
The ClearSpeed Accelerated DFT Library

Reference Manual

06-RM-1337 1.36.1.6

Company Confidential: This document contains proprietary and confidential information of ClearSpeed Technology plc, and is not to be copied in whole or in part.

ClearSpeed Technology, Inc.

3031 Tisch Way, Suite 200
San Jose, CA 95128

Tel: 408-557-2067

Fax: 408-557-9054

Email: info@clearspeed.com

Web: www.clearspeed.com

ClearSpeed Technology plc

3110 Great Western Court
Hunts Ground Road
Bristol BS34 8HP
United Kingdom

Tel: +44 (0)117 317 2000

Fax: +44 (0)117 317 2002

Conventions

Convention	Description
commands	This typeface means that the command must be entered exactly as shown in the text and the [Return] or [Enter] key pressed.
Screen displays	This typeface represents information as it appears on the screen and is generally enclosed within a bounding box.
[Key] names	Key names appear in the text written with brackets. For example [Return] or [F7] . If it is necessary to press more than one simultaneously, the key names are linked with a plus (+) sign: Press [Ctrl]+[Alt]+[Del]
Bold-face text	Signal name, instruction or register name. Selections made via the menu hierarchy of a software application.
Words in <i>italicized</i> type	Italics emphasize a point, concept or denote new terms.
	This symbol indicates important information or instructions.

1. Information and data contained in this document, together with the information contained in any and all associated ClearSpeed documents including without limitation, data sheets, application notes and the like ("Information") is provided in connection with ClearSpeed products and is provided for information only. Quoted figures in the Information, which may be performance, size, cost, power and the like are estimates based upon analysis and simulations of current designs and are liable to change.
2. Such Information does not constitute an offer of, or an invitation by or on behalf of ClearSpeed, or any ClearSpeed affiliate to supply any product or provide any service to any party having access to this Information. Except as provided in ClearSpeed Terms and Conditions of Sale for ClearSpeed products, ClearSpeed assumes no liability whatsoever.
3. ClearSpeed products are not intended for use, whether directly or indirectly, in any medical, life saving and/ or life sustaining systems or applications.
4. The worldwide intellectual property rights in the Information and data contained therein is owned by ClearSpeed. No license whether express or implied either by estoppel or otherwise to any intellectual property rights is granted by this document or otherwise. You may not download, copy, adapt or distribute this Information except with the consent in writing of ClearSpeed.
5. The system vendor remains solely responsible for any and all design, functionality and terms of sale of any product which incorporates a ClearSpeed product including without limitation, product liability, intellectual property infringement, warranty including conformance to specification and or performance.
6. Any condition, warranty or other term which might but for this paragraph have effect between ClearSpeed and you or which would otherwise be implied into or incorporated into the Information (including without limitation, the implied terms of satisfactory quality, merchantability or fitness for purpose), whether by statute, common law or otherwise are hereby excluded.
7. ClearSpeed reserves the right to make changes to the Information or the data contained therein at any time without notice.

© Copyright ClearSpeed Technology plc 2006. All rights reserved.

Advance, ClearSpeed, ClearConnect and the ClearSpeed logo are trade marks or registered trade marks of ClearSpeed Technology plc. All other brands and names are the property of their respective owners.

Contents

1 Introduction

1.1 Requirements

1.2 Patterns of use

1.3 Environment variables

1.4 Data formats

1.5 Memory descriptor

1.6 Plans

1.7 Flags

2 ClearSpeed DFT Host Library

CSDFT_create_plan_1d, CSDFT_create_plan_2d, CSDFT_create_plan_3d — Create a 1D, 2D or 3D plan for a DFT	12
CSDFT_execute_dft — Synchronous Execution of a DFT	14
CSDFT_create_user_function — Create user function	15
CSDFT_execute_user_function — Synchronous Execution of a user function	16
CSDFT_create_convolution_plan_2d — Create a 2D plan for a convolution operation	17
CSDFT_execute_convolution — Synchronous Execution of a convolution	18
CSDFT_free_plan — Free a DFT plan	19
CSDFT_get_status — Returns the current status for the CSDft library	20
CSDFT_return_error_message — Returns a string describing a status value of the CSDft library	22
CSDFT_get_null_descriptor — Return a special null memory descriptor	23
CSDFT_malloc_host — Allocates a memory descriptor encapsulating a data buffer allocated on the host	24
CSDFT_memory_descriptor_to_system — Given a host based memory descriptor returns a pointer to a buffer in system memory	25
CSDFT_system_to_memory_descriptor — Returns a host based memory descriptor encapsulating user provided source data.	26
CSDFT_system_split_array_to_memory_descriptor — Creates a host based memory descriptor encapsulating two user provided source data arrays	27
CSDFT_processor_to_memory_descriptor — Returns a board based memory descriptor encapsulating user provided source data.	28
CSDFT_memory_descriptor_copy — Copies a memory descriptor	29

<code>CSDFT_free</code> — Deallocates a memory descriptor created by the CSDft library	30
<code>CSDFT_ilog2</code> — Returns the integer log to the base two value of the input <code>CSDFT_bitreverse_1D_c</code> — Returns a bitreversed version of the single precision input data <code>CSDFT_bitreverse_1D_z</code> — Returns a bitreversed version of the double precision input data	
<code>CSDFT_optimal_to_natural_2D_c</code> — Transforms the input from board optimal order to natural order for single precision data <code>CSDFT_optimal_to_natural_2D_z</code> — Transforms the input from board optimal order to natural order for double precision data	
<code>CSDFT_natural_to_optimal_2D_c</code> — Transforms the input from natural order to board optimal order for single precision data <code>CSDFT_natural_to_optimal_2D_z</code> — Transforms the input from natural order to board optimal order for double precision data	31
<code>CSDFT_get_symbol_value</code> — Obtains the Board memory address of a specified symbol (on the first device of an Advance Board)	33
<code>CSDFT_get_csapi_handle</code> — Obtains the CSAPI handle for the first instance of an Advance Board in a system used by the CSDft library	34

3 ClearSpeed DFT Board Library

<code>CSDFT_create_plan_1d</code> , <code>CSDFT_create_plan_2d</code> , <code>CSDFT_create_plan_3d</code> — Create a 1D, 2D or 3D plan for a DFT	36
<code>CSDFT_execute_dft</code> — Synchronous Execution of a DFT	38
<code>CSDFT_create_convolution_plan_1d</code> , <code>CSDFT_create_convolution_plan_2d</code> — Create a 1D or 2D plan for a convolution operation	39
<code>CSDFT_execute_convolution</code> — Synchronous Execution of a convolution	40
<code>CSDFT_free_plan</code> — Free a DFT plan	41
<code>CSDFT_get_status</code> — Returns the current status for the CSDft library	42
<code>CSDFT_board_to_memory_descriptor</code> — Returns a board based memory descriptor encapsulating user provided source data.	44
<code>CSDFT_memory_descriptor_to_board</code> — Given a board based memory descriptor returns a pointer to a buffer in board memory	45
<code>CSDFT_free</code> — Deallocates a memory descriptor created by the CSDft library	46

1 Introduction

The ClearSpeed discrete Fourier transform (DFT) library consists of two parts; a host library and a library for the CSX600. The host library is called by an application running on the host computer in order to accelerate DFT operations using the CSX600 accelerator board(s). The CSX600 board library is designed to be called by code running on the CSX600 itself. The two versions of the library have a common API and a number of functions in common.

