



Micromega Corporation


Using uM-FPU V3.1 with the PICAXE Microcontroller

Introduction

The uM-FPU V3.1 chip is a 32-bit floating point coprocessor that easily connects to the PICAXE family of microcontrollers using an I²C interface. It provides extensive support for 32-bit IEEE 754 floating point operations and 32-bit long integer operations.

This document describes how to use the uM-FPU V3.1 chip with the PICAXE Microcontroller. For a full description of the uM-FPU V3.1 chip, please refer to the *uM-FPU V3.1 Datasheet* and *uM-FPU V3.1 Instruction Reference*. Application notes are also available on the Micromega website.

uM-FPU V3.1 Pin Diagram and Pin Description



Pin	Name	Type	Description
1	/MCLR	Input	Master Clear (Reset)
2	AN0	Input	Analog Input 0
3	AN1	Input	Analog Input 1
4	CS	Input	Chip Select, Interface Select
5	EXTIN	Input	External Input
6	OSC1	Input	Oscillator Crystal (optional)
7	OSC2	Output	Oscillator Crystal (optional)
8	TSTOUT	Output	Test Output, Debug Monitor - Tx
9	TSTIN	Input	Test Input, Debug Monitor - Rx
10	OUT1	Output	Test Point 1
11	SOUT SCL	Output Input	SPI Output, Busy/Ready Status I ² C Clock


12	SIN SDA	Input In/Out	SPI Input I ² C Data
13	VSS	Power	Digital Ground
14	VDD	Power	Digital Supply Voltage
15	OUT0	Output	Test Point 0
16	SCLK	Input	SPI Clock
17	AVSS	Power	Analog Ground
18	AVDD	Power	Analog Supply Voltage

Connecting uM-FPU V3.1 to the PICAXE using I²C

The default slave address for the uM-FPU V3.1 chip is 0xC8 (LSB is the R/W bit, e.g. 0xC8 for write, 0xC9 for read). See the uM-FPU datasheet for further description of the I²C interface. See the PICAXE documentation to determine the location of the I²C pins for each different microcontroller.

PICAXE-18X I²C SDA Output 1
 I²C SCL Output 4

PICAXE Pins



Brief Overview of the uM-FPU V3.1 chip

For a full description of the uM-FPU V3.1 chip, please refer to the *uM-FPU V3.1 Datasheet, uM-FPU V3.1 Instruction Reference*. Application notes are also available on the Micromega website.

The uM-FPU V3.1 chip is a separate coprocessor with its own set of registers and instructions designed to provide microcontrollers with 32-bit floating point and long integer capabilities. The PICAXE communicates with the FPU using a SPI or I²C interface. Instructions and data are sent to the FPU, and the FPU performs the calculations. The PICAXE is free to do other tasks while the FPU performs calculations. Results can be read back to the PICAXE or stored on the FPU for later use. The uM-FPU V3.1 chip has 128 registers, numbered 0 through 127, that can hold 32-bit floating point or long integer values. Register 0 is often used as a temporary register and is modified by some of the FPU instructions. Registers 1 through 127 are available for general use.

The **SELECTA** instruction is used to select any one of the 128 registers as register A. Register A can be regarded as an accumulator or working register. Arithmetic instructions use the value in register A as an operand and store the result of the operation in register A. If an instruction requires more than one operand, the additional operand is specified by the instruction. The following example selects register 2 as register A and adds register 5 to it:

```
SELECTA, 2      select register 2 as register A
FSET, 5         register[A] = register[A] + register[5]
```

Sending Instructions to the uM-FPU

Appendix A contains a table that gives a summary of each uM-FPU V3.1 instruction, with enough information to follow the examples in this document. For a detailed description of each instruction, refer to the *uM-FPU V3.1 Instruction Reference*.

To send instructions to the FPU using the I²C interface, the `hi2out` command is used as follows:

```
hi2out 0, (FADD, 5)
```

The part inside the parentheses specifies the instructions and data to send to the FPU. The zero before the parentheses is the address of the FPU data buffer. The `hi2out` command sends 8-bit data. To send a word variable, the high byte is sent first, followed by the low byte.

All instructions have an opcode that tells the FPU which operation to perform. The following example calculates the square root of register A:

```
hi2out 0, (SQRT)
```

Some instructions require additional operands or data and are specified by the bytes following the opcode. The following example adds register 5 to register A.

```
hi2out 0, (FADD, 5)
```

Some instructions return data. The `hi2in` command is used to read 8-bit data. This example reads the lower 8 bits of register A:

```
gosub fpu_wait
hi2out 0, (LREADBYTE)
hi2in 0, (dataByte)
```

The following example adds the value in register 5 to the value in register 2.

```
hi2out 0, (SELECTA, 2, FADD, 5)
```

It's a good idea to use constant definitions to provide meaningful names for the registers. This makes your program code easier to read and understand. The same example using constant definitions would be:

```
symbol    Total      = 2    ' total amount (uM-FPU register)
symbol    Count      = 5    ' current count (uM-FPU register)

hi2out 0, (SELECTA, Total, FADD, Count)
```

Tutorial Examples

Now that we've introduced some of the basic concepts of sending instructions to the FPU, let's go through a tutorial example to get a better understanding of how it all ties together. This example takes a temperature reading from a DS1620 digital thermometer and converts it to Celsius and Fahrenheit.

Most of the data read from devices connected to the PICAXE will return some type of integer value. In this example, the interface routine for the DS1620 reads a 9-bit value and stores it in a word variable on the PICAXE called `rawTemp`. The value returned by the DS1620 is the temperature in units of 1/2 degrees Celsius. The following instructions load the `rawTemp` value to the FPU, converts it to floating point, then divides it by 2 to get degrees in Celsius. (The high byte of `rawTemp` is defined as `rawHigh` and the low byte is `rawLow`.)

```
hi2out 0, (SELECTA, DegC, LOADWORD, rawHigh, rawLow, FSET0, FDIVI, 2)
```

Description:

SELECTA, DegC	select DegC as register A
LOADWORD, rawHigh, rawLow	load rawTemp to register 0 and convert to floating point
FSET0	DegC = register[0] (i.e. the floating point value of rawTemp)
FDIVI, 2	divide by the floating point value 2.0

To get the degrees in Fahrenheit we use the formula $F = C * 1.8 + 32$. Since 1.8 is a floating point constant, it would normally be loaded once in the initialization section of the program and used later in the program. The value 1.8 can be loaded using the ATOF (ASCII to float) instruction as follows:

```
hi2out 0, (SELECTA, F1_8, ATOF, "1.8", 0, FSET0)
```

Description:

SELECTA, F1.8	select F1_8 as register A
ATOF, "1.8", 0	load the string 1.8 (note: the string must be zero terminated), convert the string to floating point, and store in register 0
FSET0	F1_8 = register[0] (i.e. 1.8)

We calculate the degrees in Fahrenheit ($DegF = DegC * 1.8 + 32$) as follows:

```
hi2out 0, (SELECTA, DegF, FSET, DegC, FMUL, F1_8, FADDI, 32)
```

Description:

SELECTA, DegF	select DegF as register A
FSET, DegC	DegF = DegC
FMUL, F1_8	DegF = DegF * 1.8
FADDI, 32	DegF = DegF + 32.0

Note: this tutorial example is intended to show how to perform a familiar calculation, but the FCNV instruction could be used to perform unit conversions in one step. See the *uM-FPU V3.1 Instruction Reference* for a full list of conversions.

