

NAG Library Routine Document

F02SDF

Note: before using this routine, please read the Users' Note for your implementation to check the interpretation of ***bold italicised*** terms and other implementation-dependent details.

1 Purpose

F02SDF finds the eigenvector corresponding to a given real eigenvalue for the generalized problem $Ax = \lambda Bx$, or for the standard problem $Ax = \lambda x$, where A and B are real band matrices.

2 Specification

```
SUBROUTINE F02SDF (N, MA1, MB1, A, LDA, B, LDB, SYM, RELEP, RMU, VEC, D,      &
                  IWORK, WORK, LWORK, IFAIL)

INTEGER          N, MA1, MB1, LDA, LDB, IWORK(N), LWORK, IFAIL
REAL (KIND=nag_wp) A(LDA,N), B(LDB,N), RELEP, RMU, VEC(N), D(30),      &
                  WORK(LWORK)
LOGICAL          SYM
```

3 Description

Given an approximation μ to a real eigenvalue λ of the generalized eigenproblem $Ax = \lambda Bx$, F02SDF attempts to compute the corresponding eigenvector by inverse iteration.

F02SDF first computes lower and upper triangular factors, L and U , of $A - \mu B$, using Gaussian elimination with interchanges, and then solves the equation $Ux = e$, where $e = (1, 1, 1, \dots, 1)^T$ – this is the first half iteration.

There are then three possible courses of action depending on the input value of D(1).

1. D(1) = 0.

This setting should be used if λ is an ill-conditioned eigenvalue (provided the matrix elements do not vary widely in order of magnitude). In this case it is essential to accept only a vector found after one half iteration, and μ must be a very good approximation to λ . If acceptable growth is achieved in the solution of $Ux = e$, then the normalized x is accepted as the eigenvector. If not, columns of an orthogonal matrix are tried in turn in place of e . If none of these give acceptable growth, the routine fails, indicating that μ was not a sufficiently good approximation to λ .

2. D(1) > 0.

This setting should be used if μ is moderately close to an eigenvalue which is not ill-conditioned (provided the matrix elements do not differ widely in order of magnitude). If acceptable growth is achieved in the solution of $Ux = e$, the normalized x is accepted as the eigenvector. If not, inverse iteration is performed. Up to 30 iterations are allowed to achieve a vector and a correction to μ which together give acceptably small residuals.

3. D(1) < 0.

This setting should be used if the elements of A and B vary widely in order of magnitude. Inverse iteration is performed, but a different convergence criterion is used.

See Section 9.3 for further details.

Note that the bandwidth of the matrix A must not be less than the bandwidth of B . If this is not so, either A must be filled out with zeros, or matrices A and B may be reversed and $1/\mu$ supplied as an approximation to the eigenvalue $1/\lambda$. Also it is assumed that A and B each have the same number of subdiagonals as superdiagonals. If this is not so, they must be filled out with zeros. If A and B are **both** symmetric, only the upper triangles need be supplied.

4 References

Peters G and Wilkinson J H (1979) Inverse iteration, ill-conditioned equations and Newton's method *SIAM Rev.* **21** 339–360

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

Wilkinson J H (1972) Inverse iteration in theory and practice *Symposia Mathematica Volume X* 361–379 Istituto Nazionale di Alta Matematica, Monograf, Bologna

Wilkinson J H (1974) Notes on inverse iteration and ill-conditioned eigensystems *Acta Univ. Carolin. Math. Phys.* **1–2** 173–177

Wilkinson J H (1979) Kronecker's canonical form and the *QZ* algorithm *Linear Algebra Appl.* **28** 285–303

5 Arguments

1: N – INTEGER *Input*

On entry: n , the order of the matrices A and B .

Constraint: $N \geq 1$.

2: MA1 – INTEGER *Input*

On entry: the value $m_A + 1$, where m_A is the number of nonzero lines on each side of the diagonal of A . Thus the total bandwidth of A is $2m_A + 1$.

Constraint: $1 \leq \text{MA1} \leq N$.

