



CUFFT LIBRARY USER'S GUIDE

DU-06707-001_v9.2 | April 2018



TABLE OF CONTENTS

Chapter 1. Introduction.....	1
Chapter 2. Using the cuFFT API.....	3
2.1. Accessing cuFFT.....	4
2.2. Fourier Transform Setup.....	5
2.2.1. Free memory requirement.....	6
2.3. Fourier Transform Types.....	6
2.3.1. Half precision cuFFT Transforms.....	6
2.4. Data Layout.....	7
2.5. Multidimensional Transforms.....	8
2.6. Advanced Data Layout.....	9
2.7. Streamed cuFFT Transforms.....	10
2.8. Multiple GPU cuFFT Transforms.....	10
2.8.1. Plan Specification and Work Areas.....	11
2.8.2. Helper Functions.....	12
2.8.3. Multiple GPU 2D and 3D Transforms on Permuted Input.....	12
2.8.4. Supported Functionality.....	13
2.9. cuFFT Callback Routines.....	14
2.9.1. Overview of the cufft Callback Routine Feature.....	14
2.9.2. Specifying Load and Store Callback Routines.....	15
2.9.3. Callback Routine Function Details.....	16
2.9.4. Coding Considerations for the cuFFT Callback Routine Feature.....	17
2.10. Thread Safety.....	18
2.11. Static Library and Callback Support.....	18
2.11.1. Static library without callback support.....	20
2.12. Accuracy and Performance.....	20
2.13. Caller Allocated Work Area Support.....	21
Chapter 3. cuFFT API Reference.....	23
3.1. Return value cufftResult.....	23
3.2. cuFFT Basic Plans.....	24
3.2.1. Function cufftPlan1d().....	24
3.2.2. Function cufftPlan2d().....	24
3.2.3. Function cufftPlan3d().....	25
3.2.4. Function cufftPlanMany().....	26
3.3. cuFFT Extensible Plans.....	27
3.3.1. Function cufftCreate().....	27
3.3.2. Function cufftMakePlan1d().....	27
3.3.3. Function cufftMakePlan2d().....	28
3.3.4. Function cufftMakePlan3d().....	29
3.3.5. Function cufftMakePlanMany().....	30
3.3.6. Function cufftMakePlanMany64().....	32

3.3.7. Function <code>cufftXtMakePlanMany()</code>	33
3.4. cuFFT Estimated Size of Work Area.....	35
3.4.1. Function <code>cufftEstimate1d()</code>	35
3.4.2. Function <code>cufftEstimate2d()</code>	36
3.4.3. Function <code>cufftEstimate3d()</code>	36
3.4.4. Function <code>cufftEstimateMany()</code>	37
3.5. cuFFT Refined Estimated Size of Work Area.....	38
3.5.1. Function <code>cufftGetSize1d()</code>	38
3.5.2. Function <code>cufftGetSize2d()</code>	39
3.5.3. Function <code>cufftGetSize3d()</code>	40
3.5.4. Function <code>cufftGetSizeMany()</code>	41
3.5.5. Function <code>cufftGetSizeMany64()</code>	42
3.5.6. Function <code>cufftXtGetSizeMany()</code>	43
3.6. Function <code>cufftGetSize()</code>	44
3.7. cuFFT Caller Allocated Work Area Support.....	45
3.7.1. Function <code>cufftSetAutoAllocation()</code>	45
3.7.2. Function <code>cufftSetWorkArea()</code>	46
3.7.3. Function <code>cufftXtSetWorkAreaPolicy()</code>	46
3.8. Function <code>cufftDestroy()</code>	47
3.9. cuFFT Execution.....	47
3.9.1. Functions <code>cufftExecC2C()</code> and <code>cufftExecZ2Z()</code>	47
3.9.2. Functions <code>cufftExecR2C()</code> and <code>cufftExecD2Z()</code>	48
3.9.3. Functions <code>cufftExecC2R()</code> and <code>cufftExecZ2D()</code>	49
3.9.4. Function <code>cufftXtExec()</code>	49
3.9.5. Function <code>cufftXtExecDescriptor()</code>	50
3.10. cuFFT and Multiple GPUs.....	51
3.10.1. Function <code>cufftXtSetGPUs()</code>	51
3.10.2. Function <code>cufftXtSetWorkArea()</code>	52
3.10.3. cuFFT Multiple GPU Execution.....	52
3.10.3.1. Functions <code>cufftXtExecDescriptorC2C()</code> and <code>cufftXtExecDescriptorZ2Z()</code>	52
3.10.4. Memory Allocation and Data Movement Functions.....	53
3.10.4.1. Function <code>cufftXtMalloc()</code>	53
3.10.4.2. Function <code>cufftXtFree()</code>	54
3.10.4.3. Function <code>cufftXtMemcpy()</code>	54
3.10.5. General Multiple GPU Descriptor Types.....	55
3.10.5.1. <code>cudaXtDesc</code>	56
3.10.5.2. <code>cudaLibXtDesc</code>	56
3.11. cuFFT Callbacks.....	56
3.11.1. Function <code>cufftXtSetCallback()</code>	56
3.11.2. Function <code>cufftXtClearCallback()</code>	57
3.11.3. Function <code>cufftXtSetCallbackSharedSize()</code>	57
3.12. Function <code>cufftSetStream()</code>	58
3.13. Function <code>cufftGetVersion()</code>	58

3.14. Function <code>cufftGetProperty()</code>	59
3.15. cuFFT Types.....	59
3.15.1. Parameter <code>cufftType</code>	59
3.15.2. Parameters for Transform Direction.....	59
3.15.3. Type definitions for callbacks.....	60
3.15.4. Other cuFFT Types.....	60
3.15.4.1. <code>cufftHandle</code>	60
3.15.4.2. <code>cufftReal</code>	61
3.15.4.3. <code>cufftDoubleReal</code>	61
3.15.4.4. <code>cufftComplex</code>	61
3.15.4.5. <code>cufftDoubleComplex</code>	61
3.16. Common types.....	61
3.16.1. <code>cudaDataType</code>	61
3.16.2. <code>libraryPropertyType</code>	62
Chapter 4. cuFFT Code Examples.....	63
4.1. 1D Complex-to-Complex Transforms.....	64
4.2. 1D Real-to-Complex Transforms.....	65
4.3. 2D Complex-to-Real Transforms.....	66
4.4. 3D Complex-to-Complex Transforms.....	67
4.5. 2D Advanced Data Layout Use.....	68
4.6. 3D Complex-to-Complex Transforms using Two GPUs.....	69
4.7. 1D Complex-to-Complex Transforms using Two GPUs with Natural Order.....	70
4.8. 1D Complex-to-Complex Convolution using Two GPUs.....	71
Chapter 5. Multiple GPU Data Organization.....	73
5.1. Multiple GPU Data Organization for Batched Transforms.....	73
5.2. Multiple GPU Data Organization for Single 2D and 3D Transforms.....	73
5.3. Multiple-GPU Data Organization for Single 1D Transforms.....	74
Chapter 6. FFTW Conversion Guide.....	79
Chapter 7. FFTW Interface to cuFFT.....	80
Chapter 8. Deprecated Functionality.....	83

Chapter 1.

INTRODUCTION

This document describes cuFFT, the NVIDIA® CUDA™ Fast Fourier Transform (FFT) product. It consists of two separate libraries: cuFFT and cuFFTW. The cuFFT library is designed to provide high performance on NVIDIA GPUs. The cuFFTW library is provided as a porting tool to enable users of FFTW to start using NVIDIA GPUs with a minimum amount of effort.

The FFT is a divide-and-conquer algorithm for efficiently computing discrete Fourier transforms of complex or real-valued data sets. It is one of the most important and widely used numerical algorithms in computational physics and general signal processing. The cuFFT library provides a simple interface for computing FFTs on an NVIDIA GPU, which allows users to quickly leverage the floating-point power and parallelism of the GPU in a highly optimized and tested FFT library.

