



Micromega Corporation

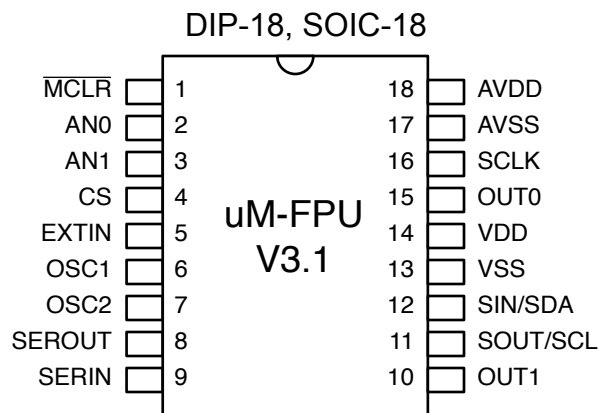
Using uM-FPU V3.1 with the Javelin Stamp™

Introduction

The uM-FPU V3.1 chip is a 32-bit floating point coprocessor that can be easily interfaced with the Javelin Stamp™ to provide support for 32-bit IEEE 754 floating point operations and 32-bit long integer operations. The uM-FPU V3.1 chip supports both I²C and SPI connections.

This document describes how to use the uM-FPU V3.1 chip with the Javelin Stamp. 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	SEROUT	Output	Serial Output, Debug Monitor - Tx
9	SERIN	Input	Serial Input, Debug Monitor - Rx
10	OUT1	Output	Digital Output 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	Digital Output 0
16	SCLK	Input	SPI Clock
17	AVSS	Power	Analog Ground
18	AVDD	Power	Analog Supply Voltage

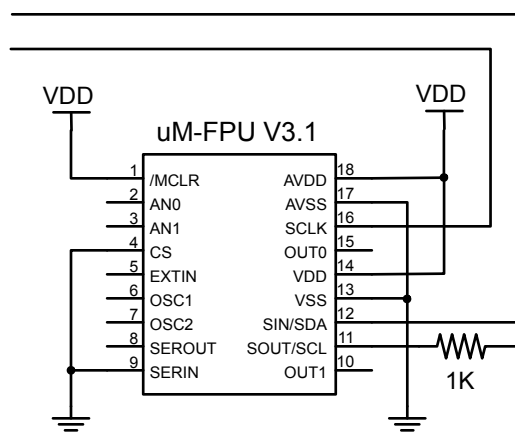
Connecting the Javelin Stamp using 2-wire SPI

Only two pins are required for interfacing to the Javelin Stamp to the uM-FPU V3.1 chip using a 2-wire SPI interface. The communication uses a bidirectional serial interface that requires a clock pin and a data pin. The default settings for these pins are shown below (they can be changed to suit your application):

```
final static int DATA_PIN = CPU.pin14;
final static int CLOCK_PIN = CPU.pin15;
```

Javelin Stamp Pins

CPU.pin14 (default) DATA_PIN
CPU.pin15 (default) CLOCK_PIN



Brief Overview of the uM-FPU V3.1 Floating Point Coprocessor

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 Javelin Stamp communicates with the FPU using an SPI or I²C interface. Instructions and data are sent to the FPU, and the FPU performs the calculations. The Javelin Stamp is free to do other tasks while the FPU performs calculations. Results can be read back to the Javelin Stamp 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 uM-FPU V3.1 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 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 the `Fpu.startWrite`, `Fpu.write`, and `Fpu.stop` methods are used as follows:

```
Fpu.startWrite();
Fpu.write(Fpu.FADD, 5);
Fpu.stop();
```

The `Fpu.startWrite` and `Fpu.stop` methods are used to indicate the start and end of a write transfer. A write transfer will often consist of several instructions and data. Up to 256 bytes can be sent in a single write transfer. If more than 256 bytes are required, the `Fpu.wait` method must be called to wait for the FPU to be ready before starting another write transfer and sending more instructions and data.

The `Fpu.write` method can have up to eight parameters. Each parameter is an 8-bit value that represents an instruction or data to be sent to the FPU. The `Fpu` class contains definitions for all of the uM-FPU V3.1 opcodes.

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

```
Fpu.write(Fpu.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.

