NAG Fortran Library Routine Document

E04JYF

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

E04JYF is an easy-to-use quasi-Newton algorithm for finding a minimum of a function $F(x_1, x_2, \ldots, x_n)$, subject to fixed upper and lower bounds of the independent variables $x_1, x_2, \ldots, x_n$, using function values only.

It is intended for functions which are continuous and which have continuous first and second derivatives (although it will usually work even if the derivatives have occasional discontinuities).

2 Specification

```
SUBROUTINE E04JYF(N, IBOUND, FUNCT1, BL, BU, X, F, IW, LIW, W, LW,
1                  IUSER, USER, IFAIL)
INTEGER           N, IBOUND, IW(LIW), LIW, LW, IUSER(*), IFAIL
real              BL(N), BU(N), X(N), F, W(LW), USER(*)
EXTERNAL          FUNCT1
```

3 Description

This routine is applicable to problems of the form:

$$\text{Minimize} \quad F(x_1, x_2, \ldots, x_n) \quad \text{subject to} \quad l_j \le x_j \le u_j, \quad j = 1, 2, \ldots, n$$

when derivatives of $F(x)$ are unavailable.

Special provision is made for problems which actually have no bounds on the x_j, problems which have only non-negativity bounds and problems in which $l_1 = l_2 = \cdots = l_n$ and $u_1 = u_2 = \cdots = u_n$. The user must supply a subroutine to calculate the value of $F(x)$ at any point x.

From a starting point supplied by the user there is generated, on the basis of estimates of the gradient and the curvature of $F(x)$, a sequence of feasible points which is intended to converge to a local minimum of the constrained function. An attempt is made to verify that the final point is a minimum.

A typical iteration starts at the current point x where n_z (say) variables are free from both their bounds. The projected gradient vector g_z, whose elements are finite-difference approximations to the derivatives of $F(x)$ with respect to the free variables, is known. A unit lower triangular matrix L and a diagonal matrix D (both of dimension n_z), such that LDL^T is a positive-definite approximation of the matrix of second derivatives with respect to the free variables (i.e., the projected Hessian) are also held. The equations

$$LDL^T p_z = -g_z$$

are solved to give a search direction p_z, which is expanded to an n-vector p by an insertion of appropriate zero elements. Then α is found such that $F(x + \alpha p)$ is approximately a minimum (subject to the fixed bounds) with respect to α; x is replaced by $x + \alpha p$, and the matrices L and D are updated so as to be consistent with the change produced in the estimated gradient by the step αp. If any variable actually reaches a bound during the search along p, it is fixed and n_z is reduced for the next iteration. Most iterations calculate g_z using forward differences, but central differences are used when they seem necessary.

There are two sets of convergence criteria – a weaker and a stronger. Whenever the weaker criteria are satisfied, the Lagrange-multipliers are estimated for all the active constraints. If any Lagrange-multiplier estimate is significantly negative, then one of the variables associated with a negative Lagrange-multiplier estimate is released from its bound and the next search direction is computed in the extended subspace (i.e., n_z is increased). Otherwise minimization continues in the current subspace provided that this is practicable. When it is not, or when the stronger convergence criteria are already satisfied, then, if one or

more Lagrange-multiplier estimates are close to zero, a slight perturbation is made in the values of the corresponding variables in turn until a lower function value is obtained. The normal algorithm is then resumed from the perturbed point.

If a saddle point is suspected, a local search is carried out with a view to moving away from the saddle point. A local search is also performed when a point is found which is thought to be a constrained minimum.

4   References

Gill P E and Murray W (1976) Minimization subject to bounds on the variables *NPL Report NAC 72* National Physical Laboratory

5   Parameters

1:      N – INTEGER   *Input*

   On entry: the number n of independent variables.

   Constraint: $N \geq 1$.

2:      IBOUND – INTEGER   *Input*

   On entry: indicates whether the facility for dealing with bounds of special forms is to be used.

   It must be set to one of the following values:

   IBOUND = 0

       if the user will be supplying all the l_j and u_j individually.

   IBOUND = 1

       if there are no bounds on any x_j.

   IBOUND = 2

       if all the bounds are of the form $0 \leq x_j$.

   IBOUND = 3

       if $l_1 = l_2 = \ldots = l_n$ and $u_1 = u_2 = \ldots = u_n$.

3:      FUNCT1 – SUBROUTINE, supplied by the user.   *External Procedure*

   This routine must be supplied by the user to calculate the value of the function $F(x)$ at any point x. It should be tested separately before being used with E04JYF (see the E04 Chapter Introduction).

