

# NAG Library Routine Document

## E04STF

**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.

**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

E04STF, an interior point method optimization solver, based on the IPOPT software package, is a solver for the NAG optimization modelling suite and is suitable for large scale nonlinear programming (NLP) problems.

### 2 Specification

```

SUBROUTINE E04STF (HANDLE, OBJFUN, OBJGRD, CONFUN, CONGRD, HESS, MON,      &
                  NVAR, X, NNZU, U, RINFO, STATS, IUSER, RUSER, CPUSER,  &
                  IFAIL)
INTEGER          NVAR, NNZU, IUSER(*), IFAIL
REAL (KIND=nag_wp) X(NVAR), U(NNZU), RINFO(32), STATS(32), RUSER(*)
EXTERNAL        OBJFUN, OBJGRD, CONFUN, CONGRD, HESS, MON
TYPE (C_PTR)    HANDLE, CPUSER

```

### 3 Description

E04STF will typically be used for nonlinear programming problems (NLP)

$$\begin{aligned}
 & \underset{x \in \mathbb{R}^n}{\text{minimize}} && f(x) && \text{(a)} \\
 & \text{subject to} && l_g \leq g(x) \leq u_g && \text{(b)} \\
 & && l_B \leq Bx \leq u_B && \text{(c)} \\
 & && l_x \leq x \leq u_x && \text{(d)}
 \end{aligned} \tag{1}$$

where

$n$  is the number of the decision variables,

$m_g$  is the number of the nonlinear constraints and  $g(x)$ ,  $l_g$  and  $u_g$  are  $m_g$ -dimensional vectors,

$m_B$  is the number of the linear constraints and  $B$  is a  $m_B$  by  $n$  matrix,  $l_B$  and  $u_B$  are  $m_B$ -dimensional vectors,

there are  $n$  box constraints and  $l_x$  and  $u_x$  are  $n$ -dimensional vectors.

The objective  $f(x)$  can be specified in a number of ways: E04REF for a dense linear function, E04RFF for a sparse linear or quadratic function and E04RGF for a general nonlinear function. In the last case, OBJFUN and OBJGRD will be used to compute values and gradients of the objective function. Variable box bounds  $l_x, u_x$  can be specified with E04RHF. The special case of linear constraints  $l_B, B, u_B$  is handled by E04RJF while general nonlinear constraints  $l_g, g(x), u_g$  are specified by E04RKF (both can be specified). Again, in the last case, CONFUN and CONGRD will be used to compute values and gradients of the nonlinear constraint functions.

Finally, if the user is willing to calculate second derivatives, the sparsity structure of the second partial derivatives of a nonlinear objective and/or of any nonlinear constraints is specified by E04RLF while the values of these derivatives themselves will be computed by user-supplied HESS. While there is an option (see **Hessian Mode**) that forces internal approximation of second derivatives, no such option

exists for first derivatives which must be computed accurately. If E04RLF has been called and HESS is used to calculate values for second derivatives, both the objective and all the constraints must be included; it is not possible to provide a subset of these. If E04RLF is not called, then internal approximation of second derivatives will take place.

### 3.1 Structure of the Lagrange Multipliers

For a problem consisting of  $n$  variable bounds,  $m_B$  linear constraints and  $m_g$  nonlinear constraints (as specified in NVAR, NCLIN and NCNLN of E04RHF, E04RJF and E04RKF, respectively), the number of Lagrange multipliers, and consequently the correct value for NNZU, will be  $q = 2 * n + 2 * m_B + 2 * m_g$ . The order these will be found in the U array is

$$z_{1L}, z_{1U}, z_{2L}, z_{2U} \dots z_{nL}, z_{nU}, \lambda_{1L}, \lambda_{1U}, \lambda_{2L}, \lambda_{2U} \dots \lambda_{m_{BL}}, \lambda_{m_{BU}}, \lambda_{(m_B+1)L}, \lambda_{(m_B+1)U}, \lambda_{(m_B+2)L}, \lambda_{(m_B+2)U} \dots \lambda_{(m_B+m_g)L}, \lambda_{(m_B+m_g)U}$$

where the  $L$  and  $U$  subscripts refer to lower and upper bounds, respectively, and the variable bound constraint multipliers come first (if present, i.e., if E04RHF was called), followed by the linear constraint multipliers (if present, i.e., if E04RJF was called) and the nonlinear constraint multipliers (if present, i.e., if E04RKF was called).

Significantly nonzero values for any of these, after the solver has terminated, indicates that the corresponding constraint is active. Significance is judged in the first instance by the relative scale of any value compared to the smallest among them.

## 4 References

- Byrd R H, Gilbert J Ch and Nocedal J (2000) A trust region method based on interior point techniques for nonlinear programming *Mathematical Programming* **89** 149–185
- Byrd R H, Liu G and Nocedal J (1997) On the local behavior of an interior point method for nonlinear programming *Numerical Analysis* (eds D F Griffiths and D J Higham) Addison–Wesley
- Conn A R, Gould N I M, Orban D and Toint Ph L (2000) A primal-dual trust-region algorithm for non-convex nonlinear programming *Mathematical Programming* **87 (2)** 215–249
- Conn A R, Gould N I M and Toint Ph L (2000) *Trust Region Methods* SIAM, Philadelphia
- Fiacco A V and McCormick G P (1990) *Nonlinear Programming: Sequential Unconstrained Minimization Techniques* SIAM, Philadelphia
- Gould N I M, Orban D, Sartenaer A and Toint Ph L (2001) Superlinear convergence of primal-dual interior point algorithms for nonlinear programming *SIAM Journal on Optimization* **11 (4)** 974–1002
- Hock W and Schittkowski K (1981) *Test Examples for Nonlinear Programming Codes. Lecture Notes in Economics and Mathematical Systems* **187** Springer–Verlag
- Hogg J D and Scott J A (2011) HSL MA97: a bit-compatible multifrontal code for sparse symmetric systems *RAL Technical Report. RAL-TR-2011-024*
- Wächter A and Biegler L T (2006) On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming *Mathematical Programming* **106(1)** 25–57
- Williams P and Lang B (2013) A framework for the  $MR^3$  algorithm: theory and implementation *SIAM J. Sci. Comput.* **35** 740–766
- Yamashita H (1998) A globally convergent primal-dual interior-point method for constrained optimization *Optimization Methods and Software* **10** 443–469

## 5 Arguments

- 1: HANDLE – TYPE (C\_PTR) Input
- On entry:* the handle to the problem. It needs to be initialized by E04RAF and the problem formulated by some of the routines E04REF, E04RFF, E04RGF, E04RHF, E04RJF, E04RKF and E04RLF. It **must not** be changed between calls to the NAG optimization modelling suite.

- 2: OBJFUN – SUBROUTINE, supplied by the NAG Library or the user. *External Procedure*

OBJFUN must calculate the value of the nonlinear objective function  $f(x)$  at a specified value of the  $n$ -element vector of  $x$  variables. If there is no nonlinear objective (e.g., E04REF or E04RFF was called to define a linear or quadratic objective function), OBJFUN will never be called by E04STF and OBJFUN may be the dummy routine E04STV. (E04STV is included in the NAG Library.)

The specification of OBJFUN is:

```
SUBROUTINE OBJFUN (NVAR, X, FX, INFORM, IUSER, RUSER, CPUSER)
INTEGER          NVAR, INFORM, IUSER(*)
REAL (KIND=nag_wp) X(NVAR), FX, RUSER(*)
TYPE (C_PTR)    CPUSER
```

1: 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.

2: X(NVAR) – REAL (KIND=nag\_wp) array *Input*

*On entry:* the vector  $x$  of variable values at which the objective function is to be evaluated.

3: FX – REAL (KIND=nag\_wp) *Output*

*On exit:* the value of the objective function at  $x$ .

4: INFORM – INTEGER *Input/Output*

*On entry:* a non-negative value.

*On exit:* must be set to a value describing the action to be taken by the solver on return from OBJFUN. Specifically, if the value is negative then the value of FX will be discarded and the solver will either attempt to find a different trial point or terminate immediately with IFAIL = 25 (the same will happen if FX is Infinity or NaN); otherwise, the solver will proceed normally.

5: IUSER(\*) – INTEGER array *User Workspace*

6: RUSER(\*) – REAL (KIND=nag\_wp) array *User Workspace*

7: CPUSER – TYPE (C\_PTR) *User Workspace*

OBJFUN is called with the arguments IUSER, RUSER and CPUSER as supplied to E04STF. You should use the arrays IUSER, RUSER and 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 E04STF is called. Arguments denoted as *Input* must **not** be changed by this procedure.

- 3: OBJGRD – SUBROUTINE, supplied by the NAG Library or the user. *External Procedure*

OBJGRD must calculate the values of the nonlinear objective function gradients  $\frac{\partial f}{\partial x}$  at a specified value of the  $n$ -element vector of  $x$  variables. If there is no nonlinear objective (e.g., E04REF or E04RFF was called to define a linear or quadratic objective function), OBJGRD will never be called by E04STF and OBJGRD may be the dummy subroutine E04STW. (E04STW is included in the NAG Library.)

