ Ax_{s+2} &= \lambda_i x_{s+2} + x_{s+1}, & (A - \lambda_i I)x_{s+2} &= x_{s+1}, \\ &\cdots\cdots\cdots\cdots & &\cdots\cdots\cdots\cdots \\ Ax_t &= \lambda_i x_t + x_{t-1}, & (A - \lambda_i I)x_t &= x_{t-1}, \end{aligned} \tag{1.5}$$

where, here and later, we shall denote the ith column of a matrix X (say) by x_i. The first of these relations implies that x_s is an eigenvector corresponding to λ_i. The remaining equations imply that

$$\begin{aligned} &(A - \lambda_i I)^2 x_{s+1} = 0, \ (A - \lambda_i I)^3 x_{s+2} = 0, \\ &\qquad \cdots, (A - \lambda_i I)^{t-s+1} x_t \equiv (A - \lambda_i I)^r x_t = 0. \end{aligned} \tag{1.6}$$

Notice that in general the x_{s+i} satisfy the relations

$$(A - \lambda_i I)^{p-1} x_{s+p-1} = x_s \neq 0 \quad \text{and} \quad (A - \lambda_i I)^p x_{s+p-1} = 0. \tag{1.7}$$

We shall refer to *any vector* x such that $(A - \lambda I)^{p-1} x \neq 0$, $(A - \lambda I)^p x = 0$, as a *vector of grade p*, and for uniformity, an eigenvector becomes a vector of grade 1. It is evident, for example, that

$$\begin{aligned} &(A - \lambda I)^2(\alpha_2 x_{s+2} + \alpha_1 x_{s+1} + \alpha_0 x_s) = \alpha_2 x_s, \\ &(A - \lambda I)^3(\alpha_2 x_{s+2} + \alpha_1 x_{s+1} + \alpha_0 x_s) = 0, \end{aligned} \tag{1.8}$$

so that $\alpha_2 x_{s+2} + \alpha_1 x_{s+1} + \alpha_0 x_s$ is a vector of grade 3 for all α_i provided $\alpha_2 \neq 0$. The vectors x_{s+i} arising in the Jordan canonical reduction are special in that they satisfy the *chain* relations (1.5). We shall refer to the vectors of grades 1, 2, 3, $\cdots$ associated with a Jordan block as *principal vectors of grades* 1, 2, 3, $\cdots$.

Clearly $\det(\lambda I - J_r(\lambda_i)) = (\lambda - \lambda_i)^r$, and we may associate such a polynomial with each of the blocks in the J.c.f. These polynomials are called the *elementary*

divisors of A. An enumeration of the elementary divisors gives a unique specification of the J.c.f. Corresponding to a Jordan block of dimension unity the elementary divisor is $(\lambda - \lambda_i)$, i.e., it is linear. If all the Jordan blocks in the J.c.f. are of dimension unity, then the J.c.f. is *strictly* diagonal, the matrix has n independent eigenvectors given by the columns of X and all the elementary divisors are linear. These four properties are fully equivalent to each other. Notice that if there are n distinct λ_i, then all the blocks are necessarily of dimension unity. Departure from strict diagonal form can occur only if there is at least one multiple eigenvalue, though even in this case the J.c.f. can be diagonal.

A matrix is said to be *defective* if the J.c.f. is not strictly diagonal. In this case, at least one elementary divisor is nonlinear and the number of independent eigenvectors is less than n; the remaining columns of X are principal vectors of the appropriate grades.

A matrix is said to be *derogatory* if there is at least one λ_i which is associated with more than one diagonal block in the J.c.f. If such a λ_i is associated with k different blocks, then there are precisely k independent eigenvectors associated with λ_i.

It should be emphasized that a matrix may be defective without being derogatory and vice versa, or it can be both defective and derogatory. If the λ_i are distinct, it cannot be either. If A is normal (including Hermitian, skew Hermitian or unitary), then its J.c.f. is always strictly diagonal, and the X producing the J.c.f. may be chosen to be unitary. A normal matrix with a multiple eigenvalue is therefore derogatory but not defective.

We do not report on numerical experiments in this paper, although many of the algorithms described have been implemented with success. It is the aim of this paper to emphasize the problems associated with computing invariant subspaces and to stimulate research in this area. We have not attempted to be encyclopedic (despite the length of the paper) but state those principles which we feel are of importance in this area.

2. Linear differential equations and the J.c.f. The practical significance of the J.c.f. of a matrix A is that it provides the general solution of the associated system of linear differential equations with constant coefficients defined by

$$\frac{du}{dt} = Au, \tag{2.1}$$

where u is a vector of order n. Under the linear transformation $u = Xv$, the equation becomes

$$X\frac{dv}{dt} = AXv \quad \text{or} \quad \frac{dv}{dt} = X^{-1}AXv = Jv. \tag{2.2}$$

Hence the J.c.f. gives a simplified version of the original system. If J is strictly diagonal (i.e., A is not defective), the transformed system is

$$\frac{dv_i}{dt} = \lambda_i v_i, \tag{2.3}$$

and in terms of variables v_i, the equations are completely decoupled. The general

solution is

$$v_i = v_i^{(0)} e^{\lambda_i t}, \qquad u = \sum v_i^{(0)} x_i e^{\lambda_i t}, \tag{2.4}$$

and is therefore directly expressible in terms of the n independent eigenvectors x_i and n independent constants $v_i^{(0)}$, the initial values of the v_i. Notice that the analysis is not affected by any multiplicities in the λ_i provided J is strictly diagonal. An eigenvalue λ_i of multiplicity r is then associated with r independent eigenvectors and r arbitrary $v_j^{(0)}$. When A is defective, the linear transformation does not give a complete decoupling of the equations, but there is a decoupling of those equations involving the v_i associated with each specific block from those associated with all other v_j. The general solution is most readily exposed in terms of the concept of the "exponential" of a matrix. We define $\exp(B)$ by the relation

$$\exp(B) = I + \frac{1}{1!}B + \frac{1}{2!}B^2 + \cdots + \frac{1}{r!}B^r + \cdots, \tag{2.5}$$

the matrix series being convergent for all B. The solution of (2.1) such that $u = u^{(0)}$ when $t = 0$ is given by

$$u = \exp(At)u^{(0)}. \tag{2.6}$$

From the series expansion it will readily be verified that

$$\exp(XBX^{-1}t) = X \exp(Bt)X^{-1}, \tag{2.7}$$

and hence the solution of (2.1) is

$$u = X \exp(Jt)X^{-1}u^{(0)}$$

or

$$v = \exp(Jt)v^{(0)}, \quad \text{where } v = X^{-1}u. \tag{2.8}$$

If $J_r(\lambda_i)$ is a typical block in J, then $\exp(Jt)$ has the same block structure, with $\exp(J_r(\lambda_i)t)$ in place of each $J_r(\lambda_i)$, and the form of $\exp(J_r(\lambda_i)t)$ is fully illustrated by the relation

$$\exp(J_4(\lambda_i)t) = \exp(\lambda_i t)\begin{bmatrix} 1 & t/1! & t^2/2! & t^3/3! \\ & 1 & t/1! & t^2/2! \\ & & 1 & t/1! \\ & & & 1 \end{bmatrix} \tag{2.9}$$

Hence on transforming back from the v-coordinates to the u-coordinates, the solution corresponding to the initial problem is again given in terms of the vectors x_i but corresponding to a Jordan block $J_r(\lambda_i)$, terms involving $\exp(\lambda_i t)t^s/s!$ $(s = 0, \cdots, r-1)$ arise.

This discussion gives the impression that the theoretical significance of the J.c.f. is fully matched by its practical importance since it is precisely because of its relationship to the solution of systems of linear differential equations that the

algebraic eigenvalue problem occupies such a prominent position in practical applied mathematics. The principal objective of the remainder of this paper is to show the basic limitations of the J.c.f. from the point of view of practical computation and, indeed, to cast doubt on the advisability of trying to determine it.

Before proceeding, it is useful to consider the degree of arbitrariness in the matrix X involved in the reduction to J.c.f. If the λ_i are distinct, J is diagonal and the x_i are the unique eigenvectors. The only degree of arbitrariness is in the scaling of the x_i. We have

$$D^{-1}X^{-1}AXD = D^{-1}JD = J, \tag{2.10}$$

where D is a nonsingular diagonal matrix.

