

NAG Library Routine Document

e04fff

Note: this routine uses **optional parameters** to define choices in the problem specification and in the details of the algorithm. If you wish to use default settings for all of the optional parameters, you need only read Sections 1 to 10 of this document. If, however, you wish to reset some or all of the settings please refer to Section 11 for a detailed description of the algorithm and to Section 12 for a detailed description of the specification of the optional parameters.

1 Purpose

e04fff is a derivative free solver from the NAG optimization modelling suite for small to medium-scale nonlinear least squares problems with bound constraints.

2 Specification

Fortran Interface

```
Subroutine e04fff (handle, objfun, mon, nvar, x, nres, rx, rinfo, stats,      &
                 iuser, ruser, cpuser, ifail)
```

```
Integer, Intent (In)           :: nvar, nres
Integer, Intent (Inout)        :: iuser(*), ifail
Real (Kind=nag_wp), Intent (Inout) :: x(nvar), ruser(*)
Real (Kind=nag_wp), Intent (Out)  :: rx(nres), rinfo(100), stats(100)
Type (c_ptr), Intent (In)       :: handle, cpuser
External                       :: objfun, mon
```

3 Description

e04fff serves as a solver for compatible problems stored as a handle. The handle points to an internal data structure which defines the problem and serves as a means of communication for routines in the suite.

e04fff is aimed at minimizing a sum of a squares objective function subject to bound constraints:

$$\begin{aligned} \underset{x \in R^n}{\text{minimize}} \quad & \sum_{i=1}^{m_r} r_i(x)^2 & \text{(a)} \\ \text{subject to} \quad & l_x \leq x \leq u_x, & \text{(b)} \end{aligned} \tag{1}$$

Here the $r_i(x)$ are smooth nonlinear functions called residuals and l_x and u_x are n -dimensional vectors defining bounds on the variables. Typically, in a calibration or data fitting context, the residuals will be defined as the difference between a data point and a nonlinear model (see Section 2.2.3 in the E04 Chapter Introduction)

To define a compatible problem handle, you must call **e04raf** followed by **e04rmf** to initialize it and optionally call **e04rhf** to define bounds on the variables. If **e04rhf** is not called, all the variables will be considered free by the solver. It should be noted that **e04fff** always assumes that the Jacobian of the residuals is dense, therefore defining a sparse structure for the residuals in the call to **e04rmf** will have no effect.

It is possible to fix some variables with the definition of the bounds. However, some constraints must be met in order to be able to call the solver:

the number of non-fixed variables n_r has to be at least 2

for all non-fixed variable x_i , the value of $u_x(i) - l_x(i)$ has to be at least twice the starting trust region radius (see the consistency constraint of the optional parameter **DFLS Starting Trust Region**).

The solver is based on a derivative free trust region framework. This type of method is well suited for small to medium-scale problems (around 100 variables) for which the derivatives are unavailable or not easy to compute and/or for which the function evaluations are expensive or noisy. For a detailed description of the algorithm see Section 11. The algorithm behaviour and solver strategy can be modified by various optional parameters (see Section 12) which can be set by **e04zmf** and **e04zpf** anytime between the initialization of the handle by **e04raf** and a call to the solver. The default values for these optional parameters are chosen to work well in the general case but it is recommended to tune them to your particular problem. In particular, if the objective function is noisy, it is highly recommended to set the optional parameter **DFLS Trust Region Update** to SLOW to improve convergence. Once the solver has finished, options may be modified for the next solve. The solver may be called repeatedly with various starting points and/or optional parameters.

4 References

Powell M J D (2009) The BOBYQA algorithm for bound constrained optimization without derivatives *Report DAMTP 2009/NA06* University of Cambridge http://www.damtp.cam.ac.uk/user/na/NA_papers/NA2009_06.pdf

Zhang H, CONN A R and Scheinberg k (2010) A Derivative-Free Algorithm for Least-Squares Minimization *SIAM J. Optim.* **20(6)** 3555–3576

5 Arguments

1: **handle** – Type (c_ptr) *Input*

On entry: the handle to the problem. It needs to be initialized by **e04raf** and the objective must be declared as nonlinear least squares by a call to the routine **e04rmf**. The routine **e04rhf** can optionally be called to define box bounds. It **must not** be changed between calls to the NAG optimization modelling suite.

2: **objfun** – Subroutine, supplied by the user. *External Procedure*

objfun must evaluate the value of the nonlinear residuals $r_i(x)$ at a specified point x .

The specification of **objfun** is:

Fortran Interface

```
Subroutine objfun (nvar, x, nres, rx, inform, iuser, ruser,      &
                  cpuser)
```

```
Integer, Intent (In)           :: nvar, nres
Integer, Intent (Inout)        :: inform, iuser(*)
Real (Kind=nag_wp), Intent (In) :: x(nvar)
Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
Real (Kind=nag_wp), Intent (Out) :: rx(nres)
Type (c_ptr), Intent (In)      :: cpuser
```

1: **nvar** – Integer *Input*

On entry: n , the number of variables in the problem, as set during the initialization of the handle by **e04raf**.