The cuFFT product supports a wide range of FFT inputs and options efficiently on NVIDIA GPUs. This version of the cuFFT library supports the following features:

- ▶ Algorithms highly optimized for input sizes that can be written in the form $2^a \times 3^b \times 5^c \times 7^d$. In general the smaller the prime factor, the better the performance, i.e., powers of two are fastest.
- ▶ An $O(n \log n)$ algorithm for every input data size
- ▶ Half-precision (16-bit floating point), single-precision (32-bit floating point) and double-precision (64-bit floating point). Transforms of lower precision have higher performance.
- ▶ Complex and real-valued input and output. Real valued input or output require less computations and data than complex values and often have faster time to solution. Types supported are:
 - ▶ C2C - Complex input to complex output
 - ▶ R2C - Real input to complex output
 - ▶ C2R - Symmetric complex input to real output
- ▶ 1D, 2D and 3D transforms
- ▶ Execution of multiple 1D, 2D and 3D transforms simultaneously. These batched transforms have higher performance than single transforms.
- ▶ In-place and out-of-place transforms
- ▶ Arbitrary intra- and inter-dimension element strides (strided layout)

- ▶ FFTW compatible data layout
- ▶ Execution of transforms across multiple GPUs
- ▶ Streamed execution, enabling asynchronous computation and data movement

The cuFFTW library provides the FFTW3 API to facilitate porting of existing FFTW applications.

Chapter 2.

USING THE CUFFT API

This chapter provides a general overview of the cuFFT library API. For more complete information on specific functions, see [cuFFT API Reference](#). Users are encouraged to read this chapter before continuing with more detailed descriptions.

The Discrete Fourier transform (DFT) maps a complex-valued vector x_k (*time domain*) into its *frequency domain representation* given by:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{kn}{N}}$$

where X_k is a complex-valued vector of the same size. This is known as a *forward* DFT. If the sign on the exponent of e is changed to be positive, the transform is an *inverse* transform. Depending on N , different algorithms are deployed for the best performance.

The cuFFT API is modeled after [FFTW](#), which is one of the most popular and efficient CPU-based FFT libraries. cuFFT provides a simple configuration mechanism called a *plan* that uses internal building blocks to optimize the transform for the given configuration and the particular GPU hardware selected. Then, when the *execution* function is called, the actual transform takes place following the plan of execution. The advantage of this approach is that once the user creates a plan, the library retains whatever state is needed to execute the plan multiple times without recalculation of the configuration. This model works well for cuFFT because different kinds of FFTs require different thread configurations and GPU resources, and the plan interface provides a simple way of reusing configurations.

Computing a number **BATCH** of one-dimensional DFTs of size **NX** using cuFFT will typically look like this:

```
#define NX 256
#define BATCH 10
#define RANK 1
...
{
    cufftHandle plan;
    cufftComplex *data;
    ...
    cudaMalloc((void**)&data, sizeof(cufftComplex)*NX*BATCH);
    cufftPlanMany(&plan, RANK, NX, &iembed, istride, idist,
        &oembed, ostride, odist, CUFFT_C2C, BATCH);
    ...
    cufftExecC2C(plan, data, data, CUFFT_FORWARD);
    cudaDeviceSynchronize();
    ...
    cufftDestroy(plan);
    cudaFree(data);
}
```

2.1. Accessing cuFFT

The cuFFT and cuFFTW libraries are available as shared libraries. They consist of compiled programs ready for users to incorporate into applications with the compiler and linker. cuFFT can be downloaded from <http://developer.nvidia.com/cufft>. By selecting **Download CUDA Production Release** users are all able to install the package containing the CUDA Toolkit, SDK code samples and development drivers. The CUDA Toolkit contains cuFFT and the samples include **simplecuFFT**.

The Linux release for **simplecuFFT** assumes that the root install directory is **/usr/local/cuda** and that the locations of the products are contained there as follows. Modify the Makefile as appropriate for your system.

Product	Location and name	Include file
nvcc compiler	/bin/nvcc	
cuFFT library	{lib, lib64}/libcufft.so	inc/cufft.h
cuFFT library with Xt functionality	{lib, lib64}/libcufft.so	inc/cufftXt.h
cuFFTW library	{lib, lib64}/libcufftw.so	inc/cufftw.h

The most common case is for developers to modify an existing CUDA routine (for example, **filename.cu**) to call cuFFT routines. In this case the include file **cufft.h** or **cufftXt.h** should be inserted into **filename.cu** file and the library included in the link line. A single compile and link line might appear as

```
► /usr/local/cuda/bin/nvcc [options] filename.cu ... -I/usr/local/cuda/inc -L/usr/local/cuda/lib -lcufft
```

Of course there will typically be many compile lines and the compiler **g++** may be used for linking so long as the library path is set correctly.

Users of the FFTW interface (see [FFTW Interface to cuFFT](#)) should include `cufftw.h` and link with both cuFFT and cuFFTW libraries.

Functions in the cuFFT and cuFFTW library assume that the data is in GPU visible memory. This means any memory allocated by `cudaMalloc`, `cudaMallocHost` and `cudaMallocManaged` or registered with `cudaHostRegister` can be used as input, output or plan work area with cuFFT and cuFFTW functions. For the best performance input data, output data and plan work area should reside in device memory.

cuFFTW library also supports input data and output data that is not GPU visible.

2.2. Fourier Transform Setup

The first step in using the cuFFT Library is to create a plan using one of the following:

- ▶ `cufftPlan1D()` / `cufftPlan2D()` / `cufftPlan3D()` - Create a simple plan for a 1D/2D/3D transform respectively.
- ▶ `cufftPlanMany()` - Creates a plan supporting batched input and strided data layouts.
- ▶ `cufftXtMakePlanMany()` - Creates a plan supporting batched input and strided data layouts for any supported precision.

Among the plan creation functions, `cufftPlanMany()` allows use of more complicated data layouts and batched executions. Execution of a transform of a particular size and type may take several stages of processing. When a plan for the transform is generated, cuFFT derives the internal steps that need to be taken. These steps may include multiple kernel launches, memory copies, and so on. In addition, all the intermediate buffer allocations (on CPU/GPU memory) take place during planning. These buffers are released when the plan is destroyed. In the worst case, the cuFFT Library allocates space for $8 * \text{batch} * n[0] * \dots * n[\text{rank}-1]$ `cufftComplex` or `cufftDoubleComplex` elements (where `batch` denotes the number of transforms that will be executed in parallel, `rank` is the number of dimensions of the input data (see [Multidimensional Transforms](#)) and `n[]` is the array of transform dimensions) for single and double-precision transforms respectively. Depending on the configuration of the plan, less memory may be used. In some specific cases, the temporary space allocations can be as low as $1 * \text{batch} * n[0] * \dots * n[\text{rank}-1]$ `cufftComplex` or `cufftDoubleComplex` elements. This temporary space is allocated separately for each individual plan when it is created (i.e., temporary space is not shared between the plans).

The next step in using the library is to call an execution function such as `cufftExecC2C()` (see [Parameter cufftType](#)) which will perform the transform with the specifications defined at planning.

One can create a cuFFT plan and perform multiple transforms on different data sets by providing different input and output pointers. Once the plan is no longer needed, the `cufftDestroy()` function should be called to release the resources allocated for the plan.

2.2.1. Free memory requirement

The first program call to any cuFFT function causes the initialization of the cuFFT kernels. This can fail if there is not enough free memory on the GPU. It is advisable to initialize `cufft` first (e.g. by creating a plan) and then allocating memory.

2.3. Fourier Transform Types

Apart from the general complex-to-complex (C2C) transform, cuFFT implements efficiently two other types: real-to-complex (R2C) and complex-to-real (C2R). In many practical applications the input vector is real-valued. It can be easily shown that in this case the output satisfies Hermitian symmetry ($X_k = X_{N-k}^*$ where the star denotes complex conjugation). The converse is also true: for complex-Hermitian input the inverse transform will be purely real-valued. cuFFT takes advantage of this redundancy and works only on the first half of the Hermitian vector.

Transform execution functions for single and double-precision are defined separately as:

- ▶ `cufftExecC2C()` / `cufftExecZ2Z()` - complex-to-complex transforms for single/double precision.
- ▶ `cufftExecR2C()` / `cufftExecD2Z()` - real-to-complex forward transform for single/double precision.
- ▶ `cufftExecC2R()` / `cufftExecZ2D()` - complex-to-real inverse transform for single/double precision.

