

NAG Fortran Library Routine Document

D02HAF

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

D02HAF solves the two-point boundary-value problem for a system of ordinary differential equations, using a Runge–Kutta–Merson method and a Newton iteration in a shooting and matching technique.

2 Specification

```
SUBROUTINE D02HAF(U, V, N, A, B, TOL, FCN, SOLN, M1, W, IW, IFAIL)
INTEGER          N, M1, IW, IFAIL
real           U(N,2), V(N,2), A, B, TOL, SOLN(N,M1), W(N,IW)
EXTERNAL        FCN
```

3 Description

D02HAF solves the two-point boundary-value problem for a system of n ordinary differential equations in the range a, b . The system is written in the form:

$$y'_i = f_i(x, y_1, y_2, \dots, y_n), \quad i = 1, 2, \dots, n \quad (1)$$

and the derivatives f_i are evaluated by a subroutine FCN supplied by the user. Initially, n boundary values of the variables y_i must be specified, some of which will be specified at a and some at b . The user must supply estimates of the remaining n boundary values (called parameters below), and the subroutine corrects them by a form of Newton iteration. It also calculates the complete solution on an equispaced mesh if required.

Starting from the known and estimated values of y_i at a , the subroutine integrates the equations from a to b (using a Runge–Kutta–Merson method). The differences between the values of y_i at b from integration and those specified initially should be zero for the true solution. (These differences are called residuals below.) The subroutine uses a generalized Newton method to reduce the residuals to zero, by calculating corrections to the estimated boundary values. This process is repeated iteratively until convergence is obtained, or until the routine can no longer reduce the residuals. See Hall and Watt (1976) for a simple discussion of shooting and matching techniques.

4 References

Hall G and Watt J M (ed.) (1976) *Modern Numerical Methods for Ordinary Differential Equations* Clarendon Press, Oxford

5 Parameters

1: U(N,2) – *real* array *Input/Output*

On entry: U(i , 1) must be set to the known or estimated value of y_i at a , and U(i , 2) must be set to the known or estimated value of y_i at b , for $i = 1, 2, \dots, n$.

On exit: the known values unaltered, and corrected values of the estimates, unless an error has occurred. If an error has occurred, U contains the known values and the latest values of the estimates.

- 2: $V(N,2)$ – *real* array *Input*
On entry: $V(i, j)$ must be set to 0.0 if $U(i, j)$ is a known value, and 1.0 if $U(i, j)$ is an estimated value to be corrected.
Constraint: precisely n of the $V(i, j)$ must be set to 0.0, i.e., precisely n of the $U(i, j)$ must be known values, and these must not be all at a or all at b .
- 3: N – INTEGER *Input*
On entry: the number of equations, n .
Constraint: $N \geq 2$.
- 4: A – *real* *Input*
On entry: the initial point of the interval of integration, a .
- 5: B – *real* *Input*
On entry: the final point of the interval of integration, b .
- 6: TOL – *real* *Input*
On entry: TOL must be set to a small quantity suitable for
 (a) testing the local error in y_i during integration,
 (b) testing for the convergence of y_i at b ,
 (c) calculating the perturbation in estimated boundary values for y_i , which are used to obtain the approximate derivatives of the residuals for use in the Newton iteration.
 The user is advised to check his results by varying TOL .
Constraint: $TOL > 0.0$.
- 7: FCN – SUBROUTINE, supplied by the user. *External Procedure*
 FCN must evaluate the functions f_i (i.e., the derivatives y'_i), for $i = 1, 2, \dots, n$, at a general point x .
 Its specification is:

```

SUBROUTINE FCN(X, Y, F)
  real X, Y(n), F(n)
  where n is the actual value of N on the call of D02HAF.
1: X – real Input
   On entry: the value of the argument,  $x$ .
2: Y(n) – real array Input
   On entry: the value of the argument  $y_i$ , for  $i = 1, 2, \dots, n$ .
3: F(n) – real array Output
   On exit: the value of  $f_i$ ,  $x$ , for  $i = 1, 2, \dots, n$ .