There are support routines provided for printing floating point and long integer numbers. `Print_float` prints an unformatted floating point value and displays up to eight digits of precision. `Print_floatFormat` prints a formatted floating point number. We'll use `Print_floatFormat` to display the results. The `format` variable is used to select the desired format, with the tens digit specifying the total number of characters to display, and the ones digit specifying the number of digits after the decimal point. The DS1620 has a maximum temperature of 125°

Celsius and one decimal point of precision, so we'll use a format of 51. Before calling the print routine the FPU register is selected and the `format` variable is set. The following example prints the temperature in degrees Celsius and Fahrenheit.

```
hi2out 0, (SELECTA, DegC)
format = 51
gosub print_floatFormat

hi2out 0, (SELECTA, DegF)
format = 51
gosub print_floatFormat
```

Sample code for this tutorial and a wiring diagram for the DS1620 are shown at the end of this document. The file *demo1.bas* is also included with the support software. There is a second file called *demo2.bas* that extends this demo to include minimum and maximum temperature calculations. If you have a DS1620 you can wire up the circuit and try out the demos.

uM-FPU Support Software for the PICAXE Microcontroller

A template file contains all of the definitions and support code required for communicating with the uM-FPU V3.1 chip:

template.bas provides support for an I²C connection to the uM-FPU V3.1 chip

This file can be used directly as the starting point for a new program, or the definitions and support code can be copied to another program. It contains the following:

- pin definitions for the uM-FPU V3.1
- opcode definitions for all uM-FPU V3.1 instructions
- various definitions for the byte variable used by the support routines
- a sample program with a place to insert your application code
- the support routines described below

fpu_reset

To ensure that the PICAXE and the FPU coprocessor are synchronized, a reset call must be done at the start of every program. The `fpu_reset` routine resets the FPU, confirms communications, and returns the sync character 0x5C if the reset is successful. A sample reset call is included in the *template.bas* file.

fpu_wait

The uM-FPU V3.1 chip must have completed all instructions in the instruction buffer, and be ready to return data, before sending an instruction to read data from the FPU. The `Fpu_wait` routine checks the status of the FPU and waits until it is ready. The print routines check the ready status, so it isn't necessary to call `Fpu_wait` before calling a print routine, but if your program reads directly from the FPU using the `hi2in` command, a call to `Fpu_wait` must be made prior to sending the read instruction. An example of reading a byte value is as follows:

```
gosub fpu_wait
hi2out 0, (LREADBYTE)
hi2in 0, (dataByte)
```

Description:

- wait for the FPU to be ready
- send the LREADBYTE instruction
- read a byte value

The uM-FPU V3.1 chip has a 256 byte instruction buffer. In most cases, data will be read back before 256 bytes have been sent to the FPU. If a long calculation is done that requires more than 256 bytes to be sent to the FPU, an `Fpu_wait` call should be made at least every 256 bytes to ensure that the instruction buffer doesn't overflow.

fpu_readStatus

The current status byte is read from the FPU and returned in the `statusByte` variable.

print_float

The value in register A is displayed on the PC screen as a floating point value using the `sertxd` command. Up to eight significant digits will be displayed if required. Very large or very small numbers are displayed in exponential notation. The length of the displayed value is variable and can be from 3 to 12 characters in length. The special cases of NaN (Not a Number), +Infinity, -Infinity, and -0.0 are handled. Examples of the display format are as follows:

1.0	NaN	0.0
1.5e20	Infinity	-0.0
3.1415927	-Infinity	1.0

-52.333334 -3.5e-5 0.01

print_floatFormat

The value in register A is displayed on the PC screen as a formatted floating point value using the `sertxd` command. The `format` byte is used to specify the desired format. The tens digit specifies the total number of characters to display and the ones digit specifies the number of digits after the decimal point. If the value is too large for the format specified, then asterisks will be displayed. If the number of digits after the decimal points is zero, no decimal point will be displayed. Examples of the display format are as follows:

Value in A register	format	Display format
123.567	61 (6.1)	123.6
123.567	62 (6.2)	123.57
123.567	42 (4.2)	*.*
0.9999	20 (2.0)	1
0.9999	31 (3.1)	1.0

print_long

The value in register A is displayed on the PC screen as a signed long integer using the `sertxd` command. The displayed value can range from 1 to 11 characters in length. Examples of the display format are as follows:

1
500000
-3598390

print_longFormat

The value in register A is displayed on the PC screen as a formatted long integer using the `sertxd` command. The `format` byte is used to specify the desired format. A value between 0 and 15 specifies the width of the display field for a signed long integer. The number is displayed right justified. If 100 is added to the format value the value is displayed as an unsigned long integer. If the value is larger than the specified width, asterisks will be displayed. If the width is specified as zero, the length will be variable. Examples of the display format are as follows:

Value in register A	format	Display format
-1	10 (signed 10)	-1
-1	110 (unsigned 10)	4294967295
-1	4 (signed 4)	-1
-1	104 (unsigned 4)	****
0	4 (signed 4)	0
0	0 (unformatted)	0
1000	6 (signed 6)	1000

print_version

The uM-FPU V3.1 version string is displayed on the PC screen using the `sertxd` command.

print_fpuString

The contents of the uM-FPU V3.1 string buffer are displayed to the PC screen using the `sertxd` command. This call is made at the end of the `print_float`, `print_floatFormat`, `print_long` and `print_longFormat` routines so is not normally called directly by the user unless string instructions are being used directly.

Writing Data Values to the FPU

Most of the data read from devices connected to the PICAXE will return some type of integer value. There are several ways to load integer values to the FPU and convert them to 32-bit floating point or long integer values.

8-bit Integer to Floating Point

The `FSETI`, `FADDI`, `FSUBI`, `FSUBRI`, `FMULI`, `FDIVI`, `FDIVRI`, `FPOWI`, and `FCMPI` instructions read the byte following the opcode as an 8-bit signed integer, convert the value to floating point, and then perform the operation. It's a convenient way to work with constants or data values that are signed 8-bit values. The following example stores the byte variable `dataByte` to the `Result` register on the FPU.

```
hi2out 0, (SELECTA, Result, FSETI, dataByte)
```

The `LOADBYTE` instruction reads the byte following the opcode as an 8-bit signed integer, converts the value to floating point, and stores the result in register 0.