Each of those functions demands different input data layout (see [Data Layout](#) for details).

Functions `cufftXtExec()` and `cufftXtExecDescriptor()` can perform transforms on any of the supported types.

2.3.1. Half precision cuFFT Transforms

Half precision transforms have the following limitations:

- ▶ Minimum GPU architecture is SM_53
- ▶ Sizes are restricted to powers of two only
- ▶ Strides on the real part of real-to-complex and complex-to-real transforms are not supported
- ▶ More than one GPU is not supported
- ▶ Transforms spanning more than 4 billion elements are not supported

Please refer to `cufftXtMakePlanMany` function for plan creation details.

CUDA Toolkit provides `cuda_fp16.h` header with types and intrinsic functions for handling half precision arithmetic.

2.4. Data Layout

In the cuFFT Library, data layout depends strictly on the configuration and the transform type. In the case of general complex-to-complex transform both the input and output data shall be a `cuFFTComplex/cuFFTDoubleComplex` array in single- and double-precision modes respectively. In C2R mode an input array $(x_1, x_2, \dots, x_{\lfloor \frac{N}{2} \rfloor + 1})$ of only non-redundant complex elements is required. The output array (X_1, X_2, \dots, X_N) consists of `cuFFTReal/cuFFTDouble` elements in this mode. Finally, R2C demands an input array (X_1, X_2, \dots, X_N) of real values and returns an array $(x_1, x_2, \dots, x_{\lfloor \frac{N}{2} \rfloor + 1})$ of non-redundant complex elements.

In real-to-complex and complex-to-real transforms the size of input data and the size of output data differ. For out-of-place transforms a separate array of appropriate size is created. For in-place transforms the user should use **padded** data layout. This layout is FFTW compatible.

In the **padded** layout output signals begin at the same memory addresses as the input data. Therefore input data for real-to-complex and output data for complex-to-real must be padded.

Expected sizes of input/output data for 1-d transforms are summarized in the table below:

FFT type	input data size	output data size
C2C	\times <code>cuFFTComplex</code>	\times <code>cuFFTComplex</code>
C2R	$\lfloor \frac{N}{2} \rfloor + 1$ <code>cuFFTComplex</code>	\times <code>cuFFTReal</code>
R2C*	\times <code>cuFFTReal</code>	$\lfloor \frac{N}{2} \rfloor + 1$ <code>cuFFTComplex</code>

The real-to-complex transform is implicitly a forward transform. For an in-place real-to-complex transform where FFTW compatible output is desired, the input size must be padded to $(\lfloor \frac{N}{2} \rfloor + 1)$ complex elements. For out-of-place transforms, input and output sizes match the logical transform size N and the non-redundant size $\lfloor \frac{N}{2} \rfloor + 1$, respectively.

The complex-to-real transform is implicitly inverse. For in-place complex-to-real FFTs where FFTW compatible output is selected (default padding mode), the input size is assumed to be $\lfloor \frac{N}{2} \rfloor + 1$ `cuFFTComplex` elements. Note that in-place complex-to-real FFTs may **overwrite** arbitrary imaginary input point values when non-unit input and output strides are chosen. Out-of-place complex-to-real FFT will **overwrite** input buffer if custom strides are set by the user. For out-of-place transforms, input and output sizes match the logical transform non-redundant size $\lfloor \frac{N}{2} \rfloor + 1$ and size N , respectively.

2.5. Multidimensional Transforms

Multidimensional DFT map a d -dimensional array $x_{\mathbf{n}}$, where $\mathbf{n} = (n_1, n_2, \dots, n_d)$ into its frequency domain array given by:

$$X_{\mathbf{k}} = \sum_{n=0}^{N-1} x_{\mathbf{n}} e^{-2\pi i \frac{\mathbf{k}\mathbf{n}}{N}}$$

where $\frac{\mathbf{n}}{N} = (\frac{n_1}{N_1}, \frac{n_2}{N_2}, \dots, \frac{n_d}{N_d})$, and the summation denotes the set of nested summations

$$\sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} \dots \sum_{n_d=0}^{N_d-1}$$

cuFFT supports one-dimensional, two-dimensional and three-dimensional transforms, which can all be called by the same **cuFFTExec*** functions (see [Fourier Transform Types](#)).

Similar to the one-dimensional case, the frequency domain representation of real-valued input data satisfies Hermitian symmetry, defined as: $x_{(n_1, n_2, \dots, n_d)} = x_{(N_1-n_1, N_2-n_2, \dots, N_d-n_d)}^*$

C2R and R2C algorithms take advantage of this fact by operating only on half of the elements of signal array, namely on: $x_{\mathbf{n}}$ for

$$\mathbf{n} \in \{1, \dots, N_1\} \times \dots \times \{1, \dots, N_{d-1}\} \times \{1, \dots, \lfloor \frac{N_d}{2} \rfloor + 1\}.$$

The general rules of data alignment described in [Data Layout](#) apply to higher-dimensional transforms. The following table summarizes input and output data sizes for multidimensional DFTs:

Dims	FFT type	Input data size	Output data size
1D	C2C	N_1 cuFFTComplex	N_1 cuFFTComplex
1D	C2R	$\lfloor \frac{N_1}{2} \rfloor + 1$ cuFFTComplex	N_1 cuFFTReal
1D	R2C	N_1 cuFFTReal	$\lfloor \frac{N_1}{2} \rfloor + 1$ cuFFTComplex
2D	C2C	$N_1 N_2$ cuFFTComplex	$N_1 N_2$ cuFFTComplex
2D	C2R	$N_1 (\lfloor \frac{N_2}{2} \rfloor + 1)$ cuFFTComplex	$N_1 N_2$ cuFFTReal
2D	R2C	$N_1 N_2$ cuFFTReal	$N_1 (\lfloor \frac{N_2}{2} \rfloor + 1)$ cuFFTComplex
3D	C2C	$N_1 N_2 N_3$ cuFFTComplex	$N_1 N_2 N_3$ cuFFTComplex
3D	C2R	$N_1 N_2 (\lfloor \frac{N_3}{2} \rfloor + 1)$ cuFFTComplex	$N_1 N_2 N_3$ cuFFTReal
3D	R2C	$N_1 N_2 N_3$ cuFFTReal	$N_1 N_2 (\lfloor \frac{N_3}{2} \rfloor + 1)$ cuFFTComplex

For example, static declaration of a three-dimensional array for the output of an out-of-place real-to-complex transform will look like this:

```
cufftComplex odata[N1][N2][N3/2+1];
```

2.6. Advanced Data Layout

The advanced data layout feature allows transforming only a subset of an input array, or outputting to only a portion of a larger data structure. It can be set by calling function:

```
cufftResult cufftPlanMany(cufftHandle *plan, int rank, int *n, int *inembed,
    int istride, int idist, int *onembed, int ostride,
    int odist, cufftType type, int batch);
```

Passing **inembed** or **onembed** set to **NULL** is a special case and is equivalent to passing **n** for each. This is same as the basic data layout and other advanced parameters such as **istride** are ignored.

If the advanced parameters are to be used, then all of the advanced interface parameters must be specified correctly. Advanced parameters are defined in units of the relevant data type (**cufftReal**, **cufftDoubleReal**, **cufftComplex**, or **cufftDoubleComplex**).

Advanced layout can be perceived as an additional layer of abstraction above the access to input/output data arrays. An element of coordinates [**z**] [**y**] [**x**] in signal number **b** in the batch will be associated with the following addresses in the memory:

► 1D

```
input[ b * idist + x * istride ]
```

```
output[ b * odist + x * ostride ]
```

► 2D

```
input[ b * idist + (x * inembed[1] + y) * istride ]
```

```
output[ b * odist + (x * onembed[1] + y) * ostride ]
```

► 3D

```
input[ b * idist + ((x * inembed[1] + y) * inembed[2] + z) * istride ]
```

```
output[ b * odist + ((x * onembed[1] + y) * onembed[2] + z) * ostride ]
```

The **istride** and **ostride** parameters denote the distance between two successive input and output elements in the least significant (that is, the innermost) dimension respectively. In a single 1D transform, if every input element is to be used in the transform, **istride** should be set to 1; if every other input element is to be used in the transform, then **istride** should be set to 2. Similarly, in a single 1D transform, if it is desired to output final elements one after another compactly, **ostride** should be set to 1; if spacing is desired between the least significant dimension output data, **ostride** should be set to the distance between the elements.

