

1 SUMMARY

To find the minimum of a function $F(\mathbf{x})$ of n variables \mathbf{x} . It is assumed that the function is differentiable, although it is not necessary to supply a formula for the derivatives. The method used is a quasi-Newton method in which derivatives are estimated by differences and is described in R. Fletcher, *FORTTRAN subroutines for minimization by quasi-Newton methods*, AERE Report R7125 (1972). The subroutine complements VA04 but some comparisons (R. Fletcher, *loc. cit*) indicate that VA04 is less efficient than VA10 and more affected by round off error. VA04 also uses twice as much storage as VA10. It is therefore suggested that VA10 be used in the first instance on any problem. VA10 should not be used when explicit expressions are available for derivatives (use VA09) nor when the function is a sum of squares (use VA05, VA02 or one of the NS routines as appropriate).

ATTRIBUTES — **Version:** 1.0.0. **Types:** VA10A, VA10AD **Calls:** MC11, FD05. **Original date:** April 1972. **Origin:** R. Fletcher, Harwell.

2 HOW TO USE THE PACKAGE

2.1 The argument list and calling sequence

The single precision version

```
CALL VA10A(FUNCT, N, X, F, G, H, W, DFN, XM, HH, EPS,  
          MODE, MAXFN, IPRINT, IEXIT)
```

The double precision version

```
CALL VA10AD(FUNCT, N, X, F, G, H, W, DFN, XM, HH, EPS,  
          MODE, MAXFN, IPRINT, IEXIT)
```

FUNCT is the name of the subroutine provided by the user which is considered in section 3. It must be declared in an EXTERNAL statement.

N is an INTEGER to be set to the number of variables ($N \geq 2$).

X is a REAL (DOUBLE PRECISION in the D version) array in which the solution is stored. An initial approximation must be set in X on entry to VA10 and the best estimate obtained will be returned on exit.

F is a REAL (DOUBLE PRECISION in the D version) number in which the best value of $F(\mathbf{x})$ corresponding to X above will be returned.

G is a REAL (DOUBLE PRECISION in the D version) array of N elements which is used to store an estimate of the gradient vector $\partial F(\mathbf{x})/\partial \mathbf{x}$. This array need not be set on entry.

H is a REAL (DOUBLE PRECISION in the D version) array of dimension $N(N+1)/2$ elements in which an estimate of the Hessian matrix $\partial^2 F/(\partial x_i \partial x_j)$ is stored. The matrix is represented in the product form \mathbf{LDL}^T where \mathbf{L} is a lower triangular matrix with unit diagonals and \mathbf{D} is a diagonal matrix. The lower triangle of \mathbf{L} is stored by columns in H excepting that the unit diagonal elements are replaced by the corresponding elements of \mathbf{D} . The setting of H on entry is controlled by the parameter MODE (*q.v.*).

A is a REAL (DOUBLE PRECISION in the D version) array of $3N$ elements used as working space.

DFN is a REAL (DOUBLE PRECISION in the D version) number which must be set so as to give VA10 an estimate of the likely reduction to be obtained in $F(\mathbf{x})$. DFN is used only on the first iteration so an order of magnitude estimate will suffice. The information can be provided in different ways depending upon the sign of DFN which should be set in one of the following ways:

if $DFN > 0$ the setting of DFN itself will be taken as the likely reduction to be obtained in $F(\mathbf{x})$.

if $DFN=0$ it will be assumed that an estimate of the minimum value of $F(\mathbf{x})$ has been set in argument F , and the likely reduction in $F(\mathbf{x})$ will be computed according to the initial function value.

if $DFN < 0$ a multiple $|DFN|$ of the modulus of the initial function value will be taken as an estimate of the likely reduction.

XM is a REAL (DOUBLE PRECISION in the D version) array of N elements to be set on entry so that $XM(I) > 0$ contains an indication of the magnitude of $X(I)$. This quantity need not be set precisely as it is merely used in scaling the problem.