Turning now to the case where J has a single block of dimension r, we see that there is already a wide freedom of choice in X. Suppose, for illustration, that there is a block of order 4 associated with λ_i; then from (1.5) we see, writing $B \equiv A - \lambda_i I$, that

$$\begin{aligned} B(ax_{s+3} + bx_{s+2} + cx_{s+1} + dx_s) &= ax_{s+2} + bx_{s+1} + cx_s, \\ B(ax_{s+2} + bx_{s+1} + cx_s) &= ax_{s+1} + bx_s, \\ B(ax_{s+1} + bx_s) &= ax_s, \\ B(ax_s) &= 0, \end{aligned} \tag{2.11}$$

where the a, b, c, d are arbitrary, but $a \neq 0$. Hence the chain of vectors x_{s+3}, x_{s+2}, x_{s+1}, x_s may be replaced by the chain of vectors given in (2.11) and on this account X may be replaced by XP, where

$$P = \left[\begin{array}{c|cccc|c} I & & & & & \\ \hline & a & b & c & d & \\ & & a & b & c & \\ & & & a & b & \\ & & & & a & \\ \hline & & & & & I \end{array}\right] \tag{2.12}$$

The derogatory case, i.e., the case when there is more than one block associated with a given λ_i, may be illustrated by the case when there are blocks of orders 2 and 3 starting in positions s and t, respectively. From the two chains

$$\begin{aligned} Bx_s &= 0, & Bx_t &= 0, \\ Bx_{s+1} &= x_s, & & \\ Bx_{s+2} &= x_{s+1}, & & \end{aligned} \tag{2.13}$$

the two generalized chains defined by

$$\begin{aligned} &B(ax_{s+2} + bx_{s+1} + cx_s + dx_{t+1} + ex_t) = ax_{s+1} + bx_s + dx_t, \\ &B(ax_{s+1} + bx_s + dx_t) = ax_s, \\ &B(ax_s) = 0, \end{aligned} \tag{2.14a}$$

and

$$\begin{aligned} &B(fx_{s+1} + gx_s + hx_{t+1} + ix_t) = fx_s + hx_t, \\ &B(fx_s + hx_t) = 0, \end{aligned} \tag{2.14b}$$

may be derived, where the $a, b, \cdots, i$ are arbitrary, except that $a \neq 0$, $h \neq 0$, and X may be varied correspondingly.

3. Sensitivity of the eigenvalues of a defective matrix. Blocks of dimension greater than unity in the J.c.f. can emerge, if at all, only as the result of the presence of multiple eigenvalues. In the classical theory there is a clear-cut distinction between equal and unequal eigenvalues. In practice, the situation is very different since a matrix may not be representable exactly in the computer and, in any case, rounding errors are, in general, involved in computing transformations. Let us consider the effect of small perturbations on the eigenvalues of an elementary Jordan block $J_r(\lambda_i)$. If the zero element in position $(r, 1)$ is replaced by ε, the characteristic equation

$$(\lambda - \lambda_i)^r = \varepsilon \tag{3.1}$$

and the multiple eigenvalue λ_i is replaced by r distinct eigenvalues $\lambda_i + \varepsilon^{1/r}(\cos(2s\pi/r) + i\sin(2s\pi/r))$ $(s = 0, \cdots, r-1)$. Suppose λ_i is of order unity, $r = 10$ and $\varepsilon = 10^{-10}$. Then the separation of the perturbed roots is of order 10^{-1} and they cannot in any reasonable sense be regarded as "close".

In practice, we have to diagnose multiplicities and the degree of defectiveness from computed eigenvalues. When these are determined by a very stable algorithm, we cannot rely on any of them being recognizably "close", even when the given A really does have some multiple eigenvalues. When A has an elementary divisor of high degree, this danger appears to be particularly severe.

However, even this remark somewhat oversimplifies the situation. One tends to be seduced by the simplicity of the J.c.f. and as a result to attach too much significance to every detail of it. When attempting to construct "difficult" matrices for practical experiments, it is common to take a nondiagonal J.c.f., subject it to some exact similarity transformation and then to regard the resulting matrix as wholly typical of a defective matrix.

But this is to attach too much significance to the unity elements in the Jordan blocks. If $D = \operatorname{diag}(d_i)$ is any nonsingular diagonal matrix, then from (1.2) we have

$$D^{-1}X^{-1}AXD = D^{-1}JD. \tag{3.2}$$

Hence if J has a unity element in position $(p, p+1)$, the matrix $D^{-1}JD$ has d_p^{-1}, d_{p+1} in this position; by a suitable choice of the d_i the unity elements may be given arbitrary values. The choice of the unity elements in the J.c.f. is purely for notational convenience. However, in classical mathematics we can make a sharp

distinction between zero and nonzero elements, a luxury we are denied in practical computation. We refer to a matrix as being in *quasi-J.c.f.* if the only difference from strict J.c.f. is that some of the super-diagonals have values other than unity.

It is *possible* for a matrix A to be highly defective without its eigenvalues being unduly sensitive. Suppose, for example, that A is such that there is an orthogonal matrix X for which

$$X^{-1}AX = \tilde{J}, \tag{3.3}$$

where $\tilde{J}$ is of quasi-J.c.f. in which nonzero super-diagonal elements are all 10^{-10}. Perturbations of order 10^{-10} in J (which correspond to perturbations of order 10^{-10} in A since X is orthogonal) produce perturbations of order 10^{-10} at most in the eigenvalues. If $\|A_2\|$ is of the order of unity, then *from the point of view of 10-digit decimal computation*, the eigenvalues of A are not at all sensitive. One cannot even *rely* on defectiveness being characterized by sensitivity of the corresponding eigenvalues. Nevertheless it is true that $\partial\lambda_i/\partial\varepsilon = O(\varepsilon^{1/r-1})$ for some perturbations when J has a block of order r, and hence, $\partial\lambda_i/\partial\varepsilon \to \infty$ as $\varepsilon \to 0$. This means that if we are prepared to extend the precision of computation indefinitely, we shall ultimately gain only one figure of accuracy for r extra figures of precision.

At this stage, one might ask what is the "natural" quasi-J.c.f. for computational purposes. A reasonable definition is that it is the $\tilde{J}$ for which the corresponding $\|X\|_2\|X^{-1}\|_2 = \varkappa(X)$ is a minimum. If this $\tilde{J}$ has super-diagonal elements which are all small relative to $\|\tilde{J}\|_2$, the matrix A will not have sensitive eigenvalues.

As a final result relating small eigenvalues and small singular values, we note the following theorem (for the definition of singular values, see § 7).

THEOREM. *Let A be an $n \times n$ matrix with $\lambda_n = \varepsilon$ and $|\lambda_n| \leqq |\lambda_j|$ and such that there are p Jordan blocks of dimensions $k_1, k_2, \cdots, k_p$, with $k_1 \leqq k_2 \leqq \cdots \leqq k_p$, associated with λ_n. Then if $A = XJX^{-1}$,*

$$\sigma_{n-j+1}(A) \leqq \|X\|_2\|X^{-1}\|_2|\varepsilon|^{k_p-j+1} + O(|\varepsilon|^{k_p-j+1+2}), \qquad j = 1, 2, \cdots, p. \tag{3.4}$$

Proof.

$$\begin{aligned}\sigma_{n-j+1}(A) = \sigma_{n-j+1}(XJX^{-1}) &\leqq \sigma_1(X)\sigma_{n-j+1}(JX^{-1}) \\ &\leqq \sigma_1(X)\sigma_1(X^{-1})\sigma_{n-j+1}(J).\end{aligned} \tag{3.5}$$

Since the singular values of J are given by $[\lambda_i(JJ^T)]^{1/2}$, it is obvious that they are singular values of the elementary Jordan blocks. Consider the $k \times k$ block

$$K = \begin{bmatrix} \varepsilon & 1 & & & \\ & \varepsilon & 1 & & \\ & & \cdot & & \\ & & & \cdot & 1 \\ & & & & \varepsilon \end{bmatrix} \tag{3.6}$$

From the form of KK^T, $k - 1$ of the singular values are close to unity and since

their product is ε^k, the remaining singular value is $O(\varepsilon^k)$. In fact,

$$\sigma_k(K) = \min_{x \neq 0} \frac{\|Kx\|_2}{\|x\|_2} \tag{3.7}$$

and taking $\tilde{x}^T = (1, -\varepsilon, \varepsilon^2, \cdots, (-1)^{k-1}\varepsilon^{k-1})$, we have

$$\sigma_k(K) = |\varepsilon|^k + O(|\varepsilon|^{k+2}). \tag{3.8}$$

The result is thus established. Note that although we have shown that the singular values are small, we have not shown and cannot show that the elements of the corresponding singular vectors are correspondingly small.