3: MB1 – INTEGER *Input*

On entry: if $\text{MB1} \leq 0$, B is assumed to be the unit matrix. Otherwise MB1 must specify the value $m_B + 1$, where m_B is the number of nonzero lines on each side of the diagonal of B . Thus the total bandwidth of B is $2m_B + 1$.

Constraint: $\text{MB1} \leq \text{MA1}$.

4: A(LDA,N) – REAL (KIND=nag_wp) array *Input/Output*

On entry: the n by n band matrix A . The m_A subdiagonals must be stored in the first m_A rows of the array; the diagonal in the $(m_A + 1)$ th row; and the m_A superdiagonals in rows $m_A + 2$ to $2m_A + 1$. Each row of the matrix must be stored in the corresponding column of the array. For example, if $n = 6$ and $m_A = 2$ the storage scheme is:

*	*	a_{31}	a_{42}	a_{53}	a_{64}
*	a_{21}	a_{32}	a_{43}	a_{54}	a_{65}
a_{11}	a_{22}	a_{33}	a_{44}	a_{55}	a_{66}
a_{12}	a_{23}	a_{34}	a_{45}	a_{56}	*
a_{13}	a_{24}	a_{35}	a_{46}	*	*

Elements of the array marked * need not be set. The following code assigns the matrix elements within the band to the correct elements of the array:

```
DO 20 J = 1, N
  DO 10 I = MAX(1,J-MA1+1), MIN(N,J+MA1-1)
    A(I-J+MA1,J) = matrix(J,I)
  10 CONTINUE
20 CONTINUE
```

If $\text{SYM} = \text{.TRUE.}$ (i.e., both A and B are symmetric), only the lower triangle of A need be stored in the first MA1 rows of the array.

On exit: details of the factorization of $A - \bar{\lambda}B$, where $\bar{\lambda}$ is an estimate of the eigenvalue.

- 5: LDA – INTEGER *Input*
On entry: the first dimension of the array A as declared in the (sub)program from which F02SDF is called.
Constraint: $LDA \geq 2 \times MA1 - 1$.
- 6: B(LDB, N) – REAL (KIND=nag_wp) array *Input/Output*
On entry: if $MB1 > 0$, B must contain the n by n band matrix B , stored in the same way as A . If $SYM = .TRUE.$, only the lower triangle of B need be stored in the first $MB1$ rows of the array. If $MB1 \leq 0$, the array is not used.
On exit: elements in the top-left corner, and in the bottom right corner if $SYM = .FALSE.$, are set to zero; otherwise the array is unchanged.
- 7: LDB – INTEGER *Input*
On entry: the first dimension of the array B as declared in the (sub)program from which F02SDF is called.
Constraints:
 if $SYM = .FALSE.$, $LDB \geq 2 \times MB1 - 1$;
 if $SYM = .TRUE.$, $LDB \geq MB1$.
- 8: SYM – LOGICAL *Input*
On entry: if $SYM = .TRUE.$, both A and B are assumed to be symmetric and only their upper triangles need be stored. Otherwise SYM must be set to $.FALSE.$.
- 9: RELEP – REAL (KIND=nag_wp) *Input*
On entry: the relative error of the coefficients of the given matrices A and B . If the value of RELEP is less than the **machine precision**, the **machine precision** is used instead.
- 10: RMU – REAL (KIND=nag_wp) *Input*
On entry: μ , an approximation to the eigenvalue for which the corresponding eigenvector is required.
- 11: VEC(N) – REAL (KIND=nag_wp) array *Output*
On exit: the eigenvector, normalized so that the largest element is unity, corresponding to the improved eigenvalue $RMU + D(30)$.
- 12: D(30) – REAL (KIND=nag_wp) array *Input/Output*
On entry: $D(1)$ must be set to indicate the type of problem (see Section 3):
 $D(1) > 0.0$
 Indicates a well-conditioned eigenvalue.
 $D(1) = 0.0$
 Indicates an ill-conditioned eigenvalue.
 $D(1) < 0.0$
 Indicates that the matrices have elements varying widely in order of magnitude.
On exit: if $D(1) \neq 0.0$ on entry, the successive corrections to μ are given in $D(i)$, for $i = 1, 2, \dots, k$, where $k + 1$ is the total number of iterations performed. The final correction is also given in the last position, $D(30)$, of the array. The remaining elements of D are set to zero.
 If $D(1) = 0.0$ on entry, no corrections to μ are computed and $D(i)$ is set to 0.0, for $i = 1, 2, \dots, 30$. Thus in all three cases the best available approximation to the eigenvalue is $RMU + D(30)$.