The **inembed** and **onembed** parameters define the number of elements in each dimension in the input array and the output array respectively. The **inembed[rank-1]** contains the number of elements in the least significant (innermost) dimension of the input data excluding the **istride** elements; the number of total elements in the least significant dimension of the input array is then **istride*inembed[rank-1]**. The

`inembed[0]` or `onembed[0]` corresponds to the most significant (that is, the outermost) dimension and is effectively ignored since the `idist` or `odist` parameter provides this information instead. Note that the size of each dimension of the transform should be less than or equal to the `inembed` and `onembed` values for the corresponding dimension, that is $n[i] \leq \text{inembed}[i]$, $n[i] \leq \text{onembed}[i]$, where $i \in \{0, \dots, \text{rank} - 1\}$.

The `idist` and `odist` parameters indicate the distance between the first element of two consecutive batches in the input and output data.

2.7. Streamed cuFFT Transforms

Every cuFFT plan may be associated with a CUDA stream. Once so associated, all launches of the internal stages of that plan take place through the specified stream. Streaming of cuFFT execution allows for potential overlap between transforms and memory copies. (See the *NVIDIA CUDA Programming Guide* for more information on streams.) If no stream is associated with a plan, launches take place in `stream(0)`, the default CUDA stream. Note that many plan executions require multiple kernel launches.

`cufftSetStream()` returns an error in the multiple GPU case as multiple GPU plans perform operations in their own streams.

Please note that in order to overlap plans using single plan handle user needs to manage work area buffers. Each concurrent plan execution needs its exclusive work area. Work area can be set by `cufftSetWorkArea` function.

2.8. Multiple GPU cuFFT Transforms

cuFFT supports using up to eight GPUs connected to a CPU to perform Fourier Transforms whose calculations are distributed across the GPUs. An API has been defined to allow users to write new code or modify existing code to use this functionality.

Some existing functions such as the creation of a plan using `cufftCreate()` also apply in the multiple GPU case. Multiple GPU routines contain `xt` in their name.

The memory on the GPUs is managed by helper functions `cufftXtMalloc()` / `cufftXtFree()` and `cufftXtMemcpy()` using the `cudaLibXtDesc` descriptor.

Performance is a function of the bandwidth between the GPUs, the computational ability of the individual GPUs, and the type and number of FFT to be performed. The highest performance is obtained using NVLink interconnect (<http://www.nvidia.com/object/nvlink.html>). The second best option is using PCI Express 3.0 between the GPUs and ensuring that both GPUs are on the same switch. Note that multiple GPU execution is not guaranteed to solve a given size problem in a shorter time than single GPU execution.

The multiple GPU extensions to cuFFT are built on the extensible cuFFT API. The general steps in defining and executing a transform with this API are:

- ▶ `cufftCreate()` - create an empty plan, as in the single GPU case

- ▶ **cufftXtSetGPUs ()** - define which GPUs are to be used
- ▶ Optional: **cufftEstimate{1d,2d,3d,Many} ()** - estimate the sizes of the work areas required. These are the same functions used in the single GPU case although the definition of the argument **workSize** reflects the number of GPUs used.
- ▶ **cufftMakePlan{1d,2d,3d,Many} ()** - create the plan. These are the same functions used in the single GPU case although the definition of the argument **workSize** reflects the number of GPUs used.
- ▶ Optional: **cufftGetSize{1d,2d,3d,Many} ()** - refined estimate of the sizes of the work areas required. These are the same functions used in the single GPU case although the definition of the argument **workSize** reflects the number of GPUs used.
- ▶ Optional: **cufftGetSize ()** - check workspace size. This is the same function used in the single GPU case although the definition of the argument **workSize** reflects the number of GPUs used.
- ▶ Optional: **cufftXtSetWorkArea ()** - do your own workspace allocation.
- ▶ **cufftXtMalloc ()** - allocate descriptor and data on the GPUs
- ▶ **cufftXtMemcpy ()** - copy data to the GPUs
- ▶ **cufftXtExecDescriptorC2C () / cufftXtExecDescriptorZ2Z ()** - execute the plan
- ▶ **cufftXtMemcpy ()** - copy data from the GPUs
- ▶ **cufftXtFree ()** - free any memory allocated with **cufftXtMalloc ()**
- ▶ **cufftDestroy ()** - free cuFFT plan resources

2.8.1. Plan Specification and Work Areas

In the single GPU case a plan is created by a call to **cufftCreate ()** followed by a call to **cufftMakePlan* ()**. For multiple GPUs, the GPUs to use for execution are identified by a call to **cufftXtSetGPUs ()** and this must occur after the call to **cufftCreate ()** and prior to the call to **cufftMakePlan* ()**.

Note that when **cufftMakePlan* ()** is called for a single GPU, the work area is on that GPU. In a multiple GPU plan, the returned work area has multiple entries; one value per GPU. That is **workSize** points to a **size_t** array, one entry per GPU. Also the strides and batches apply to the entire plan across all GPUs associated with the plan.

Once a plan is locked by a call to **cufftMakePlan* ()**, different descriptors may be specified in calls to **cufftXtExecDescriptor* ()** to execute the plan on different data sets, but the new descriptors must use the same GPUs in the same order.

As in the single GPU case, **cufftEstimateSize{Many,1d,2d,3d} ()** and **cufftGetSize{Many,1d,2d,3d} ()** give estimates of the work area sizes required for a multiple GPU plan and in this case **workSize** points to a **size_t** array, one entry per GPU.

Similarly the actual work size returned by **cufftGetSize ()** is a **size_t** array, one entry per GPU in the multiple GPU case.

2.8.2. Helper Functions

Multiple GPU cuFFT execution functions assume a certain data layout in terms of what input data has been copied to which GPUs prior to execution, and what output data resides in which GPUs post execution. cuFFT provides functions to assist users in manipulating data on multiple GPUs. These must be called after the call to **cufftMakePlan*()**.

On a single GPU users may call **cudaMalloc()** and **cudaFree()** to allocate and free GPU memory. To provide similar functionality in the multiple GPU case, cuFFT includes **cufftXtMalloc()** and **cufftXtFree()** functions. The function **cufftXtMalloc()** returns a descriptor which specifies the location of these memories.

On a single GPU users may call **cudaMemcpy()** to transfer data between host and GPU memory. To provide similar functionality in the multiple GPU case, cuFFT includes **cufftXtMemcpy()** which allows users to copy between host and multiple GPU memories or even between the GPU memories.

All single GPU cuFFT FFTs return output the data in natural order, that is the ordering of the result is the same as if a DFT had been performed on the data. Some Fast Fourier Transforms produce intermediate results where the data is left in a permutation of the natural output. When batch is one, data is left in the GPU memory in a permutation of the natural output.

When **cufftXtMemcpy()** is used to copy data from GPU memory back to host memory, the results are in natural order regardless of whether the data on the GPUs is in natural order or permuted. Using **CUFFT_COPY_DEVICE_TO_DEVICE** allows users to copy data from the permuted data format produced after a single transform to the natural order on GPUs.

2.8.3. Multiple GPU 2D and 3D Transforms on Permuted Input

For single 2D or 3D transforms on multiple GPUs, when **cufftXtMemcpy()** distributes the data to the GPUs, the array is divided on the X axis. E.G. for two GPUs half of the X dimension points, for all Y (and Z) values, are copied to each of the GPUs. When the transform is computed, the data are permuted such that they are divided on the Y axis. I.E. half of the Y dimension points, for all X (and Z) values are on each of the GPUs.