This manual has a chapter for each version of the library.

This library is still under active development and should be considered to be a pre-alpha product. There are many features still to be added and more comprehensive testing needs to be undertaken.

1.1 Requirements

In addition to the ClearSpeed runtime and drivers package, the ClearSpeed SDK should also be installed in order to be able to build the user examples included in the package or to develop code for the *Advance* board using the CSDFT board library.

1.2 Patterns of use

The ClearSpeed DFT Library consists of two similar APIs, one of which runs on the host and one on the processors on a ClearSpeed *Advance* board.

The CSDFT host API, while providing a similar interface to that of the CSDFT *Advance* board API, has additional heuristics which try and efficiently utilize both CSX600 processors on a board. To this end when a plan representing more than one DFT is executed the host library will automatically try and load-balance the work between the CSX600 processors. The host API is mostly intended to be the target for developers working to port existing host based applications, which may already use one of the many third party DFT libraries available (such as FFTW, MKL and ACML), to use the ClearSpeed *Advance* board.

The CSDFT board API is very similar to the host API but the mode of execution is different. Calls made to the board API from within a Cn or assembly language program are run on the single CSX600 which made the call.

The CSDFT board API can also be thought of as an embedded version of the API. This means that any Cn or assembly language program can call CSDFT routines. This also means that you can construct a .csx file which uses the board CSDFT library but where the data operated on comes from the host, via the usual CSAPI calls, without having to use the host CSDFT library. This means it is more suitable for users working to port full applications to run more or less entirely on the *Advance* board rather than simply trying to off load the DFT portions of their application to the *Advance* board.

1.3 Environment variables

There are three environment variables that can be set by a user which affect the behaviour of the CSDFT host library.

The first is `CS_CSAPI_DEBUGGER`, which will need to be enabled when debugging with `csgdb` (you will also need to specify `CS_CSAPI_DEBUGGER_ATTACH` as well, please see the SDK documentation for further information about `csgdb` environment variables). This environment variable affects which version of the board side executable that the host library will search for and load (release vs. debug). In this case it will be the debug versions of either

the default CSDFT library CSX files or any user built CSX files. If they are not present then the library will signal an error.

Another environment variable that may be set is CS_FFT_CSX which is used when the user needs to override the default board executable with one which contains their user routines. The environment variable needs a full path to the relevant csx file and must include the root of the csx filename, but without any appended _0, _1 or _debug. Be careful on Windows as the default installation location for ClearSpeed software is under the "Program Files" directory which has a space in it. This can cause windows to add "" to any command completed string and this will cause the host library to fail to locate the csx file.

The last environment variable that a user may set is CS_USE_SIM. When set this variable signals the host library that rather than connecting to an *Advance* board, the library is actually connecting to one of ClearSpeed's software simulators. The CSX executable can take some time to load (minutes).

1.4 Data formats

The two sample library data formats assumed are *interleaved* and *split*. In interleaved mode, real and imaginary data values are next to each other in one contiguous piece of memory as shown in Figure 1.1.

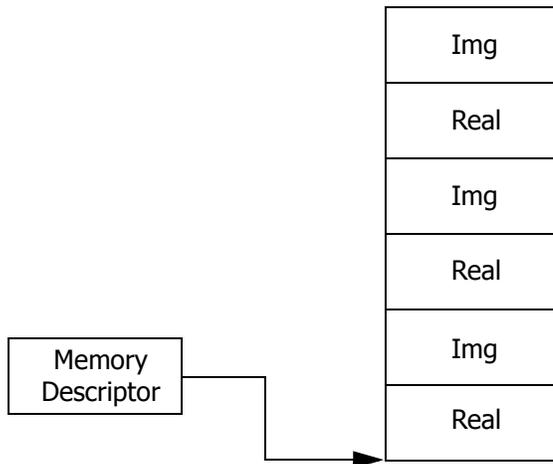


Figure 1.1 Interleaved mode

In split mode, the memory descriptor is used to point to two separate memory blocks, one for real, the other for imaginary data as shown in Figure 1.2.

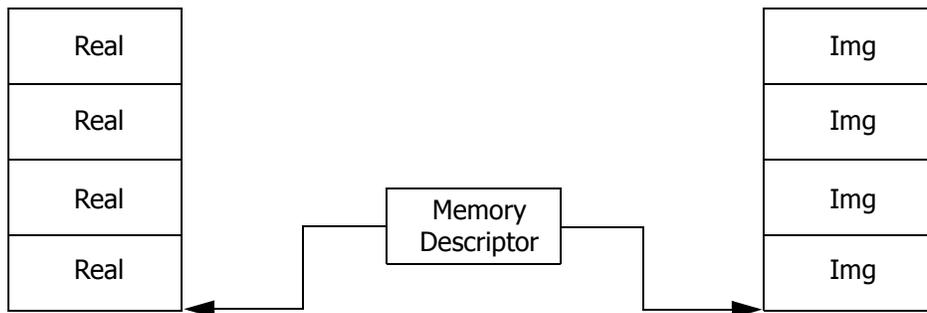


Figure 1.2 Split mode

The CSDFT library provides platform-independent single and double-precision structures, `FloatComplex` and `DoubleComplex`, which implement the interleaved memory layout. Developers may also use their own structures to implement interleaved or split mode data layout. Alternatively, if using Fortran or a C99-compatible compiler, developers may wish to use the complex data types built into these languages.

Data is also arranged in a row-major order (also known as 'C' format), which means that as you move through adjoining memory locations, the last dimension's index varies most quickly, and the first dimension's index varies most slowly.

1.5 Memory descriptor

The memory descriptor structure, `CSDftMemoryDescriptor`, provides an abstract datatype which can describe memory on the host or board to the CSDFT API.

The memory descriptor is essentially an intelligent pointer. As well as a pointer to the data, the memory descriptor also contains the following information:

- How the memory was allocated (allocated as part of the memory descriptor or referenced);
- The location on the host or board;
- The size of the memory;
- Whether there is one contiguous block of memory or it is split into two (necessary for split mode).

The memory descriptor may be used to reference input and output data in contiguous memory for all DFT functions.

Through the use of these memory descriptors, developers do not need to concentrate on where data is being stored (on the host or on the board) when executing plans with data referenced by these descriptors. All the logic required to move the data to the board to be executed and on to the destination to be stored is handled by the API.

The memory descriptor can be used to maximize library efficiency by ensuring that data is moved from the host to the board, multiple operations performed on it and then returned to the host. Instead of moving the data from the host to the board and back for each operation.