2: **x(nvar)** – Real (Kind=nag_wp) array *Input*

On entry: x , the vector of variable values at which the residuals r_i are to be evaluated.

3: **nres** – Integer *Input*

On entry: m_r , the number of residuals in the problem, as set during the initialization of the handle by **e04rmf**.

4:	rx(nres) – Real (Kind=nag_wp) array <i>On exit:</i> the value of the residuals $r_i(x)$ at x .	<i>Output</i>
5:	inform – Integer <i>On entry:</i> a non-negative value. <i>On exit:</i> may be used to request the solver to stop immediately. Specifically, if inform < 0 then the value of rx will be discarded and the solver will terminate immediately with ifail = 20 otherwise, the solver will proceed normally.	<i>Input/Output</i>
6:	iuser(*) – Integer array	<i>User Workspace</i>
7:	ruser(*) – Real (Kind=nag_wp) array	<i>User Workspace</i>
8:	cpuser – Type (c_ptr)	<i>User Workspace</i>

objfun is called with the arguments **iuser**, **ruser** and **cpuser** as supplied to **e04fff**. You should use the arrays **iuser** and **ruser** and the data handle **cpuser** to supply information to **objfun**.

objfun must either be a module subprogram USED by, or declared as EXTERNAL in, the (sub) program from which **e04fff** is called. Arguments denoted as *Input* must **not** be changed by this procedure.

Note: **objfun** should not return floating-point NaN (Not a Number) or infinity values, since these are not handled by **e04fff**. If your code inadvertently **does** return any NaNs or infinities, **e04fff** is likely to produce unexpected results.

3: **mon** – Subroutine, supplied by the NAG Library or the user. *External Procedure*

mon is provided to enable you to monitor the progress of the optimization and, if necessary, to halt the optimization process using argument **inform**.

If no monitoring is required, **mon** may be the dummy subroutine e04ffu supplied in the NAG Library.

mon is called at the end of every i^{th} step where i is controlled by the optional parameter **DFLS Monitor Frequency** (default value 0, **mon** is never called).

The specification of mon is:		
Fortran Interface		
Subroutine mon (nvar, x, inform, rinfo, stats, iuser, ruser, cpuser) &		
Integer, Intent (In)	:: nvar	
Integer, Intent (Inout)	:: inform, iuser(*)	
Real (Kind=nag_wp), Intent (In)	:: x(nvar), rinfo(100), stats(100)	&
Real (Kind=nag_wp), Intent (Inout)	:: ruser(*)	
Type (c_ptr), Intent (In)	:: cpuser	
1:	nvar – Integer <i>On entry:</i> n , the number of variables in the problem.	<i>Input</i>
2:	x(nvar) – Real (Kind=nag_wp) array <i>On entry:</i> the current best point.	<i>Input</i>
3:	inform – Integer <i>On entry:</i> a non-negative value.	<i>Input/Output</i>

On exit: may be used to request the solver to stop immediately. Specifically, if **inform** < 0 then the value of **rx** will be discarded and the solver will terminate immediately with **ifail** = 20 otherwise, the solver will proceed normally.

4: **rinto**(100) – Real (Kind=nag_wp) array *Input*

On entry: best objective value computed and various indicators (the values are as described in the main argument **rinto**).

5: **stats**(100) – Real (Kind=nag_wp) array *Input*

On entry: solver statistics at the end of the current iteration (the values are as described in the main argument **stats**).

6: **iuser**(*) – Integer array *User Workspace*

7: **ruser**(*) – Real (Kind=nag_wp) array *User Workspace*

8: **cpuser** – Type (c_ptr) *User Workspace*

mon is called with the arguments **iuser**, **ruser** and **cpuser** as supplied to **e04fff**. You should use the arrays **iuser** and **ruser** and the data handle **cpuser** to supply information to **mon**.

mon must either be a module subprogram USED by, or declared as EXTERNAL in, the (sub) program from which **e04fff** is called. Arguments denoted as *Input* must **not** be changed by this procedure.

4: **nvar** – Integer *Input*

On entry: n , the number of variables in the problem. It must be unchanged from the value set during the initialization of the handle by **e04raf**.

Constraint: $nvar \geq 2$.

5: **x**(**nvar**) – Real (Kind=nag_wp) array *Input/Output*

On entry: x_0 , the initial estimates of the variables x .

On exit: the final values of the variables x .

6: **nres** – Integer *Input*

On entry: m_r , the number of residuals in the problem. It must be unchanged from the value set during the definition of the objective structure by **e04rmf**.

7: **rx**(**nres**) – Real (Kind=nag_wp) array *Output*

On exit: the values of the residuals at the final point given in **x**.

8: **rinto**(100) – Real (Kind=nag_wp) array *Output*

On exit: optimal objective value and various indicators at the end of the final iteration as given in the table below:

1 objective function value $f(x)$ (sum of the squared residuals);

2 ρ , the size of trust region at the end of the algorithm;

3 the number of interpolation points used by the solver.

4 – 100 reserved for future use.

