Nikhef-10-002
arXiv:1005.1481

A Portable Collection of Fortran Utility Routines

MBUTIL version 5.0

M. Botje*
Nikhef, Science Park 105, 1098XG Amsterdam, the Netherlands

January 27, 2016

Abstract

The MBUTIL library is a well documented collection of FORTRAN routines. Some of these are taken from the public libraries CERNLIB and NETLIB while others are privately developed for use in the QCD evolution program QCDNUM. The aim of this collection is to make QCDNUM independent of the availability of external libraries, and also to give easy access to a large set of general-purpose QCDNUM routines. Of particular interest are routines that partition a linear store into multidimensional arrays—these are the basis of the simple but effective and fast QCDNUM dynamic memory management—and routines for fast interpolation. Also included is a string formatter that can be used to process free-format datacards.

*email m.botje@nikhef.nl

Contents

1 Introduction

The MBUTIL package is an integral part of the QCDNUM distribution[1] and contains a pool of FORTRAN utility routines. Some of these are developed privately, some are taken from CERNLIB and some are picked-up from public source code repositories such as NETLIB[2].

The routines in MBUTIL are grouped into general purpose utilities (Section 2), fast solution of triangular and band systems (Section 3), pointer arithmetic in a linear store (Section 4), fast interpolation (Section 5), bitwise operations (Section 6) character string manipulations (Section 7) and a string formatter (Section 8).

It is worthwhile to have a look at Section 4 where routines are described that dynamically partition a linear store into multi-dimensional arrays, and allow you to achieve fast indexing in these arrays. Indeed, QCDNUM owes much of its flexibility and speed to these routines. Of interest may also be Section 5 (fast interpolation) and Section 8 where it is described how to use a string formatter to process free-format datacards.

The syntax of the MBUTIL calls is as follows

```
xMB_name ( arguments )
```

where x = S for subroutines and x = L, I, R or D for logical, integer, real and double-precision functions, respectively. The functions must be explicitly typed in the calling routine unless this is taken care of by an implicit type declaration, thus:

```
logical  lval, lmb_function
lval = LMB_FUNCTION ( arguments )
```

Unless otherwise stated all floating point calculations are done in double precision. This implies that actual floating point arguments should be given in double precision format:

```
dval = dmb_gamma ( 3.D0 )      ! ok
dval = dmb_gamma ( 3.0  )      ! wrong!
```

2 Utility Routines

In this section we describe the MBUTIL utility programs given in Table 1.

`dval = DMB_GAMMA ( x )`

Calculate the gamma function

$$\Gamma(x) = \int_0^\infty e^{-t}\, t^{\,x-1}\, dt \quad (x > 0).$$

The function `dmb_gamma` as well as `x` and `dval` should be declared `double precision` in the calling routine. Code taken from CERNLIB C302 (`dgamma`).

[1] `http://www.nikhef.nl/user/h24/qcdnum`
[2] `http://www.netlib.org`

Table 1: Utility routines in MBUTIL. Output variables are marked by an asterisk (*) and in-out variables by an exclamation mark (!). CERNLIB routines are identified in the first column.

CERNLIB	Subroutine or function	Description
C302	`DMB_GAMMA ( x )`	Gamma function
C332	`DMB_DILOG ( x )`	Dilogarithm
D401	`SMB_DERIV ( f, x, !del, *dfdx, *derr )`	Differentiation
D103	`DMB_GAUSS ( f, a, b, e )`	Gauss integration
F010	`SMB_DMINV ( n, !a, m, ir, *ierr )`	Matrix inversion
F010	`SMB_DMEQN ( n, a, m, ir, *ierr, k, b )`	Linear equations
F012	`SMB_DSINV ( n, !a, m, *ierr )`	Invert symmetric matrix
M103	`SMB_RSORT ( !rarr, n )`	Sort real array
	`SMB_ASORT ( !rarr, n, *m )`	Sort and weed real array
	`RMB_URAND ( !iy )`	Uniform random numbers

```
dval = DMB_DILOG ( x )
```

Calculate the dilogarithm

$$\mathrm{Li}_2(x) = -\int_0^x \frac{\ln|1-t|}{t}\,dt.$$

The function `dmb_dilog` as well as `x` and `dval` should be declared `double precision` in the calling routine. Code taken from CERNLIB C332 (`ddilog`).

```
call SMB_DERIV ( fun, x, !del, *dfdx, *erel )
```

Calculate the first derivative $f'(x)$. The derivative of f should exist at and in the neighborhood of x. This is the responsibility of the user: output will be misleading if the function f is not well behaved. Code taken from CERNLIB D401 (`dderiv`).

`fun` User supplied double precision function of one argument (x). Should be declared `external` in the calling routine.

`x` Value of x where the derivative is calculated.

`del` Scaling factor. Can be set to 1 on input and contains the last value of this factor on output (see the CERNLIB write-up).

`dfdx` Estimate of f' on exit. Set to zero if the routine fails.

`erel` Estimate of the relative error on f'. Set to one if the routine fails.

```
dval = DMB_GAUSS ( fun, a, b, epsi )
```

Calculate by Gauss quadrature the integral

$$I = \int_a^b f(x)\,dx.$$

In the calling routine the function `dmb_gauss`, all its arguments and `dval` should be declared `double precision`. Code taken from CERNLIB D103 (`dgauss`).