   Its specification is:

```
      SUBROUTINE FUNCT1(N, XC, FC, IUSER, USER)
      INTEGER           N, IUSER(*)
      real              XC(N), FC, USER(*)
```

   1:    N – INTEGER   *Input*

      On entry: the number n of variables.

   2:    XC(N) – *real* array   *Input*

      On entry: the point x at which the function value is required.

   3:    FC – *real*   *Output*

      On exit: the value of the function F at the current point x.

| 4: | IUSER(*) – INTEGER array | *User Workspace* |
| 5: | USER(*) – ***real*** array | *User Workspace* |

> FUNCT1 is called from E04JYF with the parameters IUSER and USER as supplied to E04JYF. The user is free to use the arrays IUSER and USER to supply information to FUNCT1 as an alternative to using COMMON.

FUNCT1 must be declared as EXTERNAL in the (sub)program from which E04JYF is called. Parameters denoted as *Input* must **not** be changed by this procedure.

4:      BL(N) – ***real*** array   *Input/Output*

On entry: the lower bounds l_j.

If IBOUND is set to 0, the user must set BL(j) to l_j, for $j = 1, 2, \ldots, n$. (If a lower bound is not specified for a particular x_j, the corresponding BL(j) should be set to -10^6.)

If IBOUND is set to 3, the user must set BL(1) to l_1; E04JYF will then set the remaining elements of BL equal to BL(1).

On exit: the lower bounds actually used by E04JYF.

5:      BU(N) – ***real*** array   *Input/Output*

On entry: the upper bounds u_j.

If IBOUND is set to 0, the user must set BU(j) to u_j, for $j = 1, 2, \ldots, n$. (If an upper bound is not specified for a particular x_j, the corresponding BU(j) should be set to 10^6.)

If IBOUND is set to 3, the user must set BU(1) to u_1; E04JYF will then set the remaining elements of BU equal to BU(1).

On exit: the upper bounds actually used by E04JYF.

6:      X(N) – ***real*** array   *Input/Output*

On entry: X(j) must be set to an estimate of the jth component of the position of the minimum, for $j = 1, 2, \ldots, n$.

On exit: the lowest point found during the calculations. Thus, if IFAIL $= 0$ on exit, X(j) is the jth component of the position of the minimum.

7:      F – ***real***   *Output*

On exit: the value of $F(x)$ corresponding to the final point stored in X.

8:      IW(LIW) – INTEGER array   *Output*

On exit: if IFAIL $= 0$, 3 or 5, the first N elements of IW contain information about which variables are currently on their bounds and which are free. Specifically, if x_i is

(a)   fixed on its upper bound, IW(i) is -1;

(b)   fixed on its lower bound, IW(i) is -2;

(c)   effectively a constant (i.e., $l_j = u_j$), IW(i) is -3;

(d)   free, IW(i) gives its position in the sequence of free variables.

In addition, IW(N + 1) contains the number of free variables (i.e., n_z). The rest of the array is used as workspace.

9:      LIW – INTEGER   *Input*

On entry: the dimension of the array IW as declared in the (sub)program from which E04JYF is called.

Constraint: LIW $\geq$ N + 2.

10: W(LW) – ***real*** array *Output*

On exit: if IFAIL = 0, 3 or 5, $W(i)$ contains a finite-difference approximation to the ith element of the projected gradient vector g_z, for $i = 1, 2, \ldots, N$. In addition, $W(N + 1)$ contains an estimate of the condition number of the projected Hessian matrix (i.e., k). The rest of the array is used as workspace.

11: LW – INTEGER *Input*

On entry: the dimension of the array W as declared in the (sub)program from which E04JYF is called.

Constraint: $LW \geq \max(N \times (N - 1)/2 + 12 \times N, 13)$.

12: IUSER(∗) – INTEGER array *User Workspace*

Note: the dimension of the array IUSER must be at least 1.

IUSER is not used by E04JYF, but is passed directly to FUNCT1 and may be used to pass information to that routine.

13: USER(∗) – ***real*** array *User Workspace*

Note: the dimension of the array USER must be at least 1.

USER is not used by E04JYF, but is passed directly to FUNCT1 and may be used to pass information to that routine.

14: IFAIL – INTEGER *Input/Output*

On entry: IFAIL must be set to 0, −1 or 1. Users who are unfamiliar with this parameter should refer to Chapter P01 for details.