HH is a REAL (DOUBLE PRECISION in the D version) number to be set so that $HH \cdot XM(I)$ contains a step length to be used in calculating $G(I)$ by differences. Set HH equal to $2^{-t/2}$ where t is the number of significant binary digits in the calculation of F . If F contains only small errors the setting $HH=1.E-3$ is appropriate for VA10A and $HH=1.0D-6$ for VA10AD.

EPS is a REAL (DOUBLE PRECISION in the D version) number to be set on entry so that the accuracy required in $X(I)$ is $EPS \cdot XM(I)$ for all I , ($EPS > 0$).

MODE is an INTEGER which controls the setting of the initial estimate of the Hessian matrix in the parameter H . The following settings of **MODE** are permitted:

if **MODE**=1 an estimate corresponding to a unit matrix is set in H by VA10A.

if **MODE**=2 VA10A/AD assumes that the Hessian matrix itself has been set in H by columns of its lower triangle, and the conversion to \mathbf{LDL}^T form is carried out by VA10. The Hessian matrix must be positive definite.

if **MODE**=3 VA10A/AD assumes that the Hessian matrix has been set in H in product form. This is convenient when using the \mathbf{H} matrix from one problem as an initial estimate for another, in which case the contents of H are passed on unchanged.

MAXFN is an INTEGER set to the maximum number of calls of **FUNCT** permitted. Up to $2N$ more calls may be taken if the limit is exceeded whilst evaluating a gradient vector by differences.

IPRINT An is an INTEGER controlling printing. Printing occurs every $|IPRINT|$ iterations and also on exit, in the form

Iteration No. No of calls of **FUNCT**, **IEXIT** (on exit only)

Function only $X(1), X(2), \dots, X(N)$ 8 to a line (5 in VA10AD)

$G(1), G(2), \dots, G(N)$ 8 to a line (5 in VA10AD)

The values of X and G can be suppressed on intermediate iterations by setting $IPRINT < 0$. All intermediate printing can be suppressed by setting $IPRINT=MAXFN+1$. All printing can be suppressed by setting $IPRINT=0$.

IEXIT is an INTEGER giving the reason for exit from VA10A/AD. This will be set by VA10A/AD as follows

if **IEXIT**=0 (**MODE**=2 only) the estimate of the Hessian matrix is not positive definite.

if **IEXIT**=1 a normal exit has been made in which $|DX(I)| < EPS(I)$ for all $I=1, 2, \dots, N$, where $DX(I)$ is the change in X on an iteration.

if **IEXIT**=2 $\mathbf{G}^T \mathbf{DX} > 0$. This is an error exit, either due to rounding errors because EPS is set too small for the computer word length, or because the truncation error in the finite difference formula for \mathbf{G} is dominant.

if **IEXIT**=3 **FUNCT** has been called **MAXFN** times.

3 USER SUBROUTINE

The user must provide a subroutine headed

The single precision version

```
SUBROUTINE FUNCT(N,X,F)
REAL X(N)
```

The double precision version

```
SUBROUTINE FUNCT(N,X,F)
DOUBLE PRECISION X(N)
```

here FUNCT is an identifier chosen by the user. This subroutine should use the variable x supplied in $X(1), X(2), \dots, X(N)$ to evaluate the function and place it in F. XXX must be passed to VA10A/AD as VA10A/AD's first argument, see section 2, and appear in an EXTERNAL statement in the program that calls VA10AD.

N is an INTEGER equal to the number of variables

X is a REAL (DOUBLE PRECISION in the D version) array of dimension N giving the values of \mathbf{x} at which the function is to be evaluated.

F is a REAL (DOUBLE PRECISION in the D version) variable in which the value of $F(\mathbf{x})$ corresponding to X above must be returned.

4 GENERAL

Use of common: : none

Workspace: : N words + $N(N+1)/2$ words + 4N words provided by the user in G, H and W.

Input/output: : controlled by the user through IPRINT. All output is on stream 6.

Restrictions: : none