An example of the `CSDftMemoryDescriptor` is shown in the code fragment below:

```
//Declare memory descriptors to describe input and output data
CSDftMemoryDescriptor data_in;
CSDftMemoryDescriptor data_out;

// Declare other variables
int n=128, num_ffts=3;

//Create input and output data memory descriptors referencing memory on the host
data_in = CSDFT_system_to_memory_descriptor((void *) source_data_pointer,
                                           sizeof(DoubleComplex) * n,
                                           num_ffts);
data_out = CSDFT_system_to_memory_descriptor((void *) output_data_pointer,
                                             sizeof(DoubleComplex) * n,
                                             num_ffts);

//Create the plan
plan = CSDFT_create_plan_1d(...);

//Execute plan
status = CSDFT_execute_dft(plan, data_in, data_out);
```

1.6 Plans

The CSDFT library uses the concept of a *plan* to define a DFT before it is executed. A plan contains information about:

- the direction
- the precision
- the data format
 - whether data is complex or real
 - whether data is natural or optimal
- the number of samples
- the scale factor to use

All but the last two in the list above are set by using flags. The required flags are ORed together and passed into the “create plan” functions as an unsigned integer – a list of the flags, and an explanation of their function is provided in section 1.7, *Flags*, on page 9.

Data manipulated on the board and returned to the host can be in one of two formats, natural order (CSDFT_NATURAL_ORDER flag) or a more optimal order (CSDFT_OPTIMAL_ORDER) used by the CSDft library. The default is natural order, though specifying this will incur a performance penalty on the board (This mode is not supported directly for 1D on the board but uses the host to perform the transform from optimal to natural order). Note that it is possible to work only in optimal order and to this end the library provides some support routines to move data between these two data formats.

To achieve no apparent scaling in a forward DFT, the scale factor should be 1.0. In a backwards DFT however, the scale factor should be $1/(n^{DIMENSION})$. For example, a 1D DFT of size 128 would have a scale factor of 1/128, while a 2D DFT of size 128x128 would have a scale factor of 1/(128*128)

Example:

```
unsigned int flags;
CSDftPlan forward;
CSDftStatus status;

/* Create the flags for a forward, optimal, double and complex DFT */
flags = CSDFT_FORWARD | CSDFT_OPTIMAL_ORDER | CSDFT_DOUBLE | CSDFT_COMPLEX;

/* Create the 1D plan - 5 ffts each of size 128 */
forward = CSDFT_create_plan_1d(flags, 256, 5, 1.0);

/* Execute the plan */
status = CSDFT_execute_dft(forward, ...);
```

1.7 Flags

A number of functions use flag parameters to provide information about the operation to be performed. These are described in Table 1.1.

Flag	Description	Applicable
CSDFT_FORWARD	Perform a forward DFT	All
CSDFT_BACKWARD	Perform a backward DFT	All except convolution
CSDFT_NATURAL_ORDER	Assume order of input and output is natural. This is the default	All except 1D board function
CSDFT_OPTIMAL_ORDER	Assume the order of input or output is optimal for the particular DFT	All
CSDFT_SRC_SINGLE	Source is a single precision floating point format	All
CSDFT_SRC_DOUBLE	Source is in double precision floating point format	All
CSDFT_SRC_REAL	Source is of real type (Only valid for a forward Real to Complex forward transform)	Real to Complex forward transform
CSDFT_SRC_COMPLEX	Source is of complex type. This is the default	All
CSDFT_SRC_INTERLEAVED	Source is an interleaved array of data. This is the default.	All
CSDFT_SRC_SPLIT_ARRAY	Source is two arrays of data, one representing the real component, the other imaginary. Only valid for complex data. Not valid for optimal format data (for instance with a backward DFT).	All except optimal format data and 1D DFT.
CSDFT_DST_SINGLE	Destination is a single precision floating point format	All
CSDFT_DST_DOUBLE	Destination is a double precision floating point format	All
CSDFT_DST_REAL	Destination is of real type (Only valid for a Complex to Real backward transform)	Complex to real backward transform
CSDFT_DST_COMPLEX	Destination is of complex type. This is the default	All
CSDFT_DST_INTERLEAVED	Destination is an interleaved array of data. This is the default	All
CSDFT_DST_SPLIT_ARRAY	Destination is two arrays of data, one representing the real component, the other imaginary. Only valid for complex data. Not valid for optimal format data (for instance with a backward DFT).	All except optimal format data
CSDFT_SINGLE	Source and destination are in single precision floating point format	All
CSDFT_DOUBLE	Source and destination are in double precision floating point format	All
CSDFT_REAL	Source and destination are of real type	All
CSDFT_COMPLEX	Source and destination are of complex type. This is the default	All
CSDFT_TRANSFORM_KERNEL	The kernel will undergo a DFT before being multiplied with the transformed source	Convolution

Table 1.1 Flag parameters

2 ClearSpeed DFT Host Library

This section documents the ClearSpeed DFT host interface.

CSDFT_create_plan_1d, CSDFT_create_plan_2d, CSDFT_create_plan_3d — Create a 1D, 2D or 3D plan for a DFT

Synopsis

```
#include "csdft/csdft.h"
CSDftPlan CSDFT_create_plan_1d( CSDftFlags flags, unsigned int size, unsigned int
num_ffts, double scale );
CSDftPlan CSDFT_create_plan_2d( CSDftFlags flags, unsigned int size_x, unsigned int
size_y, unsigned int num_ffts, double scale );
CSDftPlan CSDFT_create_plan_3d( CSDftFlags flags, unsigned int size_x, unsigned int
size_y, unsigned int size_z, unsigned int num_ffts, double scale );
```

Description

These functions create plans for 1, 2 and 3 dimensional DFTs, which can then be executed using the CSDFT_execute_dft routine.

This library supports:

1D DFT

- Power of two sizes from 128 to 8192
- Single and Double precision
- Complex to Complex (forward and backward)
- Interleaved
- Natural and Optimal order else
- Optimal

2D DFT

- Power of two sizes from 128 to 2048 (must be square)
- Single and Double precision
- Real and Complex types:
 - Complex to Complex (forward and backward)
 - Real to Complex (forward)
 - Complex to Real (backward)
- Interleaved and Split Array inputs
- Natural and Optimal order

3D DFT

- 128 size only
- Single and Double precision
- Complex to Complex (forward and backward)
- Interleaved

- Natural order only

The complex data generated from a real to complex forward 2D DFT is a partial representation of the full complex output. This is due to the symmetry of the resulting data. Only the first $(n/2)+1$ columns are represented; in natural mode this means each row has only $(n/2)+1$ samples, with no padding at the end of the row. In optimal mode, only the first $(n/2)+1$ columns are output. When doing a backward complex to real DFT, this half-complex format is assumed.

It is possible to create multiple plans and relate them to multiple DFT computations.

The ClearSpeed DFT library performs the following forward 1D FFT:

$$Y[k]=\text{Sum}\{j=0,n-1\}(X[j]*\exp(-2*\pi*j*k*\text{sqrt}(-1)/n))$$

The ClearSpeed DFT library performs the following backward 1D DFT:

$$Y[k]=\text{Sum}\{j=0,n-1\}(X[j]*\exp(2*\pi*j*k*\text{sqrt}(-1)/n))$$

All the various configuration parameters related to a DFT are contained in a plan object in the CSDft API. A descriptor is created by the following API calls for 1d, 2d and 3D versions of `CSDFT_create_plan_<1D|2D|3D>`.

Returns

This function returns a plan if there are no errors found in the supplied flags or parameters, otherwise it returns NULL. If there is an error, then the global status will be set to indicate an error, which can be retrieved by using `CSDFT_get_status()`.

