

D02BHF – NAG Fortran Library Routine Document

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

D02BHF integrates a system of first-order ordinary differential equations over an interval with suitable initial conditions, using a Runge–Kutta–Merson method, until a user-specified function of the solution is zero.

2 Specification

```

SUBROUTINE D02BHF(X, XEND, N, Y, TOL, IRELAB, HMAX, FCN, G, W,
1          IFAIL)
  INTEGER      N, IRELAB, IFAIL
  real        X, XEND, Y(N), TOL, HMAX, G, W(N,7)
  EXTERNAL    FCN, G

```

3 Description

The routine advances the solution of a system of ordinary differential equations

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

from $x = X$ towards $x = XEND$ using a Merson form of the Runge–Kutta method. The system is defined by a subroutine FCN supplied by the user, which evaluates f_i in terms of x and y_1, y_2, \dots, y_n (see Section 5), and the values of y_1, y_2, \dots, y_n must be given at $x = X$.

As the integration proceeds, a check is made on the function $g(x, y)$ specified by the user, to determine an interval where it changes sign. The position of this sign change is then determined accurately by interpolating for the solution and its derivative. It is assumed that $g(x, y)$ is a continuous function of the variables, so that a solution of $g(x, y) = 0$ can be determined by searching for a change in sign in $g(x, y)$.

The accuracy of the integration and, indirectly, of the determination of the position where $g(x, y) = 0$, is controlled by the parameter TOL.

For a description of Runge–Kutta methods and their practical implementation see Hall and Watt [1].

4 References

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

5 Parameters

1: X — *real* *Input/Output*

On entry: X must be set to the initial value of the independent variable x .

On exit: the point where $g(x, y) = 0.0$ unless an error has occurred, when it contains the value of x at the error. In particular, if $g(x, y) \neq 0.0$ anywhere on the range X to $XEND$, it will contain $XEND$ on exit.

2: $XEND$ — *real* *Input*

On entry: the final value of the independent variable x .

If $XEND < X$ on entry, integration proceeds in a negative direction.

- 3:** N — INTEGER *Input*
On entry: the number of differential equations, n .
Constraint: $N > 0$.
- 4:** Y(N) — *real* array *Input/Output*
On entry: the initial values of the solution y_1, y_2, \dots, y_n .
On exit: the computed values of the solution at the final point $x = X$.
- 5:** TOL — *real* *Input/Output*
On entry: TOL must be set to a **positive** tolerance for controlling the error in the integration and in the determination of the position where $g(x, y) = 0.0$.

D02BHF has been designed so that, for most problems, a reduction in TOL leads to an approximately proportional reduction in the error in the solution obtained in the integration. The relation between changes in TOL and the error in the determination of the position where $g(x, y) = 0.0$ is less clear, but for TOL small enough the error should be approximately proportional to TOL. However, the actual relation between TOL and the accuracy cannot be guaranteed. The user is strongly recommended to call D02BHF with more than one value for TOL and to compare the results obtained to estimate their accuracy. In the absence of any prior knowledge the user might compare results obtained by calling D02BHF with $TOL = 10.0^{-p}$ and $TOL = 10.0^{-p-1}$ if p correct decimal digits in the solution are required.

Constraint: $TOL > 0.0$.

On exit: normally unchanged. However if the range from $x = X$ to the position where $g(x, y) = 0.0$ (or to the final value of x if an error occurs) is so short that a small change in TOL is unlikely to make any change in the computed solution, then TOL is returned with its sign changed. To check results returned with $TOL < 0.0$, D02BHF should be called again with a positive value of TOL whose magnitude is considerably smaller than that of the previous call.

- 6:** IRELAB — INTEGER *Input*
On entry: IRELAB determines the type of error control. At each step in the numerical solution an estimate of the local error, EST, is made. For the current step to be accepted the following condition must be satisfied:

$$\begin{aligned} \text{IRELAB} = 0 & \qquad \qquad \qquad \text{EST} \leq \text{TOL} \times \max\{1.0, |y_1|, |y_2|, \dots, |y_n|\}; \\ \text{IRELAB} = 1 & \\ \text{EST} \leq \text{TOL}; & \\ \text{IRELAB} = 2 & \qquad \text{EST} \leq \text{TOL} \times \max\{\epsilon, |y_1|, |y_2|, \dots, |y_n|\}, \text{ where } \epsilon \text{ is } \mathbf{machine\ precision}. \end{aligned}$$

If the appropriate condition is not satisfied, the step size is reduced and the solution recomputed on the current step.

