Science and Technology Facilities Council

EA22

1 SUMMARY

Given a **real symmetric matrix** $\mathbf{A}=\{a_{ij}\}_{n\times n}$, this subroutine **calculates the** k **eigenvalues,** $\lambda_1, \lambda_2, ..., \lambda_k$ of largest modulus $|\lambda_1| \geq |\lambda_2| \geq ,..., \geq |\lambda_k|$ and their associated eigenvectors $\mathbf{x}_i$, i=1,2,...,k, where $\mathbf{A}\mathbf{x}_i = \lambda_i \mathbf{x}_i$.

The subroutine uses the method of simultaneous iteration and also allows the user to take advantage of sparsity. Eigensolutions from other parts of the spectrum can be obtained by using shifts of the form $\mathbf{A}-\beta\mathbf{I}$ optionally combined with inversion.

ATTRIBUTES — **Version:** 1.2.0. (26 March 2013) **Types:** Real (single, double). **Calls:** FA14. **Original date:** May 2001. **Remark:** The EA22 entries are threadsafe versions of EA12. **Origin:** I. S. Duff, Harwell.

2 HOW TO USE THE PACKAGE

2.1 Overall control

Rather than ask the user to furnish a subroutine to perform the matrix-vector multiplication **Aw** inside the algorithm, control is returned at each iteration so that code for computing **Aw** can be written by the user in line. On exit from each successful call to EA22A/AD the vector **w** is contained in the first N locations of the work array W and the user must place the resultant vector **Aw** in the array U, but no other change should be made to the arguments of EA22A/AD before the next call to EA22A/AD. When EA22A/AD is called again, the parameter IPOS determines the place in the code from which the computation proceeds. If EA22A/AD returns with a value of IPOS equal to 1 then the algorithm has converged, while a return with a negative value of IPOS indicates an error has occurred (see Section 2.4).

A simple example to illustrate the calling sequence is given in Section 5.

2.2 Argument list

Initialization

The EA22I/ID entry must be called prior to the first call to the EA22A/AD entry to initialize the control array and private workspace.

The single precision version

```
CALL EA22I(ICNTL,KEEP,RKEEP)
```

The double precision version

```
CALL  EA22ID(ICNTL,KEEP,RKEEP)
```

ICNTL is an INTEGER array of length 10, see Section 2.5.

KEEP is an INTEGER array of length 30 used by EA22 as private workspace and should not be altered by the user. The element KEEP(21) holds the random number seed used by FA14, which is used by EA22 to generate a random start vector. In exceptional circumstances the user may wish to reset it after the EA22I/ID call to to obtain a different random sequence.

RKEEP is a REAL (DOUBLE PRECISION in the D version) array of length 20 used by EA22 as private workspace and should not be altered by the user.

To calculate the eigenvalues and eigenvectors

The single precision version

```
CALL EA22A(N,NUMEIG,NUMCOL,EPS,X,NX,D,U,W,LW,IW,KMULT,IPOS,ICNTL,KEEP,RKEEP)
```

The double precision version

```
CALL EA22AD(N,NUMEIG,NUMCOL,EPS,X,NX,D,U,W,LW,IW,KMULT,IPOS,ICNTL,KEEP,RKEEP)
```

N is an INTEGER variable which must be set by the user to the order of the matrix. Unchanged on exit. **Restriction**: N ≥ 1.

NUMEIG is an INTEGER variable which must be set by the user to the number of eigensolutions required. On each exit, NUMEIG is equal to the number of eigensolutions found so far. **Restriction**: NUMEIG ≥ 1.

NUMCOL is an INTEGER variable which must be set by the user to the number of vectors on which simultaneous iteration will be performed. It must exceed NUMEIG unless an estimate of the (NUMCOL+1)-th eigenvalue has been specified in D(1), in which case it may equal NUMEIG. Unchanged on exit. **Restriction**: NUMCOL ≤ N and NUMCOL ≥ NUMEIG.