- 13: IWORK(N) – INTEGER array *Workspace*
 14: WORK(LWORK) – REAL (KIND=nag_wp) array *Workspace*
 15: LWORK – INTEGER *Input*

On entry: the dimension of the array WORK as declared in the (sub)program from which F02SDF is called.

Constraints:

if $D(1) \neq 0.0$, $LWORK \geq N \times (MA1 + 1)$;
 if $D(1) = 0.0$, $LWORK \geq 2 \times N$.

- 16: IFAIL – INTEGER *Input/Output*

On entry: IFAIL must be set to 0, -1 or 1. If you are unfamiliar with this argument you should refer to Section 3.4 in How to Use the NAG Library and its Documentation for details.

For environments where it might be inappropriate to halt program execution when an error is detected, the value -1 or 1 is recommended. If the output of error messages is undesirable, then the value 1 is recommended. Otherwise, if you are not familiar with this argument, the recommended value is 0. **When the value -1 or 1 is used it is essential to test the value of IFAIL on exit.**

On exit: IFAIL = 0 unless the routine detects an error or a warning has been flagged (see Section 6).

6 Error Indicators and Warnings

If on entry IFAIL = 0 or -1, explanatory error messages are output on the current error message unit (as defined by X04AAF).

Errors or warnings detected by the routine:

IFAIL = 1

On entry, $N < 1$,
 or $MA1 < 1$,
 or $MA1 > N$,
 or $LDA < 2 \times MA1 - 1$,
 or $LDB < MB1$ when $SYM = .TRUE.$,
 or $LDB < 2 \times MB1 - 1$ when $SYM = .FALSE.$ (LDB is not checked if $MB1 \leq 0$).

IFAIL = 2

On entry, $MA1 < MB1$. Either fill out A with zeros, or reverse the roles of A and B, and replace RMU by its reciprocal, i.e., solve $Bx = \lambda^{-1}Ax$.

IFAIL = 3

On entry, $LWORK < 2 \times N$ when $D(1) = 0.0$,
 or $LWORK < N \times (MA1 + 1)$ when $D(1) \neq 0.0$.

IFAIL = 4

A is null. If B is nonsingular, all the eigenvalues are zero and any set of N orthogonal vectors forms the eigensolution.

IFAIL = 5

B is null. If A is nonsingular, all the eigenvalues are infinite, and the columns of the unit matrix are eigenvectors.

IFAIL = 6

On entry, A and B are both null. The eigensolution is arbitrary.

IFAIL = 7

$D(1) \neq 0.0$ on entry and convergence is not achieved in 30 iterations. Either the eigenvalue is ill-conditioned or RMU is a poor approximation to the eigenvalue. See Section 9.3.

IFAIL = 8

$D(1) = 0.0$ on entry and no eigenvector has been found after $\min(N, 5)$ back-substitutions. RMU is not a sufficiently good approximation to the eigenvalue.

IFAIL = 9

$D(1) < 0.0$ on entry and RMU is too inaccurate for the solution to converge.

IFAIL = -99

An unexpected error has been triggered by this routine. Please contact NAG.

See Section 3.9 in How to Use the NAG Library and its Documentation for further information.

IFAIL = -399

Your licence key may have expired or may not have been installed correctly.

See Section 3.8 in How to Use the NAG Library and its Documentation for further information.

IFAIL = -999

Dynamic memory allocation failed.

See Section 3.7 in How to Use the NAG Library and its Documentation for further information.

7 Accuracy

The eigensolution is exact for some problem

$$(A + E)x = \mu(B + F)x,$$

where $\|E\|, \|F\|$ are of the order of $\eta(\|A\| + \mu\|B\|)$, where η is the value used for RELEP.

