# NAG Library Routine Document E04UCF/E04UCA 


#### Abstract

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 9 of this document. If, however, you wish to reset some or all of the settings please refer to Section 10 for a detailed description of the algorithm, to Section 11 for a detailed description of the specification of the optional parameters and to Section 12 for a detailed description of the monitoring information produced by the routine.


## 1 Purpose

E04UCF/E04UCA is designed to minimize an arbitrary smooth function subject to constraints (which may include simple bounds on the variables, linear constraints and smooth nonlinear constraints) using a sequential quadratic programming (SQP) method. As many first derivatives as possible should be supplied by you; any unspecified derivatives are approximated by finite differences. It is not intended for large sparse problems.
E04UCF/E04UCA may also be used for unconstrained, bound-constrained and linearly constrained optimization.

E04UCF/E04UCA uses forward communication for evaluating the objective function, the nonlinear constraint functions, and any of their derivatives.
E04UCA is a version of E04UCF that has additional parameters in order to make it safe for use in multithreaded applications (see Section 5). The initialization routine E04WBF must have been called before calling E04UCA.

## 2 Specification

### 2.1 Specification for E04UCF

```
SUBROUTINE EO4UCF (N, NCLIN, NCNLN, LDA, LDCJ, LDR, A, BL, BU, CONFUN,
INTEGER N, NCLIN, NCNLN, LDA, LDCJ, LDR, ITER, ISTATE (N+NCLIN+NCNLN) , IWORK (LIWORK) , LIWORK, LWORK,
REAL (KIND=nag_wp) A (LDA, *) , BL (N+NCLIN+NCNLN), BU (N+NCLIN+NCNLN)

\subsection*{2.2 Specification for E04UCA}
```

SUBROUTINE EO4UCA (N, NCLIN, NCNLN, LDA, LDCJ, LDR, A, BL, BU, CONFUN, \&
OBJFUN, ITER, ISTATE, C, CJAC, CLAMDA, OBJF, OBJGRD, R, \&
X, IWORK, LIWORK, WORK, LWORK, IUSER, RUSER, LWSAV, \&
IWSAV, RWSAV, IFAIL)
INTEGER N, NCLIN, NCNLN, LDA, LDCJ, LDR, ITER
IUSER(*), IWSAV(610), IFAIL
REAL (KIND=nag_wp) A(LDA,*), BL(N+NCLIN+NCNLN), BU(N+NCLIN+NCNLN),

EXTERNAL CONFUN, OBJFUN
Before calling E04UCA, or either of the option setting routines E04UDA or E04UEA, E04WBF must be called. The specification for E04WBF is:

```
SUBROUTINE EO4WBF (RNAME, CWSAV, LCWSAV, LWSAV, LLWSAV, IWSAV, LIWSAV,
    RWSAV, LRWSAV, IFAIL)
INTEGER LCWSAV, LLWSAV, IWSAV(LIWSAV), LIWSAV, LRWSAV, IFAIL
REAL (KIND=nag_wp) RWSAV(LRWSAV)
LOGICAL LWSAV(LLWSAV)
CHARACTER(*) RNAME
CHARACTER(80) CWSAV(LCWSAV)
```

E04WBF should be called with RNAME = 'E04UCA'. LCWSAV, LLWSAV, LIWSAV and LRWSAV, the declared lengths of CWSAV, LWSAV, IWSAV and RWSAV respectively, must satisfy:

$$
\begin{aligned}
& \text { LCWSAV } \geq 1 ; \\
& \text { LLWSAV } \geq 120 ; \\
& \text { LIWSAV } \geq 610 ; \\
& \text { LRWSAV } \geq 475 .
\end{aligned}
$$

The contents of the arrays CWSAV, LWSAV, IWSAV and RWSAV must not be altered between calling routines E04UCA, E04UDA, E04UEA and E04WBF.

## 3 Description

E04UCF/E04UCA is designed to solve the nonlinear programming problem - the minimization of a smooth nonlinear function subject to a set of constraints on the variables. The problem is assumed to be stated in the following form:

$$
\underset{x \in R^{z}}{\operatorname{minimize}} F(x) \quad \text { subject to } \quad l \leq\left(\begin{array}{c}
x  \tag{1}\\
A_{L} x \\
c(x)
\end{array}\right) \leq u,
$$

where $F(x)$ (the objective function) is a nonlinear function, $A_{L}$ is an $n_{L}$ by $n$ constant matrix, and $c(x)$ is an $n_{N}$ element vector of nonlinear constraint functions. (The matrix $A_{L}$ and the vector $c(x)$ may be empty.) The objective function and the constraint functions are assumed to be smooth, i.e., at least twicecontinuously differentiable. (The method of E04UCF/E04UCA will usually solve (1) if there are only isolated discontinuities away from the solution.)
Note that although the bounds on the variables could be included in the definition of the linear constraints, we prefer to distinguish between them for reasons of computational efficiency. For the same reason, the linear constraints should not be included in the definition of the nonlinear constraints. Upper and lower bounds are specified for all the variables and for all the constraints. An equality constraint can be specified by setting $l_{i}=u_{i}$. If certain bounds are not present, the associated elements of $l$ or $u$ can be set to special values that will be treated as $-\infty$ or $+\infty$. (See the description of the optional parameter Infinite Bound Size.)

If there are no nonlinear constraints in (1) and $F$ is linear or quadratic, then it will generally be more efficient to use one of E04MFF/E04MFA, E04NCF/E04NCA or E04NFF/E04NFA, or E04NKF/E04NKA if the problem is large and sparse. If the problem is large and sparse and does have nonlinear constraints, then E04UGF/E04UGA should be used, since E04UCF/E04UCA treats all matrices as dense.
You must supply an initial estimate of the solution to (1), together with subroutines that define $F(x), c(x)$ and as many first partial derivatives as possible; unspecified derivatives are approximated by finite differences.
The objective function is defined by OBJFUN, and the nonlinear constraints are defined by CONFUN. On every call, these subroutines must return appropriate values of the objective and nonlinear constraints. You should also provide the available partial derivatives. Any unspecified derivatives are approximated by finite differences; see Section 11.1 for a discussion of the optional parameter Derivative Level. Note that if there are any nonlinear constraints then the first call to CONFUN will precede the first call to OBJFUN.

For maximum reliability, it is preferable for you to provide all partial derivatives (see Chapter 8 of Gill et al. (1981), for a detailed discussion). If all gradients cannot be provided, it is similarly advisable to provide as many as possible. While developing OBJFUN and CONFUN, the optional parameter Verify should be used to check the calculation of any known gradients.
The method used by E04UCF/E04UCA is described in detail in Section 10.
E04UFF/E04UFA is an alternative routine which uses exactly the same method, but uses reverse communication for evaluating the objective and constraint functions.

## 4 References

Dennis J E Jr and Moré J J (1977) Quasi-Newton methods, motivation and theory SIAM Rev. 19 46-89
Dennis J E Jr and Schnabel R B (1981) A new derivation of symmetric positive-definite secant updates nonlinear programming (eds O L Mangasarian, R R Meyer and S M Robinson) 4 167-199 Academic Press

Dennis J E Jr and Schnabel R B (1983) Numerical Methods for Unconstrained Optimization and Nonlinear Equations Prentice-Hall
Fletcher R (1987) Practical Methods of Optimization (2nd Edition) Wiley
Gill P E, Hammarling S, Murray W, Saunders M A and Wright M H (1986) Users' guide for LSSOL (Version 1.0) Report SOL 86-1 Department of Operations Research, Stanford University
Gill P E, Murray W, Saunders M A and Wright M H (1984a) Procedures for optimization problems with a mixture of bounds and general linear constraints ACM Trans. Math. Software 10 282-298

Gill P E, Murray W, Saunders M A and Wright M H (1984b) Users' guide for SOL/QPSOL version 3.2 Report SOL 84-5 Department of Operations Research, Stanford University
Gill P E, Murray W, Saunders M A and Wright M H (1986a) Some theoretical properties of an augmented Lagrangian merit function Report SOL 86-6R Department of Operations Research, Stanford University
Gill P E, Murray W, Saunders M A and Wright M H (1986b) Users' guide for NPSOL (Version 4.0): a Fortran package for nonlinear programming Report SOL 86-2 Department of Operations Research, Stanford University
Gill P E, Murray W and Wright M H (1981) Practical Optimization Academic Press
Hock W and Schittkowski K (1981) Test Examples for Nonlinear Programming Codes. Lecture Notes in Economics and Mathematical Systems 187 Springer-Verlag

Powell M J D (1974) Introduction to constrained optimization Numerical Methods for Constrained Optimization (eds P E Gill and W Murray) 1-28 Academic Press
Powell M J D (1983) Variable metric methods in constrained optimization Mathematical Programming: the State of the Art (eds A Bachem, M Grötschel and B Korte) 288-311 Springer-Verlag

## 5 Parameters

1: N - INTEGER Input
On entry: $n$, the number of variables.
Constraint: $\mathrm{N}>0$.
2: NCLIN - INTEGER Input
On entry: $n_{L}$, the number of general linear constraints.
Constraint: NCLIN $\geq 0$.

3: NCNLN - INTEGER
On entry: $n_{N}$, the number of nonlinear constraints.
Constraint: NCNLN $\geq 0$.
4: LDA - INTEGER
Input
On entry: the first dimension of the array A as declared in the (sub)program from which E04UCF/ E04UCA is called.

Constraint: LDA $\geq \max (1$, NCLIN $)$.
5: LDCJ - INTEGER
Input
On entry: the first dimension of the array CJAC as declared in the (sub)program from which E04UCF/E04UCA is called.

Constraint: LDCJ $\geq \max (1, \mathrm{NCNLN})$.
6: LDR - INTEGER
Input
On entry: the first dimension of the array R as declared in the (sub)program from which E04UCF/ E04UCA is called.

Constraint: $\mathrm{LDR} \geq \mathrm{N}$.
7: A(LDA,*) - REAL (KIND=nag_wp) array
Input
Note: the second dimension of the array A must be at least N if NCLIN $>0$, and at least 1 otherwise.

On entry: the $i$ th row of A contains the $i$ th row of the matrix $A_{L}$ of general linear constraints in (1). That is, the $i$ th row contains the coefficients of the $i$ th general linear constraint, for $i=1,2, \ldots$, NCLIN.
If $\mathrm{NCLIN}=0$, the array A is not referenced.

$$
\begin{array}{ll}
\text { BL }(\mathrm{N}+\mathrm{NCLIN}+\mathrm{NCNLN})-\text { REAL }\left(\mathrm{KIND=nag} \mathrm{\left.\_wp\right)} \mathrm{array}\right. & \text { Input } \\
\mathrm{BU}(\mathrm{~N}+\mathrm{NCLIN}+\mathrm{NCNLN})-\text { REAL (KIND=nag_wp) array } & \text { Input }
\end{array}
$$

On entry: BL must contain the lower bounds and BU the upper bounds for all the constraints in the following order. The first $n$ elements of each array must contain the bounds on the variables, the next $n_{L}$ elements the bounds for the general linear constraints (if any) and the next $n_{N}$ elements the bounds for the general nonlinear constraints (if any). To specify a nonexistent lower bound (i.e., $l_{j}=-\infty$ ), set $\mathrm{BL}(j) \leq-$ bigbnd, and to specify a nonexistent upper bound (i.e., $u_{j}=+\infty$ ), set $\mathrm{BU}(j) \geq$ bigbnd; the default value of bigbnd is $10^{20}$, but this may be changed by the optional parameter Infinite Bound Size. To specify the $j$ th constraint as an equality, set $\mathrm{BL}(j)=\mathrm{BU}(j)=\beta$, say, where $|\beta|<$ bigbnd .
Constraints:

$$
\begin{aligned}
& \mathrm{BL}(j) \leq \mathrm{BU}(j), \text { for } j=1,2, \ldots, \mathrm{~N}+\mathrm{NCLIN}+\mathrm{NCNLN} ; \\
& \text { if } \mathrm{BL}(j)=\mathrm{BU}(j)=\beta,|\beta|<\text { bigbnd. }
\end{aligned}
$$

10: CONFUN - SUBROUTINE, supplied by the NAG Library or the user. External Procedure CONFUN must calculate the vector $c(x)$ of nonlinear constraint functions and (optionally) its Jacobian $\left(=\frac{\partial c}{\partial x}\right)$ for a specified $n$-element vector $x$. If there are no nonlinear constraints (i.e., NCNLN $=0$ ), CONFUN will never be called by E04UCF/E04UCA and CONFUN may be the dummy routine E04UDM. (E04UDM is included in the NAG Library.) If there are nonlinear constraints, the first call to CONFUN will occur before the first call to OBJFUN.