The specification of OBJGRD is:

```
SUBROUTINE OBJGRD (NVAR, X, NNZFD, FDX, INFORM, IUSER, RUSER, CPUSER) &
```

```
INTEGER NVAR, NNZFD, INFORM, IUSER(*)
REAL (KIND=nag_wp) X(NVAR), FDX(NNZFD), RUSER(*)
TYPE (C_PTR) CPUSER
```

1: 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.

2: X(NVAR) – REAL (KIND=nag\_wp) array *Input*

*On entry:* the vector  $x$  of variable values at which the objective function gradient is to be evaluated.

3: NNZFD – INTEGER *Input*

*On entry:* the number of nonzero elements in the sparse gradient vector of the objective function, as was set in a previous call to E04RGF.

4: FDX(NNZFD) – REAL (KIND=nag\_wp) array *Output*

*On exit:* the values of the nonzero elements in the sparse gradient vector of the objective function, in the order specified by IDXFD in a previous call to E04RGF. FDX( $i$ ) will be the gradient  $\frac{\partial f}{\partial x_{\text{IDXFD}(i)}}$ .

5: INFORM – INTEGER *Input/Output*

*On entry:* a non-negative value.

*On exit:* must be set to a value describing the action to be taken by the solver on return from OBJGRD. Specifically, if the value is negative the solution of the current problem will terminate immediately with IFAIL = 25 (the same will happen if FDX contains Infinity or NaN); otherwise, computations will continue.

6: IUSER(\*) – INTEGER array *User Workspace*

7: RUSER(\*) – REAL (KIND=nag\_wp) array *User Workspace*

8: CPUSER – TYPE (C\_PTR) *User Workspace*

OBJGRD is called with the arguments IUSER, RUSER and CPUSER as supplied to E04STF. You should use the arrays IUSER, RUSER and CPUSER to supply information to OBJGRD.

OBJGRD must either be a module subprogram USED by, or declared as EXTERNAL in, the (sub)program from which E04STF is called. Arguments denoted as *Input* must **not** be changed by this procedure.

4: CONFUN – SUBROUTINE, supplied by the NAG Library or the user. *External Procedure*

CONFUN must calculate the values of the  $m_g$ -element vector  $g_i(x)$  of nonlinear constraint functions at a specified value of the  $n$ -element vector of  $x$  variables. If no nonlinear constraints were registered in this HANDLE, CONFUN will never be called by E04STF and CONFUN may be the dummy subroutine E04STX. (E04STX is included in the NAG Library.)

The specification of CONFUN is:

```
SUBROUTINE CONFUN (NVAR, X, NCNLN, GX, INFORM, IUSER, RUSER, CPUSER) &
```

	INTEGER	NVAR, NCNLN, INFORM, IUSER(*)	
	REAL (KIND=nag_wp)	X(NVAR), GX(NCNLN), RUSER(*)	
	TYPE (C_PTR)	CPUSER	
1:	NVAR – INTEGER		<i>Input</i>
	<i>On entry:</i> $n$ , the number of variables in the problem. It must be unchanged from the value set during the initialization of the handle by E04RAF.		
2:	X(NVAR) – REAL (KIND=nag_wp) array		<i>Input</i>
	<i>On entry:</i> the vector $x$ of variable values at which the constraint functions are to be evaluated.		
3:	NCNLN – INTEGER		<i>Input</i>
	<i>On entry:</i> $m_g$ , the number of nonlinear constraints, as specified in an earlier call to E04RKF.		
4:	GX(NCNLN) – REAL (KIND=nag_wp) array		<i>Output</i>
	<i>On exit:</i> the $m_g$ values of the nonlinear constraint functions at $x$ .		
5:	INFORM – INTEGER		<i>Input/Output</i>
	<i>On entry:</i> a non-negative value.		
	<i>On exit:</i> must be set to a value describing the action to be taken by the solver on return from CONFUN. Specifically, if the value is negative then the value of GX will be discarded and the solver will either attempt to find a different trial point or terminate immediately with IFAIL = 25 (the same will happen if GX contains Infinity or NaN); otherwise, the solver will proceed normally.		
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>
	CONFUN is called with the arguments IUSER, RUSER and CPUSER as supplied to E04STF. You should use the arrays IUSER, RUSER and CPUSER to supply information to CONFUN.		

CONFUN must either be a module subprogram USED by, or declared as EXTERNAL in, the (sub)program from which E04STF is called. Arguments denoted as *Input* must **not** be changed by this procedure.

- 5: CONGRD – SUBROUTINE, supplied by the NAG Library or the user. *External Procedure*
- CONGRD must calculate the nonzero values of the sparse Jacobian of the nonlinear constraint functions  $\frac{\partial g_i}{\partial x}$  at a specified value of the  $n$ -element vector of  $x$  variables. If there are no nonlinear constraints (e.g., E04RKF was never called with the same HANDLE or it was called with NCNLN = 0), CONGRD will never be called by E04STF and CONGRD may be the dummy subroutine E04STY. (E04STY is included in the NAG Library.)

The specification of CONGRD is:

```

SUBROUTINE CONGRD (NVAR, X, NNZGD, GDX, INFORM, IUSER, RUSER,      &
                  CPUSER)

INTEGER              NVAR, NNZGD, INFORM, IUSER(*)
REAL (KIND=nag_wp) X(NVAR), GDX(NNZGD), RUSER(*)
TYPE (C_PTR)        CPUSER

```

1:	NVAR – INTEGER	<i>Input</i>
	<i>On entry:</i> $n$ , the number of variables in the problem. It must be unchanged from the value set during the initialization of the handle by E04RAF.	
2:	X(NVAR) – REAL (KIND=nag_wp) array	<i>Input</i>
	<i>On entry:</i> the vector $x$ of variable values at which the Jacobian of the constraint functions is to be evaluated.	
3:	NNZGD – INTEGER	<i>Input</i>
	<i>On entry:</i> is the number of nonzero elements in the sparse Jacobian of the constraint functions, as was set in a previous call to E04RKF.	
4:	GDX(NNZGD) – REAL (KIND=nag_wp) array	<i>Output</i>
	<i>On exit:</i> the nonzero values of the Jacobian of the nonlinear constraints, in the order specified by IROWGD and ICOLGD in an earlier call to E04RKF. $GDX(i)$ will be the gradient $\frac{\partial g_{IROWGD(i)}}{\partial x_{ICOLGD(i)}}$ .	
5:	INFORM – INTEGER	<i>Input/Output</i>
	<i>On entry:</i> a non-negative value.	
	<i>On exit:</i> must be set to a value describing the action to be taken by the solver on return from CONGRD. Specifically, if the value is negative the solution of the current problem will terminate immediately with IFAIL = 25 (the same will happen if GDX contains Infinity or NaN); otherwise, computations will continue.	
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>
	CONGRD is called with the arguments IUSER, RUSER and CPUSER as supplied to E04STF. You should use the arrays IUSER, RUSER and CPUSER to supply information to CONGRD.	

CONGRD must either be a module subprogram USED by, or declared as EXTERNAL in, the (sub)program from which E04STF is called. Arguments denoted as *Input* must **not** be changed by this procedure.

6: HESS – SUBROUTINE, supplied by the NAG Library or the user. *External Procedure*

HESS must calculate the nonzero values of one of a set of second derivative quantities:

the Hessian of the Lagrangian function  $\sigma \nabla^2 f + \sum_{i=1}^{m_g} \lambda_i \nabla^2 g_i$

the Hessian of the objective function  $\nabla^2 f$

the Hessian of the constraint functions  $\nabla^2 g_i$

The value of argument IDF determines which one of these is to be computed and this, in turn, is determined by earlier calls to E04RLF, when the nonzero sparsity structure of these Hessians was registered. Please note that it is not possible to only supply a subset of the Hessians (see IFAIL = 6). If there were no calls to E04RLF, HESS will never be called by E04STF and HESS may be the dummy subroutine E04STZ (E04STZ is included in the NAG Library). In this case, the Hessian of the Lagrangian will be approximated by a limited-memory quasi-Newton method (L-BFGS).

The specification of HESS is:

```
SUBROUTINE HESS (NVAR, X, NCNLN, IDF, SIGMA, LAMBDA, NNZH, HX,      &
                 INFORM, IUSER, RUSER, CPUSER)
```

```
INTEGER          NVAR, NCNLN, IDF, NNZH, INFORM, IUSER(*)
REAL (KIND=nag_wp) X(NVAR), SIGMA, LAMBDA(NCNLN), HX(NNZH),      &
                 RUSER(*)
TYPE (C_PTR)     CPUSER
```

1: 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.

2: X(NVAR) – REAL (KIND=nag\_wp) array *Input*

*On entry:* the vector  $x$  of variable values at which the Hessian functions are to be evaluated.

3: NCNLN – INTEGER *Input*

*On entry:*  $m_g$ , the number of nonlinear constraints, as specified in an earlier call to E04RKF.

4: IDF – INTEGER *Input*

*On entry:* specifies the quantities to be computed in HX.

IDF = -1

The values of the Hessian of the Lagrangian will be computed in HX. This will be the case if E04RLF has been called with IDF of the same value.

IDF = 0

The values of the Hessian of the objective function will be computed in HX. This will be the case if E04RLF has been called with IDF of the same value.

IDF > 0

The values of the Hessian of the IDFth constraint function will be computed in HX. This will be the case if E04RLF has been called with IDF of the same value.