CSDFT_execute_dft — Synchronous Execution of a DFT

Synopsis

```
#include "csdft/csdft.h"
CSDftStatus CSDFT_execute_dft( const CSDftPlan plan, CSDftMemoryDescriptor source,
CSDftMemoryDescriptor destination );
```

Description

This function performs the DFT as described by the plan parameter, using data specified by the source and destination memory descriptors. This function is a synchronous interface and will only return when the DFT has completed.

Returns

Returns a status code of CSDFT_NO_ERROR if the function is successful, CSDFT_INVALID_PLAN if the plan passed as a parameter is not valid (such as the handle being NULL), CSDFT_INVALID_MEMORY_DESCRIPTOR if the handles for the memory descriptors are invalid or CSDFT_BOARD_ERROR if the library cannot connect to an Advance Card.

CSDFT_create_user_function — Create user function

Synopsis

```
#include "csdft/csdft.h"
CSDftUserFunction CSDFT_create_user_function(const char *func, unsigned int
common_size, unsigned int block_size_src, unsigned int block_size_dst, unsigned int
num_blocks);
```

Description

This function creates a handle to a 'user function', which is in the ClearSpeed executable (CSX) on the ClearSpeed Advance board. This function needs to be built into 2 CSX files (one for each chip) which the user has compiled. These CSX files need to have been linked with the board CSDFT library. They need to be of the form name_0.csx and name_1.csx. If debugging is necessary, then `_debug` should be added after the digit e.g. `mycsx_0_debug.csx`. In order to load this CSX file the environment variable `CS_FFT_CSX` must be set to be the name part of the CSX without the suffixes e.g. `/home/csx/mycsx`

The user function must be of the following format:

```
void funcname(CSDftMemoryDescriptor common, CSDftMemoryDescriptor in, CSDftMemoryDescriptor out);
```

This first parameter in `CSDFT_create_user_function()`, `func`, is the name of the function to be called in the CSX executable on the board. `block_size_src` and `block_size_dst` are the size in bytes of the source and the destination respectively, and `num_blocks` is the number of blocks in the source and destination.

The number of blocks indicates how many times the user function will be called. Each time, an amount of data equal to `block_size_src` will be given to the user function, and an amount of data equal to `block_size_dst` will be written to the destination. The common data parameter is broadcast once to each CSX600 running the user function.

To actually execute the board function, `CSDFT_execute_user_function()` must be called.

EXAMPLES

```
/*
 * To create a handle to a user function called myfunc
 * /
CSDftUserFunction func;
/*
 * Send 1kb of data to myfunc() at a time; repeat 10 times.
 * /
func = CSDFT_create_user_function("myfunc", 0, 1024, 1024, 10);
```

Returns

This function returns a `CSDftUserFunction` if there are no errors, other it returns `NULL`. If there is an error, the global status will be set to indicate an error and can be retrieved by using `CSDFT_get_status()`.

CSDFT_execute_user_function — Synchronous Execution of a user function

Synopsis

```
#include "csdft/csdft.h"
CSDftStatus CSDFT_execute_user_function(CSDftUserFunction func, CSDftMemoryDescriptor
common, CSDftMemoryDescriptor in, CSDftMemoryDescriptor out, CSDftProcessorID
processor );
```

Description

This function executes a user function residing in a CSX executable on a ClearSpeed Advance Board. The first parameter, *func*, is the handle to the user function, and is created using *CSDFT_create_user_function()*. *in* and *out* are memory descriptors describing the source and destination data locations respectively. *common* is a memory descriptor pointing to data to be sent to all relevant chips on the Advance boards (i.e. will act like a broadcast).

processor describes to which board and processor this data should be sent to. To choose a specific processor on a specific board, *CSDFT_BOARD_N* should be ORed with *CSDFT_PROCESSOR_M*, where N is the board number (starting at 0 and up to 6) and M is the processor number (starting at 0 and up to 3). Alternatively, developers can choose to send data to any free processor in sequence - to do this, *CSDFT_BOARD_SEQUENCE* should be ORed with *CSDFT_PROCESSOR_SEQUENCE*. If a specific processor is chosen, currently only one block of data may be sent (i.e the *num_blocks* parameter of *CSDFT_create_user_function()* must be 1).

The input blocks of data will be read contiguously from the input descriptor and written to the output descriptor.

EXAMPLES

```
/*
 * To create a handle to a user function called myfunc
 */
CSDftUserFunction func;
/*
 * Send 10kb of data to myfunc(), 1kb at a time.
 */
func = CSDFT_create_user_function("myfunc", 0, 1024, 1024, 10);
/*
 * Execute the function, sending data to whichever processor is free
 */
CSDFT_execute_user_function(func, md_null, md_in, md_out, CSDFT_BOARD_SEQUENCE |
CSDFT_PROCESSOR_SEQUENCE)
```

Returns

This function returns a *CSDftUserFunction* if there are no errors, other it returns *NULL*. If there is an error, the status will be set to indicate an error.

CSDFT_create_convolution_plan_2d — Create a 2D plan for a convolution operation

Synopsis

```
#include "csdft/csdft.h"
CSDftPlan CSDFT_create_convolution_plan_2d(CSDftFlags flags, unsigned int size_x,
unsigned int size_y, unsigned int num_ffts, double scale);
```

Description

Performing a convolution is similar to any other DFT - a plan should be created, specifying the correct flags and then executed. As usual, flags are ORed together - however one flag of specific note is CSDFT_TRANSFORM_KERNEL. If set, the kernel will undergo a DFT before being convolved with the transformed source. Otherwise, the kernel will be convolved "as is" with the transformed source. By default, the kernel is assumed to not need to be pre-transformed, and, in this implementation, the kernel is assumed to be the result of a single DFT. Therefore, the same kernel will be reused if *num_ffts* is greater than 1.

To perform a convolution with no apparent scaling, the scale parameter should be $1/(\text{size} \times \text{size})$. For example, scale would be $1/(512 \times 512)$ when performing a 512x512 DFT.

The complex data generated from a real to complex forward 2D DFT is a partial representation of the full complex output. This is due to the symmetry of the resulting data. Only the first $(n/2)+1$ columns are represented; in natural mode this means each row has only $(n/2)+1$ samples, with no padding at the end of the row. In optimal mode, only the first $n/2+1$ columns are output. When performing a backward complex to real DFT, this half-complex format is assumed. The kernel MUST be in optimal order when performing convolutions.

Returns

This function returns a plan if there are no errors found in the supplied flags or parameters, otherwise it returns NULL. If there is an error, then the global status will be set to indicate an error, which can be retrieved by using *CSDFT_get_status()*.

EXAMPLES

```
/*
To create a plan for a 2D, forward, double precision, optimal, complex convolution of
size 128:
*/
CSDftPlan plan;
unsigned int flags = 0;
flags = CSDFT_FORWARD | CSDFT_SRC_DOUBLE | CSDFT_OPTIMAL_ORDER | CSDFT_SRC_COMPLEX;
plan = CSDFT_create_convolution_plan_2d(flags, 128, 1, 1/(128*128));
```

CSDFT_execute_convolution — Synchronous Execution of a convolution

Synopsis

```
#include "csdft/csdft.h"
CSDftStatus CSDFT_execute_convolution( const CSDftPlan plan, CSDftMemoryDescriptor
kernel, CSDftMemoryDescriptor source, CSDftMemoryDescriptor destination ) ;
```

Description

This function uses memory descriptors to specify the memory locations of the kernel, source and destination data. A memory descriptor is an opaque data structure which represents data, both in host and board space. More information on this structure is available in the Introduction section.

