

NAG Fortran Library Chapter Introduction

F08 – Least-squares and Eigenvalue Problems (LAPACK)

Contents

| | | |
|----------|--|-----------|
| 1 | Scope of the Chapter | 3 |
| 2 | Background to the Problems | 3 |
| 2.1 | Linear Least-squares Problems | 3 |
| 2.2 | Orthogonal Factorizations and Least-squares Problems | 4 |
| 2.2.1 | <i>QR</i> factorization | 4 |
| 2.2.2 | <i>LQ</i> factorization | 5 |
| 2.2.3 | <i>QR</i> factorization with column pivoting | 5 |
| 2.2.4 | Complete orthogonal factorization | 6 |
| 2.2.5 | Other factorizations | 6 |
| 2.3 | The Singular Value Decomposition | 6 |
| 2.4 | The Singular Value Decomposition and Least-squares Problems | 7 |
| 2.5 | Generalized Linear Least-squares Problems | 7 |
| 2.6 | Generalized Orthogonal Factorization and Generalized Linear Least-squares Problems | 8 |
| 2.6.1 | Generalized <i>QR</i> Factorization | 8 |
| 2.6.2 | Generalized <i>RQ</i> Factorization | 9 |
| 2.6.3 | Generalized Singular Value Decomposition (GSVD) | 10 |
| 2.7 | Symmetric Eigenvalue Problems | 11 |
| 2.8 | Generalized Symmetric-definite Eigenvalue Problems | 12 |
| 2.9 | Packed Storage for Symmetric Matrices | 13 |
| 2.10 | Band Matrices | 13 |
| 2.11 | Nonsymmetric Eigenvalue Problems | 13 |
| 2.12 | Generalized Nonsymmetric Eigenvalue Problem | 14 |
| 2.13 | The Sylvester Equation | 15 |
| 2.14 | Error and Perturbation Bounds and Condition Numbers | 15 |
| 2.14.1 | Least-squares problems | 16 |
| 2.14.2 | The singular value decomposition | 17 |
| 2.14.3 | The symmetric eigenproblem | 18 |
| 2.14.4 | The generalized symmetric-definite eigenproblem | 19 |
| 2.14.5 | The nonsymmetric eigenproblem | 19 |
| 2.14.6 | Balancing and condition for the nonsymmetric eigenproblem | 20 |
| 2.14.7 | The generalized nonsymmetric eigenvalue problem | 20 |
| 2.14.8 | Balancing the generalized eigenvalue problem | 21 |
| 2.14.9 | Other problems | 21 |
| 2.15 | Block Partitioned Algorithms | 21 |
| 3 | Recommendations on Choice and Use of Available Routines | 22 |
| 3.1 | Available Routines | 22 |
| 3.1.1 | Driver routines | 22 |
| 3.1.1.1 | Linear least-squares problems (LLS) | 22 |

| | | |
|----------|---|-----------|
| 3.1.1.2 | Generalized linear least-squares problems (LSE and GLM) | 22 |
| 3.1.1.3 | Symmetric eigenvalue problems (SEP) | 22 |
| 3.1.1.4 | Nonsymmetric eigenvalue problem (NEP) | 23 |
| 3.1.1.5 | Singular value decomposition (SVD) | 23 |
| 3.1.1.6 | Generalized symmetric definite eigenvalue problems (GSEP) | 23 |
| 3.1.1.7 | Generalized nonsymmetric eigenvalue problem (GNEP) | 23 |
| 3.1.1.8 | Generalized singular value decomposition (GSVD) | 24 |
| 3.1.2 | Computational routines | 24 |
| 3.1.2.1 | Orthogonal factorizations | 24 |
| 3.1.2.2 | Singular value problems | 24 |
| 3.1.2.3 | Symmetric eigenvalue problems | 25 |
| 3.1.2.4 | Generalized symmetric-definite eigenvalue problems | 27 |
| 3.1.2.5 | Nonsymmetric eigenvalue problems | 28 |
| 3.1.2.6 | Generalized nonsymmetric eigenvalue problems | 29 |
| 3.1.2.7 | Sylvester's equation | 30 |
| 3.2 | NAG Names and LAPACK Names | 30 |
| 3.3 | Matrix Storage Schemes | 31 |
| 3.3.1 | Conventional storage | 31 |
| 3.3.2 | Packed storage | 32 |
| 3.3.3 | Band storage | 33 |
| 3.3.4 | Tridiagonal and bidiagonal matrices | 33 |
| 3.3.5 | Real diagonal elements of complex matrices | 34 |
| 3.3.6 | Representation of orthogonal or unitary matrices | 34 |
| 3.4 | Parameter Conventions | 34 |
| 3.4.1 | Option parameters | 34 |
| 3.4.2 | Problem dimensions | 34 |
| 3.4.3 | Length of work arrays | 35 |
| 3.4.4 | Error-handling and the diagnostic parameter INFO | 35 |
| 4 | Decision Trees | 35 |
| 4.1 | General Purpose Routines (eigenvalues and eigenvectors) | 36 |
| 4.2 | General Purpose Routines (singular value decomposition) | 42 |
| 5 | Index | 42 |
| 6 | Routines Withdrawn or Scheduled for Withdrawal | 47 |
| 7 | References | 47 |

1 Scope of the Chapter

This chapter provides routines for the solution of linear least-squares problems, eigenvalue problems and singular value problems, as well as associated computations. It provides routines for:

- solution of linear least-squares problems
- solution of symmetric eigenvalue problems
- solution of nonsymmetric eigenvalue problems
- solution of singular value problems
- solution of generalized symmetric-definite eigenvalue problems
- solution of generalized nonsymmetric eigenvalue problems
- solution of generalized singular value problems
- solution of generalized linear least-squares problems
- matrix factorizations associated with the above problems
- estimating condition numbers of eigenvalue and eigenvector problems
- estimating the numerical rank of a matrix
- solution of the Sylvester matrix equation

Routines are provided for both *real* and *complex* data.

For a general introduction to the solution of linear least-squares problems, you should turn first to Chapter F04. The decision trees, at the end of Chapter F04, direct you to the most appropriate routines in Chapters F04 or F08. Chapters F04 and F08 contain *Black Box* (or *driver*) routines which enable standard linear least-squares problems to be solved by a call to a single routine.

For a general introduction to eigenvalue and singular value problems, you should turn first to Chapter F02. The decision trees, at the end of Chapter F02, direct you to the most appropriate routines in Chapters F02 or F08. Chapters F02 and F08 contain *Black Box* (or *driver*) routines which enable standard types of problem to be solved by a call to a single routine. Often routines in Chapter F02 call Chapter F08 routines to perform the necessary computational tasks.

The routines in this chapter (F08) handle only *dense*, *band*, *tridiagonal* and *Hessenberg* matrices (not matrices with more specialized structures, or general sparse matrices). The tables in Section 3 and the decision trees in Section 4 direct you to the most appropriate routines in Chapter F08.

The routines in this chapter have all been derived from the LAPACK project (see Anderson *et al.* (1999)). They have been designed to be efficient on a wide range of high-performance computers, without compromising efficiency on conventional serial machines.

It is not expected that you will need to read all of the following sections, but rather you will pick out those sections relevant to your particular problem.

2 Background to the Problems

This section is only a brief introduction to the numerical solution of linear least-squares problems, eigenvalue and singular value problems. Consult a standard textbook for a more thorough discussion, for example Golub and Van Loan (1996).

2.1 Linear Least-squares Problems

The *linear least-squares problem* is

$$\underset{x}{\text{minimize}} \|b - Ax\|_2, \quad (1)$$

where A is an m by n matrix, b is a given m element vector and x is an n element solution vector.

In the most usual case $m \geq n$ and $\text{rank}(A) = n$, so that A has *full rank* and in this case the solution to problem (1) is unique; the problem is also referred to as finding a *least-squares solution* to an *overdetermined* system of linear equations.

When $m < n$ and $\text{rank}(A) = m$, there are an infinite number of solutions x which exactly satisfy $b - Ax = 0$. In this case it is often useful to find the unique solution x which minimizes $\|x\|_2$, and the problem is referred to as finding a *minimum-norm solution* to an *underdetermined* system of linear equations.

In the general case when we may have $\text{rank}(A) < \min(m, n)$ – in other words, A may be *rank-deficient* – we seek the *minimum-norm least-squares* solution x which minimizes both $\|x\|_2$ and $\|b - Ax\|_2$.

This chapter (F08) contains driver routines to solve these problems with a single call, as well as computational routines that can be combined with routines in Chapter F07 to solve these linear least-squares problems. Chapter F04 also contains Black Box routines to solve these linear least-squares problems in standard cases. The next two sections discuss the factorizations that can be used in the solution of linear least-squares problems.

2.2 Orthogonal Factorizations and Least-squares Problems

A number of routines are provided for factorizing a general rectangular m by n matrix A , as the product of an *orthogonal* matrix (*unitary* if complex) and a *triangular* (or possibly trapezoidal) matrix.

A real matrix Q is *orthogonal* if $Q^T Q = I$; a complex matrix Q is *unitary* if $Q^H Q = I$. Orthogonal or unitary matrices have the important property that they leave the 2-norm of a vector invariant, so that

$$\|x\|_2 = \|Qx\|_2,$$

if Q is orthogonal or unitary. They usually help to maintain numerical stability because they do not amplify rounding errors.

Orthogonal factorizations are used in the solution of linear least-squares problems. They may also be used to perform preliminary steps in the solution of eigenvalue or singular value problems, and are useful tools in the solution of a number of other problems.

2.2.1 QR factorization

The most common, and best known, of the factorizations is the *QR factorization* given by

$$A = Q \begin{pmatrix} R \\ 0 \end{pmatrix}, \quad \text{if } m \geq n,$$

where R is an n by n upper triangular matrix and Q is an m by m orthogonal (or unitary) matrix. If A is of full rank n , then R is non-singular. It is sometimes convenient to write the factorization as

$$A = (Q_1 Q_2) \begin{pmatrix} R \\ 0 \end{pmatrix}$$

which reduces to

$$A = Q_1 R,$$

where Q_1 consists of the first n columns of Q , and Q_2 the remaining $m - n$ columns.

If $m < n$, R is trapezoidal, and the factorization can be written

$$A = Q(R_1 R_2), \quad \text{if } m < n,$$

where R_1 is upper triangular and R_2 is rectangular.

The *QR* factorization can be used to solve the linear least-squares problem (1) when $m \geq n$ and A is of full rank, since

$$\|b - Ax\|_2 = \|Q^T b - Q^T A x\|_2 = \left\| \begin{pmatrix} c_1 - Rx \\ c_2 \end{pmatrix} \right\|_2,$$

where

$$c \equiv \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} Q_1^T b \\ Q_2^T b \end{pmatrix} = Q^T b;$$

and c_1 is an n element vector. Then x is the solution of the upper triangular system

$$Rx = c_1.$$

The residual vector r is given by

$$r = b - Ax = Q \begin{pmatrix} 0 \\ c_2 \end{pmatrix}.$$

The residual sum of squares $\|r\|_2^2$ may be computed without forming r explicitly, since

$$\|r\|_2 = \|b - Ax\|_2 = \|c_2\|_2.$$

2.2.2 LQ factorization

The *LQ factorization* is given by

$$A = \begin{pmatrix} L & 0 \end{pmatrix} Q = \begin{pmatrix} L & 0 \end{pmatrix} \begin{pmatrix} Q_1 \\ Q_2 \end{pmatrix} = LQ_1, \quad \text{if } m \leq n,$$

where L is m by m lower triangular, Q is n by n orthogonal (or unitary), Q_1 consists of the first m rows of Q , and Q_2 the remaining $n - m$ rows.

The *LQ* factorization of A is essentially the same as the *QR* factorization of A^T (A^H if A is complex), since

$$A = \begin{pmatrix} L & 0 \end{pmatrix} Q \Leftrightarrow A^T = Q^T \begin{pmatrix} L^T \\ 0 \end{pmatrix}.$$

The *LQ* factorization may be used to find a minimum norm solution of an underdetermined system of linear equations $Ax = b$ where A is m by n with $m < n$ and has rank m . The solution is given by

$$x = Q^T \begin{pmatrix} L^{-1}b \\ 0 \end{pmatrix}.$$

2.2.3 QR factorization with column pivoting

To solve a linear least-squares problem (1) when A is not of full rank, or the rank of A is in doubt, we can perform either a *QR* factorization with column pivoting or a singular value decomposition.

The *QR factorization with column pivoting* is given by

$$A = Q \begin{pmatrix} R \\ 0 \end{pmatrix} P^T, \quad m \geq n,$$

where Q and R are as before and P is a (real) permutation matrix, chosen (in general) so that

$$|r_{11}| \geq |r_{22}| \geq \cdots \geq |r_{nn}|$$

and moreover, for each k ,

$$|r_{kk}| \geq \|R_{k:j,j}\|_2, \quad j = k + 1, \dots, n.$$

If we put

$$R = \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix}$$

where R_{11} is the leading k by k upper triangular submatrix of R then, in exact arithmetic, if $\text{rank}(A) = k$, the whole of the submatrix R_{22} in rows and columns $k + 1$ to n would be zero. In numerical computation, the aim must be to determine an index k , such that the leading submatrix R_{11} is well-conditioned, and R_{22} is negligible, so that

$$R = \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix} \simeq \begin{pmatrix} R_{11} & R_{12} \\ 0 & 0 \end{pmatrix}.$$

Then k is the effective rank of A . See Golub and Van Loan (1996) for a further discussion of numerical rank determination.

The so-called basic solution to the linear least-squares problem (1) can be obtained from this factorization as

$$x = P \begin{pmatrix} R_{11}^{-1} \hat{c}_1 \\ 0 \end{pmatrix},$$

where \hat{c}_1 consists of just the first k elements of $c = Q^T b$.

2.2.4 Complete orthogonal factorization

The QR factorization with column pivoting does not enable us to compute a *minimum norm* solution to a rank-deficient linear least-squares problem, unless $R_{12} = 0$. However, by applying for further orthogonal (or unitary) transformations from the right to the upper trapezoidal matrix $(R_{11} \ R_{12})$, R_{12} can be eliminated:

$$(R_{11} \ R_{12})Z = (T_{11} \ 0).$$

This gives the **complete orthogonal factorization**

$$AP = Q \begin{pmatrix} T_{11} & 0 \\ 0 & 0 \end{pmatrix} Z^T$$

from which the minimum norm solution can be obtained as

$$x = PZ \begin{pmatrix} T_{11}^{-1} \hat{c}_1 \\ 0 \end{pmatrix}.$$

2.2.5 Other factorizations

The QL and RQ factorizations are given by

$$A = Q \begin{pmatrix} 0 \\ L \end{pmatrix}, \quad \text{if } m \geq n,$$

and

$$A = (0 \ R)Q, \quad \text{if } m \leq n.$$

The factorizations are less commonly used than either the QR or LQ factorizations described above, but have applications in, for example, the computation of generalized QR factorizations.

2.3 The Singular Value Decomposition

The *singular value decomposition* (SVD) of an m by n matrix A is given by

$$A = U \Sigma V^T, \quad (A = U \Sigma V^H \text{ in the complex case})$$

where U and V are orthogonal (unitary) and Σ is an m by n diagonal matrix with real diagonal elements, σ_i , such that

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{\min(m,n)} \geq 0.$$

The σ_i are the *singular values* of A and the first $\min(m,n)$ columns of U and V are the *left* and *right singular vectors* of A . The singular values and singular vectors satisfy

$$Av_i = \sigma_i u_i \quad \text{and} \quad A^T u_i = \sigma_i v_i \quad (\text{or } A^H u_i = \sigma_i v_i)$$

where u_i and v_i are the i th columns of U and V respectively.

The computation proceeds in the following stages.

1. The matrix A is reduced to bidiagonal form $A = U_1 B V_1^T$ if A is real ($A = U_1 B V_1^H$ if A is complex), where U_1 and V_1 are orthogonal (unitary if A is complex), and B is real and upper bidiagonal when $m \geq n$ and lower bidiagonal when $m < n$, so that B is nonzero only on the main diagonal and either on the first super-diagonal (if $m \geq n$) or the first sub-diagonal (if $m < n$).
2. The SVD of the bidiagonal matrix B is computed as $B = U_2 \Sigma V_2^T$, where U_2 and V_2 are orthogonal and Σ is diagonal as described above. The singular vectors of A are then $U = U_1 U_2$ and $V = V_1 V_2$.

If $m \gg n$, it may be more efficient to first perform a QR factorization of A , and then compute the SVD of the n by n matrix R , since if $A = QR$ and $R = U \Sigma V^T$, then the SVD of A is given by $A = (QU) \Sigma V^T$.

Similarly, if $m \ll n$, it may be more efficient to first perform an LQ factorization of A .

2.4 The Singular Value Decomposition and Least-squares Problems

The SVD may be used to find a minimum norm solution to a (possibly) rank-deficient linear least-squares problem (1). The effective rank, k , of A can be determined as the number of singular values which exceed a suitable threshold. Let $\hat{\Sigma}$ be the leading k by k submatrix of Σ , and \hat{V} be the matrix consisting of the first k columns of V . Then the solution is given by

$$x = \hat{V} \hat{\Sigma}^{-1} \hat{c}_1,$$

where \hat{c}_1 consists of the first k elements of $c = U^T b = U_2^T U_1^T b$.

2.5 Generalized Linear Least-squares Problems

The simple type of linear least-squares problem described in Section 2.1 can be generalized in various ways.

1. Linear least-squares problems with **equality constraints**:

$$\text{find } x \text{ to minimize } S = \|c - Ax\|_2^2 \quad \text{subject to} \quad Bx = d,$$

where A is m by n and B is p by n , with $p \leq n \leq m + p$. The equations $Bx = d$ may be regarded as a set of equality constraints on the problem of minimizing S . Alternatively the problem may be regarded as solving an overdetermined system of equations

$$\begin{pmatrix} A \\ B \end{pmatrix} x = \begin{pmatrix} c \\ d \end{pmatrix},$$

where some of the equations (those involving B) are to be solved exactly, and the others (those involving A) are to be solved in a least-squares sense. The problem has a unique solution on the assumptions that B has full row rank p and the matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ has full column rank n . (For linear least-squares problems with **inequality constraints**, refer to Chapter E04.)

2. **General Gauss–Markov linear model problems**:

$$\text{minimize } \|y\|_2 \quad \text{subject to} \quad d = Ax + By,$$

where A is m by n and B is m by p , with $n \leq m \leq n + p$. When $B = I$, the problem reduces to an ordinary linear least-squares problem. When B is square and nonsingular, it is equivalent to a **weighted linear least-squares problem**:

$$\text{find } x \text{ to minimize } \|B^{-1}(d - Ax)\|_2.$$

The problem has a unique solution on the assumptions that A has full column rank n , and the matrix (A, B) has full row rank m . Unless B is diagonal, for numerical stability it is generally preferable to solve a weighted linear least-squares problem as a general Gauss–Markov linear model problem.

2.6 Generalized Orthogonal Factorization and Generalized Linear Least-squares Problems

2.6.1 Generalized QR Factorization

The **generalized QR (GQR) factorization** of an n by m matrix A and an n by p matrix B is given by the pair of factorizations

$$A = QR \quad \text{and} \quad B = QTZ,$$

where Q and Z are respectively n by n and p by p orthogonal matrices (or unitary matrices if A and B are complex). R has the form

$$R = \begin{matrix} & m \\ n-m & \begin{pmatrix} R_{11} \\ 0 \end{pmatrix} \end{matrix}, \quad \text{if } n \geq m,$$

or

$$R = \begin{matrix} n & m-n \\ n & \begin{pmatrix} R_{11} & R_{12} \end{pmatrix} \end{matrix}, \quad \text{if } n < m,$$

where R_{11} is upper triangular. T has the form

$$T = \begin{matrix} p-n & n \\ n & \begin{pmatrix} 0 & T_{12} \end{pmatrix} \end{matrix}, \quad \text{if } n \leq p,$$

or

$$T = \begin{matrix} n-p & p \\ p & \begin{pmatrix} T_{11} \\ T_{21} \end{pmatrix} \end{matrix}, \quad \text{if } n > p,$$

where T_{12} or T_{21} is upper triangular.

Note that if B is square and nonsingular, the GQR factorization of A and B implicitly gives the QR factorization of the matrix $B^{-1}A$:

$$B^{-1} = Z^T(T^{-1}R)$$

without explicitly computing the matrix inverse B^{-1} or the product $B^{-1}A$.