5: SIGMA – REAL (KIND=nag\_wp) *Input*

*On entry:* if IDF = -1, the value of the  $\sigma$  quantity in the definition of the Hessian of the Lagrangian. Otherwise, SIGMA should not be referenced.

6: LAMBDA(NCNLN) – REAL (KIND=nag\_wp) array *Input*

*On entry:* if IDF = -1, the values of the  $\lambda_i$  quantities in the definition of the Hessian of the Lagrangian. Otherwise, LAMBDA should not be referenced.

7: NNZH – INTEGER *Input*

*On entry:* the number of nonzero elements in the Hessian to be computed.

8: HX(NNZH) – REAL (KIND=nag\_wp) array *Output*

*On exit:* the nonzero values of the requested Hessian evaluated at  $x$ . For each value of IDF, the ordering of nonzeros must follow the sparsity structure registered in the HANDLE by earlier calls to E04RLF through the arguments IROWH and ICOLH.

9: INFORM – INTEGER *Input/Output*

*On entry:* a non-negative value.

*On exit:* must be set to a value describing the action to be taken by the solver on return from HESS. Specifically, if the value is negative the solution of the current problem will terminate immediately with IFAIL = 25 (the same will happen if HX contains Infinity or NaN); otherwise, computations will continue.

- |     |                                     |                       |
|-----|-------------------------------------|-----------------------|
| 10: | IUSER(*) – INTEGER array            | <i>User Workspace</i> |
| 11: | RUSER(*) – REAL (KIND=nag_wp) array | <i>User Workspace</i> |
| 12: | CPUSER – TYPE (C_PTR)               | <i>User Workspace</i> |

HESS is called with the arguments IUSER, RUSER and CPUSER as supplied to E04STF. You should use the arrays IUSER, RUSER and CPUSER to supply information to HESS.

HESS must either be a module subprogram USED by, or declared as EXTERNAL in, the (sub) program from which E04STF is called. Arguments denoted as *Input* must **not** be changed by this procedure.

- 7: 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. A facility is provided to halt the optimization process if necessary, using parameter INFORM.

MON may be the dummy subroutine E04STU (E04STU is included in the NAG Library).

The specification of MON is:

```
SUBROUTINE MON (NVAR, X, NNZU, U, INFORM, RINFO, STATS, IUSER,      &
                RUSER, CPUSER)
```

```
INTEGER          NVAR, NNZU, INFORM, IUSER(*)
REAL (KIND=nag_wp) X(NVAR), U(NNZU), RINFO(32), STATS(32),      &
                RUSER(*)
TYPE (C_PTR)     CPUSER
```

- |    |                |              |
|----|----------------|--------------|
| 1: | NVAR – INTEGER | <i>Input</i> |
|----|----------------|--------------|

*On entry:*  $n$ , the number of variables in the problem.

- |    |                                    |              |
|----|------------------------------------|--------------|
| 2: | X(NVAR) – REAL (KIND=nag_wp) array | <i>Input</i> |
|----|------------------------------------|--------------|

*On entry:*  $x^i$ , the values of the decision variables  $x$  at the current iteration.

- |    |                |              |
|----|----------------|--------------|
| 3: | NNZU – INTEGER | <i>Input</i> |
|----|----------------|--------------|

*On entry:* the dimension of array U.

- |    |                                    |              |
|----|------------------------------------|--------------|
| 4: | U(NNZU) – REAL (KIND=nag_wp) array | <i>Input</i> |
|----|------------------------------------|--------------|

*On entry:* if  $NNZU > 0$ , U holds the values at the current iteration of Lagrange multipliers (dual variables) for the constraints. See Section 3.1 for layout information.

- |    |                  |                     |
|----|------------------|---------------------|
| 5: | INFORM – INTEGER | <i>Input/Output</i> |
|----|------------------|---------------------|

*On entry:* a non-negative value.

*On exit:* must be set to a value describing the action to be taken by the solver on return from MON. Specifically, if the value is negative the solution of the current problem will terminate immediately with IFAIL = 20; otherwise, computations will continue.

- |    |                                      |              |
|----|--------------------------------------|--------------|
| 6: | RINFO(32) – REAL (KIND=nag_wp) array | <i>Input</i> |
|----|--------------------------------------|--------------|

*On entry:* error measures and various indicators at the end of the current iteration as described in Section 9.1.



7:	STATS(32) – REAL (KIND=nag_wp) array	<i>Input</i>
	<i>On entry:</i> solver statistics at the end of the current iteration as described in Section 9.1.	
8:	IUSER(*) – INTEGER array	<i>User Workspace</i>
9:	RUSER(*) – REAL (KIND=nag_wp) array	<i>User Workspace</i>
10:	CPUSER – TYPE (C_PTR)	<i>User Workspace</i>
	MON is called with the arguments IUSER, RUSER and CPUSER as supplied to E04STF. You should use the arrays IUSER, RUSER and 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 E04STF is called. Arguments denoted as *Input* must **not** be changed by this procedure.

- 8: 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.
- 9: 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  $x$ .
- 10: NNZU – INTEGER *Input*  
*On entry:* the number of Lagrange multipliers that are to be returned in array U.  
 If NNZU = 0, U will not be referenced; otherwise it needs to match the dimension  $q$  as explained in Section 3.1.  
*Constraints:*  

$$\text{NNZU} \geq 0;$$
 if NNZU > 0, NNZU =  $q$ .
- 11: U(NNZU) – REAL (KIND=nag\_wp) array *Output*  
**Note:** if NNZU > 0, U holds Lagrange multipliers (dual variables) for the constraints. See Section 3.1 for layout information. If NNZU = 0, U will not be referenced.  
*On exit:* the final value of Lagrange multipliers  $z, \lambda$ .
- 12: RINFO(32) – REAL (KIND=nag\_wp) array *Output*  
*On exit:* error measures and various indicators at the end of the final iteration as given in the table below:
- |   |   |
|---|---|
| 1 | objective function value $f(x)$                 |
| 2 | constraint violation (primal infeasibility) (8) |
| 3 | dual infeasibility (7)                          |
| 4 | complementarity                                 |
| 5 | Karush–Kuhn–Tucker error                        |
- 13: STATS(32) – 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 the iterations           |
| 3 | number of backtracking trial steps |

4	number of Hessian evaluations	
5	number of objective gradient evaluations	
8	total wall clock time elapsed	
19	number of objective function evaluations	
20	number of constraint function evaluations	
21	number of constraint Jacobian evaluations	
14:	IUSER(*) – INTEGER array	<i>User Workspace</i>
15:	RUSER(*) – REAL (KIND=nag_wp) array	<i>User Workspace</i>
16:	CPUSER – TYPE (C_PTR)	<i>User Workspace</i>

IUSER, RUSER and CPUSER are not used by E04STF, but are passed directly to OBJFUN, OBJGRD, CONFUN, CONGRD, HESS and MON and should be used to pass information to these routines.

17:	IFAIL – INTEGER	<i>Input/Output</i>
-----	-----------------	---------------------

*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:** E04STF 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.

This solver does not support matrix inequality constraints.

IFAIL = 3

A different solver from the suite has already been used. Initialize a new HANDLE using E04RAF.

IFAIL = 4

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

On entry, NNZU =  $\langle value \rangle$ .

Constraint: NNZU =  $\langle value \rangle$  or 0.

On entry, NNZU =  $\langle value \rangle$ .

Constraint: no constraints present, so NNZU must be 0.

IFAIL = 6

*Either all of the constraint and objective Hessian structures must be defined or none (in which case, the Hessians will be approximated by a limited-memory quasi-Newton L-BFGS method).*

On entry, a nonlinear objective function has been defined but no objective Hessian sparsity structure has been defined through E04RLF.

On entry, a nonlinear constraint function has been defined but no constraint Hessian sparsity structure has been defined through E04RLF, for constraint number  $\langle value \rangle$ .

IFAIL = 7

The dummy CONFUN routine was called but the problem requires these values. Please provide a proper CONFUN routine.

The dummy CONGRD routine was called but the problem requires these derivatives. Please provide a proper CONGRD routine.

The dummy HESS routine was called but the problem requires these derivatives. Either change the option **Hessian Mode** or provide a proper HESS routine.

The dummy OBJFUN routine was called but the problem requires these values. Please provide a proper OBJFUN routine.

The dummy OBJGRD routine was called but the problem requires these derivatives. Please provide a proper OBJGRD routine.

IFAIL = 20

User requested termination during a monitoring step. INFORM was set to a negative value in MON.

IFAIL = 22

Maximum number of iterations exceeded.

IFAIL = 23

The solver terminated after an error in the step computation. *This message is printed if the solver is unable to compute a search direction, despite several attempts to modify the iteration matrix. Usually, the value of the regularization parameter then becomes too large. One situation where this can happen is when values in the Hessian are invalid (NaN or Infinity). You can check whether this is true by using the **Verify Derivatives** option.*

The solver terminated after failure in the restoration phase. *This indicates that the restoration phase failed to find a feasible point that was acceptable to the filter line search for the original problem. This could happen if the problem is highly degenerate, does not satisfy the constraint qualification, or if your NLP code provides incorrect derivative information.*

The solver terminated after the maximum time allowed was exceeded. *Maximum number of seconds exceeded. Use option **Time Limit** to reset the limit.*

The solver terminated due to an invalid option. Please contact NAG with details of the call to E04STF.