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

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, because for this routine the values of the output parameters may be useful even if IFAIL ≠ 0 on exit, the recommended value is −1. **When the value −1 or 1 is used it is essential to test the value of IFAIL 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

On entry, N < 1,
or   IBOUND < 0,
or   IBOUND > 3,
or   IBOUND = 0 and $BL(j) > BU(j)$ for some j,
or   IBOUND = 3 and $BL(1) > BU(1)$,
or   LIW < N + 2,
or $LW < \max(13, 12 \times N + N \times (N - 1)/2)$.

IFAIL = 2

There have been $400 \times n$ function evaluations, yet the algorithm does not seem to be converging. The calculations can be restarted from the final point held in X. The error may also indicate that $F(x)$ has no minimum.

IFAIL = 3

> The conditions for a minimum have not all been met but a lower point could not be found and the algorithm has failed.

IFAIL = 4

> An overflow has occurred during the computation. This is an unlikely failure, but if it occurs the user should restart at the latest point given in X.

IFAIL = 5
IFAIL = 6
IFAIL = 7
IFAIL = 8

> There is some doubt about whether the point x found by E04JYF is a minimum. The degree of confidence in the result decreases as IFAIL increases. Thus, when IFAIL = 5 it is probable that the final x gives a good estimate of the position of a minimum, but when IFAIL = 8 it is very unlikely that the routine has found a minimum.

IFAIL = 9

> In the search for a minimum, the modulus of one of the variables has become very large ($\sim 10^6$). This indicates that there is a mistake in FUNCT1, that the user's problem has no finite solution, or that the problem needs rescaling (see Section 8).

IFAIL = 10

> The computed set of forward-difference intervals (stored in $W(9 * N + 1), W(9 * N + 2), \ldots,$ $W(10 * N)$) is such that $X(i) + W(9 * N + i) \leq X(i)$ for some i.
>
> This is an unlikely failure, but if it occurs the user should attempt to select another starting point.

If the user is dissatisfied with the result (e.g., because IFAIL = 5, 6, 7 or 8), it is worth restarting the calculations from a different starting point (not the point at which the failure occurred) in order to avoid the region which caused the failure. If persistent trouble occurs and the gradient can be calculated, it may be advisable to change to a routine which uses gradients (see the E04 Chapter Introduction).

7 Accuracy

A successful exit (IFAIL = 0) is made from E04JYF when (B1, B2 and B3) or B4 hold, and the local search confirms a minimum, where

$$B1 \equiv \alpha^{(k)} \times \|p^{(k)}\| < (x_{tol} + \sqrt{\epsilon}) \times (1.0 + \|x^{(k)}\|)$$

$$B2 \equiv |F^{(k)} - F^{(k-1)}| < (x_{tol}^2 + \epsilon) \times (1.0 + |F^{(k)}|)$$

$$B3 \equiv \|g_z^{(k)}\| < (\epsilon^{1/3} + x_{tol}) \times (1.0 + |F^{(k)}|)$$

$$B4 \equiv \|g_z^{(k)}\| < 0.01 \times \sqrt{\epsilon}.$$

(Quantities with superscript k are the values at the kth iteration of the quantities mentioned in Section 3, $x_{tol} = 100\sqrt{\epsilon}$, ϵ is the **machine precision** and $\|.\|$ denotes the Euclidean norm. The vector g_z is returned in the array W.)

If IFAIL = 0, then the vector in X on exit, x_{sol}, is almost certainly an estimate of the position of the minimum, x_{true}, to the accuracy specified by x_{tol}.

If IFAIL = 3 or 5, x_{sol} may still be a good estimate of x_{true}, but the following checks should be made. Let k denote an estimate of the condition number of the projected Hessian matrix at x_{sol}. (The value of k is returned in $W(N + 1)$). If

(i)   the sequence $\{F(x^{(k)})\}$ converges to $F(x_{sol})$ at a superlinear or a fast linear rate,

(ii) $\|g_z(x_{xol})\|^2 < 10.0 \times \epsilon$, and

(iii) $k < 1.0 / \| g_z(x_{sol}) \|$,

then it is almost certain that x_{sol} is a close approximation to the position of a minimum. When (ii) is true, then usually $F(x_{sol})$ is a close approximation to $F(x_{true})$.

When a successful exit is made then, for a computer with a mantissa of t decimals, one would expect to get about $t/2 - 1$ decimals accuracy in x and about $t - 1$ decimals accuracy in F, provided the problem is reasonably well scaled.