The GQR factorization can be used to solve the general (Gauss–Markov) linear model problem (GLM) (see Section 2.5). Using the GQR factorization of A and B , we rewrite the equation $d = Ax + By$ as

$$\begin{aligned} Q^T d &= Q^T Ax + Q^T By \\ &= Rx + TZy. \end{aligned}$$

We partition this as

$$\begin{pmatrix} d_1 \\ d_2 \end{pmatrix} = \begin{matrix} m \\ n-m & \begin{pmatrix} R_{11} \\ 0 \end{pmatrix} \end{matrix} x + \begin{matrix} p-n+m & n-m \\ n-m & \begin{pmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{pmatrix} \end{matrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}$$

where

$$\begin{pmatrix} d_1 \\ d_2 \end{pmatrix} \equiv Q^T d, \quad \text{and} \quad \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \equiv Zy.$$

The GLM problem is solved by setting

$$y_1 = 0 \quad \text{and} \quad y_2 = T_{22}^{-1} d_2$$

from which we obtain the desired solutions

$$x = R_{11}^{-1}(d_1 - T_{12}y_2) \quad \text{and} \quad y = Z^T \begin{pmatrix} 0 \\ y_2 \end{pmatrix}.$$

2.6.2 Generalized RQ Factorization

The **generalized RQ (GRQ) factorisation** of an m by n matrix A and a p by n matrix B is given by the pair of factorizations

$$A = RQ, \quad B = ZTQ$$

where Q and Z are respectively n by n and p by p orthogonal matrices (or unitary matrices if A and B are complex. (R has the form

$$R = \begin{matrix} & n-m & m \\ m & \begin{pmatrix} & \\ 0 & R_{12} \end{pmatrix} \end{matrix}, \quad \text{if } m \leq n,$$

or

$$R = \begin{matrix} & n \\ m-n & \begin{pmatrix} R_{11} \\ & \\ & R_{21} \end{pmatrix} \\ n & \end{matrix}, \quad \text{if } m > n,$$

where R_{12} or R_{21} is upper triangular. T has the form

$$T = \begin{matrix} & n \\ n & \begin{pmatrix} T_{11} \\ & \\ & 0 \end{pmatrix} \\ p-n & \end{matrix}, \quad \text{if } p \geq n,$$

or

$$T = \begin{matrix} & p & n-p \\ p & \begin{pmatrix} T_{11} & T_{12} \end{pmatrix} \end{matrix}, \quad \text{if } p < n,$$

where T_{11} is upper triangular.

Note that if B is square and nonsingular, the GRQ factorization of A and B implicitly gives the RQ factorization of the matrix AB^{-1} :

$$AB^{-1} = (RT^{-1})Z^T$$

without explicitly computing the matrix B^{-1} or the product AB^{-1} .

The GRQ factorization can be used to solve the linear equality-constrained least-squares problem (LSE) (see Section 2.5). We use the GRQ factorization of B and A (note that B and A have swapped roles), written as

$$B = TQ \quad \text{and} \quad A = ZRQ.$$

We write the linear equality constraints $Bx = d$ as

$$TQx = d,$$

which we partition as:

$$\begin{matrix} & n-p & p \\ p & \begin{pmatrix} & \\ 0 & T_{12} \end{pmatrix} \end{matrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = d \quad \text{where} \quad \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \equiv Qx.$$

Therefore x_2 is the solution of the upper triangular system

$$T_{12}x_2 = d.$$

Furthermore,

$$\begin{aligned} \|Ax - c\|_2 &= \|Z^T Ax - Z^T c\|_2 \\ &= \|RQx - Z^T c\|_2. \end{aligned}$$

We partition this expression as:

$$P + m - n \begin{pmatrix} n-p & p \\ R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - \begin{pmatrix} c_1 \\ c_2 \end{pmatrix},$$

where $\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \equiv Z^T c$.

To solve the LSE problem, we set

$$R_{11}x_1 + R_{12}x_2 - c_1 = 0$$

which gives x_1 as the solution of the upper triangular system

$$R_{11}x_1 = c_1 - R_{12}x_2.$$

Finally, the desired solution is given by

$$x = Q^T \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

2.6.3 Generalized Singular Value Decomposition (GSVD)

The **generalized (or quotient) singular value decomposition** of an m by n matrix A and a p by n matrix B is given by the pair of factorizations

$$A = U \Sigma_1 [0, R] Q^T \quad \text{and} \quad B = V \Sigma_2 [0, R] Q^T.$$

The matrices in these factorizations have the following properties:

- U is m by m , V is p by p , Q is n by n , and all three matrices are orthogonal. If A and B are complex, these matrices are unitary instead of orthogonal, and Q^T should be replaced by Q^H in the pair of factorizations.
- R is r by r , upper triangular and nonsingular. $[0, R]$ is r by n (in other words, the 0 is an r by $n-r$ zero matrix). The integer r is the rank of $\begin{pmatrix} A \\ B \end{pmatrix}$, and satisfies $r \leq n$.
- Σ_1 is m by r , Σ_2 is p by r , both are real, nonnegative and diagonal, and $\Sigma_1^T \Sigma_1 + \Sigma_2^T \Sigma_2 = I$. Write $\Sigma_1^T \Sigma_1 = \text{diag}(\alpha_1^2, \dots, \alpha_r^2)$ and $\Sigma_2^T \Sigma_2 = \text{diag}(\beta_1^2, \dots, \beta_r^2)$, where α_i and β_i lie in the interval from 0 to 1. The ratios $\alpha_1/\beta_1, \dots, \alpha_r/\beta_r$ are called the **generalized singular values** of the pair A, B . If $\beta_i = 0$, then the generalized singular value α_i/β_i is **infinite**.

Σ_1 and Σ_2 have the following detailed structures, depending on whether $m-r \geq 0$ or $m-r < 0$. In the first case, $m-r \geq 0$, then

$$\Sigma_1 = \begin{matrix} & k & l \\ & \begin{pmatrix} I & 0 \\ 0 & C \end{pmatrix} \\ m-k-l & \begin{pmatrix} 0 & 0 \end{pmatrix} \end{matrix} \quad \text{and} \quad \Sigma_2 = \begin{matrix} & k & l \\ & \begin{pmatrix} 0 & S \\ 0 & 0 \end{pmatrix} \\ p-l & \begin{pmatrix} 0 & 0 \end{pmatrix} \end{matrix}.$$

Here l is the rank of B , $k = r - l$, C and S are diagonal matrices satisfying $C^2 + S^2 = I$, and S is nonsingular. We may also identify $\alpha_1 = \dots = \alpha_k = 1$, $\alpha_{k+i} = c_{ii}$ for $i = 1, \dots, l$, $\beta_1 = \dots = \beta_k = 0$, and $\beta_{k+i} = s_{ii}$ for $i = 1, \dots, l$. Thus, the first k generalized singular values $\alpha_1/\beta_1, \dots, \alpha_k/\beta_k$ are infinite, and the remaining l generalized singular values are finite.

In the second case, when $m-r < 0$,

$$\Sigma_1 = \begin{matrix} & k & m-k & k+l-m \\ & \begin{pmatrix} I & 0 & 0 \\ 0 & C & 0 \end{pmatrix} \\ m-k & \begin{pmatrix} 0 & 0 & 0 \end{pmatrix} \end{matrix}$$

and

$$\Sigma_2 = \begin{matrix} & \begin{matrix} k & m-k & k+l-m \end{matrix} \\ \begin{matrix} m-k \\ k+l-m \\ p-l \end{matrix} & \begin{pmatrix} 0 & S & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{pmatrix} \end{matrix}.$$

Again, l is the rank of B , $k = r - l$, C and S are diagonal matrices satisfying $C^2 + S^2 = I$, and S is nonsingular, and we may identify $\alpha_1 = \dots = \alpha_k = 1$, $\alpha_{k+i} = c_{ii}$ for $i = 1, \dots, m - k$, $\alpha_{m+1} = \dots = \alpha_r = 0$, $\beta_1 = \dots = \beta_k = 0$, $\beta_{k+i} = s_{ii}$ for $i = 1, \dots, m - k$ and $\beta_{m+1} = \dots = \beta_r = 1$. Thus, the first k generalized singular values $\alpha_1/\beta_1, \dots, \alpha_k/\beta_k$ are infinite, and the remaining l generalized singular values are finite.

Here are some important special case of the generalized singular value decomposition. First, if B is square and nonsingular, then $r = n$ and the generalized singular value decomposition of A and B is equivalent to the singular value decomposition of AB^{-1} , where the singular values of AB^{-1} are equal to the generalized singular values of the pair A, B :

$$AB^{-1} = (U\Sigma_1 RQ^T)(V\Sigma_2 RQ^T)^{-1} = U(\Sigma_1 \Sigma_2^{-1})V^T.$$

Second, if the columns of $(A^T B^T)^T$ are orthonormal, then $r = n$, $R = I$ and the generalized singular value decomposition of A and B is equivalent to the CS (Cosine-Sine) decomposition of $(A^T B^T)^T$:

$$\begin{pmatrix} A \\ B \end{pmatrix} = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} \Sigma_1 \\ \Sigma_2 \end{pmatrix} Q^T.$$

Third, the generalized eigenvalues and eigenvectors of $A^T A - \lambda B^T B$ can be expressed in terms of the generalized singular value decomposition: Let

$$X = Q \begin{pmatrix} I & 0 \\ 0 & R^{-1} \end{pmatrix}.$$

Then

$$X^T A^T A X = \begin{pmatrix} 0 & 0 \\ 0 & \Sigma_1^T \Sigma_1 \end{pmatrix} \quad \text{and} \quad X^T B^T B X = \begin{pmatrix} 0 & 0 \\ 0 & \Sigma_2^T \Sigma_2 \end{pmatrix}.$$

Therefore, the columns of X are the eigenvectors of $A^T A - \lambda B^T B$, and ‘nontrivial’ eigenvalues are the squares of the generalized singular values (see also Section 2.8). ‘Trivial’ eigenvalues are those corresponding to the leading $n - r$ columns of X , which span the common null space of $A^T A$ and $B^T B$. The ‘trivial eigenvalues’ are not well defined.

2.7 Symmetric Eigenvalue Problems

The *symmetric eigenvalue problem* is to find the *eigenvalues*, λ , and corresponding *eigenvectors*, $z \neq 0$, such that

$$Az = \lambda z, \quad A = A^T, \quad \text{where } A \text{ is real.}$$

For the *Hermitian eigenvalue problem* we have

$$Az = \lambda z, \quad A = A^H, \quad \text{where } A \text{ is complex.}$$

For both problems the eigenvalues λ are real.

When all eigenvalues and eigenvectors have been computed, we write

$$A = \Lambda Z Z^T \quad (\text{or } A = \Lambda Z Z^H \text{ if complex}),$$

where Λ is a diagonal matrix whose diagonal elements are the eigenvalues, and Z is an orthogonal (or unitary) matrix whose columns are the eigenvectors. This is the classical *spectral factorization* of A .

The basic task of the symmetric eigenproblem routines is to compute values of λ and, optionally, corresponding vectors z for a given matrix A . This computation proceeds in the following stages.

1. The real symmetric or complex Hermitian matrix A is reduced to *real tridiagonal form* T . If A is real symmetric this decomposition is $A = QTQ^T$ with Q orthogonal and T symmetric tridiagonal. If A is complex Hermitian, the decomposition is $A = QTQ^H$ with Q unitary and T , as before, *real symmetric tridiagonal*.
2. Eigenvalues and eigenvectors of the real symmetric tridiagonal matrix T are computed. If all eigenvalues and eigenvectors are computed, this is equivalent to factorizing T as $T = SAS^T$, where S is orthogonal and A is diagonal. The diagonal entries of A are the eigenvalues of T , which are also the eigenvalues of A , and the columns of S are the eigenvectors of T ; the eigenvectors of A are the columns of $Z = QS$, so that $A = ZAZ^T$ (ZAZ^H when A is complex Hermitian).

This chapter supports three primary algorithms for computing eigenvalues and eigenvectors of real symmetric matrices and complex Hermitian matrices. They are:

- (i) the divide-and-conquer algorithm;
- (ii) the QR algorithm;
- (iii) bisection followed by inverse iteration.

The divide-and-conquer algorithm is generally more efficient than the traditional QR algorithm and is recommended for computing all eigenvalues and eigenvectors.

2.8 Generalized Symmetric-definite Eigenvalue Problems

This section is concerned with the solution of the generalized eigenvalue problems $Az = \lambda Bz$, $ABz = \lambda z$, and $BAz = \lambda z$, where A and B are real symmetric or complex Hermitian and B is positive-definite. Each of these problems can be reduced to a standard symmetric eigenvalue problem, using a Cholesky factorization of B as either $B = LL^T$ or $B = U^T U$ (LL^H or $U^H U$ in the Hermitian case).

With $B = LL^T$, we have

$$Az = \lambda Bz \Rightarrow (L^{-1}AL^{-T})(L^T z) = \lambda(L^T z).$$

Hence the eigenvalues of $Az = \lambda Bz$ are those of $Cy = \lambda y$, where C is the symmetric matrix $C = L^{-1}AL^{-T}$ and $y = L^T z$. In the complex case C is Hermitian with $C = L^{-1}AL^{-H}$ and $y = L^H z$.

Table 1 summarizes how each of the three types of problem may be reduced to standard form $Cy = \lambda y$, and how the eigenvectors z of the original problem may be recovered from the eigenvectors y of the reduced problem. The table applies to real problems; for complex problems, transposed matrices must be replaced by conjugate-transposes.

| | Type of problem | Factorization of B | Reduction | Recovery of eigenvectors |
|----|-------------------|-----------------------------|--|----------------------------------|
| 1. | $Az = \lambda Bz$ | $B = LL^T$, $B = U^T U$ | $C = L^{-1}AL^{-T}$, $C = U^{-T}AU^{-1}$ | $z = L^{-T}y$, $z = U^{-1}y$ |
| 2. | $ABz = \lambda z$ | $B = LL^T$, $B = U^T U$ | $C = L^T AL$, $C = UAU^T$ | $z = L^{-T}y$, $z = U^{-1}y$ |
| 3. | $BAz = \lambda z$ | $B = LL^T$, $B = U^T U$ | $C = L^T AL$, $C = UAU^T$ | $z = Ly$, $z = U^T y$ |

Table 1

Reduction of generalized symmetric-definite eigenproblems to standard problems

When the generalized symmetric-definite problem has been reduced to the corresponding standard problem $Cy = \lambda y$, this may then be solved using the routines described in the previous section. No special routines are needed to recover the eigenvectors z of the generalized problem from the eigenvectors y of the standard

problem, because these computations are simple applications of Level 2 or Level 3 BLAS (see Chapter F06).

2.9 Packed Storage for Symmetric Matrices

Routines which handle symmetric matrices are usually designed so that they use either the upper or lower triangle of the matrix; it is not necessary to store the whole matrix. If either the upper or lower triangle is stored conventionally in the upper or lower triangle of a two-dimensional array, the remaining elements of the array can be used to store other useful data. However, that is not always convenient, and if it is important to economize on storage, the upper or lower triangle can be stored in a one-dimensional array of length $n(n+1)/2$; that is, the storage is almost halved.

This storage format is referred to as *packed storage*; it is described in Section 3.3.

Routines designed for packed storage are usually less efficient, especially on high-performance computers, so there is a trade-off between storage and efficiency.

2.10 Band Matrices

A *band* matrix is one whose elements are confined to a relatively small number of sub-diagonals or super-diagonals on either side of the main diagonal. Algorithms can take advantage of bandedness to reduce the amount of work and storage required. The storage scheme for band matrices is described in Section 3.3.

If the problem is the generalized symmetric definite eigenvalue problem $Az = \lambda Bz$ and the matrices A and B are additionally banded, the matrix C as defined in Section 2.8 is, in general, full. We can reduce the problem to a banded standard problem by modifying the definition of C thus:

$$C = X^T A X, \quad \text{where} \quad X = U^{-1}Q \quad \text{or} \quad L^{-T}Q,$$

where Q is an orthogonal matrix chosen to ensure that C has bandwidth no greater than that of A .

A further refinement is possible when A and B are banded, which halves the amount of work required to form C . Instead of the standard Cholesky factorization of B as $U^T U$ or LL^T , we use a *split Cholesky* factorization $B = S^T S$, where

$$S = \begin{pmatrix} U_{11} & \\ M_{21} & L_{22} \end{pmatrix}$$

with U_{11} upper triangular and L_{22} lower triangular of order approximately $n/2$; S has the same bandwidth as B .

2.11 Nonsymmetric Eigenvalue Problems

The *nonsymmetric eigenvalue problem* is to find the *eigenvalues*, λ , and corresponding *eigenvectors*, $v \neq 0$, such that

$$Av = \lambda v.$$

More precisely, a vector v as just defined is called a *right eigenvector* of A , and a vector $u \neq 0$ satisfying

$$u^T A = \lambda u^T \quad (u^H A = \lambda u^H \quad \text{when } u \text{ is complex})$$

is called a *left eigenvector* of A .

A real matrix A may have complex eigenvalues, occurring as complex conjugate pairs.

This problem can be solved via the *Schur factorization* of A , defined in the real case as

$$A = ZT Z^T,$$

where Z is an orthogonal matrix and T is an upper quasi-triangular matrix with 1 by 1 and 2 by 2 diagonal blocks, the 2 by 2 blocks corresponding to complex conjugate pairs of eigenvalues of A . In the complex case, the Schur factorization is

$$A = ZT Z^H,$$

where Z is unitary and T is a complex upper triangular matrix.

The columns of Z are called the *Schur vectors*. For each k ($1 \leq k \leq n$), the first k columns of Z form an orthonormal basis for the *invariant subspace* corresponding to the first k eigenvalues on the diagonal of T . Because this basis is orthonormal, it is preferable in many applications to compute Schur vectors rather than eigenvectors. It is possible to order the Schur factorization so that any desired set of k eigenvalues occupy the k leading positions on the diagonal of T .

The two basic tasks of the nonsymmetric eigenvalue routines are to compute, for a given matrix A , all n values of λ and, if desired, their associated right eigenvectors v and/or left eigenvectors u , and the Schur factorization.

These two basic tasks can be performed in the following stages.

1. A general matrix A is reduced to *upper Hessenberg form* H which is zero below the first sub-diagonal. The reduction may be written $A = QHQ^T$ with Q orthogonal if A is real, or $A = QHQ^H$ with Q unitary if A is complex.
2. The upper Hessenberg matrix H is reduced to Schur form T , giving the Schur factorization $H = STS^T$ (for H real) or $H = STS^H$ (for H complex). The matrix S (the Schur vectors of H) may optionally be computed as well. Alternatively S may be postmultiplied into the matrix Q determined in stage 1, to give the matrix $Z = QS$, the Schur vectors of A . The eigenvalues are obtained from the diagonal elements or diagonal blocks of T .
3. Given the eigenvalues, the eigenvectors may be computed in two different ways. Inverse iteration can be performed on H to compute the eigenvectors of H , and then the eigenvectors can be multiplied by the matrix Q in order to transform them to eigenvectors of A . Alternatively the eigenvectors of T can be computed, and optionally transformed to those of H or A if the matrix S or Z is supplied.

The accuracy with which eigenvalues can be obtained can often be improved by *balancing* a matrix. This is discussed further in Section 2.14.6 below.

2.12 Generalized Nonsymmetric Eigenvalue Problem

The *generalized nonsymmetric eigenvalue problem* is to find the eigenvalues, λ , and corresponding eigenvectors, $v \neq 0$, such that

$$Av = \lambda Bv.$$

More precisely, a vector v as just defined is called a *right eigenvector* of the matrix pair (A, B) , and a vector $u \neq 0$ satisfying

$$u^T A = \lambda u^T B \quad (u^H A = \lambda u^H B \text{ when } u \text{ is complex})$$

is called a *left eigenvector* of the matrix pair (A, B) .

If B is singular then the problem has one or more *infinite eigenvalues* $\lambda = \infty$, corresponding to $Bv = 0$. Note that if A is non-singular, then the equivalent problem $\mu Av = Bv$ is perfectly well defined and an infinite eigenvalue corresponds to $\mu = 0$. To deal with both finite (including zero) and infinite eigenvalues, the routines in this chapter do not compute λ explicitly, but rather return a pair of numbers (α, β) such that if $\beta \neq 0$

$$\lambda = \alpha/\beta$$

and if $\alpha \neq 0$ and $\beta = 0$ then $\lambda = \infty$. β is always returned as real and non-negative. Of course, computationally an infinite eigenvalue may correspond to a small β rather than an exact zero.

