

=====  
Contents  
=====

Installing the FM package

Troubleshooting

User's guide for the FM package

    Functions available

    Converting a Fortran program to use the FM package

    FMZM Interface Notes

        Initialization and internal names

        Conversion functions

        Inquiry functions

        Multiple precision versions of Fortran operations and functions

        Formatting functions

        Integer functions

        Special functions

        Operations defined using subroutines

        Array operations

        Array functions

FM.f95 Notes -- low-level routines

=====  
Installing the FM package  
=====

The files for version 1.3:

1. FMSAVE.f95           Module for FM internal global variables
2. FM.f95               Subroutine library for multiple-precision operations
3. FMZM90.f95          Modules for interfaces and definitions of derived-types

4. TestFM.f95            Test program for the FM routines
5. SampleFM.f95        Small sample program using FM

The first three files form the FM library routines that are used by an application program. They need to be compiled once, then any application program using FM can be compiled and linked to these library object files.

TestFM.f95 is a large program that does thousands of tests and calls all the library routines. It checks each operation and should be run once after compiling the FM library routines, to make sure FM is properly installed.

SampleFM.f95 is a small program with several examples showing how a typical user's program would use FM for some multiple precision real, integer, and complex calculations. It should be helpful as a model for getting started with FM.

Here are two example sets of compiler/linker commands for building the programs using the gfortran compiler (free -- click on the download link "Binaries for Windows, Linux, MacOS and much more" from this page: [gcc.gnu.org/wiki/GFortran](http://gcc.gnu.org/wiki/GFortran)).

The first three files are compiled as object code libraries FMSAVE.o, FM.o, FMZM90.o, and then each program that uses FM is compiled and linked to those three libraries. Some compilers name these object files FMSAVE.obj, etc.

Most compilers also produce files FMVALS.mod, FMZM.mod, etc., containing module information from the first three files.

1. For Windows, after installing the compiler, run in a PowerShell command prompt window:

```
gfortran fmsave.f95 -c -O3
```

```
gfortran FM.f95 -c -O3
```

```
gfortran FMZM90.f95 -c -O3
```

```
gfortran TestFM.f95 -c -O3
```

```
gfortran fmsave.o FM.o FMZM90.o TestFM.o -o TestFM.exe
```

```
./TestFM
```

```
gfortran SampleFM.f95 -c -O3
```

```
gfortran fmsave.o FM.o FMZM90.o SampleFM.o -o SampleFM.exe
```

```
./SampleFM
```

2. For a Mac, after installing the compiler, run in a Terminal window:

```
gfortran fmsave.f95 -c -O3

gfortran FM.f95 -c -O3

gfortran FMZM90.f95 -c -O3

gfortran TestFM.f95 -c -O3

gfortran fmsave.o FM.o FMZM90.o TestFM.o -o TestFM

./TestFM

gfortran SampleFM.f95 -c -O3

gfortran fmsave.o FM.o FMZM90.o SampleFM.o -o SampleFM

./SampleFM
```

The compiler options used in the examples were:

```
-c      compile to object code -- don't make executable
-O3     optimization level 3
-o      output file name for the executable program
```

Most other compilers use options very similar to these. The g95 and NAG compilers use the same commands ("g95 fmsave.f95 -c -O3" and "nagfor fmsave.f95 -c -O3" respectively), and Lahey's compiler uses

```
lf95 fmsave.f95 -ap -c -o1
...
lf95 fmsave.obj FM.obj FMZM90.obj TestFM.obj -out TestFM.exe
```

Older versions of Lahey's compiler needed the -ap option (see below) for all files -- the latest version I have tested is 7.3 in mid-2017.

Some compilers have options that might improve the speed of the program beyond the basic -O optimization. For example, "gfortran fm.f95 -c -O3 -funroll-loops". This runs slightly faster when precision is above 100 digits.

Some programs need 64-bit integers, and the easiest way to get them is often by using command-line compiler directives to change all integer constants and variables in a program.

Prior to June, 2015, the FM package was not compatible with 64-bit integers, since a few routines assumed that any integer value could be represented exactly when converted to double precision. But 64-bit integers can have more than 16 decimal digits, causing errors when converted to double precision. With the current version, FM can be used with either 32-bit or 64-bit integers.

With the gfortran compiler the command "gfortran prog.f95 -fdefault-integer-8 -o prog" will make 64-bit integers the default. When using 64-bit integers, all FM files and all files from the user's program should be compiled with the -fdefault-integer-8 option.

From the user's point of view, the only difference in results that come from the previous version and this one is that using 64-bit integers allows the range between FM's underflow and overflow thresholds to be greater.