8 Further Comments

The number of iterations required depends on the number of variables, the behaviour of $F(x)$ and the distance of the starting point from the solution. The number of operations performed in an iteration of E04JYF is roughly proportional to n^2. In addition, each iteration makes at least $m + 1$ calls of FUNCT1, where m is the number of variables not fixed on bounds. So, unless $F(x)$ can be evaluated very quickly, the run time will be dominated by the time spent in FUNCT1.

Ideally the problem should be scaled so that at the solution the value of $F(x)$ and the corresponding values of $x_1, x_2, \ldots, x_n$ are each in the range $(-1, +1)$, and so that at points a unit distance away from the solution, F is approximately a unit value greater than at the minimum. It is unlikely that the user will be able to follow these recommendations very closely, but it is worth trying (by guesswork), as sensible scaling will reduce the difficulty of the minimization problem, so that E04JYF will take less computer time.

9 Example

To minimize

$$F = (x_1 + 10x_2)^2 + 5(x_3 - x_4)^2 + (x_2 - 2x_3)^4 + 10(x_1 - x_4)^4$$

subject to

$$
\begin{aligned}
1 &\le x_1 \le 3 \\
-2 &\le x_2 \le 0 \\
1 &\le x_4 \le 3,
\end{aligned}
$$

starting from the initial guess $(3, -1, 0, 1)$.

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.

```
*       E04JYF Example Program Text.
*       Mark 18 Release. NAG Copyright 1997.
*       .. Parameters ..
        INTEGER           N, LIW, LW
        PARAMETER         (N=4,LIW=N+2,LW=N*(N-1)/2+12*N)
        INTEGER           NOUT
        PARAMETER         (NOUT=6)
*       .. Local Scalars ..
        real              F
        INTEGER           IBOUND, IFAIL, J
*       .. Local Arrays ..
        real              BL(N), BU(N), USER(N), W(LW), X(N)
        INTEGER           IUSER(N), IW(6)
*       .. External Subroutines ..
        EXTERNAL          E04JYF, FUNCT1
*       .. Executable Statements ..
        WRITE (NOUT,*) 'E04JYF Example Program Results'
        X(1) = 3.0e0
        X(2) = -1.0e0
        X(3) = 0.0e0
        X(4) = 1.0e0
        IBOUND = 0
```

```
      BL(1) = 1.0e0
      BU(1) = 3.0e0
      BL(2) = -2.0e0
      BU(2) = 0.0e0
*
*     X(3) is unconstrained, so we set BL(3) to a large negative
*     number and BU(3) to a large positive number.
*
      BL(3) = -1.0e6
      BU(3) = 1.0e6
      BL(4) = 1.0e0
      BU(4) = 3.0e0
      IFAIL = 1
*
      CALL E04JYF(N,IBOUND,FUNCT1,BL,BU,X,F,IW,LIW,W,LW,IUSER,USER,
     +            IFAIL)
*
      IF (IFAIL.NE.0) THEN
         WRITE (NOUT,*)
         WRITE (NOUT,99999) 'Error exit type', IFAIL,
     +    ' - see routine document'
      END IF
      IF (IFAIL.NE.1) THEN
         WRITE (NOUT,*)
         WRITE (NOUT,99998) 'Function value on exit is ', F
         WRITE (NOUT,99997) 'at the point', (X(J),J=1,N)
      END IF
      STOP
*
99999 FORMAT (1X,A,I3,A)
99998 FORMAT (1X,A,F8.4)
99997 FORMAT (1X,A,4F9.4)
      END
*
      SUBROUTINE FUNCT1(N,XC,FC,IUSER,USER)
*     Routine to evaluate objective function.
*     .. Scalar Arguments ..
      real               FC
      INTEGER            N
*     .. Array Arguments ..
      real               USER(*), XC(N)
      INTEGER            IUSER(*)
*     .. Local Scalars ..
      real               X1, X2, X3, X4
*     .. Executable Statements ..
      X1 = XC(1)
      X2 = XC(2)
      X3 = XC(3)
      X4 = XC(4)
      FC = (X1+10.0e0*X2)**2 + 5.0e0*(X3-X4)**2 + (X2-2.0e0*X3)**4 +
     +     10.0e0*(X1-X4)**4
      RETURN
      END
```

9.2 Program Data

None.

9.3 Program Results

```
 E04JYF Example Program Results

 Error exit type  5 - see routine document

 Function value on exit is   2.4338
 at the point   1.0000  -0.0852   0.4093   1.0000
```