

NAG Fortran Library Routine Document

E04XAF/E04XAA

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

E04XAF/E04XAA computes an approximation to the gradient vector and/or the Hessian matrix for use in conjunction with, or following the use of an optimization routine (such as E04UCF/E04UCA).

E04XAA is a version of E04XAF that has additional parameters in order to make it safe for use in multithreaded applications (see Section 5 below). The initialisation routine E04WBF **must** have been called prior to calling E04XAA.

2 Specifications

2.1 Specification for E04XAF

```

SUBROUTINE E04XAF(MSGLVL, N, EPSRF, X, MODE, OBJFUN, LHES, HFORW, OBJF,
1              OBJGRD, HCNTRL, HESIAN, IWARN, WORK, IUSER, USER,
2              INFO, IFAIL)
  INTEGER      MSGLVL, N, MODE, LHES, IWARN, IUSER(*), INFO(N), IFAIL
  real        EPSRF, X(N), HFORW(N), OBJF, OBJGRD(N), HCNTRL(N),
1              HESIAN(LHES,*), WORK(*), USER(*)
  EXTERNAL    OBJFUN

```

2.2 Specification for E04XAA

```

SUBROUTINE E04XAA(MSGLVL, N, EPSRF, X, MODE, OBJFUN, LHES, HFORW, OBJF,
1              OBJGRD, HCNTRL, HESIAN, IWARN, WORK, IUSER, USER,
2              INFO, LWSAV, IWSAV, RWSAV, IFAIL)
  INTEGER      MSGLVL, N, MODE, LHES, IWARN, IUSER(*), INFO(N),
1              IWSAV(610), IFAIL
  real        EPSRF, X(N), HFORW(N), OBJF, OBJGRD(N), HCNTRL(N),
1              HESIAN(LHES,*), WORK(*), USER(*), RWSAV(475)
  LOGICAL      LWSAV(120)
  EXTERNAL    OBJFUN

```

Before calling E04XAA, routine E04WBF **must** be called. The specification for E04WBF is:

```

SUBROUTINE E04WBF(RNAME, CWSAV, LCWSAV, LWSAV, LLWSAV, IWSAV, LIWSAV,
1              RWSAV, LRWSAV, IFAIL)
  INTEGER      LCWSAV, LLWSAV, IWSAV(LIWSAV), LIWSAV, LRWSAV, IFAIL
  real        RWSAV(LRWSAV)
  LOGICAL      LWSAV(LLWSAV)
  CHARACTER*6  RNAME
  CHARACTER*80 CWSAV(LCWSAV)

```

E04WBF should be called with RNAME = 'E04XAA'. LCWSAV, LLWSAV, LIWSAV and LRWSAV, the declared lengths of CWSAV, LWSAV, IWSAV and RWSAV respectively, must satisfy:

$$LCWSAV \geq 1$$

$$LLWSAV \geq 120$$

$$LIWSAV \geq 610$$

$$LRWSAV \geq 475$$

The contents of the arrays CWSAV, LWSAV, IWSAV and RWSAV **must not** be altered between calling routines E04WBF and E04XAA.

3 Description

This routine is similar to routine FDCALC described in Gill *et al.* (1983). It should be noted that this routine aims to compute sufficiently accurate estimates of the derivatives for use with an optimization algorithm. Users who require more accurate estimates should refer to Chapter D04.

E04XAF/E04XAA computes finite-difference approximations to the gradient vector and the Hessian matrix for a given function. The simplest approximation involves the forward-difference formula, in which the derivative $f'(x)$ of a univariate function $f(x)$ is approximated by the quantity

$$\rho_F(f, h) = \frac{f(x+h) - f(x)}{h}$$

for some interval $h > 0$, where the subscript 'F' denotes 'forward-difference' (see Gill *et al.* (1981)).

To summarize the procedure used by E04XAF/E04XAA (for the case when the objective function is available and the user requires estimates of gradient values and Hessian matrix diagonal values, i.e., $\text{MODE} = 0$) consider a univariate function f at the point x . (In order to obtain the gradient of a multivariate function $F(x)$, where x is an n -vector, the procedure is applied to each component of x , keeping the other components fixed.) Roughly speaking, the method is based on the fact that the bound on the relative truncation error in the forward-difference approximation tends to be an increasing function of h , while the relative condition error bound is generally a decreasing function of h , hence changes in h will tend to have opposite effects on these errors (see Gill *et al.* (1981)).