The solver terminated due to an invalid problem definition. Please contact NAG with details of the call to E04STF.

The solver terminated with not enough degrees of freedom. *This indicates that your problem, as specified, has too few degrees of freedom. This can happen if you have too many equality constraints, or if you fix too many variables.*

IFAIL = 24

The solver terminated after the search direction became too small. *This indicates that the solver is calculating very small step sizes and is making very little progress. This could happen if the problem has been solved to the best numerical accuracy possible given the current NLP scaling.*

IFAIL = 25

Invalid number detected in user function. *Either INFORM was set to a negative value within the user-supplied functions OBJFUN, OBJGRD, CONFUN, CONGRD or HESS, or an Infinity or NaN was detected in values returned from them.*

IFAIL = 50

The solver reports NLP solved to acceptable level. *This indicates that the algorithm did not converge to the desired tolerances, but that it was able to obtain a point satisfying the acceptable tolerance level. This may happen if the desired tolerances are too small for the current problem.*

IFAIL = 51

The solver detected an infeasible problem. *The restoration phase converged to a point that is a minimizer for the constraint violation (in the  $\ell_1$ -norm), but is not feasible for the original problem. This indicates that the problem may be infeasible (or at least that the algorithm is stuck at a locally infeasible point). The returned point (the minimizer of the constraint violation) might help you to find which constraint is causing the problem. If you believe that the NLP is feasible, it might help to start the optimization from a different point.*

IFAIL = 54

The solver terminated due to diverging iterates. *The max-norm of the iterates has become larger than a preset value. This can happen if the problem is unbounded below and the iterates are diverging.*

IFAIL = -199

This routine is not available in this implementation.

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 accuracy of the solution is driven by optional parameter **Stop Tolerance 1**.

If IFAIL = 0 on the final exit, the returned point satisfies Karush–Kuhn–Tucker (KKT) conditions to the requested accuracy (under the default settings close to  $\sqrt{\epsilon}$  where  $\epsilon$  is the *machine precision*) and

thus it is a good estimate of a local solution. If `IFAIL = 50`, some of the convergence conditions were not fully satisfied but the point still seems to be a reasonable estimate and should be usable. Please refer to Section 11.1 and the description of the particular options.

## 8 Parallelism and Performance

E04STF is not threaded in any implementation.

## 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 streams (files) which are set by optional parameters **Print File** and **Monitoring File**. Optional parameters **Print Level** and **Monitoring Level** determine the exposed level of detail. This allows, for example, the generation of a detailed log in a file while the condensed information is displayed on the screen. This section also describes what kind of information is made available to the monitoring routine MON via RINFO and STATS.

There are four sections printed to the primary output with the default settings (level 2): a derivative check, a header, an iteration log and a summary. At higher levels more information will be printed, including any internal IPOPT options that have been changed from their default values.

#### Derivative Check

If **Verify Derivatives** is set, then information will appear about any errors detected in the user-supplied derivative routines OBJGRD, CONGRD or HESS. It may look like this:

```
Starting derivative checker for first derivatives.
* grad_f[          1] = -2.000000E+00   ~ 2.455000E+01 [ 1.081E+00]
* jac_g [    1,    4] = 4.700969E+01 v ~ 5.200968E+01 [ 9.614E-02]
Starting derivative checker for second derivatives.
*          obj_hess[    1,    1] = 1.881000E+03 v ~ 1.882000E+03 [
5.314E-04]
*    1-th constr_hess[    1,    3] = 2.988964E+00 v ~ -1.103543E-02 [
3.000E+00]

Derivative checker detected 3 error(s).
```

The first line indicates that the value for the partial derivative of the objective with respect to the first variable as returned by OBJGRD (the first one printed) differs sufficiently from a finite difference estimation derived from OBJFUN (the second one printed). The number in square brackets is the relative difference between these two numbers.

The second line reports on a discrepancy for the partial derivative of the first constraint with respect to the fourth variable. If the indicator `v` is absent, the discrepancy refers to a component that had not been included in the sparsity structure, in which case the nonzero structure of the derivatives should be corrected. Mistakes in the first derivatives should be corrected before attempting to correct mistakes in the second derivatives.

The third line reports on a discrepancy in a second derivative of the objective function, differentiated with respect to the first variable, twice.

The fourth line reports on a discrepancy in a second derivative of the first constraint, differentiated with respect to the first and third variables.

#### Header

If **Print Level**  $\geq 1$ , the header will contain statistics about the size of the problem how the solver sees it, i.e., it reflects any changes imposed by preprocessing and problem transformations. The header may look like:

```
Number of nonzeros in equality constraint Jacobian...:      4Number of
nonzeros in inequality constraint Jacobian.:      8Number of nonzeros in
```

```

Lagrangian Hessian.....:      10Total number of vari-
ables.....:      4          variables with
only lower bounds:      4          variables with lower and upper
bounds:      0          variables with only upper bounds:
0Total number of equality constraints.....:      1Total number
of inequality constraints.....:      2          inequality
constraints with only lower bounds:      2          inequality constraints with
lower and upper bounds:      0          inequality constraints with only upper
bounds:      0

```

It summarises what is known about the variables and the constraints. Simple bounds are set by E04RHF and standard equalities and inequalities by E04RJF.

### Iteration log

If **Print Level** = 2, the status of each iteration is condensed to one line. The line shows:

<code>iter</code>	The current iteration count. This includes regular iterations and iterations during the restoration phase. If the algorithm is in the restoration phase, the letter <b>r</b> will be appended to the iteration number. The iteration number 0 represents the starting point. This quantity is also available as <code>STATS(1)</code> of <code>MON</code> .
<code>objective</code>	The unscaled objective value at the current point (given the current NLP scaling). During the restoration phase, this value remains the unscaled objective value for the original problem. This quantity is also available as <code>RINFO(1)</code> of <code>MON</code> .
<code>inf_pr</code>	The unscaled constraint violation at the current point (given the current NLP scaling). This quantity is the infinity-norm (max) of the (unscaled) constraints $g_i$ . During the restoration phase, this value remains the constraint violation of the original problem at the current point. This quantity is also available as <code>RINFO(2)</code> of <code>MON</code> .
<code>inf_du</code>	The scaled dual infeasibility at the current point (given the current NLP scaling). This quantity measure the infinity-norm (max) of the internal dual infeasibility, $\lambda_i$ of Eq. (4a) in the implementation paper Wächter and Biegler (2006), including inequality constraints reformulated using slack variables and NLP scaling. During the restoration phase, this is the value of the dual infeasibility for the restoration phase problem. This quantity is also available as <code>RINFO(3)</code> of <code>MON</code> .
<code>lg(mu)</code>	$\log_{10}$ of the value of the barrier parameter $\mu$ . $\mu$ itself is also available as <code>RINFO(4)</code> of <code>MON</code> .
<code>  d  </code>	The infinity norm (max) of the primal step (for the original variables $x$ and the internal slack variables $s$ ). During the restoration phase, this value includes the values of additional variables, $\bar{p}$ and $\bar{n}$ (see Eq. (30) in Wächter and Biegler (2006)). This quantity is also available as <code>RINFO(5)</code> of <code>MON</code> .
<code>lg(rg)</code>	$\log_{10}$ of the value of the regularization term for the Hessian of the Lagrangian in the augmented system ( $\delta_w$ of Eq. (26) and Section 3.1 in Wächter and Biegler (2006)). A dash (–) indicates that no regularization was done. The regularization term itself is also available as <code>RINFO(6)</code> of <code>MON</code> .
<code>alpha_du</code>	The stepsize for the dual variables ( $\alpha_k^z$ of Eq. (14c) in Wächter and Biegler (2006)). This quantity is also available as <code>RINFO(7)</code> of <code>MON</code> .
<code>alpha_pr</code>	The stepsize for the primal variables ( $\alpha_k$ of Eq. (14a) in Wächter and Biegler (2006)). This quantity is also available as <code>RINFO(8)</code> of <code>MON</code> . The number is usually followed by a character for additional diagnostic information regarding the step acceptance criterion.

ff-type iteration in the filter method without second order correction  
Ff-type iteration in the filter method with second order correction  
hh-type iteration in the filter method without second order correction  
Hh-type iteration in the filter method with second order correction

kpenalty value unchanged in merit function method without second order correction

Kpenalty value unchanged in merit function method with second order correction

npenalty value updated in merit function method without second order correction

Npenalty value updated in merit function method with second order correction

RRestoration phase just started

win watchdog procedure

sstep accepted in soft restoration phase

t/Ttiny step accepted without line search

rsume previous iterate restored

ls                    The number of backtracking line search steps (does not include second order correction steps). This quantity is also available as STATS(2) of MON.

Note that the step acceptance mechanisms in IPOPT consider the barrier objective function (5) which is usually different from the value reported in the objective column. Similarly, for the purposes of the step acceptance, the constraint violation is measured for the internal problem formulation, which includes slack variables for inequality constraints and potentially NLP scaling of the constraint functions. This value, too, is usually different from the value reported in `inf_pr`. As a consequence, a new iterate might have worse values both for the objective function and the constraint violation as reported in the iteration output, seemingly contradicting globalization procedure.

Note that all these values are also available in `RINFO(1), ..., RINFO(8)` and `STATS(1), ..., STATS(2)` of the monitoring routine MON.