For a given pair (A, B) the set of all the matrices of the form $(A - \lambda B)$ is called a matrix *pencil* and λ and v are said to be an eigenvalue and eigenvector of the pencil $(A - \lambda B)$. If A and B are both singular and share a common null-space then

$$\det(A - \lambda B) \equiv 0$$

so that the pencil $(A - \lambda B)$ is *singular* for all λ . In other words any λ can be regarded as an eigenvalue. In exact arithmetic a singular pencil will have $\alpha = \beta = 0$ for some (α, β) . Computationally if some pair (α, β) is small then the pencil is singular, or nearly singular, and no reliance can be placed on any of the computed eigenvalues. Singular pencils can also manifest themselves in other ways; see, in particular, Sections 2.3.5.2 and 4.11.1.4 of Anderson *et al.* (1999) for further details.

The generalized eigenvalue problem can be solved via the *generalized Schur factorization* of the pair (A, B) defined in the real case as

$$A = QSZ^T, \quad B = QTZ^T,$$

where Q and Z are orthogonal, T is upper triangular with non-negative diagonal elements and S is upper quasi-triangular with 1 by 1 and 2 by 2 diagonal blocks, the 2 by 2 blocks corresponding to complex conjugate pairs of eigenvalues. In the complex case, the generalized Schur factorization is

$$A = QSZ^H, \quad B = QTZ^H,$$

where Q and Z are unitary and S and T are upper triangular, with T having real non-negative diagonal elements. The columns of Q and Z are called respectively the *left* and *right generalized Schur vectors* and span pairs of *deflating subspaces* of A and B , which are a generalization of invariant subspaces.

The two basic tasks of the generalized nonsymmetric eigenvalue routines are to compute, for a given pair (A, B) , all n values of λ and, if desired, their associated right eigenvectors v and/or left eigenvectors u , and the generalized Schur factorization.

These two basic tasks can be performed in the following stages.

1. The matrix pair (A, B) is reduced to *generalized upper Hessenberg form* (H, R) , where H is upper Hessenberg (zero below the first sub-diagonal) and R is upper triangular. The reduction may be written as $A = Q_1 H Z_1^T, B = Q_1 R Z_1^T$ in the real case with Q_1 and Z_1 orthogonal, and $A = Q_1 H Z_1^H, B = Q_1 R Z_1^H$ in the complex case with Q_1 and Z_1 unitary.
2. The generalized upper Hessenberg form (H, R) is reduced to the generalized Schur form (S, T) using the generalized Schur factorization $H = Q_2 S Z_2^T, R = Q_2 T Z_2^T$ in the real case with Q_2 and Z_2 orthogonal, and $H = Q_2 S Z_2^H, R = Q_2 T Z_2^H$ in the complex case. The generalized Schur vectors of (A, B) are given by $Q = Q_1 Q_2, Z = Z_1 Z_2$. The eigenvalues are obtained from the diagonal elements (or blocks) of the pair (S, T) .
3. Given the eigenvalues, the eigenvectors of the pair (S, T) can be computed, and optionally transformed to those of (H, R) or (A, B) .

The accuracy with which eigenvalues can be obtained can often be improved by *balancing* a matrix pair. This is discussed further in Section 2.14.8 below.

2.13 The Sylvester Equation

The Sylvester equation is a matrix equation of the form

$$AX + XB = C,$$

where A, B , and C are given matrices with A being m by m , B an n by n matrix and C , and the solution matrix X , m by n matrices. The solution of a special case of this equation occurs in the computation of the condition number for an invariant subspace, but a combination of routines in this chapter allows the solution of the general Sylvester equation.

2.14 Error and Perturbation Bounds and Condition Numbers

In this section we discuss the effects of rounding errors in the solution process and the effects of uncertainties in the data, on the solution to the problem. A number of the routines in this chapter return information, such as condition numbers, that allow these effects to be assessed. First we discuss some notation used in the error bounds of later sections.

The bounds usually contain the factor $p(n)$ (or $p(m, n)$), which grows as a function of the matrix dimension n (or matrix dimensions m and n). It measures how errors can grow as a function of the matrix dimension, and represents a potentially different function for each problem. In practice, it usually grows just linearly; $p(n) \leq 10n$ is often true, although generally only much weaker bounds can be actually proved. We normally describe $p(n)$ as a ‘modestly growing’ function of n . For detailed derivations of various $p(n)$, see Golub and Van Loan (1996) and Wilkinson (1965).

For linear equation (see Chapter F07) and least-squares solvers, we consider bounds on the relative error $\|x - \hat{x}\|/\|x\|$ in the computed solution \hat{x} , where x is the true solution. For eigenvalue problems we

consider bounds on the error $|\lambda_i - \hat{\lambda}_i|$ in the i th computed eigenvalue $\hat{\lambda}_i$, where λ_i is the true i th eigenvalue. For singular value problems we similarly consider bounds $|\sigma_i - \hat{\sigma}_i|$.

Bounding the error in computed eigenvectors and singular vectors \hat{v}_i is more subtle because these vectors are not unique: even though we restrict $\|\hat{v}_i\|_2 = 1$ and $\|v_i\|_2 = 1$, we may still multiply them by arbitrary constants of absolute value 1. So to avoid ambiguity we bound the *angular difference* between \hat{v}_i and the true vector v_i , so that

$$\begin{aligned}\theta(v_i, \hat{v}_i) &= \text{acute angle between } v_i \text{ and } \hat{v}_i \\ &= \arccos|v_i^H \hat{v}_i|.\end{aligned}\tag{2}$$

Here $\arccos(\theta)$ is in the standard range: $0 \leq \arccos(\theta) < \pi$. When $\theta(v_i, \hat{v}_i)$ is small, we can choose a constant α with absolute value 1 so that $\|\alpha v_i - \hat{v}_i\|_2 \approx \theta(v_i, \hat{v}_i)$.

In addition to bounds for individual eigenvectors, bounds can be obtained for the spaces spanned by collections of eigenvectors. These may be much more accurately determined than the individual eigenvectors which span them. These spaces are called *invariant subspaces* in the case of eigenvectors, because if v is any vector in the space, Av is also in the space, where A is the matrix. Again, we will use angle to measure the difference between a computed space \hat{S} and the true space S :

$$\begin{aligned}\theta(S, \hat{S}) &= \text{acute angle between } S \text{ and } \hat{S} \\ &= \max_{\substack{s \in S \\ s \neq 0}} \min_{\substack{\hat{s} \in \hat{S} \\ \hat{s} \neq 0}} \theta(s, \hat{s}) \quad \text{or} \quad \max_{\substack{\hat{s} \in \hat{S} \\ \hat{s} \neq 0}} \min_{\substack{s \in S \\ s \neq 0}} \theta(s, \hat{s})\end{aligned}\tag{3}$$

$\theta(S, \hat{S})$ may be computed as follows. Let S be a matrix whose columns are orthonormal and span S . Similarly let \hat{S} be an orthonormal matrix with columns spanning \hat{S} . Then

$$\theta(S, \hat{S}) = \arccos \sigma_{\min}(S^H \hat{S}).$$

Finally, we remark on the accuracy of the bounds when they are large. Relative errors like $\|\hat{x} - x\|/\|x\|$ and angular errors like $\theta(\hat{v}_i, v_i)$ are only of interest when they are much less than 1. Some stated bounds are not strictly true when they are close to 1, but rigorous bounds are much more complicated and supply little extra information in the interesting case of small errors. These bounds are indicated by using the symbol \lesssim , or ‘approximately less than’, instead of the usual \leq . Thus, when these bounds are close to 1 or greater, they indicate that the computed answer may have no significant digits at all, but do not otherwise bound the error.

A number of routines in this chapter return error estimates and/or condition number estimates directly. In other cases Anderson *et al.* (1999) gives code fragments to illustrate the computation of these estimates, and a number of the Chapter F08 example programs, for the driver routines, implement these code fragments.

2.14.1 Least-squares problems

The conventional error analysis of linear least-squares problems goes as follows. The problem is to find the x minimizing $\|Ax - b\|_2$. Let \hat{x} be the solution computed using one of the methods described above. We discuss the most common case, where A is overdetermined (i.e., has more rows than columns) and has full rank.

Then the computed solution \hat{x} has a small normwise backward error. In other words \hat{x} minimizes $\|(A + E)\hat{x} - (b + f)\|_2$, where

$$\max\left(\frac{\|E\|_2}{\|A\|_2}, \frac{\|f\|_2}{\|b\|_2}\right) \leq p(n)\epsilon$$

and $p(n)$ is a modestly growing function of n and ϵ is the *machine precision*. Let $\kappa_2(A) = \sigma_{\max}(A)/\sigma_{\min}(A)$, $\rho = \|Ax - b\|_2$, and $\sin(\theta) = \rho/\|b\|_2$. Then if $p(n)\epsilon$ is small enough, the error $\hat{x} - x$ is bounded by

$$\frac{\|\hat{x} - x\|_2}{\|x\|_2} \lesssim p(n)\epsilon \left\{ \frac{2\kappa_2(A)}{\cos(\theta)} + \tan(\theta)\kappa_2^2(A) \right\}.$$

If A is rank-deficient, the problem can be *regularized* by treating all singular values less than a user-specified threshold as exactly zero. See Golub and Van Loan (1996) for error bounds in this case, as well as for the underdetermined case.

The solution of the overdetermined, full-rank problem may also be characterized as the solution of the linear system of equations

$$\begin{pmatrix} I & A \\ A^T & 0 \end{pmatrix} \begin{pmatrix} r \\ x \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix}.$$

By solving this linear system (see Chapter F07) componentwise error bounds can also be obtained Arioli *et al.* (1989).

2.14.2 The singular value decomposition

The usual error analysis of the SVD algorithm is as follows (see Golub and Van Loan (1996)).

The computed SVD, $\hat{U}\hat{\Sigma}\hat{V}^T$, is nearly the exact SVD of $A + E$, i.e., $A + E = (\hat{U} + \delta\hat{U})\hat{\Sigma}(\hat{V} + \delta\hat{V})$ is the true SVD, so that $\hat{U} + \delta\hat{U}$ and $\hat{V} + \delta\hat{V}$ are both orthogonal, where $\|E\|_2/\|A\|_2 \leq p(m, n)\epsilon$, $\|\delta\hat{U}\| \leq p(m, n)\epsilon$, and $\|\delta\hat{V}\| \leq p(m, n)\epsilon$. Here $p(m, n)$ is a modestly growing function of m and n and ϵ is the **machine precision**. Each computed singular value $\hat{\sigma}_i$ differs from the true σ_i by an amount satisfying the bound

$$|\hat{\sigma}_i - \sigma_i| \leq p(m, n)\epsilon\sigma_1.$$

Thus large singular values (those near σ_1) are computed to high relative accuracy and small ones may not be.

The angular difference between the computed left singular vector \hat{u}_i and the true u_i satisfies the approximate bound

$$\theta(\hat{u}_i, u_i) \lesssim \frac{p(m, n)\epsilon\|A\|_2}{\text{gap}_i}$$

where

$$\text{gap}_i = \min_{j \neq i} |\sigma_i - \sigma_j|$$

is the *absolute gap* between σ_i and the nearest other singular value. Thus, if σ_i is close to other singular values, its corresponding singular vector u_i may be inaccurate. The same bound applies to the computed right singular vector \hat{v}_i and the true vector v_i . The gaps may be easily obtained from the computed singular values.

Let $\hat{\mathcal{S}}$ be the space spanned by a collection of computed left singular vectors $\{\hat{u}_i, i \in \mathbf{I}\}$, where \mathbf{I} is a subset of the integers from 1 to n . Let \mathcal{S} be the corresponding true space. Then

$$\theta(\hat{\mathcal{S}}, \mathcal{S}) \lesssim \frac{p(m, n)\epsilon\|A\|_2}{\text{gap}_{\mathbf{I}}}.$$

where

$$\text{gap}_{\mathbf{I}} = \min\{|\sigma_i - \sigma_j| \mid i \in \mathbf{I}, j \notin \mathbf{I}\}$$

is the absolute gap between the singular values in \mathbf{I} and the nearest other singular value. Thus, a cluster of close singular values which is far away from any other singular value may have a well determined space $\hat{\mathcal{S}}$ even if its individual singular vectors are ill-conditioned. The same bound applies to a set of right singular vectors $\{\hat{v}_i, i \in \mathbf{I}\}$.

In the special case of bidiagonal matrices, the singular values and singular vectors may be computed much more accurately (see Demmel and Kahan (1990)). A bidiagonal matrix B has nonzero entries only on the main diagonal and the diagonal immediately above it (or immediately below it). Reduction of a dense matrix to bidiagonal form B can introduce additional errors, so the following bounds for the bidiagonal case do not apply to the dense case.

Using the routines in this chapter, each computed singular value of a bidiagonal matrix is accurate to nearly full relative accuracy, no matter how tiny it is, so that

$$|\hat{\sigma}_i - \sigma_i| \leq p(m, n)\epsilon\sigma_i.$$

The computed left singular vector \hat{u}_i has an angular error at most about

$$\theta(\hat{u}_i, u_i) \lesssim \frac{p(m, n)\epsilon}{\text{relgap}_i}$$

where

$$\text{relgap}_i = \min_{j \neq i} |\sigma_i - \sigma_j| / (\sigma_i + \sigma_j)$$

is the *relative gap* between σ_i and the nearest other singular value. The same bound applies to the right singular vector \hat{v}_i and v_i . Since the relative gap may be much larger than the absolute gap, this error bound may be much smaller than the previous one. The relative gaps may be easily obtained from the computed singular values.

2.14.3 The symmetric eigenproblem

The usual error analysis of the symmetric eigenproblem is as follows (see Parlett (1998)).

The computed eigendecomposition $\hat{Z}\hat{A}\hat{Z}^T$ is nearly the exact eigendecomposition of $A + E$, i.e., $A + E = (\hat{Z} + \delta\hat{Z})\hat{A}(\hat{Z} + \delta\hat{Z})^T$ is the true eigendecomposition so that $\hat{Z} + \delta\hat{Z}$ is orthogonal, where $\|E\|_2/\|A\|_2 \leq p(n)\epsilon$ and $\|\delta\hat{Z}\|_2 \leq p(n)\epsilon$ and $p(n)$ is a modestly growing function of n and ϵ is the **machine precision**. Each computed eigenvalue $\hat{\lambda}_i$ differs from the true λ_i by an amount satisfying the bound

$$|\hat{\lambda}_i - \lambda_i| \leq p(n)\epsilon\|A\|_2.$$

Thus large eigenvalues (those near $\max_i |\lambda_i| = \|A\|_2$) are computed to high relative accuracy and small ones may not be.

The angular difference between the computed unit eigenvector \hat{z}_i and the true z_i satisfies the approximate bound

$$\theta(\hat{z}_i, z_i) \lesssim \frac{p(n)\epsilon\|A\|_2}{\text{gap}_i}$$

if $p(n)\epsilon$ is small enough, where

$$\text{gap}_i = \min_{j \neq i} |\lambda_i - \lambda_j|$$

is the *absolute gap* between λ_i and the nearest other eigenvalue. Thus, if λ_i is close to other eigenvalues, its corresponding eigenvector z_i may be inaccurate. The gaps may be easily obtained from the computed eigenvalues.

Let $\hat{\mathcal{S}}$ be the invariant subspace spanned by a collection of eigenvectors $\{\hat{z}_i, i \in \mathbf{I}\}$, where \mathbf{I} is a subset of the integers from 1 to n . Let \mathcal{S} be the corresponding true subspace. Then

$$\theta(\hat{\mathcal{S}}, \mathcal{S}) \lesssim \frac{p(n)\epsilon\|A\|_2}{\text{gap}_{\mathbf{I}}}$$

where

$$\text{gap}_{\mathbf{I}} = \min\{|\lambda_i - \lambda_j| \mid \text{for } i \in \mathbf{I}, j \notin \mathbf{I}\}$$

is the absolute gap between the eigenvalues in \mathbf{I} and the nearest other eigenvalue. Thus, a cluster of close eigenvalues which is far away from any other eigenvalue may have a well determined invariant subspace $\hat{\mathcal{S}}$ even if its individual eigenvectors are ill-conditioned.

In the special case of a real symmetric tridiagonal matrix T , routines in this chapter can compute the eigenvalues and eigenvectors much more accurately. See Anderson *et al.* (1999) for further details.

2.14.4 The generalized symmetric-definite eigenproblem

The three types of problem to be considered are $A - \lambda B$, $AB - \lambda I$ and $BA - \lambda I$. In each case A and B are real symmetric (or complex Hermitian) and B is positive-definite. We consider each case in turn, assuming that routines in this chapter are used to transform the generalized problem to the standard symmetric problem, followed by the solution of the symmetric problem. In all cases

$$\text{gap}_i = \min_{j \neq i} |\lambda_i - \lambda_j|$$

is the *absolute gap* between λ_i and the nearest other eigenvalue.

1. $A - \lambda B$. The computed eigenvalues $\hat{\lambda}_i$ can differ from the true eigenvalues λ_i by an amount

$$|\hat{\lambda}_i - \lambda_i| \lesssim p(n)\epsilon \|B^{-1}\|_2 \|A\|_2.$$

The angular difference between the computed eigenvector \hat{z}_i and the true eigenvector z_i is

$$\theta(\hat{z}_i, z_i) \lesssim \frac{p(n)\epsilon \|B^{-1}\|_2 \|A\|_2 (\kappa_2(B))^{1/2}}{\text{gap}_i}.$$

2. $AB - \lambda I$ or $BA - \lambda I$. The computed eigenvalues $\hat{\lambda}_i$ can differ from the true eigenvalues λ_i by an amount

$$|\hat{\lambda}_i - \lambda_i| \lesssim p(n)\epsilon \|B\|_2 \|A\|_2.$$

The angular difference between the computed eigenvector \hat{z}_i and the true eigenvector z_i is

$$\theta(\hat{z}_i, z_i) \lesssim \frac{q(n)\epsilon \|B\|_2 \|A\|_2 (\kappa_2(B))^{1/2}}{\text{gap}_i}.$$

These error bounds are large when B is ill-conditioned with respect to inversion ($\kappa_2(B)$ is large). It is often the case that the eigenvalues and eigenvectors are much better conditioned than indicated here. One way to get tighter bounds is effective when the diagonal entries of B differ widely in magnitude, as for example with a *graded matrix*.

1. $A - \lambda B$. Let $D = \text{diag}(b_{11}^{-1/2}, \dots, b_{nn}^{-1/2})$ be a diagonal matrix. Then replace B by DBD and A by DAD in the above bounds.
2. $AB - \lambda I$ or $BA - \lambda I$. Let $D = \text{diag}(b_{11}^{-1/2}, \dots, b_{nn}^{-1/2})$ be a diagonal matrix. Then replace B by DBD and A by $D^{-1}AD^{-1}$ in the above bounds.

Further details can be found in Anderson *et al.* (1999).

2.14.5 The nonsymmetric eigenproblem

The nonsymmetric eigenvalue problem is more complicated than the symmetric eigenvalue problem. In this section, we just summarize the bounds. Further details can be found in Anderson *et al.* (1999).

We let $\hat{\lambda}_i$ be the i th computed eigenvalue and λ_i the i th true eigenvalue. Let \hat{v}_i be the corresponding computed right eigenvector, and v_i the true right eigenvector (so $Av_i = \lambda_i v_i$). If \mathbf{I} is a subset of the integers from 1 to n , we let $\lambda_{\mathbf{I}}$ denote the average of the selected eigenvalues: $\lambda_{\mathbf{I}} = \left(\sum_{i \in \mathbf{I}} \lambda_i \right) / \left(\sum_{i \in \mathbf{I}} 1 \right)$, and similarly for $\hat{\lambda}_{\mathbf{I}}$. We also let $\mathbf{S}_{\mathbf{I}}$ denote the subspace spanned by $\{v_i, i \in \mathbf{I}\}$; it is called a right invariant subspace because if v is any vector in $\mathbf{S}_{\mathbf{I}}$ then Av is also in $\mathbf{S}_{\mathbf{I}}$. $\hat{\mathbf{S}}_{\mathbf{I}}$ is the corresponding computed subspace.

The algorithms for the nonsymmetric eigenproblem are normwise backward stable: they compute the exact eigenvalues, eigenvectors and invariant subspaces of slightly perturbed matrices $A + E$, where $\|E\| \leq p(n)\epsilon \|A\|$. Some of the bounds are stated in terms of $\|E\|_2$ and others in terms of $\|E\|_F$; one may use $p(n)\epsilon$ for either quantity.

