

NAG Toolbox for MATLAB

e04uf

1 Purpose

e04uf 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.

e04uf may also be used for unconstrained, bound-constrained and linearly constrained optimization.

e04uf uses **reverse communication** for evaluating the objective function, the nonlinear constraint functions and any of their derivatives.

2 Syntax

```
[irevcm, iter, istate, c, cjac, clamda, objf, objgrd, r, x, needc,
iwork, work, cwsav, lwsav, iwsav, rwsav, ifail] = e04uf(irevcm, nclin,
a, bl, bu, iter, istate, c, cjac, clamda, objf, objgrd, r, x, iwork,
work, cwsav, lwsav, iwsav, rwsav, 'n', n, 'ncnln', ncnln)
```

Before calling e04uf, or either of the option setting functions e04ue, e04wb **must** be called.

3 Description

e04uf 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:

$$\underset{x \in R^n}{\text{Minimize } F(x)} \quad \text{subject to} \quad l \leq \begin{pmatrix} x \\ A_L x \\ c(x) \end{pmatrix} \leq u, \quad (1)$$

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 twice-continuously differentiable. (The method of e04uf 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 e04mf, e04nc or e04nf, or e04nq if the problem is large and sparse. If the problem is large and sparse and does have nonlinear constraints, e04ug should be used, since e04uf treats all matrices as dense.

e04uf uses reverse communication for evaluating $F(x)$, $c(x)$ and as many of their first partial derivatives as possible; any remaining derivatives are approximated by finite differences. See the description of the optional parameter **Derivative Level**.

On initial entry, you must supply an initial estimate of the solution to (1).

On intermediate exits, the calling program must compute appropriate values for the objective function, the nonlinear constraints or their derivatives, as specified by the parameter **irevcm**, and then re-enter the function. Just before an intermediate exit when derivatives are required, each element of the current array of derivatives **objgrd** or **cjac** is initialized to a special value. On re-entry to the function, any element that retains this value is estimated by finite differences.

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 they cannot all be provided, it is advisable to provide as many as possible. While developing code to evaluate the objective function and the constraints, the optional parameter **Verify** should be used to check the calculation of any known derivatives.

The method used by e04uf is described in detail in Section 10.

e04wd is an alternative function which uses a similar method, but with **forward communication**: that is, the objective and constraint functions are evaluated by (sub)programs, supplied as parameters to the function.

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* (ed 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 1986a 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 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 1984b 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 1986b 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 1986c 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
- Murtagh B A and Saunders M A 1983 MINOS 5.0 user's guide *Report SOL 83-20* Department of Operations Research, Stanford University
- Powell M J D 1974 Introduction to constrained optimization *Numerical Methods for Constrained Optimization* (ed 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* (ed A Bachem, M Grötschel and B Korte) 288–311 Springer–Verlag

5 Parameters

Note: this function uses **reverse communication**. Its use involves an initial entry, intermediate exits and re-entries, and a final exit, as indicated by the **parameter irevcm**. Between intermediate exits and re-entries, **all parameters must remain unchanged except those specified by the value of irevcm**.

5.1 Compulsory Input Parameters

1: **irevcm** – int32 scalar

On initial entry: must be set to 0.

On intermediate re-entry: **must remain unchanged**, unless you wish to terminate the solution to the current problem. In this case **irevcm** may be set to a negative value and then e04uf will take a final exit with **ifail** set to this value of **irevcm**.

Constraint: $\text{irevcm} \leq 6$.

2: **nclin** – int32 scalar

On initial entry: n_L , the number of general linear constraints.

Constraint: $\text{nclin} \geq 0$.

3: **a(lda,*)** – double array

The first dimension of the array **a** must be at least $\max(1, \text{nclin})$

The second dimension of the array must be at least **n** if $\text{nclin} > 0$, and at least 1 otherwise

On initial entry: the i th row of the array **a** must contain 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 $i = 1, 2, \dots, \text{nclin}$.

If $\text{nclin} = 0$, the array **a** is not referenced.

4: **bl(n + nclin + ncnn)** – double array

5: **bu(n + nclin + ncnn)** – double array

On initial 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 $\text{bl}(j) \leq -\text{bigbnd}$, and to specify a nonexistent upper bound (i.e., $u_j = +\infty$), set $\text{bu}(j) \geq \text{bigbnd}$; 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 $\text{bl}(j) = \text{bu}(j) = \beta$, say, where $|\beta| < \text{bigbnd}$.

Constraints:

$$\begin{aligned} \text{bl}(j) &\leq \text{bu}(j), \text{ for } j = 1, 2, \dots, n + \text{nclin} + \text{ncnn}; \\ \text{if } \text{bl}(j) &= \text{bu}(j) = \beta, |\beta| < \text{bigbnd}. \end{aligned}$$

6: **iter** – int32 scalar

On intermediate re-entry: must remain unchanged from a previous call to e04uf.

7: **istate(n + nclin + ncnn)** – int32 array

On initial 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 **istate**(j) are as follows:

istate(j)	Meaning
0	The corresponding constraint is <i>not</i> 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 bl(j) = bu(j) .

The values -2 , -1 and 4 are also acceptable but will be modified by the function. If e04uf has been called previously with the same values of **n**, **nclin** and **ncnln**, **istate** already contains satisfactory information. (See also the description of the optional parameter **Warm Start**.) The function also adjusts (if necessary) the values supplied in **x** to be consistent with **istate**.

Constraint: $-2 \leq \text{istate}(j) \leq 4$, for $j = 1, 2, \dots, \mathbf{n} + \mathbf{nclin} + \mathbf{ncnln}$.