```
Fpu.write(Fpu.FADD, 5);
```

Some instructions return data. This example reads the lower 8 bits of register A:

```
Fpu.write(Fpu.LREADBYTE);  
Fpu.wait();  
dataByte = Fpu.read(Fpu.dataByte);
```

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

```
Fpu.write(Fpu.SELECTA, 2, Fpu.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:

```
final static int Total = 2 // total amount (uM-FPU register)  
final static int Count = 5 // current count (uM-FPU register)  
  
Fpu.startWrite();  
Fpu.write(Fpu.SELECTA, Total, Fpu.FADD, Count);  
Fpu.stop();
```

Tutorial Examples

Now that we've introduced some of the basic concepts of sending instructions to the uM-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 Javelin Stamp 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 variable on the Javelin Stamp 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 uM-FPU, convert it to floating point, then divide it by 2 to get degrees in Celsius.

```
Fpu.write(Fpu.SELECTA, DegC, Fpu.LOADWORD); // select Dec C as register A
Fpu.writeWord(rawTemp); // load rawTemp and convert to float
Fpu.write(Fpu.FSET0, Fpu.FDIVI, 2); // store value and divide by 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 often be loaded once in the initialization section of the program and used later in the main program. The value 1.8 can be loaded using the ATOF instruction as follows:

```
Fpu.write(Fpu.SELECTA, F1_8, Fpu.ATOF); // select F1_8 as register A
Fpu.writeString("1.8"); // load string and convert to float
Fpu.write(Fpu.FSET0); // store value
```

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

```
Fpu.write(Fpu.SELECTA, DegF, Fpu.FSET, DegC); // set DegF to DegC value
Fpu.write(Fpu.FMUL, F1_8, Fpu.FADDI, 32); // multiply by 1.8 and add 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.

Now we print the results. The `Fpu.floatFormat` method is used to convert a floating point value to a formatted string. The first parameter selects the FPU register, and the second parameter specifies the desired format. The tens digit is the total number of characters to display, and the ones digit is 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. The following example prints the temperature in degrees Fahrenheit.

```
System.out.println(Fpu.floatFormat(DegF, 51));
```

Sample code for this tutorial and a wiring diagram for the DS1620 are shown at the end of this document. The file `demo1.java` is also included with the support software. There is a second file called `demo2.java` 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.

Using the uM-FPU Javelin Stamp Packages

Two packages are provided to handle the communication between the Javelin Stamp and the uM-FPU V3.1 floating point coprocessor, using either a SPI or I²C interface. They are located as follows:

```
~\lib\com\micromegacorp\math\v3-spi    SPI interface
~\lib\com\micromegacorp\math\v3-i2c    I2C interface
```

All of the default methods are the same for both the SPI and I²C interfaces, so user programs can be developed to be compatible with either interface. The desired package can be selected with the `package` statement at the start of any file that uses the `Fpu` object. One of the following statements should be included at the start of the file.

```
package com.micromegacorp.math.v3-spi;
    or
package com.micromegacorp.math.v3-i2c;
```

All of the files in the support packages have been commented to provide support generating API documentation using `javadoc`.

Fpu.reset

In order to ensure that the Javelin Stamp and the uM-FPU coprocessor are synchronized, a reset call must be done at the start of every program. All FPU registers are reset to the special value NaN (Not a Number) which is equal to the hexadecimal value 7FC00000. The `Fpu.reset` method resets the FPU, confirms communications, and returns true if successful, or false if the reset fails. An example of a typical reset is as follows:

```
if (!Fpu.reset()) {
    System.out.println("uM-FPU not detected.");
    return;
}
```

The version number of the FPU chip can be displayed with the following statement:

```
System.out.println(Fpu.version());
```

Note: An alternate version of the `Fpu.reset` method is available if that supports user-defined pins for the interface.

Fpu.startWrite

This method is called to start all write transfers. It's required when using an I²C interface, but is optional when using a 2-wire SPI interface. If code is being written to be compatible with both types of interface, it should be included.

Fpu.startRead

This method is called to start all read transfers. It's required when using an I²C interface, but is optional when using a 2-wire SPI interface. If code is being written to be compatible with both types of interface, it should be included.

Fpu.stop

This method is called to stop a write or read transfer. It's required when using an I²C interface, but is optional when using a 2-wire SPI interface. If code is being written to be compatible with both types of interface, it should be included. If a read transfer begins immediately after a write transfer, the `Fpu.stop` is not required. The `Fpu.wait`, `Fpu.floatFormat`, and `Fpu.longFormat` methods call `Fpu.stop` internally.