`fun` User supplied double precision function of one argument (x). Should be declared `external` in the calling routine.

`a`,`b` Integration limits.

`epsi` Required accuracy of the numerical integration.

```
call SMB_DMINV ( n, arr, idim, ir, *ierr )
```

Calculate the inverse of an $n \times n$ matrix A. Code taken from CERNLIB F010 (`dinv`).

`n` Dimension of the square matrix to be inverted.

`arr` Array, declared in the calling routine as `double precision arr(idim,jdim)` with both `idim` $\geq$ `n` and `jdim` $\geq$ `n`. On entry the first $n \times n$ elements of `arr` should contain the matrix A. On exit these elements will correspond to A^{-1}, provided that A is not found to be singular (as signalled by the flag `ierr`).

`idim` First dimension of `arr`.

`ir` Integer array of at least `n` elements (working space).

`ierr` Set to -1 if A is found to be singular, to 0 otherwise.

```
call SMB_DMEQN ( n, arr, idim, ir, *ierr, k, b )
```

Solve the matrix equation $Ax = b$ with multiple right-hand sides. Code taken from CERNLIB F010 (`deqn`).

`n` Dimension of the square matrix A.

`arr` Array, declared in the calling routine as `double precision arr(idim,jdim)` with both `idim` $\geq$ `n` and `jdim` $\geq$ `n`. On entry the first $n \times n$ elements of `arr` should contain the matrix A. On exit, A is destroyed.

`idim` First dimension of `arr`.

`ir` Integer array of at least `n` elements (working space).

`ierr` Set to -1 if A is found to be singular, to 0 otherwise.

`k` Second dimension of array `b`.

`b` Array, dimensioned `b(idim,k)` that contains on entry a set of k right-hand side vectors of dimension n, and on exit the set of k solution vectors x.

```
call SMB_DSINV ( n, arr, idim, *ierr )
```

Calculate the inverse of an $n \times n$ symmetric positive definite matrix A (i.e. a covariance matrix). Code taken from CERNLIB F012 (`dsinv`).

`n` Dimension of the square matrix to be inverted.

`arr` Array, declared in the calling routine as `double precision arr(idim,jdim)` with both `idim` $\geq$ `n` and `jdim` $\geq$ `n`. On entry the first $n \times n$ elements of `arr` should contain the matrix A. On exit these elements will correspond to A^{-1}, provided that A is found to be positive definite (as signalled by the flag `ierr`).

`idim` First dimension of `arr`.

`ierr` Set to 0 if A is found to be positive definite, to -1 otherwise.

```
call SMB_RSORT ( !rarr, n )
```

Sort the first n elements of the real array `rarr` in ascending order onto itself. On exit we thus have

$$A_1 \leq A_2 \leq \cdots \leq A_{n-1} \leq A_n.$$

Note that `rarr` should be declared `real` and not `double precision` in the calling routine. Code taken from CERNLIB M103 (`flpsor`).

```
call SMB_ASORT ( !rarr, n, *m )
```

Sort the first n elements of the real array `rarr` in ascending order onto itself but discard equal terms. On exit `rarr` contains a list of $m \leq n$ terms such that

$$A_1 < A_2 < \cdots < A_{m-1} < A_m.$$

The remaining $n - m$ elements are undefined. Notice that `rarr` should be declared `real` and not `double precision` in the calling routine.

```
rval = RMB_URAND ( !iy )
```

Return a (real) uniform random number in the interval $(0, 1)$. The integer `iy` should be initialized to an arbitrary value before the first call to `rmb_urand` but should not be altered by the calling routine in between subsequent calls. Note that `rmb_urand` must be declared `real` in the calling routine. Code taken from NETLIB.

3 Triangular and Diagonal Band Equations

In QCDNUM there are lower triangular and lower diagonal band systems to be solved. For reasons of efficient storage and speed, we provide a set of routines optimized for lower and upper systems in different storage schemes.

In what follows we will use the characters `U` and `L` to denote upper and lower triangular or band matrices. A second index labels the storage scheme: (`M`) FORTRAN matrix; (`L`) linear storage; (`T`) packed triangular storage and (`B`) band storage. Thus, the following routine solves a lower diagonal system (`L`) in the FORTRAN matrix storage scheme (`M`).

```
call SMB_LMEQS ( A, na, m, *x, b, n, *ierr )
```

Solve, by forward substitution, the lower diagonal band system $Ax = b$ of dimension n and bandwidth m. When $n = m$, the system is lower triangular.

`A` 2-dim square array declared in the calling routine as `double precision A(na,na)`. The $n \times n$ sub-matrix of `A` should be filled with the lower triangular or lower diagonal band matrix.

`na` Dimension of `A` as declared in the calling routine.

`m` Bandwidth $\leq n$. To be set to n if the matrix is triangular.

`x` Contains the solution vector x on exit. Should be dimensioned to at least n in the calling routine.

`b` Right-hand side vector b. Should be dimensioned to at least n in the calling routine.

`n` Dimension of the triangular system to be solved.

`ierr` Set to a non-zero value if `A` is singular. The output vector `x` is then undefined.