The 'best' interval h is given by

$$h_F = 2\sqrt{\frac{(1 + |f(x)|)e_R}{|\Phi|}} \quad (1)$$

where Φ is an estimate of $f''(x)$, and e_R is an estimate of the relative error associated with computing the function (see Chapter 8 of Gill *et al.* (1981)). Given an interval h , Φ is defined by the second-order approximation

$$\Phi = \frac{f(x+h) - 2f(x) + f(x-h)}{h^2}.$$

The decision as to whether a given value of Φ is acceptable involves $\hat{c}(\Phi)$, the following bound on the relative condition error in Φ :

$$\hat{c}(\Phi) = \frac{4e_R(1 + |f|)}{h^2|\Phi|}$$

(When Φ is zero, $\hat{c}(\Phi)$ is taken as an arbitrary large number.)

The procedure selects the interval h_ϕ (to be used in computing Φ) from a sequence of trial intervals (h_k) . The initial trial interval is taken as $10\bar{h}$, where

$$\bar{h} = 2(1 + |x|)\sqrt{e_R}$$

unless the user specifies the initial value to be used.

The value of $\hat{c}(\Phi)$ for a trial value h_k is defined as 'acceptable' if it lies in the interval $[0.001, 0.1]$. In this case h_ϕ is taken as h_k , and the current value of Φ is used to compute h_F from (1). If $\hat{c}(\Phi)$ is unacceptable, the next trial interval is chosen so that the relative condition error bound will either decrease or increase, as required. If the bound on the relative condition error is too large, a larger interval is used as the next trial value in an attempt to reduce the condition error bound. On the other hand, if the relative condition error bound is too small, h_k is reduced.

The procedure will fail to produce an acceptable value of $\hat{c}(\Phi)$ in two situations. Firstly, if $f''(x)$ is extremely small, then $\hat{c}(\Phi)$ may never become small, even for a very large value of the interval. Alternatively, $\hat{c}(\Phi)$ may never exceed 0.001, even for a very small value of the interval. This usually implies that $f''(x)$ is extremely large, and occurs most often near a singularity.

As a check on the validity of the estimated first derivative, the procedure provides a comparison of the forward-difference approximation computed with h_F (as above) and the central-difference approximation computed with h_ϕ . Using the central-difference formula the first derivative can be approximated by

$$\rho_c(f, h) = \frac{f(x+h) - f(x-h)}{2h}$$

where $h > 0$. If the values h_F and h_ϕ do not display some agreement, neither can be considered reliable.

When both function and gradients are available and the user requires the Hessian matrix (i.e., MODE = 1) E04XAF/E04XAA follows a similar procedure to the case above with the exception that the gradient function $g(x)$ is substituted for the objective function and so the forward-difference interval for the first derivative of $g(x)$ with respect to variable x_j is computed. The j th column of the approximate Hessian matrix is then defined as in Chapter 2 of Gill *et al.* (1981), by

$$\frac{g(x + h_j e_j) - g(x)}{h_j}$$

where h_j is the best forward-difference interval associated with the j th component of g and e_j is the vector with unity in the j th position and zeros elsewhere.

When only the objective function is available and the user requires the gradients and Hessian matrix (i.e., MODE = 2) E04XAF/E04XAA again follows the same procedure as the case for MODE = 0 except that this time the value of $\hat{c}(\Phi)$ for a trial value h_k is defined as acceptable if it lies in the interval $[0.0001, 0.01]$ and the initial trial interval is taken as

$$\bar{h} = 2(1 + |x|)\sqrt[4]{e_R}.$$

The approximate Hessian matrix G is then defined as in Chapter 2 of Gill *et al.* (1981), by

$$G_{ij}(x) = \frac{1}{h_i h_j} (f(x + h_i e_i + h_j e_j) - f(x + h_i e_i) - f(x + h_j e_j) + f(x)).$$

4 References

Gill P E, Murray W, Saunders M A and Wright M H (1983) Documentation for FDCALC and FDCORE *Technical Report SOL 83-6* Stanford University

Gill P E, Murray W, Saunders M A and Wright M H (1981) Computing forward-difference intervals for numerical optimization *SIAM J. Sci. Statist. Comput.* **4** 310–321

Gill P E, Murray W and Wright M H (1981) *Practical Optimization* Academic Press

5 Parameters

1: MSGLVL – INTEGER *Input*

On entry: MSGLVL must indicate the amount of intermediate output desired (see Section 8.1 for a description of the printed output). All output is written on the current advisory message unit (see X04ABF).