The `LOADUBYTE` instruction reads the byte following the opcode as an 8-bit unsigned integer, converts the value to floating point, and stores the result in register 0.

16-bit Integer to Floating Point

The `LOADWORD` instruction reads the two bytes following the opcode as a 16-bit signed integer (MSB first), converts the value to floating point, and stores the result in register 0. The following example adds the word variable `dataWord` (bytes `dataHigh` and `dataLow`) to the `Result` register on the FPU.

```
hi2out 0, (SELECTA, Result, LOADWORD, dataHigh, dataLow, FADD0)
```

The `LOADUWORD` instruction reads the two bytes following the opcode as a 16-bit unsigned integer (MSB first), converts the value to floating point, and stores the result in register 0.

Floating Point

The `FWRITE`, `FWRITEA`, `FWRITEEX`, and `FWRITE0` instructions read the four bytes following the opcode as a 32-bit floating point value and stores it in the specified register. This is one of the more efficient ways to load floating point constants, but requires knowledge of the internal representation for floating point numbers (see Appendix B). The *uM-FPU V3 IDE* can be used to easily generate the 32-bit values. This example sets `Angle = 20.0` (the floating point representation for 20.0 is hex 41A00000).

```
hi2out 0, (FWRITE, Angle, $41, $A0, $00, $00)
```

ASCII string to Floating Point

The `ATOF` instruction is used to convert strings to floating point values. The instruction reads the bytes following the opcode (until a zero terminator is read), converts the string to floating point, and stores the result in register 0. For example, to set `Angle = 1.5885`:

```
hi2out 0, (SELECTA, Angle, ATOF, "1.5885", 0, FSET0)
```

8-bit Integer to Long Integer

The `LSETI`, `LADDI`, `LSUBI`, `LMULI`, `LDIVI`, `LCMPI`, `LUDIVI`, `LUCMPI`, and `LTSTI` instructions read the byte following the opcode as an 8-bit signed integer, convert the value to long integer, and then perform the operation. It's a convenient way to work with constants or data values that are signed 8-bit values. The following example adds the byte variable `dataByte` to the `Total` register on the FPU.

```
hi2out 0, (SELECTA, Total, LADDI, dataByte)
```

The LONGBYTE instruction reads the byte following the opcode as an 8-bit signed integer, converts the value to long integer, and stores the result in register 0.

The LONGUBYTE instruction reads the byte following the opcode as an 8-bit unsigned integer, converts the value to long integer, and stores the result in register 0.

16-bit Integer to Long Integer

The LONGWORD instruction reads the two bytes following the opcode as a 16-bit signed integer (MSB first), converts the value to long integer, and stores the result in register 0. The following example adds the word variable `dataWord` to the `Total` register on the FPU.

```
hi2out 0, (SELECTA, Total, LONGWORD, dataHigh, dataLow, LADD0)
```

The LONGUWORD instruction reads the two bytes following the opcode as a 16-bit unsigned integer (MSB first), converts the value to long integer, and stores the result in register 0.

Long Integer

The LWRITE, LWRITEA, LWRITEX, and LWRITE0 instructions read the four bytes following the opcode as a 32-bit long integer value and stores it in the specified register. This is used to load integer values greater than 16 bits. The *uM-FPU V3 IDE* can be used to easily generate the 32-bit values. For example, to set `Total = 500000`:

```
hi2out 0, (LWRITE, Total, $00, $07, $A1, $20)
```

ASCII string to Long Integer

The ATOL instruction is used to convert strings to long integer values. The instruction reads the bytes following the opcode (until a zero terminator is read), converts the string to long integer, and stores the result in register 0. For example, to set `Total = 500000`:

```
hi2out 0, (SELECTA, Total, ATOL, "5000000", 0, FSET0)
```

The fastest operations occur when the FPU registers are already loaded with values. In time critical portions of code floating point constants should be loaded beforehand to maximize the processing speed in the critical section. With 128 registers available on the FPU, it's often possible to pre-load all of the required constants. In non-critical sections of code, data and constants can be loaded as required.

Reading Data Values from the FPU

The uM-FPU V3.1 chip has a 256 byte instruction buffer which allows data transmission to continue while previous instructions are being executed. Before reading data, you must check to ensure that the previous commands have completed, and the FPU is ready to send data. The `Fpu_wait` routine is used to wait until the FPU is ready, then a read command is sent, and the `hi2in` command is used to read data.

Floating Point

The FREAD, FREADA, FREADX, and FREAD0 instructions read four bytes from the FPU as a 32-bit floating point value. The following example reads the 32-bit floating point value from register A and stores it in byte variables `byte0`, `byte1`, `byte2`, and `byte3`.

```
gosub Fpu_wait
hi2out 0, (FREADA)
hi2in 0, (byte0, byte1, byte2, byte3)
```

Floating Point to ASCII string

The FTOA instruction can be used to convert floating point values to an ASCII string. The `print_float` and `print_floatFormat` routines use this instruction to read the floating point value from register A and display it on the PC screen.

8-bit Integer

The LREADBYTE instruction reads the lower 8 bits from register A. The following example stores the lower 8 bits of register A in byte variable `dataByte`.

```
gosub Fpu_wait
hi2out 0, (LREADBYTE)
hi2in 0, (dataByte)
```

16-bit Integer

The LREADWORD instruction reads the lower 16 bits from register A. The following example stores the lower 16 bits of register A in word variable `dataWord` (bytes `dataHigh` and `dataLow`).

```
gosub Fpu_wait
hi2out 0, (LREADWORD)
hi2in 0, (dataHigh, dataLow)
```

Long Integer

The LREAD, LREADA, LREADX, and LREAD0 instructions read four bytes from the FPU as a 32-bit long integer value. The following example reads the 32-bit long integer value from register A and stores it in byte variables `byte0`, `byte1`, `byte2`, and `byte3`.

```
gosub Fpu_wait
hi2out 0, (LREADA)
hi2in 0, (byte0, byte1, byte2, byte3)
```

Long Integer to ASCII string

The LTOA instruction can be used to convert long integer values to an ASCII string. The `print_long` and `print_longFormat` routines use this instruction to read the long integer value from register A and display it on the PC screen.