If the user wishes to measure the error in the computed solution in terms of the number of correct decimal places, then IRELAB should be given the value 1 on entry, whereas if the error requirement is in terms of the number of correct significant digits, then IRELAB should be given the value 2. Where there is no preference in the choice of error test, $\text{IRELAB} = 0$ will result in a mixed error test. It should be borne in mind that the computed solution will be used in evaluating $g(x, y)$.

Constraint: $0 \leq \text{IRELAB} \leq 2$.

- 7:** HMAX — *real* *Input*
On entry: if $\text{HMAX} = 0.0$, no special action is taken.

If $\text{HMAX} \neq 0.0$, a check is made for a change in sign of $g(x, y)$ at steps not greater than $|\text{HMAX}|$. This facility should be used if there is any chance of ‘missing’ the change in sign by checking too infrequently. For example, if two changes of sign of $g(x, y)$ are expected within a distance h , say, of each other, then a suitable value for HMAX might be $\text{HMAX} = h/2$. If only one change of sign in $g(x, y)$ is expected on the range X to XEND, then the choice $\text{HMAX} = 0.0$ is most appropriate.

8: FCN — SUBROUTINE, supplied by the user. *External Procedure*

FCN must evaluate the functions f_i (i.e., the derivatives y_i') for given values of its arguments x, y_1, \dots, y_n .

Its specification is:

```
SUBROUTINE FCN(X, Y, F)
  real          X, Y(n), F(n)
```

where n is the actual value of N in the call of D02BHF.

- | | | |
|-----------|---|---------------|
| 1: | X — <i>real</i>
<i>On entry:</i> the value of the argument x . | <i>Input</i> |
| 2: | Y(n) — <i>real</i> array
<i>On entry:</i> the value of the argument y_i , for $i = 1, 2, \dots, n$. | <i>Input</i> |
| 3: | F(n) — <i>real</i> array
<i>On exit:</i> the value of f_i , for $i = 1, 2, \dots, n$. | <i>Output</i> |

FCN must be declared as EXTERNAL in the (sub)program from which D02BHF is called. Parameters denoted as *Input* must **not** be changed by this procedure.

9: G — *real* FUNCTION, supplied by the user. *External Procedure*

G must evaluate the function $g(x, y)$ at a specified point.

Its specification is:

```
real FUNCTION G(X, Y)
  real          X, Y(n)
```

where n is the actual value of N in the call of D02BHF.

- | | | |
|-----------|--|--------------|
| 1: | X — <i>real</i>
<i>On entry:</i> the value of the independent variable x . | <i>Input</i> |
| 2: | Y(n) — <i>real</i> array
<i>On entry:</i> the value of y_i , for $i = 1, 2, \dots, n$. | <i>Input</i> |

G must be declared as EXTERNAL in the (sub)program from which D02BHF is called. Parameters denoted as *Input* must **not** be changed by this procedure.

10: W(N,7) — *real* array *Workspace*

11: IFAIL — INTEGER *Input/Output*

On entry: IFAIL must be set to 0, -1 or 1. For users not familiar with this parameter (described in Chapter P01) the recommended value is 0.

On exit: IFAIL = 0 unless the routine detects an error (see Section 6).

6 Error Indicators and Warnings

Errors detected by the routine:

IFAIL = 1

On entry, TOL \leq 0.0,
or N \leq 0,
or IRELAB \neq 0, 1 or 2.

IFAIL = 2

With the given value of TOL, no further progress can be made across the integration range from the current point $x = X$, or dependence of the error on TOL would be lost if further progress across the integration range were attempted (see Section 8 for a discussion of this error exit). The components $Y(1), Y(2), \dots, Y(n)$ contain the computed values of the solution at the current point $x = X$. No point at which $g(x, y)$ changes sign has been located up to the point $x = X$.