Routines are provided so that, for each $(\hat{\lambda}_i, \hat{v}_i)$ pair the two values s_i and sep_i , or for a selected subset \mathbf{I} of eigenvalues the values $s_{\mathbf{I}}$ and $sep_{\mathbf{I}}$ can be obtained, for which the error bounds in Table 2 are true for sufficiently small $\|E\|$, (which is why they are called asymptotic):

| | |
|--------------------|---|
| Simple eigenvalue | $ \hat{\lambda}_i - \lambda_i \lesssim \ E\ _2/s_i$ |
| Eigenvalue cluster | $ \hat{\lambda}_{\mathbf{I}} - \lambda_{\mathbf{I}} \lesssim \ E\ _2/s_{\mathbf{I}}$ |
| Eigenvector | $\theta(\hat{v}_i, v_i) \lesssim \ E\ _F/sep_i$ |
| Invariant subspace | $\theta(\hat{S}_{\mathbf{I}}, S_{\mathbf{I}}) \lesssim \ E\ _F/sep_{\mathbf{I}}$ |

Table 2

Asymptotic error bounds for the nonsymmetric eigenproblem

If the problem is ill-conditioned, the asymptotic bounds may only hold for extremely small $\|E\|$. The global error bounds of Table 3 are guaranteed to hold for all $\|E\|_F < s \times sep/4$:

| | | |
|--------------------|--|---|
| Simple eigenvalue | $ \hat{\lambda}_i - \lambda_i \leq n\ E\ _2/s_i$ | Holds for all E |
| Eigenvalue cluster | $ \hat{\lambda}_{\mathbf{I}} - \lambda_{\mathbf{I}} \leq 2\ E\ _2/s_{\mathbf{I}}$ | Requires $\ E\ _F < s_{\mathbf{I}} \times sep_{\mathbf{I}}/4$ |
| Eigenvector | $\theta(\hat{v}_i, v_i) \leq \arctan(2\ E\ _F/(sep_i - 4\ E\ _F/s_i))$ | Requires $\ E\ _F < s_i \times sep_i/4$ |
| Invariant subspace | $\theta(\hat{S}_{\mathbf{I}}, S_{\mathbf{I}}) \leq \arctan(2\ E\ _F/(sep_{\mathbf{I}} - 4\ E\ _F/s_{\mathbf{I}}))$ | Requires $\ E\ _F < s_{\mathbf{I}} \times sep_{\mathbf{I}}/4$ |

Table 3

Global error bounds for the nonsymmetric eigenproblem

2.14.6 Balancing and condition for the nonsymmetric eigenproblem

There are two preprocessing steps one may perform on a matrix A in order to make its eigenproblem easier. The first is *permutation*, or reordering the rows and columns to make A more nearly upper triangular (closer to Schur form): $A' = PAP^T$, where P is a permutation matrix. If A' is permutable to upper triangular form (or close to it), then no floating-point operations (or very few) are needed to reduce it to Schur form. The second is *scaling* by a diagonal matrix D to make the rows and columns of A' more nearly equal in norm: $A'' = DA'D^{-1}$. Scaling can make the matrix norm smaller with respect to the eigenvalues, and so possibly reduce the inaccuracy contributed by roundoff (see Chapter, II/11 of Wilkinson and Reinsch (1971)). We refer to these two operations as *balancing*.

Permuting has no effect on the condition numbers or their interpretation as described previously. Scaling, however, does change their interpretation and further details can be found in Anderson *et al.* (1999).

2.14.7 The generalized nonsymmetric eigenvalue problem

The algorithms for the generalized nonsymmetric eigenvalue problem are normwise backward stable: they compute the exact eigenvalues (as the pairs (α, β)), eigenvectors and deflating subspaces of slightly perturbed pairs $(A + E, B + F)$, where

$$\|(E, F)\|_F \leq p(n)\epsilon\|(A, B)\|_F.$$

2.14.8 Balancing the generalized eigenvalue problem

As with the standard nonsymmetric eigenvalue problem, there are two preprocessing steps one may perform on a matrix pair (A, B) in order to make its eigenproblem easier; permutation and scaling, which together are referred to as balancing, as indicated in the following two steps.

1. The balancing routine first attempts to permute A and B to block upper triangular form by a similarity transformation:

$$PAP^T = F = \begin{pmatrix} F_{11} & F_{12} & F_{13} \\ & F_{22} & F_{23} \\ & & F_{33} \end{pmatrix},$$

$$PBP^T = G = \begin{pmatrix} G_{11} & G_{12} & G_{13} \\ & G_{22} & G_{23} \\ & & G_{33} \end{pmatrix},$$

where P is a permutation matrix, F_{11} , F_{33} , G_{11} and G_{33} are upper triangular. Then the diagonal elements of the matrix (F_{11}, G_{11}) and (G_{33}, H_{33}) are generalized eigenvalues of (A, B) . The rest of the generalized eigenvalues are given by the matrix pair (F_{22}, G_{22}) . Subsequent operations to compute the eigenvalues of (A, B) need only be applied to the matrix (F_{22}, G_{22}) ; this can save a significant amount of work if (F_{22}, G_{22}) is smaller than the original matrix pair (A, B) . If no suitable permutation exists (as is often the case), then there is no gain in efficiency or accuracy.

2. The balancing routine applies a diagonal similarity transformation to (F, G) , to make the rows and columns of (F_{22}, G_{22}) as close as in norm as possible:

$$DFD^{-1} = \begin{pmatrix} I & & \\ & D_{22} & \\ & & I \end{pmatrix} \begin{pmatrix} F_{11} & F_{12} & F_{13} \\ & F_{22} & F_{23} \\ & & F_{33} \end{pmatrix} \begin{pmatrix} I & & \\ & D_{22}^{-1} & \\ & & I \end{pmatrix},$$

$$DGD^{-1} = \begin{pmatrix} I & & \\ & D_{22} & \\ & & I \end{pmatrix} \begin{pmatrix} G_{11} & G_{12} & G_{13} \\ & G_{22} & G_{23} \\ & & G_{33} \end{pmatrix} \begin{pmatrix} I & & \\ & D_{22}^{-1} & \\ & & I \end{pmatrix}.$$

This transformation usually improves the accuracy of computed generalized eigenvalues and eigenvectors. However, there are exceptional occasions when this transformation increases the norm of the pencil; in this case accuracy could be lower with diagonal balancing. See Anderson *et al.* (1999) for further details.

2.14.9 Other problems

Error bounds for other problems such as the generalized linear least-squares problem and generalized singular value decomposition can be found in Anderson *et al.* (1999).

2.15 Block Partitioned Algorithms

A number of the routines in this chapter use what is termed a *block partitioned algorithm*. This means that at each major step of the algorithm a *block* of rows or columns is updated, and much of the computation is performed by matrix-matrix operations on these blocks. The matrix-matrix operations are performed by calls to the Level 3 BLAS (see Chapter F06), which are the key to achieving high performance on many modern computers. In the case of the *QR* algorithm for reducing an upper Hessenberg matrix to Schur form, a multishift strategy is used in order to improve performance. See Golub and Van Loan (1996) or Anderson *et al.* (1999) for more about block partitioned algorithms and the multishift strategy.

The performance of a block partitioned algorithm varies to some extent with the **block size** – that is, the number of rows or columns per block. This is a machine-dependent parameter, which is set to a suitable value when the library is implemented on each range of machines. Users of the library do not normally need to be aware of what value is being used. Different block sizes may be used for different routines. Values in the range 16 to 64 are typical.

On more conventional machines there is often no advantage from using a block partitioned algorithm, and then the routines use an *unblocked* algorithm (effectively a block size of 1), relying solely on calls to the Level 2 BLAS (see Chapter F06 again).

The only situation in which you need some awareness of the block size is when it affects the amount of workspace to be supplied to a particular routine. This is discussed in Section 3.4.3.

3 Recommendations on Choice and Use of Available Routines

Note: please refer to the Users' Note for your implementation to check that a routine is available.

3.1 Available Routines

The tables in the following sub-sections show the routines which are provided for performing different computations on different types of matrices. Each entry in the table gives the NAG routine name and the LAPACK double precision name (see Section 3.2).

Black box (or driver) routines are provided for the solution of most problems. In a number of cases there are *simple drivers*, which just return the solution to the problem, as well as *expert drivers*, which return additional information, such as condition number estimates, and may offer additional facilities such as balancing. The following sub-sections give tables for the driver routines.

3.1.1 Driver routines

3.1.1.1 Linear least-squares problems (LLS)

| Operation | real | complex |
|--|-----------------|-----------------|
| solve LLS using <i>QR</i> or <i>LQ</i> factorization | F08AAF (DGELS) | F08ANF (ZGELS) |
| solve LLS using complete orthogonal factorization | F08BAF (DGELSY) | F08BNF (ZGELSY) |
| solve LLS using SVD | F08KAF (DGELSS) | F08KNF (ZGELSS) |
| solve LLS using divide-and-conquer SVD | F08KCF (DGELSD) | F08KQF (ZGELSD) |

3.1.1.2 Generalized linear least-squares problems (LSE and GLM)

| Operation | real | complex |
|-----------------------------|-----------------|-----------------|
| solve LSE problem using GRQ | F08ZAF (DGGLSE) | F08ZNF (ZGGLSE) |
| solve GLM problem using GQR | F08ZBF (DGGGLM) | F08ZPF (ZGGGLM) |

3.1.1.3 Symmetric eigenvalue problems (SEP)

| Function and storage scheme | real | complex |
|--|---|---|
| simple driver divide-and-conquer driver expert driver RRR driver | F08FAF (DSYEV) F08FCF (DSYEVD) F08FBF (DSYEVX) F08FDF (DSYEVX) | F08FNF (ZHEEV) F08FQF (ZHEEVD) F08FPF (ZHEEVX) F08FRF (ZHEEVR) |
| packed storage simple driver divide-and-conquer driver expert driver | F08GAF (DSPEV) F08GCF (DSPEVD) F08GBF (DSPEVX) | F08GNF (ZHPEV) F08GQF (ZHPEVD) F08GPF (ZHPEVX) |
| band matrix simple driver divide-and-conquer driver expert driver | F08HAF (DSBEV) F08HCF (DSBEVD) F08HBF (DSBEVX) | F08HNF (ZHBEV) F08HQF (ZHBEVD) F08HPF (ZHBEVX) |

| | | |
|--|---|--|
| tridiagonal matrix simple driver divide-and-conquer driver expert driver RRP driver | F08JAF (DSTEV) F08JCF (DSTEVD) F08JBF (DSTEVX) F08JDF (DSTEVX) | |
|--|---|--|

3.1.1.4 Nonsymmetric eigenvalue problem (NEP)

| Function and storage scheme | real | complex |
|--|--|--|
| simple driver for Schur factorization expert driver for Schur factorization simple driver for eigenvalues/vectors expert driver for eigenvalues/vectors | F08PAF (DGEES) F08PBF (DGEESX) F08NAF (DGEEV) F08NBF (DGEEVX) | F08PNF (ZGEES) F08PPF (ZGEESX) F08NNF (ZGEEV) F08NPF (ZGEEVX) |

3.1.1.5 Singular value decomposition (SVD)

| Function and storage scheme | real | complex |
|--|------------------------------------|------------------------------------|
| simple driver divide-and-conquer driver | F08KBF (DGESVD) F08KDF (DGESDD) | F08KPF (ZGESVD) F08KRF (ZGESDD) |

3.1.1.6 Generalized symmetric definite eigenvalue problems (GSEP)

| Function and storage scheme | real | complex |
|--|--|--|
| simple driver divide-and-conquer driver expert driver | F08SAF (DSYGV) F08SCF (DSYGVD) F08SBF (DSYGVX) | F08SNF (ZHEGV) F08SQF (ZHEGVD) F08SPF (ZHEGVX) |
| packed storage simple driver divide-and-conquer driver expert driver | F08TAF (DSPGV) F08TCF (DSPGVD) F08TBF (DSPGVX) | F08TNF (ZHPGV) F08TQF (ZHPGVD) F08TPF (ZHPGVX) |
| band matrix simple driver divide-and-conquer driver expert driver | F08UAF (DSBGV) F08UCF (DSBGVD) F08UBF (DSBGVX) | F08UNF (ZHBGV) F08UQF (ZHBGVD) F08UPF (ZHBGVX) |

3.1.1.7 Generalized nonsymmetric eigenvalue problem (GNEP)

| Function and storage scheme | real | complex |
|--|--|--|
| simple driver for Schur factorization expert driver for Schur factorization simple driver for eigenvalues/vectors expert driver for eigenvalues/vectors | F08XAF (DGGES) F08XBF (DGGESX) F08WAF (DGGEV) F08WBF (DGGEVX) | F08XNF (ZGGES) F08XPF (ZGGESX) F08WNF (ZGGEV) F08WPF (ZGGEVX) |

3.1.1.8 Generalized singular value decomposition (GSVD)

| Function and storage scheme | real | complex |
|-----------------------------|-----------------|-----------------|
| singular values/vectors | F08VAF (DGGSDV) | F08VNF (ZGGSDV) |

3.1.2 Computational routines

It is possible to solve problems by calling two or more routines in sequence. Some common sequences of routines are indicated in the tables in the following sub-sections; an asterisk (*) against a routine name means that the sequence of calls is illustrated in the example program for that routine.

It should be noted that all the LAPACK computational routines from Release 3 are included in the Library and can be called by their LAPACK name*, although not all of these routines are currently documented in Chapters F07 and F08.

3.1.2.1 Orthogonal factorizations

Routines are provided for QR factorization (with and without column pivoting), and for LQ factorization (without pivoting only), of a general real or complex rectangular matrix.

The factorization routines do not form the matrix Q explicitly, but represent it as a product of elementary reflectors (see Section 3.3.6). Additional routines are provided to generate all or part of Q explicitly if it is required, or to apply Q in its factored form to another matrix (specifically to compute one of the matrix products QC , $Q^T C$, CQ or CQ^T with Q^T replaced by Q^H if C and Q are complex).

| | Factorize without pivoting | Factorize with pivoting | Generate Matrix Q | Apply matrix Q |
|--------------------------------------|----------------------------------|----------------------------|------------------------|---------------------|
| QR factorization, real matrices | F08AEF (DGEQRF) | F08BEF (DGEQPF) | F08AFF (DORGQR) | F08AGF (DORMQR) |
| LQ factorization, real matrices | F08AHF (DGELQF) | | F08AJF (DORGLQ) | F08AKF (DORMLQ) |
| QR factorization, complex matrices | F08ASF (ZGEQRF) | F08BSF (ZGEQPF) | F08ATF (ZUNGQR) | F08AUF (ZUNMQR) |
| LQ factorization, complex matrices | F08AVF (ZGELQF) | | F08AWF (ZUNGLQ) | F08AXF (ZUNMLQ) |

To solve linear least-squares problems, as described in Sections 2.2.1 or 2.2.3, routines based on the QR factorization can be used:

| | |
|--------------------------------------|-------------------------|
| real data, full-rank problem | F08AEF*, F08AGF, F06YJF |
| complex data, full-rank problem | F08ASF*, F08AUF, F06ZJF |
| real data, rank-deficient problem | F08BEF*, F08AGF, F06YJF |
| complex data, rank-deficient problem | F08BSF*, F08AUF, F06ZJF |

To find the minimum norm solution of under-determined systems of linear equations, as described in Section 2.2.2, routines based on the LQ factorization can be used:

| | |
|---------------------------------|-------------------------|
| real data, full-rank problem | F08AHF*, F06YJF, F08AKF |
| complex data, full-rank problem | F08AVF*, F06ZJF, F08AXF |

3.1.2.2 Singular value problems

Routines are provided to reduce a general real or complex rectangular matrix A to real bidiagonal form B by an orthogonal transformation $A = QBP^T$ (or by a unitary transformation $A = QBP^H$ if A is complex).

Different routines allow a full matrix A to be stored conventionally (see Section 3.3.1), or a band matrix to use band storage (see Section 3.3.3).

The routines for reducing full matrices do not form the matrix Q or P explicitly; additional routines are provided to generate all or part of them, or to apply them to another matrix, as with the routines for orthogonal factorizations. Explicit generation of Q or P is required before using the bidiagonal QR algorithm to compute left or right singular vectors of A .

The routines for reducing band matrices have options to generate Q or P if required.

Further routines are provided to compute all or part of the singular value decomposition of a real bidiagonal matrix; the same routines can be used to compute the singular value decomposition of a real or complex matrix that has been reduced to bidiagonal form.

| | Reduce to bidiagonal form | Generate matrix Q or P^T | Apply matrix Q or P | Reduce band matrix to bidiagonal form | SVD of bidiagonal form (QR algorithm) |
|------------------|---------------------------------|------------------------------------|-------------------------------|--|---|
| real matrices | F08KEF (DGEHRD) | F08KFF (DORGBR) | F08KGF (DORMBR) | F08LEF (DGBHRD) | F08MEF (DBDSQR) |
| complex matrices | F08KSF (ZGEHRD) | F08KTF (ZUNGBR) | F08KUF (ZUNMBR) | F08LSF (ZGBHRD) | F08MSF (ZBDSQR) |

To compute the singular values and vectors of a rectangular matrix, as described in Section 2.3, use the following sequence of calls:

Rectangular matrix (standard storage)

real matrix, singular values and vectors

F08KEF, F08KFF*, F08MEF

complex matrix, singular values and vectors

F08KSF, F08KTF*, F08MSF

Rectangular matrix (banded)

real matrix, singular values and vectors

F08LEF

complex matrix, singular values and vectors

F08LSF

To use the singular value decomposition to solve a linear least-squares problem, as described in Section 2.4, the following routines are required:

real data:

F08KEF, F08KGF, F08KFF,
F08MEF, F06YAF

complex data:

F08KSF, F08KUF, F08KTF,
F08MSF, F06ZAF

3.1.2.3 Symmetric eigenvalue problems

Routines are provided to reduce a real symmetric or complex Hermitian matrix A to real tridiagonal form T by an orthogonal similarity transformation $A = QTQ^T$ (or by a unitary transformation $A = QTQ^H$ if A is complex). Different routines allow a full matrix A to be stored conventionally (see Section 3.3.1) or in packed storage (see Section 3.3.2); or a band matrix to use band storage (see Section 3.3.3).

The routines for reducing full matrices do not form the matrix Q explicitly; additional routines are provided to generate Q , or to apply it to another matrix, as with the routines for orthogonal factorizations. Explicit generation of Q is required before using the QR algorithm to find all the eigenvectors of A ; application of Q to another matrix is required after eigenvectors of T have been found by inverse iteration, in order to transform them to eigenvectors of A .

The routines for reducing band matrices have an option to generate Q if required.

| | Reduce to tridiagonal form | Generate matrix Q | Apply matrix Q |
|---|----------------------------|---------------------|------------------|
| real symmetric matrices | F08FEF (DSYTRD) | F08FFF (DORGTR) | F08FGF (DORMTR) |
| real symmetric matrices (packed storage) | F08GEF (DSPTRD) | F08GFF (DOPGTR) | F08GGF (DOPMTR) |
| real symmetric band matrices | F08HEF (DSBTRD) | | |
| complex Hermitian matrices | F08FSF (ZHETRD) | F08FTF (ZUNGTR) | F08FUF (ZUNMTR) |
| complex Hermitian matrices (packed storage) | F08GSF (ZHPTRD) | F08GTF (ZUPGTR) | F08GUF (ZUPMTR) |
| complex Hermitian band matrices | F08HSF (ZHBTRD) | | |

A variety of routines are provided to compute eigenvalues and eigenvectors of the real symmetric tridiagonal matrix T , some computing all eigenvalues and eigenvectors, some computing selected eigenvalues and eigenvectors. The same routines can be used to compute eigenvalues and eigenvectors of a real symmetric or complex Hermitian matrix which has been reduced to tridiagonal form.

Eigenvalues and eigenvectors of real symmetric tridiagonal matrices:

The original (non-reduced) matrix is Real or Complex Hermitian

| | |
|---|--------|
| all eigenvalues (root-free QR algorithm) | F08JFF |
| all eigenvalues (root-free QR algorithm called by divide-and-conquer) | F08JCF |
| selected eigenvalues (bisection) | F08JJF |

The original (non-reduced) matrix is Real

| | |
|---|--------|
| all eigenvalues and eigenvectors (QR algorithm) | F08JEF |
| all eigenvalues and eigenvectors (divide-and-conquer) | F08JCF |
| all eigenvalues and eigenvectors (positive-definite case) | F08JGF |
| selected eigenvectors (inverse iteration) | F08JKF |

The original (non-reduced) matrix is Complex Hermitian

| | |
|---|--------|
| all eigenvalues and eigenvectors (QR algorithm) | F08JSF |
| all eigenvalues and eigenvectors (positive-definite case) | F08JUF |
| selected eigenvectors (inverse iteration) | F08JXF |

The following sequences of calls may be used to compute various combinations of eigenvalues and eigenvectors, as described in Section 2.7.