4. Ill-conditioned eigenvalues. Since in practice it will usually be impossible to determine whether a matrix has exactly equal eigenvalues, it is necessary to consider the problem of the sensitivity of a *simple* eigenvalue with respect to perturbations in A. If J is the J.c.f., we have

$$AX = XJ, \quad ZA = JZ, \quad Z = X^{-1} \tag{4.1}$$

When λ_1 is a simple eigenvalue, x_1 is the corresponding right-hand eigenvector and

$$Ax_1 = \lambda_1 x_1. \tag{4.2}$$

If z_1^T is the first row of Z, then

$$z_1^T A = z_1^T \lambda_1. \tag{4.3}$$

It is customary to define the left-hand eigenvector y_1 of A corresponding to λ_1 as the vector satisfying

$$y_1^H A = y_1^H \lambda_1, \tag{4.4}$$

and hence if we write $Y = Z^H$, the first column of Y gives this eigenvector and

$$Y^H X = I. \tag{4.5}$$

Consider now the corresponding eigenvalue $\lambda_1(\varepsilon)$ and right-hand eigenvector $x_1(\varepsilon)$ of $A + \varepsilon B$, where $\|B\|_2 = 1$. For sufficiently small ε, it is easy to show that $\lambda_1(\varepsilon)$ and $x_1(\varepsilon)$ may be expanded as convergent power series

$$\lambda_1(\varepsilon) = \lambda_1 + p_1\varepsilon + p_2\varepsilon^2 + \cdots, \qquad x_1(\varepsilon) = x_1 + v_1\varepsilon + v_2\varepsilon^2 + \cdots, \tag{4.6}$$

where the v_i lie in the space spanned by $x_2, \cdots, x_n$. (Note that in general these x_i will include principal vectors which are not eigenvectors.) Equating coefficients of ε in the relation

$$(A + \varepsilon B)(x_1 + v_1\varepsilon + \cdots) = (\lambda_1 + p_1\varepsilon + \cdots)(x_1 + v_1\varepsilon + \cdots) \tag{4.7}$$

gives

$$Bx_1 + Av_1 = \lambda_1 v_1 + p_1 x_1. \tag{4.8}$$

Now both v_1 and Av_1 lie in the space spanned by $x_2, \cdots, x_n$, and from (4.5),

$y_1^H x_i = 0$ $(i = 2, \cdots, n)$. Hence premultiplying (4.8) by y_1^H, we obtain

$$p_1 = y_1^H B x_1 / y_1^H x_1. \tag{4.9}$$

As derived above, $y_1^H x_1 = 1$, but clearly in (4.9), x_1 and y_1 may be arbitrarily scaled and it is convenient computationally to have $\|x_1\|_2 = \|y_1\|_2 = 1$. In this case, $y_1^H x_1 = s_1$ (in the notation of [25]), where s_1 is the cosine of the angle between x_1 and y_1. From (4.9),

$$\left|\frac{\partial \lambda_1}{\partial \varepsilon}\right|_{\varepsilon=0} = |p_1| \leqq \frac{\|y_1\|_2 \|B\|_2 \|x_1\|_2}{s_1} = \frac{1}{|s_1|}. \tag{4.10}$$

The derivative is finite for any "direction" of B. This is in contrast to the case where λ_i is associated with a defective matrix when $|\partial \lambda_i / \partial \varepsilon|_{\varepsilon=0} = \infty$. This latter result is in agreement with (4.10) since the left-hand and right-hand eigenvectors are orthogonal corresponding to a "defective" λ_i. The bound in (4.10) is attained when $B = y_1 x_1^H$, since then

$$y_1^H B x_1 = y_1^H y_1 x_1^H x_1 = 1. \tag{4.11}$$

Further, taking $B = e^{i\theta} y_1 x_1^H$, we can make $(\partial \lambda_1 / \partial \varepsilon)_{\varepsilon=0}$ have any required phase. There is one very unsatisfactory feature of the above analysis. The quantity s_i is not invariant with respect to diagonal similarity transformation. Consider the matrix

$$A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}, \tag{4.12}$$

with

$$\lambda_1 = 3, \quad \lambda_2 = 1, \quad y_1^H = \frac{[1, 1]}{2^{1/2}}, \quad x_1^H = \frac{[1, 1]}{2^{1/2}}, \quad s_1 = 1. \tag{4.13}$$

The eigenvalue λ_1 is therefore very well-conditioned, as indeed are all eigenvalues of all normal matrices. However, we have

$$D^{-1} A D = \begin{bmatrix} 2 & \alpha \\ \alpha^{-1} & 2 \end{bmatrix}, \quad \text{where } D = \begin{bmatrix} 1 & \\ & \alpha \end{bmatrix}, \tag{4.14}$$

and now

$$y_1^H = \frac{[1, \alpha]}{(1 + \alpha^2)^{1/2}}, \quad x_1^H = \frac{[\alpha, 1]}{(1 + \alpha^2)^{1/2}}, \quad s_1 = \frac{2\alpha}{(1 + \alpha^2)}. \tag{4.15}$$

Hence we may make s_1 arbitrarily small by taking α sufficiently large or sufficiently small. It is clear that a small s_i induced in this way is a very artificial phenomenon. In this example, when s_1 is small, $\|D^{-1} A D\|_2 \gg \|A\|_2$. In practice, the relevant values of s_i are those for $D^{-1} A D$, where D has been chosen so that $\|D^{-1} A D\|_2$ is a minimum. Reducing this norm to a true minimum is not vital, and in practice, the process of *balancing* described by Parlett and Reinsch in [12] is usually adequate.

High sensitivity of an eigenvalue λ_i has now been encountered in two different contexts, first when λ_i is associated with defectiveness and secondly when a value of s_i is small. We now show that when an s_i is small, A is necessarily relatively close to a matrix with a multiple eigenvalue. Let

$$Ax_1 = \lambda_1 x_1, \quad y_1^H A = y_1^H \lambda_1, \quad s_1 = y_1^H x, \quad \text{with } \|x_1\|_2 = \|y_1\|_2 = 1, \tag{4.16}$$

and suppose P is a unitary matrix such that $Px_1 = e_1$, where $e_1^T = (1, 0, \cdots, 0)$. Then

$$PAP^H Px_1 = \lambda_1 Px_1, \qquad (PAP^H)e_1 = \lambda_1 e_1, \tag{4.17}$$

and $B = PAP^H$ must be of the form

$$B = \left[\begin{array}{c|c} \lambda_1 & b_1^H \\ \hline 0 & B_1 \end{array}\right]. \tag{4.18}$$

Further,

$$s_1 = y_1^H x_1 = (y_1^H P^H)(Px_1) = (Py_1)^H e_1, \tag{4.19}$$

and writing $Py_1 = p_1$, we have

$$p_1^H B = p_1^H PAP^H = y_1^H AP^H = \lambda_1 (y_1^H P^H) = \lambda_1 p_1^H, \tag{4.20}$$

while

$$s_1 = p_1^H e_1 = \bar{p}_{11}. \tag{4.21}$$

Hence if we write $p_1^H = (\bar{p}_{11} | v^H)$, where v is of order $n - 1$,

$$\bar{p}_{11} b_1^H + v^H B_1 = \lambda_1 v^H, \qquad v^H \left(B_1 - \lambda_1 I + \bar{p}_{11} \frac{v b_1^H}{v^H v} \right) = 0, \tag{4.22}$$

i.e., the matrix $B_1 + \bar{p}_{11}(vb_1^H/v^H v)$ has λ_1 as an eigenvalue and v as a left-hand eigenvector. Now

$$\left\| \bar{p}_{11} \left(\frac{v b_1^H}{v^H v} \right) \right\| \leqq |\bar{p}_{11}| \frac{\|v\| \, \|b_1\|}{\|v\|^2} = \frac{|s_1| \, \|b_1\|}{(1 - s_1^2)^{1/2}} \leqq \frac{s_1 \|B\|_2}{(1 - s_1^2)^{1/2}} = \frac{s_1 \|A\|_2}{(1 - s_1^2)^{1/2}}, \tag{4.23}$$

and when s_1 is small, a small relative perturbation in B converts λ_1 into an eigenvalue of multiplicity at least two. Since the l_2-norm is invariant with respect to unitary transformations, the same remark is true of A. By a similar argument, Kahan in an unpublished paper has shown that the denominator $(1 - s^2)^{1/2}$ may be replaced by 1 in the final bound. However, the above argument shows that the relevant bound is $|s_1| \, \|b_1\|_2/(1 - s_1^2)^{1/2}$ and in replacing $\|b_1\|_2$ by $\|B\|_2$ and hence by $\|A\|_2$, the result is weakened. When A is normal, B is also normal and $b_1 = 0$. Hence if $|s_1| < 1$ for a normal matrix, λ_1 must already be a multiple eigenvalue. This is otherwise obvious, since if λ_1 is a simple eigenvalue of a normal matrix, $y_1 = x_1$ and $s_1 = 1$. The bound we have given is, in general, a considerable improvement on the bound given by Ruhe [16].

