C28x Floating Point Unit Library

Module User's Guide

C28x Foundation Software

V1.00 Beta 1

January 7, 2008

IMPORTANT NOTICE

Texas Instruments and its subsidiaries (TI) reserve the right to make changes to their products or to discontinue any product or service without notice, and advise customers to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgement, including those pertaining to warranty, patent infringement, and limitation of liability.

TI warrants performance of its semiconductor products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Customers are responsible for their applications using TI components.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards must be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance or customer product design. TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such products or services might be or are used. TI's publication of information regarding any third party's products or services does not constitute TI's approval, license, warranty or endorsement thereof.

Reproduction of information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations and notices. Representation or reproduction of this information with alteration voids all warranties provided for an associated TI product or service, is an unfair and deceptive business practice, and TI is not responsible or liable for any such use.

Resale of TI's products or services with *statements different from or beyond the parameters* stated by TI for that products or service voids all express and any implied warranties for the associated TI product or service, is an unfair and deceptive business practice, and TI is not responsible nor liable for any such use.

Also see: Standard Terms and Conditions of Sale for Semiconductor Products.
www.ti.com/sc/docs/stdterms.htm

Mailing Address:
Texas Instruments
Post Office Box 655303
Dallas, Texas 75265

Contents

1.Introduction

The Texas Instruments TMS320C28x Floating Point Unit (FPU) Library is collection of highly optimized application functions written for the C28x+FPU. These functions enable C/C++ programmers to take full advantage of the performance potential of the C28x+FPU. This document provides a description of each function included within the library.

Note:

The first release of the library contains an optimized real FFT algorithm. Future releases will include additional highly optimized algorithms.

2.Installing the Library

2.1. Where the Files are Located (Directory Structure)

As installed, the *C28x FPU Library* is partitioned into a well-defined directory structure. By default, the library and source code is installed into the c:\tidcs\c28\C28x_FPU_Lib\<version> directory. Table 1 describes the contents of the main directories used by library:

Table 1. C28x FPU Library Directory Structure

Directory	Description
<base>	Base install directory. By default this is c:\tidcs\c28\C28x_FPU_Lib\beta1 For the rest of this document <base> will be omitted from the directory names.
<base>\doc	Documentation including the revision history from the previous release.
<base>\lib	The built library.
<base>\include	Include files for the library functions. These include function prototypes and structure definitions.
<base>\source	Source files for the library. This also includes a Code Composer Studio project that can be used to re-build the library if required.

2.2. Build options used to build the library

The beta1 library is built with C28x codegen tools V5.0 Beta3 with the following options:

-g -o3 -d"_DEBUG" -d"LARGE_MODEL" -ml -v28 --float_support=fpu32

2.3. Header Files

A library header file is supplied in the <base>/include folder. This file contains structure definitions and function prototypes. The header file also includes the C28x data type definitions shown below:

```
#ifndef DSP28_DATA_TYPES
#define DSP28_DATA_TYPES
typedef int                 int16;
typedef long                int32;
typedef long long           int64;
typedef unsigned int        Uint16;
typedef unsigned long       Uint32;
typedef unsigned long long  Uint64;
typedef float               float32;
typedef long double         float64;
#endif
```

2.4. A note about C functions and IQMath

Most of the functions contained in the C28x FPU library are c-callable assembly. A few functions may be written in C. These C functions are written using the IQMath pre-processor notation. This allows these functions to be easily ported from fixed point to floating-point math. The included IQMath header file, IQmathLib.h, controls whether the code is built for fixed point or floating-point.

In this installation, the file IQmathLib.h file is configured to generate floating-point code. i.e. the MATH_TYPE in the file is defined as FLOAT_MATH. For more information on the IQMath notation, please refer to ***C28x™ IQMath Library - A Virtual Floating Point Engine*** (SPRC087) which can be downloaded from TI's website.

3. Function Summary

3.1. FPU Function Summary

This release is for the real FFT function. Other functions will be added in future releases.