9: **stats**(100) – Real (Kind=nag_wp) array *Output*

On exit: solver statistics at the end of the final iteration as given in the table below:

- 1 number of calls to the objective function;
- 2 if **Stats Time** is activated, total time spent in the solver (including time spent evaluating the objective);
- 3 if **Stats Time** is activated, total time spent evaluating the objective function;
- 4 number of steps.
- 5 – 100 reserved for future use.
- 10: **iuser**(*) – Integer array *User Workspace*
- 11: **ruser**(*) – Real (Kind=nag_wp) array *User Workspace*
- 12: **cpuser** – Type (c_ptr) *User Workspace*

iuser, **ruser** and **cpuser** are not used by **e04fff**, but are passed directly to **objfun** and **mon** and may be used to pass information to these routines. If you do not need to reference **cpuser**, it should be initialized to **c_null_ptr**.

- 13: **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, because for this routine the values of the output arguments may be useful even if **ifail** \neq 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.**

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**).

Note: **e04fff** may return useful information for one or more of the following detected errors or warnings.

Errors or warnings detected by the routine:

ifail = 1

The supplied **handle** does not define a valid handle to the data structure for the NAG optimization modelling suite. It has not been initialized by **e04raf** or it has been corrupted.

ifail = 2

The problem is already being solved.

The solver does not support the model defined in the handle.

It supports only nonlinear least squares problems with bound constraints.

ifail = 4

The information supplied does not match with that previously stored.

On entry, **nres** = $\langle value \rangle$ must match that given during the definition of the objective in the **handle**, i.e., $\langle value \rangle$.

The information supplied does not match with that previously stored.

On entry, **nvar** = $\langle value \rangle$ must match that given during initialization of the **handle**, i.e., $\langle value \rangle$.

ifail = 5

Inconsistent optional arguments **DFLS Trust Region Tolerance** ρ_{end} and **DFLS Starting Trust Region** ρ_{beg} .

Constraint: $\rho_{\text{end}} < \rho_{\text{beg}}$.

Use **e04zmf** to set compatible option values.

Inconsistent optional arguments **DFLS Trust Region Tolerance** ρ_{end} and **DFLS Trust Region Slow Tol** ρ_{tol} .

Constraint: $\rho_{\text{end}} < \rho_{\text{tol}}$.

Use **e04zmf** to set compatible option values.

Optional argument **DFLS Starting Trust Region** $\rho_{\text{beg}} = \langle \text{value} \rangle$, $l_x(i) = \langle \text{value} \rangle$, $u_x(i) = \langle \text{value} \rangle$ and $i = \langle \text{value} \rangle$.

Constraint: if $l_x(i) \neq u_x(i)$ in coordinate i , then $u_x(i) - l_x(i) \geq 2 \times \rho_{\text{beg}}$.

Use **e04zmf** to set compatible option values.

ifail = 6

There were $n_r = \langle \text{value} \rangle$ unequal bounds.

Constraint: $n_r \geq 2$.

There were $n_r = \langle \text{value} \rangle$ unequal bounds and the optional argument **DFLS Number Interp Points** $n_{\text{pt}} = \langle \text{value} \rangle$

Constraint: $n_r + 2 \leq n_{\text{pt}} \leq \frac{(n_r+1) \times (n_r+2)}{2}$.

Use **e04zmf** to set compatible option values.

ifail = 18

The predicted reduction in a trust region step was non-positive. Check your specification of **objfun** and whether the function needs rescaling. Try a different initial **x**.

ifail = 19

A rescue procedure has been called in order to correct damage from rounding errors when computing an update to a quadratic approximation of F , but no further progress could be made. Check your specification of **objfun** and whether the function needs rescaling. Try a different initial **x**.

ifail = 20

User requested termination after a call to the objective function. **inform** was set to a negative value within the user-supplied function **objfun**.

User requested termination during a monitoring step. **inform** was set to a negative value within the user-supplied function **mon**.

ifail = 21

Maximum number of function evaluations exceeded.

ifail = 23

The solver terminated after the maximum time allowed was exceeded.

*Maximum number of seconds exceeded. Use option **Time Limit** to reset the limit.*

ifail = 24

No progress, the solver was stopped after $\langle \text{value} \rangle$ consecutive slow steps.

Use the optional argument **DFLS Maximum Slow Steps** to modify the maximum number of slow steps accepted.

*The solver stopped after $5 \times$ **DFLS Maximum Slow Steps** consecutive slow steps and a trust region above the tolerance set by **DFLS Trust Region Slow Tol**.*

ifail = 50

The problem was solved to an acceptable level after $\langle value \rangle$ consecutive slow iterations.

Use the optional argument **DFLS Maximum Slow Steps** to modify the maximum number of slow steps accepted.

*The solver stopped after **DFLS Maximum Slow Steps** consecutive slow steps and a trust region below the tolerance set by **DFLS Trust Region Slow Tol**.*

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 solver can declare convergence on two conditions:

- (i) The trust region radius is below the tolerance ρ_{end} set by the optional parameter **DFLS Trust Region Tolerance**. When this condition is met, the corresponding solution will generally be at a distance lower than $10 \times \rho_{\text{end}}$ of a local minimum.
- (ii) The sum of the square of the residuals is below the tolerance set by the optional parameter **DFLS Small Residuals Tol**. In a data fitting context, this condition means that the error between the observed data and the model is smaller than the requested tolerance.