If the convolution is performed correctly and no errors occur, the solution will be in the memory space referenced by the destination memory descriptors.

Returns

Returns a status code of CSDFT_NO_ERROR if the function is successful, CSDFT_INVALID_PLAN if the plan passed as a parameter is not valid (such as the handle being NULL), CSDFT_INVALID_MEMORY_DESCRIPTOR if the handles for the memory descriptors are invalid or CSDFT_BOARD_ERROR if the library cannot connect to an Advance Board.

CSDFT_free_plan — Free a DFT plan

Synopsis

```
#include "csdft/csdft.h"
CSDftStatus CSDFT_free_plan( CSDftPlan plan );
```

Description

This function frees up the previously created plan.

Returns

This function frees all associated resources for a plan and returns a CSDftStatus result of CSDFT_NO_ERROR or CSDFT_INVALID_MEMORY_DESCRIPTOR if *plan* is a null pointer.

CSDFT_get_status — Returns the current status for the CSDft library

Synopsis

```
#include "csdft/csdft.h"
CSDftStatus CSDFT_get_status( void );
```

Description

This function returns the current status code for the CSDft library. Specifically it returns the status code of the last executed CSDft library function. All Memory Descriptor, Plan and Execute functions will either set the global status code or return their status code as a return value as their prototype dictates.

The following are the valid status codes which will be returned or set by CSDft library functions.

- CSDFT_NO_ERROR No current error detected.
- CSDFT_INVALID_MEMORY_DESCRIPTOR When set by one of the Execute functions or one of the memory manipulation functions this code indicates that one of the pointer parameters to the function were NULL.
- CSDFT_MEMORY_ALLOCATION_ERROR When set by one of the Memory Descriptor functions that allocate memory this code indicates that a memory allocation error has been encountered either on the host or on the Advance board.
- CSDFT_BOARD_ERROR This code can be set as a result of the CSDft library not being able to find an Advance board to connect to. If there is definitely a board present in the computer, typing "csreset -v" into your command prompt should identify or fix the problem. Alternatively, this error is set due to an internal, unrecoverable error on the board or in the driver.
- CSDFT_INVALID_BOARD This code is returned when an operation is adjudged to be sent to a board that does not exist.
- CSDFT_INVALID_PROCESSOR This code is returned when an operation is adjudged to be sent to a processor that does not exist on a board that does exist.
- CSDFT_INVALID_PLAN When set by one of the planner functions this code indicates that one or more of the flags of a requested DFT was invalid. When returned by one of the execute functions it indicates that the plan was somehow invalid or not defined (NULL).
- CSDFT_INVALID_SYMBOL Set as a result of one of the User Planner functions not finding the symbol in the loaded CSX file on the Advance board. Usually this is the result of the wrong path being set in the CS_FF_CSX environment variable.
- CSDFT_INVALID_SIZE Set as a result of an unsupported size being specified as part of the plan. This is most likely the result of a size being set that is not currently supported by the CSDFT library (currently a power of 2).
- CSDFT_INVALID_FLAG Set as a result of incompatible options being chosen when creating a plan. An example of this is choosing different source and destination formats.
- CSDFT_INVALID_PARAMETER Set as a result of incompatible parameters being specified when calling a function.
- CSDFT_INVALID_USER_FUNCTION Set as a result of an invalid user function handle being specified when calling a function. An example of this is passing a null handle to *CSDFT_execute_user_function()*.

This function can be used to test for reasons behind unexpected behaviour and when used in conjunction with the *CSDFT_return_error_message()* can provide textual messages for the user at runtime.

Returns

This function returns a status of CSDFT_NO_ERROR if has been no error in the last performed function, otherwise it returns another of the listed status codes.

CSDFT_return_error_message — Returns a string describing a status value of the CSDft library

Synopsis

```
#include "csdft/csdft.h"
const char* CSDFT_return_error_message( CSDftStatus status );
```

Description

This function returns a string representing a textual description of the the current status code for the CSDft library.

The following are the status strings returned by the routine.

- "No error" Corresponds to the CSDFT_NO_ERROR status code.
- "Error: Allocating memory" Corresponds to the CSDFT_MEMORY_ALLOCATION status code.
- "Error: Invalid plan" Corresponds to the CSDFT_INVALID_PLAN status code.
- "Error: Invalid symbol" Corresponds to the CSDFT_INVALID_SYMBOL status code.
- "Error: Board not found or unrecoverable" Corresponds to the CSDFT_BOARD_ERROR status code.
- "Error: Invalid board" Corresponds to the CSDFT_INVALID_BOARD status code.
- "Error: Invalid flag" Corresponds to the CSDFT_INVALID_FLAG status code.
- "Error: Invalid size" Corresponds to the CSDFT_INVALID_SIZE status code.
- "Error: Invalid processor" Corresponds to the CSDFT_INVALID_PROCESSOR status code.
- "Error: Invalid memory descriptor" Corresponds to the CSDFT_INVALID_MEMORY_DESCRIPTOR status code.

Returns

This function returns a const char * string which corresponds to the status code used as an argument.

CSDFT_get_null_descriptor — Return a special null memory descriptor

Synopsis

```
#include "csdft/csdft.h"  
CSDftMemoryDescriptor CSDFT_get_null_descriptor( void ) ;
```

Description

This function returns a null memory descriptor which is a special type of memory descriptor that indicates to a routine to which it is passed that the parameter is deliberately set to null. This helps with parameter error checking for the library.

For example, a null memory descriptor may be used when executing a user function, when one of the parameter sizes is zero.

Returns

This function returns a handle to a CSDftMemoryDescriptor on success and a NULL handle on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ALLOCATION_ERROR as required.

CSDFT_malloc_host — Allocates a memory descriptor encapsulating a data buffer allocated on the host

Synopsis

```
#include "csdft/csdft.h"
CSDftMemoryDescriptor CSDFT_malloc_host( unsigned int size_in_bytes ) ;
```

Description

This function returns a CSDftMemoryDescriptor, where the data buffer is allocated on the host by the routine. The parameter is the size in bytes of the buffer required. If the function cannot allocate memory for the buffer for any reason it returns a NULL handle and sets the global status code to be CSDFT_MEMORY_ALLOCATION_ERROR. In order to safely access the buffer in the case where the *CSDFT_malloc_host()* function has created the buffer, you should use the *CSDFT_memory_descriptor_to_system()* call to provide a pointer to the buffer.

Returns

This function returns a handle to a CSDftMemoryDescriptor on success and a NULL handle on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ALLOCATION_ERROR as required.

CSDFT_memory_descriptor_to_system — Given a host based memory descriptor returns a pointer to a buffer in system memory

Synopsis

```
#include "csdft/csdft.h"
void * CSDFT_memory_descriptor_to_system( CSDftMemoryDescriptor handle ) ;
```

Description

This function returns a void * pointer to the data buffer that is encapsulated by the CSDftMemoryDescriptor. The parameter is a handle to a CSDftMemoryDescriptor. If the function cannot return a valid pointer to system allocated data, for example if a board based memory descriptor was passed in, then it returns a NULL pointer and sets the global status code to be CSDFT_MEMORY_ERROR.