When cuFFT creates a 2D or 3D plan for a single transform on multiple GPUs, it actually creates two plans. One plan expects input to be divided on the X axis. The other plan expects data to be divided on the Y axis. This is done because many algorithms compute a forward FFT, then perform some point-wise operation on the result, and then compute the inverse FFT. A memory copy to restore the data to the original order would be expensive. To avoid this, **cufftXtMemcpy** and **cufftXtExecDescriptor()** keep track of the data ordering so that the correct operation is used.

The ability of cuFFT to process data in either order makes the following sequence possible.

- ▶ **cufftCreate()** - create an empty plan, as in the single GPU case

- ▶ `cufftXtSetGPUs()` - define which GPUs are to be used
- ▶ `cufftMakePlan{1d,2d,3d,Many}()` - create the plan.
- ▶ `cufftXtMalloc()` - allocate descriptor and data on the GPUs
- ▶ `cufftXtMemcpy()` - copy data to the GPUs
- ▶ `cufftXtExecDescriptorC2C()/cufftXtExecDescriptorZ2Z()` - compute the forward FFT
- ▶ `userFunction()` - modify the data in the frequency domain
- ▶ `cufftXtExecDescriptorC2C()/cufftXtExecDescriptorZ2Z()` - compute the inverse FFT
- ▶ Note that it was not necessary to copy/permute the data between execute calls
- ▶ `cufftXtMemcpy()` - copy data to the host
- ▶ `cufftXtFree()` - free any memory allocated with `cufftXtMalloc()`
- ▶ `cufftDestroy()` - free cuFFT plan resources

2.8.4. Supported Functionality

Since version 7.0 only a subset of single GPU functionality is supported for multiple GPU execution.

Supported functionality:

- ▶ Plans operating on two, four or eight GPUs are supported.
- ▶ Up to eight GPUs are supported for 2D and 3D transforms or when number of batches is greater than 1.
- ▶ All GPUs must have the same CUDA architecture level.
- ▶ The GPUs must support the Unified Virtual Address Space.
- ▶ On Windows, the GPU boards must be operating in Tesla Compute Cluster (TCC) mode.
- ▶ Running cuFFT on multiple GPUs is not compatible with an application that uses the CUDA Driver API.
- ▶ Strided input and output are not supported.
- ▶ When the number of batches is 1:
 - ▶ Only **c2c** and **z2z** transform types are supported.
 - ▶ Only in-place transforms are supported.
 - ▶ For 1D transforms, the dimension must be a power of 2 greater than 64. For eight GPUs transform size must be greater than 128.
 - ▶ For 2D and 3D transforms, the dimensions must factor into primes less than or equal to 127. The X and Y dimensions of the transform must be greater than or equal to 32.

General guidelines are:

- ▶ Parameter **whichGPUs** of `cufftXtSetGPUs()` function determines ordering of the GPUs with respect to data decomposition (first data chunk is placed on GPU denoted by first element of **whichGPUs**)
- ▶ The data for the entire transform must fit within the memory of the GPUs assigned to it.
- ▶ For batch size **m** on **n** GPUs :

- ▶ The first $m \% n$ GPUs execute $\lfloor \frac{m}{n} \rfloor + 1$ transforms.
- ▶ The remaining GPUs execute $\lfloor \frac{m}{n} \rfloor$ transforms.

Batch size output differences:

Single GPU cuFFT results are always returned in natural order. When multiple GPUs are used to perform more than one transform, the results are also returned in natural order. When multiple GPUs are used to perform a single transform the results are returned in a permutation of the normal results to reduce communication time.

Number of GPUs	Number of transforms	Output Order on GPUs
One	One or multiple transforms	Natural order
Multiple	One	Permuted results
Multiple	Multiple	Natural order

To produce natural order results in GPU memory in the 1D single transform case, requires calling `cufftXtMemcpy()` with `CUFFT_COPY_DEVICE_TO_DEVICE`.

2D and 3D transforms support execution of a transform given permuted order results as input. After execution in this case, the output will be in natural order. It is also possible to use `cufftXtMemcpy()` with `CUFFT_COPY_DEVICE_TO_DEVICE` to return 2D or 3D data to natural order.

The only supported multiple GPU configurations are 2 or 4 GPUs, all with the same CUDA architecture level.

See the cuFFT Code Examples section for single GPU and multiple GPU examples.

2.9. cuFFT Callback Routines

Callback routines are user-supplied kernel routines that cuFFT will call when loading or storing data. They allow the user to do data pre- or post- processing without additional kernel calls.

2.9.1. Overview of the cuFFT Callback Routine Feature

cuFFT provides a set of APIs that allow the cuFFT user to provide CUDA functions that re-direct or manipulate the data as it is loaded prior to processing the FFT, or stored once the FFT has been done. For the load callback, cuFFT passes the callback routine the address of the input data and the offset to the value to be loaded from device memory, and the callback routine returns the value it wishes cuFFT to use instead. For the store callback, cuFFT passes the callback routine the value it has computed, along with the address of the output data and the offset to the value to be written to device memory, and the callback routine modifies the value and stores the modified result.

In order to provide a callback to cuFFT, a plan is created and configured normally using the extensible plan APIs. After the call to `cufftCreate` and `cufftMakePlan`, the user may associate a load callback routine, or a store callback routine, or both, with the plan, by calling `cufftXtSetCallback`. The caller also has the option to specify a device

pointer to an opaque structure they wish to associate with the plan. This pointer will be passed to the callback routine by the cuFFT library. The caller may use this structure to remember plan dimensions and strides, or have a pointer to auxiliary data, etc.

With some restrictions, the callback routine is allowed to request shared memory for its own use. If the requested amount of shared memory is available, cuFFT will pass a pointer to it when it calls the callback routine.

CUFFT allows for 8 types of callback routine, one for each possible combination of: load or store, real or complex, single precision or double. **It is the caller's responsibility to provide a routine that matches the function prototype for the type of routine specified.** If there is already a callback of the specified type associated with the plan, the set callback function will replace it with the new one.

The callback routine extensions to cuFFT are built on the extensible cuFFT API. The general steps in defining and executing a transform with callbacks are:

- ▶ `cufftCreate()` - create an empty plan, as in the single GPU case
- ▶ `cufftMakePlan{1d,2d,3d,Many}()` - create the plan. These are the same functions used in the single GPU case.
- ▶ `cufftXtSetCallback()` - called for load and/or store callback for this plan
- ▶ `cufftExecC2C()` *etc.* - execute the plan
- ▶ `cufftDestroy()` - free cuFFT plan resources

Callback functions are not supported on transforms with a dimension size that does not factor into primes smaller than 127. Callback functions on plans whose dimensions' prime factors are limited to 2, 3, 5, and 7 can safely call `__syncthreads()`. On other plans, results are not defined.

NOTE:The callback API is available in the statically linked cuFFT library only, and only on 64 bit LINUX operating systems.

2.9.2. Specifying Load and Store Callback Routines

In order to associate a callback routine with a plan, it is necessary to obtain a device pointer to the callback routine.

As an example, if the user wants to specify a load callback for an R2C transform, they would write the device code for the callback function, and define a global device variable that contains a pointer to the function:

```
__device__ cufftReal myOwnCallback(void *dataIn,
                                   size_t offset,
                                   void *callerInfo,
                                   void *sharedPtr) {
    cufftReal ret;
    // use offset, dataIn, and optionally callerInfo to
    // compute the return value
    return ret;
}
__device__ cufftCallbackLoadR myOwnCallbackPtr = myOwnCallback;
```

From the host side, the user then has to get the address of the callback routine, which is stored in `myOwnCallbackPtr`. This is done with `cudaMemcpyFromSymbol`, as follows:

```

cufftCallbackLoadR hostCopyOfCallbackPtr;

cudaMemcpyFromSymbol (&hostCopyOfCallbackPtr,
                      myOwnCallbackPtr,
                      sizeof(hostCopyOfCallbackPtr));

```

`hostCopyOfCallbackPtr` then contains the device address of the callback routine, that should be passed to `cufftXtSetCallback`. Note that, for multi-GPU transforms, `hostCopyOfCallbackPtr` will need to be an array of pointers, and the `cudaMemcpyFromSymbol` will have to be invoked for each GPU. Please note that `__managed__` variables are not suitable to pass to `cufftSetCallback` due to restrictions on variable usage (See the *NVIDIA CUDA Programming Guide* for more information about `__managed__` variables).