Comparing and Testing Floating Point Values

Floating point values can be zero, positive, negative, infinite, or Not a Number (which occurs if an invalid operation is performed on a floating point value). The status byte is read using the `Fpu_ReadStatus` routine. It waits for the FPU to be ready before sending the READSTATUS instruction and reading the status from the FPU. The current status is returned in the `statusByte` variable. Definitions for the status bits are provided as follows:

symbol	IS_ZERO	= 0x81	' positive zero
symbol	IS_NZERO	= 0x83	' negative zero
symbol	IS_NEGATIVE	= 0x82	' negative
symbol	IS_NAN	= 0x84	' NaN (Not-a-Number)
symbol	IS_PINF	= 0x88	' positive infinity
symbol	IS_NINF	= 0x8A	' negative infinity

The FSTATUS and FSTATUSA instructions are used to set the status byte to the floating point status of the selected register. The following example checks the floating point status of register A:

```
hi2out 0, (FSTATUSA)
gosub fpu_readStatus
```

```

IF statusByte = IS_ZERO or statusByte = IS_NZERO then zeroValue
IF statusByte = IS_NEGATIVE then negativeValue
    sertxd("value is positive")
...
negativeValue:
    sertxd("value is negative")
...
zeroValue:
    sertxd("value is zero")

```

The FCMP, FCMP0, and FCMP1 instructions are used to compare two floating point values. The status bits are set for the result of register A minus the operand (the selected registers are not modified). For example, to compare register A to the value 10.0:

```

hi2out 0, (FCMP1, 10)
gosub fpu_readStatus
IF statusByte = IS_ZERO then sameAs
IF statusByte = IS_NEGATIVE then lessThan
    sertxd("A > 10")
...
lessThan:
    sertxd("A < 10")
...
sameAs:
    sertxd("A = 10")

```

The FCMP2 instruction compares two floating point registers. The status bits are set for the result of the first register minus the second register (the selected registers are not modified). For example, to compare registers Value1 and Value2:

```

hi2out 0, (FCMP2, Value1, Value2)
gosub fpu_readStatus
...

```

Comparing and Testing Long Integer Values

A long integer value can be zero, positive, or negative. The status byte is read using the `Fpu_ReadStatus` routine. It waits for the FPU to be ready before sending the `READSTATUS` command and reading the status. The current status is returned in the `statusByte` variable. Definitions for the status bits are provided as follows:

```

symbol    IS_ZERO      = 0x81 ' zero
symbol    IS_NEGATIVE  = 0x82 ' negative

```

The LSTATUS and LSTATUSA instructions are used to set the status byte to the long integer status of the selected register. The following example checks the long integer status of register A:

```

hi2out 0, (LSTATUSA)
IF statusByte = IS_ZERO then zeroValue
IF statusByte = IS_NEGATIVE then negativeValue
    sertxd("value is positive")
...
negativeValue:
    sertxd("value is negative")
...
zeroValue:
    sertxd("value is zero")

```

The LCMP, LCMP0, and LCMP1 instructions are used to do a signed comparison of two long integer values. The status bits are set for the result of register A minus the operand (the selected registers are not modified). For example, to compare register A to the value 10:

```
hi2out 0, (LCMP1, 10)
gosub fpu_readStatus
IF statusByte = IS_ZERO then sameAs
IF statusByte = IS_NEGATIVE then lessThan
    sertxd("A > 10")
...
lessThan:
    sertxd("A < 10")
...
sameAs:
    sertxd("A = 10")
```

The LCMP2 instruction does a signed compare of two long integer registers. The status bits are set for the result of the first register minus the second register (the selected registers are not modified). For example, to compare registers Value1 and Value2:

```
hi2out 0, (LCMP2, Value1, Value2)
gosub fpu_readStatus
...
```

The LUCMP, LUCMP0, and LUCMP1 instructions are used to do an unsigned comparison of two long integer values. The status bits are set for the result of register A minus the operand (the selected registers are not modified).

The LUCMP2 instruction does an unsigned compare of two long integer registers. The status bits are set for the result of the first register minus the second register (the selected registers are not modified).

The LTST, LTST0 and LTST1 instructions are used to do a bit-wise compare of two long integer values. The status bits are set for the logical AND of register A and the operand (the selected registers are not modified).


Further Information

The following documents are also available:

<i>uM-FPU V3.1 Datasheet</i>	provides hardware details and specifications
<i>uM-FPU V3.1 Instruction Reference</i>	provides detailed descriptions of each instruction
<i>uM-FPU V3.1 Application Notes</i>	various application notes and examples

Check the Micromega website at www.micromegacorp.com for up-to-date information.

DS1620 Connections for Demo 1



Sample Code for Tutorial (Demo1.bas)

```
' This program demonstrates the use of the uM-FPU V3.1 floating point coprocessor
' with the PICAXE microcontroller using an I2C interface. It takes temperature
' readings from a DS1620 digital thermometer, converts them to floating point
' and displays them in degrees Celsius and degrees Fahrenheit.
```

```
'=====
'----- uM-FPU V3.1 I2C definitions ----- 2011-10-17
'=====
```

```
symbol fpuID          = 0xC8 ' uM-FPU I2C address
```

```
'----- uM-FPU V3.1 opcode definitions -----
```

```
symbol SELECTA       = 0x01 ' Select register A
symbol ATOF           = 0x1E ' Convert ASCII to float, store in reg[0]
symbol FTOA           = 0x1F ' Convert float to ASCII
symbol FSET           = 0x20 ' reg[A] = reg[nn]
symbol FMUL           = 0x24 ' reg[A] = reg[A] * reg[nn]
symbol FSET0          = 0x29 ' reg[A] = reg[0]
symbol FADDI          = 0x33 ' reg[A] = reg[A] + float(bb)
symbol FDIVI          = 0x37 ' reg[A] = reg[A] / float(bb)
symbol LOADWORD       = 0x5B ' reg[0] = float(signed word)
symbol LOADWORD       = 0x5B ' reg[0] = float(signed word)
symbol SYNC           = 0xF0 ' Get synchronization byte
symbol READSTATUS     = 0xF1 ' Read status byte
symbol READSTR        = 0xF2 ' Read string from string buffer
symbol VERSION        = 0xF3 ' Copy version string to string buffer
```

```
symbol SYNC_CHAR     = 0x5C ' sync character
```

```
'----- uM-FPU variables -----
```

```
symbol dataByte      = b13      ' data byte
symbol format        = b13      ' format (same as dataByte)
symbol statusByte    = b13      ' status byte (same as dataByte)
```

```
'===== end of uM-FPU V3.1 I2C definitions =====
```

```
'----- DS1620 pin definitions -----
```

```
symbol DS_RST        = output7   ' DS1620 reset/enable
symbol DS_CLK        = output6   ' DS1620 clock
symbol DS_DATAOUT    = output5   ' DS1620 data out
```

```

symbol DS_DATAIN    = input7      ' DS1620 data in

'----- uM-FPU register definitions -----

symbol DegC         = 1           ' degrees Celsius
symbol DegF         = 2           ' degrees Fahrenheit
symbol F1_8         = 3           ' constant 1.8

'----- variables -----

symbol rawTemp      = W0           ' raw temperature reading
symbol rawHigh      = B1           ' high byte of raw temperature
symbol rawLow       = B0           ' low byte of raw temperature
symbol bitcnt       = B2           ' bit count
symbol tmp          = B3           ' temporary variable

'=====
'----- main routine -----
'=====

main:
    pause 500                      ' wait for terminal window
    sertxd(13, 10, 13, 10, "demo1", 13, 10)

    gosub fpu_reset                 ' reset the uM-FPU
    if statusByte = SYNC_CHAR then main2
    sertxd (13, 10, "uM-FPU not detected.")
    end

main2:
    gosub print_version             ' display the uM-FPU version number

    ' initialize DS1620
    '-----
    gosub init_DS1620

    ' load floating point constant
    '-----
    hi2out 0, (SELECTA, F1_8, ATOF, "1.8", 0, FSET0)

displayLoop:
    ' get temperature reading from DS1620
    '-----
    gosub read_DS1620

    ' send rawTemp to uM-FPU V3.1
    ' convert to floating point
    ' divide by 2 to get degrees Celsius
    '-----
    hi2out 0, (SELECTA, DegC, LOADWORD, rawHigh, rawLow, FSET0, FDIVI, 2)

    ' degF = degC * 1.8 + 32
    '-----
    hi2out 0, (SELECTA, DegF, FSET, DegC, FMUL, F1_8, FADDI, 32)

    ' display degrees Celsius
    '-----