An upper triangular or upper diagonal band system is solved by

```
call SMB_UMEQS ( A, na, m, *x, b, n, *ierr ).
```

For these routines the matrix A is stored in a 2-dimensional FORTRAN array such as

$$A_{ij} = \begin{pmatrix} A_{11} & & & \\ A_{21} & A_{22} & & \\ A_{31} & A_{32} & A_{33} & \\ A_{41} & A_{42} & A_{43} & A_{44} \end{pmatrix} \qquad \text{or} \qquad A_{ij} = \begin{pmatrix} A_{11} & A_{12} & & \\ & A_{22} & A_{23} & \\ & & A_{33} & A_{34} \\ & & & A_{44} \end{pmatrix}.$$

In this scheme $n \times n$ words of storage are used but only $n(n+1)/2$ words are occupied by a triangular matrix and even less by a band matrix. For better use of memory and CPU we provide a set of routines which employ more efficient storage schemes and fast addressing. As already mentioned above, these additional storage schemes are labeled by: (`L`) linear storage; (`T`) packed triangular and (`B`) band storage.

Linear storage: The address $k(i,j)$ in a linear store (column-wise storage) is given by

$$k(i,j) = i + (j-1)n \qquad A_{ij} = \begin{pmatrix} 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \\ 4 & 8 & 12 & 16 \end{pmatrix}.$$

This storage model takes as much space as a normal FORTRAN array ($n \times n$ words) but the address arithmetic in the substitution loops is faster (factor of two, roughly).

Triangular storage: The storage model for lower triangular matrices is

$$k(i,j) = i(i-1)/2 + j \qquad (i \geq j) \qquad A^{\mathrm{LT}}_{ij} = \begin{pmatrix} 1 & & & \\ 2 & 3 & & \\ 4 & 5 & 6 & \\ 7 & 8 & 9 & 10 \end{pmatrix}$$

and for upper triangular matrices

$$k(i,j) = (n+1-i)(n-i)/2 + n + 1 - j \qquad (i \leq j) \qquad A^{\mathrm{UT}}_{ij} = \begin{pmatrix} 10 & 9 & 8 & 7 \\ & 6 & 5 & 4 \\ & & 3 & 2 \\ & & & 1 \end{pmatrix}.$$

Note that the address k is *not* linear in i but linear in j allowing for fast indexing in the forward or backward substitution loop. These storage schemes occupy $n(n+1)/2$ words and are fully efficient for triangular but not for band matrices.

Band storage: Storage of band matrices in $n \times m$ words is achieved by

$$k(i,j) = (i-j)n+i \qquad (i \geq j) \qquad A_{ij}^{\mathrm{LB}} = \begin{pmatrix} 1 & & & \\ 6 & 2 & & \\ & 7 & 3 & \\ & & 8 & 4 \end{pmatrix}$$

$$k(i,j) = (j-i)n+j \qquad (i \leq j) \qquad A_{ij}^{\mathrm{UB}} = \begin{pmatrix} 1 & 6 & & \\ & 2 & 7 & \\ & & 3 & 8 \\ & & & 4 \end{pmatrix}.$$

These schemes are not fully efficient since $m(m-1)/2$ words are wasted but they are better than the triangular schemes when $nm < n(n+1)/2$, that is, when the bandwidth $m < (n+1)/2$. Note that, again, the indexing is linear in j allowing for fast addressing in the substitution loop.

In the table below we list the routines to solve the triangular or banded equations and the corresponding functions to address the elements of the associated matrix. The

Table 2: Routines to solve triangular systems.

Equation Solver	Address Arithmetic	Size of `A`	Scheme
`SMB_LLEQS(A,m,x,b,n,ierr)`	`IMB_LLADR(i,j,m,n)`	$n \times n$	Linear
`SMB_ULEQS(A,m,x,b,n,ierr)`	`IMB_ULADR(i,j,m,n)`		
`SMB_LTEQS(A,m,x,b,n,ierr)`	`IMB_LTADR(i,j,m,n)`	$n(n+1)/2$	Triangular
`SMB_UTEQS(A,m,x,b,n,ierr)`	`IMB_UTADR(i,j,m,n)`		
`SMB_LBEQS(A,m,x,b,n,ierr)`	`IMB_LBADR(i,j,m,n)`	$n \times m$	Band
`SMB_UBEQS(A,m,x,b,n,ierr)`	`IMB_UBADR(i,j,m,n)`		

arguments of these routines are as follows: `A` is a one-dimensional array containing the input matrix (the required size of A is given in Table 2); `m` ($\leq$ `n`) is the bandwidth of the system; `x` is the output solution vector; `b` is the input right-hand side vector and `n` is the dimension of the system to be solved. The error flag `ierr` is set to a non-zero value if `A` is singular in which case `x` is undefined.

Because the structure of the matrix is completely specified by the dimension n and the bandwidth m and because these two parameters are passed to the address functions it is possible to do array boundary checks in these functions: an address of zero is returned when (i,j) does not correspond to an element of the triangular or band matrix. The array boundary check is illustrated in the following example:

```
parameter (n=5, m=2)
dimension s(n,n), a(n*m), x(n), b(n)
```

```
C--   Pack lower band matrix
      do i = 1,100                    ! We can loop over as many i,j
        do j = -50,50                 ! as we like. It does not matter
          k  = imb_LBadr(i,j,m,n)     ! because k = 0 when (i,j) is not
          if(k.ne.0) a(k) = s(i,j)    ! an element of the band matrix ...
        enddo
      enddo
C--   Solve lower diagonal band system
      call smb_LBeqs(a,m,x,b,n,ierr)
```

4 Pointer Arithmetic in a Linear Store

In this section we describe a pointer arithmetic which maps multi-dimensional arrays onto linear storage so that it is possible to declare one large linear store at compilation time and dynamically partition it at run time. This simple method of dynamic memory management already provides many advantages compared to using fixed FORTRAN arrays.