5. Almost linearly dependent eigenvectors. The perturbation analysis described above can be used to give the first order perturbation of x_1 resolved in the directions

$x_2, \cdots, x_n$. In the case when A is nondefective, this leads to

$$x_1(\varepsilon) = x_1 + \varepsilon\left\{\sum_{i=2}^{n}\left(\frac{y_i^H B x_1}{s_i(\lambda_i - \lambda_1)}\right)x_i\right\} + O(\varepsilon^2), \tag{5.1}$$

and the coefficient of x_i is bounded by $1/|s_i(\lambda_i - \lambda_1)|$. Hence we obtain a large perturbation *in the direction of* x_i if s_i or $\lambda_i - \lambda_1$ is small. However, this analysis is rather unsatisfactory. When A has an ill-conditioned eigenvalue problem, the set of x_i will be almost linearly dependent, as we show below. The fact that some of the x_i have large coefficients need not necessarily mean that the perturbation as a whole is large.

The left-hand eigenvector y_1 is orthogonal to $x_2, \cdots, x_n$, and hence x_1 may be expanded in terms of $y_1, x_2, \cdots, x_n$. In fact,

$$x_1 = s_1 y_1 + \sum_{i=2}^{n} \alpha_i x_i \tag{5.2}$$

since $y_1^H x_1 = s_1$ and $y_1^H x_i = 0$ $(i = 2, \cdots, n)$. Equation (5.2) may be expressed in the form

$$\sum_{i=1}^{n} \beta_i x_i = s_1 y_1/(1 + \textstyle\sum \alpha_i^2)^{1/2} \tag{5.3}$$

where

$$\beta_1 = 1/(1 + \textstyle\sum \alpha_i^2)^{1/2}, \quad \beta_i = -\alpha_i/(1 + \textstyle\sum \alpha_i^2)^{1/2}, \quad \|\beta\|_2 = 1. \tag{5.4}$$

Hence we have a unit vector β so that

$$\|X\beta\|_2 = |s_1|/(1 + \textstyle\sum \alpha_i^2)^{1/2} < |s_1|, \tag{5.5}$$

and when s_1 is small, the vectors x_i are "almost linearly dependent". (Note that in general, the x_i $(i = 2, \cdots, n)$ will include principal vectors which are not eigenvectors.) Anticipating § 7, we note that (5.5) implies that $\sigma_n(X) < |s_1|$. Conversely, if a set of the x_i are almost linearly dependent, then at least one of the associated s_i is small and A has an ill-conditioned eigenvalue. Suppose, for example,

$$\sum_{i=1}^{p} \alpha_i x_i = u, \quad \text{where } \|u\|_2 = \varepsilon, \quad \sum_{i=1}^{p} \alpha_i^2 = 1. \tag{5.6}$$

Then if the vectors y_i are the normalized columns of $(X^{-1})^H$, we have

$$\alpha_i y_i^H x_i = y_i^H u, \qquad s_i = y_i^H u/\alpha_i, \quad |s_i| \leqq \varepsilon/|\alpha_i|. \tag{5.7}$$

Since at least one α_i is such that $|\alpha_i| > p^{-1/2}$, this means that at least one s_i is small. In fact, it is obvious that at least two of the s_i must be small, since otherwise just one of the eigenvalues would be sensitive and the remainder insensitive; as the trace is obviously not sensitive, this is impossible.

This result emphasizes one very unsatisfactory feature of ill-conditioned eigensystems. Suppose we have managed (in spite of the practical difficulties) to obtain *correctly rounded* versions of a set of ill-conditioned eigenvectors $x_1, \cdots, x_p$. We may now wish to determine an accurate orthogonal basis for this subspace of dimension p. However, since the vectors $x_1, \cdots, x_p$ are almost linearly dependent,

when we perform the Schmidt orthogonalization process on these x_i, the orthogonal basis is bound to be poorly-determined. In fact, information about the last of the orthogonal vectors will be completely vitiated by the rounding errors which will usually be inherent in the representation of the x_i in the computer.

This casts doubt on the advisability of attempting to determine the x_i themselves and suggests that it might be better to determine directly an orthogonal basis for the subspace corresponding to such vectors.

6. Orthogonal bases for invariant subspaces. The eigenvectors of A corresponding to λ are the solutions of the equation $(A - \lambda I)x = 0$. If $A - \lambda I$ is of nullity n_1 (rank $= n - n_1$), then there will be n_1 independent eigenvectors. These vectors span a subspace P_1, the nullspace of $A - \lambda I$. Let $x_1^{(1)}, x_2^{(1)}, \cdots, x_{n_1}^{(1)}$ be an orthogonal basis of this subspace P_1.

Turning now to the solutions of the equation $(A - \lambda I)^2 x$, we can clearly see that they include any vector in P_1, since if $(A - \lambda I)x = 0$, then certainly $(A - \lambda I)^2 x = 0$. The nullity of $(A - \lambda I)^2$ may therefore be denoted by $n_1 + n_2$, where $n_2 \geqq 0$. If the nullspace is denoted by P_2, then $P_2 \supset P_1$ and the basis $x_1^{(1)}, x_2^{(1)}, \cdots, x_{n_1}^{(1)}$ may be extended to an orthogonal basis of P_2 by the addition of further orthogonal vectors $x_1^{(2)}, x_2^{(2)}, \cdots, x_{n_2}^{(2)}$. These additional vectors satisfy the relations

$$u_i = (A - \lambda I)x_i^{(2)} \neq 0, \quad (A - \lambda I)^2 x_i^{(2)} = 0, \quad i = 1, \cdots, n_2, \tag{6.1}$$

and hence they are vectors of grade 2.

We now show that $n_2 \leqq n_1$, for the vectors u_i are nonnull and satisfy the relation $(A - \lambda I)u_i = 0$. Hence they lie in P_1, and if $n_2 > n_1$,

$$\sum \alpha_i u_i = 0, \quad \text{i.e., } (A - \lambda I) \sum \alpha_i x_i^{(2)} = 0, \tag{6.2}$$

which means that $\sum \alpha_i x_i^{(2)} \subset P_1$. But $\sum \alpha_i x_i^{(2)}$ is orthogonal to the $x_i^{(1)}$ by the construction, and hence we have a contradiction.

Continuing in this way by considering the nullities of $(A - \lambda I)^3, (A - \lambda I)^4, \cdots$, we obtain numbers $n_3, n_4, \cdots$ such that $n_{i+1} \leqq n_i$ and orthogonal bases of subspaces P_i such that $P_{i+1} \supset P_i$. The subspace P_i is of dimension $m_i = n_1 + \cdots + n_i$. In general, the orthogonal vectors $x_i^{(s)}$ are such that $(A - \lambda I)^{s-1} x_i^{(s)} \neq 0$ but $(A - \lambda I)^s x_i^{(s)} = 0$.

The sequence comes to an end with $(A - \lambda I)^k$, where $(A - \lambda I)^{k+1}$ is of the same nullity as $(A - \lambda I)^k$.

Comparing these spaces with those spanned by the chains of vectors associated with λ in the J.c.f., we see that P_1 is the space spanned by the principal vectors of grade 1, P_2 that spanned by principal vectors of grades 1 and 2, etc. Notice, though, that the space spanned by $x_1^{(2)}, \cdots, x_{n_2}^{(2)}$ is *not* in general the same as that spanned by the principal vectors of grade 2 in the Jordan chains.

n_1 is equal to the number of blocks associated with λ in the J.c.f., and, in general, n_s is the number of those blocks which are of dimension not less than s.