In previous versions of FM,  $\exp(1.0e+8) = 1.5500e+43429448$  did not overflow, while  $\exp(1.0e+9)$  overflowed. Using 64-bit integers with this version,  $\exp(1.0e+15) = 6.7244e+434294481903251$  does not overflow, while  $\exp(1.0e+16)$  overflows.

FM has been run using many different compilers, both free and commercial. Most are used in a similar way to gfortran, although many can also be used with development environments that can make the process of compiling, linking, and executing the files even easier.

Two other files define optional multiprecision operations for exact rational arithmetic (fm\_rational.f95) and interval arithmetic (fm\_interval.f95). To use these, compile and link as with the examples above. See the FM web page for testing and sample programs for each of these, analogous to TestFM.f95 and SampleFM.f95.

---

---

### ----- Troubleshooting -----

After compiling and running the programs TestFM and SampleFM, each should say "no errors were found" at the end. If there were problems compiling the programs or some errors were found when they ran, read the rest of this section for possible fixes.

1. After downloading the files, if the compiler gives many error messages or it appears to see no code in the file at all, check that the lines in the file have the proper end-of-line characters for your system.

For a PC, this means each line should end with both a carriage return `<cr>` character (ascii 13) and a line feed `<lf>` character (ascii 10). If the file appears to be one huge line when viewed in an editor, one of these two characters is probably missing and should be added to each line.

To use FM on a Unix system, lines end with `<lf>`, and for a Mac system they might end with `<cr>`, but should end in `<lf>` because the Terminal window is really Unix. On these systems, failing to fix the end-of-line characters might mean the file seems to have twice the expected number of lines, with a blank line between each line of code when viewed in an editor, or the compiler might give an error message like "Unexpected end of file".

For a Mac with `<cr>` lines, a file like fmsave.f95 can be fixed by changing its name to fmsave0.f95 and then giving this command in Terminal:

```
tr \r \n < fmsave0.f95 > fmsave.f95
```

Then the new file fmsave.f95 should have the proper format.

Some good text editors will recognize a foreign end-of-line format and automatically fix each file the first time it is opened. Many recent Fortran compilers also automatically convert input source files to the proper type.

2. The compiler gives an "out of memory" error message or crashes during compile of one or more of the files.

It might be necessary to break the file into smaller pieces or split it into separate files for each routine or module. This could be caused by lack of system memory, lack of virtual memory, or a bug (memory leak) in the compiler.

Some compilers have an option (e.g., `-split`) to do this automatically.

This problem is less likely on more recent compilers and computers. With several recent compilers, each file can compile (barely) on machines with 1/2 Gb of memory, and should compile easily with 1 Gb or more.

3. Most of the routines compile, but a few fail with error messages like "symbol 120 is not the label of a branch target statement". However, looking at the code shows there is a label 120 in that routine.

This might happen in the larger routines. Some older compilers may require additional options be enabled (e.g., to force 32-bit branches or addresses to be used). Check in the compiler manual and try turning on any options that mention "long branches", "32-bit addresses", etc.

4. All files compile, but the TestFM program reports a few errors when it runs. There are other possibilities, but one thing to check is whether the compiler has any options controlling arithmetic precision of intermediate results.

Because the FM numbers are stored as integer values in double precision arrays, any sloppy rounding can cause problems. In one case, a compiler optimized an expression by leaving the result of a division in an 80-bit register and then used that result later in the calculation. Rounding the division back to double precision would have fixed the error, but using the inaccurate extended precision value caused the final result to be off by one when it was returned to an integer value.

This compiler had an option (`-ap`) to force intermediate results to not be left in registers, and that fixed the problem.

Another way to check to see if this is the problem is to create a version of FM that uses integer arrays instead of double precision. See the section titled "EFFICIENCY" below to see how to make this change. On most current machines, there is not a large speed penalty for using integer arrays as long as the precision is under 100 significant digits (i.e.,  $NDIG < 15$  or so with  $Mbase = 10^{**7}$ ).

5. TestFM or SampleFM gives an error message beginning something like this:

Element ( 1 ) of a multiple precision one-dimensional array is undefined in an expression.

This could happen with older compilers that support Fortran-90 but not Fortran-95. Version 1.3 of FM needs Fortran-95 or later. Try upgrading to the latest version of your compiler, or try the latest version of the free gfortran compiler to check to see if this is causing that error message.

The various lists of available multiple precision operations and routines have been collected here, along with some general advice on using the package.

See the program SampleFM.f95 for some examples of initializing and using the package.

-----  
 ----- Functions available -----  
 -----

The FM package provides 5 types of multiple precision operations:

1. type(FM) - floating-point
2. type(IM) - integer
3. type(ZM) - complex
4. type(FM\_RATIONAL) - rational
5. type(FM\_INTERVAL) - interval

The type definitions and interfaces for the first three are in file fmzm90.f95, and the other two are in files fm\_rational.f95 and fm\_interval.f95.