EPS is a REAL (DOUBLE PRECISION in the D version) variable. An eigensolution ($\mathbf{x}$,λ) will be regarded as satisfactory whenever its residual

$$\|\mathbf{Ax} - \lambda \mathbf{x}\|_2$$

is less than EPS. If the user asks for an accuracy greater than can be achieved, then the subroutine will increase EPS. The new value for EPS, which may be different for different eigenpairs, is output. This value will always be less than $n\varepsilon\|\mathbf{A}\|_2$, where ε is the machine precision.

X is a REAL (DOUBLE PRECISION in the D version) two-dimensional array with dimensions (NX, NUMCOL). X is used to hold successive approximates to eigenvectors of **A**. If the user wishes to supply an initial estimate it should be input in X (and ICNTL(3) set to 1, see Section 2.5); on exit, the first NUMEIG columns of X will hold orthonormalized eigenvectors corresponding to the eigenvalues of **A** of largest modulus while the other columns will hold unnormalized approximations to other eigenvectors.

NX is an INTEGER variable which must be set by the user to equal the first dimension of the variable in the calling program corresponding to the parameter X. Unchanged on exit. **Restriction**: NX ≥ N.

D is a REAL (DOUBLE PRECISION in the D version) array of length NUMCOL. On input D(1) can be set to an estimate of the modulus of the (NUMCOL+1)-th eigenvalue of the matrix, if known, and ICNTL(2) set to 1 (see Section 2.5). If the user specifies this, the convergence of the algorithm is likely to be enhanced. On exit, the first NUMEIG entries of D contain the eigenvalues of **A** of largest modulus while the rest of the array contains estimates of other eigenvalues.

U is a REAL (DOUBLE PRECISION in the D version) array of length N which is used as workspace by EA22A/AD. This array is used by the user when performing multiplications of vectors by **A**.

W is a REAL (DOUBLE PRECISION in the D version) array of length LW which is used as workspace by EA22A/AD. The first N locations of W are used by the user when performing multiplications of vectors by **A**.

LW is an INTEGER variable which must be set by the user to the length of array W. It must be at least 2*N+2*NUMCOL*NUMCOL. Unchanged on exit.

IW is an INTEGER array of length NUMCOL which is used as a work array by EA22A/AD.

KMULT is an INTEGER variable. On entry, KMULT must be set to the maximum number of matrix-vector multiplications the user wishes to allow. Unchanged on exit.

IPOS is an INTEGER variable used to control the subroutine (see Section 2.1). Its value must not be changed between calls of EA22A/AD. IPOS must be set initially to 1. On exit, a negative value of IPOS indicates an error has occurred (see Section 2.4); a value of IPOS greater than 1 indicates further iterations are to be performed (the value of IPOS determines the place in the code from which computation proceeds on the next call to EA22A/AD); a value of IPOS equal to 1 indicates that the algorithm has successfully converged.

ICNTL is an INTEGER array of length 10 whose first three elements can optionally be set by the user, see Section 2.5.

KEEP is an INTEGER array of length 30 used by EA22 as private workspace, see Section 2.1.

RKEEP is a REAL (DOUBLE PRECISION in the D version) array of length 20 used by EA22 as private workspace, see Section 2.1.

2.3 Use of EA22A/AD to obtain other parts of the spectrum

Subroutine EA22A/AD can be used to find other parts of the spectrum than that corresponding to the eigenvalues of **A** of largest modulus. If, for example, we wish to find the smallest eigenvalues of a positive definite matrix **A**, and we have an upper bound β for the largest eigenvalue of **A**, then the use of the shifted matrix $(\mathbf{A}-\beta\mathbf{I})$ in place of **A** will ensure convergence to the appropriate part of the spectrum.

If we wish a group of interior eigenvalues, say those nearest to p, then we require the matrix-vector multiplication by $(\mathbf{A}-p\,\mathbf{I})^{-1}$ which is equivalent to solving the system

$$(\mathbf{A}-p\,\mathbf{I})\,\mathbf{u}=\mathbf{w} \tag{1}$$

at each matrix-vector operation step. Naturally, if **A** is large and sparse the solution of the system (1) may itself be quite time-consuming but note that the decomposition of $\mathbf{A}-p\,\mathbf{I}$ into triangular factors needs only to be done once for the entire calculation, the only operations in the inner loop being the relatively cheap forward and backward substitutions.