To make clear how it works let us declare a 3-dimensional FORTRAN array

```
dimension A ( i1min:i1plus, i2min:i2plus, i3min:i3plus )
```

The number of words occupied by A is given by $n_A = n_1 n_2 n_3$ with $n_k = i_k^+ - i_k^- + 1$. Instead of declaring A, we now partition a linear storage B which itself is declared as

```
dimension B ( m1:m2 )
```

with $m_2 \geq m_1 + n_A - 1$ if B is to hold the data stored in A. It is easy to construct a linear pointer function $P(i_1, i_2, i_3)$ which assigns a unique address m to any possible combination of the indices:

$$m = P(i_1, i_2, i_3) = C_0 + C_1\, i_1 + C_2\, i_2 + C_3\, i_3. \tag{1}$$

The coefficients C_k are unique functions of $i_1^\pm$, $i_2^\pm$, $i_3^\pm$ and m_1, provided that a convention of 'row-wise' or 'column-wise' storage is adopted. We take in MBUTIL the FORTRAN column-wise convention where the first index 'runs fastest', that is:

$$P(i_1 + 1, i_2, i_3) \equiv P(i_1, i_2, i_3) + 1.$$

In the following we describe the routine `smb_bkmat` which defines the partition of the linear store (much like a FORTRAN dimension statement) and the function `imb_index` which calculates an address in this linear store.

```
call SMB_BKMAT ( imin, imax, *karr, n, im1, *im2 )
```

Define a partition of a linear store B such that it maps onto a multi-dimensional array $A(i_1, \ldots, i_n)$ with n indices. The definition range of each index is $i_k^- \leq i_k \leq i_k^+$.

`imin` Input integer array containing the lower index limits i_k^-. Should be dimensioned to n in the calling routine.

`imax` As above, but now containing the upper limits i_k^+.

`karr` Integer array containing, on exit, the coefficients C_k used to calculate the address in the linear storage. Should be dimensioned to `karr(0:n)` in the calling routine.

`n` Dimension of the partition.

`im1` Address in B where the first word of A should be stored.

`im2` Gives, on exit, the address in B where the last word of A will be stored. B should thus be dimensioned to at least `B(im1:im2)`.

Note that once a partition is defined there is nothing against booking another one starting at `im2+1` provided that the store B is large enough.

```
iaddr = IMB_INDEX ( iarr, karr, n )
```

Calculate an address in the linear store.[3]

`iarr` Input integer array containing the values $(i_1, \ldots, i_n)$ of the indices.

`karr` Input integer array containing the coefficients C_k to calculate the address in the linear store. This array should have been filled beforehand by a defining call to `smb_bkmat`.

`n` Dimension of the partition.

Note that this function does not perform array boundary checks so that it is your responsibility to make sure that all indices stored in `iarr` are within their respective ranges. Note also that `smb_bkmat` and `imb_index` do not *operate* on the store B but merely calculate an address in B.

The address arithmetic as given in (1) provides the possibility to do fast addressing in nested loops. To see this, take for example a 3-dimensional array `A(100,100,100)` and map it onto a linear store `B(1000000)`. Now consider the loop:

```
do i1 = 1,100
  do i2 = 1,100
    do i3 = 1,100
      A123 = B( P(i1,i2,i3) )
      ..
```

where `P(i,j,k)` is a wrapper[4] for `imb_index` that returns the address in B. The address calculation in the inner loop costs 3×10^6 additions and 3×10^6 multiplications. However, from (1) it follows that increasing or decreasing an index i_k by one unit does correspond to a unique fixed increment of the address in B. Fast addressing can then be achieved by maintaining running sums of these increments, as is illustrated below:

[3] After the call to `smb_bkmat` you could also calculate the address yourself through, for 3 indices,

```
iaddr = P(i,j,k) = karr(0) + karr(1)*i + karr(2)*j + karr(3)*k
```

[4] The sole reason to introduce here this wrapper is to make the code examples below easier to read.

```
inc1 = P(2,1,1)-P(1,1,1)    !Address increment of i1
inc2 = P(1,2,1)-P(1,1,1)    !Address increment of i2
inc3 = P(1,1,2)-P(1,1,1)    !Address increment of i3
m1   = P(1,1,1)-inc1        !Base address
do i1 = 1,100
  m1 = m1 + inc1
  m2 = m1 - inc2
  do i2 = 1,100
    m2 = m2 + inc2
    m3 = m2 - inc3
    do i3 = 1,100
      m3   = m3 + inc3
      A123 = B(m3)          !A(i1,i2,i3)
      ..
```

There are now only slightly more than 10^6 additions (no multiplications) to calculate the addresses. Note that the addressing scheme given above works for any nesting order of the loops. When the nesting is in the same order as the indices, with the first index running in the inner loop, then one walks sequentially through the store so that the code simplifies to:

```
ia = P(1,1,1)-1          !Base address
do i3 = 1,100
  do i2 = 1,100
    do i1 = 1,100
      ia   = ia + 1
      A123 = B(ia)       !A(i1,i2,i3)
      ..
```

Here is code with very fast addressing that initializes the array

```
ia = P(  1,  1,  1)      !First address
ib = P(100,100,100)      !Last  address
do i = ia,ib
  B(i) = value
enddo
```

5 Fast Interpolation

Piecewise polynomial interpolation of order n on tabulated data consists of selecting an n-point sub-grid around the interpolation point, followed by an interpolation of the data on that sub-grid. This interpolation can be written as a (nested) weighted sum with weights that do not depend on the data. Pre-calculating the weights for a set of interpolation points thus allows for fast interpolation of more than one table.