Sequences for computing eigenvalues and eigenvectors

Real Symmetric matrix (standard storage)

| | |
|---|---------------------------------|
| all eigenvalues and eigenvectors (using divide-and-conquer) | F08FCF |
| all eigenvalues and eigenvectors (using QR algorithm) | F08FEF, F08FFF*, F08JEF |
| selected eigenvalues and eigenvectors (bisection and inverse iteration) | F08FEF, F08JJF, F08JKF, F08FGF* |

Real Symmetric matrix (packed storage)

| | |
|---|---------------------------------|
| all eigenvalues and eigenvectors (using divide-and-conquer) | F08GCF |
| all eigenvalues and eigenvectors (using <i>QR</i> algorithm) | F08GEF, F08GFF*, F08JEF |
| selected eigenvalues and eigenvectors (bisection and inverse iteration) | F08GEF, F08JJF, F08JKF, F08GGF* |

Real Symmetric banded matrix

| | |
|--|-----------------|
| all eigenvalues and eigenvectors (using divide-and-conquer) | F08HCF |
| all eigenvalues and eigenvectors (using <i>QR</i> algorithm) | F08HEF*, F08JEF |

Complex Hermitian matrix (standard storage)

| | |
|---|---------------------------------|
| all eigenvalues and eigenvectors (using divide-and-conquer) | F08FQF |
| all eigenvalues and eigenvectors (using <i>QR</i> algorithm) | F08FSF, F08FTF*, F08JSF |
| selected eigenvalues and eigenvectors (bisection and inverse iteration) | F08FSF, F08JJF, F08JXF, F08FUF* |

Complex Hermitian matrix (packed storage)

| | |
|---|---------------------------------|
| all eigenvalues and eigenvectors (using divide-and-conquer) | F08GQF |
| all eigenvalues and eigenvectors (using <i>QR</i> algorithm) | F08GSF, F08GTF*, F08JSF |
| selected eigenvalues and eigenvectors (bisection and inverse iteration) | F08GSF, F08JJF, F08JXF, F08GUF* |

Complex Hermitian banded matrix

| | |
|--|-----------------|
| all eigenvalues and eigenvectors (using divide-and-conquer) | F08HQF |
| all eigenvalues and eigenvectors (using <i>QR</i> algorithm) | F08HSF*, F08JSF |

3.1.2.4 Generalized symmetric-definite eigenvalue problems

Routines are provided for reducing each of the problems $Ax = \lambda Bx$, $ABx = \lambda x$ or $BAX = \lambda x$ to an equivalent standard eigenvalue problem $Cy = \lambda y$. Different routines allow the matrices to be stored either conventionally or in packed storage. The positive-definite matrix B must first be factorized using a routine from Chapter F07. There is also a routine which reduces the problem $Ax = \lambda Bx$ where A and B are banded, to an equivalent banded standard eigenvalue problem; this uses a split Cholesky factorization for which a routine in Chapter F08 is provided.

| | Reduce to standard problem | Reduce to standard problem (packed storage) | Reduce to standard problem (band matrices) |
|----------------------------|----------------------------|---|--|
| real symmetric matrices | F08SEF (DSYGST) | F08TEF (DSPGST) | F08UEF (DSBGST) |
| complex Hermitian matrices | F08SSF (ZHEGST) | F08TSF (ZHPGST) | F08USF (ZHBGST) |

The equivalent standard problem can then be solved using the routines discussed in Section 3.1.2.3. For example, to compute all the eigenvalues, the following routines must be called:

| | |
|--|----------------------------------|
| real symmetric-definite problem | F07FDF, F08SEF*, F08FEF, F08JFF |
| real symmetric-definite problem, packed storage | F07GDF, F08TEF*, F08GEF, F08JFF |
| real symmetric-definite banded problem | F08UFF*, F08UEF*, F08HEF, F08JFF |
| complex Hermitian-definite problem | F07FRF, F08SSF*, F08FSF, F08JFF |
| complex Hermitian-definite problem, packed storage | F07GRF, F08TSF*, F08GSF, F08JFF |
| complex Hermitian-definite banded problem | F08UTF*, F08USF*, F08HSF, F08JFF |

If eigenvectors are computed, the eigenvectors of the equivalent standard problem must be transformed back to those of the original generalized problem, as indicated in Section 2.8; routines from Chapter F06 may be used for this.

3.1.2.5 Nonsymmetric eigenvalue problems

Routines are provided to reduce a general real or complex matrix A to upper Hessenberg form H by an orthogonal similarity transformation $A = QHQ^T$ (or by a unitary transformation $A = QHQ^H$ if A is complex).

These routines do not form the matrix Q explicitly; additional routines are provided to generate Q , or to apply it to another matrix, as with the routines for orthogonal factorizations. Explicit generation of Q is required before using the QR algorithm on H to compute the Schur vectors; application of Q to another matrix is needed after eigenvectors of H have been computed by inverse iteration, in order to transform them to eigenvectors of A .

Routines are also provided to balance the matrix before reducing it to Hessenberg form, as described in Section 2.14.6. Companion routines are required to transform Schur vectors or eigenvectors of the balanced matrix to those of the original matrix.

| | Reduce to Hessenbergform | Generate matrix Q | Apply matrix Q | Balance | Backtransform vectors after balancing |
|------------------|-----------------------------|---------------------|--------------------|--------------------|---|
| real matrices | F08NEF (DGEHRD) | F08NFF (DORGHR) | F08NGF (DORMHR) | F08NHF (DGEBAL) | F08NJF (DGEBAK) |
| complex matrices | F08NSF (ZGEHRD) | F08NTF (ZUNGHR) | F08NUF (ZUNMHR) | F08NVF (ZGEBAL) | F08NWF (ZGEBAK) |

Routines are provided to compute the eigenvalues and all or part of the Schur factorization of an upper Hessenberg matrix. Eigenvectors may be computed either from the upper Hessenberg form by inverse iteration, or from the Schur form by back-substitution; these approaches are equally satisfactory for computing individual eigenvectors, but the latter may provide a more accurate basis for a subspace spanned by several eigenvectors.

Additional routines estimate the sensitivities of computed eigenvalues and eigenvectors, as discussed in Section 2.14.5.

| | Eigenvalues and Schur factorization (QR algorithm) | Eigenvectors from Hessenberg form (inverse iteration) | Eigenvectors from Schur factorization | Sensitivities of eigenvalues and eigenvectors |
|------------------|--|--|---|---|
| real matrices | F08PEF (DHSEQR) | F08PKF (DHSEIN) | F08QKF (DTREVC) | F08QLF (DTRSNA) |
| complex matrices | F08PSF (ZHSEQR) | F08PXF (ZHSEIN) | F08QXF (ZTREVC) | F08QYF (ZTRSNA) |

Finally routines are provided for re-ordering the Schur factorization, so that eigenvalues appear in any desired order on the diagonal of the Schur form. The routines F08QFF (DTREXC) and F08QTF (ZTREXC) simply swap two diagonal elements or blocks, and may need to be called repeatedly to achieve a desired order. The routines F08QGF (DTRSEN) and F08QUF (ZTRSEN) perform the whole re-ordering process for the important special case where a specified cluster of eigenvalues is to appear at the top of the Schur form; if the Schur vectors are re-ordered at the same time, they yield an orthonormal basis of the invariant subspace corresponding to the specified cluster of eigenvalues. These routines can also compute the sensitivities of the cluster of eigenvalues and the invariant subspace.

| | Reorder Schur factorization | Reorder Schur factorization, find basis of invariant subspace and estimate sensitivities |
|---------------|-----------------------------|--|
| real matrices | F08QFF (DTREXC) | F08QGF (DTRSEN) |

| | | |
|------------------|-----------------|-----------------|
| complex matrices | F08QTF (ZTREXC) | F08QUF (ZTRSEN) |
|------------------|-----------------|-----------------|

The following sequences of calls may be used to compute various combinations of eigenvalues, Schur vectors and eigenvectors, as described in Section 2.11:

| | |
|---|---|
| real matrix, all eigenvalues and Schur factorization | F08NEF, F08NFF*, F08PEF |
| real matrix, all eigenvalues and selected eigenvectors | F08NEF, F08PEF, F08PKF, F08NGF |
| real matrix, all eigenvalues and eigenvectors (with balancing) | F08NHF*, F08NEF, F08NFF, F08PEF, F08PKF, F08NJF |
| complex matrix, all eigenvalues and Schur factorization | F08NSF, F08NTF*, F08PSF |
| complex matrix, all eigenvalues and selected eigenvectors | F08NSF, F08PSF, F08PXF, F08NUF* |
| complex matrix, all eigenvalues and eigenvectors (with balancing) | F08NVF*, F08NSF, F08NTF, F08PSF, F08PXF, F08NWF |

3.1.2.6 Generalized nonsymmetric eigenvalue problems

Routines are provided to reduce a real or complex matrix pair (A_1, R_1) , where A_1 is general and R_1 is upper triangular, to generalized upper Hessenberg form by orthogonal transformations $A_1 = Q_1 H Z_1^T$, $R_1 = Q_1 R Z_1^T$, (or by unitary transformations $A_1 = Q_1 H Z_1^H$, $R = Q_1 R_1 Z_1^H$, in the complex case). These routines can optionally return Q_1 and/or Z_1 . Note that to transform a general matrix pair (A, B) to the form (A_1, R_1) a QR factorization of B ($B = \tilde{Q} R_1$) should first be performed and the matrix A_1 obtained as $A_1 = \tilde{Q}^T A$ (see Section 3.1.2.1 above).

Routines are also provided to balance a general matrix pair before reducing it to generalized Hessenberg form, as described in Section 2.14.8. Companion routines are provided to transform vectors of the balanced pair to those of the original matrix pair.

| | Reduce to generalized Hessenberg form | Balance | Backtransform vectors after balancing |
|------------------|---------------------------------------|-----------------|---------------------------------------|
| real matrices | F08WEF (DGGHRD) | F08WHF (DGGBAL) | F08WJF (DGGBAK) |
| complex matrices | F08WSF (ZGGHRD) | F08WVF (ZGGBAL) | F08WWF (ZGGBAK) |

Routines are provided to compute the eigenvalues (as the pairs (α, β)) and all or part of the generalized Schur factorization of a generalized upper Hessenberg matrix pair. Eigenvectors may be computed from the generalized Schur form by back-substitution.

| | Eigenvalues and generalized Schur factorization (QZ algorithm) | Eigenvectors from generalized Schur factorization |
|------------------|---|---|
| real matrices | F08XEF (DHGEQZ) | F08YKF (DTGEVC) |
| complex matrices | F08XSF (ZHGEQZ) | F08YXF (ZTGEVC) |

The following sequences of calls may be used to compute various combinations of eigenvalues, generalized Schur vectors and eigenvectors

| | |
|---|---|
| real matrix pair, all eigenvalues (with balancing) | F08WHF, F08AEF, F08AGF, F08WEF, F08XEF* |
| real matrix pair, all eigenvalues and generalized Schur factorization | F08AEF, F08AGF, F08AFF, F08WEF, F08XEF |
| real matrix pair, all eigenvalues and eigenvectors (with balancing) | F08WHF, F08AEF, F08AGF, F06QHF, F06QFF, F08AFF, F08WEF, F08XEF, F08YKF*, F08WJF |

| | |
|--|---|
| complex matrix pair, all eigenvalues (with balancing) | F08WVF, F08ASF, F08AUF, F08WSF, F08XSF* |
| complex matrix pair, all eigenvalues and generalized Schur factorization | F08ASF, F08AUF, F08ATF, F08WSF, F08XSF |
| complex matrix pair, all eigenvalues and eigenvectors (with balancing) | F08WVF, F08ASF, F08AUF, F06THF, F06TFF, F08ATF, F08WSF, F08XSF, F08YXF*, F08WWF |

3.1.2.7 Sylvester's equation

Routines are provided to solve the real or complex Sylvester equation $AX \pm XB = C$, where A and B are upper quasi-triangular if real, or upper triangular if complex. To solve the general form of Sylvester's equation in which A and B are general square matrices, A and B must be reduced to upper (quasi-) triangular form by the Schur factorization, using routines described in Section 3.1.2.5. For more details, see the documents for the routines listed below.

| | solve Sylvester's equation |
|------------------|-----------------------------------|
| real matrices | F08QHF (DTRSYL) |
| complex matrices | F08QVF (ZTRSYL) |

3.2 NAG Names and LAPACK Names

As well as the NAG routine name (beginning F08-), the tables in Section 3.1 show the LAPACK routine names in double precision.

The routines may be called either by their NAG names or by their LAPACK names. When using a single precision implementation of the NAG Library, the single precision form of the LAPACK name must be used (beginning with S- or C-); when using a double precision implementation of the NAG Library, the double precision form of the LAPACK name must be used (beginning with D- or Z-).

References to Chapter F08 routines in the manual normally include the LAPACK double precision names, for example F08AEF (DGEQRF). The LAPACK routine names follow a simple scheme (which is similar to that used for the BLAS in Chapter F06). Each name has the structure **XYZZZ**, where the components have the following meanings:

- the initial letter **X** indicates the data type (real or complex) and precision:
 - S – real, single precision (in Fortran 77, REAL)
 - D – real, double precision (in Fortran 77, DOUBLE PRECISION)
 - C – complex, single precision (in Fortran 77, COMPLEX)
 - Z – complex, double precision (in Fortran 77, COMPLEX*16 or DOUBLE COMPLEX)
- the 2nd and 3rd letters **YY** indicate the type of the matrix A or matrix pair (A, B) (and in some cases the storage scheme):
 - BD – bidiagonal
 - DI – diagonal
 - GB – general band
 - GE – general
 - GG – general pair (B may be triangular)
 - HG – generalized upper Hessenberg
 - HS – upper Hessenberg
 - OP – (real) orthogonal (packed storage)
 - UP – (complex) unitary (packed storage)
 - OR – (real) orthogonal
 - UN – (complex) unitary
 - PT – symmetric or Hermitian positive-definite tridiagonal
 - SB – (real) symmetric band

HB – (complex) Hermitian band
 SP – symmetric (packed storage)
 HP – Hermitian (packed storage)
 ST – (real) symmetric tridiagonal
 SY – symmetric
 HE – Hermitian
 TG – triangular pair (one may be quasi-triangular)
 TR – triangular (or quasi-triangular)

- the last 3 letters **ZZZ** indicate the computation performed. For example, QRF is a QR factorization.

Thus the routine SGEQRF performs a QR factorization of a real general matrix in a single precision implementation of the Library ; the corresponding routine in a double precision implementation is DGEQRF.

Some sections of the routine documents – Section 2 (Specification) and Section 9.1 (Example program) – print the LAPACK name in ***bold italics***, according to the NAG convention of using bold italics for precision-dependent terms – for example, ***dgeqrf***, which should be interpreted as either SGEQRF (in single precision) or DGEQRF (in double precision).

3.3 Matrix Storage Schemes

In this chapter the following storage schemes are used for matrices:

- conventional storage in a two-dimensional array;
- packed storage for symmetric or Hermitian matrices;
- packed storage for orthogonal or unitary matrices;
- band storage for general, symmetric or Hermitian band matrices;
- storage of bidiagonal, symmetric or Hermitian tridiagonal matrices in two one-dimensional arrays.

These storage schemes are compatible with those used in Chapters F06 and F07, but different schemes for packed, band and tridiagonal storage are used in a few older routines in Chapters F01, F02, F03 and F04.

In the examples below, * indicates an array element which need not be set and is not referenced by the routines. The examples illustrate only the relevant leading rows and columns of the arrays ; array arguments may of course have additional rows or columns, according to the usual rules for passing array arguments in Fortran 77.

3.3.1 Conventional storage

The default scheme for storing matrices is the obvious one: a matrix A is stored in a two-dimensional array A , with matrix element a_{ij} stored in array element $A(i, j)$.

If a matrix is *triangular* (upper or lower, as specified by the argument UPLO when present), only the elements of the relevant triangle are stored; the remaining elements of the array need not be set. Such elements are indicated by * in the examples below. For example, when $n = 4$:

| UPLO | Triangular matrix A | Storage in array A |
|------------|--|--|
| UPLO = 'U' | $\begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ & a_{22} & a_{23} & a_{24} \\ & & a_{33} & a_{34} \\ & & & a_{44} \end{pmatrix}$ | $\begin{matrix} a_{11} & a_{12} & a_{13} & a_{14} \\ * & a_{22} & a_{23} & a_{24} \\ * & * & a_{33} & a_{34} \\ * & * & * & a_{44} \end{matrix}$ |
| UPLO = 'L' | $\begin{pmatrix} a_{11} & & & \\ a_{21} & a_{22} & & \\ a_{31} & a_{32} & a_{33} & \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}$ | $\begin{matrix} a_{11} & * & * & * \\ a_{21} & a_{22} & * & * \\ a_{31} & a_{32} & a_{33} & * \\ a_{41} & a_{42} & a_{43} & a_{44} \end{matrix}$ |

Similarly, if the matrix is upper Hessenberg, or if the matrix is quasi-upper triangular, elements below the first sub-diagonal need not be set.

Routines that handle *symmetric* or *Hermitian* matrices allow for either the upper or lower triangle of the matrix (as specified by UPLO) to be stored in the corresponding elements of the array; the remaining elements of the array need not be set. For example, when $n = 4$:

| UPLO | Hermitian matrix A | Storage in array A |
|------------|--|--|
| UPLO = 'U' | $\begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ \bar{a}_{12} & a_{22} & a_{23} & a_{24} \\ \bar{a}_{13} & \bar{a}_{23} & a_{33} & a_{34} \\ \bar{a}_{14} & \bar{a}_{24} & \bar{a}_{34} & a_{44} \end{pmatrix}$ | $\begin{matrix} a_{11} & a_{12} & a_{13} & a_{14} \\ * & a_{22} & a_{23} & a_{24} \\ * & * & a_{33} & a_{34} \\ * & * & * & a_{44} \end{matrix}$ |
| UPLO = 'L' | $\begin{pmatrix} a_{11} & \bar{a}_{21} & \bar{a}_{31} & \bar{a}_{41} \\ a_{21} & a_{22} & \bar{a}_{32} & \bar{a}_{42} \\ a_{31} & a_{32} & a_{33} & \bar{a}_{43} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}$ | $\begin{matrix} a_{11} & * & * & * \\ a_{21} & a_{22} & * & * \\ a_{31} & a_{32} & a_{33} & * \\ a_{41} & a_{42} & a_{43} & a_{44} \end{matrix}$ |

3.3.2 Packed storage

Symmetric and Hermitian matrices may be stored more compactly, if the relevant triangle (again as specified by UPLO) is packed *by columns* in a one-dimensional array. In Chapters F07 and F08, arrays that hold matrices in packed storage, have argument names ending in 'P'. So:

- if UPLO = 'U', a_{ij} is stored in $AP(i + j(j-1)/2)$ for $i \leq j$;
- if UPLO = 'L', a_{ij} is stored in $AP(i + (2n-j)(j-1)/2)$ for $j \leq i$.

For example:

| UPLO | Triangle of matrix A | Packed storage in array AP |
|------------|--|--|
| UPLO = 'U' | $\begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ & a_{22} & a_{23} & a_{24} \\ & & a_{33} & a_{34} \\ & & & a_{44} \end{pmatrix}$ | $a_{11} \underbrace{a_{12}a_{22}} \underbrace{a_{13}a_{23}a_{33}} \underbrace{a_{14}a_{24}a_{34}a_{44}}$ |
| UPLO = 'L' | $\begin{pmatrix} a_{11} & & & \\ a_{21} & a_{22} & & \\ a_{31} & a_{32} & a_{33} & \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}$ | $\underbrace{a_{11}a_{21}a_{31}a_{41}} \underbrace{a_{22}a_{32}a_{42}} \underbrace{a_{33}a_{43}} a_{44}$ |

Note that for symmetric matrices, packing the upper triangle by columns is equivalent to packing the lower triangle by rows; packing the lower triangle by columns is equivalent to packing the upper triangle by rows. For Hermitian matrices, packing the upper triangle by columns is equivalent to packing the conjugate of the lower triangle by rows; packing the lower triangle by columns is equivalent to packing the conjugate of the upper triangle by rows.