Fpu.wait

This method must be called before issuing any read instruction. It waits until the FPU is ready and the 256-byte instruction buffer is empty.

```
Fpu.wait(); // wait until FPU is ready
Fpu.startWrite();
Fpu.write(Fpu.SELECTA, Fpu.READWORD); // send READWORD instruction
int dataWord = Fpu.readWord(); // read a 16-bit value
```

The uM-FPU V3.1 has a 256 byte instruction buffer. In most cases, data will be read back before 256 bytes have been sent to the FPU, but if a calculation 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.write

This method is used to send instructions and data to the FPU. Up to eight 8-bit values can be passed as parameters. A `Fpu.startWrite` call must be made at the start of a write transfer, before the first `Fpu.write` call is made.

Fpu.writeWord

This method sends a 16-bit value to the FPU.

Fpu.writeString

This method sends a string to the FPU followed by a zero byte to terminate the string.

Fpu.read

This method is used to read 8 bits of data from the FPU.

Fpu.readWord

This method is used to read a 16 bits of data from the FPU.

Fpu.readString

This method is used to read a zero terminated string from the FPU. The `Fpu.floatFormat`, `Fpu.longFormat`, and `Fpu.version` methods call `Fpu.readString` internally.

Fpu.readStatus

This method is used to read the FPU status byte.

Fpu.floatFormat

The floating point value contained in a FPU register is returned as a formatted string. The format parameter 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

If the format parameter is omitted, or has a value of zero, the default format is used. 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

Fpu.longFormat

The long integer value contained in a FPU register displayed as a formatted string. The format parameter 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

If the format parameter is omitted, or has a value of zero, the default format is used. The displayed value can range from 1 to 11 characters in length. Examples of the display format are as follows:

```

1
500000
-3598390

```

Loading Data Values to the FPU

Most of the data read from devices connected to the Javelin Stamp 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 lower 8 bits of variable `dataByte` to the `Result` register on the FPU.

```
Fpu.write(Fpu.SELECTA, Result, Fpu.FSETI, dataByte);
```

The `LOADBYTE` instruction takes 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 takes 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 takes 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 lower 16 bits of variable `dataWord` to the `Result` register on the FPU.

```
Fpu.write(Fpu.SELECTA, Result, Fpu.LOADWORD);  
Fpu.writeWord(dataWord);  
Fpu.write(Fpu.FADD0);
```

The `LOADUWORD` instruction takes 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.

32-bit Floating Point Constants to Floating point

The `FWRITE`, `FWRITE0`, `FWRITEA`, and `FWRITEEX` instructions can be used to write a 32-bit floating point value. 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 in hexadecimal is 41A00000).

```
Fpu.write(Fpu.FWRITE, Angle, 0x41, 0xA0, 0x00, 0x00);
```

ASCII string to Floating Point

The `ATOF` instruction is used to convert zero-terminated strings to floating point values. The string of bytes following the opcode are converted to a floating point value (until a zero terminator is encountered). The converted value is stored in register 0. The following example sets the register `Angle` to 1.5885.

```
Fpu.write(Fpu.SELECTA, Angle, Fpu.FTOA);  
Fpu.writeString("1.5885");  
Fpu.write(Fpu.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 lower 8 bits of variable `dataByte` to the `Total` register on the FPU.

```
Fpu.write(Fpu.SELECTA, Total, Fpu.LADDI, dataByte);
```

The `LONGBYTE` instruction takes 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 takes 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 takes 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 lower 16 bits of variable `dataWord` to the `Total` register on the FPU.

```
Fpu.write(Fpu.SELECTA, Total, Fpu.LOADWORD);
Fpu.writeWord(dataWord);
Fpu.write(Fpu.LADD0);
```

The `LONGUWORD` instruction takes 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.

32-bit Integer Constants to Long Integer

The `LWRITE`, `LWRITE0`, `LWRITEA`, and `LWRITEX` instructions can be used to write a 32-bit integer value. The java compiler can handle 32-bit constants in expressions, if the `(short)` type is used, and the result of the expression is 16-bit. The following example stores the value 5000000 in register A.

```
Fpu.write(Fpu.LWRITEA);
Fpu.writeWord((short) (5000000 >> 16));
Fpu.writeWord((short) (5000000 & 0xFFFF));
```

ASCII string to Long Integer

The `ATOL` instruction is used to convert strings to long integer values. The string of bytes following the opcode are converted to a long integer value (until a zero terminator is encountered). The converted value is stored in register 0. The following example sets the register `Total` to 500000.