Finally, it is clear that the generalized eigenvalue problem

$$\mathbf{Ax}=\lambda\,\mathbf{Bx} \tag{2}$$

with **B** positive definite and **A** symmetric can be transformed into

$$(\mathbf{L}^{-1}\mathbf{AL}^{-T})\,(\mathbf{L}^{T}\mathbf{x})=\lambda\,(\mathbf{L}^{T}\mathbf{x}) \tag{3}$$

where $\mathbf{B}=\mathbf{LL}^{T}$. EA22A/AD can then be used to solve problem (3) and hence problem (2). Notice that $\mathbf{L}^{-1}\mathbf{AL}^{-T}$ is held as three factors which in the sparse case will be far sparser than their product matrix.

2.4 Error diagnostics

If EA22A/AD returns with a negative value of IPOS an error has occurred, and an error message will be output on unit ICNTL(1). Possible negative values of IPOS are given below.

–1 Value of N out of range. N < 1.

–2 Value of NUMEIG out of range. NUMEIG < 1.

–3 Value of NX out of range. NX < N.

–4 Value of NUMCOL out of range. Either NUMCOL > N or NUMCOL < NUMEIG.

–5 Algorithm required more matrix-vector multiplications than specified by KMULT. This can be corrected by, for example, increasing the value of KMULT, changing EPS the tolerance parameter, or increasing NUMCOL the number of simultaneous iterates.

–6 Pathologically slow convergence is obtained. This is caused by the ratio of the NUMCOL-th eigenvalue to one of the converged eigenvalues being 1+ε (where ε is the machine accuracy) or by a user supplied bad estimate of the (NUMCOL+1)-th eigenvalue being given in D(1). In this case the information on the first NUMEIG eigensolutions is still valid.

–7 There is an infinitesimal chance of a catastrophic failure causing termination of the algorithm. This is indicated by the message

```
UNEXPECTED ERROR IN EA22A/AD (ORTHON)
```

on unit ICNTL(1) (see Section 2.5) and EA22A/AD returns with a value of IPOS equal to –7. In this case the user should contact the Numerical Analysis Group at the Rutherford Appleton Laboratory.

2.5 The control array

The ICNTL array argument which must be of length 10 can be used to pass optional control values to the routine.

ICNTL(1) specifies the unit number to be used to output error messages. A negative value will suppress output. It has a default value of 6.

ICNTL(2) has a default value of zero, but must be set to 1 if the user takes the option of supplying the modulus of the (NUMCOL+1)-th eigenvalue in D(1).

ICNTL(3) has a default value of zero, but must be set to 1 when X contains a starting guess to the eigenvector(s) on entry to the routine.

ICNTL(4) to ICNTL(10) should not be altered and at present are not used by EA22.

2.6 Underflows

The nature of the calculations performed in this subroutine means that underflows are very likely to occur. It is quite safe to set numbers that underflow to zero, and action from the user might be required to ensure that this is done efficiently by the computing system in use.

3 GENERAL INFORMATION

Use of common: none.

Other routines called directly: FA14A/AD, and routines in the EA22 package EA22B/BD, EA22E/ED, EA22F/FD, EA22G/GD, EA22H/HD and EA22J/JD.

Input/output: error messages are printed on unit ICNTL(1), see Sections 2.4 and 2.5.

Restrictions:

N $\geq$ 1,

NUMEIG $\geq$ 1,

NX $\geq$ N,

NUMCOL $\leq$ N,

NUMCOL > NUMEIG or NUMCOL $\geq$ NUMEIG if an estimate of the (NUMCOL+1)-th eigenvalue is supplied in D(1).