IFAIL = 3

TOL is too small for the routine to take an initial step (see Section 8). X and $Y(1), Y(2), \dots, Y(n)$ retain their initial values.

IFAIL = 4

At no point in the range X to $XEND$ did the function $g(x, y)$ change sign. It is assumed that $g(x, y) = 0.0$ has no solution.

IFAIL = 5

A serious error has occurred in an internal call to C05AZF. Check all subroutine calls and array dimensions. Seek expert help.

IFAIL = 6

A serious error has occurred in an internal call to an integration routine. Check all subroutine calls and array dimensions. Seek expert help.

IFAIL = 7

A serious error has occurred in an internal call to an interpolation routine. Check all subroutine calls and array dimensions. Seek expert help.

7 Accuracy

The accuracy depends on TOL, on the mathematical properties of the differential system, on the position where $g(x, y) = 0.0$ and on the method. It can be controlled by varying TOL but the approximate proportionality of the error to TOL holds only for a restricted range of values of TOL. For TOL too large, the underlying theory may break down and the result of varying TOL may be unpredictable. For TOL too small, rounding error may affect the solution significantly and an error exit with IFAIL = 2 or IFAIL = 3 is possible.

The accuracy may also be restricted by the properties of $g(x, y)$. The user should try to code G without introducing any unnecessary cancellation errors.

8 Further Comments

The time taken by the routine depends on the complexity and mathematical properties of the system of differential equations defined by FCN, the complexity of G, on the range, the position of the solution and the tolerance. There is also an overhead of the form $a + b \times n$ where a and b are machine-dependent computing times.

For some problems it is possible that D02BHF will return IFAIL = 4 because of inaccuracy of the computed values Y , leading to inaccuracy in the computed values of $g(x, y)$ used in the search for the solution of $g(x, y) = 0.0$. This difficulty can be overcome by reducing TOL sufficiently, and if necessary,

by choosing HMAX sufficiently small. If possible, the user should choose XEND well beyond the expected point where $g(x, y) = 0.0$; for example make $|XEND - X|$ about 50% larger than the expected range. As a simple check, if, with XEND fixed, a change in TOL does not lead to a significant change in Y at XEND, then inaccuracy is not a likely source of error.

If the routine fails with IFAIL = 3, then it could be called again with a larger value of TOL if this has not already been tried. If the accuracy requested is really needed and cannot be obtained with this routine, the system may be very stiff (see below) or so badly scaled that it cannot be solved to the required accuracy.

If the routine fails with IFAIL = 2, it is likely that it has been called with a value of TOL which is so small that a solution cannot be obtained on the range X to XEND. This can happen for well-behaved systems and very small values of TOL. The user should, however, consider whether there is a more fundamental difficulty. For example:

- (a) in the region of a singularity (infinite value) of the solution, the routine will usually stop with IFAIL = 2, unless overflow occurs first. If overflow occurs using D02BHF, D02PDF can be used instead to detect the increasing solution, before overflow occurs. In any case, numerical integration cannot be continued through a singularity, and analytical treatment should be considered;
- (b) for ‘stiff’ equations, where the solution contains rapidly decaying components, the routine will compute in very small steps in x (internally to D02BHF) to preserve stability. This will usually exhibit itself by making the computing time excessively long, or occasionally by an exit with IFAIL = 2. Merson’s method is not efficient in such cases, and the user should try D02EJF which uses a Backward Differentiation Formula method. To determine whether a problem is stiff, D02PCF may be used.

For well-behaved systems with no difficulties such as stiffness or singularities, the Merson method should work well for low accuracy calculations (three or four figures). For high accuracy calculations or where FCN is costly to evaluate, Merson’s method may not be appropriate and a computationally less expensive method may be D02CJF which uses an Adams method.

For problems for which D02BHF is not sufficiently general, the user should consider D02PDF. D02PDF is a more general routine with many facilities including a more general error control criterion. D02PDF can be combined with the rootfinder C05AZF and the interpolation routine D02PXF to solve equations involving y_1, y_2, \dots, y_n and their derivatives.