```
Fpu.write(Fpu.SELECTA, Total, Fpu.LTOA);
Fpu.writeString("500000");
Fpu.write(Fpu.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 one of the read methods is used to read data.

8-bit Integer

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

```
Fpu.wait();
Fpu.startWrite();
Fpu.write(Fpu.LREBYTE);
dataByte = Fpu.read();
```

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 variable `dataWord`.

```
Fpu.wait();
Fpu.startWrite();
Fpu.write(Fpu.LREADWORD);
dataWord = Fpu.readWord();
```

Long Integer to ASCII string

The `LTOA` instruction can be used to convert long integer values to an ASCII string, and the `Fpu.readString` method can be used to read the string.

```
Fpu.wait();
Fpu.startWrite();
Fpu.write(Fpu.LTOA, 0);
str = Fpu.readString();
```

Floating Point to ASCII string

The `FTOA` instruction can be used to convert floating point values to an ASCII string, and the `Fpu.readString` method can be used to read the string.

```
Fpu.wait();
Fpu.startWrite();
Fpu.write(Fpu.FTOA, 0);
str = Fpu.readString();
```

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 FPU status byte is read using the `Fpu.readStatus` method. The current status is returned in the `status` variable. Bit definitions are provided for the status bits in the `status` variable as follows:

<code>ZERO_FLAG</code>	Zero status bit (0-not zero, 1-zero)
<code>SIGN_FLAG</code>	Sign status bit (0-positive, 1-negative)
<code>NAN_FLAG</code>	Not a Number status bit (0-valid number, 1-NaN)
<code>INFINITY_FLAG</code>	Infinity status bit (0-not infinite, 1-infinite)

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:

```
Fpu.startWrite();
Fpu.write(Fpu.FSTATUSA);
status = Fpu.readStatus();
if ((status & Fpu.ZERO_FLAG) != 0)
    System.out.println("Result is Zero");
else if ((status & Fpu.SIGN_FLAG) != 0)
    System.out.println("Result is Negative");
```

The `FCMP`, `FCMP0`, and `FCMPI` 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:

```
Fpu.write(Fpu.FCMP0, 10);
status = Fpu.readStatus();
if ((status & Fpu.ZERO_FLAG) != 0)
    System.out.println("Value1 = Value2");
else if ((status & Fpu.SIGN_FLAG) != 0)
    System.out.println("Value1 < Value2");
else
    System.out.println("Value1 > Value2");
```

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`:

```
Fpu.write(Fpu.FCMP2, Value1, Value2);
GOSUB Fpu.readStatus
```

Comparing and Testing Long Integer Values

A long integer value can be zero, positive, or negative. The FPU status byte is read using the `Fpu.readStatus` method. It waits for the uM-FPU to be ready before sending the `READSTATUS` command and reading the status. The current status is returned in the `status` variable. Bit definitions are provided for the status bits in the `status` variable as follows:

<code>ZERO_FLAG</code>	Zero status bit (0-not zero, 1-zero)
<code>SIGN_FLAG</code>	Sign status bit (0-positive, 1-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:

```
Fpu.write(Fpu.LSTATUSA);
status = Fpu.readStatus();
if ((status & Fpu.ZERO_FLAG) != 0)
    System.out.println("Result is Zero");
else if ((status & Fpu.SIGN_FLAG) != 0)
    System.out.println("Result is Negative");
```

The `LCMP`, `LCMP0`, and `LCMPI` 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:

```
Fpu.write(Fpu.LCMPI, 10);
status = Fpu.readStatus();
if ((status & Fpu.ZERO_FLAG) != 0)
    System.out.println("Value1 = Value2");
else if ((status & Fpu.SIGN_FLAG) != 0)
    System.out.println("Value1 < Value2");
else
    System.out.println("Value1 > Value2");
```

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`:

```
Fpu.write(Fpu.LCMP2, Value1, Value2);
status = Fpu.readStatus();
```

The `LUCMP`, `LUCMP0`, and `LUCMPI` 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 `LTSTI` 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).

Int32, Float32, Math32 Objects

The Javelin Stamp only supports 16-bit values as native data types. In most cases, 32-bit values are only required for calculations on the FPU, and they can be stored in the FPU registers. If an application requires that 32-bit values be stored on the Javelin Stamp, two objects have been defined to store 32-bit data.