2.9.3. Callback Routine Function Details

Below are the function prototypes, and typedefs for pointers to the user supplied callback routines that cuFFT calls to load data prior to the transform.

```

typedef cufftComplex (*cufftCallbackLoadC)(void *dataIn,
                                           size_t offset,
                                           void *callerInfo,
                                           void *sharedPointer);

typedef cufftDoubleComplex (*cufftCallbackLoadZ)(void *dataIn,
                                                  size_t offset,
                                                  void *callerInfo,
                                                  void *sharedPointer);

typedef cufftReal (*cufftCallbackLoadR)(void *dataIn,
                                        size_t offset,
                                        void *callerInfo,
                                        void *sharedPointer);

typedef cufftDoubleReal (*cufftCallbackLoadD)(void *dataIn,
                                              size_t offset,
                                              void *callerInfo,
                                              void *sharedPointer);

```

Parameters for all of the load callbacks are defined as below:

- ▶ **offset**: offset of the input element from the start of output data. This is not a byte offset, rather it is the number of elements from start of data.
- ▶ **dataIn**: device pointer to the start of the input array that was passed in the `cufftExecute` call.
- ▶ **callerInfo**: device pointer to the optional caller specified data passed in the `cufftXtSetCallback` call.
- ▶ **sharedPointer**: pointer to shared memory, valid only if the user has called `cufftXtSetCallbackSharedSize()`.

Below are the function prototypes, and typedefs for pointers to the user supplied callback routines that cuFFT calls to store data after completion of the transform. Note that the store callback functions do not return a value. This is because a store callback

function is responsible not only for transforming the data as desired, but also for writing the data to the desired location. This allows the store callback to rearrange the data, for example to shift the zero frequency result to the center of the output.

```
typedef void (*cufftCallbackStoreC)(void *dataOut,
                                   size_t offset,
                                   cufftComplex element,
                                   void *callerInfo,
                                   void *sharedPointer);

typedef void (*cufftCallbackStoreZ)(void *dataOut,
                                   size_t offset,
                                   cufftDoubleComplex element,
                                   void *callerInfo,
                                   void *sharedPointer);

typedef void (*cufftCallbackStoreR)(void *dataOut,
                                   size_t offset,
                                   cufftReal element,
                                   void *callerInfo,
                                   void *sharedPointer);

typedef void (*cufftCallbackStoreD)(void *dataOut,
                                   size_t offset,
                                   cufftDoubleReal element,
                                   void *callerInfo,
                                   void *sharedPointer);
```

Parameters for all of the store callbacks are defined as below:

- ▶ **offset**: offset of the output element from the start of output data. This is not a byte offset, rather it is the number of elements from start of data.
- ▶ **dataOut**: device pointer to the start of the output array that was passed in the `cufftExecute` call.
- ▶ **element**: the real or complex result computed by CUFFT for the element specified by the offset argument.
- ▶ **callerInfo**: device pointer to the optional caller specified data passed in the `cufftXtSetCallback` call.
- ▶ **sharedPointer**: pointer to shared memory, valid only if the user has called `cufftXtSetCallbackSharedSize()`.

2.9.4. Coding Considerations for the cuFFT Callback Routine Feature

cuFFT supports callbacks on all types of transforms, dimension, batch, stride between elements or number of GPUs. Callbacks are supported for transforms of single and double precision.

cuFFT supports a wide range of parameters, and based on those for a given plan, it attempts to optimize performance. The number of kernels launched, and for each of those, the number of blocks launched and the number of threads per block, will vary depending on how cuFFT decomposes the transform. For some configurations, cuFFT will load or store (and process) multiple inputs or outputs per thread. For some configurations, threads may load or store inputs or outputs in any order, and cuFFT does not guarantee that the inputs or outputs handled by a given thread will be

contiguous. These characteristics may vary with transform size, transform type (e.g. C2C vs C2R), number of dimensions, and GPU architecture. These variations may also change from one library version to the next.

cuFFT will call the load callback routine, for each point in the input, once and only once. Similarly it will call the store callback routine, for each point in the output, once and only once. If cuFFT is implementing a given FFT in multiple phases, it will only call the load callback routine from the first phase kernel(s), and it will only call the store callback routine from the last phase kernel(s).

When cufft is using only a single kernel, both the load and store callback routines will be called from the same kernel. In this case, if the transform is being done in-place (i.e. input data and output data are in the same memory location) the store callback can not safely write outside the confines of the specified element, unless it is writing the data to a completely separate output buffer.

When more than one kernel are used to implement a transform, the thread and block structure of the first kernel (the one that does the load) is often different from the thread and block structure of the last kernel (the one that does the store)

One common use of callbacks is to reduce the amount of data read or written to memory, either by selective filtering or via type conversions. When more than one kernel are used to implement a transform, cuFFT alternates using the workspace and the output buffer to write intermediate results. This means that the output buffer must always be large enough to accommodate the entire transform.

For multi-GPU transforms, the index passed to the callback routine is the element index from the start of data *on that GPU*, not from the start of the entire input or output data array.

For transforms whose dimensions can be factored into powers of 2, 3, 5, or 7, cuFFT guarantees that it will call the load and store callback routines from points in the kernel that is safe to call `__syncthreads` function from within callback routine. Caller is responsible for guaranteeing that the callback routine is at a point where the callback code has converged, to avoid deadlock. For plans whose dimensions are factored into higher primes, results of a callback routine calling `__syncthreads` are not defined.

2.10. Thread Safety

cuFFT APIs are thread safe as long as different host threads execute FFTs using different plans and the output data are disjoint.

2.11. Static Library and Callback Support

Starting with release 6.5, the cuFFT Libraries are also delivered in a static form as `libcufft_static.a` and `libcufftw_static.a` on Linux and Mac. Static libraries are not supported on Windows. The static cufft and cufftw libraries depend on thread abstraction layer library `libculibos.a`.

For example, on linux, to compile a small application using cuFFT against the dynamic library, the following command can be used:

```
nvcc myCufftApp.c -lcufft -o myCufftApp
```

For cufftw on Linux, to compile a small application against the dynamic library, the following command can be used:

```
nvcc myCufftwApp.c -lcufftw -lcufft -o myCufftwApp
```

Whereas to compile against the static cuFFT Library, extra steps need to be taken. Library needs to be device linked. It may happen during build and link of a simple program or as a separate step. Entire process is described in [Using Separate Compilation in CUDA](#).

For cuFFT and cufftw in version 9.0 or later any supported architecture can be used to do device link. Prior to version 9.0 proper link required specifying subset of supported architectures as shown in following commands.

Static cuFFT compilation command:

```
nvcc myCufftApp.c -lcufft_static -lcublas -o myCufftApp\  
-gencode arch=compute_20,\"code=sm_20\"\  
-gencode arch=compute_30,\"code=sm_30\"\  
-gencode arch=compute_35,\"code=sm_35\"\  
-gencode arch=compute_50,\"code=sm_50\"\  
-gencode arch=compute_60,\"code=sm_60\"\  
-gencode arch=compute_60,\"code=compute_60\"
```

Static cufftw compilation command:

```
nvcc myCufftwApp.c libcufftw_static.a libcufft_static.a libcublas.a  
-o myCufftwApp\  
-gencode arch=compute_20,\"code=sm_20\"\  
-gencode arch=compute_30,\"code=sm_30\"\  
-gencode arch=compute_35,\"code=sm_35\"\  
-gencode arch=compute_50,\"code=sm_50\"\  
-gencode arch=compute_60,\"code=sm_60\"\  
-gencode arch=compute_60,\"code=compute_60\"
```

Please note that cuFFT library might not contain code for certain architectures as long as there is code for a lower architecture that is binary compatible (ie. SM37, SM52, SM61). This is reflected in link commands above and significant when using versions prior r9.0. To determine if a specific SM is included in the cuFFT library, one may use **cuobjdump** utility. For example, if you wish to know if SM_50 is included, the command to run is **cuobjdump -arch sm_50 libcufft_static.a**. Some kernels are built only on select architectures (ie. kernels with half precision arithmetics are present only for SM53 and above). This can cause warnings at link time that architectures are missing from these kernels. These warnings can safely be ignored.

