

Source file: meps.f

```

c=====
c   Computes and reports estimate of machine epsilon.
c
c   Recall: machine epsilon is smallest positive 'eps'
c   such that
c
c           (1.0d0 + eps ) .ne. (1.0d0)
c
c   Program accepts optional argument which specifies
c   division factor: values close to 1.0 will result
c   in more accurate estimate of machine epsilon.
c=====
c   program          meps
c
c   implicit         none
c-----
c   Note use of 'r8arg', available in 'libp410f.a' which
c   works exactly like 'i4arg' except that it returns
c   a real*8 value parsed from the specified command-line
c   argument
c-----
c   real*8          r8arg
c
c   real*8          default_fac
c   parameter      ( default_fac = 2.0d0 )
c
c   real*8          eps,          neweps,          fac
c
c   fac = r8arg(1,default_fac)
c   write(0,*) 'meps: using division factor: ', fac
c
c   eps  = 1.0d0
c   neweps = 1.0d0
c   do while( 1.0d0 .ne. (1.0d0 + neweps) )
c       eps  = neweps
c       neweps = neweps / fac
c   end do
c
c   write(*,*) eps
c
c   stop
c
c   end

```

Source file: meps-sgi-output

```

#####
# Output from 'meps' on SGIs (IEEE 64-bit floating point).
#####

sgi1 1> make meps
f77 -g -64 -c meps.f
f77 -g -64 -L/usr/local/lib meps.o -lp410f -o meps

sgi1 2> meps
meps: using division factor: 2.000000000000000
2.2204460492503131E-16

sgi1 3> meps 1.01
meps: using division factor: 1.010000000000000
1.1104218387155329E-16

sgi1 4> meps 1.0001
meps: using division factor: 1.000100000000000
1.1102645224601785E-16

```

Source file: meps-pclinux-output

```

#####
# Output from 'meps' on PC Linux machine (80 bit floating pt)
#####

lnx1 1> make meps
pgf77 -g -c meps.f
pgf77 -g -L/usr/local/PGI/lib meps.o -lp410f -o meps

lnx1 2> meps
meps: using division factor: 2.000000000000000
1.0842021724855044E-019

lnx1 3> meps 1.01
meps: using division factor: 1.010000000000000
5.4364534909517435E-020

lnx1 4> meps 1.0001
meps: using division factor: 1.000100000000000
5.4212146310714582E-020

```

Source file: catprec.f

```

c=====
c   Program illustrating "catastrophic" loss of precision
c   resulting from the subtraction of two nearly equal
c   floating point values.
c=====
      program          catprec

      implicit        none

      real*8          x
      parameter      ( x = 0.2d0 )

      integer         i
      real*8         h, dsinx

      write(*,*) '      h      d(sin) approx  '//
& 'd(sin) exact  d(sin) err'
      write(*,*)

      h = 0.5d0
      do i = 1, 16

c-----
c      Algebraically, in the limit h -> 0, dsinx should
c      approach cos(x), but sin(x+h) -> sin(x) so
c      catastrophic loss of precision occurs.
c-----
          dsinx = (sin(x+h) - sin(x)) / h
          write(*,1000) h, dsinx, cos(x), dsinx - cos(x)
1000      format(1P,E12.3,2E16.8,E12.3)
          h = 0.125d0 * h
      end do

      stop

      end