```

FCN must be declared as EXTERNAL in the (sub)program from which D02HAF is called. Parameters denoted as *Input* must **not** be changed by this procedure.

- 8: $SOLN(N,M1)$ – *real* array *Output*
On exit: the solution when $M1 > 1$ (see below).

- 9: M1 – INTEGER *Input*
On entry: a value which controls output.
M1 = 1
The final solution is not evaluated.
M1 > 1
The final values of y_i at interval $(b - a)/(M1 - 1)$ are calculated and stored in the array SOLN by columns, starting with values y_i at a stored in SOLN($i, 1$), for $i = 1, 2, \dots, n$.
Constraint: $M1 \geq 1$.
- 10: W(N,IW) – *real* array *Output*
On exit: if IFAIL = 2, 3, 4 or 5, W($i, 1$) contains the solution at the point where the integration fails, for $i = 1, 2, \dots, n$, and the point of failure is returned in W(1,2).
- 11: IW – INTEGER *Input*
On entry: the second dimension of W.
Constraint: $IW \geq 3N + 17 + \max(11, N)$.
- 12: IFAIL – INTEGER *Input/Output*
For this routine, the normal use of IFAIL is extended to control the printing of error and warning messages as well as specifying hard or soft failure (see Chapter P01).
Before entry, IFAIL must be set to a value with the decimal expansion cba , where each of the decimal digits c , b and a must have a value of 0 or 1.
 $a = 0$ specifies hard failure, otherwise soft failure;
 $b = 0$ suppresses error messages, otherwise error messages will be printed (see Section 6);
 $c = 0$ suppresses warning messages, otherwise warning messages will be printed (see Section 6).
The recommended value for inexperienced users is 110 (i.e., hard failure with all messages printed).
Unless the routine detects an error (see Section 6), IFAIL contains 0 on exit.

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

One or more of the parameters V, N, M1, IW, or TOL is incorrectly set.

IFAIL = 2

The step length for the integration is too short whilst calculating the residual (see Section 8).

IFAIL = 3

No initial step length could be chosen for the integration whilst calculating the residual.

Note: IFAIL = 2 or 3 can occur due to choosing too small a value for TOL or due to choosing the wrong direction of integration. Try varying TOL and interchanging a and b . These error exits can also occur for very poor initial estimates of the unknown initial values and, in extreme cases, because this routine cannot be used to solve the problem posed.

IFAIL = 4

As for IFAIL = 2 but the error occurred when calculating the Jacobian of the derivatives of the residuals with respect to the parameters.

IFAIL = 5

As for IFAIL = 3 but the error occurred when calculating the derivatives of the residuals with respect to the parameters.

IFAIL = 6

The calculated Jacobian has an insignificant column.

Note: IFAIL = 4, 5 or 6 usually indicate a badly scaled problem. The user may vary the size of TOL or change to one of the more general routines D02HBF or D02SAF which afford more control over the calculations.

IFAIL = 7

The linear algebra routine (F02WEF) used has failed. This error exit should not occur and can be avoided by changing the estimated initial values.

IFAIL = 8

The Newton iteration has failed to converge.

Note: IFAIL = 8 can indicate poor initial estimates or a very difficult problem. Consider varying TOL if the residuals are small in the monitoring output. If the residuals are large try varying the initial estimates.

IFAIL = 9

Indicate that a serious error has occurred in D02SAZ, D02SAW, D02SAX, D02SAU or D02SAV respectively. Check all array subscripts and subroutine parameter lists in calls to D02HAF. Seek expert help.

7 Accuracy

If the process converges, the accuracy to which the unknown parameters are determined is usually close to that specified by the user; the solution, if requested, may be determined to a required accuracy by varying the parameter TOL.

8 Further Comments

The time taken by the routine depends on the complexity of the system, and on the number of iterations required. In practice, integration of the differential equations is by far the most costly process involved.