Value Definition

0 No printout
1 A summary is printed out for each variable plus any warning messages.
Other Values other than 0 and 1 should normally be used **only** at the direction of NAG.

2: N – INTEGER *Input*

On entry: the number n of independent variables.

Constraint: $N \geq 1$.

3: EPSRF – *real* *Input*

On entry: EPSRF must define e_R , which is intended to be a measure of the accuracy with which the problem function F can be computed. The value of e_R should reflect the relative precision of $1 + |F(x)|$, i.e., acts as a relative precision when $|F|$ is large, and as an absolute precision when $|F|$

is small. For example, if $F(x)$ is typically of order 1000 and the first six significant digits are known to be correct, an appropriate value for e_R would be $1.0\text{E}-6$.

A discussion of EPSRF is given in Chapter 8 of Gill *et al.* (1981). If EPSRF is either too small or too large on entry a warning will be printed if MSGVLV = 1, the parameter IWARN set to the appropriate value on exit and E04XAF/E04XAA will use a default value of $e_M^{0.9}$, where e_M is the *machine precision*.

If $\text{EPSRF} \leq 0.0$ on entry then E04XAF/E04XAA will use the default value internally. The default value will be appropriate for most simple functions that are computed with full accuracy.

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

On entry: the point x at which the derivatives are to be computed.

On exit: used as internal workspace prior to being restored and hence is unchanged.

- 5: MODE – INTEGER *Input/Output*

On entry: indicates which derivatives are required.

MODE = 0

The user requests the gradient and Hessian diagonal values having supplied the objective function via the routine OBJFUN.

MODE = 1

The user requests the Hessian matrix having supplied both the objective function and gradients via the routine OBJFUN.

MODE = 2

The user requests the gradient values and Hessian matrix having supplied the objective function via the routine OBJFUN.

On exit: MODE is changed **only** if the user sets MODE negative in OBJFUN, i.e., the user requests termination of E04XAF/E04XAA.

- 6: OBJFUN – SUBROUTINE, supplied by the user. *External Procedure*

If MODE = 0 or 2 then OBJFUN must calculate the objective function; otherwise if MODE = 1, OBJFUN must calculate the objective function and the gradients.

Its specification is:

<pre> SUBROUTINE OBJFUN(MODE, N, X, OBJF, OBJGRD, NSTATE, IUSER, USER) INTEGER MODE, N, NSTATE, IUSER(*) real X(N), OBJF, OBJGRD(N), USER(*) </pre>	
1:	<p>MODE – INTEGER <i>Input/Output</i></p> <p>MODE indicates which parameter values within OBJFUN need to be set.</p> <p><i>On entry:</i> to OBJFUN, MODE is always set to the value that the user set it to before the call to E04XAF/E04XAA.</p> <p><i>On exit:</i> its value must not be altered unless the user wishes to indicate a failure within OBJFUN, in which case it should be set to a negative value. If MODE is negative on exit from OBJFUN, the execution of E04XAF/E04XAA is terminated with IFAIL set to MODE.</p>
2:	<p>N – INTEGER <i>Input</i></p> <p><i>On entry:</i> the number n of variables as input to E04XAF/E04XAA.</p>