8 Parallelism and Performance

e04fff makes calls to BLAS and/or LAPACK routines, which may be threaded within the vendor library used by this implementation. Consult the documentation for the vendor library for further information.

Please consult the X06 Chapter Introduction for information on how to control and interrogate the OpenMP environment used within this routine. Please also consult the Users' Note for your implementation for any additional implementation-specific information.

9 Further Comments

9.1 Description of the Printed Output

The solver can print information to give an overview of the problem and of the progress of the computation. The output may be sent to two independent unit numbers which are set by optional parameters **Print File** and **Monitoring File**. Optional parameters **Print Level**, **Print Options**, **Monitoring Level** and **Print Solution** determine the exposed level of detail. This allows, for example, a detailed log file to be generated while the condensed information is displayed on the screen.

By default (**Print File** = 6, **Print Level** = 2), four sections are printed to the standard output: a header, a list of options, an iteration log and a summary.

Header

The header contains statistics about the problem. It should look like:

```
-----
E04FF, Derivative free solver for data fitting
      (nonlinear least squares problems)
-----
```

```
Problem statistics
Number of variables           10
Number of unconstrained variables 10
Number of fixed variables     0
Number of residuals          10
```

Optional parameters list

If **Print Options** is set to YES, a list of the optional parameters and their values is printed. The list shows all options of the solver, each displayed on one line. The line contains the option name, its current value and an indicator for how it was set. The options left at their defaults are noted by (d) and the ones set by the user are noted by (U). Note that the output format is compatible with the file format expected by **e04zpf**. The output looks as follows:

```
Stats Time = Yes * U
Dfls Trust Region Tolerance = 1.00000E-07 * U
Dfls Max Objective Calls = 500 * d
Dfls Starting Trust Region = 1.10000E-01 * U
Dfls Number Interp Points = 0 * d
```

Iteration log

If **Print Level** ≥ 2 , the solver will print a summary line for each step. An iteration is considered successful when it yields a decrease of the objective sufficiently close to the decrease predicted by the quadratic model. The line shows the step number, the value of the objective function, the radius of the trust region and the cumulative number of objective function evaluations. The output looks as follow:

```
-----
step | obj      rho      | nf      |
-----|-----|-----|
  1 | 3.82E+01 1.10E-01 | 13      |
  2 | 3.55E+01 1.10E-01 | 14      |
  3 | 3.05E+01 1.10E-01 | 15      |
  4 | 2.15E+01 1.10E-01 | 18      |
-----
```

Occasionally, the letter 's' is printed at the end of the line indicating that the progress is considered slow by the slow convergence detection heuristic. After a certain number of consecutive slow steps, the solver is stopped. The limit on the number of slow iterations can be controlled by the optional parameter **DFLS Maximum Slow Steps** and the tolerance on the trust region radius before the solver can be stopped is driven by **DFLS Trust Region Slow Tol**.

Summary

Once the solver finishes, a summary is produced:

```
Status: Converged, small trust region size.

Value of the objective           2.23746E-06
Number of objective function evaluations 107
Number of steps                   51
```

Optionally, if **Stats Time** is set to YES, the timings are printed:

```
Timings
Total time spent in the solver      0.056
Time spent in the objective evaluation 0.012
```

Additionally, if **Print Solution** is set to YES, the solution is printed along with the bounds:

```
Computed Solution:
idx  Lower bound      Value      Upper bound
  1   -inf          -1.00000E+00   inf
  2   -inf          -1.00000E+00   inf
  3   -inf          -1.00000E+00   inf
  4   -inf          -1.00000E+00   inf
```

10 Example

In this example, we minimize the Kowalik and Osborne function with bounds on some of the variables. In this problem, the number of variables $n = 4$ and the number of residuals $m_r = 11$. The residuals r_i are computed by

$$r_i(x) = z_i - \frac{y_i(y_i + x_2)}{y_i(y_i + x_3) + x_4} x_1 \quad (2)$$

where

$$\begin{aligned} y &= (4.0000, 2.0000, 1.0000, 0.5000, 0.2500, 0.1670, 0.1250, 0.1000, 0.0833, 0.0714, 0.0625) \\ z &= (0.1957, 0.1947, 0.1735, 0.1600, 0.0844, 0.0627, 0.0456, 0.0342, 0.0323, 0.0235, 0.0246) \end{aligned} \quad (3)$$

The following bounds are defined on the variables

$$\begin{aligned} 0.2 &\leq x_2 \leq 1.0 \\ 0.3 &\leq x_4 \end{aligned} \quad (4)$$