```

Source file: catprec-output

```

#####
# Output from 'catprec' illustrating catastrophic precision
# loss due to subtraction of nearly-equal floating point
# values.
#####

lnx1 1> make catprec
pgf77 -g -c catprec.f
pgf77 -g -L/usr/local/PGI/lib catprec.o -o catprec

lnx1 2> catprec
      h      d(sin) approx  d(sin) exact  d(sin) err
5.000E-01  8.91096713E-01  9.80066578E-01  -8.897E-02
6.250E-02  9.73222242E-01  9.80066578E-01  -6.844E-03
7.813E-03  9.79280560E-01  9.80066578E-01  -7.860E-04
9.766E-04  9.79969416E-01  9.80066578E-01  -9.716E-05
1.221E-04  9.80054450E-01  9.80066578E-01  -1.213E-05
1.526E-05  9.80065062E-01  9.80066578E-01  -1.516E-06
1.907E-06  9.80066388E-01  9.80066578E-01  -1.895E-07
2.384E-07  9.80066554E-01  9.80066578E-01  -2.368E-08
2.980E-08  9.80066575E-01  9.80066578E-01  -2.960E-09
3.725E-09  9.80066577E-01  9.80066578E-01  -3.702E-10
4.657E-10  9.80066578E-01  9.80066578E-01  -5.373E-11
5.821E-11  9.80066578E-01  9.80066578E-01  -1.701E-10
7.276E-12  9.80066577E-01  9.80066578E-01  -8.686E-10
9.095E-13  9.80066568E-01  9.80066578E-01  -1.018E-08
1.137E-13  9.80066538E-01  9.80066578E-01  -3.998E-08
1.421E-14  9.80066299E-01  9.80066578E-01  -2.784E-07

```

Source file: dmoutines.f

```

c=====
c   Implements matrix-matrix multiply
c
c   c = a b
c
c   where a, b and c are n x n (square) real*8 matrices.
c=====
      subroutine dmmult(a,b,c,n)

      implicit        none

      integer         n
      real*8         a(n,n), b(n,n), c(n,n)

      integer         i, j, k

      do j = 1, n
        do i = 1, n
          c(i,j) = 0.0d0
          do k = 1, n
            c(i,j) = c(i,j) + a(i,k) * b(k,j)
          end do
        end do
      end do

      return

      end

c=====
c   Writes a double precision matrix (two dimensional
c   array) to file 'fname'. If 'fname' is the
c   string '-', the matrix is written to standard input.
c
c   This routine is modelled on 'dvto' previously
c   discussed in class: see `phys410/f77/ex3/dvto.f
c=====
      subroutine dmto(fname,a,d1,d2)

c-----
c   Arguments:
c
c   fname: (I) File name
c   a: (I) Input matrix
c   d1: (I) First dimension of a
c   d2: (I) Second dimension of a
c-----

      implicit        none

      integer         indlnb, getu

      character*(*)   fname
      integer         d1, d2
      real*8         a(d1,d2)

      integer         ust dout
      parameter      ( ust dout = 6 )

      integer         uto, rc

c-----
c   Parse fname: either "attach" 'uto' to stdout or
c   get a unit number using 'getu', and open the
c   file 'fname' for formatted I/O via 'uto'
c-----
      if( fname .eq. '-' ) then
        uto = ust dout
      else
        uto = getu()
        open(uto,file=fname(1:indlnb(fname)),
&         form='formatted',iostat=rc)
&         if( rc .ne. 0 ) then
          write(0,*) 'dmto: Error opening ',
&                 fname(1:indlnb(fname))
&         return
        end if
      end if
      end if

```

```

c      Write dimensions, then array elements
c-----
      write(uto,*,iostat=rc) d1, d2
      if( rc .ne. 0 ) then
        write(0,*) 'dmto: Error writing dimensions'
        go to 500
      end if

      write(uto,*,iostat=rc) a
      if( rc .ne. 0 ) then
        write(0,*) 'dmto: Error reading matrix'
      end if

c-----
c      Exit: Close file and return
c-----
500    continue
      if( uto .ne. ustcout ) then
        close(uto)
      end if

      return

      end

c=====
c      Returns a double precision matrix (two dimensional
c      array) read from file 'fname'. If 'fname' is the
c      string '-', the matrix is read from standard input.
c
c      The dimensions of the matrix must precede the matrix
c      elements themselves in the file. Specifically, the
c      file should have been created using the following
c      list-directed, formatted READ statement
c      (or equivalent):
c
c      integer      d1,      d2
c      real*8       a(d1,d2)
c      integer      uout
c
c      write(uout,*) d1, d2
c      write(uout,*) a
c
c      This routine is modelled on 'dvfrom' previously
c      discussed in class: see 'phys410/f77/ex3/dvfrom.f'
c
c      Note the use of helper routine 'dmfrom1' which
c      reads actual array values once bounds have been
c      extracted from file.
c=====

      subroutine dmfrom(fname,a,d1,d2,asize)
c-----
c      Arguments:
c
c      fname:  (I)   File name
c      a:      (O)   Return matrix
c      d1:     (O)   First dimension of a
c      d2:     (O)   Second dimension of a
c      asize:  (I)   Maximum size (d1 * d2) of a
c-----
      implicit      none

      integer      indlnb,      getu

      character*(*) fname
      integer      d1,      d2,      asize
      real*8       a(d1,d2)

      integer      ustdin
      parameter    ( ustdin = 5 )

      integer      ufrom,      rc

c-----
c      Parse fname: either "attach" 'ufrom' to stdin or
c      get a unit number using 'getu', and open the
c      file 'fname' for formatted I/O via 'ufrom'
c-----
      if( fname .eq. '-' ) then
        ufrom = ustdin