3:	X(N) – <i>real</i> array	<i>Input</i>
	<i>On entry:</i> the point x at which the objective function (and gradients if $\text{MODE} = 1$) is to be evaluated.	
4:	OBJF – <i>real</i>	<i>Output</i>
	<i>On exit:</i> OBJF must be set to the value of the objective function.	
5:	OBJGRD(N) – <i>real</i> array	<i>Output</i>
	<i>On exit:</i> if $\text{MODE} = 1$, OBJGRD(j) must contain the value of the first derivative with respect to x . If $\text{MODE} \neq 1$ then OBJGRD need not be set.	
6:	NSTATE – INTEGER	<i>Input</i>
	<i>On entry:</i> NSTATE will be set to 1 on the first call of OBJFUN by E04XAF/E04XAA, and is 0 for all subsequent calls. Thus, if the user wishes, NSTATE may be tested within OBJFUN in order to perform certain calculations once only. For example the user may read data.	
7:	IUSER(*) – INTEGER array	<i>User Workspace</i>
8:	USER(*) – <i>real</i> array	<i>User Workspace</i>
	OBJFUN is called from E04XAF/E04XAA with the parameters IUSER and USER as supplied to E04XAF/E04XAA. The user is free to use arrays IUSER and USER to supply information to OBJFUN as an alternative to using COMMON.	

OBJFUN must be declared as EXTERNAL in the (sub)program from which E04XAF/E04XAA is called. Parameters denoted as *Input* must **not** be changed by this procedure.

- 7: LHES – INTEGER *Input*
- On entry:* the first dimension of the array HESIAN as declared in the (sub)program from which E04XAF/E04XAA is called.
- Constraint:* LHES \geq N.
- 8: HFORW(N) – *real* array *Input/Output*
- On entry:* the initial trial interval for computing the appropriate partial derivative to the j th variable. If $\text{HFORW}(j) \leq 0.0$, then the initial trial interval is computed by E04XAF/E04XAA (see Section 3).
- On exit:* HFORW(j) is the best interval found for computing a forward-difference approximation to the appropriate partial derivative for the j th variable.
- 9: OBJF – *real* *Output*
- On exit:* the value of the objective function evaluated at the input vector in X.
- 10: OBJGRD(N) – *real* array *Output*
- On exit:* if $\text{MODE} = 0$ or 2 , OBJGRD(j) contains the best estimate of the first partial derivative for the j th variable. If $\text{MODE} = 1$, OBJGRD(j) contains the first partial derivative for the j th variable evaluated at the input vector in X.
- 11: HCNTRL(N) – *real* array *Output*
- On exit:* HCNTRL(j) is the best interval found for computing a central-difference approximation to the appropriate partial derivative for the j th variable.

- 12: HESIAN(LHES,*) – *real* array Output
Note: the second dimension of the array HESIAN must be at least 1 if MODE = 0, or N if MODE = 1 or 2.
On exit: if MODE = 0 the estimated Hessian diagonal elements are contained in the first column of this array.
 If MODE = 1 or 2, the estimated Hessian matrix is contained in the leading n by n part of this array.
- 13: IWARN – INTEGER Output
On exit: IWARN = 0 on successful exit.
 If the value of EPSRF on entry is too small or too large then IWARN is set to 1 or 2 respectively on exit and EPSRF is set to the default value within E04XAF/E04XAA.
 If MSGGLVL > 0 then warnings will be printed if EPSRF is too small or too large.
- 14: WORK(*) – *real* array Workspace
Note: the first dimension of the array WORK must be at least N if MODE = 0 and at least $N * (N + 1)$ if MODE = 1 or 2.
- 15: IUSER(*) – INTEGER array User Workspace
 This array is not used by E04XAF/E04XAA, but is passed directly to routine OBJFUN and may be used to supply information to OBJFUN.
- 16: USER(*) – *real* array User Workspace
 This array is not used by E04XAF/E04XAA, but is passed directly to routine OBJFUN and may be used to supply information to OBJFUN.
- 17: INFO(N) – INTEGER array Output
On exit: INFO(j) represents diagnostic information on variable j . (See Section 6 for more details.)
- 18: IFAIL – INTEGER Input/Output
Note: for E04XAA, IFAIL does not occur in this position in the parameter list. See the additional parameters described below.
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, for users not familiar with this parameter the recommended value is 0. **When the value -1 or 1 is used it is essential to test the value of IFAIL on exit.**
- Note:** the following are additional parameters for specific use with E04XAA. Users of E04XAF therefore need not read the remainder of this section.
- 18: LWSAV(120) – LOGICAL array Workspace
 19: IWSAV(610) – INTEGER array Workspace
 20: RWSAV(475) – *real* array Workspace
 The arrays LWSAV, IWSAV and RWSAV **must not** be altered between calls to either of the routines E04WBF and E04XAA.
- 21: IFAIL – INTEGER Input/Output
Note: see the parameter description for IFAIL above.