The initial guess is

$$x_0 = (0.25, 0.39, 0.415, 0.39) \quad (5)$$

The expected solution is

$$x^* = (0.1813, 0.5901, 0.2569, 0.3000) \quad (6)$$

10.1 Program Text

```
! e04fff Example Program Text

! Nonlinear least square example: bounded Kowalik and Osborne function
! function

! Mark 26.1 Release. NAG Copyright 2017.

Module e04fffe_mod

! Problem data derived type to be passed to objfun through cpuser

! .. Use Statements ..
Use iso_c_binding, Only: c_f_pointer, c_ptr
Use nag_library, Only: nag_wp
! .. Implicit None Statement ..
Implicit None
! .. Accessibility Statements ..
Private
Public                                :: objfun
! .. Derived Type Definitions ..
Type, Public                            :: pdata
Integer                                  :: ny, nz
Real (Kind=nag_wp), Allocatable :: y(:), z(:)
End Type pdata

Contains

Subroutine objfun(nvar,x,nres,rx,inform,iuser,ruser,cpuser)

! .. Scalar Arguments ..
Type (c_ptr), Intent (In)                :: cpuser
Integer, Intent (Inout)                  :: inform
Integer, Intent (In)                     :: nres, nvar
! .. Array Arguments ..
Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
Real (Kind=nag_wp), Intent (Out) :: rx(nres)
Real (Kind=nag_wp), Intent (In) :: x(nvar)
Integer, Intent (Inout)                :: iuser(*)
! .. Local Scalars ..
Type (pdata), Pointer                    :: pd
Real (Kind=nag_wp)                       :: r1, r2
```

```

Integer                                :: i
! .. Executable Statements ..

! Interrupt solver if the dimensions are incorrect
If (nres/=11 .Or. nvar/=4) Then
  inform = -1
  Go To 100
End If

! extract the problem data structure from the C pointer
Call c_f_pointer(cpuser,pd)
If (pd%ny/=nres .Or. pd%nz/=nres) Then
  inform = -1
  Go To 100
End If

Do i = 1, nres
  r1 = pd%y(i)*(pd%y(i)+x(2))
  r2 = pd%y(i)*(pd%y(i)+x(3)) + x(4)
  rx(i) = pd%z(i) - x(1)*r1/r2
End Do

100 Continue
Return

End Subroutine objfun
End Module e04fffe_mod

Program e04fffe

! .. Use Statements ..
Use e04fffe_mod, Only: objfun, pdata
Use iso_c_binding, Only: c_loc, c_ptr
Use nag_library, Only: e04fff, e04ffu, e04raf, e04rhf, e04rmf, e04rzf, &
  e04zmf, nag_wp

! .. Implicit None Statement ..
Implicit None

! .. Parameters ..
Real (Kind=nag_wp), Parameter      :: infbnd = 1.0E20_nag_wp
Integer, Parameter                 :: nout = 6

! .. Local Scalars ..
Type (c_ptr)                       :: cpuser, handle
Type (pdata), Target               :: pd
Integer                             :: ifail, isparse, nnzrd, nres, nvar

! .. Local Arrays ..
Real (Kind=nag_wp), Allocatable    :: lx(:), rx(:), ux(:), x(:)
Real (Kind=nag_wp)                 :: rinfo(100), ruser(1), stats(100)
Integer                             :: icolrd(1), irowrd(1), iuser(1)

! .. Executable Statements ..

Write (nout,*) 'E04FFF Example Program Results'
Write (nout,*)
Flush (nout)

! fill the problem data structure
pd%ny = 11
pd%nz = 11
Allocate (pd%y(pd%ny),pd%z(pd%nz))

pd%y(1:11) = (/4.0E0_nag_wp,2.0E0_nag_wp,1.0E0_nag_wp,5.0E-1_nag_wp, &
  2.5E-1_nag_wp,1.67E-1_nag_wp,1.25E-1_nag_wp,1.0E-1_nag_wp, &
  8.33E-2_nag_wp,7.14E-2_nag_wp,6.25E-2_nag_wp/)
pd%z(1:11) = (/1.957E-1_nag_wp,1.947E-1_nag_wp,1.735E-1_nag_wp, &
  1.6E-1_nag_wp,8.44E-2_nag_wp,6.27E-2_nag_wp,4.56E-2_nag_wp, &
  3.42E-2_nag_wp,3.23E-2_nag_wp,2.35E-2_nag_wp,2.46E-2_nag_wp/)

nvar = 4
nres = 11

! Initialize handle

```

```

        ifail = 0
        Call e04raf(handle,nvar,ifail)

!       Define residuals structure, isparse=0 means the residual structure is
!       dense => irowrd and icolrd are not accessed
        isparse = 0
        nnzrd = 1
        Call e04rmf(handle,nres,isparse,nnzrd,irowrd,icolrd,ifail)

!       Set options
!       relax the main convergence criteria a bit
        Call e04zmf(handle,'DFLS Trust Region Tolerance = 5.0e-6',ifail)
!       Turn off option printing
        Call e04zmf(handle,'Print Options = NO',ifail)
!       Print the solution
        Call e04zmf(handle,'Print Solution = YES',ifail)

!       Define starting point
        Allocate (x(nvar),rx(nres))
        x(1:4) = (/0.25_nag_wp,0.39_nag_wp,0.415_nag_wp,0.39_nag_wp/)

!       Define bounds for the second and the fourth variable
        Allocate (lx(nvar),ux(nvar))
        lx(1:4) = (/ -infbnd,0.2_nag_wp, -infbnd,0.3_nag_wp/)
        ux = infbnd
        ux(2) = 1.0_nag_wp
        Call e04rhf(handle,nvar,lx,ux,ifail)

!       Call the solver
        ifail = -1
        cpuser = c_loc(pd)
        Call e04fff(handle,objfun,e04ffu,nvar,x,nres,rx,rinfo,stats,iuser,ruser, &
            cpuser,ifail)

!       Free the handle memory
        ifail = 0
        Call e04rzf(handle,ifail)

    End Program e04fffe

```

10.2 Program Data

None.