The specification of CONFUN is:

```
SUBROUTINE CONFUN (MODE, NCNLN, N, LDCJ, NEEDC, X, C, CJAC, NSTATE, &
    IUSER, RUSER)
INTEGER MODE, NCNLN, N, LDCJ, NEEDC(NCNLN), NSTATE, &
    IUSER(*)
REAL (KIND=nag_wp) X(N), C(NCNLN), CJAC(LDCJ,N), RUSER(*)
```

1: MODE - INTEGER
On entry: indicates which values must be assigned during each call of CONFUN. Only the following values need be assigned, for each value of $i$ such that $\operatorname{NEEDC}(i)>0$ :

$$
\begin{array}{r}
\mathrm{MODE}=0 \\
\mathrm{C}(i) .
\end{array}
$$

$\mathrm{MODE}=1$
All available elements in the $i$ th row of CJAC.
MODE $=2$
$\mathrm{C}(i)$ and all available elements in the $i$ th row of CJAC.
On exit: may be set to a negative value if you wish to terminate the solution to the current problem. In this case E04UCF/E04UCA will terminate with IFAIL set to MODE.

NCNLN - INTEGER
On entry: $n_{N}$, the number of nonlinear constraints.
3: $\quad \mathrm{N}$ - INTEGER
On entry: $n$, the number of variables.
LDCJ - INTEGER
Input
On entry: the first dimension of the array CJAC as declared in the (sub)program from which E04UCF/E04UCA is called.

5: $\quad$ NEEDC(NCNLN) - INTEGER array
Input
On entry: the indices of the elements of C and/or CJAC that must be evaluated by CONFUN. If NEEDC $(i)>0$, the $i$ th element of C and/or the available elements of the $i$ th row of CJAC (see parameter MODE) must be evaluated at $x$.

6:
$\mathrm{X}(\mathrm{N})$ - REAL (KIND=nag_wp) array
Input
On entry: $x$, the vector of variables at which the constraint functions and/or the available elements of the constraint Jacobian are to be evaluated.

7: $\quad \mathrm{C}(\mathrm{NCNLN})-$ REAL (KIND=$=$ nag_wp) array
Output
On exit: if $\operatorname{NEEDC}(i)>0$ and $\mathrm{MODE}=0$ or $2, \mathrm{C}(i)$ must contain the value of the $i$ th constraint at $x$. The remaining elements of C , corresponding to the non-positive elements of NEEDC, are ignored.

On entry: unless Derivative Level $=2$ or 3, the elements of CJAC are set to special values which enable E04UCF/E04UCA to detect whether they are changed by CONFUN.

On exit: if $\operatorname{NEEDC}(i)>0$ and $\mathrm{MODE}=1$ or 2 , the $i$ th row of CJAC must contain the available elements of the vector $\nabla c_{i}$ given by

$$
\nabla c_{i}=\left(\frac{\partial c_{i}}{\partial x_{1}}, \frac{\partial c_{i}}{\partial x_{2}}, \ldots, \frac{\partial c_{i}}{\partial x_{n}}\right)^{\mathrm{T}}
$$

where $\frac{\partial c_{i}}{\partial x_{j}}$ is the partial derivative of the $i$ th constraint with respect to the $j$ th variable, evaluated at the point $x$. See also the parameter NSTATE. The remaining rows of CJAC, corresponding to non-positive elements of NEEDC, are ignored.

If all elements of the constraint Jacobian are known (i.e., Derivative Level $=2$ or 3), any constant elements may be assigned to CJAC one time only at the start of the optimization. An element of CJAC that is not subsequently assigned in CONFUN will retain its initial value throughout. Constant elements may be loaded into CJAC either before the call to E04UCF/E04UCA or during the first call to CONFUN (signalled by the value
NSTATE $=1$ ). The ability to preload constants is useful when many Jacobian elements are identically zero, in which case CJAC may be initialized to zero and nonzero elements may be reset by CONFUN.
Note that constant nonzero elements do affect the values of the constraints. Thus, if $\operatorname{CJAC}(i, j)$ is set to a constant value, it need not be reset in subsequent calls to CONFUN, but the value CJAC $(i, j) \times \mathrm{X}(j)$ must nonetheless be added to $\mathrm{C}(i)$. For example, if $\operatorname{CJAC}(1,1)=2$ and $\operatorname{CJAC}(1,2)=-5$ then the term $2 \times \mathrm{X}(1)-5 \times \mathrm{X}(2)$ must be included in the definition of $\mathrm{C}(1)$.
It must be emphasized that, if Derivative Level $=0$ or 1 , unassigned elements of CJAC are not treated as constant; they are estimated by finite differences, at nontrivial expense. If you do not supply a value for the optional parameter Difference Interval, an interval for each element of $x$ is computed automatically at the start of the optimization. The automatic procedure can usually identify constant elements of CJAC, which are then computed once only by finite differences.

9: NSTATE - INTEGER
Input
On entry: if NSTATE $=1$ then E04UCF/E04UCA is calling CONFUN for the first time. This parameter setting allows you to save computation time if certain data must be read or calculated only once.

IUSER $(*)$ - INTEGER array User Workspace
$\operatorname{RUSER}(*)$ - REAL (KIND=nag_wp) array User Workspace
CONFUN is called with the parameters IUSER and RUSER as supplied to E04UCF/ E04UCA. You are free to use the arrays IUSER and RUSER to supply information to CONFUN as an alternative to using COMMON global variables.

CONFUN must either be a module subprogram USEd by, or declared as EXTERNAL in, the (sub)program from which E04UCF/E04UCA is called. Parameters denoted as Input must not be changed by this procedure.
CONFUN should be tested separately before being used in conjunction with E04UCF/E04UCA. See also the description of the optional parameter Verify.

11: OBJFUN - SUBROUTINE, supplied by the user.
External Procedure
OBJFUN must calculate the objective function $F(x)$ and (optionally) its gradient $g(x)=\frac{\partial F}{\partial x}$ for a specified $n$-vector $x$.

The specification of OBJFUN is:

```
SUBROUTINE OBJFUN (MODE, N, X, OBJF, OBJGRD, NSTATE, IUSER, RUSER)
INTEGER MODE, N, NSTATE, IUSER(*)
REAL (KIND=nag_wp) X(N), OBJF,OBJGRD(N), RUSER(*)
```

1: MODE - INTEGER
On entry: indicates which values must be assigned during each call of OBJFUN. Only the following values need be assigned:
MODE $=0$
OBJF.
MODE $=1$
All available elements of OBJGRD.
$\mathrm{MODE}=2$
OBJF and all available elements of OBJGRD.
On exit: may be set to a negative value if you wish to terminate the solution to the current problem. In this case E04UCF/E04UCA will terminate with IFAIL set to MODE.

2: N - INTEGER Input
On entry: $n$, the number of variables.
3: $\quad \mathrm{X}(\mathrm{N})$ - REAL (KIND=nag_wp) array Input

On entry: $x$, the vector of variables at which the objective function and/or all available elements of its gradient are to be evaluated.

OBJF - REAL (KIND=nag_wp)
Output
On exit: if MODE $=0$ or 2 , OBJF must be set to the value of the objective function at $x$.
5: $\quad \operatorname{OBJGRD}(\mathrm{N})-\operatorname{REAL}(\mathrm{KIND}=$ nag_wp) array
On entry: the elements of OBJGRD are set to special values which enable E04UCF/ E04UCA to detect whether they are changed by OBJFUN.

On exit: if $\mathrm{MODE}=1$ or 2 , OBJGRD must return the available elements of the gradient evaluated at $x$.

6: NSTATE - INTEGER
Input
On entry: if NSTATE $=1$ then E04UCF/E04UCA is calling OBJFUN for the first time. This parameter setting allows you to save computation time if certain data must be read or calculated only once.

IUSER $(*)$ - INTEGER array User Workspace
$\operatorname{RUSER}(*)$ - REAL (KIND=nag_wp) array User Workspace
OBJFUN is called with the parameters IUSER and RUSER as supplied to E04UCF/ E04UCA. You are free to use the arrays IUSER and RUSER to supply information to OBJFUN as an alternative to using COMMON global variables.

OBJFUN must either be a module subprogram USEd by, or declared as EXTERNAL in, the (sub)program from which E04UCF/E04UCA is called. Parameters denoted as Input must not be changed by this procedure.
OBJFUN should be tested separately before being used in conjunction with E04UCF/E04UCA. See also the description of the optional parameter Verify.

12: ITER - INTEGER
On exit: the number of major iterations performed.
13: $\operatorname{ISTATE}(\mathrm{N}+\mathrm{NCLIN}+\mathrm{NCNLN})-$ INTEGER array
On entry: need not be set if the (default) optional parameter Cold Start is used.

If the optional parameter Warm Start has been chosen, the elements of ISTATE corresponding to the bounds and linear constraints define the initial working set for the procedure that finds a feasible point for the linear constraints and bounds. The active set at the conclusion of this procedure and the elements of ISTATE corresponding to nonlinear constraints then define the initial working set for the first QP subproblem. More precisely, the first $n$ elements of ISTATE refer to the upper and lower bounds on the variables, the next $n_{L}$ elements refer to the upper and lower bounds on $A_{L} x$, and the next $n_{N}$ elements refer to the upper and lower bounds on $c(x)$. Possible values for $\operatorname{ISTATE}(j)$ are as follows:

## ISTATE $(\boldsymbol{j})$

0 The corresponding constraint is not in the initial QP working set.
1 This inequality constraint should be in the working set at its lower bound.
2 This inequality constraint should be in the working set at its upper bound.
3 This equality constraint should be in the initial working set. This value must not be specified unless $\operatorname{BL}(j)=\mathrm{BU}(j)$.

The values $-2,-1$ and 4 are also acceptable but will be modified by the routine. If E04UCF/ E04UCA has been called previously with the same values of N, NCLIN and NCNLN, ISTATE already contains satisfactory information. The routine also adjusts (if necessary) the values supplied in X to be consistent with ISTATE.
Constraint: $-2 \leq \operatorname{ISTATE}(j) \leq 4$, for $j=1,2, \ldots, \mathrm{~N}+\mathrm{NCLIN}+\mathrm{NCNLN}$.
On exit: the status of the constraints in the QP working set at the point returned in X. The significance of each possible value of $\operatorname{ISTATE}(j)$ is as follows:

## ISTATE ( $\boldsymbol{j}$ )

-2 This constraint violates its lower bound by more than the appropriate feasibility tolerance (see the optional parameters Linear Feasibility Tolerance and
Nonlinear Feasibility Tolerance). This value can occur only when no feasible point can be found for a QP subproblem.
$-1 \quad$ This constraint violates its upper bound by more than the appropriate feasibility tolerance (see the optional parameters Linear Feasibility Tolerance and
Nonlinear Feasibility Tolerance). This value can occur only when no feasible point can be found for a QP subproblem.
0 The constraint is satisfied to within the feasibility tolerance, but is not in the QP working set.

1 This inequality constraint is included in the QP working set at its lower bound.
2 This inequality constraint is included in the QP working set at its upper bound.
3 This constraint is included in the QP working set as an equality. This value of ISTATE can occur only when $\operatorname{BL}(j)=\operatorname{BU}(j)$.

On exit: if NCNLN $>0, \mathrm{C}(i)$ contains the value of the $i$ th nonlinear constraint function $c_{i}$ at the final iterate, for $i=1,2, \ldots$, NCNLN.
If $\mathrm{NCNLN}=0$, the array C is not referenced.
CJAC(LDCJ,*) - REAL (KIND=nag_wp) array
Input/Output
Note: the second dimension of the array CJAC must be at least N if $\mathrm{NCNLN}>0$, and at least 1 otherwise.

On entry: in general, CJAC need not be initialized before the call to E04UCF/E04UCA. However, if Derivative Level $=2$ or 3, you may optionally set the constant elements of CJAC (see parameter NSTATE in the description of CONFUN). Such constant elements need not be re-assigned on subsequent calls to CONFUN.

On exit: if NCNLN $>0$, CJAC contains the Jacobian matrix of the nonlinear constraint functions at the final iterate, i.e., $\operatorname{CJAC}(i, j)$ contains the partial derivative of the $i$ th constraint function with respect to the $j$ th variable, for $i=1,2, \ldots, \mathrm{NCNLN}$ and $j=1,2, \ldots, \mathrm{~N}$. (See the discussion of parameter CJAC under CONFUN.)
If $\mathrm{NCNLN}=0$, the array CJAC is not referenced.
CLAMDA $(\mathrm{N}+\mathrm{NCLIN}+\mathrm{NCNLN})-$ REAL (KIND=nag_wp) array
Input/Output
On entry: need not be set if the (default) optional parameter Cold Start is used.
If the optional parameter Warm Start has been chosen, CLAMDA $(j)$ must contain a multiplier estimate for each nonlinear constraint with a sign that matches the status of the constraint specified by the ISTATE array, for $j=\mathrm{N}+\mathrm{NCLIN}+1, \ldots, \mathrm{~N}+\mathrm{NCLIN}+\mathrm{NCNLN}$. The remaining elements need not be set. Note that if the $j$ th constraint is defined as 'inactive' by the initial value of ISTATE array (i.e., $\operatorname{ISTATE}(j)=0$ ), $\operatorname{CLAMDA}(j)$ should be zero; if the $j$ th constraint is an inequality active at its lower bound (i.e., $\operatorname{ISTATE}(j)=1), \operatorname{CLAMDA}(j)$ should be non-negative; if the $j$ th constraint is an inequality active at its upper bound (i.e., $\operatorname{ISTATE}(j)=2), \operatorname{CLAMDA}(j)$ should be non-positive. If necessary, the routine will modify CLAMDA to match these rules.