Description	Prototype
Matrix and Vector Functions	
	Coming in a future release
DSP Functions	
RFFT_f32	void RFFT_f32(RFFT_F32_STRUCT *);
RFFT_f32u	void RFFT_f32u(RFFT_F32_STRUCT *);
RFFT_f32_mag	void RFFT_f32_mag(RFFT_F32_STRUCT *);
RFFT_f32s_mag	void RFFT_f32s_mag(RFFT_F32_STRUCT *);
RFFT_f32_phase	void RFFT_f32_phase(RFFT_F32_STRUCT *);
RFFT_f32_sincostable	void RFFT_f32_sincostable(RFFT_F32_STRUCT *);
Filter Functions	
	Coming in a future release

4. Function Descriptions

RFFT_f32	*Single-Precision Real Fast Fourier Transform*

Description This module computes a 32-bit floating-point real FFT including input bit reversing. This version of the function requires input buffer memory alignment. If you do not wish to align the input buffer, then use the RFFT_f32u function.

Header File `FPU.h`

Declaration `VOID RFFT_f32 (RFFT_F32_STRUCT *)`

Usage A pointer to the following structure is passed to the RFFT_f32 function:

```
typedef struct {
  float32          *InBuf;
  float32          *OutBuf;
  float32          *CosSinBuf;
  float32          *MagBuf;
  float32          *PhaseBuf;
  Uint16           FFTSize;
  Uint16           FFTStages;
} RFFT_F32_STRUCT;
```

Item	Description	Format	Comment
InBuf	Input data	Pointer to 32-bit float array	Input buffer alignment is required. Refer to the alignment section.

OutBuf	Output buffer	Pointer to 32-bit float array	Result of RFFT_f32. This buffer is used as the input to the magnitude and phase calculations. The output order for FFTSize = N is:
			`OutBuf[0] = real[0]`
			`OutBuf[1] = real[1]`
			`OutBuf[2] = real[2]`
			……
			`OutBuf[N/2] = real[N/2]`
			`OutBuf[N/2+1] = imag[N/2-1]`
			……
			`OutBuf[N-3] = imag[3]`
			`OutBuf[N-2] = imag[2]`
			`OutBuf[N-1] = imag[1]`

CosSinBuf	Twiddle factors	Pointer to 32-bit float array	Calculate using RFFT_f32_sincostable( ).
FFTSize	FFT size	Uint16	Must be a power of 2 greater than or equal to 32.
FFTStages	Number of stages	Uint16	Stages = log2(FFTSize)
MagBuf	Magnitude buffer	Pointer to 32-bit float array	Not used.
PhaseBuf	Phase buffer	Pointer to 32-bit float array	Not used.

Alignment Requirements:

The input buffer must be aligned to a multiple of the 2*FFTSize. For example, if the FFTSize is 128 you must align the buffer corresponding to InBuf to 2*128 = 256. An alignment to 64 will not work for a 128 FFT.

To align the input buffer, use the DATA_SECTION pragma to assign the buffer to a code section and then align the buffer to the proper offset in the linker command file. In this code example the buffer is assigned to the INBUF section.

```
#define FFT_SIZE 128
#pragma DATA_SECTION(Inbuffer, "INBUF");
float32 Inbuffer [N];
```

In the project's linker command file, the INBUF section is then aligned to a multiple of the FFTSize as shown below:

```
INBUF   ALIGN( 256 ) : { } >   RAML6 PAGE 1
```

Note:

If the input buffer is not properly aligned, then the output will be unpredictable.

Note:

If you do not wish to align the input buffer, then you must use the RFFT_f32u function. This version of the function does not have any input buffer alignment requirements. Using RFFT_f32u will, however, result in a lower cycle performance.

Example: The following sample code obtains the FFT of the real input.

```
#include FPU.h
#define FFT_SIZE   128        /* 32, 64, 128, 256, etc        */
#define FFT_STAGES   7        /* log2(FFT_SIZE)               */
/* Align the INBUF section to 2*FFT_SIZE in the linker file   */
#pragma DATA_SECTION(Inbuf, "INBUF");
float32 InBuffer[FFT_SIZE];
float32 OutBuffer[FFT_SIZE];
float32 TwiddleBuffer[FFT_SIZE];
RFFT_F32_STRUCT fft;
main()
{
       fft.InBuf = InBuffer;          /* Input data buffer      */
       fft.OutBuf = OutBuffer;        /* FFT output buffer      */
       fft.CosSinBuf = TwiddleBuffer;/* Twiddle factor buffer  */
       fft.FFTSize = FFT_SIZE;        /* FFT length             */
       fft.FFTStages = FFT_STAGES;    /* FFT Stages             */
       ……….
       RFFT_f32_sincostable(&fft) /* Initialize twiddle buffer */
       RFFT_f32(&fft);            /* Calculate output          */
}
```