Returns

This function returns a pointer to system allocated data on success and a NULL pointer on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_INVALID_MEMORY_DESCRIPTOR as required.

CSDFT_system_to_memory_descriptor — Returns a host based memory descriptor encapsulating user provided source data.

Synopsis

```
#include "csdft/csdft.h"
CSDftMemoryDescriptor CSDFT_system_to_memory_descriptor( void * source_data, unsigned
int size_in_bytes ) ;
```

Description

This function returns a CSDftMemoryDescriptor where the data buffer was pre-allocated on the host. The parameters are: a pointer to the array of data already allocated on the host and the size in bytes of the buffer. If the function cannot for some reason create a Memory Descriptor it returns a NULL handle and sets the global status code to be CSDFT_MEMORY_ERROR.

Returns

This function returns a handle to a CSDftMemoryDescriptor on success and a NULL handle on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ALLOCATION_ERROR as required.

CSDFT_system_split_array_to_memory_descriptor — **Creates a host based memory descriptor encapsulating two user provided source data arrays**

Synopsis

```
#include "csdft/csdft.h"
CSDftMemoryDescriptor CSDFT_system_split_array_to_memory_descriptor( void *
source_data_real, void * source_data_imag, unsigned int size_in_bytes )
```

Description

This function returns a CSDftMemoryDescriptor encapsulating the separate real and imaginary data buffers which are pre-allocated on the host. The parameters are: two pointers to arrays of data already allocated on the host (representing real and imaginary parts) and the size in bytes of the data. If the function cannot for some reason create a Memory Descriptor it returns a NULL handle and sets the global status code to be CSDFT_MEMORY_ALLOCATION_ERROR.

Returns

This function returns a handle to a CSDftMemoryDescriptor on success and a NULL handle on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ALLOCATION_ERROR as required.

CSDFT_processor_to_memory_descriptor — Returns a board based memory descriptor encapsulating user provided source data.

Synopsis

```
#include "csdft.h"
CSDftMemoryDescriptor CSDFT_processor_to_memory_descriptor( unsigned int source_data,
unsigned int size_in_bytes, unsigned int board_instance, unsigned int
processor_instance );
```

Description

This function returns a CSDftMemoryDescriptor where the data buffer was pre-allocated on the board. The parameters are: a pointer to the array of data already allocated on the board, the size in bytes of the buffer, the board number (if only one board is in use, this will be 0) and the processor the memory should be allocated on. If the function cannot for some reason create a Memory Descriptor it returns a NULL handle and sets the global status code to be CSDFT_MEMORY_ALLOCATION_ERROR.

Returns

This function returns a handle to a CSDftMemoryDescriptor on success and a NULL handle on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ALLOCATION_ERROR as required.

CSDFT_memory_descriptor_copy — Copies a memory descriptor

Synopsis

```
#include "csdft/csdft.h"
CSDFTStatus CSDFT_memory_descriptor_copy( CSDftMemoryDescriptor source,
CSDftMemoryDescriptor destination, unsigned int size_in_bytes );
```

Description

This function copies the *size_in_bytes* amount of data from the memory descriptor *source* to the memory descriptor *destination*. If *destination* is not the same or bigger size as *size_in_bytes*, a memory descriptor error will be returned. This function is especially useful when copying to and from the board, as the function will ensure previous operations have completed.

Returns

This function returns a CSDFTStatus code.

CSDFT_free — Deallocates a memory descriptor created by the CSDft library

Synopsis

```
#include "csdft/csdft.h"  
CSDftStatus CSDFT_free( CSDftMemoryDescriptor handle ) ;
```

Description

This function deallocates a memory descriptor created by one of a number of routines provided by the CSDft library. The function is passed a handle to a CSDftMemoryDescriptor and returns a CSDftStatus code indicating success or failure. In the case of the Memory Descriptor object being deallocated without any problems this will be CSDFT_NO_ERROR and in the case where there was a problem, such as trying to deallocate an already freed object, the status code will be CSDFT_MEMORY_ERROR.

Returns

This function returns a status of CSDFT_NO_ERROR or CSDFT_MEMORY_ERROR.

CSDFT_ilog2 — Returns the integer log to the base two value of the input

CSDFT_bitreverse_1D_c — Returns a bitreversed version of the single precision input data

CSDFT_bitreverse_1D_z — Returns a bitreversed version of the double precision input data

CSDFT_optimal_to_natural_2D_c — Transforms the input from board optimal order to natural order for single precision data

CSDFT_optimal_to_natural_2D_z — Transforms the input from board optimal order to natural order for double precision data

CSDFT_natural_to_optimal_2D_c — Transforms the input from natural order to board optimal order for single precision data

CSDFT_natural_to_optimal_2D_z — Transforms the input from natural order to board optimal order for double precision data

Synopsis

```
#include "csdft/csdft_support.h"
unsigned int CSDFT_ilog2          ( unsigned int n ) ;
CSDftStatus CSDFT_bitreverse_1D_c ( void      *data,
                                     unsigned int x_size,
                                     unsigned int number_of_arrays ) ;
CSDftStatus CSDFT_bitreverse_1D_z ( void      *data,
                                     unsigned int x_size,
                                     unsigned int number_of_arrays ) ;
CSDftStatus CSDFT_optimal_to_natural_2D_c( void      *data,
                                     unsigned int x_size,
                                     unsigned int y_size,
                                     unsigned int number_of_arrays ) ;
CSDftStatus CSDFT_optimal_to_natural_2D_z( void      *data,
                                     unsigned int x_size,
                                     unsigned int y_size,
                                     unsigned int number_of_arrays ) ;
CSDftStatus CSDFT_natural_to_optimal_2D_c( void      *data,
                                     unsigned int x_size,
                                     unsigned int y_size,
                                     unsigned int number_of_arrays ) ;
CSDftStatus CSDFT_natural_to_optimal_2D_z( void      *data,
                                     unsigned int x_size,
                                     unsigned int y_size,
                                     unsigned int number_of_arrays ) ;
```

Description

These functions perform bitreversal and conversion to and from natural and optimal board format for the input data. Currently supports conversion only for 1D and 2D complex interleaved arrays.

For any dft there exists an optimal input or output.

- 1D forward output and 1D backward input:- values are in bit reversed order
- 2D forward output and 2D backward input:- values are in natural column major order, with the contents of each column in bit reversed order. (i.e. column 0, column 1, column 2)

For the 1D bitreverse routines the `x_size` parameter represents the number of elements within each array, and the `number_of_arrays` parameter describes the number of individual arrays of data to be operated upon.

For the 2D routines, the `x_size` parameter represents the number of columns and the `y_size` parameter represents the number of rows.

Returns

This function returns a status of `CSDFT_NO_ERROR` if has been no error in the last performed function, otherwise it returns another of the library status codes.