On exit: the values of the QP multipliers from the last QP subproblem. CLAMDA $(j)$ should be non-negative if $\operatorname{ISTATE}(j)=1$ and non-positive if $\operatorname{ISTATE}(j)=2$.

17: OBJF - REAL (KIND=nag_wp)
Output
On exit: the value of the objective function at the final iterate.
18: $\operatorname{OBJGRD}(\mathrm{N})$ - REAL (KIND=nag_wp) array
Output
On exit: the gradient of the objective function at the final iterate (or its finite difference approximation).

19: $\quad \mathrm{R}(\mathrm{LDR}, \mathrm{N})$ - REAL (KIND=nag_wp) array
Input/Output
On entry: need not be initialized if the (default) optional parameter Cold Start is used.
If the optional parameter Warm Start has been chosen, R must contain the upper triangular Cholesky factor $R$ of the initial approximation of the Hessian of the Lagrangian function, with the variables in the natural order. Elements not in the upper triangular part of R are assumed to be zero and need not be assigned.
On exit: if Hessian $=$ NO, R contains the upper triangular Cholesky factor $R$ of $Q^{\mathrm{T}} \tilde{H} Q$, an estimate of the transformed and reordered Hessian of the Lagrangian at $x$ (see (6) in Section 10.1). If Hessian = YES, R contains the upper triangular Cholesky factor $R$ of $H$, the approximate (untransformed) Hessian of the Lagrangian, with the variables in the natural order.

20: $\quad \mathrm{X}(\mathrm{N})$ - REAL (KIND=nag_wp) array
On entry: an initial estimate of the solution.
On exit: the final estimate of the solution.
21: IWORK(LIWORK) - INTEGER array Workspace
22: LIWORK - INTEGER Input
On entry: the dimension of the array IWORK as declared in the (sub)program from which E04UCF/ E04UCA is called.

Constraint: LIWORK $\geq 3 \times \mathrm{N}+$ NCLIN $+2 \times$ NCNLN.
WORK(LWORK) - REAL (KIND=nag_wp) array
Workspace
LWORK - INTEGER
Input
On entry: the dimension of the array WORK as declared in the (sub)program from which E04UCF/ E04UCA is called.

## Constraints:

$$
\begin{aligned}
& \text { if NCNLN }=0 \text { and NCLIN }=0, \text { LWORK } \geq 20 \times \mathrm{N} \text {; } \\
& \text { if } N C N L N ~=0 \text { and NCLIN }>0, \text { LWORK } \geq 2 \times \mathrm{N}^{2}+20 \times \mathrm{N}+11 \times \text { NCLIN; } \\
& \text { if NCNLN }>0 \text { and NCLIN } \geq 0, \text { LWORK } \geq 2 \times \mathrm{N}^{2}+\mathrm{N} \times \text { NCLIN }+2 \times \mathrm{N} \times \\
& \text { NCNLN }+20 \times \mathrm{N}+11 \times \mathrm{NCLIN}+21 \times N C N L N
\end{aligned}
$$

The amounts of workspace provided and required are (by default) output on the current advisory message unit (as defined by X04ABF). As an alternative to computing LIWORK and LWORK from the formulas given above, you may prefer to obtain appropriate values from the output of a preliminary run with LIWORK and LWORK set to 1. (E04UCF/E04UCA will then terminate with IFAIL $=9$.)

| 25: | $\operatorname{IUSER}(*)$ - INTEGER array | User Workspace |
| :--- | :--- | :--- |
| 26: | $\operatorname{RUSER}(*)$ - REAL (KIND=nag_wp) array | User Workspace |

26: $\operatorname{RUSER}(*)$ - REAL (KIND=nag_wp) array
User Workspace
IUSER and RUSER are not used by E04UCF/E04UCA, but are passed directly to CONFUN and OBJFUN and may be used to pass information to these routines as an alternative to using COMMON global variables.

27: IFAIL - INTEGER

## Input/Output

Note: for E04UCA, IFAIL does not occur in this position in the parameter list. See the additional parameters described below.

On entry: IFAIL must be set to $0,-1$ or 1 . If you are unfamiliar with this parameter you should refer to Section 3.3 in the Essential Introduction 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 parameters 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).

E04UCF/E04UCA returns with IFAIL $=0$ if the iterates have converged to a point $x$ that satisfies the first-order Kuhn-Tucker (see Section 10.1) conditions to the accuracy requested by the optional parameter Optimality Tolerance (default value $=\epsilon_{R}^{0.8}$, where $\epsilon_{r}$ is the value of the optional parameter Function Precision (default value $=\epsilon^{0.9}$, where $\epsilon$ is the machine precision)), i.e., the projected gradient and active constraint residuals are negligible at $x$.
You should check whether the following four conditions are satisfied:
(i) the final value of Norm Gz (see Section 8.1) is significantly less than that at the starting point;
(ii) during the final major iterations, the values of Step and Mnr (see Section 8.1) are both one;
(iii) the last few values of both Norm Gz and Violtn (see Section 8.1) become small at a fast linear rate; and
(iv) Cond Hz (see Section 8.1) is small.

If all these conditions hold, $x$ is almost certainly a local minimum of (1).
Note: the following are additional parameters for specific use with E04UCA. Users of E04UCF therefore need not read the remainder of this description.

```
LWSAV(120) - LOGICAL array
    Communication Array
IWSAV(610) - INTEGER array Communication Array
RWSAV(475) - REAL (KIND=nag_wp) array Communication Array
```

The arrays LWSAV, IWSAV and RWSAV must not be altered between calls to any of the routines E04UCA, E04UDA or E04UEA.

Note: see the parameter description for IFAIL above.

## 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: E04UCF/E04UCA may return useful information for one or more of the following detected errors or warnings.
Errors or warnings detected by the routine:
IFAIL $<0$
A negative value of IFAIL indicates an exit from E04UCF/E04UCA because you set MODE $<0$ in OBJFUN or CONFUN. The value of IFAIL will be the same as your setting of MODE.

IFAIL $=1$
The final iterate $x$ satisfies the first-order Kuhn-Tucker conditions (see Section 10.1) to the accuracy requested, but the sequence of iterates has not yet converged. E04UCF/E04UCA was terminated because no further improvement could be made in the merit function (see Section 8.1).
This value of IFAIL may occur in several circumstances. The most common situation is that you ask for a solution with accuracy that is not attainable with the given precision of the problem (as specified by the optional parameter Function Precision (default value $=\epsilon^{0.9}$, where $\epsilon$ is the machine precision)). This condition will also occur if, by chance, an iterate is an 'exact' KuhnTucker point, but the change in the variables was significant at the previous iteration. (This situation often happens when minimizing very simple functions, such as quadratics.)
If the four conditions listed in Section 5 for IFAIL $=0$ are satisfied, $x$ is likely to be a solution of (1) even if $\operatorname{IFAIL}=1$.

IFAIL $=2$
E04UCF/E04UCA has terminated without finding a feasible point for the linear constraints and bounds, which means that either no feasible point exists for the given value of the optional parameter Linear Feasibility Tolerance (default value $=\sqrt{\epsilon}$, where $\epsilon$ is the machine precision), or no feasible point could be found in the number of iterations specified by the optional parameter Minor Iteration Limit (default value $=\max \left(50,3\left(n+n_{L}+n_{N}\right)\right)$ ). You should check that there are no constraint redundancies. If the data for the constraints are accurate only to an absolute precision $\sigma$, you should ensure that the value of the optional parameter Linear Feasibility Tolerance is greater than $\sigma$. For example, if all elements of $A_{L}$ are of order unity and are accurate to only three decimal places, Linear Feasibility Tolerance should be at least $10^{-3}$.

IFAIL $=3$
No feasible point could be found for the nonlinear constraints. The problem may have no feasible solution. This means that there has been a sequence of QP subproblems for which no feasible point could be found (indicated by I at the end of each line of intermediate printout produced by the major iterations; see Section 8.1). This behaviour will occur if there is no feasible point for the nonlinear constraints. (However, there is no general test that can determine whether a feasible point exists for a set of nonlinear constraints.) If the infeasible subproblems occur from the very first major iteration, it is highly likely that no feasible point exists. If infeasibilities occur when earlier subproblems have been feasible, small constraint inconsistencies may be present. You should check the validity of constraints with negative values of ISTATE. If you are convinced that a feasible point does exist, E04UCF/E04UCA should be restarted at a different starting point.

IFAIL $=4$
The limiting number of iterations (as determined by the optional parameter Major Iteration Limit $\left(\right.$ default value $\left.\left.=\max \left(50,3\left(n+n_{L}\right)+10 n_{N}\right)\right)\right)$ has been reached.
If the algorithm appears to be making satisfactory progress, then Major Iteration Limit may be too small. If so, either increase its value and rerun E04UCF/E04UCA or, alternatively, rerun E04UCF/ E04UCA using the optional parameter Warm Start. If the algorithm seems to be making little or no progress however, then you should check for incorrect gradients or ill-conditioning as described under $\operatorname{IFAIL}=6$.

Note that ill-conditioning in the working set is sometimes resolved automatically by the algorithm, in which case performing additional iterations may be helpful. However, ill-conditioning in the Hessian approximation tends to persist once it has begun, so that allowing additional iterations without altering $R$ is usually inadvisable. If the quasi-Newton update of the Hessian approximation was reset during the latter major iterations (i.e., an $R$ occurs at the end of each line of intermediate printout; see Section 8.1), it may be worthwhile to try a Warm Start at the final point as suggested above.

IFAIL $=5$
Not used by this routine.
IFAIL $=6$
$x$ does not satisfy the first-order Kuhn-Tucker conditions (see Section 10.1) and no improved point for the merit function (see Section 8.1) could be found during the final linesearch.

This sometimes occurs because an overly stringent accuracy has been requested, i.e., the value of the optional parameter Optimality Tolerance (default value $=\epsilon_{R}^{0.8}$, where $\epsilon_{r}$ is the value of the optional parameter Function Precision (default value $=\epsilon^{0.9}$, where $\epsilon$ is the machine precision)) is too small. In this case you should apply the four tests outlined in the description of the parameter IFAIL to determine whether or not the final solution is acceptable (see Gill et al. (1981), for a discussion of the attainable accuracy).

If many iterations have occurred in which essentially no progress has been made and E04UCF/ E04UCA has failed completely to move from the initial point then user-supplied subroutines OBJFUN and/or CONFUN may be incorrect. You should refer to comments under IFAIL $=7$ and check the gradients using the optional parameter Verify (default value $=0$ ). Unfortunately, there may be small errors in the objective and constraint gradients that cannot be detected by the verification process. Finite difference approximations to first derivatives are catastrophically affected by even small inaccuracies. An indication of this situation is a dramatic alteration in the iterates if the finite difference interval is altered. One might also suspect this type of error if a switch is made to central differences even when Norm Gz and Violtn (see Section 8.1) are large.

Another possibility is that the search direction has become inaccurate because of ill-conditioning in the Hessian approximation or the matrix of constraints in the working set; either form of illconditioning tends to be reflected in large values of Mnr (the number of iterations required to solve each QP subproblem; see Section 8.1).
If the condition estimate of the projected Hessian (Cond Hz; see Section 8.1) is extremely large, it may be worthwhile rerunning E04UCF/E04UCA from the final point with the optional parameter Warm Start. In this situation, ISTATE and CLAMDA should be left unaltered and $R$ should be reset to the identity matrix.
If the matrix of constraints in the working set is ill-conditioned (i.e., Cond $T$ is extremely large; see Section 12), it may be helpful to run E04UCF/E04UCA with a relaxed value of the Feasibility Tolerance (default value $=\sqrt{\epsilon}$, where $\epsilon$ is the machine precision). (Constraint dependencies are often indicated by wide variations in size in the diagonal elements of the matrix $T$, whose diagonals will be printed if Major Print Level $\geq 30$ ).

## IFAIL $=7$

The user-supplied derivatives of the objective function and/or nonlinear constraints appear to be incorrect.

Large errors were found in the derivatives of the objective function and/or nonlinear constraints. This value of IFAIL will occur if the verification process indicated that at least one gradient or Jacobian element had no correct figures. You should refer to the printed output to determine which elements are suspected to be in error.

As a first-step, you should check that the code for the objective and constraint values is correct - for example, by computing the function at a point where the correct value is known. However, care should be taken that the chosen point fully tests the evaluation of the function. It is remarkable how often the values $x=0$ or $x=1$ are used to test function evaluation procedures, and how often the special properties of these numbers make the test meaningless.
Special care should be used in this test if computation of the objective function involves subsidiary data communicated in COMMON storage. Although the first evaluation of the function may be correct, subsequent calculations may be in error because some of the subsidiary data has accidentally been overwritten.

Gradient checking will be ineffective if the objective function uses information computed by the constraints, since they are not necessarily computed before each function evaluation.
Errors in programming the function may be quite subtle in that the function value is 'almost' correct. For example, the function may not be accurate to full precision because of the inaccurate calculation of a subsidiary quantity, or the limited accuracy of data upon which the function depends. A common error on machines where numerical calculations are usually performed in double precision is to include even one single precision constant in the calculation of the function; since some compilers do not convert such constants to double precision, half the correct figures may be lost by such a seemingly trivial error.