The derivation of these orthogonal bases is in some ways more satisfactory than that of the Jordan chains themselves, and though the chains may be derived from the orthogonal bases, there will in general be a loss of digital information in this process.

7. The singular values. In the previous section it was shown that in the solution of the complete eigenvalue problem, we are concerned with the determination of the nullities or ranks of sequences of matrices. Rank determination is a notoriously dangerous numerical problem, and in practice the only reliable way of doing it is via the singular value decomposition (S.V.D). Accordingly we now give a brief review of the S.V.D. and the properties of singular values.

For our purposes, the singular values of a complex $m \times n$ matrix A may be defined to be the nonnegative square roots of the eigenvalues of the matrix A^HA. Clearly A^HA is an $n \times n$ nonnegative definite Hermitian matrix, and its eigenvalues may be denoted by σ_i^2 $(i = 1, \cdots, n)$; the σ_i are the singular values of A. Although apparently a more sophisticated *concept* than the eigenvalues, the *determination* of the singular values is more satisfactory from the computational point of view. The σ_i are defined in terms of the eigenvalues of a Hermitian matrix, and these are always insensitive to small perturbations in elements of that matrix. We shall assume that the σ_i are ordered so that $\sigma_1 \geqq \sigma_2 \geqq \cdots \geqq \sigma_n$. The σ_i^2 may be defined via the min-max properties of $(x^HA^HAx)/x^Hx$, i.e., of

$$\frac{\|Ax\|_2^2}{\|x\|_2^2}$$

$$\sigma_1(X) = \max \frac{\|Ax\|_2}{\|x\|_2} = \|A\|_2, \qquad \sigma_n(X) = \min \frac{\|Ax\|_2}{\|x\|_2} \tag{7.1}$$

and

$$\begin{aligned} \sigma_r(A) - \|B\|_2 &= \sigma_r(A) - \sigma_1(B) \\ &\leqq \sigma_r(A + B) \leqq \sigma_r(A) + \sigma_1(B) \leqq \sigma_r(A) + \|B\|_2. \end{aligned} \tag{7.2}$$

From the last of these relations, the well-conditioned nature of the σ_r is well exposed.

Although we have defined the σ_i via A^HA, they should not be determined in this way. In practice, they are computed via the S.V.D., which is defined as follows.

Any $m \times n$ complex matrix A may be factorized in the form

$$A = U\Sigma V^H \tag{7.3}$$

where U and V are $m \times m$ and $n \times n$ unitary matrices, respectively, and Σ is an $m \times n$ matrix with $\Sigma_{ii} = \sigma_i$ and $\Sigma_{ij} = 0$ otherwise. Golub and Reinsch [4] have described an extremely efficient and stable method for determining the S.V.D. and hence the σ_i. The computed U and V^H are almost orthogonal to the working accuracy, and the computed σ_i correspond to those of some $(A + E)$, where $\|E\|_2/\|A\|_2$ is a modest multiple of the machine precision. Since the σ_r are insensitive to E, this is very satisfactory.

Clearly, from (7.3),

$$A^HA = V\Sigma^H\Sigma V^H, \qquad A^HAV = V\Sigma^H\Sigma, \tag{7.4}$$

so that the columns of V are orthogonal eigenvectors of A^HA. Similarly

$$AA^H = U\Sigma\Sigma^HU^H, \qquad AA^HU = U\Sigma\Sigma^H, \tag{7.5}$$

and the columns of U are orthogonal eigenvectors of AA^H.

Turning now to the case when A is $n \times n$, we have, from the definitions of the eigenvalues and of the singular values,

$$\prod \lambda_i = \det(A), \quad \prod (\sigma_i^2) = \det(A^H A) = |\det(A)|^2, \tag{7.6}$$

and hence

$$\prod |\lambda_i| = \prod \sigma_i. \tag{7.7}$$

We have the fundamental result that $\lambda_n = 0$ iff $\sigma_n = 0$ and both imply that A is singular. The three properties are fully equivalent.

From this it is intuitively obvious that if A is "nearly" singular, λ_n and σ_n are "small" with appropriate determination of the terms "nearly" singular and "small". As a measure of the proximity of A to singularity we shall take $\|E\|_2/\|A\|_2 = \varepsilon$, where E is the matrix of minimum norm such that $A + E$ is singular. Since $A + E$ is singular, there exists a y such that

$$(A + E)y = 0. \tag{7.8}$$

Hence

$$\sigma_n = \min \frac{\|Ax\|}{\|x\|} \leqq \frac{\|Ay\|}{\|y\|} = \frac{\|-Ey\|}{\|y\|} \leqq \|E\|_2 = \varepsilon \|A\|_2. \tag{7.9}$$

On the other hand, since $\min(\|Ax\|/\|x\|)$ is attained for some unit vector, y (say),

$$\sigma_n = \|Ay\|, \qquad Ay = \sigma_n z \quad \text{with } \|z\|_2 = 1. \tag{7.10}$$

Hence $(A - \sigma_n z y^H)y = 0$, and $A - \sigma_n z y^H$ must be singular. But $\|\sigma_n z y^H\| = \sigma_n$ and $\varepsilon = \min(\|E\|_2/\|A\|_2) \leqq \sigma_n/\|A\|_2$; hence $\sigma_n/\|A\|_2 = \varepsilon$.

Turning now to λ_n, we have

$$Ay = \lambda_n y \quad \text{for some } \|y\|_2 = 1$$

and

$$\sigma_n = \min \frac{\|Ax\|}{\|x\|} \leqq \frac{\|Ay\|}{\|y\|} = |\lambda_n|. \tag{7.11}$$

On the other hand, from (7.7),

$$|\lambda_n|^n \leqq \sigma_n \sigma_1^{n-1}, \tag{7.12}$$

$$|\lambda_n/\sigma_1|^n \leqq \sigma_n/\sigma_1 = \sigma_n/\|A\|_2 = \varepsilon, \tag{7.13}$$

giving

$$|\lambda_n| \leqq \sigma_1 \varepsilon^{1/n}. \tag{7.14}$$

This last relation is disappointing, but unfortunately it is a best possible result, as is illustrated by the matrices K_n typified by

$$K_4 = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \varepsilon & 0 & 0 & 0 \end{bmatrix}. \tag{7.15}$$

In general, $|\lambda_i| = \varepsilon^{1/n}$ $(i = 1, \cdots, n)$, but $\sigma_1 = \cdots = \sigma_{n-1} = 1$ and $\sigma_n = \varepsilon$. All extreme examples are of this kind, since we have equality in (7.12) only if $|\lambda_i| = |\lambda_n|$ (all n) and $\sigma_1 = \sigma_2 = \cdots = \sigma_{n-1}$. In practice, then, we may well have a matrix which is singular to working accuracy and therefore has a negligible singular value but which has no eigenvalues which can be regarded as in any sense small.

The practical consequences of this theorem are very serious. The most stable algorithms for computing eigenvalues can guarantee only that each computed eigenvalue λ_i' is exact for some $A + E_i$, where $\|E_i\|_2/\|A\|_2$ is a modest multiple of the machine precision, and it is difficult to conceive how such algorithms can be improved upon, except, of course, by working to higher accuracy at least in some significant part of the computation. This means that $(A + E_i - \lambda_i' I)$ is exactly singular and hence that $A - \lambda_i' I$ is within $\|E_i\|_2$ of a singular matrix. Hence $A - \lambda_i' I$ has a *singular value* bounded by $\|E_i\|_2$, but the bound for the smallest eigenvalue of $A - \lambda_i' I$ involves $\|E_i\|_2^{1/n}$. All that we can guarantee a priori is that each computed λ_i' will have an error which involves the factor $\|E_i\|_2^{1/n}$, and this may be far from small.

For a normal matrix, $|\lambda_r| = \sigma_r$, and hence this weakness disappears. If λ_i is an eigenvalue of A, then $A + E_i$ has an eigenvalue λ_i' such that

$$|\lambda_i - \lambda_i'| \leqq \|E_i\|. \tag{7.16}$$

Unfortunately, the realization that this result is true has tended to lead to an overconfidence when dealing with real symmetric and Hermitian matrices, which are the commonest examples of normal matrices.

8. Factorizations of almost-singular matrices. If B is an exactly singular matrix and $B = XY$ is a factorization of B, then either X or Y (or both) is exactly singular. Most of the common factorizations used in practice ensure that one of the factors is certainly not singular, and hence with exactly singular B and exact factorization, the other factor must be singular.