Wherever it occurs in the routine, the error parameter TOL is used in 'mixed' form; that is TOL always occurs in expressions of the form $TOL \times (1 + |y_i|)$. Though not ideal for every application, it is expected that this mixture of absolute and relative error testing will be adequate for most purposes.

The user is strongly recommended to set IFAIL to obtain self-explanatory error messages, and also monitoring information about the course of the computation. The user may select the channel numbers on which this output is to appear by calls of X04AAF (for error messages) or X04ABF (for monitoring information) – see Section 9 for an example. Otherwise the default channel numbers will be used, as specified in the implementation document. The monitoring information produced at each iteration includes the current parameter values, the residuals and two norms: a basic norm and a current norm. At each iteration the aim is to find parameter values which make the current norm less than the basic norm. Both these norms should tend to zero as should the residuals. (They would all be zero if the exact parameters were used as input.) For more details, the user may consult the specification of D02SAF, and especially the description of the parameter MONIT there.

The computing time for integrating the differential equations can sometimes depend critically on the quality of the initial estimates. If it seems that too much computing time is required and, in particular, if the values of the residuals printed by the monitoring routine are much larger than the expected values of the solution at b , then the coding of the subroutine FCN should be checked for errors. If no errors can be found, an independent attempt should be made to improve the initial estimates. In practical problems it is not uncommon for the differential equation to have a singular point at one or both ends of the range. Suppose a is a singular point; then the derivatives y'_i in (1) (in Section 3) cannot be evaluated at a , usually because one or more of the expressions for f_i give overflow. In such a case it is necessary for the user to take a a short distance away from the singularity, and to find values for y_i at the new value of a (e.g., use the first one or two terms of an analytical (power series) solution). The user should experiment with the new position of a ; if it is taken too close to the singular point, the derivatives f_i will be inaccurate, and the routine may sometimes fail with IFAIL = 2 or 3 or, in extreme cases, with an overflow condition. A more general treatment of singular solutions is provided by the subroutine D02HBF.

Another difficulty which often arises in practice is the case when one end of the range, b say, is at infinity. The user must approximate the end-point by taking a finite value for b , which is obtained by estimating where the solution will reach its asymptotic state. The estimate can be checked by repeating the calculation with a larger value of b . If b is very large, and if the matching point is also at b , the numerical solution may suffer a considerable loss of accuracy in integrating across the range, and the program may fail with IFAIL = 6 or 8. (In the former case, solutions from all initial values at a are tending to the same curve at infinity.) The simplest remedy is to try to solve the equations with a smaller value of b , and then to increase b in stages, using each solution to give boundary value estimates for the next calculation. For problems where some terms in the asymptotic form of the solution are known, D02HBF will be more successful.

If the unknown quantities are not boundary values, but are eigenvalues or the length of the range or some other parameters occurring in the differential equations, D02HBF may be used.

9 Example

To find the angle at which a projectile must be fired for a given range.

The differential equations are:

$$\begin{aligned} y' &= \tan \phi \\ v' &= \frac{-0.032 \tan \phi}{v} - \frac{0.02v}{\cos \phi} \\ \phi' &= \frac{-0.032}{v^2}, \end{aligned}$$

with the following boundary conditions:

$$\begin{aligned} y = 0, \quad v = 0.5 \quad \text{at} \quad x = 0, \\ y = 0 \quad \quad \quad \text{at} \quad x = 5. \end{aligned}$$

The remaining boundary conditions are estimated as:

$$\begin{aligned} \phi = 1.15 \quad \quad \quad \text{at} \quad x = 0, \\ \phi = 1.2, \quad v = 0.46 \quad \text{at} \quad x = 5. \end{aligned}$$

We write $y = Z(1)$, $v = Z(2)$, $\phi = Z(3)$. To check the accuracy of the results the problem is solved twice with TOL = 5.0E-3 and 5.0E-4 respectively. Note the call to X04ABF before the call to D02HAF.