4 METHOD

The method used is simultaneous iteration and is an adaptation of a procedure of Rutishauser (Numer. Math. **16** (1970), pp.205-223). Chebyshev acceleration is used to speed convergence if the eigenvalues are closely spaced, as is likely to happen if the shifted matrix ($\mathbf{A} - \beta\mathbf{I}$) is used to find the smallest eigenvalues of the positive-definite matrix **A** (see Section 2.3).

5 EXAMPLE OF USE

The following program illustrates the use of EA22A/AD. We wish to calculate the largest eigenvalue and corresponding eigenvector of the matrix

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 & 4 \\ 0 & 2 & 3 & 0 \\ 0 & 3 & 5 & 1 \\ 4 & 0 & 1 & 2 \end{pmatrix}.$$

```
C CODE TO ILLUSTRATE THE USE OF EA22
      INTEGER NX,NMAXE,LWMAX
      PARAMETER(NX=5, NMAXE=4, LWMAX=40)
      INTEGER NZ,I,K,J,L
      INTEGER N,NUMEIG,NUMCOL,LW,KMULT,IPOS,
     *        IW(NMAXE),IRN(NMAXE*NMAXE),ICN(NMAXE*NMAXE)
      INTEGER ICNTL(10),KEEP(30)
      DOUBLE PRECISION RKEEP(20)
      DOUBLE PRECISION EPS, X(NX,NMAXE),D(NMAXE),U(NMAXE),
     *                 A(NMAXE*NMAXE),W(LWMAX)
C  INITIALISE
      CALL EA22ID(ICNTL,KEEP,RKEEP)
C
C  READ IN DATA
C
      READ(5,*) N, NUMEIG, NUMCOL
      READ(5,*) EPS
      READ(5,*) KMULT
C NZ IS THE NUMBER OF NONZERO ENTRIES IN THE MATRIX A.
C ENTRY A(I) OF THE MATRIX A IS IN ROW IRN(I) AND COLUMN ICN(I).
      READ(5,*) NZ
      DO 10 I=1,NZ
        READ(5,*)IRN(I),ICN(I),A(I)
  10  CONTINUE
C
      LW = 2*N + 2*NUMCOL*NUMCOL
      IPOS = 1
      DO 30 K=1,KMULT
         CALL EA22AD(N,NUMEIG,NUMCOL,EPS,X,NX,D,U,W,LW,IW,KMULT,IPOS,
     +               ICNTL,KEEP,RKEEP)
         IF(IPOS.EQ.1) GO TO 40
         IF(IPOS.LT.0) GO TO 100
         DO 20 I=1,N
           U(I) = 0.0D0
  20     CONTINUE
         DO 25 L=1,NZ
           I=IRN(L)
           J=ICN(L)
           U(I) = U(I) + A(L)*W(J)
  25     CONTINUE
  30  CONTINUE
  40  CONTINUE
C
C ALGORITHM HAS TERMINATED SUCCESSFULLY
C
      WRITE(6,910) K
      WRITE(6,900) (D(I),I=1,NUMEIG)
      WRITE(6,905) ((X(J,I),J=1,N),I=1,NUMEIG)
 100  CONTINUE
900   FORMAT(/' CALCULATED LARGEST EIGENVALUE IS: ',1P,D12.4)
905   FORMAT(/' CORRESPONDING EIGENVECTOR IS: ',/(1X,1P,4D12.4))
910   FORMAT(/' NUMBER OF MATRIX-VECTOR MULTIPLICATIONS USED WAS: ',I2)
      STOP
      END
```

The input data used is

```
   4  1  3
   1.0D-02
  25
  10
   4       4            2.
   1       1            1.
   4       3            1.
   3       4            1.
   2       2            2.
   1       4            4.
   3       2            3.
   4       1            4.
   3       3            5.
   2       3            3.
```

This gives the following output:

```
  NUMBER OF MATRIX-VECTOR MULTIPLICATIONS USED WAS: 15

  CALCULATED LARGEST EIGENVALUE IS:   7.1436D+00

  CORRESPONDING EIGENVECTOR IS:
   -2.0522D-01 -4.6681D-01 -8.0038D-01 -3.1520D-01
```