8 Parallelism and Performance

F02SDF is not threaded in any implementation.

9 Further Comments

9.1 Timing

The time taken by F02SDF is approximately proportional to $n(2m_A + 1)^2$ for factorization, and to $n(2m_A + 1)$ for each iteration.

9.2 Storage

The storage of the matrices A and B is designed for efficiency on a paged machine.

F02SDF will work with full matrices but it will do so inefficiently, particularly in respect of storage requirements.

9.3 Algorithmic Details

Inverse iteration is performed according to the rule

$$(A - \mu B)y_{r+1} = Bx_r$$

$$x_{r+1} = \frac{1}{\alpha_{r+1}} y_{r+1}$$

where α_{r+1} is the element of y_{r+1} of largest magnitude.

Thus:

$$(A - \mu B)x_{r+1} = \frac{1}{\alpha_{r+1}} Bx_r.$$

Hence the residual corresponding to x_{r+1} is very small if $|\alpha_{r+1}|$ is very large (see Peters and Wilkinson (1979)). The first half iteration, $Uy_1 = e$, corresponds to taking $L^{-1}PBx_0 = e$.

If μ is a very accurate eigenvalue, then there should always be an initial vector x_0 such that one half iteration gives a small residual and thus a good eigenvector. If the eigenvalue is ill-conditioned, then second and subsequent iterated vectors may not be even remotely close to an eigenvector of a neighbouring problem (see pages 374–376 of Wilkinson (1972) and Wilkinson (1974)). In this case it is essential to accept only a vector obtained after one half iteration.

However, for well-conditioned eigenvalues, there is no loss in performing more than one iteration (see page 376 of Wilkinson (1972)), and indeed it will be necessary to iterate if μ is not such a good approximation to the eigenvalue. When the iteration has converged, y_{r+1} will be some multiple of x_r , $y_{r+1} = \beta_{r+1}x_r$, say.

Therefore

$$(A - \mu B)\beta_{r+1}x_r = Bx_r,$$

giving

$$\left(A - \left(\mu + \frac{1}{\beta_{r+1}}\right)B\right)x_r = 0.$$

Thus $\mu + \frac{1}{\beta_{r+1}}$ is a better approximation to the eigenvalue. β_{r+1} is obtained as the element of y_{r+1} which corresponds to the element of largest magnitude, $+1$, in x_r . The routine terminates when $\left\| \left(A - \left(\mu + \frac{1}{\beta_r}\right)B\right)x_r \right\|$ is of the order of the *machine precision* relative to $\|A\| + |\mu|\|B\|$.

If the elements of A and B vary widely in order of magnitude, then $\|A\|$ and $\|B\|$ are excessively large and a different convergence test is required. The routine terminates when the difference between successive corrections to μ is small relative to μ .

In practice one does not necessarily know if the given problem is well-conditioned or ill-conditioned. In order to provide some information on the condition of the eigenvalue or the accuracy of μ in the event of failure, successive values of $\frac{1}{\beta_r}$ are stored in the vector D when D(1) is nonzero on input. If these values appear to be converging steadily, then it is likely that μ was a poor approximation to the eigenvalue and it is worth trying again with RMU + D(30) as the initial approximation. If the values in D vary considerably in magnitude, then the eigenvalue is ill-conditioned.

A discussion of the significance of the singularity of A and/or B is given in relation to the QZ algorithm in Wilkinson (1979).

10 Example

Given the generalized eigenproblem $Ax = \lambda Bx$ where

$$A = \begin{pmatrix} 1 & 1 & 2 & & \\ -1 & 2 & 1 & 2 & \\ & -1 & 3 & 1 & 2 \\ & & -1 & 4 & 1 \\ & & & -1 & 5 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 5 & 1 & & & \\ 1 & 4 & 2 & & \\ & 2 & 3 & 2 & \\ & & 2 & 2 & 1 \\ & & & 1 & 1 \end{pmatrix}$$

find the eigenvector corresponding to the approximate eigenvalue -12.33 .