CSDFT_get_symbol_value — Obtains the Board memory address of a specified symbol (on the first device of an Advance Board)

Synopsis

```
#include "csdft/csdft.h"
CSDftStatus CSDFT_get_symbol_value( const char *symbol, unsigned int
*symbol_address, unsigned int board_instance, unsigned int processor_instance );
```

Description

This function is used to obtain the address of a symbol in a CSX file. This address is then used when creating a memory descriptor on the board, using *CSDFT_processor_to_memory_descriptor()*. The parameters are: *symbol* is the name of the symbol within the CSX executable, *symbol_address* is a pointer to an integer which will be populated with the symbol address, *board_instance* is the board number (starting at 0), and *processor_instance* is the processor number (again, starting at 0).

EXAMPLES

In the CSX executable on the ClearSpeed Advance board:

```
double blank_fft[1024][1024][2];
```

In the host executable:

```
unsigned int address = 0;
CSDFTMemoryDescriptor mem_desc;
/* Get the address of the symbol */
CSDFT_get_symbol_value("blank_fft", &int,0,0);
/* Allocate memory on board 0, processor 0 */
mem_desc = CSDFT_processor_to_memory_descriptor(int, sizeof(Double)*1024*1024*2, 0,0);
```

Returns

This function returns a status of CSDFT_NO_ERROR if has been no error in the last performed function, otherwise it returns another of the library status codes.

CSDFT_get_csapi_handle — Obtains the CSAPI handle for the first instance of an Advance Board in a system used by the CSDft library

Synopsis

```
#include "csdft/csdft_support.h"
CSDftStatus CSDFT_get_csapi_handle      ( struct CSAPIState  **handle ) ;
CSDftStatus CSDFT_get_csapi_handle_board( struct CSAPIState  **handle,
                                         unsigned int  board_instance ) ;
```

Description

These functions allow the user to obtain a valid CSAPI handle to a particular Advance Board. These functions should be used if the developer wishes to access the low level CSAPI (which is not recommended), and otherwise should be avoided.

Returns

This function returns a status of CSDFT_NO_ERROR if has been no error in the last performed function, otherwise it returns CSDFT_INVALID_BOARD.

3 ClearSpeed DFT Board Library

This section documents the ClearSpeed DFT board-side interface.

CSDFT_create_plan_1d, CSDFT_create_plan_2d, CSDFT_create_plan_3d — Create a 1D, 2D or 3D plan for a DFT

Synopsis

```
#include "csdft.h"
CSDftPlan CSDFT_create_plan_1d( CSDftFlags flags, unsigned int size, unsigned int
num_ffts, double scale );
CSDftPlan CSDFT_create_plan_2d( CSDftFlags flags, unsigned int size_x, unsigned int
size_y, unsigned int num_ffts, double scale );
CSDftPlan CSDFT_create_plan_3d( CSDftFlags flags, unsigned int size_x, unsigned int
size_y, unsigned int size_z, unsigned int num_ffts, double scale );
```

Description

These functions create plans for 1, 2 and 3 dimensional DFTs, which can then be executed using the CSDFT_execute_dft routine.

This library supports:

1D DFT

- Power of two sizes from 128 to 8192
- Single and Double precision
- Complex to Complex (forward and backward)
- Interleaved

2D DFT

- Power of two sizes from 128 to 2048 (must be square)
- Single and Double precision
- Real and Complex types:
 - Complex to Complex (forward and backward)
 - Real to Complex (forward)
 - Complex to Real (backward)
- Interleaved and Split Array inputs
- Natural and Optimal order

3D DFT

- 128 size only
- Single and Double precision
- Complex to Complex (forward and backward)
- Interleaved
- Natural order only

The complex data generated from a real to complex forward 2D DFT is a partial representation of the full complex output. This is due to the symmetry of the resulting data. Only the first $(n/2)+1$ columns are represented; in

natural mode this means each row has only $(n/2)+1$ samples, with no padding at the end of the row. In optimal mode, only the first $(n/2)+1$ columns are output. When doing a backward complex to real DFT, this half-complex format is assumed.

It is possible to create multiple plans and relate them to multiple DFT computations.

The ClearSpeed DFT library performs the following forward 1D FFT:

$$Y[k]=\text{Sum}\{j=0,n-1\}(X[j]*\exp(-2*\pi*j*k*\text{sqrt}(-1)/n))$$

The ClearSpeed DFT library performs the following backward 1D DFT:

$$Y[k]=\text{Sum}\{j=0,n-1\}(X[j]*\exp(2*\pi*j*k*\text{sqrt}(-1)/n))$$

All the various configuration parameters related to a DFT are contained in a plan object in the CSDft API. A descriptor is created by the following API calls for 1d, 2d and 3D versions of `CSDFT_create_plan_<1D|2D|3D>`.

Returns

This function returns a plan if there are no errors found in the supplied flags or parameters, otherwise it returns NULL. If there is an error, then the global status will be set to indicate an error, which can be retrieved by using `CSDFT_get_status()`.

CSDFT_execute_dft — Synchronous Execution of a DFT

Synopsis

```
#include "csdft.h"
CSDftStatus CSDFT_execute_dft( const CSDftPlan plan, CSDftMemoryDescriptor source,
CSDftMemoryDescriptor destination );
```

Description

This function performs the DFT as described by the plan parameter, using data specified by the source and destination data pointers. This function is a synchronous interface and will only return when the DFT has completed.

DFTs within the source and destination data are expected to be stored contiguously.

Returns

Returns a status code of CSDFT_NO_ERROR if the function is successful, CSDFT_INVALID_PLAN if the plan passed as a parameter is not valid (such as the handle being NULL), CSDFT_INVALID_MEMORY_DESCRIPTOR if the handles for the memory descriptors are invalid or CSDFT_BOARD_ERROR if the library cannot connect to an Advance Card.

CSDFT_create_convolution_plan_1d, CSDFT_create_convolution_plan_2d — Create a 1D or 2D plan for a convolution operation

Synopsis

```
#include "csdft.h"
CSDftPlan CSDFT_create_convolution_plan_1d(CSDftFlags flags, unsigned int size,
unsigned int num_ffts, double scale);
CSDftPlan CSDFT_create_convolution_plan_2d(CSDftFlags flags, unsigned int size_x,
unsigned int size_y, unsigned int num_ffts, double scale);
```

Description

Performing a convolution is similar to any other DFT - a plan should be created, specifying the correct flags and then executed. As usual, flags are ORed together - however one flag of specific note is CSDFT_TRANSFORM_KERNEL. If set, the kernel will undergo a DFT before being convolved with the transformed source. Otherwise, the kernel will be convolved "as is" with the transformed source. By default, the kernel is assumed to not need to be pre-transformed, and, in this implementation, the kernel is assumed to be the result of a single DFT. Therefore, the same kernel will be reused if *num_ffts* is greater than 1. The kernel must be pretransformed when the convolution is on the board.

To perform a convolution with no apparent scaling, the scale parameter should be $1/n$ for both 1D and 2D, where 'n' is the size of the DFT to the power of the dimension. For example, scale would be $1/256$ when performing a 1D dft with 256 samples and $1/(256*256)$ when performing a 256x256 2D.

The complex data generated from a real to complex forward 2D DFT is a partial representation of the full complex output. This is due to the symmetry of the resulting data. Only the first $(n/2)+1$ columns are represented; in natural mode this means each row has only $(n/2)+1$ samples, with no padding at the end of the row. In optimal mode, only the first $n/2+1$ columns are output. When performing a backward complex to real DFT, this half-complex format is assumed. The kernel MUST be in optimal order when performing convolutions.