We will restrict ourselves here, as in QCDNUM, to the orders $n = 1$ (point value), 2 (linear interpolation) and 3 (quadratic interpolation). The interpolation sub-grids of

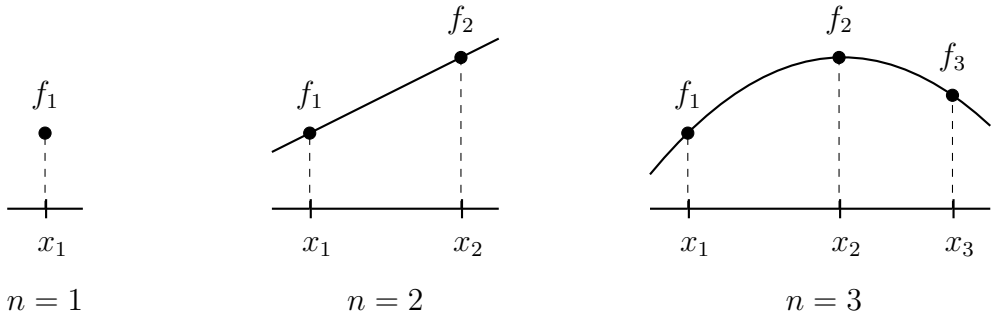

Figure 1: Interpolation function for sub-grids of size $n = 1$, 2 and 3.

size $n = (1, 2, 3)$ are shown in Figure 1. For these sub-grids, the algorithm reads

$n = 1$	$f(x) = w_1(x)\, f_1$	$w_1(x) = 1$	
$n = 2$	$f(x) = w_1(x)\, f_1 + w_2(x)\, f_2$	$w_1(x) = (x - x_1)/(x_2 - x_1)$	$w_2 = 1 - w_1$
$n = 3$	$g(x) = w_1(x)\, f_1 + w_2(x)\, f_2$	$w_1(x) = (x - x_1)/(x_2 - x_1)$	$w_2 = 1 - w_1$
	$h(x) = w_3(x)\, f_2 + w_4(x)\, f_3$	$w_3(x) = (x - x_2)/(x_3 - x_2)$	$w_4 = 1 - w_3$
	$f(x) = w_5(x)\, g \;\; + w_6(x)\, h$	$w_5(x) = (x - x_1)/(x_3 - x_1)$	$w_6 = 1 - w_5$.

It is seen that the interpolation is, for a given interpolation point x, a (nested) weighted sum of the function values f_i, with $(1, 2, 6)$ different weights for interpolation at order $n = (1, 2, 3)$. The weights depend on n, x and the grid points x_i, but not on f.

We extend the algorithm to an $n_x \times n_y$ 2-dimensional interpolation mesh simply by performing n_y interpolations in x and one interpolation in y.

When we have to interpolate more than one table it clearly makes sense to first calculate the weights and then interpolate each table. For this we provide the routine `smb_polwgt` to pre-compute the weights which can then be fed into the interpolation functions `dmb_polin1` and `dmb_polin2` for 1- or 2-dimensional interpolation, respectively.

```
call SMB_POLWGT ( x, xi, n, *w )
```

Compute the weights for interpolation on a $1, 2, 3$-point interpolation grid.

`x` Interpolation point (irrelevant when `n` = `1`). Should be inside the range of the interpolation grid to avoid extrapolation and the corresponding loss of accuracy.

`xi` Input array, dimensioned to at least `n` in the calling routine, filled with the `n` interpolation grid points x_i (see Figure 1).

`n` Number of points in the interpolation grid [`1`–`3`].

`w` Output weight array, dimensioned to at least (`1`,`2`,`6`) for `n` = (`1`,`2`,`3`).

```
val = DMB_POLIN1 ( w, fi, n )
```

One-dimensional interpolation on a $1, 2, 3$-point interpolation grid.

`w` Input weight array filled by an upstream call to `smb_polwgt`.

`fi` Input array, dimensioned to at least `n` in the calling routine, filled with `n` function values f_i (see Figure 1).

`n` Interpolation order [1–3] as set in the upstream call to `dmb_polwgt`.

```
val = DMB_POLIN2 ( wx, nx, wy, ny, fij, m )
```

Two-dimensional interpolation on an $n_x \times n_y$ interpolation mesh.

`wx` Input weight array filled by an upstream call to `smb_polwgt`.

`nx` Interpolation order in x [1–3] as set in the upstream call to `smb_polwgt`.

`wy`, `ny` As above, but now for the interpolation in y.

`fij` Input array, dimensioned to at least `fij(nx,ny)` in the calling routine, filled with the function values f_{ij} to be interpolated.

`m` First dimension of `fij` as declared in the calling routine.