```

```

sertxd(13, 10, 13, 10, "Degrees C: ")
hi2out 0, (SELECTA, DegC)
format = 51
gosub print_floatFormat

'display degrees Fahrenheit
'-----
sertxd(13, 10, "Degrees F: ")
hi2out 0, (SELECTA, DegF)
format = 51
gosub print_floatFormat

'delay, then get the next reading
'-----
pause 2000
goto displayLoop
end

'----- init_DS1620 -----

init_DS1620:
    low DS_RST                ' initialize pin states
    high DS_CLK
    pause 100

    high DS_RST              ' configure for CPU control
    dataByte = $0C
    gosub write_DS1620
    dataByte = $02
    gosub write_DS1620
    low DS_RST
    pause 100

    high DS_RST              ' start temperature conversions
    dataByte = $EE
    gosub write_DS1620
    low DS_RST
    pause 1000              ' wait for first conversion
    return

'----- read_DS1620 -----

read_DS1620:
    high DS_RST              ' read temperature value
    dataByte = $AA
    gosub write_DS1620

    for bitcnt = 1 to 8      ' read byte from DS1620 (LSB first)
        low DS_CLK
        rawLow = rawLow / 2
        if DS_DATAIN = 0 then read2
        rawLow = rawLow + 128
read2:    high DS_CLK
        next bitcnt

    low DS_CLK
    rawHigh = 0              ' read 9th bit and extend sign
    if DS_DATAIN = 0 then read3

```



```

        rawHigh = $FF

read3:
    high DS_CLK
    low DS_RST
    return

'----- write_DS1620 -----

write_DS1620:
    for bitcnt = 1 to 8      ' write byte to DS1620 (LSB first)
        tmp = dataByte & 1
        low DS_DATAOUT
        if tmp = 0 then write2
        high DS_DATAOUT
write2: pulsout DS_CLK, 1    ' pulse clock for 10us
        dataByte = dataByte / 2
    next bitcnt
    return

'=====
'----- uM-FPU V3.1 I2C support routines ----- 2011-10-17
'=====

fpu_reset:
    hi2csetup i2cmaster, fpuID, i2cfast, i2cbyte
    hi2out 1, (0)           ' reset the uM-FPU
    pause 10                ' wait for reset to complete

    hi2out 0, (SYNC)       ' check for synchronization
    goto fpu_readStatus2

fpu_wait:
    hi2in 0, (statusByte)  ' wait for ready status
    if statusByte <> 0 then fpu_wait
    return

fpu_readStatus:
    gosub fpu_wait         ' read status byte
    hi2out 0, (READSTATUS)

fpu_readStatus2:
    hi2in 0, (statusByte)  ' read status byte
    return

print_version:
    hi2out 0, (VERSION)    ' get the uM-FPU version string
    goto print_fpuString   ' print it

print_float:
    format = 0             ' set for free format
                            ' (fall through to print_floatFormat)

print_floatFormat:
    hi2out 0, (FTOA, format) ' convert floating point to formatted ASCII
    goto print_fpuString   ' print the string

print_long:
    format = 0             ' set for free format

```

```

' (fall through to print_longFormat)
print_longFormat:
    hi2out 0, (LTOA, format) ' convert long integer to formatted ASCII
' (fall through to print_fpuString)
print_fpuString:
    gosub fpu_wait          ' wait until uM-FPU is ready
    hi2out 0, (READSTR)

print_string2:
    hi2in 0, (dataByte)    ' display zero terminated string
    if dataByte = 0 then Print_String3
    if dataByte > 127 then Print_String3
    sertxd(dataByte)
    goto Print_String2

print_string3:
    return                ' end of string

'===== end of uM-FPU V3.1 I2C support routines =====

```

Appendix A

uM-FPU V3.1 Instruction Summary

Instruction	Opcode	Arguments	Returns	Description
NOP	00			No Operation
SELECTA	01	nn		Select register A
SELECTX	02	nn		Select register X
CLR	03	nn		reg[nn] = 0
CLRA	04			reg[A] = 0
CLR _X	05			reg[X] = 0, X = X + 1
CLR ₀	06			reg[nn] = reg[0]
COPY	07	mm, nn		reg[nn] = reg[mm]
COPY _A	08	nn		reg[nn] = reg[A]
COPY _X	09	nn		reg[nn] = reg[X], X = X + 1
LOAD	0A	nn		reg[0] = reg[nn]
LOAD _A	0B			reg[0] = reg[A]
LOAD _X	0C			reg[0] = reg[X], X = X + 1
ALOAD _X	0D			reg[A] = reg[X], X = X + 1
XSAVE	0E	nn		reg[X] = reg[nn], X = X + 1
XSAVE _A	0F			reg[X] = reg[A], X = X + 1
COPY ₀	10	nn		reg[nn] = reg[0]
COPY _I	11	bb, nn		reg[nn] = long(unsigned byte bb)
SWAP	12	nn, mm		Swap reg[nn] and reg[mm]
SWAP _A	13	nn		Swap reg[nn] and reg[A]
LEFT	14			Left parenthesis
RIGHT	15			Right parenthesis
FWRITE	16	nn, b1, b2, b3, b4		Write 32-bit floating point to reg[nn]
FWRITE _A	17	b1, b2, b3, b4		Write 32-bit floating point to reg[A]
FWRITE _X	18	b1, b2, b3, b4		Write 32-bit floating point to reg[X]
FWRITE ₀	19	b1, b2, b3, b4		Write 32-bit floating point to reg[0]
FREAD	1A	nn	b1, b2, b3, b4	Read 32-bit floating point from reg[nn]
FREAD _A	1B		b1, b2, b3, b4	Read 32-bit floating point from reg[A]
FREAD _X	1C		b1, b2, b3, b4	Read 32-bit floating point from reg[X]
FREAD ₀	1C		b1, b2, b3, b4	Read 32-bit floating point from reg[0]
A _T O _F	1E	aa...00		Convert ASCII to floating point
F _T O _A	1F	bb		Convert floating point to ASCII
FSET	20	nn		reg[A] = reg[nn]
FADD	21	nn		reg[A] = reg[A] + reg[nn]
FSUB	22	nn		reg[A] = reg[A] - reg[nn]
FSUB _R	23	nn		reg[A] = reg[nn] - reg[A]
FMUL	24	nn		reg[A] = reg[A] * reg[nn]
FDIV	25	nn		reg[A] = reg[A] / reg[nn]
FDIV _R	26	nn		reg[A] = reg[nn] / reg[A]
FPOW	27	nn		reg[A] = reg[A] ** reg[nn]
FCMP	28	nn		Compare reg[A], reg[nn], Set floating point status
FSET ₀	29			reg[A] = reg[0]
FADD ₀	2A			reg[A] = reg[A] + reg[0]