The output might look as follows:

iter	objective	inf_pr	inf_du	lg(mu)	d	lg(rg)	alpha_du	alpha_pr
ls								
0	2.6603500E+05	1.55E+02	3.21E+01	-1.0	0.00E+00	-	0.00E+00	0.00E+00
1	1.5053889E+05	7.95E+01	1.43E+01	-1.0	1.16E+00	-	4.55E-01	1.00E+00f
2	8.9745785E+04	3.91E+01	6.45E+00	-1.0	3.07E+01	-	5.78E-03	1.00E+00f
3	3.9878595E+04	1.63E+01	3.47E+00	-1.0	5.19E+00	0.0	2.43E-01	1.00E+00f
4	2.7780042E+04	1.08E+01	1.64E+00	-1.0	3.66E+01	-	7.24E-01	8.39E-01f
5	2.6194274E+04	1.01E+01	1.49E+00	-1.0	1.07E+01	-	1.00E+00	1.05E-01f
6	1.5422960E+04	4.75E+00	6.82E-01	-1.0	1.74E+01	-	1.00E+00	1.00E+00f
7	1.1975453E+04	3.14E+00	7.26E-01	-1.0	2.83E+01	-	1.00E+00	5.06E-01f
8	8.3508421E+03	1.34E+00	2.04E-01	-1.0	3.96E+01	-	9.27E-01	1.00E+00f
9	7.0657495E+03	4.85E-01	9.22E-02	-1.0	5.32E+01	-	1.00E+00	1.00E+00f
10	6.8359393E+03	1.17E-01	1.28E-01	-1.7	4.69E+01	-	8.21E-01	1.00E+00h
11	6.6508917E+03	1.52E-02	1.52E-02	-2.5	1.87E+01	-	1.00E+00	1.00E+00h
12	6.4123213E+03	8.77E-03	1.49E-01	-3.8	1.85E+01	-	7.49E-01	1.00E+00f
13	6.3157361E+03	4.33E-03	1.90E-03	-3.8	2.07E+01	-	1.00E+00	1.00E+00f
14	6.2989280E+03	1.12E-03	4.06E-04	-3.8	1.54E+01	-	1.00E+00	1.00E+00h

```

1      15  6.2996264E+03  9.90E-05  2.05E-04  -5.7  5.35E+00  -  9.63E-01  1.00E+00h
1      16  6.2998436E+03  0.00E+00  1.86E-07  -5.7  4.55E-01  -  1.00E+00  1.00E+00h
1      17  6.2998424E+03  0.00E+00  6.18E-12  -8.2  2.62E-03  -  1.00E+00  1.00E+00h
1

```

If **Print Level** > 2, each iteration produces significantly more detailed output comprising detailed error measures and output from internal operations. The output is reasonably self-explanatory so it is not featured here in detail.

### Summary

Once the solver finishes, a detailed summary is produced if **Print Level** ≥ 1. An example is shown below:

```

Number of Iterations.....: 6

Objective.....:          (scaled)          (unscaled)
+00              7.8692659500479623E-01    6.2324586324379867E
Dual infeasibility.....:  7.9744615766675617E-10  6.3157735687207093E-09
Constraint violation.....:  8.3555384833289281E-12  8.3555384833289281E-12
Complementarity.....:    0.0000000000000000E+00    0.0000000000000000E
+00
Overall NLP error.....:  7.9744615766675617E-10  6.3157735687207093E-09

Number of objective function evaluations      = 7
Number of objective gradient evaluations     = 7
Number of equality constraint evaluations     = 7
Number of inequality constraint evaluations   = 0
Number of equality constraint Jacobian evaluations = 7
Number of inequality constraint Jacobian evaluations = 0
Number of Lagrangian Hessian evaluations    = 6
Total CPU secs in IPOPT (w/o function evaluations) = 0.724
Total CPU secs in NLP function evaluations   = 0.343

EXIT: Optimal Solution Found.

```

It starts with the total number of iterations the algorithm went through. Then, five quantities are printed, all evaluated at the termination point: the value of the objective function, the dual infeasibility, the constraint violation, the complementarity and the NLP error.

This is followed by some statistics on the number of calls to user-supplied functions and CPU time taken in user-supplied functions and the main algorithm. Lastly, status at exit is indicated by a short message. Detailed timings of the algorithm are displayed only if **Stats Time** is set.

## 9.2 Additional Licensor

Parts of the code for E04STF are distributed according to terms imposed by another licensor. Please refer to the list of Library licensors available on the NAG Website for further details.

## 10 Example

This example is based on Problem 73 in Hock and Schittkowski (1981) and involves the minimization of the linear function

$$f(x) = 24.55x_1 + 26.75x_2 + 39.00x_3 + 40.50x_4$$

subject to the bounds



$$\begin{aligned} 0 &\leq x_1, \\ 0 &\leq x_2, \\ 0 &\leq x_3, \\ 0 &\leq x_4, \end{aligned}$$

to the nonlinear constraint

$$12x_1 + 11.9x_2 + 41.8x_3 + 52.1x_4 - 21 - 1.645\sqrt{0.28x_1^2 + 0.19x_2^2 + 20.5x_3^2 + 0.62x_4^2} \geq 0$$

and the linear constraints

$$\begin{aligned} 2.3x_1 + 5.6x_2 + 11.1x_3 + 1.3x_4 &\geq 5, \\ x_1 + x_2 + x_3 + x_4 - 1 &= 0. \end{aligned}$$

The initial point, which is infeasible, is

$$x_0 = (1, 1, 1, 1)^T$$

and  $f(x_0) = 130.8$ . The optimal solution (to five significant figures) is

$$x^* = (0.63552, 0.0, 0.31270, 0.051777)^T,$$

## 10.1 Program Text

```
! E04STF Example Program Text
! Mark 26 Release. NAG Copyright 2016.

! NLP example : Linear objective + Linear constraint + Non-Linear constraint

Module e04stfe_mod

! .. Use Statements ..
Use nag_library, Only: nag_wp
! .. Implicit None Statement ..
Implicit None
! .. Accessibility Statements ..
Private
Public          :: confun, congrd, hess, mon, objfun,   &
                  objgrd

Contains
Subroutine objfun(nvar,x,fx,inform,iuser,ruser,cpuser)

! .. Use Statements ..
Use, Intrinsic          :: iso_c_binding, Only: c_ptr
Use nag_library, Only: f06eaf
! .. Scalar Arguments ..
Type (c_ptr), Intent (Inout)  :: cpuser
Real (Kind=nag_wp), Intent (Out) :: fx
Integer, Intent (Inout)      :: inform
Integer, Intent (In)         :: nvar
! .. Array Arguments ..
Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
Real (Kind=nag_wp), Intent (In)  :: x(nvar)
Integer, Intent (Inout)         :: iuser(*)
! .. Executable Statements ..
fx = f06eaf(4,x,1,ruser,1)
Return
End Subroutine objfun
Subroutine objgrd(nvar,x,nnzfd,fdx,inform,iuser,ruser,cpuser)

! .. Use Statements ..
Use, Intrinsic          :: iso_c_binding, Only: c_ptr
! .. Scalar Arguments ..
Type (c_ptr), Intent (Inout)  :: cpuser
Integer, Intent (Inout)      :: inform
Integer, Intent (In)         :: nnzfd, nvar
! .. Array Arguments ..
```

```

Real (Kind=nag_wp), Intent (Out) :: fdx(nnzfd)
Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
Real (Kind=nag_wp), Intent (In) :: x(nvar)
Integer, Intent (Inout) :: iuser(*)
!
.. Executable Statements ..
Continue
fdx(1:nnzfd) = ruser(1:nnzfd)
Return
End Subroutine objgrd
Subroutine confun(nvar,x,ncnln,gx,inform,iuser,ruser,cpuser)

!
.. Use Statements ..
Use, Intrinsic :: iso_c_binding, Only: c_ptr
!
.. Scalar Arguments ..
Type (c_ptr), Intent (Inout) :: cpuser
Integer, Intent (Inout) :: inform
Integer, Intent (In) :: ncnln, nvar
!
.. Array Arguments ..
Real (Kind=nag_wp), Intent (Out) :: gx(ncnln)
Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
Real (Kind=nag_wp), Intent (In) :: x(nvar)
Integer, Intent (Inout) :: iuser(*)
!
.. Intrinsic Procedures ..
Intrinsic :: sqrt
!
.. Executable Statements ..
Continue
gx(1) = 12.0_nag_wp*x(1) + 11.9_nag_wp*x(2) + 41.8_nag_wp*x(3) +      &
52.1_nag_wp*x(4) - 1.645_nag_wp*sqrt(.28_nag_wp*x(1)**2+.19_nag_wp*x &
(2)**2+20.5_nag_wp*x(3)**2+.62_nag_wp*x(4)**2)
Return
End Subroutine confun
Subroutine congrd(nvar,x,nnzgd,gdx,inform,iuser,ruser,cpuser)

!
.. Use Statements ..
Use, Intrinsic :: iso_c_binding, Only: c_ptr
!
.. Scalar Arguments ..
Type (c_ptr), Intent (Inout) :: cpuser
Integer, Intent (Inout) :: inform
Integer, Intent (In) :: nnzgd, nvar
!
.. Array Arguments ..
Real (Kind=nag_wp), Intent (Out) :: gdx(nnzgd)
Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
Real (Kind=nag_wp), Intent (In) :: x(nvar)
Integer, Intent (Inout) :: iuser(*)
!
.. Local Scalars ..
Real (Kind=nag_wp) :: tmp
!
.. Intrinsic Procedures ..
Intrinsic :: sqrt
!
.. Executable Statements ..
Continue
tmp = sqrt(0.62_nag_wp*x(4)**2+20.5_nag_wp*x(3)**2+      &
0.19_nag_wp*x(2)**2+0.28_nag_wp*x(1)**2)
gdx(1) = (12.0_nag_wp*tmp-0.4606_nag_wp*x(1))/tmp
gdx(2) = (11.9_nag_wp*tmp-0.31255_nag_wp*x(2))/tmp
gdx(3) = (41.8_nag_wp*tmp-33.7225_nag_wp*x(3))/tmp
gdx(4) = (52.1_nag_wp*tmp-1.0199_nag_wp*x(4))/tmp
Return
End Subroutine congrd
Subroutine hess(nvar,x,ncnln,idf,sigma,lambda,nnzh,hx,inform,iuser,      &
ruser,cpuser)

!
.. Use Statements ..
Use, Intrinsic :: iso_c_binding, Only: c_ptr
!
.. Scalar Arguments ..
Type (c_ptr), Intent (Inout) :: cpuser
Real (Kind=nag_wp), Intent (In) :: sigma
Integer, Intent (In) :: idf, ncnln, nnzh, nvar
Integer, Intent (Inout) :: inform
!
.. Array Arguments ..
Real (Kind=nag_wp), Intent (Out) :: hx(nnzh)
Real (Kind=nag_wp), Intent (In) :: lambda(ncnln), x(nvar)