Some multiple precision functions take input arguments that are intrinsic types. In the table below, the types of arguments allowed for each function are abbreviated as:

FM, IM, ZM, RAT, IVL, for the 5 types above, and int, sp, dp, spz, dpz, str for the intrinsic Fortran types integer, single precision real, double precision real, single precision complex, double precision complex, character string.

Further description of many of these functions can be found later in this file.

ABS	FM	IM	ZM	RAT	IVL	
ACOS	FM		ZM		IVL	
ACOSH	FM		ZM		IVL	
AIMAG			ZM			
AINT	FM		ZM		IVL	
ANINT	FM		ZM		IVL	
ASIN	FM		ZM		IVL	
ASINH	FM		ZM		IVL	
ATAN	FM		ZM		IVL	
ATANH	FM		ZM		IVL	
ATAN2	FM				IVL	
BERNOULLI						int
BESSEL_J0	FM				IVL	
BESSEL_J1	FM				IVL	

BESSEL_JN	FM				IVL	
BESSEL_Y0	FM				IVL	
BESSEL_Y1	FM				IVL	
BESSEL_YN	FM				IVL	
BETA	FM				IVL	
BINOMIAL	FM	IM			IVL	int
BTEST		IM				
CEILING	FM	IM	ZM	RAT	IVL	
CPLX	FM	IM				
CONJG			ZM			
COS	FM		ZM		IVL	
COSH	FM		ZM		IVL	
COS_INTEGRAL	FM				IVL	
COSH_INTEGRAL	FM				IVL	
DBLE	FM	IM	ZM		IVL	
DIGITS	FM	IM	ZM		IVL	
DIM	FM	IM		RAT	IVL	
DINT	FM		ZM		IVL	
DOT_PRODUCT	FM	IM	ZM	RAT	IVL	
EPSILON	FM				IVL	
ERF	FM				IVL	
ERFC	FM				IVL	
ERFC_SCALED	FM					
EXP	FM		ZM		IVL	
EXPONENT	FM				IVL	
EXP_INTEGRAL_EI	FM				IVL	
EXP_INTEGRAL_EN	FM				IVL	
FACTORIAL	FM	IM			IVL	int
FLOOR	FM	IM	ZM	RAT	IVL	
FRACTION	FM		ZM		IVL	
FRESNEL_C	FM				IVL	
FRESNEL_S	FM				IVL	
GAMMA	FM				IVL	
GCD		IM				
HUGE	FM	IM	ZM		IVL	
HYPOT	FM					
INCOMPLETE_BETA	FM				IVL	
INCOMPLETE_GAMMA1	FM				IVL	
INCOMPLETE_GAMMA2	FM				IVL	
INT	FM	IM	ZM	RAT	IVL	
IS_OVERFLOW	FM	IM	ZM		IVL	
IS_UNDERFLOW	FM	IM	ZM		IVL	
IS_UNKNOWN	FM	IM	ZM	RAT	IVL	
LEFT_ENDPOINT					IVL	
LOG	FM		ZM		IVL	
LOG10	FM		ZM		IVL	
LOG_ERFC	FM				IVL	
LOG_GAMMA	FM				IVL	
LOG_INTEGRAL	FM				IVL	
MATMUL	FM	IM	ZM	RAT	IVL	
MAX	FM	IM		RAT	IVL	
MAXEXPONENT	FM				IVL	
MAXLOC	FM	IM	ZM	RAT		
MAXVAL	FM	IM	ZM	RAT	IVL	
MIN	FM	IM		RAT	IVL	
MINEXPONENT	FM				IVL	