9.1 Program Text

Note: the listing of the example program presented below uses *bold italicised* terms to denote precision-dependent details. Please read the Users' Note for your implementation to check the interpretation of these terms. As explained in the Essential Introduction to this manual, the results produced may not be identical for all implementations.

```

*      D02HAF Example Program Text
*      Mark 14 Revised.  NAG Copyright 1989.
*      .. Parameters ..
*      N.B the definition of IW must be changed for N.GT.11
      INTEGER          NOUT
      PARAMETER       (NOUT=6)
      INTEGER          N, IW, M1
      PARAMETER       (N=3,IW=3*N+17+11,M1=6)
*      .. Local Scalars ..
      real            TOL, X, X1
      INTEGER          I, IFAIL, J, L
*      .. Local Arrays ..
      real            U(N,2), V(N,2), W(N,IW), Y(N,M1)
*      .. External Subroutines ..
      EXTERNAL        D02HAF, DERIV, X04ABF
*      .. Executable Statements ..
      WRITE (NOUT,*) 'D02HAF Example Program Results'
      CALL X04ABF(1,NOUT)
      DO 40 L = 3, 4
          TOL = 5.0e0*10.0e0**(-L)
          WRITE (NOUT,*)
          WRITE (NOUT,99999) 'Results with TOL = ', TOL
          U(1,1) = 0.0e0
          V(1,1) = 0.0e0
          U(1,2) = 0.0e0
          V(1,2) = 0.0e0
          U(2,1) = 0.5e0
          V(2,1) = 0.0e0
          U(2,2) = 0.46e0
          V(2,2) = 1.0e0
          U(3,1) = 1.15e0
          V(3,1) = 1.0e0
          U(3,2) = -1.2e0
          V(3,2) = 1.0e0
          X = 0.0e0
          X1 = 5.0e0
*      * Set IFAIL to 111 to obtain monitoring information *
          IFAIL = 11
*
          CALL D02HAF(U,V,N,X,X1,TOL,DERIV,Y,M1,W,IW,IFAIL)
*
          WRITE (NOUT,*)
          IF (IFAIL.EQ.0) THEN
              WRITE (NOUT,*) ' X-value and final solution'
              DO 20 I = 1, M1
                  WRITE (NOUT,99998) I - 1, (Y(J,I),J=1,N)
20          CONTINUE
          ELSE
              WRITE (NOUT,99997) ' IFAIL =', IFAIL
          END IF
40 CONTINUE
      STOP
*
99999 FORMAT (1X,A,e10.3)
99998 FORMAT (1X,I3,3F10.4)
99997 FORMAT (1X,A,I4)
      END
*
      SUBROUTINE DERIV(X,Z,G)
*      .. Parameters ..
      INTEGER          N
      PARAMETER       (N=3)
*      .. Scalar Arguments ..
      real            X
*      .. Array Arguments ..

```

```

      real          G(N), Z(N)
*      .. Intrinsic Functions ..
      INTRINSIC    COS, TAN
*      .. Executable Statements ..
      G(1) = TAN(Z(3))
      G(2) = -0.032e0*TAN(Z(3))/Z(2) - 0.02e0*Z(2)/COS(Z(3))
      G(3) = -0.032e0/Z(2)**2
      RETURN
      END

```

9.2 Program Data

None.

9.3 Program Results

D02HAF Example Program Results

Results with TOL = 0.500E-02

X-value and final solution

0	0.0000	0.5000	1.1680
1	1.9172	0.3343	0.9746
2	2.9293	0.2067	0.4916
3	2.9762	0.1956	-0.4214
4	2.0177	0.3099	-0.9756
5	-0.0088	0.4602	-1.2020

Results with TOL = 0.500E-03

X-value and final solution

0	0.0000	0.5000	1.1681
1	1.9177	0.3343	0.9749
2	2.9280	0.2070	0.4929
3	2.9769	0.1955	-0.4194
4	2.0210	0.3095	-0.9752
5	-0.0000	0.4597	-1.2014