6 Error Indicators and Warnings

On exit from E04XAF/E04XAA both diagnostic parameters INFO and IFAIL should be tested. IFAIL represents an overall diagnostic indicator, whereas the integer array INFO represents diagnostic information on each variable.

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:

$IFAIL < 0$

A negative value of IFAIL indicates an exit from E04XAF/E04XAA because the user set MODE negative in routine OBJFUN. The value of IFAIL will be the same as the user's setting of MODE.

$IFAIL = 1$

On entry, one or more of the following conditions are satisfied: $N < 1$, $LHES < N$ or MODE is invalid.

$IFAIL = 2$

One or more variables have a non-zero INFO value. This may not necessarily represent an unsuccessful exit – see diagnostic information on INFO.

Diagnostic information returned via INFO is as follows:

$INFO(i) = 1$

The appropriate function appears to be constant. $HFORW(i)$ is set to the initial trial interval value (see Section 3) corresponding to a well-scaled problem and Error est. in the printed output is set to zero. This value occurs when the estimated relative condition error in the first derivative approximation is unacceptably large for every value of the finite-difference interval. If this happens when the function is not constant the initial interval may be too small; in this case, it may be worthwhile to rerun E04XAF/E04XAA with larger initial trial interval values supplied in HFORW (see Section 3). This error may also occur if the function evaluation includes an inordinately large constant term or if EPSRF is too large.

$INFO(i) = 2$

The appropriate function appears to be linear or odd. $HFORW(i)$ is set to the smallest interval with acceptable bounds on the relative condition error in the forward- and backward-difference estimates. In this case, the estimated relative condition error in the second derivative approximation remained large for every trial interval, but the estimated error in the first derivative approximation was acceptable for at least one interval. If the function is not linear or odd the relative condition error in the second derivative may be decreasing very slowly, it may be worthwhile to rerun E04XAF/E04XAA with larger initial trial interval values supplied in HFORW (see Section 3).

$INFO(i) = 3$

The second derivative of the appropriate function appears to be so large that it cannot be reliably estimated (i.e., near a singularity). $HFORW(i)$ is set to the smallest trial interval.

This value occurs when the relative condition error estimate in the second derivative remained very small for every trial interval.

If the second derivative is not large the relative condition error in the second derivative may be increasing very slowly. It may be worthwhile to rerun E04XAF/E04XAA with smaller initial trial interval values supplied in HFORW (see Section 3). This error may also occur when the given value of EPSRF is not a good estimate of a bound on the absolute error in the appropriate function (i.e., EPSRF is too small).

$\text{INFO}(i) = 4$

The algorithm terminated with an apparently acceptable estimate of the second derivative. However the forward-difference estimates of the appropriate first derivatives (computed with the final estimate of the ‘optimal’ forward-difference interval) and the central difference estimates (computed with the interval used to compute the final estimate of the second derivative) do not agree to half a decimal place. The usual reason that the forward- and central-difference estimates fail to agree is that the first derivative is small.

If the first derivative is not small, it may be helpful to execute the procedure at a different point.

7 Accuracy

IFAIL contains 0 on exit if the algorithm terminated successfully, i.e., the forward-difference estimates of the appropriate first derivatives (computed with the final estimate of the ‘optimal’ forward-difference interval h_F) and the central-difference estimates (computed with the interval h_ϕ used to compute the final estimate of the second derivative) agree to at least half a decimal place.

In short word length implementations when computing the full Hessian matrix given function values only (i.e., $\text{MODE} = 2$) the elements of the computed Hessian will have at best 1 – 2 figures of accuracy.

8 Further Comments

To evaluate an acceptable set of finite-difference intervals for a well-scaled problem, the routine will require around two function evaluations per variable; in a badly scaled problem however, as many as six function evaluations per variable may be needed.

If the user requests the full Hessian matrix supplying both function and gradients (i.e., $\text{MODE} = 1$) or function only (i.e., $\text{MODE} = 2$) then a further N or $3 \times N \times (N + 1)/2$ function evaluations respectively are required.

8.1 Description of the Printed Output

The following is a description of the printed output from E04XAF/E04XAA as controlled by the parameter MSGLVL.