10.3 Program Results

E04FFF Example Program Results

```

-----
E04FF, Derivative free solver for data fitting
      (nonlinear least-squares problems)
-----

```

```

Problem statistics
  Number of variables           4
  Number of unconstrained variables  2
  Number of fixed variables      0
  Number of residuals           11

```

```

-----
step |   obj   rho   |   nf   |
-----|-----|-----|
  1 | 1.89E-03 1.00E-01 |    7   |
  2 | 5.77E-04 1.00E-01 |    8   |
  3 | 4.23E-04 1.00E-01 |    9   |
  4 | 4.05E-04 1.00E-02 |   10   |
  5 | 4.02E-04 1.00E-02 |   11   |
  6 | 4.02E-04 1.00E-03 |   16   |
  7 | 4.02E-04 1.00E-03 |   18   |
  8 | 4.02E-04 7.07E-05 |   21   |
-----

```

```

     9 | 4.02E-04 7.07E-05 | 22 |
    10 | 4.02E-04 5.00E-06 | 30 |
-----
Status: Converged, small trust region size.

Value of the objective          4.02423E-04
Number of objective function evaluations 30
Number of steps                  10

Computed Solution:
  idx  Lower bound      Value      Upper bound
   1   -inf            1.81300E-01      inf
   2  2.00000E-01      5.90128E-01      1.00000E+00
   3   -inf            2.56927E-01      inf
   4  3.00000E-01      3.00000E-01      inf

```

11 Algorithmic Details

This section contains a short description of the algorithm used in **e04fff** which is based on Powell's method BOBYQA Powell (2009) and the work of Zhang *et al.* (2010). It is based on a model-based derivative free trust region framework adapted to exploit least squares problems structure.

11.1 Derivative free trust region algorithm

In this section, we are interested in generic problems of the form

$$\underset{x \in R^n}{\text{minimize}} \quad f(x) \quad (7)$$

where the derivatives of the objective function f are not easily available. A model-based derivative free optimization (DFO) algorithm maintains a set of points Y_k centred on an iterate x_k to build quadratic interpolation models of the objective

$$f(x_k + s) \approx \phi_k(s) = f(x_k) + g_k^T s + \frac{1}{2} s^T H_k s \quad (8)$$

where g_k and H_k are built with the interpolation conditions

$$y \in Y_k, \phi_k(y - x_k) = f(y) \quad (9)$$

Note that if the number of interpolation points n_{pt} is smaller than $\frac{(n_r+1) \times (n_r+2)}{2}$, the model chosen is the one for which the hessian H_k is the closest to H_{k-1} in the Frobenius norm sense. This model is iteratively optimized over a trust region, updated and moved around the new computed points. More precisely, it can be described as:

DFO Algorithm

1. Initialization

Choose an initial interpolation set Y_0 , trust region radius ρ_{beg} and build the first quadratic model ϕ_0 .

2. Iteration k

- (i) Minimize the model in the trust region to obtain a step s_k .
- (ii) If the step is too small, adjust the geometry of the interpolation set and the trust region size ρ_k and restart the iteration.
- (iii) Evaluate the objective at the new point $x_k + s_k$.
- (iv) Replace a far away point from Y_k by $x_k + s_k$ to create Y_{k+1} .
- (v) If the decrease of the objective is sufficient (successful step), choose $x_{k+1} = x_k + s_k$, else choose $x_{k+1} = x_k$.
- (vi) Choose ρ_{k+1} and adjust the geometry of Y_{k+1} , if necessary.
- (vii) Build ϕ_{k+1} using the new interpolation set.
- (viii) Stop the algorithm if ρ_{k+1} is below the chosen tolerance ρ_{end} .

In the rest of this documentation page, we call an iteration successful when the trial point $x_k + s_k$ is accepted as the next iterate.

11.2 Bounds on the variables

The bounds on the variables are handled during the model optimization step (step 2(i) of DFO Algorithm) with an active set method. If a bound is hit, it is fixed and step 2(i) is restarted. The set of active constraints is kept throughout the optimization, progressively fixing the corresponding variables.