A factorization which is frequently used is $B = QR$, where Q is unitary and R is upper triangular. Clearly Q is nonsingular, and hence if B is singular, R must also be singular and therefore have a zero eigenvalue and a zero singular value. But the eigenvalues of R are its diagonal elements and hence at least one r_{ii} must be zero, indeed r_{nn} unless B is "special".

Now consider the case when B is *almost* singular and let us assume for simplicity that B is factorized exactly. We have $\sigma_i(R) = \sigma_i(B)$ since the σ_i are invariant with respect to unitary transformations. Hence R must still have a negligible singular value. However, we can no longer guarantee that any r_{ii} is pathologically small since the r_{ii} are merely the eigenvalues, the bound for which involves $(\sigma_n(B))^{1/n}$.

This result is important in practice because many algorithms for solving the complete eigenproblem of a matrix first compute the eigenvalues and then attempt to determine the eigenvectors from them. If λ is an eigenvalue given by a stable algorithm, $(A + E - \lambda I)$ will be exactly singular with $\|E\|/\|A\|$ small, and hence $B = A - \lambda I$ will be almost singular. The situation appears particularly favorable when A is normal since the computed λ will then have an error which is small relative to $\|A\|_2$, i.e., to $|\lambda_1|$. Unfortunately, although B is normal, the same is

not true of R, and hence we still cannot guarantee that R will have any pathologically small r_{ii}. Now the weak bound for λ_n is attained only when B is extremely pathological, and hence one might expect that failure of R to have a small diagonal element would be rare. Unfortunately, this is far from true. Attempts were made to construct an algorithm based on this factorization in the case where A is a symmetric tridiagonal matrix. For such matrices, a particularly satisfactory algorithm is known for the determination of the λ's. Nevertheless, it was found in practice that when the QR factorization of $A - \lambda I$ was performed for each of the n computed λ in turn, almost invariably some of the R were such that they had no small r_{ii}, and all algorithms based on a search for a negligible r_{ii} failed disastrously.

The LL^T factorization of a positive definite matrix A is known to be extremely stable, and it might be thought that when such an A was near to singularity, this would be bound to reveal itself in the corresponding L. That this is not true is illustrated by the matrices $A = L_n L_n^T$, where L_n is of the form illustrated by

$$L_4 = \begin{bmatrix} 1 & & & \\ -1 & 1 & & \\ -1 & -1 & 1 & \\ -1 & -1 & -1 & 1 \end{bmatrix} \tag{8.1}$$

It is easy to show that $\sigma_n(A_n) = \lambda_n(A_n) = O(4^{-n})$, and hence for quite modest values of n, the matrix A_n is almost singular. Yet there is no obvious indication of this in the factor L_n since all of its diagonal elements are unity.

Finally, we consider the factorization given by Gaussian elimination with complete pivoting. This, too, would appear to be quite favorable, and yet it can fail quite catastrophically. Indeed, if A_n is of the form illustrated by

$$A_4 = \begin{bmatrix} 1 & -1 & -1 & -1 \\ & 1 & -1 & -1 \\ & & 1 & -1 \\ & & & 1 \end{bmatrix} \tag{8.2}$$

then it can be shown that $\sigma_n(A_n) = O(2^{-n})$, and hence A_n is almost singular for quite modest n. Yet the factorization given by Gaussian elimination with complete pivoting is

$$A_n = I \times A_n, \tag{8.3}$$

i.e., A_n is itself the upper triangular factor, and its diagonal elements are all unity.

These examples illustrate the fact that the determination of singularity, much less than rank, by means of simple factorizations is not a practical proposition. On the other hand, the S.V.D. is extremely reliable, and since the computed σ_i correspond to $A + E$ where $\|E\|_2/\|A\|_2$ is of the order of the machine precision, it provides an excellent means of determining the numerical rank.

9. Vectors by matrix powering. In the next three sections, we discuss some of the algorithms which have been designed to find bases for the successive nullspaces of powers of $(A - \lambda I)$ corresponding to an eigenvalue λ.

For simplicity of notation, we shall work throughout with $B = A - \lambda I$. We shall not for the moment discuss numerical stability, but knowing that most simple factorizations are numerically unreliable for finding the rank of a matrix, we shall use only the S.V.D. for this purpose. Let the S.V.D. of B be denoted by

$$B_1 \equiv B = U_1 \Sigma_1 V_1^H, \tag{9.1}$$

where U_1 and V_1 are $n \times n$ unitary matrices. Since λ is an eigenvalue, B is a singular matrix. If it is of nullity n_1, then B_1 will have n_1 zero singular values, and we may write

$$BV_1 = U_1 \Sigma_1 = [\underbrace{A_2}_{n-n_1} \mid \underbrace{0}_{n_1}]. \tag{9.2}$$

For consistency with later stages, we write $W_1 \equiv V_1$, and the last n_1 columns of W_1 clearly give an orthogonal basis for the principal vectors of grade 1, while the matrix A_2 has orthogonal columns.

Proceeding to the nullspace of B^2, we have

$$B^2 V_1 \equiv B^2 W_1 = [BA_2 \mid \underbrace{0}_{n_1}] = [B_2 \mid \underbrace{0}_{n_1}], \tag{9.3}$$

the zero columns obviously persisting. We now compute the S.V.D. of B_2:

$$B_2 = U_2 \Sigma_2 V_2^H, \tag{9.4}$$

where U_2 is an $n \times n$ unitary matrix and V_2 an $(n - n_1) \times (n - n_1)$ unitary matrix. Writing

$$\tilde{V}_2 = \begin{bmatrix} \underbrace{V_2}_{n-n_1} & \\ & \underbrace{I}_{n_1} \end{bmatrix}, \qquad W_2 = W_1 \tilde{V}_2, \tag{9.5}$$

we have

$$B^2 W_2 = [\underbrace{U_2 \Sigma_2}_{n-n_1} \mid \underbrace{0}_{n_1}]. \tag{9.6}$$

Since the nullity of B^2 is $n_1 + n_2$, B_2 will have n_2 zero singular values, and we have

$$B^2 W_2 = [A_3 \mid \underbrace{0}_{n_2} \mid \underbrace{0}_{n_1}]. \tag{9.7}$$

Writing $\sum_{i=1}^{s} n_i = m_s$, the matrix A_3 has $n - m_2$ orthogonal columns. The last m_2 columns of W_2 give an orthogonal basis for vectors of grade 2 and grade 1. The last n_1 of these columns are those of W_1 having been unaltered by this second step.

The general step is then as follows:

$$B^s W_s = [\underbrace{A_{s+1}}_{n-m_s} | \underbrace{0}_{n_s} | \cdots | \underbrace{0}_{n_1}], \tag{9.8}$$

$$B^{s+1} W_s = [B_{s+1} | 0 | \cdots | 0], \quad \text{where } B_{s+1} = BA_{s+1}, \tag{9.9}$$

$$B_{s+1} = U_{s+1}\Sigma_{s+1}V_{s+1}^H, \tag{9.10}$$

where U_{s+1} is an $n \times n$ unitary matrix and V_{s+1} an $(n - m_s) \times (n - m_s)$ unitary matrix. B_{s+1} has n_{s+1} zero singular values and writing

$$\tilde{V}_{s+1} = \begin{bmatrix} V_{s+1} & \\ & I \end{bmatrix}, \qquad W_{s+1} = W_s \tilde{V}_{s+1}, \tag{9.11}$$

$$B_{s+1}W_{s+1} = [U_{s+1}\Sigma_{s+1} | 0 | \cdots | 0] \tag{9.12}$$

$$= [A_{s+2} | \underbrace{0}_{n_{s+1}} | \cdots | \underbrace{0}_{n_1}]. \tag{9.13}$$

The process terminates when A_{s+1} is of full rank.

The main weakness of this algorithm is the difficulty of recognizing which of the elements of σ_i may be treated as zero. This is well illustrated when A and therefore B is normal. If such a matrix were inserted into this algorithm, then at the first step, the singular values would be $|\lambda_1|, |\lambda_2| \cdots |\lambda_n|$, of which n_1 would be treated as zero. For a normal matrix, the process should terminate here since all vectors are of grade 1. However, if one continues, the singular values in the second step would be $|\lambda_1|^2, |\lambda_2|^2, \cdots, |\lambda_n|^2$, and some of these might well be regarded as negligible. The algorithm can be modified to limit this shortening, but even then it compares unfavorably in most respects with the algorithm of the next section.