IFAIL $=8$
Not used by this routine.
IFAIL $=9$
An input parameter is invalid.

## Overflow

If the printed output before the overflow error contains a warning about serious ill-conditioning in the working set when adding the $j$ th constraint, it may be possible to avoid the difficulty by increasing the magnitude of the optional parameter Linear Feasibility Tolerance and/or the optional parameter Nonlinear Feasibility Tolerance and rerunning the program. If the message recurs even after this change then the offending linearly dependent constraint (with index ' $j$ ') must be removed from the problem. If overflow occurs in one of the user-supplied subroutines (e.g., if the nonlinear functions involve exponentials or singularities), it may help to specify tighter bounds for some of the variables (i.e., reduce the gap between the appropriate $l_{j}$ and $u_{j}$ ).

## 7 Accuracy

If IFAIL $=0$ on exit, then the vector returned in the array $X$ is an estimate of the solution to an accuracy of approximately Optimality Tolerance (default value $=\epsilon^{0.8}$, where $\epsilon$ is the machine precision).

## 8 Further Comments

### 8.1 Description of the Printed Output

This section describes the intermediate printout and final printout produced by E04UCF/E04UCA. The intermediate printout is a subset of the monitoring information produced by the routine at every iteration (see Section 12). You can control the level of printed output (see the description of the optional parameter

Major Print Level). Note that the intermediate printout and final printout are produced only if Major Print Level $\geq 10$ (the default for E04UCF, by default no output is produced by E04UCA).

The following line of summary output ( $<80$ characters) is produced at every major iteration. In all cases, the values of the quantities printed are those in effect on completion of the given iteration.
Maj is the major iteration count.
$\mathrm{Mnr} \quad$ is the number of minor iterations required by the feasibility and optimality phases of the QP subproblem. Generally, Mnr will be 1 in the later iterations, since theoretical analysis predicts that the correct active set will be identified near the solution (see Section 10).

Note that Mnr may be greater than the optional parameter Minor Iteration Limit if some iterations are required for the feasibility phase.

Step is the step $\alpha_{k}$ taken along the computed search direction. On reasonably wellbehaved problems, the unit step (i.e., $\alpha_{k}=1$ ) will be taken as the solution is approached.

Merit Function is the value of the augmented Lagrangian merit function (12) at the current iterate. This function will decrease at each iteration unless it was necessary to increase the penalty parameters (see Section 10.3). As the solution is approached, Merit Function will converge to the value of the objective function at the solution.

If the QP subproblem does not have a feasible point (signified by I at the end of the current output line) then the merit function is a large multiple of the constraint violations, weighted by the penalty parameters. During a sequence of major iterations with infeasible subproblems, the sequence of Merit Function values will decrease monotonically until either a feasible subproblem is obtained or E04UCF/ E04UCA terminates with IFAIL $=3$ (no feasible point could be found for the nonlinear constraints).

If there are no nonlinear constraints present (i.e., $\mathrm{NCNLN}=0$ ) then this entry contains Objective, the value of the objective function $F(x)$. The objective function will decrease monotonically to its optimal value when there are no nonlinear constraints.

Norm $\mathrm{Gz} \quad$ is $\left\|Z^{\mathrm{T}} g_{\mathrm{FR}}\right\|$, the Euclidean norm of the projected gradient (see Section 10.2). Norm Gz will be approximately zero in the neighbourhood of a solution.

Violtn is the Euclidean norm of the residuals of constraints that are violated or in the predicted active set (not printed if NCNLN is zero). Violtn will be approximately zero in the neighbourhood of a solution.
Cond $\mathrm{Hz} \quad$ is a lower bound on the condition number of the projected Hessian approximation $H_{Z}\left(H_{Z}=Z^{\mathrm{T}} H_{\mathrm{FR}} Z=R_{Z}^{\mathrm{T}} R_{Z}\right.$; see (6)). The larger this number, the more difficult the problem.

M

I
is printed if the quasi-Newton update has been modified to ensure that the Hessian approximation is positive definite (see Section 10.4).
is printed if the QP subproblem has no feasible point.
C is printed if central differences have been used to compute the unspecified objective and constraint gradients. If the value of Step is zero then the switch to central differences was made because no lower point could be found in the linesearch. (In this case, the QP subproblem is resolved with the central difference gradient and Jacobian.) If the value of Step is nonzero then central differences were computed because Norm Gz and Violtn imply that $x$ is close to a Kuhn-Tucker point (see Section 10.1 in E04UFF/E04UFA).
$\mathrm{L} \quad$ is printed if the linesearch has produced a relative change in $x$ greater than the value defined by the optional parameter Step Limit. If this output occurs frequently
during later iterations of the run, optional parameter Step Limit should be set to a larger value.

R is printed if the approximate Hessian has been refactorized. If the diagonal condition estimator of $R$ indicates that the approximate Hessian is badly conditioned then the approximate Hessian is refactorized using column interchanges. If necessary, $R$ is modified so that its diagonal condition estimator is bounded.

The final printout includes a listing of the status of every variable and constraint. The following describes the printout for each variable. A full stop (.) is printed for any numerical value that is zero.
$\begin{array}{ll}\text { Varbl } & \text { gives the name (V) and index } j \text {, for } j=1,2, \ldots, n \text {, of the variable. } \\ \text { State } & \text { gives the state of the variable (FR if neither bound is in the working set, EQ if a } \\ \text { fixed variable, LL if on its lower bound, UL if on its upper bound, TF if temporarily } \\ \text { fixed at its current value). If Value lies outside the upper or lower bounds by more } \\ \text { than the Feasibility Tolerance, State will be }++ \text { or -- respectively. (The latter } \\ \text { situation can occur only when there is no feasible point for the bounds and linear }\end{array}$ constraints.)

A key is sometimes printed before State.
A Alternative optimum possible. The variable is active at one of its bounds, but its Lagrange multiplier is essentially zero. This means that if the variable were allowed to start moving away from its bound then there would be no change to the objective function. The values of the other free variables might change, giving a genuine alternative solution. However, if there are any degenerate variables (labelled D), the actual change might prove to be zero, since one of them could encounter a bound immediately. In either case the values of the Lagrange multipliers might also change.

D Degenerate. The variable is free, but it is equal to (or very close to) one of its bounds.

I Infeasible. The variable is currently violating one of its bounds by more than the Feasibility Tolerance.
Value is the value of the variable at the final iteration.
Lower Bound is the lower bound specified for the variable. None indicates that $\mathrm{BL}(j) \leq-$ bigbnd.
Upper Bound is the upper bound specified for the variable. None indicates that $\mathrm{BU}(j) \geq$ bigbnd.

Lagr Mult

Slack is the difference between the variable Value and the nearer of its (finite) bounds $\operatorname{BL}(j)$ and $\mathrm{BU}(j)$. A blank entry indicates that the associated variable is not bounded (i.e., $\mathrm{BL}(j) \leq-b i g b n d$ and $\mathrm{BU}(j) \geq$ bigbnd $)$.

The meaning of the printout for linear and nonlinear constraints is the same as that given above for variables, with 'variable' replaced by 'constraint', $\mathrm{BL}(j)$ and $\mathrm{BU}(j)$ are replaced by $\mathrm{BL}(n+j)$ and $\mathrm{BU}(n+j)$ respectively, and with the following changes in the heading:
L Con gives the name (L) and index $j$, for $j=1,2, \ldots, n_{L}$, of the linear constraint.
N Con gives the name (N) and index $\left(j-n_{L}\right)$, for $j=n_{L}+1, \ldots, n_{L}+n_{N}$, of the nonlinear constraint.
Note that movement off a constraint (as opposed to a variable moving away from its bound) can be interpreted as allowing the entry in the Slack column to become positive.

Numerical values are output with a fixed number of digits; they are not guaranteed to be accurate to this precision.

## 9 Example

This is based on Problem 71 in Hock and Schittkowski (1981) and involves the minimization of the nonlinear function

$$
F(x)=x_{1} x_{4}\left(x_{1}+x_{2}+x_{3}\right)+x_{3}
$$

subject to the bounds

$$
\begin{aligned}
& 1 \leq x_{1} \leq 5 \\
& 1 \leq x_{2} \leq 5 \\
& 1 \leq x_{3} \leq 5 \\
& 1 \leq x_{4} \leq 5
\end{aligned}
$$

to the general linear constraint

$$
x_{1}+x_{2}+x_{3}+x_{4} \leq 20
$$

and to the nonlinear constraints

$$
\begin{aligned}
x_{1}^{2}+x_{2}^{2}+x_{3}^{2}+x_{4}^{2} & \leq 40, \\
x_{1} x_{2} x_{3} x_{4} & \geq 25
\end{aligned}
$$

The initial point, which is infeasible, is

$$
x_{0}=(1,5,5,1)^{\mathrm{T}},
$$

and $F\left(x_{0}\right)=16$.
The optimal solution (to five figures) is

$$
x^{*}=(1.0,4.7430,3.8211,1.3794)^{\mathrm{T}},
$$

and $F\left(x^{*}\right)=17.014$. One bound constraint and both nonlinear constraints are active at the solution.
The document for E04UDF/E04UDA includes an example program to solve the same problem using some of the optional parameters described in Section 11.

### 9.1 Program Text

the following program illustrates the use of E04UCF. An equivalent program illustrating the use of E04UCA is available with the supplied Library and is also available from the NAG web site.

```
EO4UCF Example Program Text
Mark 24 Release. NAG Copyright 2012.
Module eO4ucfe_mod
    EO4UCF Example Program Module:
                Parameters and User-defined Routines
    .. Use Statements ..
    Use nag_library, Only: nag_wp
    .. Implicit None Statement ..
    Implicit None
    .. Parameters ..
        Integer, Parameter :: nin = 5, nout = 6
    Contains
        Subroutine objfun(mode,n,x,objf,objgrd,nstate,iuser,ruser)
! Routine to evaluate objective function and its 1st derivatives.
! .. Scalar Arguments ..
            Real (Kind=nag_wp), Intent (Out) :: objf
            Integer, Intent (Inout) :: mode
            Integer, Intent (In) :: n, nstate
! .. Array Arguments ..
            Real (Kind=nag_wp), Intent (Inout) :: objgrd(n), ruser(*)
            Real (Kind=nag_wp), Intent (In) :: x(n)
            Integer, Intent (Inout) :: iuser(*)
! .. Executable Statements ..
```

```
    If (mode==0 .Or. mode==2) Then
        objf = x(1)*x(4)*(x(1)+x(2)+x(3)) + x(3)
    End If
    If (mode==1 .Or. mode==2) Then
        objgrd(1) = x(4)*(2.OEO_nag_wp*x(1)+x(2)+x(3))
        objgrd(2) = x(1)*x(4)
        objgrd(3) = x(1)*x(4) + 1.0EO_nag_wp
        objgrd(4) = x(1)*(x(1)+x(2)+x(3))
End If
Return
End Subroutine objfun
Subroutine confun(mode,ncnln,n,ldcj,needc,x,c,cjac,nstate,iuser,ruser)
    Routine to evaluate the nonlinear constraints and their 1st
    derivatives.
    .. Scalar Arguments ..
    Integer, Intent (In) :: ldcj, n, ncnln, nstate
    Integer, Intent (Inout) :: mode
    .. Array Arguments ..
    Real (Kind=nag_wp), Intent (Out) :: c(ncnln)
    Real (Kind=nag_wp), Intent (Inout) :: cjac(ldcj,n), ruser(*)
    Real (Kind=nag_wp), Intent (In) :: x(n)
    Integer, Intent (Inout) :: iuser(*)
    Integer, Intent (In) :: needc(ncnln)
    .. Executable Statements ..
    If (nstate==1) Then
        First call to CONFUN. Set all Jacobian elements to zero.
        Note that this will only work when 'Derivative Level = 3'
        (the default; see Section 11.2).
        cjac(1:ncnln,1:n) = 0.0EO_nag_wp
    End If
    If (needc(1)>0) Then
        If (mode==0 .Or. mode==2) Then
            c(1) = x(1)**2 + x(2)**2 + x (3)**2 + x (4)**2
        End If
        If (mode==1 .Or. mode==2) Then
            cjac(1,1) = 2.0EO_nag_wp*x(1)
            cjac(1,2) = 2.0EO_nag_wp*x(2)
            cjac(1,3) = 2.OEO_nag_wp*x(3)
            cjac(1,4) = 2.0EO_nag_wp*x(4)
        End If
    End If
    If (needc(2)>0) Then
        If (mode==0 .Or. mode==2) Then
            c(2) = x (1)*x(2)*x(3)*x(4)
        End If
        If (mode==1 .Or. mode==2) Then
            cjac(2,1)= x(2)*x(3)*x(4)
            cjac(2,2) = x(1)*x(3)*x(4)
            cjac(2,3) = x(1)*x(2)*x(4)
            cjac(2,4) = x(1)*x(2)*x(3)
        End If
    End If
    Return
    End Subroutine confun
End Module e04ucfe_mod
```

```
    Program e04ucfe
        EO4UCF Example Main Program
    .. Use Statements ..
    Use nag_library, Only: e04ucf, nag_wp
    Use eO4ucfe_mod, Only: confun, nin, nout, objfun
    .. Implicit None Statement ..
    Implicit None
    .. Local Scalars ..
    Real (Kind=nag_wp) :: objf
    Integer :: i, ifail, iter, lda, ldcj, ldr, &
        liwork, lwork, n, nclin, ncnln, &
        sda, sdcjac
    .. Local Arrays ..
    Real (Kind=nag_wp), Allocatable :: a(:,:), bl(:), bu(:), c(:), &
        cjac(:,:), clamda(:), objgrd(:), &
        r(:,:), work(:), x(:)
    Real (Kind=nag_wp)
    :: ruser(1)
    Integer, Allocatable
    :: istate(:), iwork(:)
    Integer :: iuser(1)
! .. Intrinsic Procedures ..
    Intrinsic :: max
    .. Executable Statements ..
    Write (nout,*) 'EO4UCF Example Program Results'
    Flush (nout)
    Skip heading in data file
    Read (nin,*)
    Read (nin,*) n, nclin, ncnln
    liwork = 3*n + nclin + 2*ncnln
    lda = max(1,nclin)
    If (nclin>0) Then
        sda = n
Else
    sda = 1
End If
ldcj = max(1,ncnln)
If (ncnln>0) Then
    sdcjac = n
Else
    sdcjac = 1
End If
ldr = n
If (ncnln==0 .And. nclin>0) Then
    lwork = 2*n**2 + 20*n + 11*nclin
Else If (ncnln>0 .And. nclin>=0) Then
    lwork = 2*n**2 + n*nclin + 2*n*ncnln + 20*n + 11*nclin + 21*ncnln
Else
    lwork = 20*n
End If
Allocate (istate(n+nclin+ncnln),iwork(liwork),a(lda,sda), &
    bl(n+nclin+ncnln),bu(n+nclin+ncnln),c(max(1, &
    ncnln)),cjac(ldcj,sdcjac),clamda(n+nclin+ncnln),objgrd(n),r(ldr,n), &
    x(n),work(lwork))
If (nclin>0) Then
    Read (nin,*)(a(i,1:sda),i=1,nclin)
End If
Read (nin,*) bl(1:(n+nclin+ncnln))
Read (nin,*) bu(1:(n+nclin+ncnln))
Read (nin,*) x(1:n)
```

```
    ifail = 0
    Call e04ucf(n,nclin,ncnln,lda,ldcj,ldr,a,bl,bu,confun,objfun,iter, &
        istate,c,cjac,clamda,objf,objgrd,r,x,iwork,liwork,work,lwork,iuser, &
        ruser,ifail)
End Program e04ucfe
```