11.3 Adaptation to nonlinear least squares problems

In the specific case where f is a sum of square $f(x) = \sum_{i=1}^{m_r} r_i(x)^2$, a good approximation of the hessian of the objective can be

$$\nabla^2 f(x) \approx J(x)^T J(x) \quad (10)$$

where J is the m_r by n first derivative matrix of f . This approximation is the main idea behind the Gauss–Newton and Levenberg–Marquardt methods. Following the work of Zhang *et al.* (2010), it is possible to adapt it to the DFO framework. In **e04fff**, one quadratic model is built for each residual r_i

$$r_i(x + s) \approx r_i(x) + g^{(i)T} s + \frac{1}{2} s^T H^{(i)} s \quad (11)$$

We call $J = (g^{(1)}, g^{(2)}, \dots)^T$. To build the model of the objective f , we then choose

$$f(x + s) \approx \phi(s) = f(x) + g_f^T s + \frac{1}{2} s^T H_f s \quad (12)$$

where g_f is chosen as

$$g_f = J^T f(x) \quad (13)$$

and H_f as

$$H_f = J^T J + \begin{cases} 0 & \text{if } \|g_f\| \geq \kappa_1 \\ \kappa_3 \|f(x)\| I & \text{if } \|g_f\| < \kappa_1 \text{ and } \frac{1}{2} f(x) < \kappa_2 \|g_f\| \\ \sum_{i=1}^{m_r} r_i(x) H^{(i)} & \text{otherwise} \end{cases} \quad (14)$$

The constants κ_1 , κ_2 and κ_3 are chosen as proposed in Zhang *et al.* (2010). The first expression amounts to making a Gauss–Newton approximation when we are far from a stationary point, the second to a Levenberg–Marquardt approximation when we are close to a stationary point with small residuals while the third takes the full hessian into account.

e04fff integrates this method of building models into the framework presented in the algorithm DFO Algorithm.

12 Optional Parameters

Several optional parameters in **e04fff** define choices in the problem specification or the algorithm logic. In order to reduce the number of formal arguments of **e04fff** these optional parameters have associated *default values* that are appropriate for most problems. Therefore, you need only specify those optional parameters whose values are to be different from their default values.

The remainder of this section can be skipped if you wish to use the default values for all optional parameters.

The optional parameters can be changed by calling **e04zmf** anytime between the initialization of the handle by **e04raf** and the call to the solver. Modification of the arguments during intermediate monitoring stops is not allowed. Once the solver finishes, the optional parameters can be altered again for the next solve.

The option values may be retrieved by **e04znf**.

The following is a list of the optional parameters available. A full description of each optional parameter is provided in Section 12.1.

Defaults

DFLS Maximum Slow Steps

DFLS Max Objective Calls

DFLS Monitor Frequency

DFLS Number Interp Points

DFLS Print Frequency

DFLS Small Residuals Tol

DFLS Starting Trust Region

DFLS Trust Region Slow Tol

DFLS Trust Region Tolerance

DFLS Trust Region Update

Infinite Bound Size

Monitoring File

Monitoring Level

Print File

Print Level

Print Options

Print Solution

Stats Time

Time Limit

12.1 Description of the Optional Parameters

For each option, we give a summary line, a description of the optional parameter and details of constraints.

The summary line contains:

the keywords, where the minimum abbreviation of each keyword is underlined;

a parameter value, where the letters a , i and r denote options that take character, integer and real values respectively.

the default value, where the symbol ϵ is a generic notation for *machine precision* (see **x02ajf**).

All options accept the value DEFAULT to return single options to their default states.

Keywords and character values are case and white space insensitive.

Defaults

This special keyword may be used to reset all optional parameters to their default values. Any argument value given with this keyword will be ignored.

DFLS Maximum Slow Steps i Default = 20

If **DFLS Maximum Slow Steps** > 0 , this argument defines the maximum number of consecutive slow iterations n_{slow} allowed. Set it to 0 to deactivate the slow iteration detection. The algorithm can stop in two situations:

$n_{\text{slow}} > \text{DFLS Maximum Slow Steps}$ and $\rho < \text{DFLS Trust Region Slow Tol}$ with **ifail** = 50

$n_{\text{slow}} > 5 \times \text{DFLS Maximum Slow Steps}$ with **ifail** = 24

Constraint: **DFLS Maximum Slow Steps** ≥ 0 .

DFLS Max Objective Calls i Default = 500

A limit on the number of objective function evaluations the solver is allowed to compute. If the limit is reached, the solver stops with **ifail** = 21.

Constraint: **DFLS Max Objective Calls** ≥ 1 .

DFLS Monitor Frequency i Default = 0

If **DFLS Monitor Frequency** > 0 , the solver calls the user defined monitoring function **mon** at the end of every i th step.

Constraint: **DFLS Monitor Frequency** ≥ 0 .

DFLS Number Interp Points i Default = 0

The number of interpolation points in Y_k (9) used to build the quadratic models. If **DFLS Number Interp Points** = 0, the number of points is chosen to be $n_r + 2$ where n_r is the number of non-fixed variables.

Constraint: **DFLS Number Interp Points** ≥ 0 .

Consistency constraint, the solver stops with **ifail** = 6 if not met:

$$n_r + 2 \leq \text{DFLS Number Interp Points} \leq \frac{(n_r+1) \times (n_r+2)}{2}.$$

DFLS Print Frequency i Default = 1

If **DFLS Print Frequency** > 0 , the solver prints the iteration log to the appropriate units at the end of every i th step.

Constraint: **DFLS Print Frequency** ≥ 0 .