MINLOC	FM	IM	ZM	RAT							
MINVAL	FM	IM	ZM	RAT	IVL						
MOD	FM	IM		RAT	IVL						
MODULO	FM	IM		RAT	IVL						
MULTIPLY_MOD		IM									
NEAREST	FM				IVL						
NINT	FM	IM	ZM	RAT	IVL						
NORM2	FM										
POCHHAMMER	FM				IVL						
POLYGAMMA	FM				IVL						
POWER_MOD		IM									
PRECISION	FM		ZM		IVL						
PRODUCT	FM	IM	ZM	RAT	IVL						
PSI	FM				IVL						
RADIX	FM	IM	ZM		IVL						
RANGE	FM	IM	ZM		IVL						
RATIONAL_DENOMINATOR				RAT							
RATIONAL_NUMERATOR				RAT							
REAL	FM	IM	ZM		IVL						
RIGHT_ENDPOINT					IVL						
RRSPACING	FM				IVL						
SCALE	FM		ZM		IVL						
SETEXPONENT	FM				IVL						
SIGN	FM	IM			IVL						
SIN	FM		ZM		IVL						
SINH	FM		ZM		IVL						
SIN_INTEGRAL	FM				IVL						
SINH_INTEGRAL	FM				IVL						
SPACING	FM				IVL						
SQRT	FM		ZM		IVL						
SUM	FM	IM	ZM	RAT	IVL						
TAN	FM		ZM		IVL						
TANH	FM		ZM		IVL						
TO_DP	FM	IM	ZM								
TO_DPZ	FM	IM	ZM								
TO_FM	FM	IM	ZM	RAT	IVL	int	sp	dp	spz	dpz	str
TO_FM_INTERVAL	FM	IM				int	sp	dp			str
TO_FM_RATIONAL		IM				int					str
TO_IM	FM	IM	ZM	RAT	IVL	int	sp	dp	spz	dpz	str
TO_INT	FM	IM	ZM								
TO_SP	FM	IM	ZM								
TO_SPZ	FM	IM	ZM								
TO_ZM	FM	IM	ZM	RAT	IVL	int	sp	dp	spz	dpz	str
TINY	FM	IM	ZM		IVL						
TRANSPOSE	FM	IM	ZM	RAT							

-----  
 ----- Converting a program to use the FM package -----  
 -----

0. If you want to write a program from scratch that uses FM, instead of converting an existing



double precision version, consider writing a d.p. version first anyway. It is very useful to have a working d.p. version to compare the FM results and quickly locate large errors that might be caused by mistakes in the conversion.

1. Before any variable declarations in each PROGRAM, SUBROUTINE, MODULE, or FUNCTION that will use multiple precision variables, insert

```
USE FMZM
```

This module contains all the rules needed by the compiler for doing the multiple precision operations.

2. In all routines using multiple precision variables,  
change real or double precision declarations to TYPE (FM)  
change complex or complex d.p. declarations to TYPE (ZM)  
if any integers need to be multiple precision, declare as TYPE (IM)

For example, if the original main program had these declarations,

```
REAL (KIND(1.0D0)) :: X, Y, A(50)  
COMPLEX (KIND(1.0D0)) :: C, Z(20)
```

change them to this for the FM version.

```
TYPE (FM) :: X, Y, A(50)  
TYPE (ZM) :: C, Z(20)
```

Local multiple precision variables in a subprogram are those that are not input arguments, module variables, or function names. These local variables should be saved, because a compiler might create a new instance of the local variable each time the routine is called. That would not cause an error, but it would leak memory and possibly cause the program to fail later for lack of memory. So if the original code was part of a subroutine,

```
SUBROUTINE SOLVE(A,Z,C)  
REAL (KIND(1.0D0)) :: X, Y, A(50)  
COMPLEX (KIND(1.0D0)) :: C, Z(20)
```

then the only local variables were X and Y, so the new declaration becomes

```
TYPE (FM) :: X, Y, A(50)  
TYPE (ZM) :: C, Z(20)  
SAVE :: X, Y
```

or equivalently,

```
TYPE (FM), SAVE :: X, Y  
TYPE (FM) :: A(50)  
TYPE (ZM) :: C, Z(20)
```

Using SAVE for variables in the main program is ok, but not required.

3. Variables that were initialized in the declarations of the original program must be initialized separately as FM variables.

```
DOUBLE PRECISION :: X = 1.2D0
```

becomes

```
TYPE (FM), SAVE :: X
...
X = TO_FM('1.2')
```

If this is in a subroutine and the value might change during one call and need to be remembered in a subsequent call, something like this can be done:

```
TYPE (FM), SAVE :: X
LOGICAL, SAVE :: FIRST_CALL = .TRUE.
...
IF (FIRST_CALL) THEN
  X = TO_FM('1.2')
  FIRST_CALL = .FALSE.
ENDIF
```

4. At the beginning of the main program, call FM\_SET to set the FM precision. For example, to get 50 significant digits,

```
CALL FM_SET(50)
```

Since increasing FM's internal precision level by one gives several extra base 10 significant digits, this call will actually set the user's precision to slightly more than 50 digits.

5. At the beginning of every function subprogram that returns a multiple precision value insert

```
CALL FM_ENTER_USER_FUNCTION(your_function_name)
```

where "your\_function\_name" is replaced by the actual name of the function.

Before each RETURN statement and before the END FUNCTION statement if it is not immediately preceded by a RETURN, insert

```
CALL FM_EXIT_USER_FUNCTION(your_function_name)
```

These are needed so that FM does not discard temporary variables too soon when they are created by the FMZM interface routines. The function name itself is used as a variable name to return the function value, so it is one of these temporary variables.