```

```

      Real (Kind=nag_wp), Intent (Inout) :: ruser(*)
      Integer, Intent (Inout)          :: iuser(*)
! .. Local Scalars ..
      Real (Kind=nag_wp)               :: tmp
! .. Intrinsic Procedures ..
      Intrinsic                        :: sqrt
! .. Executable Statements ..
      inform = -1
      hx = 0.0_nag_wp
      Select Case (idf)
      Case (0)
         inform = 0
      Case (1,-1)
         tmp = sqrt(0.62_nag_wp*x(4)**2+20.5_nag_wp*x(3)**2+
                   0.19_nag_wp*x(2)**2+0.28_nag_wp*x(1)**2) &
         tmp = tmp*(x(4)**2+33.064516129032258064_nag_wp*x(3)**2+
                   0.30645161290322580645_nag_wp*x(2)**2+
                   0.45161290322580645161_nag_wp*x(1)**2) &
!         1,1..4
         hx(1) = (-0.4606_nag_wp*x(4)**2-15.229516129032258064_nag_wp*x(3)**2 &
                  -0.14115161290322580645_nag_wp*x(2)**2)/tmp
         hx(2) = (0.14115161290322580645_nag_wp*x(1)*x(2))/tmp
         hx(3) = (15.229516129032258064_nag_wp*x(1)*x(3))/tmp
         hx(4) = (0.4606_nag_wp*x(1)*x(4))/tmp
!         2,2..4
         hx(5) = (-0.31255_nag_wp*x(4)**2-10.334314516129032258_nag_wp*x(3)** &
                  2-0.14115161290322580645_nag_wp*x(1)**2)/tmp
         hx(6) = (10.334314516129032258_nag_wp*x(2)*x(3))/tmp
         hx(7) = (0.31255_nag_wp*x(2)*x(4))/tmp
!         3,3..4
         hx(8) = (-33.7225_nag_wp*x(4)**2-10.334314516129032258_nag_wp*x(2)** &
                  2-15.229516129032258065_nag_wp*x(1)**2)/tmp
         hx(9) = (33.7225_nag_wp*x(3)*x(4))/tmp
!         4,4
         hx(10) = (-33.7225_nag_wp*x(3)**2-0.31255_nag_wp*x(2)**2-
                   0.4606_nag_wp*x(1)**2)/tmp &
         If (idf== -1) Then
            hx = lambda(1)*hx
         End If
         inform = 0
      Case Default
      End Select
      Return
End Subroutine hess
Subroutine mon(nvar,x,nnzu,u,inform,rinfo,stats,iuser,ruser,cpuser)

! .. Use Statements ..
Use, Intrinsic                        :: iso_c_binding, Only: c_ptr
! .. Scalar Arguments ..
Type (c_ptr), Intent (Inout)         :: cpuser
Integer, Intent (Inout)              :: inform
Integer, Intent (In)                 :: nnzu, nvar
! .. Array Arguments ..
Real (Kind=nag_wp), Intent (In)      :: rinfo(32), stats(32), u(nnzu), &
                                     x(nvar)
Real (Kind=nag_wp), Intent (Inout)   :: ruser(*)
Integer, Intent (Inout)              :: iuser(*)
! .. Local Scalars ..
Integer                               :: i, io_unit
! .. Executable Statements ..
Continue
io_unit = iuser(2)
Write (io_unit,99999)
Write (io_unit,99998) x
Write (io_unit,99997)
Write (io_unit,99998) u
Write (io_unit,99996)
Write (io_unit,99995)(i,rinfo(i),i=1,32)
Write (io_unit,99994)
Write (io_unit,99995)(i,stats(i),i=1,32)
99999 Format ('Monitoring... ',/, ' X(*)')

```

```

99998  Format (1P,E14.6)
99997  Format ('  U(*)')
99996  Format ('  RINFO(32)')
99995  Format (4(I2,1P,E14.6,1X))
99994  Format ('  STATS(32)')
      Return
      End Subroutine mon
      End Module e04stfe_mod

Program e04stfe

!      .. Use Statements ..
      Use nag_library, Only: e04raf, e04ref, e04rgf, e04rhf, e04rjf, e04rkf, &
                             e04rlf, e04ryf, e04rzf, e04stf, e04zmf, nag_wp, &
                             x04acf
      Use e04stfe_mod, Only: confun, congrd, hess, mon, objfun, objgrd
      Use, Intrinsic          :: iso_c_binding, Only: c_null_ptr, &
                             c_ptr

!      .. Implicit None Statement ..
      Implicit None

!      .. Parameters ..
      Real (Kind=nag_wp), Parameter :: solve_timeout = 5._nag_wp
      Integer, Parameter           :: nout = 6

!      .. Local Scalars ..
      Type (c_ptr)                 :: cpuser, handle
      Real (Kind=nag_wp)           :: bigbnd
      Integer                      :: i, idlc, idx, ifail, ilinear, j, &
                             nclin, ncnln, nnzu, nvar, ry_ifail
      Character (60)               :: opt_s

!      .. Local Arrays ..
      Real (Kind=nag_wp)           :: b(8), bl(4), bu(4), linbl(2), &
                             linbu(2), nlnbl(1), nlnbu(1), &
                             rinfo(32), ruser(4), stats(32), x(4)

      Real (Kind=nag_wp), Allocatable :: u(:)
      Integer                      :: icolb(8), icolgd(4), icolh(10), &
                             idxfd(4), irowb(8), irowgd(4), &
                             irowh(10), iuser(2)

!      .. Intrinsic Procedures ..
      Intrinsic                    :: int

!      .. Executable Statements ..

      Write (nout,*) 'E04STF Example Program Results'
      Write (nout,*)
      Flush (nout)

      ifail = 0
      Call x04acf(66,'e04stf.out',1,ifail)
      Call x04acf(67,'e04stf.mon',1,ifail)
      Call x04acf(68,'e04stf.umon',1,ifail)

      Do ilinear = 0, 1
         handle = c_null_ptr
         bigbnd = 1.0E40_nag_wp
         nvar = 4
         nnzu = 0
         Call e04raf(handle,nvar,ifail)

         Write (opt_s,99992) bigbnd
         Call e04zmf(handle,opt_s,ifail)

!      Add simple bounds (x_i >= 0).
         bl(1:4) = 0.0_nag_wp
         bu(1:4) = bigbnd
         nnzu = nnzu + 2*nvar
         Call e04rhf(handle,nvar,bl,bu,ifail)

         iuser(1) = ilinear

!      Add linear objective
         ruser(1:4) = (/24.55_nag_wp,26.75_nag_wp,39.00_nag_wp,40.50_nag_wp/)
         If (ilinear==1) Then