```

Source file: tdm.f

```

c=====
c   Test program for subroutine 'dmfrom', 'dmto' and
c   'dmmult' (see 'dmroutines.f')
c
c   Program expects one argument, the name of a file which
c   contains a real*8 square matrix written as described
c   in the documentation for 'dmfrom' in 'dmroutines.f'
c   Use '-' to read from stdin. Program then computes
c   square of matrix and outputs result to stdout.
c=====

      program          tdm

      implicit        none

      integer         iargc

      character*256   fname

c-----
c   Maximum size for input and output arrays (matrices).
c-----
      integer         maxsize
      parameter      ( maxsize = 100 000 )
      real*8          a(maxsize), asq(maxsize)
      integer         dia,          d2a

      if( iargc() .ne. 1 ) go to 900
      call getarg(1,fname)

c-----
c   Read matrix ...
c-----
      call dmfrom(fname,a,dia,d2a,maxsize)
      if( dia .gt. 0 .and. d2a .gt. 0 ) then
         if( dia .eq. d2a ) then

c-----
c           Compute square ...
c-----
            call dmmult(a,a,asq,dia,dia)

c-----
c           ... and output.
c-----
            call dmto('-',asq,dia,dia)
         else
            write(0,*) 'tdm: Input array not square'
         end if
      else
         write(0,*) 'tdm: dmfrom() failed'
      end if

      stop

900 continue
      write(0,*) 'usage: tdm <file name>'
      write(0,*)
      write(0,*) '      Use ''tdm -'' to read ',
&
      write(0,*) '      'from standard input'

      stop

      end

```

Source file: tdm-output

```

#####
# Building 'tdm' and sample output
#####

lnx1 1> make tdm
pgf77 -g -c tdm.f
pgf77 -g -c dmroutines.f
pgf77 -g -L/usr/local/PGI/lib tdm.o dmroutines.o -lp410f -o tdm

lnx1 2> tdm
usage: tdm <file name>

      Use 'tdm -' to read from standard input

```

```

lnx1 3> tdm -
2 2
1 2 3 4
           2           2
       7.000000000000000    10.000000000000000    15.000000000000000
       22.000000000000000
lnx1 4> tdm -
2 3
1 2 3 4 5 6
tdm: Input array not square

```

Source file: Makefile

```

.IGNORE:

F77_COMPILE = $(F77) $(F77FLAGS) $(F77CFLAGS)
F77_LOAD     = $(F77) $(F77FLAGS) $(F77LFLAGS)

.f.o:
    $(F77_COMPILE) $.f

EXECUTABLES = meps catprec tdm

all: $(EXECUTABLES)

meps: meps.o
    $(F77_LOAD) meps.o -lp410f -o meps

catprec: catprec.o
    $(F77_LOAD) catprec.o -o catprec

tdm: tdm.o dmroutines.o
    $(F77_LOAD) tdm.o dmroutines.o -lp410f -o tdm

clean:
    rm *.o
    rm $(EXECUTABLES)
    rm core

```