If a function subprogram uses multiple precision variables but returns a value that is not multiple precision, then the function name is not a temporary multiple precision object, but FM still needs to be notified not to discard temporaries within the function until that function is done. In this case, use

```
CALL FM_ENTER_USER_ROUTINE
```

and

```
CALL FM_EXIT_USER_ROUTINE
```

6. Check constants that are now part of multiple precision expressions and convert them.

`X = Y/3` need not be converted (since integers are exact in binary), but

`X = Y/3.7` should become `X = Y/TO_FM('3.7')`

Since 3.7 is not represented exactly in the machine's single or double precision, leaving the statement as `X = Y/3.7` would give X accurate only the machine's single precision, even though X and Y are multiple precision.

Also, constants in routine argument lists that now refer to multiple precision variables must be converted.

```
CALL SUB(A,B,2.6D0,X)
```

becomes

```
C = TO_FM('2.6')
CALL SUB(A,B,C,X)
```

Arithmetic expressions in subroutine calls should also be removed if they will use multiple precision in the new version, since they also create temporary FM variables.

```
CALL SUB(A,B,C+3*D,X)
```

becomes

```
T = C+3*D
CALL SUB(A,B,T,X)
```

If many such argument expressions occur in a program, an alternate way to convert them is by leaving the temporary object in the argument list, like

```
CALL SUB(A,B,TO_FM('2.6'),X)
CALL SUB(A,B,C+3*D,X)
```

and then inserting calls to `FM_ENTER_USER_ROUTINE` and `FM_EXIT_USER_ROUTINE` into subroutine `SUB` as described in section 5 above.

7. If a routine uses allocatable type FM, ZM, or IM arrays and allocates and deallocates with each call, then a call to `FM_DEALLOCATE` should be made before actually deallocating each array, to avoid leaking memory in FM's internal variable database. For example:

```
DEALLOCATE(T)
```

becomes:

```
CALL FM_DEALLOCATE(T)
DEALLOCATE(T)
```

8. Multiple precision variables in `WRITE` statements can be handled in several ways:

(a) If we use higher precision arithmetic for the calculations, but we only need to see the

final output at double precision, the simplest option is to convert the multiple precision variables back to double for printing. Then no changes to formats are needed.

```
WRITE (*,"(' Step size = ',F15.6,' tolerance = ',E15.7)'),H,T
```

becomes

```
WRITE (*,"(' Step size = ',F15.6,' tolerance = ',E15.7)'),TO_DP(H),TO_DP(T)
```

now that H and T are TYPE (FM) variables.

If the FM automatic tracing option is on (see NTRACE below), some Fortran-95 compilers might generate a "recursive write" error message here, since another write statement would be executed during the TO\_DP call. A fix is to turn the tracing off before this write statement. This can also happen if an FM error message is written during the TO\_DP call.

(b) Format the writes for multiple precision.

```
WRITE (*,"(' Step size = ',F15.6,' tolerance = ',E15.7)'),H,T
```

becomes

```
WRITE (*,"(' Step size = ',A,' tolerance = ',A)'), &  
      TRIM(FM_FORMAT('F15.6',H)),TRIM(FM_FORMAT('E15.7',T))
```

FM\_FORMAT is a formatting function used when the number of digits being shown is small enough to fit on one line.

Often after converting to multiple precision, we want to see more digits, so here F15.6 and E15.7 might become F35.20 and E35.25 in the FM version.

(c) Subroutine FM\_FORM does similar formatting, but we supply a character string for the formatted result. After declaring the strings at the top of the routine, as with

```
CHARACTER(80) :: ST1,ST2
```

the WRITE above could become

```
CALL FM_FORM('F15.6',H,ST1)  
CALL FM_FORM('E15.7',T,ST2)  
WRITE (*,"(' Step size = ',A,' tolerance = ',A)") TRIM(ST1),TRIM(ST2)
```

FM\_FORM must be used instead of FM\_FORMAT when there are more than 200 characters in the formatted string. These longer numbers usually need to be broken into several lines.

FM\_FORM should also be used when the FM trace option is on, since some compilers may generate an error message about a "recursive I/O reference" if a trace write executes from within another write statement via FM\_FORMAT.