In the example below we interpolate three functions on a 3×2 interpolation mesh.

```
dimension xi(3), wx(6), yi(2), wy(2), fij(3,2), gij(3,2), hij(3,2)
..
fill the arrays xi, yi and fij, gij, hij (code not shown)
..
call smb_polwgt( x, xi, 3, wx )          ! weights in x
call smb_polwgt( y, yi, 2, wy )          ! weights in y
f  = dmb_polin2( wx, 3, wy, 2, fij, 3 )  ! f(x,y)
g  = dmb_polin2( wx, 3, wy, 2, gij, 3 )  ! g(x,y)
h  = dmb_polin2( wx, 3, wy, 2, hij, 3 )  ! h(x,y)
```

6 Bitwise Operations

In this section we describe a few routines and functions to manipulate bits in 32-bit integers. The bit numbering runs from 1 (least significant bit) to 32 (most significant bit). The routines presented below rely on the following machine representation of the integer value +1 which, as far as we know, is standard on all platforms:

```
bit 32                          bit 1
  |                               |
  00000000000000000000000000000001
```

Please notice that the bitwise operations will *not* work for 16-bit integers.

```
call SMB_SBIT1 ( i, n )
```

Set bit n of integer i to 1.

```
call SMB_SBIT0 ( i, n )
```

Set bit n of integer i to 0.

```
ival = IMB_GBITN ( i, n )
```

Give the value (0 or 1) of bit n of integer i.

```
ival = IMB_SBITS ( cpatt )
```

Store a bit pattern in an integer i. The bit pattern is given in the input string `cpatt` containing a sequence of 32 characters '1' and '0'.

Here is an example which sets the value of i to 7:

```
character*32 cpatt
data cpatt /'00000000000000000000000000000111'/

i = imb_sbits ( cpatt )      ! i = 7
```

```
call SMB_GBITS ( i, *cpatt )
```

Store the bit pattern of an integer i in a character string. The output string `cpatt` should be declared `character*32` in the calling routine.

```
ierr = IMB_TEST0 ( mask, i )
```

Verify that a selected set of bits in i are all set to zero.

`mask` Input 32-bit integer. If bit n of `mask` is set to 1 then the corresponding bit of i will be checked. Otherwise bit n will not be checked.

`i` Input 32-bit integer variable to be checked.

`ierr` Non-zero if the test fails. Thus `ierr` = `0` means that all checked bits in i are 0.

```
ierr = IMB_TEST1 ( mask, i )
```

Verify that a selected set of bits in i are all set to one.

`mask` Input 32-bit integer. If bit n of `mask` is set to 1 then the corresponding bit of i will be checked. Otherwise bit n will not be checked.

`i` Input 32-bit integer variable to be checked.

`ierr` Non-zero if the test fails. Thus `ierr` = `0` means that all checked bits in i are 1.

7 Character String Manipulations

In this section we describe a few routines which perform elementary character string manipulations. It is recommended to explicitly initialize strings to a series of blank characters at program start-up. This is easily done by using `smb_cfill`.

```
call SMB_CFILL ( char, string )
```

Fill the character variable `string` with the character `char`.

`char` Input one-character string.

`string` Character string declared `character*n` in the calling routine. On exit all n characters of `string` will be set to `char`.

```
call SMB_CLEFT ( string )
```

Left adjust the characters in `string`, padding blanks to the right.

```
call SMB_CRGHT ( string )
```

Right adjust the characters in `string`, padding blanks to the left.

```
call SMB_CUTOL ( string )
```

Translate the character variable `string` to lower case.

```
call SMB_CLTOU ( string )
```

Translate the character variable `string` to upper case.

```
leng = IMB_LENOC ( string )
```

Returns the position of the rightmost non-blank character in `string`. This function measures the actual length of a string unlike the FORTRAN function `len()` which returns for a `character*n` variable the number `n`.

```
ipos = IMB_FRSTC ( string )
```

Returns the position of the leftmost non-blank character in `string`.

```
lval = LMB_COMPC ( stra, strb, n1, n2 )
```

Case independent comparison of the character substrings `stra(n1:n2)` and `strb(n1:n2)`.

`stra` Input character string declared in the calling routine as `character*na stra`.

`strb` Input character string declared in the calling routine as `character*nb strb`.

`n1`,`n2` Range of characters to be compared. It is required that $1 \leq n_1 \leq n_2 \leq \min(n_a, n_b)$; the comparison will yield `.false.` if this is not the case.

`lval` Set to `.true.` (`.false.`) if `stra(n1:n2)` and `strb(n1:n2)` do (do not) match. Both `lval` and `lmb_compc` should be declared `logical` in the calling routine.

Note that a case dependent comparison by the FORTRAN statement

```
character*10 line1, line2

if ( line1 .eq. line2 ) ...
```

is much faster than the case independent comparison provided by `lmb_compc`.

```
lvar = LMB_MATCH ( string, substr, cwild )
```

Verify that the character string `substr` is contained in `string`. The string `substr` may, or may not, contain a wild character `cwild` which will match any character in `string`. The matching is case insensitive. Notice that, before processing, `string` is internally converted to upper case with trailing blanks stripped off. Likewise `substr` is converted to upper case but here *both* leading and trailing blanks are stripped off.

`string` Input character string of length n_i, including leading but not trailing blanks.

`substr` Input character string of length n_s, not including leading and trailing blanks.