### 9.2 Program Data

```
EO4UCF Example Program Data
\begin{tabular}{lrrrrlllll}
4 & 1 & 2 & & & & & & :Values of \(N\), NCLIN and NCNLN \\
1.0 & 1.0 & 1.0 & 1.0 & & & : End of matrix A \\
1.0 & 1.0 & 1.0 & 1.0 & \(-1.0 E+25\) & \(-1.0 E+25\) & 25.0 & : End of BL \\
5.0 & 5.0 & 5.0 & 5.0 & 20.0 & 40.0 & \(1.0 E+25\) & : End of BU \\
1.0 & 5.0 & 5.0 & 1.0 & & & & :End of X
\end{tabular}
```


### 9.3 Program Results

EO4UCF Example Program Results

```
*** E04UCF
```

Parameters


Verification of the objective gradients.
------------------------------------------

The objective gradients seem to be ok.

| Directional derivative of the objective | $8.15250000 \mathrm{E}-01$ |
| :--- | :--- |
| Difference approximation | $8.15249734 \mathrm{E}-01$ |


| Maj | Mnr | Step Merit Function | Norm Gz | Violtn Cond Hz |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0 | 4 | $0.0 \mathrm{E}+00$ | $1.738281 \mathrm{E}+01$ | $7.1 \mathrm{E}-01$ | $1.2 \mathrm{E}+01$ | $1.0 \mathrm{E}+00$ |
| 1 | 1 | $1.0 \mathrm{E}+00$ | $1.703169 \mathrm{E}+01$ | $4.6 \mathrm{E}-02$ | $1.9 \mathrm{E}+00$ | $1.0 \mathrm{E}+00$ |



Note: the remainder of this document is intended for more advanced users. Section 10 contains a detailed description of the algorithm which may be needed in order to understand Sections 11 and 12 . Section 11 describes the optional parameters which may be set by calls to E04UDF/E04UDA and/or E04UEF/ E04UEA. Section 12 describes the quantities which can be requested to monitor the course of the computation.

## 10 Algorithmic Details

This section contains a detailed description of the method used by E04UCF/E04UCA.

### 10.1 Overview

E04UCF/E04UCA is essentially identical to the subroutine NPSOL described in Gill et al. (1986b).
At a solution of (1), some of the constraints will be active, i.e., satisfied exactly. An active simple bound constraint implies that the corresponding variable is fixed at its bound, and hence the variables are partitioned into fixed and free variables. Let $C$ denote the $m$ by $n$ matrix of gradients of the active general linear and nonlinear constraints. The number of fixed variables will be denoted by $n_{\mathrm{FX}}$, with $n_{\mathrm{FR}}$ ( $n_{\mathrm{FR}}=n-n_{\mathrm{FX}}$ ) the number of free variables. The subscripts ' FX ' and ' FR ' on a vector or matrix will denote the vector or matrix composed of the elements corresponding to fixed or free variables.
A point $x$ is a first-order Kuhn-Tucker point for (1) (see Powell (1974)) if the following conditions hold:
(i) $x$ is feasible;
(ii) there exist vectors $\xi$ and $\lambda$ (the Lagrange multiplier vectors for the bound and general constraints) such that

$$
\begin{equation*}
g=C^{\mathrm{T}} \lambda+\xi \tag{2}
\end{equation*}
$$

where $g$ is the gradient of $F$ evaluated at $x$, and $\xi_{j}=0$ if the $j$ th variable is free.
(iii) The Lagrange multiplier corresponding to an inequality constraint active at its lower bound must be non-negative, and non-positive for an inequality constraint active at its upper bound.

Let $Z$ denote a matrix whose columns form a basis for the set of vectors orthogonal to the rows of $C_{\mathrm{FR}}$; i.e., $C_{\mathrm{FR}} Z=0$. An equivalent statement of the condition (2) in terms of $Z$ is

$$
Z^{\mathrm{T}} g_{\mathrm{FR}}=0
$$

The vector $Z^{\mathrm{T}} g_{\mathrm{FR}}$ is termed the projected gradient of $F$ at $x$. Certain additional conditions must be satisfied in order for a first-order Kuhn-Tucker point to be a solution of (1) (see Powell (1974)).
E04UCF/E04UCA implements a sequential quadratic programming (SQP) method. For an overview of SQP methods, see, for example, Fletcher (1987), Gill et al. (1981) and Powell (1983).

The basic structure of E04UCF/E04UCA involves major and minor iterations. The major iterations generate a sequence of iterates $\left\{x_{k}\right\}$ that converge to $x^{*}$, a first-order Kuhn-Tucker point of (1). At a typical major iteration, the new iterate $\bar{x}$ is defined by

$$
\begin{equation*}
\bar{x}=x+\alpha p \tag{3}
\end{equation*}
$$

where $x$ is the current iterate, the non-negative scalar $\alpha$ is the step length, and $p$ is the search direction. (For simplicity, we shall always consider a typical iteration and avoid reference to the index of the iteration.) Also associated with each major iteration are estimates of the Lagrange multipliers and a prediction of the active set.
The search direction $p$ in (3) is the solution of a quadratic programming subproblem of the form

$$
\underset{p}{\operatorname{minimize}} g^{\mathrm{T}} p+\frac{1}{2} p^{\mathrm{T}} H p \quad \text { subject to } \quad \bar{l} \leq\left\{\begin{array}{c}
p  \tag{4}\\
A_{L} p \\
A_{N} p
\end{array}\right\} \leq \bar{u}
$$

where $g$ is the gradient of $F$ at $x$, the matrix $H$ is a positive definite quasi-Newton approximation to the Hessian of the Lagrangian function (see Section 10.4), and $A_{N}$ is the Jacobian matrix of $c$ evaluated at $x$. (Finite difference estimates may be used for $g$ and $A_{N}$; see the optional parameter Derivative Level.) Let $l$ in (1) be partitioned into three sections: $l_{B}, l_{L}$ and $l_{N}$, corresponding to the bound, linear and nonlinear constraints. The vector $\bar{l}$ in (4) is similarly partitioned, and is defined as

$$
\bar{l}_{B}=l_{B}-x, \quad \bar{l}_{L}=l_{L}-A_{L} x, \quad \text { and } \quad \bar{l}_{N}=l_{N}-c,
$$

where $c$ is the vector of nonlinear constraints evaluated at $x$. The vector $\bar{u}$ is defined in an analogous fashion.

The estimated Lagrange multipliers at each major iteration are the Lagrange multipliers from the subproblem (4) (and similarly for the predicted active set). (The numbers of bounds, general linear and nonlinear constraints in the QP active set are the quantities Bnd, Lin and N1n in the monitoring file output of E04UCF/E04UCA; see Section 12.) In E04UCF/E04UCA, (4) is solved using E04NCF/E04NCA. Since solving a quadratic program is itself an iterative procedure, the minor iterations of E04UCF/E04UCA are the iterations of E04NCF/E04NCA. (More details about solving the subproblem are given in Section 10.2.)
Certain matrices associated with the QP subproblem are relevant in the major iterations. Let the subscripts 'FX' and 'FR' refer to the predicted fixed and free variables, and let $C$ denote the $m$ by $n$ matrix of gradients of the general linear and nonlinear constraints in the predicted active set. First, we have available the $T Q$ factorization of $C_{\mathrm{FR}}$ :

$$
C_{\mathrm{FR}} Q_{\mathrm{FR}}=\left(\begin{array}{ll}
0 & T \tag{5}
\end{array}\right)
$$

where $T$ is a nonsingular $m$ by $m$ reverse-triangular matrix (i.e., $t_{i j}=0$ if $i+j<m$ ), and the nonsingular $n_{\mathrm{FR}}$ by $n_{\mathrm{FR}}$ matrix $Q_{\mathrm{FR}}$ is the product of orthogonal transformations (see Gill et al. (1984b)). Second, we have the upper triangular Cholesky factor $R$ of the transformed and reordered Hessian matrix

$$
\begin{equation*}
R^{\mathrm{T}} R=H_{Q} \equiv Q^{\mathrm{T}} \tilde{H} Q \tag{6}
\end{equation*}
$$

where $\tilde{H}$ is the Hessian $H$ with rows and columns permuted so that the free variables are first, and $Q$ is the $n$ by $n$ matrix

$$
Q=\left(\begin{array}{cc}
Q_{\mathrm{FR}} &  \tag{7}\\
& I_{\mathrm{FX}}
\end{array}\right)
$$

with $I_{\mathrm{FX}}$ the identity matrix of order $n_{\mathrm{FX}}$. If the columns of $Q_{\mathrm{FR}}$ are partitioned so that

$$
Q_{\mathrm{FR}}=\left(\begin{array}{ll}
Z & Y
\end{array}\right),
$$

the $n_{Z}\left(n_{Z} \equiv n_{\mathrm{FR}}-m\right)$ columns of $Z$ form a basis for the null space of $C_{\mathrm{FR}}$. The matrix $Z$ is used to compute the projected gradient $Z^{\mathrm{T}} g_{\mathrm{FR}}$ at the current iterate. (The values Nz and Norm Gz printed by E04UCF/E04UCA give $n_{Z}$ and $\left\|Z^{\mathrm{T}} g_{\mathrm{FR}}\right\|$; see Section 12.)

A theoretical characteristic of SQP methods is that the predicted active set from the QP subproblem (4) is identical to the correct active set in a neighbourhood of $x^{*}$. In E04UCF/E04UCA, this feature is exploited by using the QP active set from the previous iteration as a prediction of the active set for the next QP subproblem, which leads in practice to optimality of the subproblems in only one iteration as the solution is approached. Separate treatment of bound and linear constraints in E04UCF/E04UCA also saves computation in factorizing $C_{\mathrm{FR}}$ and $H_{Q}$.

Once $p$ has been computed, the major iteration proceeds by determining a step length $\alpha$ that produces a 'sufficient decrease' in an augmented Lagrangian merit function (see Section 10.3). Finally, the approximation to the transformed Hessian matrix $H_{Q}$ is updated using a modified BFGS quasi-Newton update (see Section 10.4) to incorporate new curvature information obtained in the move from $x$ to $\bar{x}$.
On entry to E04UCF/E04UCA, an iterative procedure from E04NCF/E04NCA is executed, starting with the user-supplied initial point, to find a point that is feasible with respect to the bounds and linear constraints (using the tolerance specified by optional parameter Linear Feasibility Tolerance). If no feasible point exists for the bound and linear constraints, (1) has no solution and E04UCF/E04UCA terminates. Otherwise, the problem functions will thereafter be evaluated only at points that are feasible with respect to the bounds and linear constraints. The only exception involves variables whose bounds differ by an amount comparable to the finite difference interval (see the discussion of optional parameter Difference Interval). In contrast to the bounds and linear constraints, it must be emphasized that the nonlinear constraints will not generally be satisfied until an optimal point is reached.
Facilities are provided to check whether the user-supplied gradients appear to be correct (see the description of the optional parameter Verify). In general, the check is provided at the first point that is feasible with respect to the linear constraints and bounds. However, you may request that the check be performed at the initial point.
In summary, the method of E04UCF/E04UCA first determines a point that satisfies the bound and linear constraints. Thereafter, each iteration includes:
(a) the solution of a quadratic programming subproblem;
(b) a linesearch with an augmented Lagrangian merit function; and
(c) a quasi-Newton update of the approximate Hessian of the Lagrangian function.