(d) To use the current FM default format and handle any line breaks automatically, subroutine FM\_PRINT can be used. Calling FM\_SET to set precision at the beginning of the program also initializes this format. For example, FM\_PRINT displays 50 significant digits after CALL FM\_SET(50). See the discussion of FM's settings for JFORM1, JFORM2, and KSWIDE for changing the default format. The FM numbers will print on separate lines.

```

WRITE (*,*) ' Step size = '
CALL FM_PRINT(H)
WRITE (*,*) ' Tolerance = '
CALL FM_PRINT(T)

```

9. Multiple precision variables in READ statements can be done with FM's free-format input:

(a) Read the line as a character string then convert using TO\_FM.

```

READ (*,*) A,B,C

```

becomes this (with ST1 declared at the top of the routine as CHARACTER with length large enough to hold each input data line).

```

READ (*,"(A)") ST1
CALL FMSCAN(ST1,1,JA,JB)
A = TO_FM(ST1(JA:JB))
J1 = JB
CALL FMSCAN(ST1,J1,JA,JB)
B = TO_FM(ST1(JA:JB))
J1 = JB
CALL FMSCAN(ST1,J1,JA,JB)
C = TO_FM(ST1(JA:JB))

```

Where FMSCAN is defined by:

```

SUBROUTINE FMSCAN(STRING,JSTART,JA,JB)
! Scan STRING from position JSTART and return JA as the next non-blank and
! JB as the next blank after JA.
CHARACTER (*) :: STRING
JA = 0
JB = 0
DO J = JSTART, LEN(STRING)
  IF (JA == 0) THEN
    IF (STRING(J:J) /= ' ') JA = J
  ELSE
    IF (STRING(J:J) == ' ') THEN
      JB = J
    RETURN
  ENDOF
ENDIF
ENDDO
END SUBROUTINE FMSCAN

```

This assumes all three numbers are on one line. If they could appear on two or three lines, more code would be needed to check for that.

It also assumes that blanks separate the numbers. If input records use commas to separate numbers, repeat counts on input items, or slashes, then either the code above can be made more elaborate to handle those cases, or the data file can be edited so the simpler code works.

(b) Formatted reads can be converted directly to calls to TO\_FM without scanning to find where each number appears on the line.

```
READ (*,"(F20.15,E26.16,E20.10)") A,B,C
```

becomes this

```
READ (*,"(A)") ST1
A = TO_FM(ST1( 1:20))
B = TO_FM(ST1(21:46))
C = TO_FM(ST1(47:66))
```

(c) Declare double precision variables so the original read statement still works, then convert to multiple precision.

```
READ (*,*) A,B,C
```

becomes this

```
READ (*,*) A_DP,B_DP,C_DP
A = A_DP
B = B_DP
C = C_DP
```

where A\_DP,B\_DP,C\_DP are double precision and A,B,C are multiple precision.

A possible drawback to this method is that the values are read as double precision, so after conversion to FM they are usually accurate only to double precision. As with the TO\_FM example in section 6 above, a number such as 3.7 is not exactly representable in binary, so if that is the value being read for A in this example, reading it as a string in (b) and then converting the string gives full multiple precision accuracy for 3.7, but reading it as in (c) gives double precision accuracy.

See the page from the link "Automatic conversion of Fortran programs to use FM" on the main FM web page for a program that tries to automate most of these changes needed to convert a normal Fortran program to a multiple-precision version using FM.

---

---

FMZM Interface Notes

---

---

The FMZM module extends the definition of the basic Fortran arithmetic and function operations so they also apply to multiple precision numbers.

There are three multiple precision data types:

- FM (multiple precision real)
- IM (multiple precision integer)
- ZM (multiple precision complex)

A routine using any of these types needs this statement at the top:

```
USE FMZM
```

For some examples and general advice about using these multiple-precision data types, see the program SampleFM.f95.

Most of the functions defined in the FMZM module are multiple precision versions of standard Fortran functions. In addition, there are functions for direct conversion, formatting, and some mathematical special functions.

An attempt to use a multiple precision variable that has not been defined will be detected by the routines in FMZM and an error message printed.

----- Initialization and internal names -----

Initialization: The default precision for the multiple-precision numbers is about 50 significant digits.

To set precision to a different value, put this

```
CALL FM_SET(N)
```

in the main program before any multiple precision operations are done, with N replaced by the number of decimal digits of accuracy to be used.

Routine names: For each multiple precision operation there are several routines with related names that perform variations of that operation. For example, the addition operation has these forms:

Using the FMZM interface module to perform real (floating-point) multiple precision addition, declare the variables as FM derived types with

```
USE FMZM  
TYPE ( FM ) A,B,C
```

and then after values are assigned to A and B, doing a multiple precision addition looks the same as if the variables were real or double.

```
C = A + B
```

Normally, using the interface module avoids the need to know the name of the FM routine being called. For some operations, usually those that are not numerical Fortran functions (such as formatting a number), a direct call may be needed. For the addition above there is no reason to write it as a call in the user's program, but it could be done as

```
CALL FM_ADD(A,B,C)
```

Routines with names starting with FM\_ in the FMZM module (file FMZM90.f95) then call the low-level arithmetic routines in file FM.f95. The low-level routines are not usually called directly by the user's program, since they do not operate on the derived type variables that the user sees, but on the internal components of the types.