Although B is symmetric, A is not, so SYM must be set to .FALSE. and all the elements of B in the band must be supplied to the routine. A (as written above) has 1 subdiagonal and 2 superdiagonals, so MA1 must be set to 3 and A filled out with an additional subdiagonal of zeros. Each row of the matrices is read in as data in turn.

10.1 Program Text

```

Program f02sdfe

!      F02SDF Example Program Text

!      Mark 26 Release. NAG Copyright 2016.

!      .. Use Statements ..
      Use nag_library, Only: f02sdf, nag_wp
!      .. Implicit None Statement ..
      Implicit None
!      .. Parameters ..
      Integer, Parameter          :: nin = 5, nout = 6
!      .. Local Scalars ..
      Real (Kind=nag_wp)          :: relep, rmu
      Integer                     :: i, ifail, j, k, k1, k2, lda, ldb,      &
                                   lwork, ma, mb, n
      Logical                     :: sym
!      .. Local Arrays ..
      Real (Kind=nag_wp), Allocatable :: a(:,,:), b(:,,:), vec(:), work(:)
      Real (Kind=nag_wp)              :: d(30)
      Integer, Allocatable             :: iwork(:)
!      .. Intrinsic Procedures ..
      Intrinsic                      :: min
!      .. Executable Statements ..
      Write (nout,*) 'F02SDF Example Program Results'
!      Skip heading in data file
      Read (nin,*)
      Read (nin,*) n, ma, mb
      lda = 2*ma + 1
      ldb = 2*mb + 1
      lwork = n*(ma+2)
      Allocate (a(lda,n),b(ldb,n),vec(n),work(lwork),iwork(n))
      Do i = 1, n
         k1 = ma + 1 - min(ma,i-1)
         k2 = ma + 1 + min(ma,n-i)
         Read (nin,*)(a(k,i),k=k1,k2)
      End Do
      Do i = 1, n
         k1 = mb + 1 - min(mb,i-1)
         k2 = mb + 1 + min(mb,n-i)
         Read (nin,*)(b(k,i),k=k1,k2)
      End Do
      Read (nin,*) rmu, d(1)
      sym = .False.
      relep = 0.0E0_nag_wp

!      ifail: behaviour on error exit
!      =0 for hard exit, =1 for quiet-soft, =-1 for noisy-soft
      ifail = 1
      Call f02sdf(n,ma+1,mb+1,a,lda,b,ldb,sym,relep,rmu,vec,d,iwork,work,      &

```

```

        lwork,ifail)

Write (nout,*)
If (ifail==0) Then
    Write (nout,99999) 'Corrected eigenvalue = ', rmu + d(30)
    Write (nout,*)
    Write (nout,*) 'Eigenvector is'
    Write (nout,99998) vec(1:n)
Else If (ifail>0) Then
    Write (nout,99997) 'Error in F02SDF. IFAIL =', ifail
    If (ifail==7 .Or. ifail==9) Then
        Write (nout,*)
        Write (nout,*) 'Successive corrections to RMU were'
        Write (nout,*)
        Do j = 1, 29
            If (d(j)==0.0E0_nag_wp) Then
                Go To 100
            End If
            Write (nout,99996) d(j)
        End Do
    End If
Else
    Write (nout,99995) ifail
End If
100    Continue

99999 Format (1X,A,F8.4)
99998 Format (1X,5F9.4)
99997 Format (1X,A,I5)
99996 Format (1X,E20.4)
99995 Format (1X,' ** F02SDF returned with IFAIL = ',I5)
End Program f02sdfe

```

10.2 Program Data

F02SDF Example Program Data

```

5 2 1                                : n, ma, mb
1.0 1.0 2.0
-1.0 2.0 1.0 2.0
0.0 -1.0 3.0 1.0 2.0
0.0 -1.0 4.0 1.0
0.0 -1.0 5.0                        : matrix A
5.0 1.0
1.0 4.0 2.0
2.0 3.0 2.0
2.0 2.0 1.0
1.0 1.0                            : matrix B
-12.33 1.0                        : rmu, d(1)

```

10.3 Program Results

F02SDF Example Program Results

Corrected eigenvalue = -12.3394

Eigenvector is
 -0.0572 0.3951 -0.8427 1.0000 -0.6540