These three procedures are described in more detail in Sections 10.2 to 10.4.

### 10.2 Solution of the Quadratic Programming Subproblem

The search direction $p$ is obtained by solving (4) using E04NCF/E04NCA (see Gill et al. (1986)), which was specifically designed to be used within an SQP algorithm for nonlinear programming.
E04NCF/E04NCA is based on a two-phase (primal) quadratic programming method. The two phases of the method are: finding an initial feasible point by minimizing the sum of infeasibilities (the feasibility phase), and minimizing the quadratic objective function within the feasible region (the optimality phase). The computations in both phases are performed by the same subroutines. The two-phase nature of the algorithm is reflected by changing the function being minimized from the sum of infeasibilities to the quadratic objective function.

In general, a quadratic program must be solved by iteration. Let $p$ denote the current estimate of the solution of (4); the new iterate $\bar{p}$ is defined by

$$
\begin{equation*}
\bar{p}=p+\sigma d \tag{8}
\end{equation*}
$$

where, as in (3), $\sigma$ is a non-negative step length and $d$ is a search direction.

At the beginning of each iteration of E04NCF/E04NCA, a working set is defined of constraints (general and bound) that are satisfied exactly. The vector $d$ is then constructed so that the values of constraints in the working set remain unaltered for any move along $d$. For a bound constraint in the working set, this property is achieved by setting the corresponding element of $d$ to zero, i.e., by fixing the variable at its bound. As before, the subscripts 'FX' and 'FR' denote selection of the elements associated with the fixed and free variables.

Let $C$ denote the sub-matrix of rows of

$$
\binom{A_{L}}{A_{N}}
$$

corresponding to general constraints in the working set. The general constraints in the working set will remain unaltered if

$$
\begin{equation*}
C_{\mathrm{FR}} d_{\mathrm{FR}}=0 \tag{9}
\end{equation*}
$$

which is equivalent to defining $d_{\mathrm{FR}}$ as

$$
\begin{equation*}
d_{\mathrm{FR}}=Z d_{Z} \tag{10}
\end{equation*}
$$

for some vector $d_{Z}$, where $Z$ is the matrix associated with the $T Q$ factorization (5) of $C_{\mathrm{FR}}$.
The definition of $d_{Z}$ in (10) depends on whether the current $p$ is feasible. If not, $d_{Z}$ is zero except for an element $\gamma$ in the $j$ th position, where $j$ and $\gamma$ are chosen so that the sum of infeasibilities is decreasing along $d$. (For further details, see Gill et al. (1986).) In the feasible case, $d_{Z}$ satisfies the equations

$$
\begin{equation*}
R_{Z}^{\mathrm{T}} R_{Z} d_{Z}=-Z^{\mathrm{T}} q_{\mathrm{FR}} \tag{11}
\end{equation*}
$$

where $R_{Z}$ is the Cholesky factor of $Z^{\mathrm{T}} H_{\mathrm{FR}} Z$ and $q$ is the gradient of the quadratic objective function $(q=g+H p)$. (The vector $Z^{\mathrm{T}} q_{\mathrm{FR}}$ is the projected gradient of the QP.) With (11), $p+d$ is the minimizer of the quadratic objective function subject to treating the constraints in the working set as equalities.
If the QP projected gradient is zero, the current point is a constrained stationary point in the subspace defined by the working set. During the feasibility phase, the projected gradient will usually be zero only at a vertex (although it may vanish at non-vertices in the presence of constraint dependencies). During the optimality phase, a zero projected gradient implies that $p$ minimizes the quadratic objective function when the constraints in the working set are treated as equalities. In either case, Lagrange multipliers are computed. Given a positive constant $\delta$ of the order of the machine precision, the Lagrange multiplier $\mu_{j}$ corresponding to an inequality constraint in the working set is said to be optimal if $\mu_{j} \leq \delta$ when the $j$ th constraint is at its upper bound, or if $\mu_{j} \geq-\delta$ when the associated constraint is at its lower bound. If any multiplier is nonoptimal, the current objective function (either the true objective or the sum of infeasibilities) can be reduced by deleting the corresponding constraint from the working set.

If optimal multipliers occur during the feasibility phase and the sum of infeasibilities is nonzero, no feasible point exists. The QP algorithm will then continue iterating to determine the minimum sum of infeasibilities. At this point, the Lagrange multiplier $\mu_{j}$ will satisfy $-(1+\delta) \leq \mu_{j} \leq \delta$ for an inequality constraint at its upper bound, and $-\delta \leq \mu_{j} \leq(1+\delta)$ for an inequality at its lower bound. The Lagrange multiplier for an equality constraint will satisfy $\left|\mu_{j}\right| \leq 1+\delta$.
The choice of step length $\sigma$ in the QP iteration (8) is based on remaining feasible with respect to the satisfied constraints. During the optimality phase, if $p+d$ is feasible, $\sigma$ will be taken as unity. (In this case, the projected gradient at $\bar{p}$ will be zero.) Otherwise, $\sigma$ is set to $\sigma_{\mathrm{M}}$, the step to the 'nearest' constraint, which is added to the working set at the next iteration.

Each change in the working set leads to a simple change to $C_{\mathrm{FR}}$ : if the status of a general constraint changes, a row of $C_{\mathrm{FR}}$ is altered; if a bound constraint enters or leaves the working set, a column of $C_{\mathrm{FR}}$ changes. Explicit representations are recurred of the matrices $T, Q_{\mathrm{FR}}$ and $R$, and of the vectors $Q^{\mathrm{T}} q$ and $Q^{\mathrm{T}} g$.

### 10.3 The Merit Function

After computing the search direction as described in Section 10.2, each major iteration proceeds by determining a step length $\alpha$ in (3) that produces a 'sufficient decrease' in the augmented Lagrangian merit
function

$$
\begin{equation*}
L(x, \lambda, s)=F(x)-\sum_{i} \lambda_{i}\left(c_{i}(x)-s_{i}\right)+\frac{1}{2} \sum_{i} \rho_{i}\left(c_{i}(x)-s_{i}\right)^{2} \tag{12}
\end{equation*}
$$

where $x, \lambda$ and $s$ vary during the linesearch. The summation terms in (12) involve only the nonlinear constraints. The vector $\lambda$ is an estimate of the Lagrange multipliers for the nonlinear constraints of (1). The non-negative slack variables $\left\{s_{i}\right\}$ allow nonlinear inequality constraints to be treated without introducing discontinuities. The solution of the QP subproblem (4) provides a vector triple that serves as a direction of search for the three sets of variables. The non-negative vector $\rho$ of penalty parameters is initialized to zero at the beginning of the first major iteration. Thereafter, selected elements are increased whenever necessary to ensure descent for the merit function. Thus, the sequence of norms of $\rho$ (the printed quantity Penalty; see Section 12) is generally nondecreasing, although each $\rho_{i}$ may be reduced a limited number of times.

The merit function (12) and its global convergence properties are described in Gill et al. (1986a).

### 10.4 The Quasi-Newton Update

The matrix $H$ in (4) is a positive definite quasi-Newton approximation to the Hessian of the Lagrangian function. (For a review of quasi-Newton methods, see Dennis and Schnabel (1983).) At the end of each major iteration, a new Hessian approximation $\bar{H}$ is defined as a rank-two modification of $H$. In E04UCF/ E04UCA, the BFGS (Broyden-Fletcher-Goldfarb-Shanno) quasi-Newton update is used:

$$
\begin{equation*}
\bar{H}=H-\frac{1}{s^{\mathrm{T}} H s} H s s^{\mathrm{T}} H+\frac{1}{y^{\mathrm{T}} s} y y^{\mathrm{T}}, \tag{13}
\end{equation*}
$$

where $s=\bar{x}-x$ (the change in $x$ ).
In E04UCF/E04UCA, $H$ is required to be positive definite. If $H$ is positive definite, $\bar{H}$ defined by (13) will be positive definite if and only if $y^{\mathrm{T}} s$ is positive (see Dennis and Moré (1977)). Ideally, $y$ in (13) would be taken as $y_{L}$, the change in gradient of the Lagrangian function

$$
\begin{equation*}
y_{L}=\bar{g}-\bar{A}_{N}^{\mathrm{T}} \mu_{N}-g+A_{N}^{\mathrm{T}} \mu_{N} \tag{14}
\end{equation*}
$$

where $\mu_{N}$ denotes the QP multipliers associated with the nonlinear constraints of the original problem. If $y_{L}^{\mathrm{T}} s$ is not sufficiently positive, an attempt is made to perform the update with a vector $y$ of the form

$$
y=y_{L}+\sum_{i=1}^{m_{N}} \omega_{i}\left(a_{i}(\hat{x}) c_{i}(\hat{x})-a_{i}(x) c_{i}(x)\right),
$$

where $\omega_{i} \geq 0$. If no such vector can be found, the update is performed with a scaled $y_{L}$; in this case, M is printed to indicate that the update was modified.
Rather than modifying $H$ itself, the Cholesky factor of the transformed Hessian $H_{Q}$ (6) is updated, where $Q$ is the matrix from (5) associated with the active set of the QP subproblem. The update (13) is equivalent to the following update to $H_{Q}$ :

$$
\begin{equation*}
\bar{H}_{Q}=H_{Q}-\frac{1}{s_{Q}^{\mathrm{T}} H_{Q} s_{Q}} H_{Q} s_{Q} s_{Q}^{\mathrm{T}} H_{Q}+\frac{1}{y_{Q}^{\mathrm{T}} s_{Q}} y_{Q} y_{Q}^{\mathrm{T}} \tag{15}
\end{equation*}
$$

where $y_{Q}=Q^{\mathrm{T}} y$, and $s_{Q}=Q^{\mathrm{T}} s$. This update may be expressed as a rank-one update to $R$ (see Dennis and Schnabel (1981)).

## 11 Optional Parameters

Several optional parameters in E04UCF/E04UCA define choices in the problem specification or the algorithm logic. In order to reduce the number of formal parameters of E04UCF/E04UCA 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 following is a list of the optional parameters available. A full description of each optional parameter is provided in Section 11.1.

Central Difference Interval
Cold Start
Crash Tolerance
Defaults
Derivative Level
Difference Interval
Feasibility Tolerance
Function Precision
Hessian
Infinite Bound Size
Infinite Step Size
Iteration Limit
Iters
Itns
Linear Feasibility Tolerance
Line Search Tolerance
List
Major Iteration Limit
Major Print Level
Minor Iteration Limit
Minor Print Level
Monitoring File
Nolist
Nonlinear Feasibility Tolerance
Optimality Tolerance
Print Level
Start Constraint Check At Variable
Start Objective Check At Variable
Step Limit
Stop Constraint Check At Variable
Stop Objective Check At Variable
Verify
Verify Constraint Gradients
Verify Gradients
Verify Level
Verify Objective Gradients
Warm Start
Optional parameters may be specified by calling one, or both, of E04UDF/E04UDA and E04UEF/E04UEA before a call to E04UCF/E04UCA.

E04UDF/E04UDA reads options from an external options file, with Begin and End as the first and last lines respectively and each intermediate line defining a single optional parameter. For example,

```
Begin
    Print level = 1
End
```

The call
CALL EO4UDF (IOPTNS, INFORM)
can then be used to read the file on unit IOPTNS. INFORM will be zero on successful exit. E04UDF/ E04UDA should be consulted for a full description of this method of supplying optional parameters.

E04UEF/E04UEA can be called to supply options directly, one call being necessary for each optional parameter. For example,

```
CALL EO4UEF ('Print Level = 1')
```

E04UEF/E04UEA should be consulted for a full description of this method of supplying optional parameters.
All optional parameters not specified by you are set to their default values. Optional parameters specified by you are unaltered by E04UCF/E04UCA (unless they define invalid values) and so remain in effect for subsequent calls to E04UCF/E04UCA, unless altered by you.

### 11.1 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 (if no characters of an optional qualifier are underlined, the qualifier may be omitted);
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), and $\epsilon_{r}$ denotes the relative precision of the objective function Function Precision, and bigbnd signifies the value of Infinite Bound Size.

Keywords and character values are case and white space insensitive.

## Central Difference Interval

If the algorithm switches to central differences because the forward-difference approximation is not sufficiently accurate, the value of $r$ is used as the difference interval for every element of $x$. The switch to central differences is indicated by $C$ at the end of each line of intermediate printout produced by the major iterations (see Section 8.1). The use of finite differences is discussed further under the optional parameter Difference Interval.