3.3.3 Band storage

A general m by n band matrix with k_l sub-diagonals and k_u super-diagonals may be stored compactly in a two-dimensional array with $k_l + k_u + 1$ rows and n columns. Columns of the matrix are stored in corresponding columns of the array, and diagonals of the matrix are stored in rows of the array. This storage scheme should be used in practice only if $k_l, k_u \ll n$, although routines in Chapters F07 and F08 work correctly for all values of k_l and k_u . In Chapters F07 and F08, arrays that hold matrices in band storage have argument names ending in 'B'. So:

$$a_{ij} \text{ is stored in } AB(k_u + 1 + i - j, j) \text{ for } \max(1, j - k_u) \leq i \leq \min(m, j + k_l).$$

For example, when $m = 6$, $n = 5$, $k_l = 2$ and $k_u = 1$:

| general band matrix A | Band storage in array AB |
|--|--|
| $\begin{pmatrix} a_{11} & a_{12} & & & \\ a_{21} & a_{22} & a_{23} & & \\ a_{31} & a_{32} & a_{33} & a_{34} & \\ & a_{42} & a_{43} & a_{44} & a_{45} \\ & & a_{53} & a_{54} & a_{55} \\ & & & a_{64} & a_{65} \end{pmatrix}$ | $\begin{matrix} * & a_{12} & a_{23} & a_{34} & a_{45} \\ a_{11} & a_{22} & a_{33} & a_{44} & a_{55} \\ a_{21} & a_{32} & a_{43} & a_{54} & a_{65} \\ a_{31} & a_{42} & a_{53} & a_{64} & * \end{matrix}$ |

A symmetric or Hermitian band matrix with k sub-diagonals and super-diagonals may be stored more compactly in a two-dimensional array with $k + 1$ rows and n columns. Only the upper or lower triangle (as specified by UPLO) need to be stored. So:

$$\text{if UPLO} = 'U', a_{ij} \text{ is stored in } AB(k + 1 + i - j, j) \text{ for } \max(1, j - k) \leq i \leq j;$$

$$\text{if UPLO} = 'L', a_{ij} \text{ is stored in } AB(1 + i - j, j) \text{ for } j \leq i \leq \min(n, j + k).$$

For example, when $n = 5$ and $k = 2$:

| UPLO | Hermitian band matrix A | Band storage in array AB |
|------------|--|---|
| UPLO = 'U' | $\begin{pmatrix} a_{11} & a_{12} & a_{13} & & \\ \bar{a}_{12} & a_{22} & a_{23} & a_{24} & \\ \bar{a}_{13} & \bar{a}_{23} & a_{33} & a_{34} & a_{35} \\ & \bar{a}_{24} & \bar{a}_{34} & a_{44} & a_{45} \\ & & \bar{a}_{35} & \bar{a}_{45} & a_{55} \end{pmatrix}$ | $\begin{matrix} * & * & a_{13} & a_{24} & a_{35} \\ * & a_{12} & a_{23} & a_{34} & a_{45} \\ a_{11} & a_{22} & a_{33} & a_{44} & a_{55} \end{matrix}$ |
| UPLO = 'L' | $\begin{pmatrix} a_{11} & \bar{a}_{21} & \bar{a}_{31} & & \\ a_{21} & a_{22} & \bar{a}_{32} & \bar{a}_{42} & \\ a_{31} & a_{32} & a_{33} & \bar{a}_{43} & \bar{a}_{53} \\ & a_{42} & a_{43} & a_{44} & \bar{a}_{54} \\ & & a_{53} & a_{54} & a_{55} \end{pmatrix}$ | $\begin{matrix} a_{11} & a_{22} & a_{33} & a_{44} & a_{55} \\ a_{21} & a_{32} & a_{43} & a_{54} & * \\ a_{31} & a_{42} & a_{53} & * & * \end{matrix}$ |

3.3.4 Tridiagonal and bidiagonal matrices

A symmetric tridiagonal or bidiagonal matrix is stored in two one-dimensional arrays, one of length n containing the diagonal elements, and one of length $n - 1$ containing the off-diagonal elements. (Older routines in Chapter F02 store the off-diagonal elements in elements 2 : n of a vector of length n .)

3.3.5 Real diagonal elements of complex matrices

Complex Hermitian matrices have diagonal matrices that are by definition purely real. In addition, some complex triangular matrices computed by F08 routines are defined by the algorithm to have real diagonal elements – in QR factorization, for example.

If such matrices are supplied as input to F08 routines, the imaginary parts of the diagonal elements are not referenced, but are assumed to be zero. If such matrices are returned as output by F08 routines, the computed imaginary parts are explicitly set to zero.

3.3.6 Representation of orthogonal or unitary matrices

A real orthogonal or complex unitary matrix (usually denoted Q) is often represented in the NAG Library as a product of *elementary reflectors* – also referred to as *elementary Householder matrices* (usually denoted H_i). For example,

$$Q = H_1 H_2 \cdots H_k.$$

Most users need not be aware of the details, because routines are provided to work with this representation, either to generate all or part of Q explicitly, or to multiply a given matrix by Q or Q^T (Q^H in the complex case) without forming Q explicitly.

Nevertheless, the following further details may occasionally be useful.

An elementary reflector (or elementary Householder matrix) H of order n is a unitary matrix of the form

$$H = I - \tau v v^H \quad (4)$$

where τ is a scalar, and v is an n element vector, with $|\tau|^2 \|v\|_2^2 = 2 \times \text{Re}(\tau)$; v is often referred to as the *Householder vector*. Often v has several leading or trailing zero elements, but for the purpose of this discussion assume that H has no such special structure.

There is some redundancy in the representation (4), which can be removed in various ways. The representation used in Chapter F08 and in LAPACK (which differs from those used in some of the routines in Chapters F01, F02, F04 and F06) sets $v_1 = 1$; hence v_1 need not be stored. In real arithmetic, $1 \leq \tau \leq 2$, except that $\tau = 0$ implies $H = I$.

In complex arithmetic, τ may be complex, and satisfies $1 \leq \text{Re}(\tau) \leq 2$ and $|\tau - 1| \leq 1$. Thus a complex H is not Hermitian (as it is in other representations), but it is unitary, which is the important property. The advantage of allowing τ to be complex is that, given an arbitrary complex vector x , Hx can be computed so that

$$H^H x = \beta(1, 0, \dots, 0)^T$$

with *real* β . This is useful, for example, when reducing a complex Hermitian matrix to real symmetric tridiagonal form, or a complex rectangular matrix to real bidiagonal form.

3.4 Parameter Conventions

3.4.1 Option parameters

Most routines in this chapter have one or more option parameters, of type CHARACTER. The descriptions in Section 5 of the routine documents refer only to upper case values (for example UPLO = 'U' or UPLO = 'L'); however in every case, the corresponding lower case characters may be supplied (with the same meaning). Any other value is illegal.

A longer character string can be passed as the actual parameter, making the calling program more readable, but only the first character is significant. (This is a feature of Fortran 77.) For example:

```
CALL SSYTRD ('Upper', ...)
```

3.4.2 Problem dimensions

It is permissible for the problem dimensions (for example, M or N) to be passed as zero, in which case the computation (or part of it) is skipped. Negative dimensions are regarded as an error.

3.4.3 Length of work arrays

A number of routines implementing block algorithms require workspace sufficient to hold one block of rows or columns of the matrix if they are to achieve optimum levels of performance – for example, workspace of size $n \times nb$, where nb is the optimum block size. In such cases, the actual declared length of the work array must be passed as a separate argument LWORK, which immediately follows WORK in the argument-list.

The routine will still perform correctly when less workspace is provided: it simply uses the largest block size allowed by the amount of workspace supplied, as long as this is likely to give better performance than the unblocked algorithm. On exit, WORK(1) contains the minimum value of LWORK which would allow the routine to use the optimum block size; this value of LWORK can be used for subsequent runs.

If LWORK indicates that there is insufficient workspace to perform the unblocked algorithm, this is regarded as an illegal value of LWORK, and is treated like any other illegal parameter value (see Section 3.4.4).

If you are in doubt how much workspace to supply and are concerned to achieve optimum performance, supply a generous amount (assume a block size of 64, say), and then examine the value of WORK(1) on exit.

3.4.4 Error-handling and the diagnostic parameter INFO

Routines in this chapter do not use the usual NAG Library error-handling mechanism, involving the parameter IFAIL. Instead they have a diagnostic parameter INFO. (Thus they preserve complete compatibility with the LAPACK specification.)

Whereas IFAIL is an *Input/Output* parameter and must be set before calling a routine, INFO is purely an *Output* parameter and need not be set before entry.

INFO indicates the success or failure of the computation, as follows:

INFO = 0: successful termination;

INFO > 0: failure in the course of computation, control returned to the calling program.

If the routine document specifies that the routine may terminate with INFO > 0, then it is **essential to test INFO on exit** from the routine. (This corresponds to a *soft failure* in terms of the usual NAG error-handling terminology.) No error message is output.

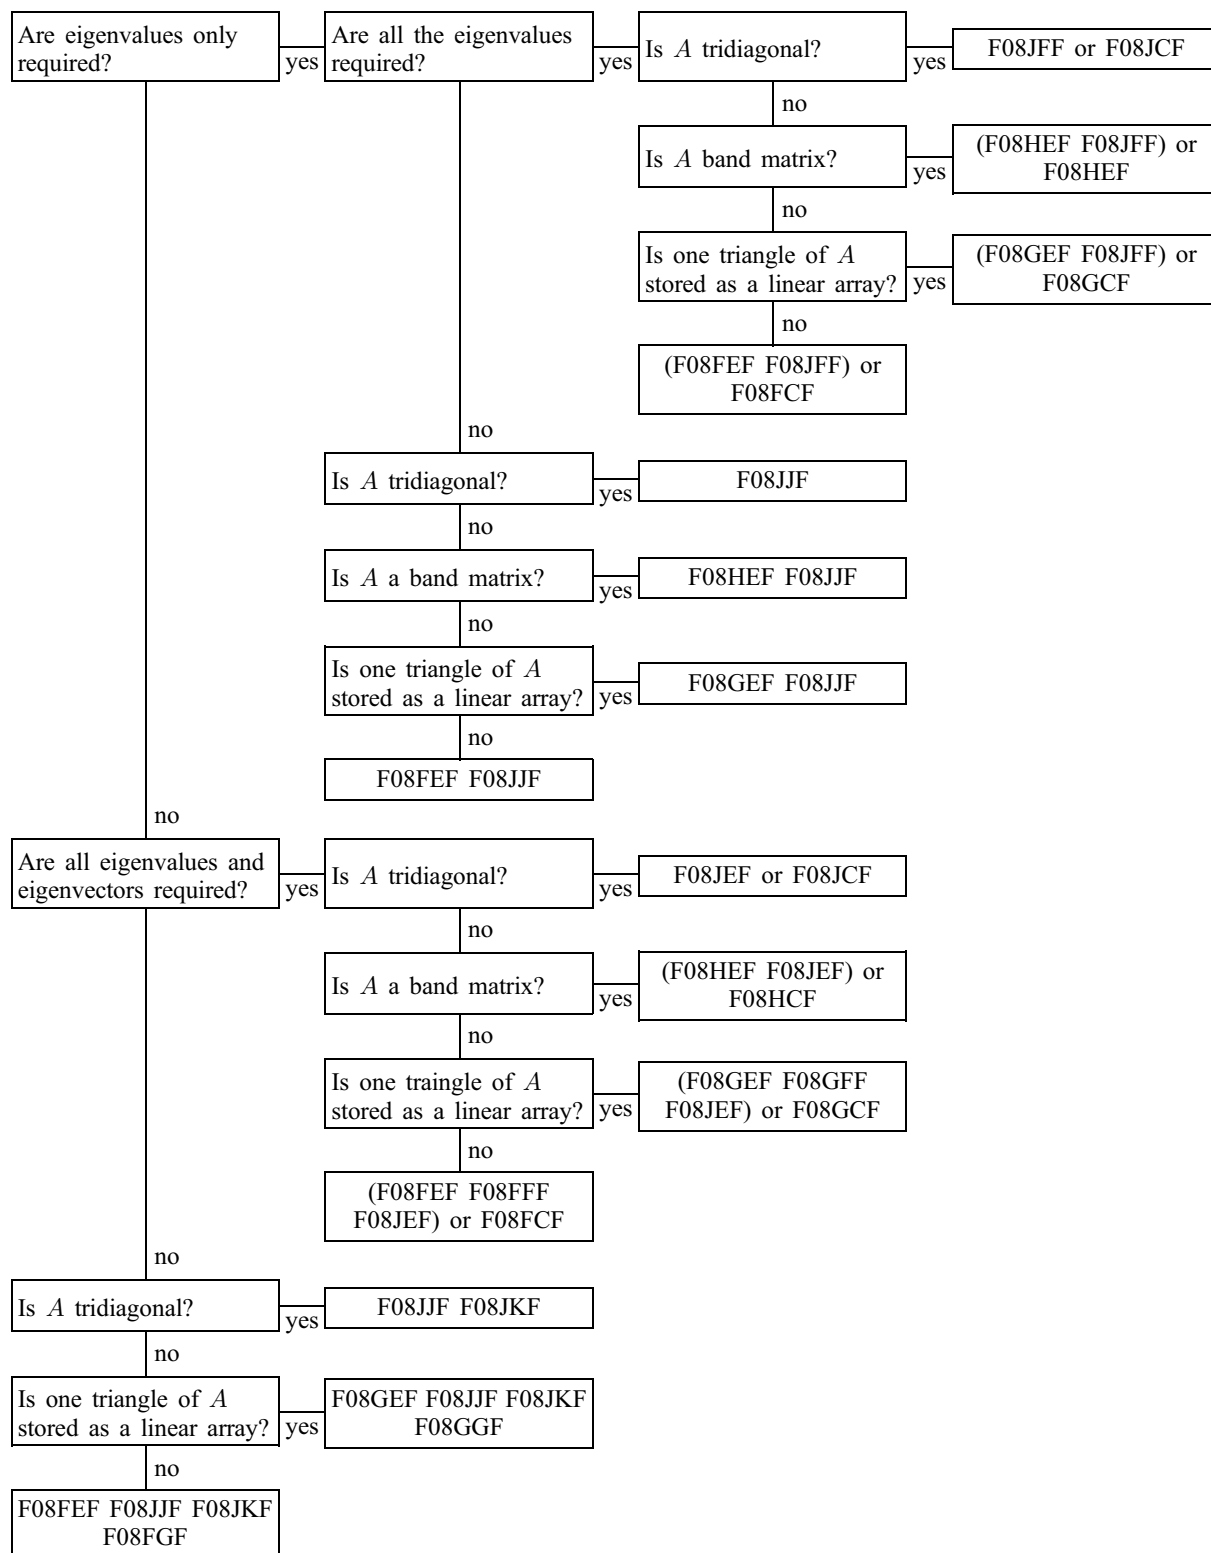
All routines check that input parameters such as N or LDA or option parameters of type CHARACTER have permitted values. If an illegal value of the i th parameter is detected, INFO is set to $-i$, a message is output, and execution of the program is terminated. (This corresponds to a *hard failure* in the usual NAG terminology.) In some implementations, especially when linking to vendor versions of LAPACK, execution of the program may continue, in which case, it is essential to test INFO on exit from the routine.

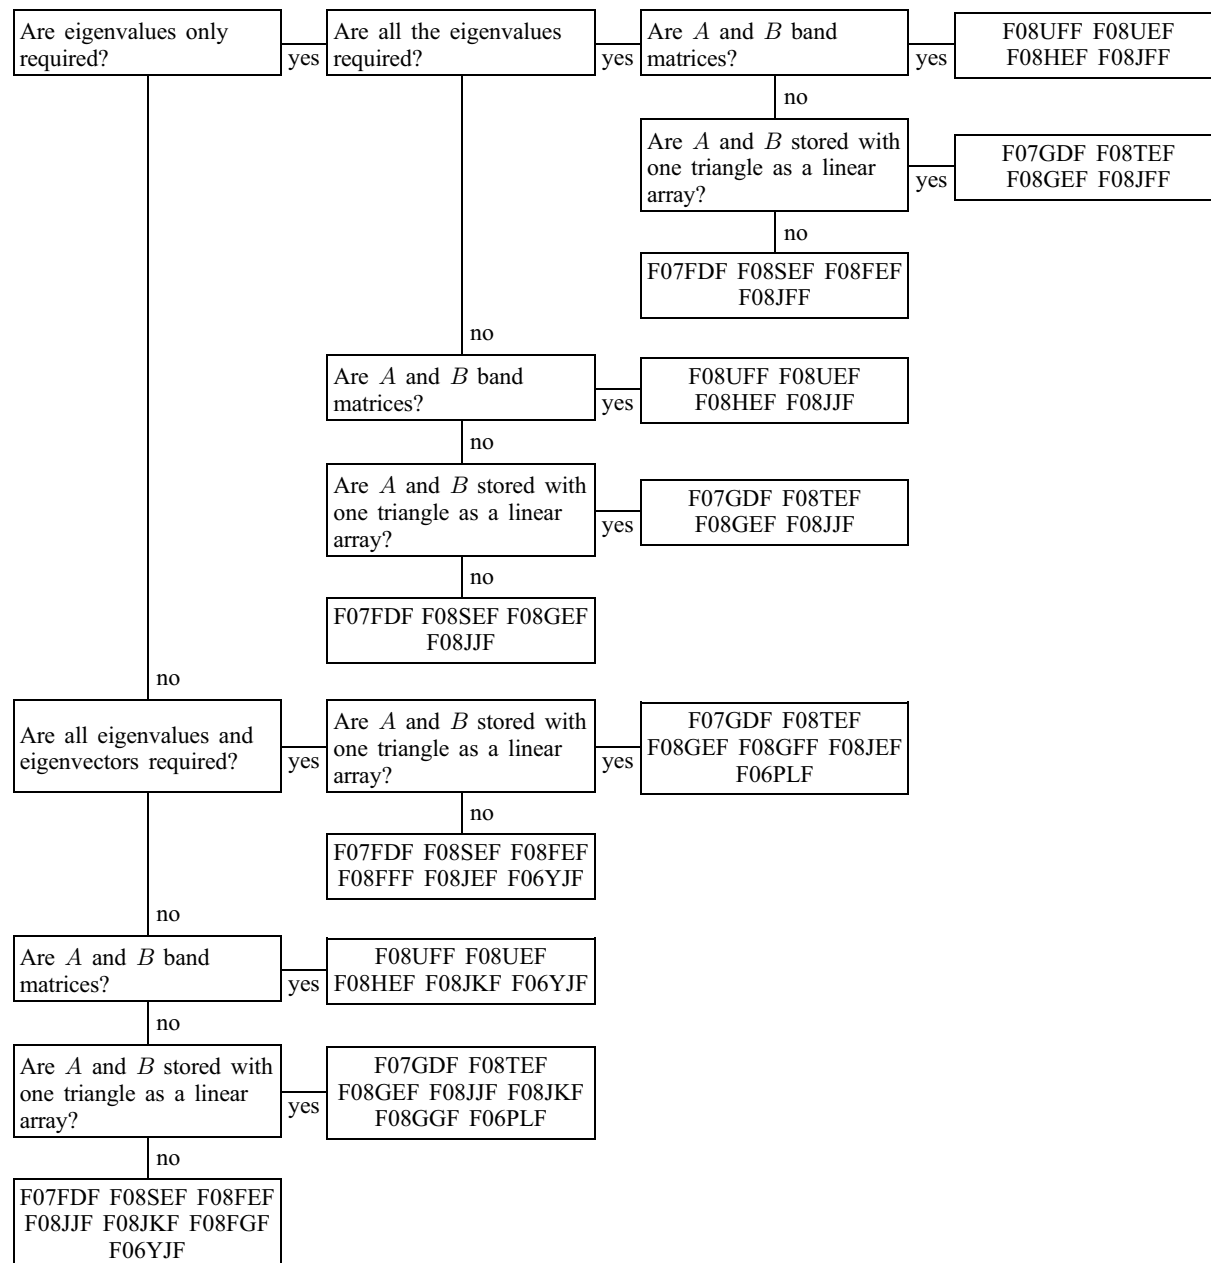
4 Decision Trees

The following decision trees are principally for the computation (general purpose) routines. See Section 3.1.1.1 for tables of the driver (black box) routines.