The low-level routines in FM.f95 usually have similar names to those in FMZM90.f95, but with no underscore after the first two letters. In this case the low-level routine is named FMADD.

For a few routines that don't have multiple precision arguments, like FM\_SET and FM\_SETVAR, the corresponding low-level names FMSET AND FMSETVAR are also available, and either form can be used.

----- Conversion functions -----

TO\_FM is a function for converting other types of numbers to type FM. Note that TO\_FM(3.12) converts the REAL constant to FM, but it is accurate only to single precision, since the number 3.12 cannot be represented exactly in binary and has already been rounded to single precision. Similarly, TO\_FM(3.12D0) agrees with 3.12 to double precision accuracy, and TO\_FM('3.12') or TO\_FM(312)/TO\_FM(100) agrees to full FM accuracy.

TO\_IM converts to type IM, and TO\_ZM converts to type ZM.

Functions are also supplied for converting the three multiple precision types to the other numeric data types:

TO\_INT converts to machine precision integer  
TO\_SP converts to single precision  
TO\_DP converts to double precision  
TO\_SPZ converts to single precision complex  
TO\_DPZ converts to double precision complex

WARNING: When multiple precision type declarations are inserted in an existing program, take care in converting functions like DBLE(X), where X has been declared as a multiple precision type. If X was single precision in the original program, then replacing the DBLE(X) by TO\_DP(X) in the new version could lose accuracy. For this reason, the Fortran type-conversion functions defined in the module assume that results should be multiple precision whenever inputs are. Examples:

DBLE(TO\_FM('1.23E+123456')) is type FM  
REAL(TO\_FM('1.23E+123456')) is type FM  
REAL(TO\_ZM('3.12+4.56i')) is type FM = TO\_FM('3.12')  
INT(TO\_FM('1.23')) is type IM = TO\_IM(1)  
INT(TO\_IM('1E+23')) is type IM  
CMPLX(TO\_FM('1.23'),TO\_FM('4.56')) is type ZM

----- Inquiry functions -----

IS\_OVERFLOW, IS\_UNDERFLOW, and IS\_UNKNOWN are logical functions for checking whether a multiple precision number is in one of the exception categories. Testing to see if a type FM number is in the +overflow category by directly using an IF can be tricky. When X is +overflow, the statement

```
IF (X == TO_FM(' +OVERFLOW ')) THEN
```

will return false, since the comparison routine cannot be sure that two different overflowed results would have been equal if the overflow threshold had been higher. Instead, use

```
IF (IS_OVERFLOW(X)) THEN
```



which will be true if X is + or - overflow.

----- Multiple precision operations and functions -----

For each of the operations =, ==, /=, <, <=, >, >=, +, -, \*, /, and \*\*, the FMZM interface module defines all mixed mode variations involving one of the three multiple precision derived types and another argument having one of these types:

{ integer, real, double, complex, complex double, FM, IM, ZM }

So mixed mode expressions such as

```
X = 12
X = X + 1
IF (ABS(X) > 1.0D-23) THEN
```

are handled correctly.

Not all the named functions are defined for all three multiple precision derived types, so the list below shows which can be used. The labels "real", "integer", and "complex" refer to types FM, IM, and ZM respectively, "string" means the function accepts character strings (e.g., TO\_FM('3.45')), and "other" means the function can accept any of the machine precision data types integer, real, double, complex, or complex double. For functions that accept two or more arguments, like ATAN2 or MAX, all the arguments must be of the same type.

TO\_ZM also has a 2-argument form: TO\_ZM(2,3) for getting 2 + 3\*i. CMPLX can be used for that, as in CMPLX( TO\_FM(2) , TO\_FM(3) ), and so can the string form, TO\_ZM(' 2 + 3 i '), but the 2-argument form is sometimes more convenient. The 2-argument form is available for machine precision integer, single and double precision real pairs. For others, such as X and Y being type(fm), just use CMPLX(X,Y).