It is also possible to use the native Host C++ compiler and perform device link as a separate step. Please consult NVCC documentation for more details. Depending on the Host Operating system, some additional libraries like **pthread** or **d1** might be needed on the linking line.

Note that in this case, the library **cuda** is not needed. The CUDA Runtime will try to open explicitly the **cuda** library if needed. In the case of a system which does not have

the CUDA driver installed, this allows the application to gracefully manage this issue and potentially run if a CPU-only path is available.

The cuFFT static library supports user supplied callback routines. The callback routines are CUDA device code, and must be separately compiled with NVCC and linked with the cuFFT library. Please refer to the NVCC documentation regarding separate compilation for details. If you specify an SM when compiling your callback functions, you must specify one of the SM's cuFFT includes.

2.11.1. Static library without callback support

Starting with cuFFT version 9.2, a new variant of the cuFFT static library, `libcufft_static_nocallback.a`, was added. This new version does not contain callback functionality and can be linked using the host compiler only.

2.12. Accuracy and Performance

A DFT can be implemented as a matrix vector multiplication that requires $O(N^2)$ operations. However, the cuFFT Library employs the [Cooley-Tukey algorithm](#) to reduce the number of required operations to optimize the performance of particular transform sizes. This algorithm expresses the DFT matrix as a product of sparse building block matrices. The cuFFT Library implements the following building blocks: radix-2, radix-3, radix-5, and radix-7. Hence the performance of any transform size that can be factored as $2^a \times 3^b \times 5^c \times 7^d$ (where a , b , c , and d are non-negative integers) is optimized in the cuFFT library. There are also radix- m building blocks for other primes, m , whose value is < 128 . When the length cannot be decomposed as multiples of powers of primes from 2 to 127, [Bluestein's algorithm](#) is used. Since the Bluestein implementation requires more computations per output point than the Cooley-Tukey implementation, the accuracy of the Cooley-Tukey algorithm is better. The pure Cooley-Tukey implementation has excellent accuracy, with the relative error growing proportionally to $\log_2(N)$, where N is the transform size in points.

For sizes handled by the Cooley-Tukey code path, the most efficient implementation is obtained by applying the following constraints (listed in order from the most generic to the most specialized constraint, with each subsequent constraint providing the potential of an additional performance improvement).

Half precision transforms might not be suitable for all kinds of problems due to limited range represented by half precision floating point arithmetics. Please note that the first element of FFT result is the sum of all input elements and it is likely to overflow for certain inputs.

Results produced by the cuFFT library are deterministic (ie, bitwise reproducible) as long as the following are kept constant between runs: plan input parameters, cuFFT version, and GPU model.

Applies to	Recommendation	Comment
All	Use single precision transforms.	Single precision transforms require less bandwidth per computation than double precision transforms.
All	Restrict the size along all dimensions to be representable as $2^a \times 3^b \times 5^c \times 7^d$.	The cuFFT library has highly optimized kernels for transforms whose dimensions have these prime factors. In general the best performance occurs when using powers of 2, followed by powers of 3, then 5, 7.
All	Restrict the size along each dimension to use fewer distinct prime factors.	A transform of size 2^n or 3^n will usually be faster than one of size $2^i \times 3^j$ even if the latter is slightly smaller, due to the composition of specialized paths.
All	Restrict the data to be contiguous in memory when performing a single transform. When performing multiple transforms make the individual datasets contiguous	The cuFFT library has been optimized for this data layout.
All	Perform multiple (i.e., batched) transforms.	Additional optimizations are performed in batched mode.
real-to-complex transforms or complex-to-real transforms	Ensure problem size of x dimension is a multiple of 4.	This scheme uses more efficient kernels to implement conjugate symmetry property.
real-to-complex transforms or complex-to-real transforms	Use <code>out-of-place</code> mode.	This scheme uses more efficient kernels than <code>in-place</code> mode.
Multiple GPU transforms	Use PCI Express 3.0 between GPUs and ensure the GPUs are on the same switch.	The faster the interconnect between the GPUs, the faster the performance.

2.13. Caller Allocated Work Area Support

cuFFT plans may use additional memory to store intermediate results. The cuFFT library offers several functions to manage this temporary memory utilization behavior:

- ▶ `cufftSetAutoAllocation`
- ▶ `cufftXtSetWorkAreaPolicy`
- ▶ `cufftEstimate1d`, `cufftEstimate2d`, `cufftEstimate3d` and `cufftEstimateMany`
- ▶ `cufftGetSize`
- ▶ `cufftXtSetWorkAreaPolicy`

The first two functions manage allocation and ownership of temporary memory. By default cuFFT always allocates its own work area in GPU memory. Each cuFFT handle allocates data separately. If multiple cuFFT plans are to be launched sequentially it is

possible to assign the same memory chunk as work area to all those plans and reduce memory overhead.

The memory assigned as work area needs to be GPU visible. In addition to the regular memory acquired with `cudaMalloc`, usage of CUDA Unified Virtual Addressing enables cuFFT to use the following types of memory as work area memory: pinned host memory, managed memory, memory on GPU other than the one performing the calculations. While this provides flexibility, it comes with a performance penalty whose magnitude depends on the available memory bandwidth.

The `cufftEstimateNd`, `cufftEstimateMany`, and `cufftGetSize` functions provide information about the required memory size for cases where the user is allocating the work space buffer.

In version 9.2 cuFFT also introduced the `cufftXtSetWorkAreaPolicy` function. This function allows fine tuning of work area memory usage.

cuFFT 9.2 version supports only the `CUFFT_WORKAREA_MINIMAL` policy, which instructs cuFFT to re-plan the existing plan without the need to use work area memory.

Also as of cuFFT 9.2, supported FFT transforms that allow for `CUFFT_WORKAREA_MINIMAL` policy are as follows:

- ▶ Transforms of type **C2C** are supported with sizes up to 4096 in any dimension.
- ▶ Transforms of type **Z2Z** are supported with sizes up to 2048 in any dimension.
- ▶ Only single GPU transforms are supported.

Depending on the FFT transform size, a different FFT algorithm may be used when the `CUFFT_WORKAREA_MINIMAL` policy is set.

Chapter 3.

CUFFT API REFERENCE

This chapter specifies the behavior of the cuFFT library functions by describing their input/output parameters, data types, and error codes. The cuFFT library is initialized upon the first invocation of an API function, and cuFFT shuts down automatically when all user-created FFT plans are destroyed.

3.1. Return value cufftResult

All cuFFT Library return values except for **CUFFT_SUCCESS** indicate that the current API call failed and the user should reconfigure to correct the problem. The possible return values are defined as follows:

```
typedef enum cufftResult_t {
    CUFFT_SUCCESS = 0, // The cuFFT operation was successful
    CUFFT_INVALID_PLAN = 1, // cuFFT was passed an invalid plan handle
    CUFFT_ALLOC_FAILED = 2, // cuFFT failed to allocate GPU or CPU memory
    CUFFT_INVALID_TYPE = 3, // No longer used
    CUFFT_INVALID_VALUE = 4, // User specified an invalid pointer or
parameter
    CUFFT_INTERNAL_ERROR = 5, // Driver or internal cuFFT library error
    CUFFT_EXEC_FAILED = 6, // Failed to execute an FFT on the GPU
    CUFFT_SETUP_FAILED = 7, // The cuFFT library failed to initialize
    CUFFT_INVALID_SIZE = 8, // User specified an invalid transform size
    CUFFT_UNALIGNED_DATA = 9, // No longer used
    CUFFT_INCOMPLETE_PARAMETER_LIST = 10, // Missing parameters in call
    CUFFT_INVALID_DEVICE = 11, // Execution of a plan was on different GPU than
plan creation
    CUFFT_PARSE_ERROR = 12, // Internal plan database error
    CUFFT_NO_WORKSPACE = 13 // No workspace has been provided prior to plan
execution
    CUFFT_NOT_IMPLEMENTED = 14, // Function does not implement functionality for
parameters given.
    CUFFT_LICENSE_ERROR = 15, // Used in previous versions.
    CUFFT_NOT_SUPPORTED = 16 // Operation is not supported for parameters
given.
} cufftResult;
```

Users are encouraged to check return values from cuFFT functions for errors as shown in [cuFFT Code Examples](#).