Output when $\text{MSGLVL} = 1$ is as follows:

J	number of variable for which the difference interval has been computed.
$X(j)$	j th variable of x as set by the user.
F. dif. int.	the best interval found for computing a forward-difference approximation to the appropriate partial derivative with respect to the j th variable.
C. dif. int.	the best interval found for computing a central-difference approximation to the appropriate partial derivative with respect to the j th variable.
Error est.	a bound on the estimated error in the final forward-difference approximation. When $\text{INFO}(j) = 1$, Error est. is set to zero.
Grad. est.	best estimate of the first partial derivative with respect to the j th variable.
Hess diag est.	best estimate of the second partial derivative with respect to the j th variable.
fun evals.	the number of function evaluations used to compute the final difference intervals for the j th variable.
$\text{info}(j)$	the value of INFO for the j th variable.

9 Example

To compute the gradient vector and the Hessian matrix of the following function:

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

at the point (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.

Note: the following program illustrates the use of E04XAF. An equivalent program illustrating the use of E04XAA is available with the supplied Library and is also available from the NAG web site.

```
*      E04XAF Example Program Text
*      Mark 14 Revised.  NAG Copyright 1989.
*      .. Parameters ..
      INTEGER          N, LHES, LWORK
      PARAMETER        (N=4,LHES=N,LWORK=N*N+N)
      INTEGER          NOUT
      PARAMETER        (NOUT=6)
*      .. Local Scalars ..
      real             EPSRF, OBJF
      INTEGER          I, IFAIL, IMODE, IWARN, J, MODE, MSGLVL
*      .. Local Arrays ..
      real             HCNTL(N), HESIAN(LHES,N), HFORW(N), OBJGRD(N),
+                     USER(1), WORK(LWORK), X(N)
      INTEGER          INFO(N), IUSER(1)
*      .. External Subroutines ..
      EXTERNAL         E04XAF, OBJFUN
*      .. Executable Statements ..
      WRITE (NOUT,*) 'E04XAF Example Program Results'
      MSGLVL = 0
*      Set the point at which the derivatives are to be estimated.
      X(1) = 3.0e0
      X(2) = -1.0e0
      X(3) = 0.0e0
      X(4) = 1.0e0
*      Take default value of EPSRF.
      EPSRF = -1.0e0
*      Illustrate the different values of MODE.
      DO 40 IMODE = 0, 2
        MODE = IMODE
        WRITE (NOUT,*)
        IF (MODE.EQ.0) THEN
          WRITE (NOUT,*)
+          'Find gradients and Hessian diagonals given function only'
          WRITE (NOUT,*) '( i.e. MODE = 0 ).'
        ELSE IF (MODE.EQ.1) THEN
          WRITE (NOUT,*)
+          'Find Hessian matrix given function and gradients'
          WRITE (NOUT,*) '( i.e. MODE = 1 ).'
        ELSE IF (MODE.EQ.2) THEN
          WRITE (NOUT,*)
+          'Find gradients and Hessian matrix given function only'
          WRITE (NOUT,*) '( i.e. MODE = 2 ).'
        END IF
*      Set HFORW(I) = -1.0 so that E04XAF computes the initial trial
*      interval.
      DO 20 I = 1, N
        HFORW(I) = -1.0e0
20    CONTINUE
      IFAIL = 1
*
      CALL E04XAF(MSGLVL,N,EPSRF,X,MODE,OBJFUN,LHES,HFORW,OBJF,
+              OBJGRD,HCNTL,HESIAN,IWARN,WORK,IUSER,USER,INFO,
+              IFAIL)
```

```

*
      IF (IFAIL.EQ.0 .OR. IFAIL.EQ.2) THEN
        WRITE (NOUT,99999) 'Function value is ', OBJF
        IF (MODE.EQ.1) THEN
          WRITE (NOUT,*) 'Gradient vector is'
          WRITE (NOUT,99998) (OBJGRD(I),I=1,N)
        ELSE
          WRITE (NOUT,*) 'Estimated gradient vector is'
          WRITE (NOUT,99998) (OBJGRD(I),I=1,N)
        END IF
        IF (MODE.EQ.0) THEN
          WRITE (NOUT,*) 'Estimated Hessian matrix diagonal is'
          WRITE (NOUT,99998) (HESIAN(I,1),I=1,N)
        ELSE
          WRITE (NOUT,*)
+           'Estimated Hessian matrix (machine dependent) is'
          WRITE (NOUT,99998) ((HESIAN(I,J),J=1,N),I=1,N)
        END IF
      ELSE
        WRITE (NOUT,*)
        WRITE (NOUT,99997) 'On exit from E04XAF IFAIL = ', IFAIL
      END IF
40  CONTINUE
      STOP
*
99999 FORMAT (1X,A,1P,e12.4)
99998 FORMAT (4(1X,1P,e12.4))
99997 FORMAT (1X,A,I2)
      END
*
      SUBROUTINE OBJFUN(MODE,N,X,OBJF,OBJGRD,NSTATE,IUSER,USER)
*
      .. Scalar Arguments ..
      real          OBJF
      INTEGER        MODE, N, NSTATE
*
      .. Array Arguments ..
      real          OBJGRD(N), USER(*), X(N)
      INTEGER        IUSER(*)
*
      .. Local Scalars ..
      real          A, B, C, D
*
      .. Executable Statements ..
      A = X(1) + 10.0e0*X(2)
      B = X(3) - X(4)
      C = X(2) - 2.0e0*X(3)
      D = X(1) - X(4)
      OBJF = A**2 + 5.0e0*B**2 + C**4 + 10.0e0*D**4
      IF (MODE.EQ.1) THEN
        OBJGRD(1) = 4.0e1*X(1)**3 + 2.0e0*X(1) - 1.2e2*X(4)*X(1)**2 +
+         1.2e2*X(1)*X(4)**2 + 2.0e1*X(2) - 4.0e1*X(4)**3
        OBJGRD(2) = 2.0e2*X(2) + 2.0e1*X(1) + 4.0e0*X(2)**3 +
+         4.8e1*X(2)*X(3)**2 - 2.4e1*X(3)*X(2)**2 -
+         32.0e0*X(3)**3
        OBJGRD(3) = 1.0e1*X(3) - 1.0e1*X(4) - 8.0e0*X(2)**3 +
+         4.8e1*X(3)*X(2)**2 - 9.6e1*X(2)*X(3)**2 +
+         6.4e1*X(3)**3
        OBJGRD(4) = 1.0e1*X(4) - 1.0e1*X(3) - 4.0e1*X(1)**3 +
+         1.2e2*X(4)*X(1)**2 - 1.2e2*X(1)*X(4)**2 +
+         4.0e1*X(4)**3
      END IF
      RETURN
      END