Benchmark Information:

FFTSize	C-Callable ASM
32	608 cycles
64	1278 cycles
128	2784 cycles
256	6170 cycles
512	13672 cycles

Note: All buffers and stack are placed in internal memory (zero-wait states in data space).

RFFT_f32u	*Single-Precision Real Fast Fourier Transform*

Description This module computes a 32-bit floating-point real FFT including input bit reversing. This version of the function does not have any buffer alignment requirements. If you can align the input buffer, then use the RFFT_f32 function for improved performance.

Header File `FPU.h`

Declaration `VOID RFFT_f32u (RFFT_F32_STRUCT *)`

Usage A pointer to the following structure is passed to the RFFT_f32 function:

```
typedef struct {
  float32          *InBuf;
  float32          *OutBuf;
  float32          *CosSinBuf;
  float32          *MagBuf;
  float32          *PhaseBuf;
  Uint16           FFTSize;
  Uint16           FFTStages;
} RFFT_F32_STRUCT;
```

Item	Description	Format	Comment
InBuf	Input data	Pointer to 32-bit float array	Input data. No alignment is required.

OutBuf	Output buffer	Pointer to 32-bit float array	Result of RFFT_f32u. This buffer is used as the input to the magnitude and phase calculations. The output order for FFTSize = N is:
			`OutBuf[0] = real[0]`
			`OutBuf[1] = real[1]`
			`OutBuf[2] = real[2]`
			……
			`OutBuf[N/2] = real[N/2]`
			`OutBuf[N/2-1] = imag[N/2]`
			……
			`OutBuf[N-3] = imag[3]`
			`OutBuf[N-2] = imag[2]`
			`OutBuf[N-1] = imag[1]`

CosSinBuf	Twiddle factors	Pointer to 32-bit float array	Calculate using RFFT_f32_sincostable( ).
FFTSize	FFT Size	Uint16	Must be a power of 2 greater than or equal to 32.
FFTStages	Number of stages	Uint16	Stages = log2(FFTSize)
MagBuf	Magnitude buffer	Pointer to 32-bit float array	Not used.
PhaseBuf	Phase	Pointer to 32-bit float array	Not used.

Note:

If you can align the input buffer to a 2*FFTSize, then consider using the RFFT_f32 function. This version of the function has input buffer alignment requirements, but it is more cycle efficient

Example: The following sample code obtains the FFT of the real input.

```
#include FPU.h
#define FFT_SIZE   128        /* 32, 64, 128, 256, etc  */
#define FFT_STAGES   7        /* log2(FFT_SIZE)         */
float32 InBuffer[FFT_SIZE];
float32 OutBuffer[FFT_SIZE];
float32 TwiddleBuffer[FFT_SIZE];
RFFT_F32_STRUCT fft;
main()
{
       fft.InBuf = InBuffer;          /* Input Buffer           */
       fft.OutBuf = OutBuffer;        /* FFT Output Buffer      */
       fft.CosSinBuf = TwiddleBuffer;/* Twiddle Buffer         */
       fft.FFTSize = FFT_SIZE;        /* FFT length             */
       fft.FFTStages = FFT_STAGES;    /* FFT Stages             */
       ……….
       RFFT_f32_sincostable(&fft) /* Initialize twiddle buffer */
       RFFT_f32u(&fft);           /* Calculate output          */
}
```

Benchmark Information:

FFTSize	C-Callable ASM
32	664 cycles
64	1390 cycles
128	3008 cycles
256	6618 cycles
512	14568 cycles

Note: All buffers and stack are placed in internal memory (zero-wait states in data space).