8: **c(*)** – double array

Note: the dimension of the array **c** must be at least $\max(1, \mathbf{ncnln})$.

On initial entry: need not be set.

On intermediate re-entry: if **irevcm** = 4 or 6 and **needc(i) > 0**, **c(i)** must contain the value of the *i*th constraint at *x*. The remaining elements of **c**, corresponding to the nonpositive elements of **needc**, are ignored.

9: **cjac(ldcj,*)** – double array

The first dimension of the array **cjac** must be at least $\max(1, \mathbf{ncnln})$

The second dimension of the array must be at least **n** if **ncnln** > 0, and at least 1 otherwise

On initial entry: in general, **cjac** need not be initialized before the call to e04uf. However, if the optional parameter **Derivative Level** = 2 or 3, you may optionally set the constant elements of **cjac**. Such constant elements need not be re-assigned on subsequent intermediate exits.

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 during an intermediate exit will retain its initial value throughout. Constant elements may be loaded into **cjac** either before the call to e04uf or during the first intermediate exit. 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 during intermediate exits.

On intermediate re-entry: if **irevcm** = 5 or 6 and **needc(i) > 0**, the *i*th row of **cjac** must contain the available elements of the vector ∇c_i given by

$$\nabla c_i = \left(\frac{\partial c_i}{\partial x_1}, \frac{\partial c_i}{\partial x_2}, \dots, \frac{\partial c_i}{\partial x_n} \right)^T,$$

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*. The remaining rows of **cjac**, corresponding to nonpositive elements of **needc**, are ignored.

Note that constant nonzero elements do affect the values of the constraints. Thus, if **cjac(i,j)** is set to a constant value, it need not be reset during subsequent intermediate exits, but the value **cjac(i,j) × x(j)** must nonetheless be added to **c(i)**. For example, if **cjac(1,1) = 2** and **cjac(1,2) = -5**, then the term $2 \times \mathbf{x}(1) - 5 \times \mathbf{x}(2)$ must be included in the definition of **c(1)**.

It must be emphasised 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.

See also the description of the optional parameter **Verify**.

10: **clamda**(**n** + **nclin** + **ncnln**) – double array

On initial 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 $j = n + nclin + 1, n + nclin + 2, \dots, n + nclin + ncnln$. The remaining elements need not be set. Note that if the *j*th constraint is defined as ‘inactive’ by the initial value of the **istate** array (i.e. **istate**(*j*) = 0), **clamda**(*j*) should be zero; if the *j*th constraint is an inequality active at its lower bound (i.e. **istate**(*j*) = 1), **clamda**(*j*) should be nonnegative; if the *j*th constraint is an inequality active at its upper bound (i.e. **istate**(*j*) = 2), **clamda**(*j*) should be nonpositive. If necessary, the function will modify **clamda** to match these rules.

11: **objf** – double scalar

On initial entry: need not be set.

On intermediate re-entry: if **irevcn** = 1 or 3, **objf** must be set to the value of the objective function at *x*.

12: **objgrd**(**n**) – double array

On initial entry: need not be set.

On intermediate re-entry: if **irevcn** = 2 or 3, **objgrd** must contain the available elements of the gradient evaluated at *x*.

See also the description of the optional parameter **Verify**.

13: **r**(**ldr**,**n**) – double array

ldr, the first dimension of the array, must be at least **n**.

On initial 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.

14: **x**(**n**) – double array

On initial entry: an initial estimate of the solution.

15: **iwork**(**liwork**) – int32 array

On initial entry:

Constraint: **liwork** $\geq 3 \times n + nclin + 2 \times ncnln$.

16: **work**(**lwork**) – double array

On initial entry:

Constraints:

if **ncnln** = 0 and **nclin** = 0, **lwork** $\geq 21 \times n + 2$;
 if **ncnln** = 0 and **nclin** > 0, **lwork** $\geq 2 \times n^2 + 21 \times n + 11 \times nclin + 2$;
 if **ncnln** > 0 and **nclin** ≥ 0 , **lwork** $\geq 2 \times n^2 + n \times nclin + 2 \times n \times ncnln + 21 \times n + 11 \times nclin + 22 \times ncnln + 1$.

The amounts of workspace provided and required may be (by default for e04uf) output on the current advisory message unit (as defined by x04ab). As an alternative to computing **liwork** and

lwork from the formulae given above, you may prefer to obtain appropriate values from the output of a preliminary run with **liwork** and **lwork** set to 1. (e04uf will then terminate with **ifail** = 9.)

- 17: **cwsav(5)** – string array
- 18: **lwsav(120)** – logical array
- 19: **iwsav(610)** – int32 array
- 20: **rwsav(475)** – double array

The arrays **lwsav**, **iwsav**, **rwsav** and **cwsav** **must not** be altered between calls to any of the functions e04wb, e04uf, e04ue.

5.2 Optional Input Parameters

- 1: **n** – int32 scalar

Default: The dimension of the arrays **objgrd**, **r**, **x**. (An error is raised if these dimensions are not equal.)

On initial entry: n , the number of variables.

Constraint: $n > 0$.

- 2: **ncnln** – int32 scalar

Default: The dimension of the array **c**.

On initial entry: n_N , the number of nonlinear constraints.

Constraint: $\text{ncnln} \geq 0$.