DFLS Small Residuals Tol r Default = $\epsilon^{0.75}$

This option defines the tolerance on the value of the residuals. Namely, the solver declares convergence if

$$f(x) = \sum_{i=1}^{m_r} r_i(x)^2 < \text{DFLS Small Residuals Tol}.$$

Constraint: **DFLS Small Residuals Tol** $> \epsilon^2$.

DFLS Starting Trust Region r Default = 0.1

ρ_{beg} , the initial trust region radius. This argument should be set to about one tenth of the greatest expected overall change to a variable: the initial quadratic model will be constructed by taking steps from the initial x of length ρ_{beg} along each coordinate direction. The default value assumes that the variables have an order of magnitude 1.

Constraint: **DFLS Starting Trust Region** $> \epsilon$.

Consistency constraints, the solver stops with **ifail** = 5 if not met:

DFLS Starting Trust Region \leq **DFLS Trust Region Tolerance**.

$$\text{DFLS Starting Trust Region} \leq \frac{1}{2} \min_i (u_x(i) - l_x(i))$$

DFLS Trust Region Tolerance r Default = $\epsilon^{0.37}$

ρ_{end} , the requested trust region radius. The algorithm declares convergence when the trust region radius reaches this limit. It should indicate the absolute accuracy that is required in the final values of the variables.

Constraint: **DFLS Trust Region Tolerance** $> \epsilon$.

Consistency constraints, the solver stops with **ifail** = 5 if not met:

DFLS Starting Trust Region > DFLS Trust Region Tolerance.

DFLS Trust Region Slow Tol r Default = $\epsilon^{0.25}$

The minimal acceptable trust region radius for the solution to be declared as acceptable. The solver stops if:

$n_{\text{slow}} > \mathbf{DFLS\ Maximum\ Slow\ Steps}$ and $\rho_k < \mathbf{DFLS\ Trust\ Region\ Slow\ Tol}$

Constraint: **DFLS Trust Region Slow Tol** $> \epsilon$.

Consistency constraints, the solver stops with **ifail** = 5 if not met:

DFLS Trust Region Slow Tol > DFLS Trust Region Tolerance

DFLS Trust Region Update a Default = FAST

Controls the speed at which the trust region is decreased after unsuccessful iterations. In smooth non-noisy cases, a fast decrease often leads to faster convergence. However, in noisy cases, a slow decrease is recommended to avoid premature stops.

Constraint: **DFLS Trust Region Update** = FAST or SLOW.

Infinite Bound Size r Default = 10^{20}

This defines the ‘infinite’ bound *bigbnd* in the definition of the problem constraints. Any upper bound greater than or equal to *bigbnd* will be regarded as $+\infty$ (and similarly any lower bound less than or equal to $-bigbnd$ will be regarded as $-\infty$). Note that a modification of this optional parameter does not influence constraints which have already been defined; only the constraints formulated after the change will be affected.

Constraint: **Infinite Bound Size** ≥ 1000 .

Monitoring File i Default = -1

If $i \geq 0$, the unit number for the secondary (monitoring) output. If set to -1 , no secondary output is provided. The information output to this unit is controlled by **Monitoring Level**.

Constraint: **Monitoring File** ≥ -1 .

Monitoring Level i Default = 4

This argument sets the amount of information detail that will be printed by the solver to the secondary output. The meaning of the levels is the same as with **Print Level**.

Constraint: $0 \leq \mathbf{Monitoring\ Level} \leq 5$.

Print File i Default = advisory message unit number

If $i \geq 0$, the unit number for the primary output of the solver. If **Print File** = -1 , the primary output is completely turned off independently of other settings. The default value is the advisory message unit number as defined by **x04abf** at the time of the optional parameters initialization, e.g., at the initialization of the handle. The information output to this unit is controlled by **Print Level**.

Constraint: **Print File** ≥ -1 .

Print Level i Default = 2

This argument defines how detailed information should be printed by the solver to the primary and secondary output.

***i* Output**

0 No output from the solver

1 The Header and Summary.

2,3,4,5 Additionally, the Iteration log.

Constraint: $0 \leq \mathbf{Print Level} \leq 5$.

Print Options *a* Default = YES

If **Print Options** = YES, a listing of optional parameters will be printed to the primary output. It is always printed to the secondary output.

Constraint: **Print Options** = YES or NO.

Print Solution *a* Default = NO

If **Print Solution** = YES, the solution will be printed to the primary and secondary output.

Constraint: **Print Solution** = NO or YES.

Stats Time *a* Default = NO

This argument turns on timings of various parts of the algorithm to give a better overview of where most of the time is spent. This might be helpful for a choice of different solving approaches. It is possible to choose between CPU and wall clock time. Choice YES is equivalent to wall clock.

Constraint: **Stats Time** = YES, NO, CPU or WALL CLOCK.

Time Limit *r* Default = 10^6

A limit on seconds that the solver can use to solve one problem. If during the convergence check this limit is exceeded, the solver will terminate with **ifail** = 23 error message.

Warning: the timings are not computed if **Stats Time** is set to NO. The solver will therefore NOT be stopped if the time limit is exceeded in such a case.

Constraint: **Time Limit** > 0.