<code>Int32</code>	stores 32-bit long integer values
<code>Float32</code>	stores 32-bit floating point values

Both objects have `set`, `read`, `write`, and `toString` methods. The `Float32` object also provides `add`, `subtract`, `multiply`, `divide`, and `toHexString` methods to support legacy applications. The `Math32` object provides methods that are compatible with the java math package and are also provided to support legacy applications. See the javadoc descriptions of these files for details.


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 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.java)

```
import com.micromegacorp.math.v3_spi.*; // (use one of the uM-FPU packages)
//import com.micromegacorp.math.v3_i2c.*;
import stamp.core.*;
import stamp.peripheral.sensor.temperature.DS1620;

/**
 * @file    Demo1.java
 * @target  Javelin Stamp
 *
 * This program demonstrates how to use the uM-FPU V3.1 floating point chip
 * connected to the Javelin Stamp using either a 2-wire SPI or 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.
 *
 * @author  Cam Thompson, Micromega Corporation
 * @version
 *   August 24, 2007
 *   - updated for uM-FPU V3.1
 *   May 30, 2005
 *   - original version for uM-FPU V2
 */

public class Demo1 {

    final static int DS_DATA = CPU.pin10; // DS1620 data pin
    final static int DS_CLK  = CPU.pin11; // DS1620 clock pin
    final static int DS_RST  = CPU.pin12; // DS1620 reset/enable pin

    //----- uM-FPU register definitions -----

    final static int DegC = 1;           // degrees Celsius
    final static int DegF = 2;           // degrees Fahrenheit
    final static int F1_8 = 3;           // constant 1.8

    //----- main routine -----

    public static void main() {
        int rawTemp;

        // display program name
```



```

System.out.println("\u0010Demol");

// reset the uM-FPU and print version string
if (!Fpu.reset()) {
    System.out.println("uM-FPU not responding.");
    return;
}
else
    System.out.println(Fpu.version());

// get a DS1620 object and initialize
DS1620 ds = new DS1620(DS_DATA, DS_CLK, DS_RST);
CPU.delay(10000);

// store constant value (1.8)
Fpu.startWrite();
Fpu.write(Fpu.SELECTA, F1_8, Fpu.ATOF);
Fpu.writeString("1.8");
Fpu.write(Fpu.FSET0);
Fpu.stop();

// loop forever, read and display temperature
while (true) {
    // get temperature reading from DS1620
    rawTemp = ds.getTempRaw();

    // send to FPU and convert to floating point
    // divide by 2 to get degrees Celsius
    Fpu.startWrite();
    Fpu.write(Fpu.SELECTA, DegC, Fpu.LOADWORD);
    Fpu.writeWord(rawTemp);
    Fpu.write(Fpu.FSET0, Fpu.FDIVI, 2);

    // degF = degC * 1.8 + 32
    Fpu.write(Fpu.SELECTA, DegF, Fpu.FSET, DegC, Fpu.FMUL, F1_8, Fpu.FADDI, 32);
    Fpu.stop();

    // display degrees Celsius
    System.out.print("\n\rDegrees C: ");
    System.out.println(Fpu.floatFormat(DegC, 51));

    // display degrees Fahrenheit
    System.out.print("Degrees F: ");
    System.out.println(Fpu.floatFormat(DegF, 51));

    // delay about 2 seconds, then get the next reading
    CPU.delay(21000);
}
}
} // end class