5.3 Input Parameters Omitted from the MATLAB Interface

lda, ldcj, ldr, liwork, lwork

5.4 Output Parameters

- 1: **irevcn** – int32 scalar

On intermediate exit: specifies what values the calling program must assign to parameters of e04uf before re-entering the function.

irevcm = 1

Set **objf** to the value of the objective function $F(x)$.

irevcm = 2

Set **objgrd**($<j$) to the value $\frac{\partial F}{\partial x_j}$ if available, for $j = 1, 2, \dots, n$.

irevcm = 3

Set **objf** and **objgrd**(j) as for **irevcm** = 1 and **irevcm** = 2.

irevcm = 4

Set **c**(i) to the value of the constraint function $c_i(x)$, for each i such that **needc**(i) > 0.

irevcm = 5

Set **cjac**(i,j) to the value $\frac{\partial c_i}{\partial x_j}$ if available, for each i such that **needc**(i) > 0 and $j = 1, 2, \dots, n$.

irevcm = 6

Set **c**(i) and **cjac**(i,j) as for **irevcm** = 4 and **irevcm** = 5.

On final exit: **irevcm** = 0.

2: **iter** – **int32** scalar

On final exit: the number of major iterations performed.

3: **istate**(**n** + **nclin** + **ncnln**) – **int32** array

On final exit: the status of the constraints in the QP working set at the point returned in **x**. The significance of each possible value of **istate**(j) is as follows:

istate (j)	Meaning
–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	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 bl (j) = bu (j).

4: **c**(*) – **double** array

Note: the dimension of the array **c** must be at least $\max(1, \text{ncnln})$.

On final exit: if **ncnln** > 0, **c**(i) contains the value of the i th nonlinear constraint function c_i at the final iterate, for $i = 1, 2, \dots, \text{ncnln}$.

If **ncnln** = 0, the array **c** is not referenced.

5: **cjac(ldcj,*)** – double array

The first dimension of the array **cjac** must be at least $\max(1, \mathbf{ncnln})$

The second dimension of the array must be at least **n** if $\mathbf{ncnln} > 0$, and at least 1 otherwise

On final exit: if $\mathbf{ncnln} > 0$, **cjac** contains the Jacobian matrix of the nonlinear constraint functions at the final iterate, i.e., **cjac**(*i*,*j*) contains the partial derivative of the *i*th constraint function with respect to the *j*th variable, for $i = 1, 2, \dots, \mathbf{ncnln}$ and $j = 1, 2, \dots, \mathbf{n}$.

If $\mathbf{ncnln} = 0$, the array **cjac** is not referenced.

6: **clamda(n + nclin + ncnln)** – double array

On final exit: the values of the QP multipliers from the last QP subproblem. **clamda**(*j*) should be nonnegative if **istate**(*j*) = 1 and nonpositive if **istate**(*j*) = 2.

7: **objf** – double scalar

On final exit: the value of the objective function at the final iterate.

8: **objgrd(n)** – double array

On final exit: the gradient of the objective function at the final iterate (or its finite difference approximation).

9: **r(ldr,n)** – double array

On final exit: if **Hessian** = No, **r** contains the upper triangular Cholesky factor R of $Q^T \tilde{H} Q$, 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.

10: **x(n)** – double array

On intermediate exit: the point x at which the objective function, constraint functions or their derivatives are to be evaluated.

On final exit: the final estimate of the solution.

11: **needc(*)** – int32 array

Note: the dimension of the array **needc** must be at least $\max(1, \mathbf{ncnln})$.

On intermediate exit: if **irevcn** ≥ 4 , **needc** specifies the indices of the elements of **c** and/or **cjac** that must be assigned. If **needc**(*i*) > 0 , then the *i*th element of **c** and/or the available elements of the *i*th row of **cjac** must be evaluated at x .

12: **iwork(liwork)** – int32 array13: **work(lwork)** – double array

The amounts of workspace provided and required may be (by default for e04uf) output on the current advisory message unit (as defined by x04ab). As an alternative to computing **liwork** and **lwork** from the formulae given above, you may prefer to obtain appropriate values from the output of a preliminary run with **liwork** and **lwork** set to 1. (e04uf will then terminate with **ifail** = 9.)

- 14: **cwsav(5)** – string array
- 15: **lwsav(120)** – logical array
- 16: **iwsav(610)** – int32 array
- 17: **rwsav(475)** – double array

The arrays **lwsav**, **iwsav**, **rwsav** and **cwsav** **must not** be altered between calls to any of the functions e04wb, e04uf, e04ue.

- 18: **ifail** – int32 scalar

0 unless the function detects an error (see Section 6).

6 Error Indicators and Warnings

Errors or warnings detected by the function:

ifail < 0

A negative value of **ifail** indicates an exit from e04uf because you set **irevcn** < 0 during an intermediate exit. The value of **ifail** will be the same as your setting of **irevcn**.

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. e04uf 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**). This condition will also occur if, by chance, an iterate is an ‘exact’ Kuhn–Tucker 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 **ifail** = 1.

ifail = 2

e04uf 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**, or no feasible point could be found in the number of iterations specified by the optional parameter **Minor Iteration Limit**. You should check that there are no constraint redundancies. If the data for the constraints are accurate only to an absolute precision σ , you should ensure that the value of the optional parameter **Linear Feasibility Tolerance** is greater than σ . 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, e04uf should be restarted at a different starting point.

ifail = 4

The limiting number of iterations (as determined by the optional parameter **Major Iteration Limit**) has been reached.

If the algorithm appears to be making satisfactory progress, then optional parameter **Major Iteration Limit** may be too small. If so, either increase its value and rerun e04uf or, alternatively, rerun e04uf 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 **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 function.