Returns

This function returns a plan if there are no errors found in the supplied flags or parameters, otherwise it returns NULL. If there is an error, then the global status will be set to indicate an error, which can be retrieved by using *CSDFT_get_status()*.

EXAMPLES

```
/*
 * To create a plan for a 1D, forward, single precision, optimal convolution of size 2048:
 */
CSDftPlan plan;
unsigned int flags = 0;
flags = CSDFT_FORWARD | CSDFT_SRC_SINGLE | CSDFT_OPTIMAL_ORDER;
plan = CSDFT_create_convolution_plan_1d(flags, 2048, 1, 1/2048);
/*
 * To create a plan for a 2D, forward, double precision, optimal, complex convolution of
 * size 128:
 */
CSDftPlan plan;
unsigned int flags = 0;
flags = CSDFT_FORWARD | CSDFT_SRC_DOUBLE | CSDFT_OPTIMAL_ORDER | CSDFT_SRC_COMPLEX;
plan = CSDFT_create_convolution_plan_2d(flags, 128, 1, 1/(128*128));
```

CSDFT_execute_convolution — Synchronous Execution of a convolution

Synopsis

```
#include "csdft.h"
CSDftStatus CSDFT_execute_convolution( const CSDftPlan plan, CSDftMemoryDescriptor
kernel, CSDftMemoryDescriptor source, CSDftMemoryDescriptor destination ) ;
```

Description

This function uses memory descriptors to specify the memory locations of the kernel, source and destination data. A memory descriptor is an opaque data structure which represents data, both in host and board space. More information on this structure is available in the Introduction section.

If the convolution is performed correctly and no errors occur, the solution will be in the memory space referenced by the destination memory descriptors.

Returns

Returns a status code of CSDFT_NO_ERROR if the function is successful, CSDFT_INVALID_PLAN if the plan passed as a parameter is not valid (such as the handle being NULL), CSDFT_INVALID_MEMORY_DESCRIPTOR if the handles for the memory descriptors are invalid or CSDFT_BOARD_ERROR if the library cannot connect to an Advance Board.

CSDFT_free_plan — Free a DFT plan

Synopsis

```
#include "csdft.h"  
CSDftStatus CSDFT_free_plan( CSDftPlan plan );
```

Description

This function frees up the previously created plan.

Returns

This function frees all associated resources for a plan and returns a CSDftStatus result of CSDFT_NO_ERROR or CSDFT_INVALID_MEMORY_DESCRIPTOR if *plan* is a null pointer.

CSDFT_get_status — Returns the current status for the CSDft library

Synopsis

```
#include "csdft.h"
CSDftStatus CSDFT_get_status( void );
```

Description

This function returns the current status code for the CSDft library. Specifically it returns the status code of the last executed CSDft library function. All Plan and Execute functions will either set the global status code or return their status code as a return value as their prototype dictates.

The following are the valid status codes which will be returned or set by CSDft library functions.

- CSDFT_NO_ERROR No current error detected.
- CSDFT_INVALID_MEMORY_DESCRIPTOR When set by one of the Execute functions or one of the memory manipulation functions this code indicates that one of the pointer parameters to the function were NULL.
- CSDFT_MEMORY_ALLOCATION_ERROR When set by one of the Memory Descriptor functions that allocate memory this code indicates that a memory allocation error has been encountered either on the host or on the Advance board.
- CSDFT_BOARD_ERROR This code can be set as a result of the CSDft library not being able to find an Advance board to connect to. If there is definitely a board present in the computer, typing "csreset -v" into your command prompt should identify or fix the problem. Alternatively, this error is set due to an internal, unrecoverable error on the board or in the driver.
- CSDFT_INVALID_BOARD This code is returned when an operation is adjudged to be sent to a board that does not exist.
- CSDFT_INVALID_PROCESSOR This code is returned when an operation is adjudged to be sent to a processor that does not exist on a board that does exist.
- CSDFT_INVALID_PLAN When set by one of the planner functions this code indicates that one or more of the flags of a requested DFT was invalid. When returned by one of the execute functions it indicates that the plan was somehow invalid or not defined (NULL).
- CSDFT_INVALID_SIZE Set as a result of an unsupported size being specified as part of the plan. This is most likely the result of a size being set that is not currently supported by the CSDFT library (currently a power of 2).
- CSDFT_INVALID_FLAG Set as a result of incompatible options being chosen when creating a plan. An example of this is choosing different source and destination formats.
- CSDFT_INVALID_PARAMETER Set as a result of incompatible parameters being specified when calling a function.
- CSDFT_INVALID_USER_FUNCTION Set as a result of an invalid user function handle being specified when calling a function. An example of this is passing a null handle to *CSDFT_execute_user_function()*.

This function can be used to test for reasons behind unexpected behaviour and when used in conjunction with the *CSDFT_return_error_message()* can provide textual messages for the user at runtime.

Returns

This function returns a status of CSDFT_NO_ERROR if has been no error in the last performed function, otherwise it returns another of the listed status codes.

CSDFT_board_to_memory_descriptor — Returns a board based memory descriptor encapsulating user provided source data.

Synopsis

```
#include "csdft.h"
CSDftMemoryDescriptor CSDFT_board_to_memory_descriptor( void *source_data, unsigned int
size_in_bytes ) ;
```

Description

This function returns a CSDftMemoryDescriptor where the data buffer was pre-allocated on the board. The parameters are: a pointer to the array of data already allocated on the host and the size in bytes of the buffer. If the function cannot for some reason create a Memory Descriptor it returns a NULL handle and sets the global status code to be CSDFT_MEMORY_ERROR.

Returns

This function returns a handle to a CSDftMemoryDescriptor on success and a NULL handle on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ERROR as required.

CSDFT_memory_descriptor_to_board — Given a board based memory descriptor returns a pointer to a buffer in board memory

Synopsis

```
#include "csdft.h"
void * CSDFT_memory_descriptor_to_board( CSDftMemoryDescriptor handle ) ;
```

Description

This function returns a void * pointer to the data buffer that is encapsulated by the CSDftMemoryDescriptor. The parameter is a handle to a CSDftMemoryDescriptor. If the function cannot for some reason return a valid pointer to system allocated data then it returns a NULL pointer and sets the global status code to be CSDFT_MEMORY_ERROR.

Returns

This function returns a pointer to system allocated data on success and a NULL pointer on failure as well as setting the global status code to CSDFT_NO_ERROR or CSDFT_MEMORY_ERROR as required.

CSDFT_free — Deallocates a memory descriptor created by the CSDft library

Synopsis

```
#include "csdft.h"  
CSDftStatus CSDFT_free( CSDftMemoryDescriptor handle ) ;
```

Description

This function deallocates a memory descriptor created by one of a number of routines provided by the CSDft library. The function is passed a handle to a CSDftMemoryDescriptor and returns a CSDftStatus code indicating success or failure. In the case of the Memory Descriptor object being deallocated without any problems this will be CSDFT_NO_ERROR and in the case where there was a problem, such as trying to deallocate an already freed object, the status code will be CSDFT_MEMORY_ERROR.

Returns

This function returns a status of CSDFT_NO_ERROR or CSDFT_MEMORY_ERROR.