D02BHF can also be used to solve an equation involving x, y_1, y_2, \dots, y_n and the derivatives of y_1, y_2, \dots, y_n . For example in Section 9, D02BHF is used to find a value of $X > 0.0$ where $Y(1) = 0.0$. It could instead be used to find a turning-point of y_1 by replacing the function $g(x, y)$ in the program by:

```

real FUNCTION G(X,Y)
real X,Y(3),F(3)
CALL FCN(X,Y,F)
G = F(1)
RETURN
END

```

This routine is only intended to locate the **first** zero of $g(x, y)$. If later zeros are required, users are strongly advised to construct their own more general root finding routines as discussed above.

9 Example

To find the value $X > 0.0$ at which $y = 0.0$, where y, v, ϕ are defined by

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

and where at $X = 0.0$ we are given $y = 0.5$, $v = 0.5$ and $\phi = \pi/5$. We write $y = Y(1)$, $v = Y(2)$ and $\phi = Y(3)$ and we set $TOL = 1.0E-4$ and $TOL = 1.0E-5$ in turn so that we can compare the solutions. We expect the solution $X \simeq 7.3$ and so we set $XEND = 10.0$ to avoid determining the solution of $y = 0.0$ too near the end of the range of integration. The value of π is obtained by using X01AAF.

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.

```

*      D02BHF Example Program Text
*      Mark 14 Revised.  NAG Copyright 1989.
*      .. Parameters ..
      INTEGER          NOUT
      PARAMETER       (NOUT=6)
      INTEGER          N
      PARAMETER       (N=3)
*      .. Local Scalars ..
      real            HMAX, PI, TOL, X, XEND
      INTEGER          I, IFAIL, IRELAB, J
*      .. Local Arrays ..
      real            W(N,7), Y(N)
*      .. External Functions ..
      real            G, X01AAF
      EXTERNAL         G, X01AAF
*      .. External Subroutines ..
      EXTERNAL         D02BHF, FCN
*      .. Executable Statements ..
      WRITE (NOUT,*) 'D02BHF Example Program Results'
      XEND = 10.0e0
      HMAX = 0.0e0
      IRELAB = 0
      PI = X01AAF(X)
      DO 20 J = 4, 5
          TOL = 10.0e0**(-J)
          WRITE (NOUT,*)
          WRITE (NOUT,99999) 'Calculation with TOL =', TOL
          X = 0.0e0
          Y(1) = 0.5e0
          Y(2) = 0.5e0
          Y(3) = 0.2e0*PI
          IFAIL = 0

*
          CALL D02BHF(X,XEND,N,Y,TOL,IRELAB,HMAX,FCN,G,W,IFAIL)
*
          WRITE (NOUT,99998) ' Root of Y(1) at', X
          WRITE (NOUT,99997) ' Solution is', (Y(I),I=1,N)
          IF (TOL.LT.0.0e0) WRITE (NOUT,*)
+           ' Over one-third steps controlled by HMAX'
      20 CONTINUE
      STOP
*
      99999 FORMAT (1X,A,e8.1)
      99998 FORMAT (1X,A,F7.4)
      99997 FORMAT (1X,A,3F13.5)
      END
*

```

```

SUBROUTINE FCN(T,Y,F)
*   .. Parameters ..
INTEGER      N
PARAMETER    (N=3)
*   .. Scalar Arguments ..
real         T
*   .. Array Arguments ..
real         F(N), Y(N)
*   .. Intrinsic Functions ..
INTRINSIC    COS, TAN
*   .. Executable Statements ..
F(1) = TAN(Y(3))
F(2) = -0.032e0*TAN(Y(3))/Y(2) - 0.02e0*Y(2)/COS(Y(3))
F(3) = -0.032e0/Y(2)**2
RETURN
END

*
real FUNCTION G(T,Y)
*   .. Parameters ..
INTEGER      N
PARAMETER    (N=3)
*   .. Scalar Arguments ..
real         T
*   .. Array Arguments ..
real         Y(N)
*   .. Executable Statements ..
G = Y(1)
RETURN
END

```

9.2 Program Data

None.

9.3 Program Results

D02BHF Example Program Results

Calculation with TOL = 0.1E-03

Root of Y(1) at 7.2884

Solution is 0.00000 0.47485 -0.76010

Calculation with TOL = 0.1E-04

Root of Y(1) at 7.2883

Solution is 0.00000 0.47486 -0.76011