FSUB0	2B			$\text{reg}[A] = \text{reg}[A] - \text{reg}[0]$
FSUBR0	2C			$\text{reg}[A] = \text{reg}[0] - \text{reg}[A]$
FMUL0	2D			$\text{reg}[A] = \text{reg}[A] * \text{reg}[0]$
FDIV0	2E			$\text{reg}[A] = \text{reg}[A] / \text{reg}[0]$
FDIVR0	2F			$\text{reg}[A] = \text{reg}[0] / \text{reg}[A]$
FPOW0	30			$\text{reg}[A] = \text{reg}[A] ** \text{reg}[0]$
FCMP0	31			Compare $\text{reg}[A]$, $\text{reg}[0]$, Set floating point status
FSETI	32	bb		$\text{reg}[A] = \text{float}(bb)$
FADDI	33	bb		$\text{reg}[A] = \text{reg}[A] - \text{float}(bb)$
FSUBI	34	bb		$\text{reg}[A] = \text{reg}[A] - \text{float}(bb)$
FSUBRI	35	bb		$\text{reg}[A] = \text{float}(bb) - \text{reg}[A]$
FMULI	36	bb		$\text{reg}[A] = \text{reg}[A] * \text{float}(bb)$
FDIVI	37	bb		$\text{reg}[A] = \text{reg}[A] / \text{float}(bb)$
FDIVRI	38	bb		$\text{reg}[A] = \text{float}(bb) / \text{reg}[A]$
FPOWI	39	bb		$\text{reg}[A] = \text{reg}[A] ** bb$
FCMPI	3A	bb		Compare $\text{reg}[A]$, $\text{float}(bb)$, Set floating point status
FSTATUS	3B	nn		Set floating point status for $\text{reg}[nn]$
FSTATUSA	3C			Set floating point status for $\text{reg}[A]$
FCMP2	3D	nn, mm		Compare $\text{reg}[nn]$, $\text{reg}[mm]$ Set floating point status
FNEG	3E			$\text{reg}[A] = -\text{reg}[A]$
FABS	3F			$\text{reg}[A] = \text{reg}[A] $
FINV	40			$\text{reg}[A] = 1 / \text{reg}[A]$
SQRT	41			$\text{reg}[A] = \text{sqrt}(\text{reg}[A])$
ROOT	42	nn		$\text{reg}[A] = \text{root}(\text{reg}[A], \text{reg}[nn])$
LOG	43			$\text{reg}[A] = \text{log}(\text{reg}[A])$
LOG10	44			$\text{reg}[A] = \text{log}_{10}(\text{reg}[A])$
EXP	45			$\text{reg}[A] = \text{exp}(\text{reg}[A])$
EXP10	46			$\text{reg}[A] = \text{exp}_{10}(\text{reg}[A])$
SIN	47			$\text{reg}[A] = \text{sin}(\text{reg}[A])$
COS	48			$\text{reg}[A] = \text{cos}(\text{reg}[A])$
TAN	49			$\text{reg}[A] = \text{tan}(\text{reg}[A])$
ASIN	4A			$\text{reg}[A] = \text{asin}(\text{reg}[A])$
ACOS	4B			$\text{reg}[A] = \text{acos}(\text{reg}[A])$
ATAN	4C			$\text{reg}[A] = \text{atan}(\text{reg}[A])$
ATAN2	4D	nn		$\text{reg}[A] = \text{atan}_{2}(\text{reg}[A], \text{reg}[nn])$
DEGREES	4E			$\text{reg}[A] = \text{degrees}(\text{reg}[A])$
RADIANS	4F			$\text{reg}[A] = \text{radians}(\text{reg}[A])$
FMOD	50	nn		$\text{reg}[A] = \text{reg}[A] \text{ MOD } \text{reg}[nn]$
FLOOR	51			$\text{reg}[A] = \text{floor}(\text{reg}[A])$
CEIL	52			$\text{reg}[A] = \text{ceil}(\text{reg}[A])$
ROUND	53			$\text{reg}[A] = \text{round}(\text{reg}[A])$
FMIN	54	nn		$\text{reg}[A] = \text{min}(\text{reg}[A], \text{reg}[nn])$
FMAX	55	nn		$\text{reg}[A] = \text{max}(\text{reg}[A], \text{reg}[nn])$
FCNV	56	bb		$\text{reg}[A] = \text{conversion}(bb, \text{reg}[A])$
FMAC	57	nn, mm		$\text{reg}[A] = \text{reg}[A] + (\text{reg}[nn] * \text{reg}[mm])$
FMSC	58	nn, mm		$\text{reg}[A] = \text{reg}[A] - (\text{reg}[nn] * \text{reg}[mm])$