```

```

    Call e04ref(handle,4,ruser(1:4),ifail)
Else
    idxfd(1:4) = (/1,2,3,4/)
    Call e04rgf(handle,4,idxfd,ifail)
End If

!    Add two linear constraints
nclin = 2
nnzu = nnzu + 2*nclin
linbl(1:2) = (/5.0_nag_wp,1.0_nag_wp/)
linbu(1:2) = (/bigbnd,1.0_nag_wp/)
irowb(1:8) = (/1,1,1,1,2,2,2,2/)
icolb(1:8) = (/1,2,3,4,1,2,3,4/)
b(1:8) = (/2.3_nag_wp,5.6_nag_wp,11.1_nag_wp,1.3_nag_wp,1.0_nag_wp, &
    1.0_nag_wp,1.0_nag_wp,1.0_nag_wp/)
idlc = 0
Call e04rjf(handle,nclin,linbl,linbu,nclin*nvar,irowb,icolb,b,idlc, &
    ifail)

!    Add one nonlinear constraint
ncnln = 1
nnzu = nnzu + 2*ncnln
nlnbl(1:1) = (/21.0_nag_wp/)
nlnbu(1:1) = (/bigbnd/)
irowgd(1:4) = (/1,1,1,1/)
icolgd(1:4) = (/1,2,3,4/)
Call e04rkf(handle,ncnln,nlnbl,nlnbu,4,irowgd,icolgd,ifail)

!    Define dense structure of the Hessian
idx = 1
Do i = 1, nvar
    Do j = i, nvar
        icolh(idx) = j
        irowh(idx) = i
        idx = idx + 1
    End Do
End Do

!    Hessian of nonlinear constraint
Call e04rlf(handle,1,idx-1,irowh,icolh,ifail)
If (ilinear/=1) Then
    Call e04rlf(handle,0,idx-1,irowh,icolh,ifail)
End If

Call e04ryf(handle,nout,'Overview',ifail)

!    call solver
x = (/1.0_nag_wp,1.0_nag_wp,1.0_nag_wp,1.0_nag_wp/)
Allocate (u(nnzu))

iuser(2) = 68
Call e04zmf(handle,'Monitoring File = 67',ifail)
Call e04zmf(handle,'Monitoring Level = 5',ifail)
Call e04zmf(handle,'Outer Iteration Limit = 26',ifail)
Call e04zmf(handle,'Print File = 66',ifail)
Call e04zmf(handle,'Print Level = 2',ifail)
Call e04zmf(handle,'Stop Tolerance 1 = 2.5e-8',ifail)
Call e04zmf(handle,'Time Limit = 60',ifail)

ifail = -1
Call e04stf(handle,objfun,objgrd,confun,congrd,hess,mon,nvar,x,nnzu,u, &
    rinfo,stats,iuser,ruser,cpuser,ifail)

ry_ifail = 0
Call e04ryf(handle,nout,'Options',ry_ifail)

If (ifail==0) Then
    Write (nout,99999)
    Write (nout,99995)(i,x(i),i=1,nvar)
    Write (nout,99998)
    Write (nout,99993)(i,u(2*i-1),i,u(2*i),i=1,nvar)

```

```

Write (nout,99997)
Write (nout,99994)(i,u(2*i-1+2*nvar),i,u(2*i+2*nvar),i=1,nclin)
Write (nout,99996)
Write (nout,99994)(i,u(2*i-1+2*nvar+2*nclin),i,                                     &
    u(2*i+2*nvar+2*nclin),i=1,ncnln)
Write (nout,99991) rinfo(1)
Write (nout,99990) rinfo(2)
Write (nout,99989) rinfo(3)
Write (nout,99988) rinfo(4)
Write (nout,99987) rinfo(5)
If (stats(8)<solve_timeout) Then
    Write (nout,99986)
Else
    Write (nout,99985) stats(8)
End If
Write (nout,99984) int(stats(1))
Write (nout,99983) int(stats(19))
Write (nout,99982) int(stats(5))
Write (nout,99981) int(stats(20))
Write (nout,99980) int(stats(21))
Write (nout,99979) int(stats(4))
End If

ifail = 0
Call e04ryf(handle,nout,'Overview',ifail)

99999 Format (/, 'Variables')
99998 Format ('Variable bound Lagrange multipliers')
99997 Format ('Linear constraints Lagrange multipliers')
99996 Format ('Nonlinear constraints Lagrange multipliers')
99995 Format (5X, 'x(', I10, ')', 17X, '=', 1P, E20.6)
99994 Format (5X, 'l+', (', I10, ')', 16X, '=', 1P, E20.6, /, 5X, 'l-(', I10, ')', 16X, '=', &
    1P, E20.6)
99993 Format (5X, 'zL(', I10, ')', 16X, '=', 1P, E20.6, /, 5X, 'zU(', I10, ')', 16X, '=', &
    1P, E20.6)
99992 Format ('Infinite Bound Size = ', 1P, E14.6)
99991 Format ('At solution, Objective minimum = ', 1P, E20.7)
99990 Format (' Constraint violation = ', 1P, E20.2)
99989 Format (' Dual infeasibility = ', 1P, E20.2)
99988 Format (' Complementarity = ', 1P, E20.2)
99987 Format (' KKT error = ', 1P, E20.2)
99986 Format ('Solved in allotted time limit')
99985 Format ('Solution took ', E10.3, ' sec, which is longer than expected')
99984 Format (' after iterations :', I11)
99983 Format (' after objective evaluations :', I11)
99982 Format (' after objective gradient evaluations :', I11)
99981 Format (' after constraint evaluations :', I11)
99980 Format (' after constraint gradient evaluations :', I11)
99979 Format (' after hessian evaluations :', I11)
Deallocate (u)
! release all memory held in the handle
Call e04rzf(handle,ifail)
Write (nout,*) '-----'
End Do
End Program e04stfe

```

## 10.2 Program Results

### E04STF Example Program Results

#### Overview

```

Status:                Problem and option settings are editable.
No of variables:       4
Objective function:    nonlinear
Simple bounds:         defined
Linear constraints:    2
Nonlinear constraints: 1
Matrix constraints:    not defined yet
Option settings
Begin of Options

```

```

Outer Iteration Limit      =                26      * U
Infinite Bound Size       =          1.00000E+40    * U
Print File                 =                66      * U
Print Level                =                 2      * U
Monitoring File           =                67      * U
Monitoring Level          =                 5      * U
Stats Time                 =                 No    * d
Stop Tolerance 1          =          2.50000E-08    * U
Hessian Mode               =                 Exact * S
Verify Derivatives        =                 Yes   * S
Time Limit                 =          6.00000E+01    * U
End of Options

Variables
  x(      1)              =          6.355216E-01
  x(      2)              =          2.066279E-10
  x(      3)              =          3.127019E-01
  x(      4)              =          5.177655E-02
Variable bound Lagrange multipliers
  zL(      1)              =          3.916168E-09
  zU(      1)              =          0.000000E+00
  zL(      2)              =          2.433326E-01
  zU(      2)              =          0.000000E+00
  zL(      3)              =          7.974843E-09
  zU(      3)              =          0.000000E+00
  zL(      4)              =          4.944607E-08
  zU(      4)              =          0.000000E+00
Linear constraints Lagrange multipliers
  l+(      1)              =          0.000000E+00
  l-(      1)              =          4.105411E-01
  l+(      2)              =          0.000000E+00
  l-(      2)              =          5.803551E-01
Nonlinear constraints Lagrange multipliers
  l+(      1)              =          0.000000E+00
  l-(      1)              =          1.837124E+01
At solution, Objective minimum =          2.9894378E+01
                  Constraint violation =          1.11E-16
                  Dual infeasibility  =          6.71E-12
                  Complementarity     =          2.56E-09
                  KKT error           =          2.56E-09
Solved in allotted time limit
  after iterations          :                8
  after objective evaluations :                9
  after objective gradient evaluations :                9
  after constraint evaluations :                9
  after constraint gradient evaluations :                9
  after hessian evaluations :                8
Overview
  Status:                  Solver finished, only options can be changed.
  No of variables:         4
  Objective function:      nonlinear
  Simple bounds:           defined
  Linear constraints:      2
  Nonlinear constraints:   1
  Matrix constraints:      not defined
-----
Overview
  Status:                  Problem and option settings are editable.
  No of variables:         4
  Objective function:      linear
  Simple bounds:           defined
  Linear constraints:      2
  Nonlinear constraints:   1
  Matrix constraints:      not defined yet
Option settings
Begin of Options
  Outer Iteration Limit    =                26      * U
  Infinite Bound Size      =          1.00000E+40    * U
  Print File               =                66      * U
  Print Level              =                 2      * U
  Monitoring File          =                67      * U

```

```

Monitoring Level           =                    5          * U
Stats Time                 =                    No          * d
Stop Tolerance 1          =          2.50000E-08          * U
Hessian Mode              =                    Exact        * S
Verify Derivatives        =                    Yes         * S
Time Limit                 =          6.00000E+01          * U
End of Options

Variables
x(          1)             =          6.355216E-01
x(          2)             =          2.066279E-10
x(          3)             =          3.127019E-01
x(          4)             =          5.177655E-02
Variable bound Lagrange multipliers
zL(          1)             =          3.916168E-09
zU(          1)             =          0.000000E+00
zL(          2)             =          2.433326E-01
zU(          2)             =          0.000000E+00
zL(          3)             =          7.974843E-09
zU(          3)             =          0.000000E+00
zL(          4)             =          4.944607E-08
zU(          4)             =          0.000000E+00
Linear constraints Lagrange multipliers
l+(          1)             =          0.000000E+00
l-(          1)             =          4.105411E-01
l+(          2)             =          0.000000E+00
l-(          2)             =          5.803551E-01
Nonlinear constraints Lagrange multipliers
l+(          1)             =          0.000000E+00
l-(          1)             =          1.837124E+01
At solution, Objective minimum =          2.9894378E+01
                  Constraint violation =          1.11E-16
                  Dual infeasibility =          6.71E-12
                  Complementarity =          2.56E-09
                  KKT error =          2.56E-09
Solved in allotted time limit
after iterations           :                    8
after objective evaluations :                    9
after objective gradient evaluations :                    9
after constraint evaluations :                    9
after constraint gradient evaluations :                    9
after hessian evaluations  :                    8
Overview
Status:                   Solver finished, only options can be changed.
No of variables:          4
Objective function:       linear
Simple bounds:            defined
Linear constraints:        2
Nonlinear constraints:     1
Matrix constraints:       not defined
-----

```

## 11 Algorithmic Details

E04STF is an implementation of IPOPT (see Wächter and Biegler (2006)) that is fully supported and maintained by NAG. It uses Harwell packages MA97 for the underlying sparse linear algebra factorization and MC68 approximate minimum degree algorithm for the ordering. Any issues relating to E04STF should be directed to NAG who assume all responsibility for the E04STF routine and its implementation.