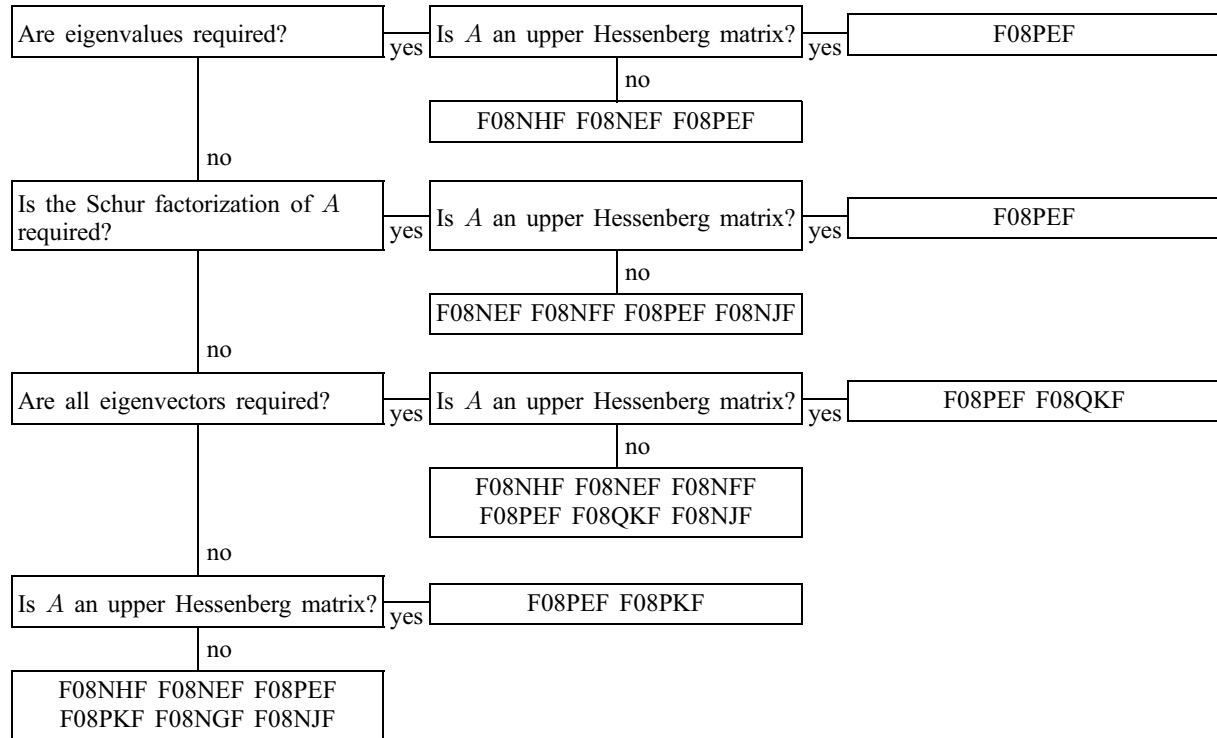
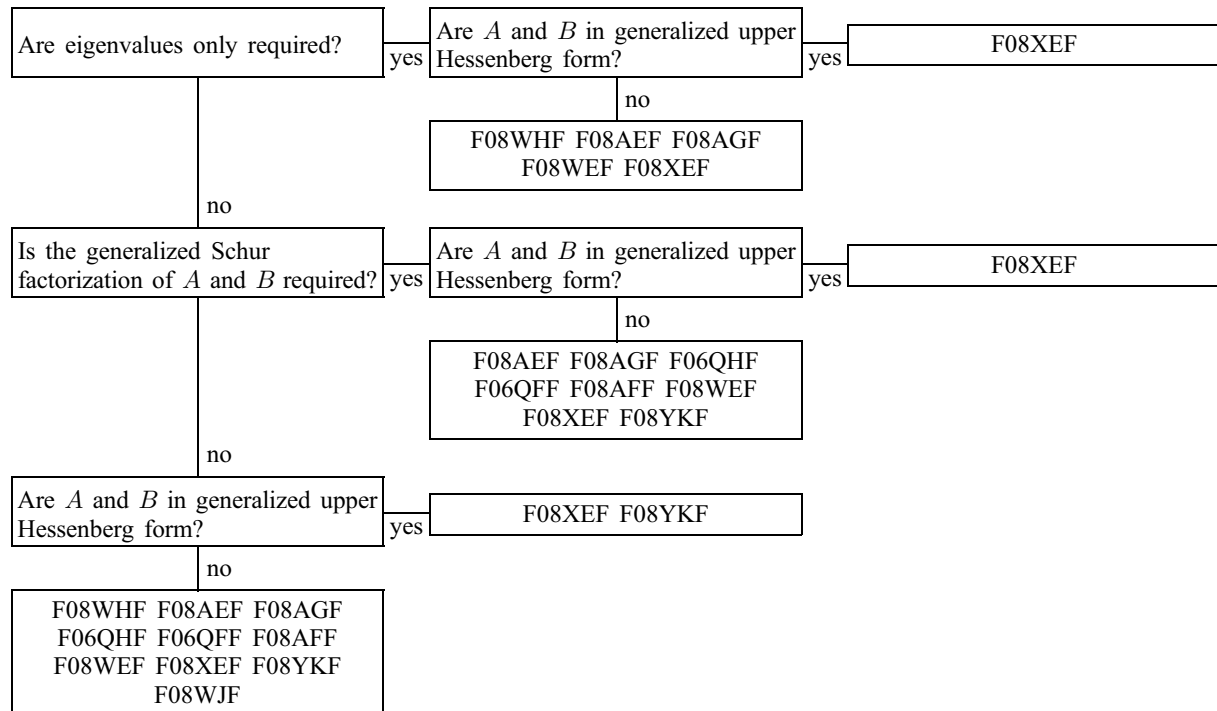
4.1 General Purpose Routines (eigenvalues and eigenvectors)

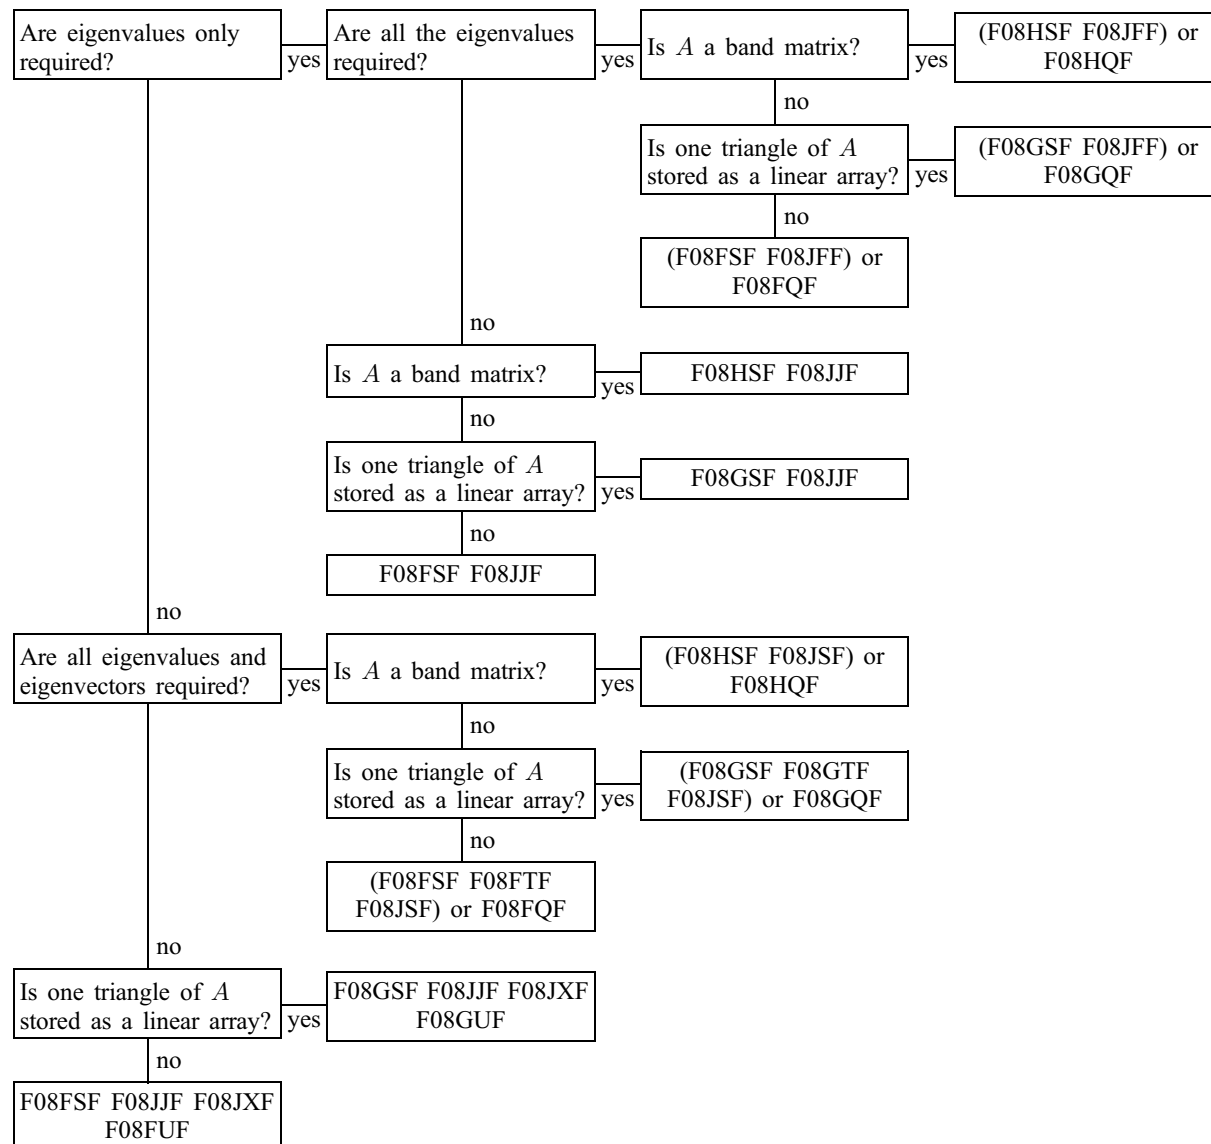
Tree 1: Real Symmetric Eigenvalue Problems

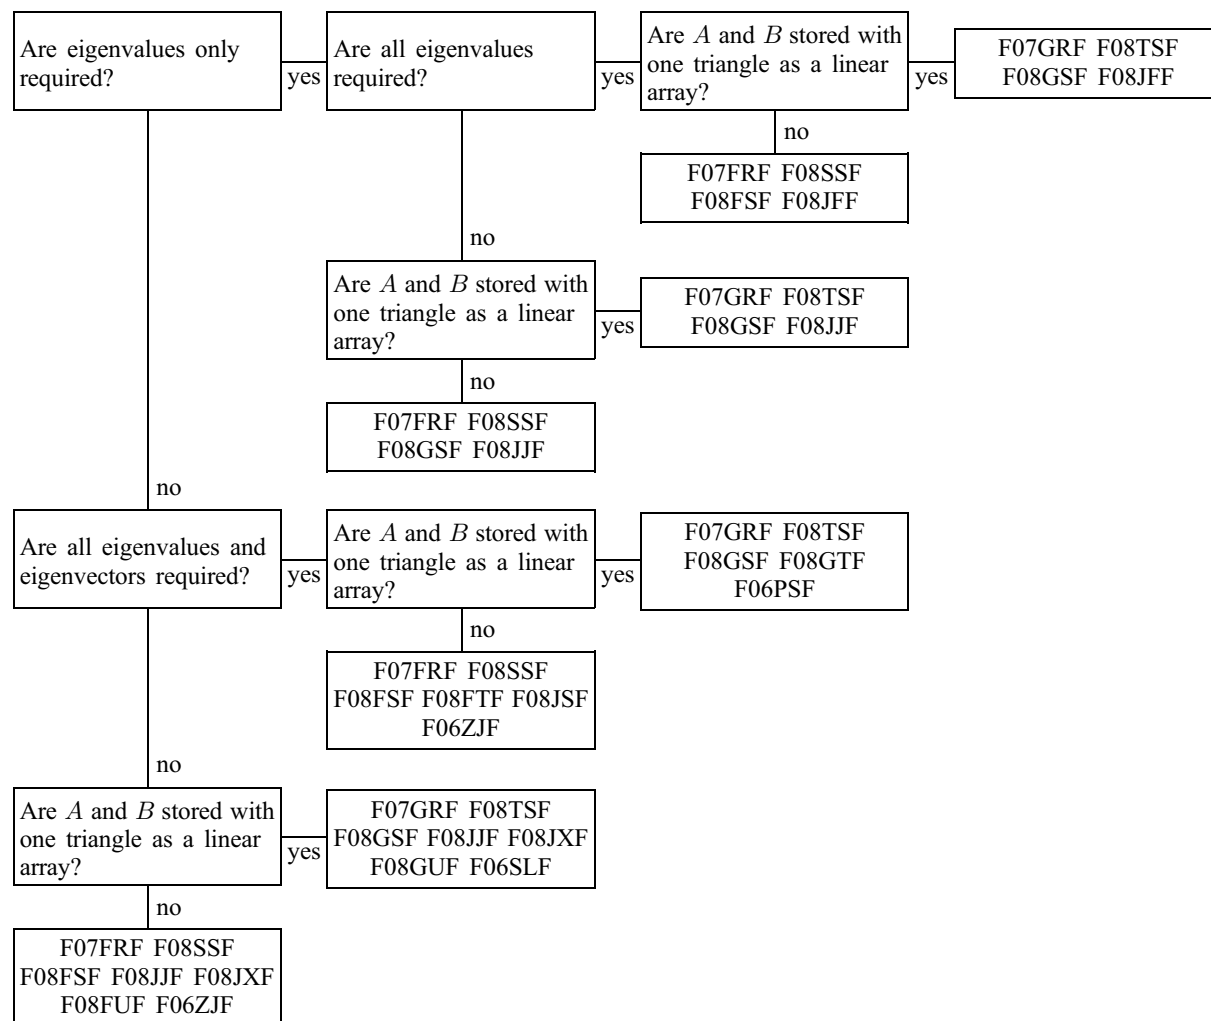


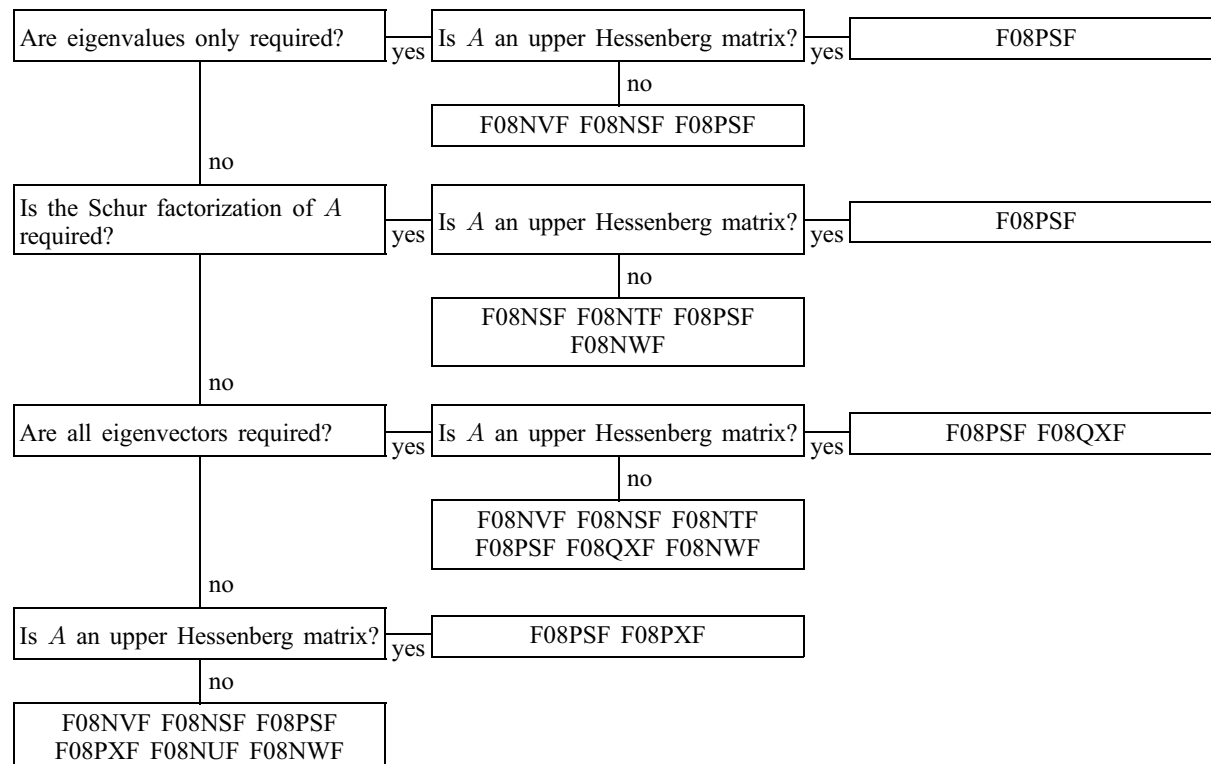
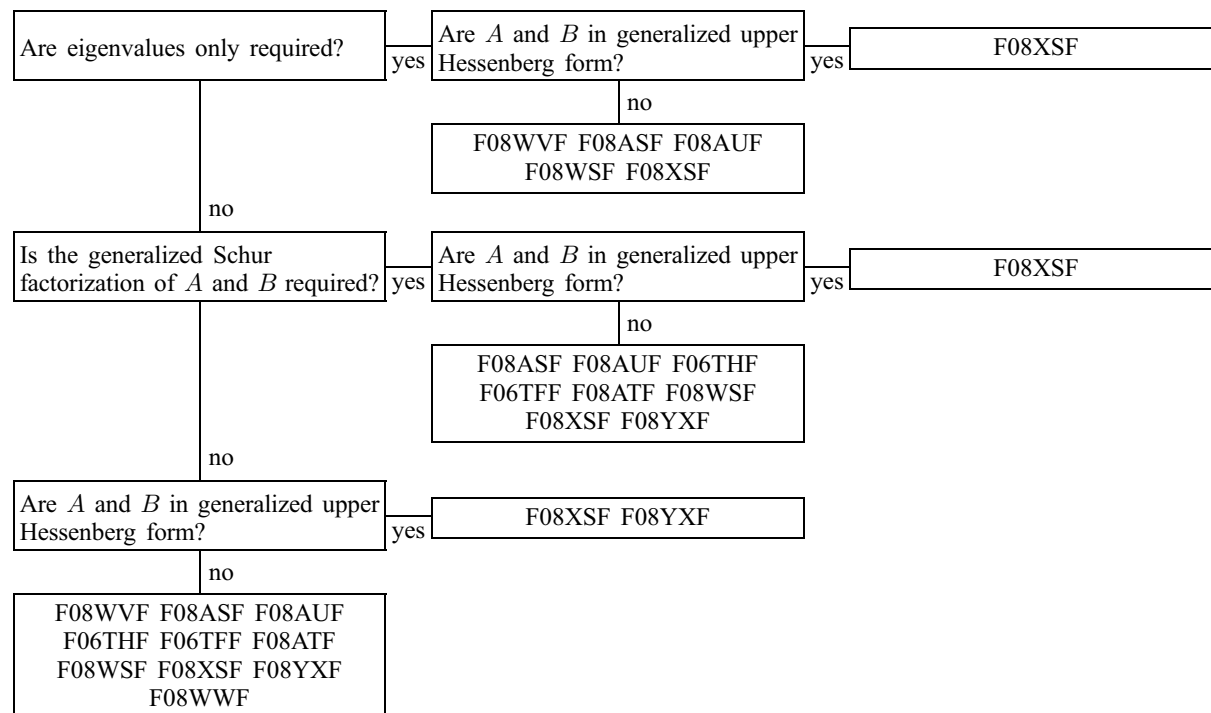
Tree 2: Real Generalized Symmetric-definite Eigenvalue Problems

Note: the routines for band matrices only handle the problem $Ax = \lambda Bx$; the other routines handle all three types of problems ($Ax = \lambda Bx$, $ABx = \lambda x$ or $BAx = \lambda x$) except that, if the problem is $BAx = \lambda x$ and eigenvectors are required, F06PHF must be used instead of F06PLF and F06YFF instead of F06YJF.