LOADBYTE	59	bb		reg[0] = float(signed bb)
LOADUBYTE	5A	bb		reg[0] = float(unsigned byte)
LOADWORD	5B	b1, b2		reg[0] = float(signed b1*256 + b2)
LOADUWORD	5C	b1, b2		reg[0] = float(unsigned b1*256 + b2)
LOADE	5D			reg[0] = 2.7182818
LOADPI	5E			reg[0] = 3.1415927
LOADCON	5F	bb		reg[0] = float constant(bb)
FLOAT	60			reg[A] = float(reg[A])
FIX	61			reg[A] = fix(reg[A])
FIXR	62			reg[A] = fix(round(reg[A]))
FRAC	63			reg[A] = fraction(reg[A])
FSPLIT	64			reg[A] = integer(reg[A]), reg[0] = fraction(reg[A])
SELECTMA	65	nn, b1, b2		Select matrix A
SELECTMB	66	nn, b1, b2		Select matrix B
SELECTMC	67	nn, b1, b2		Select matrix C
LOADMA	68	b1, b2		reg[0] = Matrix A[bb, bb]
LOADMB	69	b1, b2		reg[0] = Matrix B[bb, bb]
LOADMC	6A	b1, b2		reg[0] = Matrix C[bb, bb]
SAVEMA	6B	b1, b2		Matrix A[bb, bb] = reg[A]
SAVEMB	6C	b1, b2		Matrix B[bb, bb] = reg[A]
SAVEMC	6D	b1, b2		Matrix C[bb, bb] = reg[A]
MOP	6E	bb		Matrix/Vector operation
LOADIND	7A	nn		reg[0] = reg[reg[nn]]
SAVEIND	7B	nn		reg[reg[nn]] = reg[A]
INDA	7C	nn		Select register A using value in reg[nn]
INDX	7D	nn		Select register X using value in reg[nn]
FCALL	7E	fn		Call user-defined function in Flash
EECALL	7F	fn		Call user-defined function in EEPROM
RET	80			Return from user-defined function
BRA	81	bb		Unconditional branch
BRA, cc	82	cc, bb		Conditional branch
JMP	83	b1, b2		Unconditional jump
JMP, cc	84	cc, b1, b2		Conditional jump
TABLE	85	tc, t0...tn		Table lookup
F'TABLE	86	cc, tc, t0...tn		Floating point reverse table lookup
L'TABLE	87	cc, tc, t0...tn		Long integer reverse table lookup
POLY	88	tc, t0...tn		reg[A] = nth order polynomial
GOTO	89	nn		Computed GOTO
LWRITE	90	nn, b1, b2, b3, b4		Write 32-bit long integer to reg[nn]
LWRITEA	91	b1, b2, b3, b4		Write 32-bit long integer to reg[A]
LWRITEX	92	b1, b2, b3, b4		Write 32-bit long integer to reg[X], X = X + 1
LWRITE0	93	b1, b2, b3, b4		Write 32-bit long integer to reg[0]
LREAD	94	nn	b1, b2, b3, b4	Read 32-bit long integer from reg[nn]
LREADA	95		b1, b2, b3, b4	Read 32-bit long value from reg[A]
LREADX	96		b1, b2, b3, b4	Read 32-bit long integer from reg[X], X = X + 1
LREAD0	97		b1, b2, b3, b4	Read 32-bit long integer from reg[0]

LREBYTE	98		bb	Read lower 8 bits of reg[A]
LREWORD	99		b1, b2	Read lower 16 bits reg[A]
ATOL	9A	aa...00		Convert ASCII to long integer
LTOA	9B	bb		Convert long integer to ASCII
LSET	9C	nn		reg[A] = reg[nn]
LADD	9D	nn		reg[A] = reg[A] + reg[nn]
LSUB	9E	nn		reg[A] = reg[A] - reg[nn]
LMUL	9F	nn		reg[A] = reg[A] * reg[nn]
LDIV	A0	nn		reg[A] = reg[A] / reg[nn] reg[0] = remainder
LCMP	A1	nn		Signed compare reg[A] and reg[nn], Set long integer status
LUDIV	A2	nn		reg[A] = reg[A] / reg[nn] reg[0] = remainder
LUCMP	A3	nn		Unsigned compare reg[A] and reg[nn], Set long integer status
LTST	A4	nn		Test reg[A] AND reg[nn], Set long integer status
LSET0	A5			reg[A] = reg[0]
LADD0	A6			reg[A] = reg[A] + reg[0]
LSUB0	A7			reg[A] = reg[A] - reg[0]
LMUL0	A8			reg[A] = reg[A] * reg[0]
LDIV0	A9			reg[A] = reg[A] / reg[0] reg[0] = remainder
LCMP0	AA			Signed compare reg[A] and reg[0], set long integer status
LUDIV0	AB			reg[A] = reg[A] / reg[0] reg[0] = remainder
LUCMP0	AC			Unsigned compare reg[A] and reg[0], Set long integer status
LTST0	AD			Test reg[A] AND reg[0], Set long integer status
LSETI	AE	bb		reg[A] = long(bb)
LADDI	AF	bb		reg[A] = reg[A] + long(bb)
LSUBI	B0	bb		reg[A] = reg[A] - long(bb)
LMULI	B1	bb		reg[A] = reg[A] * long(bb)
LDIVI	B2	bb		reg[A] = reg[A] / long(bb) reg[0] = remainder
LCMPI	B3	bb		Signed compare reg[A] - long(bb), Set long integer status
LUDIVI	B4	bb		reg[A] = reg[A] / unsigned long(bb) reg[0] = remainder
LUCMPI	B5	bb		Unsigned compare reg[A] and long(bb), Set long integer status
LTSTI	B6	bb		Test reg[A] AND long(bb), Set long integer status
LSTATUS	B7	nn		Set long integer status for reg[nn]
LSTATUSA	B8			Set long integer status for reg[A]
LCMP2	B9	nn, mm		Signed long compare reg[nn], reg[mm] Set long integer status