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 ϵ_r is the value of the optional parameter **Function Precision**) is too small. In this case you should apply the four tests described under **ifail = 0** 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 e04uf has failed completely to move from the initial point, then values set by the calling program for the objective or constraint functions or their derivatives during intermediate exits may be incorrect. You should refer to comments under **ifail = 7** and check the gradients using the optional parameter **Verify**. 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 ill-conditioning 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 e04uf 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 e04uf with a relaxed value of the optional parameter **Feasibility Tolerance**. (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** ≥ 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 in such a test, and how often the special properties of these numbers make the test meaningless.

Special care should be used in the test if computation of the objective function involves subsidiary data communicated in global 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 function.

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 (sub)programs (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 final 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 ϵ is the *machine precision*).

8 Further Comments

8.1 Description of the Printed Output

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.

Mnr 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 α_k taken along the computed search direction. On reasonably well-behaved 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 e04uf terminates with ifail = 3 (no feasible point could be found for the nonlinear constraints). If there are no nonlinear constraints present (i.e., 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 $\ Z^T g_{FR}\ $, the Euclidean norm of the projected gradient (see Section 10.2). Norm Gz will be approximately zero in the neighbourhood of a solution.
Violtn	is the Euclidean norm of the residuals of constraints that are violated or in the predicted active set (not printed if ncnln is zero). Violtn will be approximately zero in the neighbourhood of a solution.
Cond Hz	is a lower bound on the condition number of the projected Hessian approximation H_Z ($H_Z = Z^T H_{FR} Z = R_Z^T R_Z$; see (6)). The larger this number, the more difficult the problem.
M	is printed if the quasi-Newton update has been modified to ensure that the Hessian approximation is positive-definite (see Section 10.4).
I	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).
L	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.

Varbl	gives the name (V) and index j , for $j = 1, 2, \dots, n$, of the variable.
State	gives the state of the variable (FR if neither bound is in the working set, EQ if a fixed variable, LL if on its lower bound, UL if on its upper bound, TF if temporarily fixed at its current value). If Value lies outside the upper or lower bounds by more than the Feasibility Tolerance , State will be ++ or -- respectively. A key is sometimes printed before State. A <i>Alternative optimum possible</i> . 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 <i>might</i> 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 <i>Degenerate</i> . The variable is free, but it is equal to (or very close to) one of its bounds. I <i>Infeasible</i> . 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 $\mathbf{bl}(j) \leq -bigbnd$.
Upper Bound	is the upper bound specified for the variable. None indicates that $\mathbf{bu}(j) \geq bigbnd$.
Lagr Mult	is the Lagrange multiplier for the associated bound. This will be zero if State is FR unless $\mathbf{bl}(j) \leq -bigbnd$ and $\mathbf{bu}(j) \geq bigbnd$, in which case the entry will be blank. If x is optimal, the multiplier should be nonnegative if State is LL and nonpositive if State is UL.
Slack	is the difference between the variable Value and the nearer of its (finite) bounds $\mathbf{bl}(j)$ and $\mathbf{bu}(j)$. A blank entry indicates that the associated variable is not bounded (i.e., $\mathbf{bl}(j) \leq -bigbnd$ and $\mathbf{bu}(j) \geq bigbnd$).

The meaning of the printout for linear and nonlinear constraints is the same as that given above for variables, with ‘variable’ replaced by ‘constraint’, $\mathbf{bl}(j)$ and $\mathbf{bu}(j)$ replaced by $\mathbf{bl}(n+j)$ and $\mathbf{bu}(n+j)$ respectively and with the following changes in the heading:

L Con	gives the name (L) and index j , for $j = 1, 2, \dots, n_L$, of the linear constraint.
N Con	gives the name (N) and index $(j - n_L)$, for $j = n_L + 1, n_L + 2, \dots, n_L + n_N$, 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

```
irevcm = int32(0);
nclin = int32(1);
a = [1, 1, 1, 1];
bl = [1;
      1;
      1;
      1;
      -9.999999999999999e+24;
      -9.999999999999999e+24;
```

```

        25];
bu = [5;
      5;
      5;
      5;
      20;
      40;
      9.999999999999999e+24];
iter = int32(0);
istate = zeros(7, 1, 'int32');
c = [0;
     0];
cjac = [0, 0, 0, 0;
        0, 0, 0, 0];
clamda = zeros(7, 1);
objf = 0;
objgrd = zeros(4, 1);
r = zeros(4, 4);
x = [1; 5; 5; 1];
iwork = zeros(100, 1, 'int32');
work = zeros(1000, 1);
[cwsav,lwsav,iwsav,rwsav,ifail] = e04wb('e04uf');
if (ifail == 0)
    [irevcm, iter, istate, c, cjac, clamda, objf, objgrd, r, x, needc, ...
     iwork, work, cwsav, lwsav, iwsav, rwsav, ifail] = ...
        e04uf(irevcm, nclin, a, bl, bu, iter, istate, c, cjac, clamda, ...
              objf, objgrd, r, x, iwork, work, cwsav, lwsav, iwsav, rwsav);

while (irevcm > 0)

    if (irevcm == 1 || irevcm == 3)
        % Evaluate the objective function.
        objf = x(1)*x(4)*(x(1)+x(2)+x(3)) + x(3);
    end

    if (irevcm == 2 || irevcm == 3)
        % Evaluate the objective gradient.
        objgrd(1) = x(4)*(2*x(1)+x(2)+x(3));
        objgrd(2) = x(1)*x(4);
        objgrd(3) = x(1)*x(4) + 1;
        objgrd(4) = x(1)*(x(1)+x(2)+x(3));
    end

    if (irevcm == 4 || irevcm == 6)
        % Evaluate the nonlinear constraint functions.
        if (needc(1) > 0)
            c(1) = x(1)^2 + x(2)^2 + x(3)^2 + x(4)^2;
        end
        if (needc(2) > 0)
            c(2) = x(1)*x(2)*x(3)*x(4);
        end
    end

    if (irevcm == 5 || irevcm == 6)
        % Evaluate the constraint Jacobian.
        if (needc(1) > 0)
            cjac(1,1) = 2*x(1);
            cjac(1,2) = 2*x(2);
            cjac(1,3) = 2*x(3);
            cjac(1,4) = 2*x(4);
        end
        if (needc(2) > 0)
            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
    end

    [irevcm, iter, istate, c, cjac, clamda, objf, objgrd, r, x, needc,