`cwild` Input character (wild character acting as placeholder).

`lvar` Set to true (false) if `substr` is (is not) contained in `string`. Both `lvar` and `lmb_match` should be declared logical in the calling routine.

It is required that the substring contains at least one non-blank character ($n_s > 0$) and that it fits inside the string ($n_s \leq n_i$). The function `lmb_match = .false.` if this is not the case. Here are some examples:

```
logical lmb_match, lvar

lvar = lmb_match ( 'Amsterdam ', ' am ', '*' )  !  .true.
lvar = lmb_match ( 'Amsterdam ', '*am ', '*' )  !  .true.
lvar = lmb_match ( 'Amsterdam ', '*am*', '*' )  !  .false.
lvar = lmb_match ( ' Amsterdam', '*am*', '*' )  !  .true.
```

```
call SMB_ITOCH ( ival, *chout, *lengout )
```

Convert an integer to a character string.

`ival` Input integer.

`chout` Character string containing, on exit, the digits of `ival`. Should be declared `character*n` in the calling routine. If `n` is smaller than the number of digits of `ival`, the string will be filled with asterisks (`*`).

`lengout` Number of characters encoded in `chout`.

With this routine you can nicely embed integers into text strings as shown below:

```
character*10 string
ierr = -12
call smb_itoch(ierr,string,leng)   !unformatted write comes next
write(lun,*) 'Error ',string(1:leng),' encountered'
jerr =  12345
call smb_itoch(jerr,string,leng)   !formatted write comes next
write(lun,'(''Error '',A,'' encountered'')') string(1:leng)
```

This code will produce the strings (note the snug fit of the numbers)

```
Error -12 encountered
Error 12345 encountered
```

8 String Formatter

A powerful feature of FORTRAN is the use of strings as an internal file. A restriction is, however, that the FORTRAN77 standard does not allow list-directed (= free-format) read from internal files. For this reason we provide the routine `smb_sfmat` that determines the format of an arbitrary string so that one can do a *formatted* read on that string.

```
call SMB_SFMAT ( stin, *stout, *fmt, *ierr )
```

Bring a free-format input string into a standard format and give the FORTRAN format descriptor of the reformatted string.

`stin` Input character string. This string will be parsed into words and each word will be classified and reformatted (if necessary) as is described below.

`stout` Output character string with the (reformatted) input words. Must be declared `character*n` in the calling routine, with `n` large enough to hold the reformatted input string (error condition if `n` too small).

`fmt` Output character string with the FORTRAN format descriptor of `stout`. Should be declared `character*m` in the calling routine, with `m` large enough to hold the format descriptor (error condition if `m` too small).

`ierr` Output error flag (note that `stout` and `fmt` will be undefined upon error):

- `0` All OK.
- `1` Empty input string.
- `2` Found unbalanced quotes in `stin`.
- `3` Wordcount exceeded (max 100 words per string).
- `4` Wordlength exceeded (max 120 characters per word).
- `5` Not enough space in `stout`.
- `6` Not enough space in `fmt`.

The input string `stin` is parsed into words (these are defined as substrings separated by one or more blanks) and each word is classified as follows.

- `L` Logical. A single `T` or `F` in upper case;
- `I` Integer. Any signed or unsigned string of digits. Examples: `12`, `-12`, `+12`;
- `F` Floating point number. Any signed or unsigned string of digits with a leading, embedded or trailing decimal point. Examples: `-.12`, `12.`, `1.2`. Note that the first two numbers will be reformatted to `-0.12` and `12.0`, respectively;
- `E` Floating point number in exponentional format. This is the character `E` or `e` preceded by a signed or unsigned integer or foating point number (mantissa) and followed by a signed or unsigned integer (exponent). Examples: `.12E1`, `1.2E0`, `12.E-1`, `12e-1`. These will all be reformatted to `0.12E1`;
- `D` As above, but now in `D`-format;
- `Q` Quoted string. Anything embedded in double single quotes, like `''foo bar''`;[5]
- `A` Character string. Any word that is not classified as `L`, `I`, `F`, `E`, `D` or `Q`.

The use of the string formatter is illustrated by the datacard processor presented below.

Suppose that we want to steer a program with datacards.

```
MYSUB   .3
MYSUB   2D-4  12.
```

The first card should call `mysub(par1,par2)` with a default value for `par2` while the second card should call that same subroutine but with both parameters taken from the card. From this example we see two desirable features of a datacard processor: (i) handling of variable length parameter lists[6] and (ii) handling of free-format input.

[5] A literal string in FORTRAN is embedded in single quotes: `foo` → `'foo'`. Quotes inside the string are represented by double quotes, thus: `foo's` → `'foo''s'`. To process quoted strings, the formatter looks for opening quotes which are either those at the beginning of a string, or those preceded by a blank. When opening quotes are detected, all following characters are classified as quoted string (`Q`) until closing quotes are found (error condition if not). Closing quotes are either those followed by a blank, or those at the end of the string. Quotes inside (quoted) strings are allowed, as long as they cannot be mistaken for a pair of opening and closing quotes. For instance, `Bayes' theorem` can be passed as a 2-word string `'Bayes'' theorem'`, but not as a quoted string `'''Bayes'' theorem'''` because that would stand for the two words `Bayes` (without apostroph) and `theorem'` (with apostroph). In any case, if you use quoted strings it is always a good idea to print the output string `stout` and the format descriptor `fmt` to check that the formatter does what is intended.