10. Vectors by orthogonal deflation. Again it is convenient to work with B, and we assume that it has an eigenvalue of multiplicity k. We write $B^{(1)} = B$ and denote the S.V.D. of $B^{(1)}$ by

$$B^{(1)} = U^{(1)}\Sigma^{(1)}(V^{(1)})^H, \tag{10.1}$$

where there will be n_1 zero singular values. Hence

$$B^{(2)} = (V^{(1)})^H B^{(1)} V^{(1)} = (V^{(1)})^H U^{(1)}\Sigma^{(1)} = W^{(1)}\Sigma^{(1)}, \tag{10.2}$$

and we may write

$$B^{(2)} = \left[\begin{array}{c|c} B_{11}^{(2)} & 0 \\ \hline B_{21}^{(2)} & 0 \end{array}\right] \begin{matrix} \}n - n_1 \\ \}n_1 \end{matrix} \tag{10.3}$$

From the orthogonality of $W^{(1)}$, the first $n - n_1$ columns of $B^{(2)}$ are orthogonal and therefore independent. Relation (10.1) shows that the last n_1 columns of $V^{(1)}$ give n_1 orthogonal eigenvectors (i.e., vectors of grade 1) of $B^{(1)}$ corresponding to $\lambda = 0$.

If $n_1 = k$, then we have dealt with all the eigenvalues. Otherwise $B_{11}^{(2)}$ will have $k - n_1$ zero eigenvalues and we can proceed to the consideration of vectors of grade 2. Let z be an arbitrary nonnull vector partitioned conformally with $B^{(2)}$

so that $z^T = [x^T \mid y^T]$. Then

$$B^{(2)}z = \begin{bmatrix} B_{11}^{(2)} \\ \hline B_{21}^{(2)} \end{bmatrix} x, \tag{10.4}$$

and when $x = 0$ and $y \neq 0$, z is a vector of grade 1. If $x \neq 0$, then it follows from the independence of the first $n - n_1$ columns of $B^{(2)}$ that $B^{(2)}z \neq 0$. However, we have

$$(B^{(2)})^2 z = \begin{bmatrix} B_{11}^{(2)} \\ \hline B_{21}^{(2)} \end{bmatrix} B_{11}^{(2)} x, \tag{10.5}$$

and from the same linear independence, z is a vector of grade 2 iff $B_{11}^{(2)}x = 0$.

Hence we may proceed as follows. Let the S.V.D. of $B_{11}^{(2)}$ be given by

$$B_{11}^{(2)} = U^{(2)}\Sigma^{(2)}(V^{(2)})^H, \tag{10.6}$$

where $\Sigma^{(2)}$ has n_2 zero diagonal elements if $B_{11}^{(2)}$ is of nullity n_2. Hence

$$(V^{(2)})^H B_{11}^{(2)} V^{(2)} = (V^{(2)})^H U^{(2)}\Sigma^{(2)} = W^{(2)}\Sigma^{(2)}, \tag{10.7}$$

and we may write

$$(V^{(2)})^H B_{11}^{(2)} V^{(2)} = \left[\begin{array}{c|c} B_{13}^{(3)} & 0 \\ \hline B_{21}^{(3)} & 0 \end{array}\right] \begin{matrix} \}n - m_2 \\ \}n_2 \end{matrix} \tag{10.8}$$

Again the first $n - m_2$ columns of $(V^{(2)})^H B_{11}^{(2)} V^{(2)}$ are orthogonal and hence independent. Introducing the unitary matrix

$$\tilde{V}^{(2)} = \left[\begin{array}{c|c} V^{(2)} & 0 \\ \hline 0 & I \end{array}\right], \tag{10.9}$$

$$B^{(3)} = (\tilde{V}^{(2)})^H B^{(2)} \tilde{V}^{(2)} = \begin{bmatrix} B_{11}^{(3)} & 0 & 0 \\ B_{21}^{(3)} & 0 & 0 \\ \underbrace{B_{31}^{(3)}}_{n-m_2} & \underbrace{B_{32}^{(3)}}_{n_2} & \underbrace{0}_{n_1} \end{bmatrix} \tag{10.10}$$

It is obvious that $n_2 < n_1$; otherwise $B^{(3)}$ and hence $B^{(1)}$ would have been of nullity greater than n_1.

Again if $m_2 = k$, the process is complete. Otherwise $B_{11}^{(3)}$ has some zero eigenvalues, and we proceed via its S.V.D., this next stage being typical. If

$$B^{(3)} = U^{(3)}\Sigma^{(3)}(V^{(3)})^H, \tag{10.11}$$

then

$$(V^{(3)})^H B_{11}^{(3)} V^{(3)} = (V^{(3)})^H U^{(3)}\Sigma^{(3)} = W^{(3)}\Sigma^{(3)} \tag{10.12}$$

and again introducing

$$\tilde{V}^{(3)} = \left[\begin{array}{c|c} V^{(3)} & 0 \\ \hline 0 & I \end{array}\right], \tag{10.13}$$

we are led to

$$B^{(4)} = (\tilde{V}^{(3)})^H B^{(3)} \tilde{V}^{(3)} = \begin{bmatrix} B_{11}^{(4)} & 0 & 0 & 0 \\ B_{21}^{(4)} & 0 & 0 & 0 \\ B_{31}^{(4)} & B_{32}^{(4)} & 0 & 0 \\ \underbrace{B_{41}^{(4)}}_{n-m_3} & \underbrace{B_{42}^{(4)}}_{n_3} & \underbrace{B_{43}^{(4)}}_{n_2} & \underbrace{0}_{n_1} \end{bmatrix}, \tag{10.14}$$

where n_3 is the nullity of $B_{11}^{(3)}$. By an argument similar to that used above, the nonnull columns of $B^{(4)}$ and of leading principal submatrices of orders $n - m_1$, $n - m_2$ are linearly independent. The process clearly terminates when $m_s = k$, at which stage $B_{11}^{(s+1)}$ is no longer singular. Since

$$B^{(s+1)} = V^H B^{(1)} V \equiv V^H B V, \tag{10.15}$$

where $V = V^{(1)}\tilde{V}^{(2)}\tilde{V}^{(s)} \cdots \tilde{V}^{(s)}$, the principal vectors of $B^{(1)}$ may be found via those of $B^{(s+1)}$. For simplicity of notation, we expose the case when $s = 3$ which is wholly typical. We may write

$$B^{(4)} = \left[\begin{array}{c|c} B_{11}^{(4)} & 0 \\ \hline P & C \end{array}\right] \begin{matrix} \}n - m_3 \\ \}m_3 \end{matrix} \tag{10.16}$$

and it is evident that

$$[B^{(4)}]^t = \begin{bmatrix} (B_{11}^{(4)})^t & 0 \\ P_t & C^t \end{bmatrix}. \tag{10.17}$$

Hence

$$[B^{(4)}]^t \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} (B_{11}^{(4)})^t x \\ P_t x + C^t y \end{bmatrix}, \tag{10.18}$$

and since $B_{11}^{(4)}$ is nonsingular, $(B_{11}^{(4)})^t x$ is not null unless $x = 0$. All vectors in the relevant invariant subspace have their first $n - m_3$ components equal to zero, and since

$$[B^{(4)}]^t \begin{bmatrix} 0 \\ y \end{bmatrix} = \begin{bmatrix} 0 \\ C^t y \end{bmatrix}, \tag{10.19}$$

it is evident that we may concentrate on the matrix C given explicitly by

$$C = \begin{bmatrix} 0 & 0 & 0 \\ B_{32}^{(4)} & 0 & 0 \\ B_{42}^{(4)} & B_{43}^{(4)} & 0 \end{bmatrix} \tag{10.20}$$

A discussion of vectors of grade 3 will be fully illustrative. Let us take any vector x

of order m_3 and partition conformally into $x^T = [x_1^T \mid x_2^T \mid x_3^T]$. If $x_1 \neq 0$, we have

$$Cx = \begin{bmatrix} 0 \\ B_{32}^{(4)} \\ B_{42}^{(4)} \end{bmatrix} x_1 + \begin{bmatrix} 0 \\ 0 \\ B_{43}^{(4)} \end{bmatrix} x_2, \tag{10.21}$$

$$C^2x = \begin{bmatrix} 0 \\ 0 \\ B_{43}^{(4)} \end{bmatrix} B_{32}^{(4)} x_1 = \begin{bmatrix} 0 \\ 0 \\ B_{43}^{(4)} \end{bmatrix} z \quad \text{(say)}. \tag{10.22}$$

But since we know the columns of $B_{32}^{(4)}$ are independent, $z \neq 0$, and since also the columns of $B_{43}^{(4)}$ are independent, $C^2x \neq 0$. On the other hand, $C^3x = 0$ for any x. The last n_1 columns of the identity matrix therefore give n_1 orthogonal vectors of grade 1, the next n_2 columns of it give vectors of grade 2 and the next n_3 columns give vectors of grade 3.