If you supply a value for this optional parameter, a small value between 0.0 and 1.0 is appropriate.

## Cold Start <br> Default Warm Start

This option controls the specification of the initial working set in both the procedure for finding a feasible point for the linear constraints and bounds and in the first QP subproblem thereafter. With a Cold Start, the first working set is chosen by E04UCF/E04UCA based on the values of the variables and constraints at the initial point. Broadly speaking, the initial working set will include equality constraints and bounds or inequality constraints that violate or 'nearly' satisfy their bounds (to within Crash Tolerance).
With a Warm Start, you must set the ISTATE array and define CLAMDA and R as discussed in Section 5. ISTATE values associated with bounds and linear constraints determine the initial working set of the procedure to find a feasible point with respect to the bounds and linear constraints. ISTATE values associated with nonlinear constraints determine the initial working set of the first QP subproblem after such a feasible point has been found. E04UCF/E04UCA will override your specification of ISTATE if necessary, so that a poor choice of the working set will not cause a fatal error. For instance, any elements of ISTATE which are set to $-2,-1$ or 4 will be reset to zero, as will any elements which are set to 3 when the corresponding elements of BL and BU are not equal. A warm start will be advantageous if a good estimate of the initial working set is available - for example, when E04UCF/E04UCA is called repeatedly to solve related problems.

This value is used in conjunction with the optional parameter Cold Start (the default value) when E04UCF/E04UCA selects an initial working set. If $0 \leq r \leq 1$, the initial working set will include (if possible) bounds or general inequality constraints that lie within $r$ of their bounds. In particular, a constraint of the form $a_{j}^{\mathrm{T}} x \geq l$ will be included in the initial working set if $\left|a_{j}^{\mathrm{T}} x-l\right| \leq r(1+|l|)$. If $r<0$ or $r>1$, the default value is used.

## Defaults

This special keyword may be used to reset all optional parameters to their default values.

## Derivative Level

$i$
Default $=3$
This parameter indicates which derivatives are provided in user-supplied subroutines OBJFUN and CONFUN. The possible choices for $i$ are the following.
$i$
Meaning
3 All elements of the objective gradient and the constraint Jacobian are provided.
2 All elements of the constraint Jacobian are provided, but some elements of the objective gradient are not specified.
1 All elements of the objective gradient are provided, but some elements of the constraint Jacobian are not specified.
0 Some elements of both the objective gradient and the constraint Jacobian are not specified.
The value $i=3$ should be used whenever possible, since E04UCF/E04UCA is more reliable (and will usually be more efficient) when all derivatives are exact.

If $i=0$ or 2 , E04UCF/E04UCA will estimate the unspecified elements of the objective gradient, using finite differences. The computation of finite difference approximations usually increases the total run-time, since a call to OBJFUN is required for each unspecified element. Furthermore, less accuracy can be attained in the solution (see Chapter 8 of Gill et al. (1981), for a discussion of limiting accuracy).

If $i=0$ or 1 , E04UCF/E04UCA will approximate unspecified elements of the constraint Jacobian. One call to CONFUN is needed for each variable for which partial derivatives are not available. For example, if the Jacobian has the form

$$
\left(\begin{array}{llll}
* & * & * & * \\
* & ? & ? & * \\
* & * & ? & * \\
* & * & * & *
\end{array}\right)
$$

where ' $*$ ' indicates an element provided by you and '?' indicates an unspecified element, E04UCF/ E04UCA will call CONFUN twice: once to estimate the missing element in column 2, and again to estimate the two missing elements in column 3. (Since columns 1 and 4 are known, they require no calls to CONFUN.)
At times, central differences are used rather than forward differences, in which case twice as many calls to OBJFUN and CONFUN are needed. (The switch to central differences is not under your control.)
If $i<0$ or $i>3$, the default value is used.
Difference Interval
Default values are computed
This option defines an interval used to estimate derivatives by finite differences in the following circumstances:
(a) For verifying the objective and/or constraint gradients (see the description of the optional parameter Verify).
(b) For estimating unspecified elements of the objective gradient or the constraint Jacobian.

In general, a derivative with respect to the $j$ th variable is approximated using the interval $\delta_{j}$, where $\delta_{j}=r\left(1+\left|\hat{x}_{j}\right|\right)$, with $\hat{x}$ the first point feasible with respect to the bounds and linear constraints. If the functions are well scaled, the resulting derivative approximation should be accurate to $O(r)$. See Gill et al. (1981) for a discussion of the accuracy in finite difference approximations.

If a difference interval is not specified by you, a finite difference interval will be computed automatically for each variable by a procedure that requires up to six calls of CONFUN and OBJFUN for each element. This option is recommended if the function is badly scaled or you wish to have E04UCF/E04UCA determine constant elements in the objective and constraint gradients (see the descriptions of CONFUN and OBJFUN in Section 5).

If you supply a value for this optional parameter, a small value between 0.0 and 1.0 is appropriate.

## Feasibility Tolerance <br> $r$ <br> $$
\text { Default }=\sqrt{\epsilon}
$$

The scalar $r$ defines the maximum acceptable absolute violations in linear and nonlinear constraints at a 'feasible' point; i.e., a constraint is considered satisfied if its violation does not exceed $r$. If $r<\epsilon$ or $r \geq 1$, the default value is used. Using this keyword sets both optional parameters Linear Feasibility Tolerance and Nonlinear Feasibility Tolerance to $r$, if $\epsilon \leq r<1$. (Additional details are given under the descriptions of these optional parameters.)

## Function Precision

$r$
Default $=\epsilon^{0.9}$
This parameter defines $\epsilon_{r}$, which is intended to be a measure of the accuracy with which the problem functions $F(x)$ and $c(x)$ can be computed. If $r<\epsilon$ or $r \geq 1$, the default value is used.
The value of $\epsilon_{r}$ should reflect the relative precision of $1+|F(x)|$; i.e., $\epsilon_{r}$ acts as a relative precision when $|F|$ is large and as an absolute precision when $|F|$ is small. For example, if $F(x)$ is typically of order 1000 and the first six significant digits are known to be correct, an appropriate value for $\epsilon_{r}$ would be $10^{-6}$. In contrast, if $F(x)$ is typically of order $10^{-4}$ and the first six significant digits are known to be correct, an appropriate value for $\epsilon_{r}$ would be $10^{-10}$. The choice of $\epsilon_{r}$ can be quite complicated for badly scaled problems; see Chapter 8 of Gill et al. (1981) for a discussion of scaling techniques. The default value is appropriate for most simple functions that are computed with full accuracy. However, when the accuracy of the computed function values is known to be significantly worse than full precision, the value of $\epsilon_{r}$ should be large enough so that E04UCF/E04UCA will not attempt to distinguish between function values that differ by less than the error inherent in the calculation.

## Hessian <br> No <br> Default $=$ NO

This option controls the contents of the upper triangular matrix $R$ (see Section 5). E04UCF/E04UCA works exclusively with the transformed and reordered Hessian $H_{Q}$ (6), and hence extra computation is required to form the Hessian itself. If Hessian $=$ NO, R contains the Cholesky factor of the transformed and reordered Hessian. If Hessian = YES, the Cholesky factor of the approximate Hessian itself is formed and stored in R. You should select Hessian = YES if a Warm Start will be used for the next call to E04UCF/E04UCA.

## Infinite Bound Size

$r$
Default $=10^{20}$
If $r>0, r$ 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 $-b i g b n d$ will be regarded as $-\infty$ ). If $r<0$, the default value is used.

## Infinite $\underline{\text { Step Size }}$

$r$
Default $=\max \left(\right.$ bigbnd, $\left.10^{20}\right)$
If $r>0, r$ specifies the magnitude of the change in variables that is treated as a step to an unbounded solution. If the change in $x$ during an iteration would exceed the value of $r$, the objective function is considered to be unbounded below in the feasible region. If $r \leq 0$, the default value is used.

## Line Search Tolerance

$r$

$$
\text { Default }=0.9
$$

The value $r(0 \leq r<1)$ controls the accuracy with which the step $\alpha$ taken during each iteration approximates a minimum of the merit function along the search direction (the smaller the value of $r$, the
more accurate the linesearch). The default value $r=0.9$ requests an inaccurate search and is appropriate for most problems, particularly those with any nonlinear constraints.

If there are no nonlinear constraints, a more accurate search may be appropriate when it is desirable to reduce the number of major iterations - for example, if the objective function is cheap to evaluate, or if a substantial number of derivatives are unspecified. If $r<0$ or $r \geq 1$, the default value is used.

| Linear Feasibility Tolerance | $r_{1}$ | Default $=\sqrt{\epsilon}$ |
| :--- | :--- | :--- |
| Nonlinear Feasibility Tolerance | $r_{2}$ | Default $=\epsilon^{0.33}$ or $\sqrt{\epsilon}$ |

The default value of $r_{2}$ is $\epsilon^{0.33}$ if Derivative Level $=0$ or 1 , and $\sqrt{\epsilon}$ otherwise.
The scalars $r_{1}$ and $r_{2}$ define the maximum acceptable absolute violations in linear and nonlinear constraints at a 'feasible' point; i.e., a linear constraint is considered satisfied if its violation does not exceed $r_{1}$, and similarly for a nonlinear constraint and $r_{2}$. If $r_{m}<\epsilon$ or $r_{m} \geq 1$, the default value is used, for $m=1,2$.

On entry to E04UCF/E04UCA, an iterative procedure is executed in order to find a point that satisfies the linear constraints and bounds on the variables to within the tolerance $r_{1}$. All subsequent iterates will satisfy the linear constraints to within the same tolerance (unless $r_{1}$ is comparable to the finite difference interval).

For nonlinear constraints, the feasibility tolerance $r_{2}$ defines the largest constraint violation that is acceptable at an optimal point. Since nonlinear constraints are generally not satisfied until the final iterate, the value of optional parameter Nonlinear Feasibility Tolerance acts as a partial termination criterion for the iterative sequence generated by E04UCF/E04UCA (see the discussion of optional parameter Optimality Tolerance).
These tolerances should reflect the precision of the corresponding constraints. For example, if the variables and the coefficients in the linear constraints are of order unity, and the latter are correct to about 6 decimal digits, it would be appropriate to specify $r_{1}$ as $10^{-6}$.

## List <br> Default for E04UCF <br> Nolist <br> Default for E04UCA

For E04UCF, normally each optional parameter specification is printed as it is supplied. Optional parameter Nolist may be used to suppress the printing and optional parameter List may be used to turn on printing.

## Major Iteration Limit <br> Iteration Limit <br> Iters <br> Itns

The value of $i$ specifies the maximum number of major iterations allowed before termination. Setting $i=0$ and Major Print Level $>0$ means that the workspace needed will be computed and printed, but no iterations will be performed. If $i<0$, the default value is used.

## Major Print Level <br> i <br> Default for $\mathrm{E} 04 \mathrm{UCF}=10$ Print Level i <br> Default for E04UCA $=0$

The value of $i$ controls the amount of printout produced by the major iterations of E04UCF/E04UCA, as indicated below. A detailed description of the printed output is given in Section 8.1 (summary output at each major iteration and the final solution) and Section 12 (monitoring information at each major iteration). (See also the description of the optional parameter Minor Print Level.)
The following printout is sent to the current advisory message unit (as defined by X04ABF):
$i$

## Output

0 No output.
For the other values described below, the parameters used by E04UCF/ E04UCA are displayed in addition to intermediate and final output.

## Output

1 The final solution only.
5 One line of summary output ( $<80$ characters; see Section 8.1) for each major iteration (no printout of the final solution).
$\geq 10$ The final solution and one line of summary output for each major iteration.
The following printout is sent to the logical unit number defined by the optional parameter Monitoring File:

## Output

$<5$ No output.
$\geq 5$ One long line of output ( $>80$ characters; see Section 12) for each major iteration (no printout of the final solution).
$\geq 20$ At each major iteration, the objective function, the Euclidean norm of the nonlinear constraint violations, the values of the nonlinear constraints (the vector $c$ ), the values of the linear constraints (the vector $A_{L} x$ ), and the current values of the variables (the vector $x$ ).
$\geq 30$ At each major iteration, the diagonal elements of the matrix $T$ associated with the $T Q$ factorization (5) (see Section 10.1) of the QP working set, and the diagonal elements of $R$, the triangular factor of the transformed and reordered Hessian (6) (see Section 10.1).
If Major Print Level $\geq 5$ and the unit number defined by the optional parameter Monitoring File is the same as that defined by X 04 ABF , then the summary output for each major iteration is suppressed.

## Minor Iteration Limit

$i$

$$
\text { Default }=\max \left(50,3\left(n+n_{L}+n_{N}\right)\right)
$$

The value of $i$ specifies the maximum number of iterations for finding a feasible point with respect to the bounds and linear constraints (if any). The value of $i$ also specifies the maximum number of minor iterations for the optimality phase of each QP subproblem. If $i \leq 0$, the default value is used.