```

9.2 Program Data

None.

9.3 Program Results

E04XAF Example Program Results

Find gradients and Hessian diagonals given function only
(i.e. MODE = 0).

Function value is 2.1500E+02

Estimated gradient vector is

3.0600E+02 -1.4400E+02 -2.0000E+00 -3.1000E+02

Estimated Hessian matrix diagonal is

4.8200E+02 2.1200E+02 5.8009E+01 4.9001E+02

Find Hessian matrix given function and gradients

(i.e. MODE = 1).

Function value is 2.1500E+02

Gradient vector is

3.0600E+02 -1.4400E+02 -2.0000E+00 -3.1000E+02

Estimated Hessian matrix (machine dependent) is

4.8200E+02 2.0000E+01 0.0000E+00 -4.8000E+02

2.0000E+01 2.1200E+02 -2.4000E+01 0.0000E+00

0.0000E+00 -2.4000E+01 5.8000E+01 -1.0000E+01

-4.8000E+02 0.0000E+00 -1.0000E+01 4.9000E+02

Find gradients and Hessian matrix given function only

(i.e. MODE = 2).

Function value is 2.1500E+02

Estimated gradient vector is

3.0600E+02 -1.4400E+02 -2.0000E+00 -3.1000E+02

Estimated Hessian matrix (machine dependent) is

4.8200E+02 2.0000E+01 0.0000E+00 -4.7999E+02

2.0000E+01 2.1200E+02 -2.4013E+01 0.0000E+00

0.0000E+00 -2.4013E+01 5.7980E+01 -1.0007E+01

-4.7999E+02 0.0000E+00 -1.0007E+01 4.9000E+02