Interpreting this result in terms of B for the general case, we see that the last n_1 columns of V give orthogonal vectors of grade 1, the next n_2 give orthogonal vectors of grade 2, etc.

When the process terminates, $B_{11}^{(s+1)}$ is nonsingular and its eigenvalues are the remaining eigenvalues of B, i.e. $B_{11}^{(s+1)} + \lambda I$ gives the remaining eigenvalues of A. We can now turn to the next eigenvalue of A and repeat this process starting from $B^{(s+1)} + \lambda I$. In this way, a canonical form is ultimately attained, which may be illustrated in the case when A has only three distinct eigenvalues $\lambda_1, \lambda_2, \lambda_3$ by

$$V^H A V = \left[\begin{array}{cc|ccc|cc} \lambda_1 I & & & & & & \\ Y_{21} & \lambda_1 I & & & & & \\ \hline X_{31} & X_{32} & \lambda_2 I & & & & \\ X_{41} & X_{42} & Y_{43} & \lambda_2 I & & & \\ X_{51} & X_{52} & Y_{53} & Y_{54} & \lambda_2 I & & \\ \hline X_{61} & X_{62} & X_{63} & X_{64} & X_{65} & \lambda_3 I & \\ X_{71} & X_{72} & X_{73} & X_{74} & X_{75} & Y_{76} & \lambda_3 I \end{array}\right] \begin{matrix} \}n_2^{(1)} \\ \}n_1^{(1)} \\ \}n_3^{(2)} \\ \}n_2^{(2)} \\ \}n_1^{(2)} \\ \}n_2^{(3)} \\ \}n_1^{(3)} \end{matrix} \tag{10.23}$$

In the example given here, there were two stages with λ_3, three stages with λ_2 and two stages with λ_1 and the integers $n_j^{(i)}$ are the nullities exposed in the successive stages of the process. The matrix V being the product of unitary matrices is itself unitary. Note that we have denoted the submatrices in the diagonal blocks by Y_{ij} and outside these blocks by X_{ij}. From the definition of the algorithm, we have $n_j^{(i)} \geqq n_{j+1}^{(i)}$, and the columns of $Y_{i+1,i}$ are linearly independent. We already know that $n_1^{(3)}, n_2^{(3)}$ give the number of vectors of grades 1 and 2, respectively, associated with λ_3, and the corresponding columns of V provide the vectors themselves. The remaining columns of V cannot, of course, give vectors corresponding to

λ_2 and λ_1 since, in general, the latter will not be orthogonal to those of λ_3. We have not yet established that $n_1^{(2)}, n_2^{(2)}, n_3^{(2)}$ gives the *number* of vectors of grades 1, 2, 3 associated with λ_2, and that $n_1^{(1)}$ and $n_2^{(1)}$ the vectors of grades 1, 2 associated with λ_1, and this we now do.

We cannot proceed further with the reduction without departing from *unitary* similarities. However, if we now admit general similarities, the submatrices denoted by the X_{ij} may be annihilated. To annihilate X_{42}, for example, we premultiply by Z_{42}^{-1} and postmultiply by Z_{42}, where Z_{42} is equal to the identity matrix with a block $X_{42}/(\lambda_1 - \lambda_2)$ in the same position as is occupied by X_{42}. The X_{ij} are eliminated in this way in the order $X_{32}, X_{31}, X_{42}, X_{41}, X_{52}, X_{51}, X_{65}, X_{64}, \cdots$. It is obvious that the Y_{ij} are unaffected by this. The final result is that we have

$$C = M^{-1}AM = \left[\begin{array}{cc|cc|cc} \lambda_1 I & & & & & \\ Y_{21} & \lambda_1 I & & 0 & & 0 \\ \hline & & \lambda_2 I & & & \\ & 0 & Y_{43} & \lambda_2 I & & 0 \\ & & Y_{53} & Y_{54} & \lambda_2 I & \\ \hline & & & & \lambda_3 I & \\ & 0 & & 0 & Y_{76} & \lambda_3 I \end{array}\right], \tag{10.24}$$

where M is no longer unitary (though its last $n_1^{(3)} + n_2^{(3)}$ columns are still orthogonal). From the properties of the $Y_{i+1,i}$ described above, it is now evident that the $n_j^{(i)}$ have the significance described above and indeed that all the columns of M give independent (but not orthogonal) vectors of the relevant grades corresponding to $\lambda_1, \lambda_2, \lambda_3$. Notice that we have now proved that the canonical form (10.23) which is achieved purely by unitary transformations gives a full specification of the J.c.f. There is no need actually to proceed to the form (10.24) in order to find the J.c.f. However, from the form C and the rank properties of the $Y_{i+1,i}$, we may proceed to a demonstration of the J.c.f. itself. It is easy to show that by further similarities (10.24) may be reduced to

$$\left[\begin{array}{cc|cc|cc} \lambda_1 I & & & & & \\ K_{21} & \lambda_1 I & & & & \\ \hline & & \lambda_2 I & & & \\ & & K_{43} & \lambda_2 I & & \\ & & & K_{54} & \lambda_2 I & \\ \hline & & & & \lambda_3 I & \\ & & & & K_{76} & \lambda_3 I \end{array}\right], \tag{10.25}$$

where $K_{i+1,i}$ is a matrix of the form

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{10.26}$$

having the same dimension as $Y_{i+1,i}$. Apart from the ordering of the rows and columns, this is the J.c.f.

It should be emphasized that we are not recommending proceeding beyond the form (10.23), and, indeed, if one requires an orthogonal basis associated with each of the λ_i, one should return to the *original matrix* with each eigenvalue in turn.

The outstanding weakness of the algorithms of this section and the previous one is that the volume of work may be excessive. To find the vectors for a matrix of order n corresponding to an eigenvalue λ_1 of multiplicity k having just one block $J_r(\lambda_1)$ in the J.c.f., one must perform an S.V.D. on matrices of orders n, $n-1, \cdots, n-r$ in succession (the last one merely to reveal that there are no more eigenvalues equal to λ_1!).

Both algorithms were suggested by Kublanovskaya [10], but not in terms of the S.V.D., and have also been described by Ruhe [14], though in different terms from those used here.

11. Economical algorithm for determination of vectors. An alternative algorithm suggested by Golub and Wilkinson is considerably more economical in general (though not necessarily superior in other respects). Again corresponding to an eigenvalue λ, one works with $B = A - \lambda I$. We first give the basic motivation. Suppose we have already determined independent vectors u_1, u_2, u_3 of grade 1, vectors v_1, v_2 of grade 2 and vectors w_1, w_2 of grade 3 (not necessarily orthogonal).

If x is any vector of grade 4, then Bx is of grade 3 and hence lies in the subspace spanned by the u_i, v_i, w_i. In fact, x must satisfy a relation

$$Bx = [u_1 \mid u_2 \mid u_3 \mid v_1 \mid v_2 \mid w_1 \mid w_2]\alpha, \tag{11.1}$$

where α is a vector of order 7. However, the totality of independent solutions of (11.1) includes v_1, v_2, w_1, w_2, which will have been obtained by previously solving

$$Bx = [u_1 \mid u_2 \mid u_3 \mid v_1 \mid v_2]\beta \quad \text{and} \quad Bx = [u_1 \mid u_2 \mid u_3]\gamma. \tag{11.2}$$

We need a procedure which will reject these previous solutions. Indeed, the solutions needed at the current stage are solutions of (11.1) which are independent of v_1, v_2, w_1, w_2. To this end, we observe that instead of solving (11.1), we may equally well solve

$$Bx = ([u_1 \mid u_2 \mid u_3 \mid v_1 \mid v_2 \mid w_1 \mid w_2]Z)\alpha, \tag{11.3}$$

where Z is any nonsingular 7×7 matrix, preferably unitary if one does not wish to sacrifice numerical information. Now B is a singular matrix, and a convenient