```

```

...
    iwork, work, cwsav, lwsav, iwsav, rwsav, ifail] = ...
    e04uf(irevcm, nclin, a, bl, bu, iter, istate, c, cjac, clamda, ...
        objf, objgrd, r, x, iwork, work, cwsav, lwsav, iwsav,
        rwsav);
    end

    if (ifail == 0)
        fprintf('\n Varbl   Istate   Value           Lagr Mult\n');
        for i=1:4
            fprintf(' V   %3d %3d %14.4f %12.4f \n', i, istate(i), x(i),
            clamda(i));
        end
        ax = a*x;
        fprintf('\n L Con   Istate   Value           Lagr Mult\n')
        for i=5:4+nclin
            j=i-4;
            fprintf(' L   %3d %3d %14.4f%12.4f\n', j, istate(i), ax(j),
            clamda(i));
        end
        fprintf('\n L Con   Istate   Value           Lagr Mult\n')
        for i=5+nclin:6+nclin
            j=i-4-nclin;
            fprintf(' N   %3d %3d %14.4f%12.4f\n', j, istate(i), c(j),
            clamda(i));
        end
        fprintf('\n Final objective value = %15.7f\n', objf)
    end
end

```

Varbl	Istate	Value	Lagr Mult
V 1	1	1.0000	1.0879
V 2	0	4.7430	0.0000
V 3	0	3.8211	0.0000
V 4	0	1.3794	0.0000

L Con	Istate	Value	Lagr Mult
L 1	0	10.9436	0.0000

L Con	Istate	Value	Lagr Mult
N 1	2	40.0000	-0.1615
N 2	1	25.0000	0.5523

Final objective value =	17.0140173
-------------------------	------------

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 e04ue. 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 e04uf.

10.1 Overview

e04uf is essentially identical to the (sub)program NPSOL described in Gill *et al.* 1986c.

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_{FX} , with n_{FR} ($n_{\text{FR}} = n - n_{\text{FX}}$) 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 ξ and λ (*the Lagrange multiplier vectors for the bound and general constraints*) such that

$$g = C^T \lambda + \xi \quad (2)$$

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 nonnegative. It is nonpositive 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_{FR} ; i.e., $C_{\text{FR}}Z = 0$. An equivalent statement of the condition (2) in terms of Z is

$$Z^T g_{\text{FR}} = 0.$$

The vector $Z^T g_{\text{FR}}$ 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).

e04uf implements a sequential quadratic programming (SQP) method. For an overview of SQP methods, see Fletcher 1987, Gill *et al.* 1981 and Powell 1983.

The basic structure of e04uf involves *major* and *minor* iterations. The major iterations generate a sequence of iterates $\{x_k\}$ 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

$$\bar{x} = x + \alpha p \quad (3)$$

where x is the current iterate, the nonnegative scalar α 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

$$\underset{p}{\text{Minimize}} \quad g^T p + \frac{1}{2} p^T H p \quad \text{subject to} \quad \bar{l} \leq \begin{Bmatrix} p \\ A_{LP} \\ A_{NP} \end{Bmatrix} \leq \bar{u}, \quad (4)$$

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 $A_{\mathbf{n}}$ is the Jacobian matrix of c evaluated at x . (Finite difference estimates may be used for g and $A_{\mathbf{n}}$; see the optional parameter **Derivative Level**.) Let l in (1) be partitioned into three sections: l_B , l_L and $l_{\mathbf{n}}$, 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 Nln in the monitoring file output of e04uf; see Section 12.) In e04uf, (4) is solved using e04nc. Since solving a quadratic program is itself an iterative procedure, the *minor* iterations of e04uf are the iterations of e04nc. (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 \mathbf{c} denote the m by n matrix of gradients of the general linear and nonlinear constraints in the predicted active set. Firstly, we have available the TQ factorization of C_{FR} :

$$C_{\text{FR}}Q_{\text{FR}} = \begin{pmatrix} 0 & T \end{pmatrix}, \quad (5)$$

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

$$R^T R = H_Q \equiv Q^T \tilde{H} Q, \quad (6)$$

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

$$Q = \begin{pmatrix} Q_{\text{FR}} & \\ & I_{\text{FX}} \end{pmatrix} \quad (7)$$

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

$$Q_{\text{FR}} = \begin{pmatrix} Z & Y \end{pmatrix},$$

then the n_Z ($n_Z \equiv n_{\text{FR}} - m$) columns of Z form a basis for the null space of C_{FR} . The matrix Z is used to compute the projected gradient $Z^T g_{\text{FR}}$ at the current iterate. (The values N_Z and $\text{Norm } G_Z$ printed by e04uf give n_Z and $\|Z^T g_{\text{FR}}\|$, 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 e04uf, 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 e04uf also saves computation in factorizing C_{FR} and H_Q .