In the remainder of this section, we repeat part of Section 2.1 of Wächter and Biegler (2006).

To simplify notation, we describe the method for the problem formulation

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad f(x) \quad (2)$$

$$\text{subject to} \quad g(x) = 0 \quad (3)$$



$$x \geq 0. \quad (4)$$

Range constraints of the form  $l \leq c(x) \leq u$  can be expressed in this formulation by introducing slack variables  $x_s \geq 0$ ,  $x_t \geq 0$  (increasing  $n$  by 2) and defining new equality constraints  $g(x, x_s) \equiv c(x) - l - x_s = 0$  and  $g(x, x_t) \equiv u - c(x) - x_t = 0$ .

E04STF, like the methods discussed in Williams and Lang (2013), Byrd *et al.* (2000), Conn *et al.* (2000) and Fiacco and McCormick (1990), computes (approximate) solutions for a sequence of barrier problems

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \varphi_\mu(x) f(x) - \mu \sum_{i=1}^n \ln(x^{(i)}) \quad (5)$$

$$\text{subject to} \quad g(x) = 0 \quad (6)$$

for a decreasing sequence of barrier parameters  $\mu$  converging to zero.

The algorithm may be interpreted as a homotopy method to the primal-dual equations,

$$\nabla f(x) + \nabla g(x)\lambda - z = 0 \quad (7)$$

$$g(x) = 0 \quad (8)$$

$$XZe - \mu e = 0 \quad (9)$$

with the homotopy parameter  $\mu$ , which is driven to zero (see e.g., Byrd *et al.* (1997) and Gould *et al.* (2001)). Here,  $X\text{diag}(x)$  for a vector  $x$  (similarly  $z\text{diag}(z)$ , etc.), and  $e$  stands for the vector of all ones for appropriate dimension, while  $\lambda \in \mathbb{R}^m$  and  $z \in \mathbb{R}^n$  correspond to the Lagrange multipliers for the equality constraints (3) and the bound constraints (4), respectively.

Note, that the equations (7), (8) and (9) for  $\mu = 0$  together with ‘ $x, z \geq 0$ ’ are the Karush–Kuhn–Tucker (KKT) conditions for the original problem (2), (3) and (4). Those are the first order optimality conditions for (2), (3) and (4) if constraint qualifications are satisfied (Conn *et al.* (2000)).

Starting from an initial point supplied in X, E04STF computes an approximate solution to the barrier problem (5) and (6) for a fixed value of  $\mu$  (by default, 0.1), then decreases the barrier parameter, and continues the solution of the next barrier problem from the approximate solution of the previous one.

A sophisticated overall termination criterion for the algorithm is used to overcome potential difficulties when the Lagrange multipliers become large. This can happen, for example, when the gradients of the active constraints are nearly linear dependent. The termination criterion is described in detail by Wächter and Biegler (2006) (also see below Section 11.1).

## 11.1 Stopping Criteria

Using the individual parts of the primal-dual equations (7), (8) and (9), we define the optimality error for the barrier problem as

$$E_\mu(x, \lambda, z) \max \left\{ \frac{\|\nabla f(x) + \nabla g(x)\lambda - z\|_\infty}{s_d}, \|g(x)\|_\infty, \frac{\|XZe - \mu e\|_\infty}{s_c} \right\} \quad (10)$$

with scaling parameters  $s_d, s_c \geq 1$  defined below (not to be confused with NLP scaling factors described in Section 11.2). By  $E_0(x, \lambda, z)$  we denote (10) with  $\mu = 0$ ; this measures the optimality error for the original problem (2), (3) and (4). The overall algorithm terminates if an approximate solution  $(\tilde{x}_*, \tilde{\lambda}_*, \tilde{z}_*)$  (including multiplier estimates) satisfying

$$E_0(\tilde{x}_*, \tilde{\lambda}_*, \tilde{z}_*) \leq \epsilon_{tol} \quad (11)$$

is found, where  $\epsilon_{tol} > 0$  is the user provided error tolerance in optional parameter **Stop Tolerance 1**.

Even if the original problem is well scaled, the multipliers  $\lambda$  and  $z$  might become very large, for example, when the gradients of the active constraints are (nearly) linearly dependent at a solution of (2), (3) and (4). In this case, the algorithm might encounter numerical difficulties satisfying the unscaled primal-dual equations (7), (8) and (9) to a tight tolerance. In order to adapt the termination criteria to

handle such circumstances, we choose the scaling factors

$$s_d \max \left\{ s_{\max}, \frac{\|\lambda\|_1 + \|z\|_1}{(m+n)} \right\} / s_{\max} \quad s_c \max \left\{ s_{\max}, \frac{\|z\|_1}{n} \right\} / s_{\max}$$

in (10). In this way, a component of the optimality error is scaled, whenever the average value of the multipliers becomes larger than a fixed number  $s_{\max} \geq 1$  ( $s_{\max} = 100$  in our implementation). Also note, in the case that the multipliers diverge,  $E_0(x, \lambda, z)$  can only become small, if a Fritz John point for (2), (3) and (4) is approached, or if the primal variables diverge as well.

## 11.2 Scaling the NLP

Ideally, the formulated problem should be scaled so that, near the solution, all function gradients (objective and constraints), when nonzero, are of a similar order of a magnitude. E04STF will compute automatic NLP scaling factors for the objective and constraint functions (but not the decision variables) and apply them if large imbalances of scale are detected. This rescaling is only computed at the starting point. References to scaled or unscaled objective or constraints in Section 9.1 and Section 11 should be understood in this context.

## 12 Optional Parameters

Several optional parameters in E04STF define choices in the problem specification or the algorithm logic. In order to reduce the number of formal arguments of E04STF 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.

If any options are set by the solver (typically those with the choice of AUTO), their value can be retrieved by E04ZNF. If the solver is called again, any such arguments are reset to their default values and the decision is made again.

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

### Defaults

#### Hessian Mode

#### Infinite Bound Size

#### Monitoring File

#### Monitoring Level

#### Outer Iteration Limit

#### Print File

#### Print Level

#### Stats Time

#### Stop Tolerance 1

#### Time Limit

#### Verify Derivatives

### 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  $\epsilon$  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.

**Hessian Mode**  $a$  Default = AUTO

This argument specifies whether the Hessian will be supplied by the user (in HX) or approximated by E04STF using a limited-memory quasi-Newton L-BFGS method. In the AUTO setting, if no Hessian structure has been registered in the problem with a call to E04RLF, and there are explicitly nonlinear user-supplied functions, then the Hessian will be approximated. Otherwise HESS will be called if and only if any of E04RGF or E04RKF have been used to define the problem. Approximating the Hessian is likely to require more iterations to achieve convergence but will reduce the time spent in user-supplied functions.

Constraint: **Hessian Mode** = AUTO, EXACT or APPROXIMATE.

**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.

It also serves as a limit for the objective function to be considered unbounded (IFAIL = 54).

Constraint: **Infinite Bound Size**  $\geq 1000$ .

**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**  $\geq -1$ .

**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:  $0 \leq$  **Monitoring Level**  $\leq 5$ .

**Outer Iteration Limit**  $i$  Default = 100

The maximum number of iterations to be performed by E04STF. Setting the option too low might lead to IFAIL = 22.

Constraint: **Outer Iteration Limit**  $\geq 0$ .

**Print File**  $i$  Default = 6

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 information output to this unit is controlled by **Print Level**.

Constraint: **Print File**  $\geq -1$ .

**Print Level** *i* Default = 2

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

***i* Output**

- 0 No output from the solver (except a one-time banner)
- 1 Additionally, derivative check information, the Header and Summary.
- 2 Additionally, the Iteration log.
- 3, 4 Additionally, details of each iteration with scalar quantities printed.
- 5 Additionally, individual components of arrays are printed resulting in large output.

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

**Stats Time** *a* Default = NO

This argument allows you to turn 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.

Constraint: **Stats Time** = YES or NO.

**Stop Tolerance 1** *r* Default =  $\max(10^{-6}, \sqrt{\epsilon})$

This option sets the value  $\epsilon_{\text{tol}}$  which is used for optimality and complementarity tests from KKT conditions See Section 11.1.

Constraint: **Stop Tolerance 1**  $> \epsilon$ .

**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 a corresponding error message.

Constraint: **Time Limit**  $> 0$ .

**Verify Derivatives** *a* Default = AUTO

This argument specifies whether the routine should perform numerical checks on the consistency of the user-supplied functions. It is recommended that such checks are enabled when first developing the formulation of the problem. Option AUTO will perform the checks unless it is determined that there are no explicitly nonlinear user-supplied functions.

Constraint: **Verify Derivatives** = AUTO, YES or NO.