```

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
CLRO	06			reg[nn] = 0
COPY	07	mm, nn		reg[nn] = reg[mm]
COPYA	08	nn		reg[nn] = reg[A]
COPYX	09	nn		reg[nn] = reg[X], X = X + 1
LOAD	0A	nn		reg[0] = reg[nn]
LOADA	0B			reg[0] = reg[A]
LOADX	0C			reg[0] = reg[X], X = X + 1
ALOADX	0D			reg[A] = reg[X], X = X + 1
XSAVE	0E	nn		reg[X] = reg[nn], X = X + 1
XSAVEA	0F			reg[X] = reg[A], X = X + 1
COPY0	10	nn		reg[nn] = reg[0]
COPYI	11	bb, nn		reg[nn] = long(unsigned byte bb)
SWAP	12	nn, mm		Swap reg[nn] and reg[mm]
SWAPA	13	nn		Swap reg[A] and reg[nn]
LEFT	14			Left parenthesis
RIGHT	15			Right parenthesis
FWRITE	16	nn, b1, b2, b3, b4		Write 32-bit floating point to reg[nn]
FWRITEA	17	b1, b2, b3, b4		Write 32-bit floating point to reg[A]
FWRITEX	18	b1, b2, b3, b4		Write 32-bit floating point to reg[X]
FWRITE0	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]
FREADA	1B		b1, b2, b3, b4	Read 32-bit floating point from reg[A]
FREADX	1C		b1, b2, b3, b4	Read 32-bit floating point from reg[X]
FREAD0	1C		b1, b2, b3, b4	Read 32-bit floating point from reg[0]
ATOF	1E	aa...00		Convert ASCII to floating point
FTOA	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]
FSUBR	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]
FDIVR	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
FSET0	29			reg[A] = reg[0]
FADD0	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			$\text{reg}[A] = \text{log10}(\text{reg}[A])$
EXP	45			$\text{reg}[A] = \text{exp}(\text{reg}[A])$
EXP10	46			$\text{reg}[A] = \text{exp10}(\text{reg}[A])$
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		$\text{reg}[A] = \text{atan2}(\text{reg}[A], \text{reg}[nn])$
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
FFT	6F	bb		Fast Fourier Transform
WRBLK	70	tc t1...tn		Write multiple 32-bit values
RDBLK	71	tc	t1...tn	Read multiple 32-bit values
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	bb		Call user-defined function in Flash
EECALL	7F	bb		Call user-defined function in EEPROM
RET	80			Return from user-defined function
BRA	81	bb		Unconditional branch
BRA	82	cc, bb		Conditional branch
JMP	83	b1, b2		Unconditional jump
JMP	84	cc, b1, b2		Conditional jump
TABLE	85	tc, t0...tn		Table lookup
FTABLE	86	cc, tc, t0...tn		Floating point reverse table lookup
LTABLE	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]
LREADBYTE	98		bb	Read lower 8 bits of reg[A]
LREADWORD	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)
SETSTATUS	CD	ss		Set status byte
SEROUT	CE	bb bb bd bb aa...00		Serial output
SERIN	CF	bb		Serial input
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] = B
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
TIMESET	D7			time = reg[0]
TIMELONG	D8			reg[0] = time (long integer)
TICKLONG	D9			reg[0] = ticks (long integer)
EESAVE	DA	mm, nn		EEPROM[nn] = reg[mm]
EESAVEA	DB	nn		EEPROM[nn] = reg[A]
EELOAD	DC	mm, nn		reg[mm] = EEPROM[nn]
EELOADA	DD	nn		reg[A] = EEPROM[nn]
EEWRITE	DE	nn, bc, b1...bn		Store bytes in EEPROM
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 selected string to floating point
STRTOL	EB			Convert selected string to long integer
READSEL	EC		aa...00	Read selected string
STRBYTE	ED	bb		Insert byte at selection point
STRINC	EE			Increment string selection point
STRDEC	EF			Decrement string selection point
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 Instruction opcode 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
bd baud rate and debug mode
cc Condition code
ee EEPROM address slot (0-255)
ch A/D channel number
bc Byte count
tc 32-bit value count
t1...tn String of 32-bit values
aa...00 Zero terminated ASCII string

In the FPUdefs.bas file, LEFT, RIGHT, READ, SIN, COS, GOTO, SEROUT, SERIN have been renamed to include an F_ prefix (e.g. F_SIN, F_COS, etc.) to avoid conflicts with reserved symbol names.


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 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 uM-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 32-bit hexadecimal constants are as follows:

\$00000000	' 0.0
\$3DCCCCCD	' 0.1
\$3F000000	' 0.5
\$3F400000	' 0.75
\$3F7FF972	' 0.9999
\$3F800000	' 1.0
\$40000000	' 2.0
\$402DF854	' 2.7182818 (e)
\$40490FDB	' 3.1415927 (pi)
\$41200000	' 10.0
\$42C80000	' 100.0
\$447A0000	' 1000.0
\$449A522B	' 1234.5678
\$49742400	' 1000000.0
\$80000000	' -0.0
\$BF800000	' -1.0
\$C1200000	' -10.0
\$C2C80000	' -100.0
\$7FC00000	' NaN (Not-a-Number)
\$7F800000	' +inf
\$FF800000	' -inf