Once p has been computed, the major iteration proceeds by determining a step length α 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 e04uf, an iterative procedure from e04nc 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 the optional parameter **Linear Feasibility Tolerance**). If no feasible point exists for the bound and linear constraints, (1) has no solution and e04uf 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 the optional parameter **Difference Interval**). In contrast to the bounds and linear constraints, it must be emphasised 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 e04uf 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 e04nc (see Gill *et al.* 1986a), which was specifically designed to be used within an SQP algorithm for nonlinear programming.

e04nc 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 (sub)programs. 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

$$\bar{p} = p + \sigma d \quad (8)$$

where, as in (3), σ is a nonnegative step length and d is a search direction.

At the beginning of each iteration of e04nc, 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

$$\begin{pmatrix} A_L \\ A_N \end{pmatrix}$$

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

$$C_{FR}d_{FR} = 0, \quad (9)$$

which is equivalent to defining d_{FR} as

$$d_{FR} = Zd_Z \quad (10)$$

for some vector d_Z , where Z is the matrix associated with the TQ factorization (5) of C_{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 γ in the j th position, where j and γ are chosen so that the sum of infeasibilities is decreasing along d . (For further details, see Gill *et al.* 1986a.) In the feasible case, d_Z satisfies the equations

$$R_Z^T R_Z d_Z = -Z^T q_{FR}, \quad (11)$$

where R_Z is the Cholesky factor of $Z^T H_{FR} Z$ and q is the gradient of the quadratic objective function ($q = g + Hp$). (The vector $Z^T q_{FR}$ 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 δ of the order of the **machine precision**, the Lagrange multiplier μ_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 μ_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 $|\mu_j| \leq 1 + \delta$.

The choice of step length σ 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, σ will be taken as unity. (In this case, the projected gradient at \bar{p} will be zero.) Otherwise, σ is set to σ_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_{FR} : if the status of a general constraint changes, a *row* of C_{FR} is altered; if a bound constraint enters or leaves the working set, a *column* of C_{FR} changes. Explicit representations are recurred of the matrices T , Q_{FR} and R , and of the vectors $Q^T q$ and $Q^T g$.

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 α in (3) that produces a ‘sufficient decrease’ in the augmented Lagrangian merit function

$$L(x, \lambda, s) = F(x) - \sum_i \lambda_i (c_i(x) - s_i) + \frac{1}{2} \sum_i \rho_i (c_i(x) - s_i)^2, \quad (12)$$

where x , λ and s vary during the linesearch. The summation terms in (12) involve only the *nonlinear* constraints. The vector λ is an estimate of the Lagrange multipliers for the nonlinear constraints of (1). The nonnegative *slack variables* $\{s_i\}$ 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 nonnegative vector ρ 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 ρ (the printed quantity `Penalty`; see Section 12) is generally nondecreasing, although each ρ_i may be reduced a limited number of times.

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

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 e04uf, the BFGS (Broyden–Fletcher–Goldfarb–Shanno) quasi-Newton update is used:

$$\bar{H} = H - \frac{1}{s^T H s} H s s^T H + \frac{1}{y^T s} y y^T, \quad (13)$$

where $s = \bar{x} - x$ (the change in x).

In e04uf, 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 $y^T s$ 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

$$y_L = \bar{g} - \bar{A}_N^T \mu_N - g + A_N^T \mu_N, \quad (14)$$

where μ_N denotes the QP multipliers associated with the nonlinear constraints of the original problem. If $y_L^T s$ is not sufficiently positive, an attempt is made to perform the update with a vector y of the form

$$y = y_L + \sum_{i=1}^{m_N} \omega_i (a_i(\hat{x}) c_i(\hat{x}) - a_i(x) c_i(x)),$$

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 :

$$\bar{H}_Q = H_Q - \frac{1}{s_Q^T H_Q s_Q} H_Q s_Q s_Q^T H_Q + \frac{1}{y_Q^T s_Q} y_Q y_Q^T, \quad (15)$$

where $y_Q = Q^T y$, and $s_Q = Q^T s$. This update may be expressed as a *rank-one* update to R (see Dennis and Schnabel 1981).

11 Optional Parameters

Several optional parameters in e04uf define choices in the problem specification or the algorithm logic. In order to reduce the number of formal parameters of e04uf 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. A complete list of optional parameters and their default values is given in Section 11.1.

Optional parameters may be specified by calling e04ue before a call to e04uf.

e04ue can be called to supply options directly, one call being necessary for each optional parameter. For example,

```
[lwsav, iwsav, rwsav, inform] = e04ue('Print Level = 1', lwsav, iwsav,
rwsav);
```

e04ue 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 e04uf (unless they define invalid values) and so remain in effect for subsequent calls to e04uf, unless altered by you.

11.1 Optional Parameter Checklist and Default Values

The following list gives the valid options. For each option, we give the keyword, any essential optional qualifiers, the default value and the definition. The minimum abbreviation of each keyword is underlined. If no characters of an optional qualifier are underlined, the qualifier may be omitted. The letter *a* denotes a phrase (character string) that qualifies an option. The letters *i* and *r* denote integer and double values required with certain options. The number ϵ is a generic notation for *machine precision* (see x02aj) and ϵ_r denotes the relative precision of the objective function (the **Function Precision**).