[6] With the restriction that optional parameters should be put at the end of the list.

The first requirement can be fulfilled by first trying to read the datacard with two parameters and then, upon error, read it again with one parameter. The second requirement can in principle be fulfilled by a FORTRAN list-directed read.

```
character*5 key
read(unit=lun,fmt=*,err=100,end=100) key, par1, par2
```

Reading the first card in this way indeed produces an error, but not because the parameter list is exhausted but because the list-directed read skips to the next record and tries to read the character string `MYSUB` into the floating point variable `par2`. This brings us to the third requirement that (iii) only one card should be read at the time. This cannot be done with a list-directed read which may span records to fulfill the parameter list.

A solution to this is to read the datacard as a character string. Thus we enclose the cards in quotes (note that the alignment of the quotes below is irrelevant)

```
'  MYSUB   .3        '
'  MYSUB   2D-4  12. '
```

and code the card reading loop as

```
      character*120 dcard
      idum = 0
      do while(idum.eq.0)
        read(unit=lun,fmt='(A)',err=100,end=100) dcard
      enddo
 100  continue
```

This code clearly does a card-by-card reading of a datacard file. It also restricts the length of a card to 120 characters[7] but this is not much of a problem since longer cards tend to violate a fourth requirement: (iv) a datacard must be easy to read on a terminal.

After reading a card, the key (here with a length of 5 characters) can be separated from the parameter list by some straight forward character string manipulation.[8]

```
character*120 dcard, upars
character*5   key5
i1    = imb_frstc(dcard)
key5  = dcard(i1:i1+4)
upars = dcard(i1+5:)
```

Now we can use the string `upars` as an internal file and fetch the parameters from that string by a list-directed read

```
read(unit=upars,fmt=*,end=10,err=10) par1, par2
```

[7] Cards of less than 120 characters are blank padded while longer cards are just truncated.

[8] To keep things simple we do not handle here empty strings, strings with only one word (no parameter list) and strings where the first word is not 5 characters long.

This probably works on your platform but list-directed read from internal files is not supported by the FORTRAN77 standard so that this call is not guaranteed to be portable.

A solution is provided by calling the routine `smb_sfmat` that takes as an input any character string (`upars`, say) and produces a re-formatted output string (`fpars`), together with a format descriptor (`fmt`), thus:

```
call smb_sfmat( upars, fpars, fmt, ierr )
```

For the example cards above this gives:

key	upars	fpars	fmt
MYSUB	.3	0.3	(1X,F3.1)
MYSUB	2D-4 12.	0.2D-3 12.0	(1X,D6.1,1X,F4.1)

Now we can write *portable* code to fetch the parameters by a *formatted* read.

```
character*120 upars, fpars, fmt
call smb_sfmat( upars, fpars, fmt, ierr )
read(unit=fpars,fmt=fmt,end=10,err=10) par1, par2
```

Here are examples of keys that have a quoted string as a parameter:

```
' LOGIC   T  ''F'' '
' VERSN ''Version 3.6 12-AUG-2014'' '
' NAMES   P.A.M.  Dirac  ''P.A.M. Dirac'' '
```

The first card has one logic (`L1`) and one character parameter (`A1`). The numbers in the second card are not classified as such since they are part of a quoted string (`A23`). The last card has three parameters: initials (`A6`), last name (`A5`) and full name (`A12`) which must be quoted because it contains an embedded blank.

Note that in some exceptional cases the format descriptor can become quite long as in

```
stout = ' X X X X '    fmt = '( 1X,A1,1X,A1,1X,A1,1X,A1 )'
```

where each item occupies 2 characters in `stout`, but 6 characters in `fmt`.[9] Because `stout` and `fmt` are declared in the calling routine it is of course easy to adjust their size in case `smb_sfmat` complains about a lack of space.

Putting it all together, we arrive at the card reading code shown in Figure 2. Note that the code, as it stands, does not handle the reading of blank lines from the input file, keys without parameters, keys that are not 5 characters long and errors from `smb_sfmat`.

[9]The formatter is not smart enough to generate repeat counts and write `fmt = '( 4(1X,A1) )'`.

```
      subroutine cardread(lun)

      character*120 dcard, fpars, fmt
      character*5   key5

      idum = 0
      do while(idum.eq.0)

        read(unit=lun,fmt='(A)',err=200,end=100) dcard
        i1    = imb_frstc(dcard)
        key5  = dcard(i1:i1+4)

        call smb_sfmat(dcard(i1+5:), fpars, fmt, ierr)

        if( key5 .eq. 'MYSUB' ) then
          read(unit=fpars,fmt=fmt,end=10,err=10) par1, par2
          call MYSUB(par1,par2)
          return
 10       read(unit=fpars,fmt=fmt,end=20,err=20) par1
          par2 = default
          call MYSUB(par1,par2)
          return
 20       stop 'Error reading MYSUB datacard'
        elseif( key5 .eq. ... ) then
          ..
          process other keys, if any
          ..
        else
          stop 'Unknown card'
        endif

      enddo

 100  return
 200  stop 'Error reading input file'

      end
```

Figure 2: Listing of a datacard reading routine showing the use of the string formatter `smb_sfmat` and the handling of a variable length parameter list.

Index