RFFT_f32_mag	*Single-Precision Real Fast Fourier Transform Magnitude*

Description

This module computes the real FFT magnitude. The output from RFFT_f32_mag matches the magnitude output from the FFT found in common mathematics software and Code Composer Studio FFT graphs.

`

If instead a normalized magnitude like that performed by the fixed-point TMS320C28x IQmath FFT library is required, then the RFFT_f32s_mag function can be used. In fixed-point algorithms scaling is performed to avoid overflowing data. Floating-point calculations do not need this scaling to avoid overflow and therefore the RFFT_f32_mag function can be used instead.

Header File

`FPU.h`

Declaration

```
VOID RFFT_f32_mag (RFFT_F32_STRUCT *)
```

Usage

A pointer to the following structure is passed to the RFFT_f32_mag function:

```
typedef struct {
  float32          *InBuf;
  float32          *OutBuf;
  float32          *CosSinBuf;
  float32          *MagBuf;
  float32          *PhaseBuf;
  Uint16           FFTSize;
  Uint16           FFTStages;
} RFFT_F32_STRUCT;
```

Item	Description	Format	Comment
InBuf	Input data	Pointer to 32-bit float array	Not used.
OutBuf	Output buffer	Pointer to 32-bit float array	Result of RFFT_f32 or RFFT_f32u. Used as the input to magnitude and phase.
CosSinBuf	Twiddle factors	Pointer to 32-bit float array	Not used.
FFTSize	FFT size	Uint16	Must be a power of 2 greater than or equal to 32.
FFTStages	Number of stages	Uint16	Stages = log2(FFTSize)
MagBuf	Magnitude buffer	Pointer to 32-bit float array	Output from the magnitude calculation. FFTSize/2 in length.
PhaseBuf	Phase	Pointer to 32-bit float array	Not used.

The following sample code obtains the FFT magnitude.

```
#include FPU.h
#define FFT_SIZE   128        /* 32, 64, 128, 256, etc        */
#define FFT_STAGES   7        /* log2(FFT_SIZE)               */
/* Align the INBUF section to 2*FFT_SIZE in the linker file   */
#pragma DATA_SECTION(Inbuf, "INBUF");
float32 InBuffer[FFT_SIZE];
float32 OutBuffer[FFT_SIZE];
float32 TwiddleBuffer[FFT_SIZE];
float32 MagBuffer[FFT_SIZE/2];
RFFT_F32_STRUCT fft;

main()
{
       fft.InBuf = InBuffer;          /* Input data buffer     */
       fft.OutBuf = OutBuffer;        /* FFT output buffer     */
       fft.CosSinBuf = TwiddleBuffer; /* Twiddle factor buffer */
       fft.FFTSize = FFT_SIZE;        /* FFT length            */
       fft.FFTStages = FFT_STAGES;    /* FFT Stages            */
       fft.MagBuf = MagBuffer;        /* Magnitude buffer      */
       ……….
       RFFT_f32_sincostable(&fft) /* Initialize Twiddle Buffer */
       RFFT_f32(&fft);            /* Calculate output          */
       RFFT_f32_mag(&fft)         /* Calculate magnitude       */
}
```

Benchmark Information:

FFTSize	C-Callable ASM
32	1301 cycles
64	2645 cycles
128	5333 cycles
256	10709 cycles
512	21462 cycles

Note: All buffers and stack are placed in internal memory (zero-wait states in data space).

Note: The magnitude calculation calls the sqrt function within the runtime-support library. The magnitude function has not been optimized at this time.

RFFT_f32s_mag	*Single-Precision Real Fast Fourier Transform Magnitude*

Description This module computes the scaled real FFT magnitude. The scaling is 1/[2^(FFT_STAGES-1)], and is done to match the normalization performed by the fixed-point TMS320C28x IQmath FFT library for overflow avoidance. Floating-point calculations do not need this scaling to avoid overflow and therefore the RFFT_f32_mag function can be used instead. The output from RFFT_f32_mag matches the magnitude output from the FFT found in common mathematics software and Code Composer Studio FFT graphs.