Optional Parameters	Default Values
Central Difference Interval	Default values are computed
Cold Start	Default
Crash Tolerance	Default = 0.01
Defaults	
Derivative Level	Default = 3
Difference Interval	Default values are computed
Feasibility Tolerance	Default = $\sqrt{\epsilon}$
Function Precision	Default = $\epsilon^{0.9}$
Hessian	Default = No
Infinite Bound Size	Default = 10^{20}
Infinite Step Size	Default = $\max(\text{bigbnd}, 10^{20})$
Iteration Limit	See Major Iteration Limit
Iters	See Major Iteration Limit
Itns	See Major Iteration Limit

<u>Line Search Tolerance</u>	Default = 0.9
<u>Linear Feasibility Tolerance</u>	Default = $\sqrt{\epsilon}$
<u>List</u>	Default for e04uf
<u>Major Iteration Limit</u>	Default = $\max(50, 3(n + n_L) + 10n_N)$
<u>Major Print Level</u>	Default for e04uf = 10
<u>Minor Iteration Limit</u>	Default = $\max(50, 3(n + n_L + n_N))$
<u>Minor Print Level</u>	Default = 0
<u>Monitoring File</u>	Default = -1
<u>Nolist</u>	Default for e04uf. See <u>List</u> .
<u>Nonlinear Feasibility Tolerance</u>	Default = $\epsilon^{0.33}$ or $\sqrt{\epsilon}$. See <u>Linear Feasibility Tolerance</u> .
<u>Optimality Tolerance</u>	Default = $\epsilon_r^{0.8}$
<u>Print Level</u>	Default for e04uf = 0. See <u>Major Print Level</u> .
<u>Start Constraint Check At Variable</u>	Default = 1. See <u>Start Objective Check At Variable</u> .
<u>Start Objective Check At Variable</u>	Default = 1
<u>Step Limit</u>	Default = 2.0
<u>Stop Constraint Check At Variable</u>	Default = n . See <u>Start Objective Check At Variable</u> .
<u>Stop Objective Check At Variable</u>	Default = n . See <u>Start Objective Check At Variable</u> .
<u>Verify</u>	See <u>Verify Level</u>
<u>Verify Constraint Gradients</u>	See <u>Verify Level</u>
<u>Verify Gradients</u>	See <u>Verify Level</u>
<u>Verify Level</u>	Default = 0
<u>Verify Objective Gradients</u>	See <u>Verify Level</u>
<u>Warm Start</u>	See <u>Cold Start</u>

11.2 Description of the Optional Parameters

Central Difference Interval r Default values are computed

If the algorithm switches to central differences because the forward-difference approximation is not sufficiently accurate then 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**.

Cold Start 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 e04uf 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. e04uf 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 e04uf is called repeatedly to solve related problems.

Crash Tolerance r

Default = 0.01

This value is used in conjunction with the optional parameter **Cold Start** (the default value) when e04uf 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 $a_j^T x \geq l$ will be included in the initial working set if $|a_j^T x - l| \leq r(1 + |l|)$. 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 by you during intermediate exits. 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 e04uf is more reliable (and will usually be more efficient) when all derivatives are exact.

If $i = 0$ or 2 , e04uf 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 an intermediate exit to the calling program 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 , e04uf will approximate unspecified elements of the constraint Jacobian. One intermediate exit is needed for each variable for which partial derivatives are not available. For example, if the Jacobian has the form

$$\begin{pmatrix} * & * & * & * \\ * & ? & ? & * \\ * & * & ? & * \\ * & * & * & * \end{pmatrix}$$

where ‘*’ indicates an element provided by you and ‘?’ indicates an unspecified element, e04uf will make an intermediate exit to the calling program 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 intermediate exits for information.)

At times, central differences are used rather than forward differences, in which case twice as many intermediate exits 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 r

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 δ_j , where $\delta_j = r(1 + |\hat{x}_j|)$, with \hat{x} the first point feasible with respect to the bounds and linear constraints. If

the functions are well scaled then 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 then a finite difference interval will be computed automatically for each variable by a procedure that requires up to six intermediate exits for each element. This option is recommended if the function is badly scaled or you wish to have e04uf determine constant elements in the objective and constraint gradients.

Feasibility Tolerance

 r 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 ϵ_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 ϵ_r should reflect the relative precision of $1 + |F(x)|$; i.e., ϵ_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 ϵ_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 ϵ_r would be 10^{-10} . The choice of ϵ_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 ϵ_r should be large enough so that e04uf will not attempt to distinguish between function values that differ by less than the error inherent in the calculation.

Hessian

Default = No

This option controls the contents of the upper triangular matrix R (see Section 5). e04uf 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, \mathbf{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 \mathbf{r} . You should select **Hessian** = Yes if a **Warm Start** will be used for the next call to e04uf.

Infinite Bound Size

 r Default = 10^{20}

If $r > 0$, r defines the ‘infinite’ bound $infbnd$ in the definition of the problem constraints. Any upper bound greater than or equal to $infbnd$ will be regarded as plus infinity (and similarly any lower bound less than or equal to $-infbnd$ will be regarded as minus infinity). If $r < 0$, the default value is used.

Infinite Step Size

 r Default = $\max(bighbnd, 10^{20})$

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

Default = 0.9

The value r ($0 \leq r < 1$) controls the accuracy with which the step α 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 . Similarly a nonlinear constraint is considered satisfied if its violation does not exceed r_2 . If $r_m < \epsilon$ or $r_m \geq 1$, the default value is used, for $m = 1$ or 2 .

On entry to e04uf, 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 e04uf (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	Default for e04uf
Nolist	Default for e04uf

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 restore printing.

Major Iteration Limit	i	Default = $\max(50, 3(n + n_L) + 10n_N)$
Iteration Limit		
Iters		
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	i	Default for e04uf = 10
Print Level		Default for e04uf = 0

The value of i controls the amount of printout produced by the major iterations of e04uf, 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 x04ab):

i	Output
0	No 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).
≥ 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**:

i	Output
< 5	No output.
≥ 5	One long line of output (> 80 characters; see Section 12) for each major iteration (no printout of the final solution).
≥ 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).
≥ 30	At each major iteration, the diagonal elements of the matrix T associated with the TQ 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** ≥ 5 and the unit number defined by the optional parameter **Monitoring File** is the same as that defined by x04ab, then the summary output for each major iteration is suppressed.