## Minor Print Level

$i$
Default $=0$
The value of $i$ controls the amount of printout produced by the minor iterations of E04UCF/E04UCA (i.e., the iterations of the quadratic programming algorithm), as indicated below. A detailed description of the printed output is given in Section 8.1 (summary output at each minor iteration and the final QP solution) and Section 12 in E04NCF/E04NCA (monitoring information at each minor iteration). (See also the description of the optional parameter Major Print Level.) The following printout is sent to the current advisory message unit (as defined by X04ABF):

## Output

0 No output.
1 The final QP solution only.
5 One line of summary output ( $<80$ characters; see Section 8.2 in E04NCF/E04NCA) for each minor iteration (no printout of the final QP solution).
$\geq 10$ The final QP solution and one line of summary output for each minor iteration.
The following printout is sent to the logical unit number defined by the optional parameter Monitoring File:

## Output

## $<5$ No output.

$\geq 5$ One long line of output ( $>80$ characters; see Section 8.2 in E04NCF/E04NCA) for each minor iteration (no printout of the final QP solution).
$\geq 20$ At each minor iteration, the current estimates of the QP multipliers, the current estimate of the QP search direction, the QP constraint values, and the status of each QP constraint.
$\geq 30$ At each minor iteration, the diagonal elements of the matrix $T$ associated with the $T Q$ factorization (5) (see Section 10.1) of the QP working set, and the diagonal elements of the Cholesky factor $R$ of the transformed Hessian (6) (see Section 10.1).
If Minor Print Level $\geq 5$ and the unit number defined by the optional parameter Monitoring File is the same as that defined by X 04 ABF , then the summary output for each minor iteration is suppressed.

## Monitoring File <br> $i$ <br> Default $=-1$

If $i \geq 0$ and Major Print Level $\geq 5$ or $i \geq 0$ and Minor Print Level $\geq 5$, monitoring information produced by E04UCF/E04UCA at every iteration is sent to a file with logical unit number $i$. If $i<0$ and/ or Major Print Level $<5$ and Minor Print Level $<5$, no monitoring information is produced.

## Optimality Tolerance <br> $r$ <br> $$
\text { Default }=\epsilon_{R}^{0.8}
$$

The parameter $r\left(\epsilon_{r} \leq r<1\right)$ specifies the accuracy to which you wish the final iterate to approximate a solution of the problem. Broadly speaking, $r$ indicates the number of correct figures desired in the objective function at the solution. For example, if $r$ is $10^{-6}$ and E04UCF/E04UCA terminates successfully, the final value of $F$ should have approximately six correct figures. If $r<\epsilon_{r}$ or $r \geq 1$, the default value is used.
E04UCF/E04UCA will terminate successfully if the iterative sequence of $x$ values is judged to have converged and the final point satisfies the first-order Kuhn-Tucker conditions (see Section 10.1). The sequence of iterates is considered to have converged at $x$ if

$$
\begin{equation*}
\alpha\|p\| \leq \sqrt{r}(1+\|x\|) \tag{16}
\end{equation*}
$$

where $p$ is the search direction and $\alpha$ the step length from (3). An iterate is considered to satisfy the firstorder conditions for a minimum if

$$
\begin{equation*}
\left\|Z^{\mathrm{T}} g_{\mathrm{FR}}\right\| \leq \sqrt{r}\left(1+\max \left(1+|F(x)|,\left\|g_{\mathrm{FR}}\right\|\right)\right) \tag{17}
\end{equation*}
$$

and

$$
\begin{equation*}
\left|r e s_{j}\right| \leq \text { ftol } \quad \text { for all } \quad j \tag{18}
\end{equation*}
$$

where $Z^{\mathrm{T}} g_{\mathrm{FR}}$ is the projected gradient (see Section 10.1 ), $g_{\mathrm{FR}}$ is the gradient of $F(x)$ with respect to the free variables, res $_{j}$ is the violation of the $j$ th active nonlinear constraint, and ftol is the Nonlinear Feasibility Tolerance.

| Start Objective Check At Variable | $i_{1}$ | Default $=1$ |
| :--- | :--- | :--- |
| Stop Objective Check At Variable | $i_{2}$ | Default $=n$ |
| Start Constraint Check At Variable | $i_{3}$ | Default $=1$ |
| Stop Constraint Check At Variable | $i_{4}$ | Default $=n$ |

These keywords take effect only if Verify Level $>0$. They may be used to control the verification of gradient elements computed by OBJFUN and/or Jacobian elements computed by CONFUN. For example, if the first 30 elements of the objective gradient appeared to be correct in an earlier run, so that only element 31 remains questionable, it is reasonable to specify Start Objective Check At Variable $=31$. If the first 30 variables appear linearly in the objective, so that the corresponding gradient elements are constant, the above choice would also be appropriate.
If $i_{2 m-1} \leq 0$ or $i_{2 m-1}>\min \left(n, i_{2 m}\right)$, the default value is used, for $m=1,2$. If $i_{2 m} \leq 0$ or $i_{2 m}>n$, the default value is used, for $m=1,2$.

## Step Limit

$$
r
$$

$$
\text { Default }=2.0
$$

If $r>0, r$ specifies the maximum change in variables at the first step of the linesearch. In some cases, such as $F(x)=a e^{b x}$ or $F(x)=a x^{b}$, even a moderate change in the elements of $x$ can lead to floating point overflow. The parameter $r$ is therefore used to encourage evaluation of the problem functions at meaningful points. Given any major iterate $x$, the first point $\tilde{x}$ at which $F$ and $c$ are evaluated during the linesearch is restricted so that

$$
\|\tilde{x}-x\|_{2} \leq r\left(1+\|x\|_{2}\right)
$$

The linesearch may go on and evaluate $F$ and $c$ at points further from $x$ if this will result in a lower value of the merit function (indicated by L at the end of each line of output produced by the major iterations; see Section 8.1). If $L$ is printed for most of the iterations, $r$ should be set to a larger value.

Wherever possible, upper and lower bounds on $x$ should be used to prevent evaluation of nonlinear functions at wild values. The default value Step Limit $=2.0$ should not affect progress on well-behaved functions, but values such as 0.1 or 0.01 may be helpful when rapidly varying functions are present. If a small value of Step Limit is selected, a good starting point may be required. An important application is to the class of nonlinear least squares problems. If $r \leq 0$, the default value is used.

| Verify Level | $i$ | Default $=0$ |
| :--- | :--- | :--- |
| Verify | $i$ |  |
| Verify Constraint Gradients | $i$ |  |
| Verify Gradients | $i$ |  |
| Verify |  |  |

These keywords refer to finite difference checks on the gradient elements computed by OBJFUN and CONFUN. The possible choices for $i$ are as follows:

Meaning
-1 No checks are performed.
0 Only a 'cheap' test will be performed.
$\geq 1$ Individual gradient elements will also be checked using a reliable (but more expensive) test.
It is possible to specify Verify Level $=0$ to 3 in several ways. For example, the nonlinear objective gradient (if any) will be verified if either Verify Objective Gradients or Verify Level $=1$ is specified. The constraint gradients will be verified if Verify = YES or Verify Level $=2$ or Verify is specified. Similarly, the objective and the constraint gradients will be verified if Verify $=$ YES or Verify Level $=3$ or Verify is specified.

If $0 \leq i \leq 3$, gradients will be verified at the first point that satisfies the linear constraints and bounds.
If $i=0$, only a 'cheap' test will be performed, requiring one call to OBJFUN and (if appropriate) one call to CONFUN.
If $1 \leq i \leq 3$, a more reliable (but more expensive) check will be made on individual gradient elements, within the ranges specified by the Start Constraint Check At Variable and Stop Constraint Check At
Variable keywords. A result of the form OK or BAD? is printed by E04UCF/E04UCA to indicate whether or not each element appears to be correct.
If $10 \leq i \leq 13$, the action is the same as for $i-10$, except that it will take place at the user-specified initial value of $x$.

If $i<-1$ or $4 \leq i \leq 9$ or $i>13$, the default value is used.
We suggest that Verify Level $=3$ be used whenever a new function routine is being developed.

## 12 Description of Monitoring Information

This section describes the long line of output ( $>80$ characters) which forms part of the monitoring information produced by E04UCF/E04UCA. (See also the description of the optional parameters Major Print Level, Minor Print Level and Monitoring File.) You can control the level of printed output.

When Major Print Level $\geq 5$ and Monitoring File $\geq 0$, the following line of output is produced at every major iteration of E04UCF/E04UCA on the unit number specified by Monitoring File. In all cases, the values of the quantities printed are those in effect on completion of the given iteration.
Maj is the major iteration count.
$\mathrm{Mnr} \quad$ is the number of minor iterations required by the feasibility and optimality phases of the QP subproblem. Generally, Mnr will be 1 in the later iterations, since theoretical
analysis predicts that the correct active set will be identified near the solution (see Section 10).

Note that Mnr may be greater than the optional parameter Minor Iteration Limit if some iterations are required for the feasibility phase.

| Step | is the step $\alpha_{k}$ taken along the computed search direction. On reasonably well- <br> behaved problems, the unit step (i.e., $\alpha_{k}=1$ ) will be taken as the solution is <br> approached. |
| :--- | :--- |
| Nfun | is the cumulative number of evaluations of the objective function needed for the |
| linesearch. Evaluations needed for the estimation of the gradients by finite |  |
| differences are not included. Nfun is printed as a guide to the amount of work |  |
| required for the linesearch. |  |
| Merit Function | is the value of the augmented Lagrangian merit function (12) at the current iterate. <br> This function will decrease at each iteration unless it was necessary to increase the <br> penalty parameters (see Section 10.3). As the solution is approached, Merit <br> Function will converge to the value of the objective function at the solution. |

If the QP subproblem does not have a feasible point (signified by I at the end of the current output line) then the merit function is a large multiple of the constraint violations, weighted by the penalty parameters. During a sequence of major iterations with infeasible subproblems, the sequence of Merit Function values will decrease monotonically until either a feasible subproblem is obtained or E04UCF/ E04UCA terminates with IFAIL $=3$ (no feasible point could be found for the nonlinear constraints).
If there are no nonlinear constraints present (i.e., $\mathrm{NCNLN}=0$ ) then this entry contains Objective, the value of the objective function $F(x)$. The objective function will decrease monotonically to its optimal value when there are no nonlinear constraints.

Norm Gz is $\left\|Z^{\mathrm{T}} g_{\mathrm{FR}}\right\|$, the Euclidean norm of the projected gradient (see Section 10.2). Norm Gz will be approximately zero in the neighbourhood of a solution.
Violtn is the Euclidean norm of the residuals of constraints that are violated or in the predicted active set (not printed if NCNLN is zero). Violtn will be approximately zero in the neighbourhood of a solution.

Nz
is the number of columns of $Z$ (see Section 10.2). The value of Nz is the number of variables minus the number of constraints in the predicted active set; i.e., $\mathrm{Nz}=n-($ Bnd $+\operatorname{Lin}+\mathrm{Nln})$.
Bnd is the number of simple bound constraints in the predicted active set.
Lin is the number of general linear constraints in the predicted working set.
$\mathrm{Nln} \quad$ is the number of nonlinear constraints in the predicted active set (not printed if NCNLN is zero).
Penalty is the Euclidean norm of the vector of penalty parameters used in the augmented Lagrangian merit function (not printed if NCNLN is zero).

Cond H is a lower bound on the condition number of the Hessian approximation $H$.
Cond $\mathrm{Hz} \quad$ is a lower bound on the condition number of the projected Hessian approximation $H_{Z}\left(H_{Z}=Z^{\mathrm{T}} H_{\mathrm{FR}} Z=R_{Z}^{\mathrm{T}} R_{Z}\right.$; see (6)). The larger this number, the more difficult the problem.
Cond T is a lower bound on the condition number of the matrix of predicted active constraints.

Conv is a three-letter indication of the status of the three convergence tests (16)-(18) defined in the description of the optional parameter Optimality Tolerance. Each letter is $T$ if the test is satisfied and $F$ otherwise. The three tests indicate whether:
(i) the sequence of iterates has converged;
(ii) the projected gradient (Norm Gz) is sufficiently small; and
(iii) the norm of the residuals of constraints in the predicted active set (Violtn) is small enough.

If any of these indicators is F when E04UCF/E04UCA terminates with $\operatorname{IFAIL}=0$, you should check the solution carefully.

C


is printed if the quasi-Newton update has been modified to ensure that the Hessian approximation is positive definite (see Section 10.4).
is printed if the QP subproblem has no feasible point.
is printed if central differences have been used to compute the unspecified objective and constraint gradients. If the value of Step is zero then the switch to central differences was made because no lower point could be found in the linesearch. (In this case, the QP subproblem is resolved with the central difference gradient and Jacobian.) If the value of Step is nonzero then central differences were computed because Norm Gz and Violtn imply that $x$ is close to a Kuhn-Tucker point (see Section 10.1 in E04UFF/E04UFA). is printed if the linesearch has produced a relative change in $x$ greater than the value defined by the optional parameter Step Limit. If this output occurs frequently during later iterations of the run, optional parameter Step Limit should be set to a larger value.
is printed if the approximate Hessian has been refactorized. If the diagonal condition estimator of $R$ indicates that the approximate Hessian is badly conditioned then the approximate Hessian is refactorized using column interchanges. If necessary, $R$ is modified so that its diagonal condition estimator is bounded.