----- Multiple precision versions of Fortran operations and functions -----

```
=
+
-
*
/
**
==
/=
<
<=
>
>=
ABS      real      integer    complex
ACOS     real
```

ACOSH	real		complex
AIMAG			complex
AIN	real		complex
ANINT	real		complex
ASIN	real		complex
ASINH	real		complex
ATAN	real		complex
ATANH	real		complex
ATAN2	real		
BTEST		integer	
CEILING	real	integer	complex
CMLX	real	integer	
CONJG			complex
COS	real		complex
COSH	real		complex
DBLE	real	integer	complex
DIGITS	real	integer	complex
DIM	real	integer	
DINT	real		complex
EPSILON	real		
EXP	real		complex
EXPONENT	real		
FLOOR	real	integer	complex
FRACTION	real		complex
HUGE	real	integer	complex
HYPOT	real		
INT	real	integer	complex
LOG	real		complex
LOG10	real		complex
MAX	real	integer	
MAXEXPONENT	real		
MIN	real	integer	
MINEXPONENT	real		
MOD	real	integer	
MODULO	real	integer	
NEAREST	real		
NINT	real	integer	complex
NORM2	real		
PRECISION	real		complex
RADIX	real	integer	complex
RANGE	real	integer	complex
REAL	real	integer	complex
RRSPACING	real		
SCALE	real		complex
SETEXPONENT	real		
SIGN	real	integer	
SIN	real		complex
SINH	real		complex
SPACING	real		
SQRT	real		complex
TAN	real		complex
TANH	real		complex
TINY	real	integer	complex

TO_FM	real	integer	complex	string	other
TO_IM	real	integer	complex	string	other
TO_ZM	real	integer	complex	string	other
TO_INT	real	integer	complex		
TO_SP	real	integer	complex		
TO_DP	real	integer	complex		
TO_SPZ	real	integer	complex		
TO_DPZ	real	integer	complex		
IS_OVERFLOW	real	integer	complex		
IS_UNDERFLOW	real	integer	complex		
IS_UNKNOWN	real	integer	complex		

----- Formatting functions -----

FM_FORMAT	real		
IM_FORMAT		integer	
ZM_FORMAT			complex

----- Integer functions -----

BINOMIAL	integer
FACTORIAL	integer
GCD	integer
MULTIPLY_MOD	integer
POWER_MOD	integer

----- Special functions -----

BERNOULLI(N)	real
BESSEL_J0(X)	real
BESSEL_J1(X)	real
BESSEL_JN(N,X)	real
BESSEL_JN(N1,N2,X)	real
BESSEL_Y0(X)	real
BESSEL_Y1(X)	real
BESSEL_YN(N,X)	real
BESSEL_YN(N1,N2,X)	real
BETA(A,B)	real
BINOMIAL(A,B)	real
COS_INTEGRAL(X)	real
COSH_INTEGRAL(X)	real
ERF(X)	real
ERFC(X)	real
ERFC_SCALED(X)	real
EXP_INTEGRAL_EI(X)	real
EXP_INTEGRAL_EN(N,X)	real
FACTORIAL(X)	real
FRESNEL_C(X)	real

FRESNEL_S(X)	real
GAMMA(X)	real
INCOMPLETE_BETA(X,A,B)	real
INCOMPLETE_GAMMA1(A,X)	real
INCOMPLETE_GAMMA2(A,X)	real
LOG_ERFC(X)	real
LOG_GAMMA(X)	real
LOG_INTEGRAL(X)	real
POCHHAMMER(X,N)	real
POLYGAMMA(N,X)	real
PSI(X)	real
SIN_INTEGRAL(X)	real
SINH_INTEGRAL(X)	real

Several of these functions are described in more detail below.

----- Subroutines that do not correspond to any function above -----

1. Type (FM). MA, MB, MC refer to type (FM) numbers.

These are subroutines instead of functions, so they are invoked as with  
 CALL FM\_COS\_SIN(MA,MB,MC)

FM\_COS\_SIN(MA,MB,MC)      MB = COS(MA),    MC = SIN(MA)  
 Faster than making two separate calls.

FM\_COSH\_SINH(MA,MB,MC)    MB = COSH(MA),    MC = SINH(MA)  
 Faster than making two separate calls.

FM\_EULER(MA)              MA = Euler's constant ( 0.5772156649... )

FM\_EQU(MA,MB,NA,NB)      MB = MA      where precision is being changed.  
                                  MA is defined with NDIG = NA digits and  
                                  MB will be defined having NB digits.  
                                  MB is rounded if NB < NA  
                                  MB is zero-padded if NB > NA

FM\_FLAG(K)                K = KFLAG    get the value of the FM condition flag -- stored in  
                                  the internal FM variable KFLAG in module FMVALS.

FM\_FORM(FORM,MA,STRING)    MA is converted to a character string using format FORM and  
                                  returned in STRING.    FORM can represent I, F, E, or ES formats.  
                                  Example:  
                                  CALL FMFORM('F60.40',MA,STRING)

FM\_FPRINT(FORM,MA)        Print MA on unit KW using FORM format.

FM\_PI(MA)                 MA = pi

FM\_PRINT(MA)              Print MA on unit KW using the current default format.

FM\_RANDOM\_NUMBER(X)      X is returned as a double precision random number, uniformly  
                                  distributed on the open interval (0,1). It is a high-quality,