Header File `FPU.h`

Declaration `VOID RFFT_f32s_mag (RFFT_F32_STRUCT *)`

Usage A pointer to the following structure is passed to the RFFT_f32_mag function:

```
typedef struct {
  float32         *InBuf;
  float32         *OutBuf;
  float32         *CosSinBuf;
  float32         *MagBuf;
  float32         *PhaseBuf;
  Uint16          FFTSize;
  Uint16          FFTStages;
} RFFT_F32_STRUCT;
```

Item	Description	Format	Comment
InBuf	Input data	Pointer to 32-bit float array	Not used.
OutBuf	Output buffer	Pointer to 32-bit float array	Result of RFFT_f32 or RFFT_f32u. Used as the input to magnitude and phase.
CosSinBuf	Twiddle factors	Pointer to 32-bit float array	Not used.
FFTSize	FFT size	Uint16	Must be a power of 2 greater than or equal to 32.
FFTStages	Number of stages	Uint16	Stages = log2(FFTSize)
MagBuf	Magnitude buffer	Pointer to 32-bit float array	Output from the magnitude calculation. FFTSize/2 in length.
PhaseBuf	Phase	Pointer to 32-bit float array	Not used.

The following sample code obtains the scaled FFT magnitude.

```
#include FPU.h
#define FFT_SIZE   128        /* 32, 64, 128, 256, etc        */
#define FFT_STAGES   7        /* log2(FFT_SIZE)               */
/* Align the INBUF section to 2*FFT_SIZE in the linker file   */
#pragma DATA_SECTION(Inbuf, "INBUF");
float32 InBuffer[FFT_SIZE];
float32 OutBuffer[FFT_SIZE];
float32 TwiddleBuffer[FFT_SIZE];
float32 MagBuffer[FFT_SIZE/2];
RFFT_F32_STRUCT fft;

main()
{
       fft.InBuf = InBuffer;          /* Input data buffer     */
       fft.OutBuf = OutBuffer;        /* FFT output buffer     */
       fft.CosSinBuf = TwiddleBuffer; /* Twiddle factor buffer */
       fft.FFTSize = FFT_SIZE;        /* FFT length            */
       fft.FFTStages = FFT_STAGES;    /* FFT Stages            */
       fft.MagBuf = MagBuffer;        /* Magnitude buffer      */
       ……….
       RFFT_f32_sincostable(&fft) /* Initialize twiddle buffer */
       RFFT_f32(&fft);            /* Calculate FFT output      */
       RFFT_f32s_mag(&fft)        /* Calculate scaled magnitude*/
}
```

Benchmark Information:

FFTSize	C-Callable ASM
32	1316 cycles
64	2676 cycles
128	5396 cycles
256	10836 cycles

Note: All buffers and stack are placed in internal memory (zero-wait states in data space).

Note: The magnitude calculation calls the sqrt function within the runtime-support library. The magnitude function has not been optimized at this time.

RFFT_f32_phase	*Single-Precision Real Fast Fourier Transform Phase*

Description This module computes FFT Phase.

Header File `FPU.h`

Declaration `VOID RFFT_f32_phase (RFFT_F32_STRUCT *)`

Usage A pointer to the following structure is passed to the RFFT_f32_phase function:

```
typedef struct {
  float32          *InBuf;
  float32          *OutBuf;
  float32          *CosSinBuf;
  float32          *MagBuf;
  float32          *PhaseBuf;
  Uint16           FFTSize;
  Uint16           FFTStages;
} RFFT_F32_STRUCT;
```

Item	Description	Format	Comment
InBuf	Input data	Pointer to 32-bit float array	Not used
OutBuf	Output buffer	Pointer to 32-bit float array	Result of RFFT_f32 or RFFT_f32u. Used as the input to magnitude and phase.
CosSinBuf	Twiddle factors	Pointer to 32-bit float array	Not used.
FFTSize	FFT Size	Uint16	Must be a power of 2 greater than or equal to 32.
FFTStage	Number of stages	Uint16	Stages = log2(FFTSize)
MagBuf	Magnitude buffer	Pointer to 32-bit float array	Not used.
PhaseBuf	Phase	Pointer to 32-bit float array	Output from the phase calculation. FFTSize/2 in length.