Tree 3: Real Nonsymmetric Eigenvalue Problems**Tree 4: Real Generalized Nonsymmetric Eigenvalue Problems**

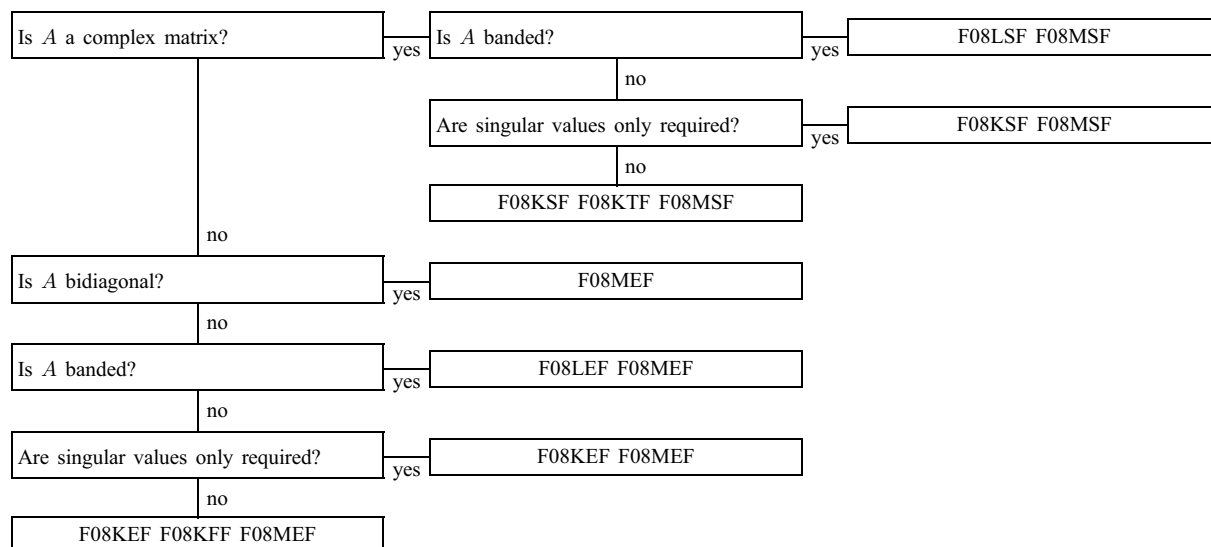
Tree 5: Complex Hermitian Eigenvalue Problems

Tree 6: Complex Generalized Hermitian-definite Eigenvalue Problems

Tree 7: Complex non-Hermitian Eigenvalue Problems**Tree 8: Complex Generalized non-Hermitian Eigenvalue Problems**

4.2 General Purpose Routines (singular value decomposition)

Tree 9



5 Index

Backtransformation of eigenvectors from those of balanced forms:

| | |
|----------------------|-----------------|
| complex matrix | F08NWF (ZGEBAK) |
| complex matrix | F08WWF (ZGGBAK) |
| real matrix | F08NJF (DGEBAK) |
| real matrix | F08WJF (DGGBAK) |

Balancing:

| | |
|------------------------------|-----------------|
| complex general matrix | F08NVF (ZGEBAL) |
| complex general matrix | F08WVF (ZGGBAL) |
| real general matrix | F08NHF (DGEBAL) |
| real general matrix | F08WHF (DGGBAL) |

Eigenvalue problems for condensed forms of matrices:

tridiagonal matrix:

| | |
|---|-----------------|
| selected eigenvalues and eigenvectors by root-free QR algorithm, inverse iteration or bisection | F08JBF (DSTEVS) |
|---|-----------------|

band matrix:

| | |
|--|-----------------|
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08HAF (DSBEV) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | F08HBF (DSBEVS) |

complex Hermitian matrix:

eigenvalues and eigenvectors:

band matrix:

| | |
|---|-----------------|
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm using packed storage | |
| all eigenvalues and eigenvectors by reduction to tridiagonal form using packed storage | F08HNF (ZHBEV) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form using packed storage | F08HPF (ZHBEVS) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm using packed storage | F08HQP (ZHBEVD) |

general matrix:

| | |
|---|----------------|
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm using packed storage | |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08FNF (ZHEEV) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form using packed storage | F08GNF (ZHPEV) |

| | |
|--|-----------------------|
| selected eigenvalues and eigenvectors by a divide-and-conquer algorithm | |
| | F08FRF (ZHEEVR) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | |
| | F08FPF (ZHEEVX) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form using packed storage | |
| | F08GPF (ZHPEVX) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08FQF (ZHEEVD) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm using packed storage | |
| | F08GQF (ZHPEVD) |
| eigenvalues only: | |
| band matrix: | |
| all eigenvalues by the Pal–Walker–Kahan variant of the QL or QR algorithm using packed storage | |
| | F08HQF (ZHBEVD) |
| general matrix: | |
| all eigenvalues by the Pal–Walker–Kahan variant of the QL or QR algorithm | |
| | F08FQF (ZHEEVD) |
| all eigenvalues by the Pal–Walker–Kahan variant of the QL or QR algorithm using packed storage | |
| | F08GQF (ZHPEVD) |
| complex upper Hessenberg matrix, reduced from complex general matrix: | |
| eigenvalues and Schur factorization | F08PSF (ZHSEQR) |
| selected right and/or left eigenvectors by inverse iteration | F08PXF (ZHSEIN) |
| general matrix: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm using packed storage | |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08FAF (DSYEV) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form using packed storage | |
| | F08GAF (DSPEV) |
| selected eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08FDF (DSYEV) |
| using packed storage | |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | F08FBF (DSYEVX) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form using packed storage | |
| | F08GBF (DSPEVX) |
| real bidiagonal matrix: | |
| singular value decomposition: | |
| after reduction from complex general matrix | F08MSF (ZBDSQR) |
| after reduction from real general matrix | F08MEF (DBDSQR) |
| real symmetric matrix: | |
| eigenvalues and eigenvectors: | |
| band matrix: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08HCF (DSBEVD) |
| general matrix: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08FCF (DSYEV) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm using packed storage | |
| | F08GCF (DSPEVD) |
| eigenvalues only: | |
| band matrix: | |
| all eigenvalues by the Pal–Walker–Kahan variant of the QL or QR algorithm | |
| | F08HCF (DSBEVD) |
| general matrix: | |
| all eigenvalues by the Pal–Walker–Kahan variant of the QL or QR algorithm | |
| | F08FCF (DSYEV) |
| all eigenvalues by the Pal–Walker–Kahan variant of the QL or QR algorithm using packed storage | |
| | F08GCF (DSPEVD) |
| real symmetric tridiagonal matrix: | |
| eigenvalues and eigenvectors: | |
| after reduction from complex Hermitian matrix: | |
| all eigenvalues and eigenvectors | F08JSF (ZSTEQR) |
| all eigenvalues and eigenvectors, | |
| positive-definite matrix | F08JUF (ZPTEQR) |

| | |
|--|-----------------|
| selected eigenvectors by inverse iteration | F08JXF (ZSTEIN) |
| all eigenvalues and eigenvectors | F08JEF (DSTEQR) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm after reduction from real symmetric matrix: | F08JCF (DSTEVD) |
| all eigenvalues and eigenvectors, positive-definite matrix | F08JGF (DPTEQR) |
| selected eigenvectors by inverse iteration | F08JKF (DSTEIN) |
| eigenvalues only: | |
| all eigenvalues by root-free <i>QR</i> algorithm | F08JFF (DSTERF) |
| all eigenvalues by the Pal–Walker–Kahan variant of the <i>QL</i> or <i>QR</i> algorithm | F08JCF (DSTEVD) |
| selected eigenvalues by bisection | F08JJF (DSTEBZ) |
| real upper Hessenberg matrix, reduced from real general matrix: | |
| eigenvalues and Schur factorization | F08PEF (DHSEQR) |
| selected right and/or left eigenvectors by inverse iteration | F08PKF (DHSEIN) |
| tridiagonal matrix: | |
| all eigenvalues and eigenvectors by root-free <i>QR</i> algorithm or inverse iteration | F08JAF (DSTEV) |
| selected eigenvalues and eigenvectors by relatively robust representations (RRR) | F08JDF (DSTEVF) |
| Eigenvalue problems for nonsymmetric matrices: | |
| complex matrix: | |
| all eigenvalues and left/right eigenvectors | F08NNF (ZGEEV) |
| all eigenvalues and left/right eigenvectors, plus balancing transformation and reciprocal condition numbers | F08NPF (ZGEEVX) |
| all eigenvalues, Schur form and Schur vectors | F08PNF (ZGEES) |
| all eigenvalues, Schur form, Schur vectors and reciprocal condition numbers | F08PPF (ZGEESX) |
| real matrix: | |
| all eigenvalues and left/right eigenvectors | F08NAF (DGEEV) |
| all eigenvalues and left/right eigenvectors, plus balancing transformation and reciprocal condition numbers | F08NBF (DGEEVX) |
| all eigenvalues, real Schur form and Schur vectors | F08PAF (DGEES) |
| all eigenvalues, real Schur form, Schur vectors and reciprocal condition numbers | F08PBF (DGEESX) |
| Eigenvalues and generalized Schur factorization, | |
| complex generalized upper Hessenberg form | F08XSF (ZHGEQZ) |
| real generalized upper Hessenberg form | F08XEF (DHGEQZ) |
| General Gauss-Markov linear model: | |
| solves a complex general Gauss-Markov linear model problem | F08ZPF (ZGGGLM) |
| solves a real general Gauss-Markov linear model problem | F08ZBF (DGGGLM) |
| Generalized eigenvalue problems for condensed forms of matrices: | |
| complex Hermitian-definite eigenproblems: | |
| banded matrices: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08UQF (ZHBGVD) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08UNF (ZHBGV) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | F08UPF (ZHBGVX) |
| general matrices: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08SQF (ZHEGVD) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm, packed storage format | F08TQF (ZHPGVD) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08SNF (ZHEGV) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form, packed storage format | F08TNF (ZHPGV) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | F08SPF (ZHEGVX) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form, packed storage format | F08TPF (ZHPGVX) |
| real symmetric-definite eigenproblems: | |
| banded matrices: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08UCF (DSBGVD) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08UAF (DSBGV) |

| | |
|---|------------------|
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | F08UBF (DSBGVX) |
| general matrices: | |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm | F08SCF (DSYGVD) |
| all eigenvalues and eigenvectors by a divide-and-conquer algorithm, packed storage format | F08TCF (DSPGVD) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form | F08SAF (DSYGV) |
| all eigenvalues and eigenvectors by reduction to tridiagonal form, packed storage format | F08TAF (DSPGV) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form | F08SBF (DSYGVX) |
| selected eigenvalues and eigenvectors by reduction to tridiagonal form, packed storage format | F08TBF (DSPGVX) |
| Generalized eigenvalue problems for nonsymmetric matrix pairs: | |
| complex nonsymmetric matrix pairs: | |
| all eigenvalues and left/right eigenvectors | F08WNF (ZGGEV) |
| all eigenvalues and left/right eigenvectors, plus the balancing transformation and reciprocal condition numbers | F08WPF (ZGGEVX) |
| all eigenvalues, generalized Schur form and Schur vectors | F08XNF (ZGGES) |
| all eigenvalues, generalized Schur form, Schur vectors and reciprocal condition numbers | F08XPF (ZGGESX) |
| real nonsymmetric matrix pairs: | |
| all eigenvalues and left/right eigenvectors | F08WAF (DGGEV) |
| all eigenvalues and left/right eigenvectors, plus the balancing transformation and reciprocal condition numbers | F08WBF (DGGEVX) |
| all eigenvalues, generalized real Schur form and left/right Schur vectors | F08XAF (DGGES) |
| all eigenvalues, generalized real Schur form and left/right Schur vectors, plus reciprocal condition numbers | F08XBF (DGGESX) |
| Generalized singular value decomposition | |
| complex matrix pair | F08VNF (ZGGSVD) |
| real matrix pair | F08VAF (DGGSD) |
| Least squares problems with linear equality constraints | |
| complex matrices: | |
| minimum norm solution subject to linear equality constraints using a generalized RQ factorization | F08ZNF (ZGGLSE) |
| real matrices: | |
| minimum norm solution subject to linear equality constraints using a generalized RQ factorization | F08ZAF (DGGGLSE) |
| Least squares problems: | |
| complex matrices: | |
| minimum norm solution using a complete orthogonal factorization | F08BNF (ZGELSY) |
| minimum norm solution using the singular value decomposition | F08KNF (ZGELSS) |
| minimum norm solution using the singular value decomposition (divide-and-conquer) | F08KQF (ZGELSD) |
| real matrices: | |
| minimum norm solution using a complete orthogonal factorization | F08BAF (DGELSY) |
| minimum norm solution using the singular value decomposition | F08KAF (DGELSS) |
| minimum norm solution using the singular value decomposition (divide-and-conquer) | F08KCF (DGELSD) |
| Left and right eigenvectors of a pair of matrices: | |
| complex upper triangular matrices | F08YXF (ZTGEVC) |
| real quasi-triangular matrices | F08YKF (DTGEVC) |
| LQ factorization and related operations: | |
| complex matrices: | |
| apply unitary matrix | F08AXF (ZUNMLQ) |
| factorization | F08AVF (ZGELQF) |
| form all or part of unitary matrix | F08AWF (ZUNGLQ) |
| real matrices: | |
| apply orthogonal matrix | F08AKF (DORMLQ) |
| factorization | F08AHF (DGELQF) |
| form all or part of orthogonal matrix | F08AJF (DORGLQ) |

Operations on Schur factorization of a general matrix:

complex matrix:

| | |
|---|-----------------|
| compute left and/or right eigenvectors | F08QXF (ZTREVC) |
| estimate sensitivities of eigenvalues and/or eigenvectors | F08QYF (ZTRSNA) |
| re-order Schur factorization | F08QTF (ZTREXC) |
| re-order Schur factorization, compute basis of invariant subspace, and estimate sensitivities | F08QUF (ZTRSEN) |

real matrix:

| | |
|---|-----------------|
| compute left and/or right eigenvectors | F08QKF (DTREVC) |
| estimate sensitivities of eigenvalues and/or eigenvectors | F08QLF (DTRSNA) |
| re-order Schur factorization | F08QFF (DTREXC) |
| re-order Schur factorization, compute basis of invariant subspace, and estimate sensitivities | F08QGF (DTRSEN) |

Overdetermined and underdetermined linear systems

complex matrices:

| | |
|--|----------------|
| solves an overdetermined or undetermined complex linear system | F08ANF (ZGELS) |
|--|----------------|

real matrices:

| | |
|---|----------------|
| solves an overdetermined or undetermined real linear system | F08AAF (DGELS) |
|---|----------------|

QR factorization and related operations:

complex matrices:

| | |
|--|-----------------|
| apply unitary matrix | F08AUF (ZUNMQR) |
| factorization | F08ASF (ZGEQRF) |
| factorization, with column pivoting | F08BSF (ZGEQPF) |
| form all or part of unitary matrix | F08ATF (ZUNGQR) |

real matrices:

| | |
|--|-----------------|
| apply orthogonal matrix | F08AGF (DORMQR) |
| factorization | F08AEF (DGEQRF) |
| factorization, with column pivoting | F08BEF (DGEQPF) |
| form all or part of orthogonal matrix | F08AFF (DORGQR) |

Reduction of a pair of general matrices to generalized upper Hessenberg form,

| | |
|---|-----------------|
| orthogonal reduction, real matrices | F08WEF (DGGHRD) |
| unitary reduction, complex matrices | F08WSF (ZGGHRD) |

Reduction of eigenvalue problems to condensed forms, and related operations:

complex general matrix to upper Hessenberg form:

| | |
|---------------------------------|-----------------|
| apply orthogonal matrix | F08NUF (ZUNMHR) |
| form orthogonal matrix | F08NTF (ZUNGHR) |
| reduce to Hessenberg form | F08NSF (ZGEHRD) |

complex Hermitian band matrix to real symmetric tridiagonal form

complex Hermitian matrix to real symmetric tridiagonal form:

| | |
|---|-----------------|
| apply unitary matrix | F08FUF (ZUNMTR) |
| apply unitary matrix, packed storage | F08GUF (ZUPMTR) |
| form unitary matrix | F08FTF (ZUNGTR) |
| form unitary matrix, packed storage | F08GTF (ZUPGTR) |
| reduce to tridiagonal form | F08FSF (ZHETRD) |
| reduce to tridiagonal form, packed storage | F08GSF (ZHPTRD) |

complex rectangular band matrix to real upper bidiagonal form

complex rectangular matrix to real bidiagonal form:

| | |
|---------------------------------|-----------------|
| apply unitary matrix | F08KUF (ZUNMBR) |
| form unitary matrix | F08KTF (ZUNGBR) |
| reduce to bidiagonal form | F08KSF (ZGEBRD) |

real general matrix to upper Hessenberg form:

| | |
|---------------------------------|-----------------|
| apply orthogonal matrix | F08NGF (DORMHR) |
| form orthogonal matrix | F08NFF (DORGHR) |
| reduce to Hessenberg form | F08NEF (DGEHRD) |

| | |
|--|-----------------|
| real rectangular band matrix to upper bidiagonal form | F08LEF (DGBBRD) |
| real rectangular matrix to bidiagonal form: | |
| apply orthogonal matrix | F08KGF (DORMBR) |
| form orthogonal matrix | F08KFF (DORGBR) |
| reduce to bidiagonal form | F08KEF (DGEBRD) |
| real symmetric band matrix to symmetric tridiagonal form | F08HEF (DSBTRD) |
| real symmetric matrix to symmetric tridiagonal form: | |
| apply orthogonal matrix | F08FGF (DORMTR) |
| apply orthogonal matrix, | |
| packed storage | F08GGF (DOPMTR) |
| form orthogonal matrix | F08FFF (DORGTR) |
| form orthogonal matrix, | |
| packed storage | F08GFF (DOPGTR) |
| reduce to tridiagonal form | F08FEF (DSYTRD) |
| reduce to tridiagonal form, | |
| packed storage | F08GEF (DSPTRD) |
| Reduction of generalized eigenproblems to standard eigenproblems: | |
| complex Hermitian-definite banded generalized eigenproblem $Ax = \lambda Bx$ | F08USF (ZHBGST) |
| complex Hermitian-definite generalized eigenproblem $Ax = \lambda Bx$, $ABx = \lambda x$ or $BAx = \lambda x$ | |
| F08SSF (ZHEGST) | |
| complex Hermitian-definite generalized eigenproblem $Ax = \lambda Bx$, $ABx = \lambda x$ or $BAx = \lambda x$, | |
| packed storage | F08TSF (ZHPGST) |
| real symmetric-definite banded generalized eigenproblem $Ax = \lambda Bx$ | F08UEF (DSBGST) |
| real symmetric-definite generalized eigenproblem $Ax = \lambda Bx$, $ABx = \lambda x$ or $BAx = \lambda x$ | |
| F08SEF (DSYGST) | |
| real symmetric-definite generalized eigenproblem $Ax = \lambda Bx$, $ABx = \lambda x$ or $BAx = \lambda x$, | |
| packed storage | F08TEF (DSPGST) |
| Singular value decomposition | |
| complex matrix: | |
| using a divide-and-conquer algorithm | F08KRF (ZGESDD) |
| using bidiagonal QR iteration | F08KPF (ZGESVD) |
| real matrix: | |
| using a divide-and-conquer algorithm | F08KDF (DGESDD) |
| using bidiagonal QR iteration | F08KBF (DGESVD) |
| Solve reduced form of Sylvester matrix equation: | |
| complex matrices | F08QVF (ZTRSYL) |
| real matrices | F08QHF (DTRSYL) |
| Split Cholesky factorization: | |
| complex Hermitian positive-definite band matrix | F08UTF (ZPBSTF) |
| real symmetric positive-definite band matrix | F08UFF (DPBSTF) |

6 Routines Withdrawn or Scheduled for Withdrawal

None.

7 References

Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J J, Du Croz J J, Greenbaum A, Hammarling S, McKenney A and Sorensen D (1999) *LAPACK Users' Guide* (3rd Edition) SIAM, Philadelphia URL: <http://www.netlib.org/lapack/lug>

Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J J, Du Croz J J, Greenbaum A, Hammarling S, McKenney A and Sorensen D (1999) *LAPACK Users' Guide* (3rd Edition) SIAM, Philadelphia

Arioli M, Duff I S and De Rijk P P M (1989) On the augmented system approach to sparse least-squares problems *Numer. Math.* **55** 667–684

Demmel J W and Kahan W (1990) Accurate singular values of bidiagonal matrices *SIAM J. Sci. Statist. Comput.* **11** 873–912

Golub G H and Van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore

Moler C B and Stewart G W (1973) An algorithm for generalized matrix eigenproblems *SIAM J. Numer. Anal.* **10** 241–256

Parlett B N (1998) *The Symmetric Eigenvalue Problem* SIAM, Philadelphia

Stewart G W and Sun J-G (1990) *Matrix Perturbation Theory* Academic Press, London

Ward R C (1981) Balancing the generalized eigenvalue problem *SIAM J. Sci. Stat. Comp.* **2** 141–152

Wilkinson J H (1965) *The Algebraic Eigenvalue Problem* Oxford University Press, Oxford

Wilkinson J H and Reinsch C (1971) *Handbook for Automatic Computation II, Linear Algebra* Springer-Verlag