LUCMP2	BA	nn, mm		Unsigned long compare reg[nn], reg[mm] Set long integer status
LNEG	BB			reg[A] = -reg[A]
LABS	BC			reg[A] = reg[A]
LINC	BD	nn		reg[nn] = reg[nn] + 1, set status
LDEC	BE	nn		reg[nn] = reg[nn] - 1, set status
LNOT	BF			reg[A] = NOT reg[A]
LAND	C0	nn		reg[A] = reg[A] AND reg[nn]
LOR	C1	nn		reg[A] = reg[A] OR reg[nn]
LXOR	C2	nn		reg[A] = reg[A] XOR reg[nn]
LSHIFT	C3	nn		reg[A] = reg[A] shift reg[nn]
LMIN	C4	nn		reg[A] = min(reg[A], reg[nn])
LMAX	C5	nn		reg[A] = max(reg[A], reg[nn])
LONGBYTE	C6	bb		reg[0] = long(signed byte bb)
LONGUBYTE	C7	bb		reg[0] = long(unsigned byte bb)
LONGWORD	C8	b1, b2		reg[0] = long(signed b1*256 + b2)
LONGUWORD	C9	b1, b2		reg[0] = long(unsigned b1*256 + b2)
LONGCON	CA	bb		reg[0] = long constant(nn)
SETOUT	D0	bb		Set OUT1 and OUT2 output pins
ADCMODE	D1	bb		Set A/D trigger mode
ADCTRIG	D2			A/D manual trigger
ADCSCALE	D3	ch		ADCscale[ch] = reg[0]
ADCLONG	D4	ch		reg[0] = ADCvalue[ch]
ADCLOAD	D5	ch		reg[0] = float(ADCvalue[ch]) * ADCscale[ch]
ADCWAIT	D6			wait for next A/D sample
TIMESSET	D7			time = reg[0]
TIMELONG	D8			reg[0] = time (long integer)
TICKLONG	D9			reg[0] = ticks (long integer)
EESAVE	DA	nn, ee		EEPROM[ee] = reg[nn]
EESAVEA	DB	ee		EEPROM[ee] = reg[A]
EELOAD	DC	nn, ee		reg[nn] = EEPROM[ee]
EELOADA	DD	ee		reg[A] = EEPROM[ee]
EEWRITE	DE	ee, bc, b1...bn		Store bytes starting at EEPROM[ee]
EXTSET	E0			external input count = reg[0]
EXTLONG	E1			reg[0] = external input counter
EXTWAIT	E2			wait for next external input
STRSET	E3	aa...00		Copy string to string buffer
STRSEL	E4	bb, bb		Set selection point
STRINS	E5	aa...00		Insert string at selection point
STRCMP	E6	aa...00		Compare string with string buffer
STRFIND	E7	aa...00		Find string and set selection point
STRFCHR	E8	aa...00		Set field separators
STRFIELD	E9	bb		Find field and set selection point
STRTOF	EA			Convert string selection to floating point
STRTOL	EB			Convert string selection to long integer
READSEL	EC		aa...00	Read string selection
SYNC	F0		5C	Get synchronization byte
READSTATUS	F1		ss	Read status byte

READSTR	F2		aa...00	Read string from string buffer
VERSION	F3			Copy version string to string buffer
IEEEMODE	F4			Set IEEE mode (default)
PICMODE	F5			Set PIC mode
CHECKSUM	F6			Calculate checksum for uM-FPU code
BREAK	F7			Debug breakpoint
TRACEOFF	F8			Turn debug trace off
TRACEON	F9			Turn debug trace on
TRACESTR	FA	aa...00		Send string to debug trace buffer
TRACEREG	FB	nn		Send register value to trace buffer
READVAR	FC	nn		Read internal register value
RESET	FF			Reset (9 consecutive FF bytes cause a reset, otherwise it is a NOP)

Notes:

Opcode	Opcode value in hexadecimal
Arguments	Additional data required by instruction
Returns	Data returned by instruction
nn	register number (0-127)
mm	register number (0-127)
fn	function number (0-63)
bb	8-bit value
b1 , b2	16-bit value (b1 is MSB)
b1 , b2 , b3 , b4	32-bit value (b1 is MSB)
b1 . . . bn	string of 8-bit bytes
ss	Status byte
cc	Condition code
ee	EEPROM address slot (0-255)
ch	A/D channel number
bc	Byte count
t1 . . . tn	String of 32-bit table values
aa . . . 00	Zero terminated ASCII string

Appendix B


Floating Point Numbers

Floating point numbers can store both very large and very small values by “floating” the window of precision to fit the scale of the number. Fixed point numbers can’t handle very large or very small numbers and are prone to loss of precision when numbers are divided. The representation of floating point numbers used by the uM-FPU V3.1 is defined by the 32-bit IEEE 754 standard. The number of significant digits for a 32-bit floating point number is slightly more than 7 digits, and the range of values that can be handled is approximately $\pm 10^{38.53}$.

32-bit IEEE 754 Floating Point Representation

IEEE 754 floating point numbers have three components: a sign, exponent, the mantissa. The sign indicates whether the number is positive or negative. The exponent has an implied base of two and a bias value. The mantissa represents the fractional part of the number.

The 32-bit IEEE 754 representation is as follows:



Sign Bit (bit 31)

The sign bit is 0 for a positive number and 1 for a negative number.

Exponent (bits 30-23)

The exponent field is an 8-bit field that stores the value of the exponent with a bias of 127 that allows it to represent both positive and negative exponents. For example, if the exponent field is 128, it represents an exponent of one ($128 - 127 = 1$). An exponent field of all zeroes is used for denormalized numbers and an exponent field of all ones is used for the special numbers +infinity, -infinity and Not-a-Number (described below).

Mantissa (bits 30-23)

The mantissa is a 23-bit field that stores the precision bits of the number. For normalized numbers there is an implied leading bit equal to one.

Special Values

Zero

A zero value is represented by an exponent of zero and a mantissa of zero. Note that +0 and -0 are distinct values although they compare as equal.

Denormalized

If an exponent is all zeros, but the mantissa is non-zero the value is a denormalized number. Denormalized numbers are used to represent very small numbers and provide for an extended range and a graceful transition towards zero on underflows. Note: The uM-FPU V3.1 chip does not support

operations using denormalized numbers.

Infinity

The values +infinity and -infinity are denoted with an exponent of all ones and a fraction of all zeroes. The sign bit distinguishes between +infinity and -infinity. This allows operations to continue past an overflow. A nonzero number divided by zero will result in an infinity value.

Not A Number (NaN)

The value NaN is used to represent a value that does not represent a real number. An operation such as zero divided by zero will result in a value of NaN. The NaN value will flow through any mathematical operation. Note: The FPU initializes all of its registers to NaN at reset, therefore any operation that uses a register that has not been previously set with a value will produce a result of NaN.

Some examples of 32-bit IEEE 754 floating point values displayed as four byte values are as follows:

\$00, \$00, \$00, \$00	'0.0
\$3D, \$CC, \$CC, \$CD	'0.1
\$3F, \$00, \$00, \$00	'0.5
\$3F, \$40, \$00, \$00	'0.75
\$3F, \$7F, \$F9, \$72	'0.9999
\$3F, \$80, \$00, \$00	'1.0
\$40, \$00, \$00, \$00	'2.0
\$40, \$2D, \$F8, \$54	'2.7182818 (e)
\$40, \$49, \$0F, \$DB	'3.1415927 (pi)
\$41, \$20, \$00, \$00	'10.0
\$42, \$C8, \$00, \$00	'100.0
\$44, \$7A, \$00, \$00	'1000.0
\$44, \$9A, \$52, \$2B	'1234.5678
\$49, \$74, \$24, \$00	'1000000.0
\$80, \$00, \$00, \$00	'-0.0
\$BF, \$80, \$00, \$00	'-1.0
\$C1, \$20, \$00, \$00	'-10.0
\$C2, \$C8, \$00, \$00	'-100.0
\$7F, \$C0, \$00, \$00	'NaN (Not-a-Number)
\$7F, \$80, \$00, \$00	'+inf
\$FF, \$80, \$00, \$00	'-inf