The following sample code obtains the FFT phase.

```
#include FPU.h
#define FFT_SIZE   128        /* 32, 64, 128, 256, etc        */
#define FFT_STAGES   7        /* ln(FFT_SIZE)/ln(2)           */
/* Align the INBUF section to 2*FFT_SIZE in the linker file   */
#pragma DATA_SECTION(Inbuf, "INBUF");
float32 InBuffer[FFT_SIZE];
float32 OutBuffer[FFT_SIZE];
float32 TwiddleBuffer[FFT_SIZE];
float32 PhaseBuffer[FFT_SIZE/2];
RFFT_F32_STRUCT fft;

main()
{
       fft.InBuf = InBuffer;          /* Input data buffer     */
       fft.OutBuf = OutBuffer;        /* FFT output buffer     */
       fft.CosSinBuf = TwiddleBuffer;/* Twiddle factor buffer */
       fft.FFTSize = FFT_SIZE;        /* FFT length            */
       fft.FFTStages = FFT_STAGES;    /* FFT Stages            */
       fft.PhaseBuf = Phasebuffer     /* Phase buffer          */
       ……….
       RFFT_f32_sincostable(&fft) /* Initialize twiddle buffer */
       RFFT_f32(&fft);            /* Calculate FFT output      */
       RFFT_f32_phase(&fft)       /* Calculate phase           */
}
```

Benchmark Information:

FFTSize	C-Callable ASM
32	9759 cycles
64	19830 cycles
128	33949 cycles
256	80225 cycles

Note: All buffers and stack are placed in internal memory (zero-wait states in data space).

Note: The phase function calls the atan2 function in the runtime-support library. The phase function has not been optimized at this time.

RFFT_f32_sincostable — *Single-Precision Real Fast Fourier Transform Twiddle Factors*

Description This module generates the twiddle factors used by the real FFT.

Header File `FPU.h`

Declaration

```
VOID RFFT_f32_sincostable (RFFT_F32_STRUCT *)
```

Usage A pointer to the following structure is passed to the RFFT_f32_sincostable function:

```
typedef struct {
  float32          *InBuf;
  float32          *OutBuf;
  float32          *CosSinBuf;
  float32          *MagBuf;
  float32          *PhaseBuf;
  Uint16           FFTSize;
  Uint16           FFTStages;
} RFFT_F32_STRUCT;
```

Item	Description	Format	Comment
Inbuf	Input data	Pointer to 32-bit float array	Not used
Outbuf	Output buffer	Pointer to 32-bit float array	Not used
CosSinbuf	Twiddle factors	Pointer to 32-bit float array	Output of the twiddle factor generation.
FFTSize	FFT Size	Uint16	Must be a power of 2 greater than or equal to 32.
FFTStage	Number of stages	Uint16	Stages = ln(FFTsize)/ln(2)
MagBuf	Magnitude buffer	Pointer to 32-bit float array	Not used.
PhaseBuf	Phase	Pointer to 32-bit float array	Not used

The following sample code obtains the FFT using the twiddle factors..

```
#define FFT_SIZE   128        /* 32, 64, 128, 256, etc       */
#define FFT_STAGES   7        /* ln(FFT_SIZE)/ln(2)          */
/* Align the INBUF section to 2*FFT_SIZE in the linker file  */
#pragma DATA_SECTION(Inbuf, "INBUF");
float32 Inbuffer[FFT_SIZE];
float32 Outbuffer[FFT_SIZE];
float32 Twiddlebuf[FFT_SIZE];
RFFT_F32_STRUCT fft;

main()
{
       fft.InBuf = Inbuffer;         /* Input data buffer     */
       fft.OutBuf = Outbuffer;       /* FFT output buffer     */
       fft.CosSinBuf = Twiddlebuf;   /* Twiddle factor buffer */
       fft.FFTSize = FFT_SIZE;       /* FFT length            */
       fft.FFTStages = FFT_STAGES;   /* FFT Stages            */
       ……….
       RFFT_f32_sincostable(&fft)  /* Initialize Twiddle Buffer*/
       RFFT_f32(&fft);             /* Calculate FFT output     */
}
```

Benchmark Information:

The RFFT_f32_cossintable function is written in C and not optimized.

5. Revision History

Beta 1 – January 7, 2008

- First release. Includes the real FFT and supporting functions.