Minor Iteration Limit i Default = $\max(50, 3(n + n_L + n_N))$

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 e04uf (i.e., the iterations of the quadratic programming algorithm), as indicated below. A detailed description of the printed output is given in Section 8.2 of the document for e04nc (summary output at each minor iteration and the final QP solution) and Section 12 of the document for e04nc) (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 x04ab):

i	Output
0	No output.
1	The final QP solution only.
5	One line of summary output (< 80 characters; see Section 8.2 of the document for e04nc) for each minor iteration (no printout of the final QP solution).
≥ 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**:

i	Output
< 5	No output.
≥ 5	One long line of output (> 80 characters; see Section 8.2 of the document for e04nc) for each minor iteration (no printout of the final QP solution).
≥ 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.
≥ 30	At each minor iteration, the diagonal elements of the matrix T associated with the TQ 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** ≥ 5 and the unit number defined by the optional parameter **Monitoring File** is the same as that defined by x04ab then the summary output for each minor iteration is suppressed.

Monitoring File i Default = -1

If $i \geq 0$ and **Major Print Level** ≥ 5 or $i \geq 0$ and **Minor Print Level** ≥ 5 , monitoring information produced by e04uf 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 r Default = $\epsilon_r^{0.8}$

The parameter r ($\epsilon_r \leq r < 1$) 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 e04uf 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.

e04uf 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

$$\alpha \|p\| \leq \sqrt{r}(1 + \|x\|), \quad (16)$$

where p is the search direction and α the step length from (3). An iterate is considered to satisfy the first-order conditions for a minimum if

$$\|Z^T g_{FR}\| \leq \sqrt{r}(1 + \max(1 + |F(x)|, \|g_{FR}\|)) \quad (17)$$

and

$$|res_j| \leq ftol \quad \text{for all } j, \quad (18)$$

where $Z^T g_{FR}$ is the projected gradient (see Section 10.1), g_{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		Default = n

These keywords take effect only if **Verify Level** > 0 . They may be used to control the verification of gradient elements and/or Jacobian elements computed by the calling program during intermediate exits. 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 $i_{2m-1} \leq 0$ or $i_{2m-1} > \min(n, i_{2m})$, the default value is used, for $m = 1$ or 2 . If $i_{2m} \leq 0$ or $i_{2m} > n$, the default value is used, for $m = 1$ or 2 .

Step Limit r 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 $F(x) = ae^{bx}$ or $F(x) = ax^b$, 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

$$\|\tilde{x} - x\|_2 \leq r(1 + \|x\|_2).$$

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 then 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		
Verify Constraint Gradients		
Verify Gradients		
Verify Objective Gradients		

The possible choices for i are as follows:

i	Meaning
-1	No checks are performed.
0	Only a ‘cheap’ test will be performed, requiring one intermediate exit for the objective function gradients and (if appropriate) one intermediate exit for the partial derivatives of the constraints.
1	In addition to the ‘cheap’ test, individual gradient elements will also be checked using a reliable (but more expensive) test.

For example, the objective gradient will be verified if **Verify**, **Verify** = Yes, **Verify Gradients**, **Verify Objective Gradients** or **Verify Level** = 1 is specified.

These keywords refer to finite difference checks on the gradient elements computed by the calling program during intermediate exits. (Unspecified gradient elements are not checked.) It is possible to specify **Verify Level** = 0 to 3 in several ways, as indicated above. For example, the nonlinear objective gradient (if any) will be verified if either **Verify Objective Gradients** or **Verify Level** = 1 is specified. Similarly, the objective and the constraint gradients will be verified if **Verify** = Yes or **Verify Level** = 3 or **Verify** is specified.

If $i = -1$, then no checking will be performed.

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 intermediate exit for the objective function gradients and (if appropriate) one intermediate exit for the partial derivatives of the constraints.

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** and **Stop** keywords. A result of the form OK or BAD? is printed by e04uf 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 calling program 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 e04uf. (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 (see the description of the optional parameter **Major Print Level**).

When **Major Print Level** ≥ 5 and **Monitoring File** ≥ 0 , the following line of output is produced at every major iteration of e04uf 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.
Mnr	<p>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).</p> <p>Note that Mnr may be greater than the optional parameter Minor Iteration Limit if some iterations are required for the feasibility phase.</p>
Step	is the step α_k taken along the computed search direction. On reasonably well-behaved problems, the unit step (i.e., $\alpha_k = 1$) will be taken as the solution is 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	<p>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.</p> <p>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 e04uf terminates with ifail = 3 (no feasible point could be found for the nonlinear constraints).</p> <p>If there are no nonlinear constraints present (i.e., 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.</p>
Norm Gz	is $\ Z^T g_{FR}\ $, 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., $Nz = n - (Bnd + Lin + Nln)$.
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.
Nln	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 Hz	is a lower bound on the condition number of the projected Hessian approximation H_Z ($H_Z = Z^T H_{FR} Z = R_Z^T R_Z$; 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	<p>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:</p> <ul style="list-style-type: none"> (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. <p>If any of these indicators is F when e04uf terminates with ifail = 0, you should check the solution carefully.</p>
M	is printed if the quasi-Newton update has been modified to ensure that the Hessian approximation is positive-definite (see Section 10.4).
I	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).
L	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.