3.2. cuFFT Basic Plans

3.2.1. Function `cufftPlan1d()`

```
cufftResult
cufftPlan1d(cufftHandle *plan, int nx, cufftType type, int batch);
```

Creates a 1D FFT plan configuration for a specified signal size and data type. The `batch` input parameter tells cuFFT how many 1D transforms to configure.

Input

<code>plan</code>	Pointer to a <code>cufftHandle</code> object
<code>nx</code>	The transform size (e.g. 256 for a 256-point FFT)
<code>type</code>	The transform data type (e.g., <code>CUFFT_C2C</code> for single precision complex to complex)
<code>batch</code>	Number of transforms of size <code>nx</code> . Please consider using <code>cufftPlanMany</code> for multiple transforms.

Output

<code>plan</code>	Contains a cuFFT 1D plan handle value
-------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	The <code>nx</code> or <code>batch</code> parameter is not a supported size.

3.2.2. Function `cufftPlan2d()`

```
cufftResult
cufftPlan2d(cufftHandle *plan, int nx, int ny, cufftType type);
```

Creates a 2D FFT plan configuration according to specified signal sizes and data type.

Input

<code>plan</code>	Pointer to a <code>cufftHandle</code> object
<code>nx</code>	The transform size in the x dimension. This is slowest changing dimension of a transform (strided in memory).
<code>ny</code>	The transform size in the y dimension. This is fastest changing dimension of a transform (contiguous in memory).

<code>type</code>	The transform data type (e.g., <code>CUFFT_C2R</code> for single precision complex to real)
-------------------	---

Output

<code>plan</code>	Contains a cuFFT 2D plan handle value
-------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	Either or both of the <code>nx</code> or <code>ny</code> parameters is not a supported size.

3.2.3. Function `cufftPlan3d()`

```
cufftResult
cufftPlan3d(cufftHandle *plan, int nx, int ny, int nz, cufftType type);
```

Creates a 3D FFT plan configuration according to specified signal sizes and data type. This function is the same as `cufftPlan2d()` except that it takes a third size parameter `nz`.

Input

<code>plan</code>	Pointer to a <code>cufftHandle</code> object
<code>nx</code>	The transform size in the x dimension. This is slowest changing dimension of a transform (strided in memory).
<code>ny</code>	The transform size in the y dimension
<code>nz</code>	The transform size in the z dimension. This is fastest changing dimension of a transform (contiguous in memory).
<code>type</code>	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)

Output

<code>plan</code>	Contains a cuFFT 3D plan handle value
-------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

<code>CUFFT_INVALID_SIZE</code>	One or more of the <code>nx</code> , <code>ny</code> , or <code>nz</code> parameters is not a supported size.
---------------------------------	---

3.2.4. Function `cufftPlanMany()`

```
cufftResult
cufftPlanMany(cufftHandle *plan, int rank, int *n, int *inembed,
              int istride, int idist, int *onembed, int ostride,
              int odist, cufftType type, int batch);
```

Creates a FFT plan configuration of dimension **rank**, with sizes specified in the array **n**. The **batch** input parameter tells cuFFT how many transforms to configure. With this function, batched plans of 1, 2, or 3 dimensions may be created.

The `cufftPlanMany()` API supports more complicated input and output data layouts via the advanced data layout parameters: **inembed**, **istride**, **idist**, **onembed**, **ostride**, and **odist**.

If **inembed** and **onembed** are set to **NULL**, all other stride information is ignored, and default strides are used. The default assumes contiguous data arrays.

All arrays are assumed to be in CPU memory.

Please note that behavior of `cufftPlanMany` function when **inembed** and **onembed** is **NULL** is different than corresponding function in FFTW library `fftw_plan_many_dft`.

Input

plan	Pointer to a <code>cufftHandle</code> object
rank	Dimensionality of the transform (1, 2, or 3).
n	Array of size rank , describing the size of each dimension, n[0] being the size of the outermost and n[rank-1] innermost (contiguous) dimension of a transform.
inembed	Pointer of size rank that indicates the storage dimensions of the input data in memory. If set to NULL all other advanced data layout parameters are ignored.
istride	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
idist	Indicates the distance between the first element of two consecutive signals in a batch of the input data
onembed	Pointer of size rank that indicates the storage dimensions of the output data in memory. If set to NULL all other advanced data layout parameters are ignored.
ostride	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
odist	Indicates the distance between the first element of two consecutive signals in a batch of the output data
type	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)
batch	Batch size for this transform

Output

<code>plan</code>	Contains a cuFFT plan handle
-------------------	------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the parameters is not a supported size.

3.3. cuFFT Extensible Plans

This API separates handle creation from plan generation. This makes it possible to change plan settings, which may alter the outcome of the plan generation phase, before the plan is actually generated.

3.3.1. Function `cufftCreate()`

```
cufftResult
cufftCreate(cufftHandle *plan);
```

Creates only an opaque handle, and allocates small data structures on the host. The `cufftMakePlan*()` calls actually do the plan generation.

Input

<code>plan</code>	Pointer to a <code>cufftHandle</code> object
-------------------	--

Output

<code>plan</code>	Contains a cuFFT plan handle value
-------------------	------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

3.3.2. Function `cufftMakePlan1d()`

```
cufftResult
cufftMakePlan1d(cufftHandle plan, int nx, cufftType type, int batch,
size_t *workSize);
```

Following a call to `cufftCreate()` makes a 1D FFT plan configuration for a specified signal size and data type. The `batch` input parameter tells cuFFT how many 1D transforms to configure.

If `cufftXtSetGPUs()` was called prior to this call with multiple GPUs, then `workSize` will contain multiple sizes. See sections on multiple GPUs for more details.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>nx</code>	The transform size (e.g. 256 for a 256-point FFT). For multiple GPUs, this must be a power of 2.
<code>type</code>	The transform data type (e.g., <code>CUFFT_C2C</code> for single precision complex to complex). For multiple GPUs this must be a complex to complex transform.
<code>batch</code>	Number of transforms of size <code>nx</code> . Please consider using <code>cufftMakePlanMany</code> for multiple transforms.
<code>*workSize</code>	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

<code>*workSize</code>	Pointer to the size(s) of the work areas.
------------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle. Handle is not valid when multi-GPU restrictions are not met.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	The <code>nx</code> or <code>batch</code> parameter is not a supported size.

3.3.3. Function `cufftMakePlan2d()`

```
cufftResult
cufftMakePlan2d(cufftHandle plan, int nx, int ny, cufftType type,
                size_t *workSize);
```

Following a call to `cufftCreate()` makes a 2D FFT plan configuration according to specified signal sizes and data type.

If `cufftXtSetGPUs()` was called prior to this call with multiple GPUs, then `workSize` will contain multiple sizes. See sections on multiple GPUs for more details.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
-------------------	---

nx	The transform size in the x dimension. This is slowest changing dimension of a transform (strided in memory). For multiple GPUs, this must be factorable into primes less than or equal to 127.
ny	The transform size in the y dimension. This is fastest changing dimension of a transform (contiguous in memory). For 2 GPUs, this must be factorable into primes less than or equal to 127.
type	The transform data type (e.g., <code>CUFFT_C2R</code> for single precision complex to real). For multiple GPUs this must be a complex to complex transform.
*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

*workSize	Pointer to the size(s) of the work areas.
------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	Either or both of the <code>nx</code> or <code>ny</code> parameters is not a supported size.

3.3.4. Function `cufftMakePlan3d()`

```

cufftResult
cufftMakePlan3d(cufftHandle plan, int nx, int ny, int nz, cufftType type,
size_t *workSize);

```

Following a call to `cufftCreate()` makes a 3D FFT plan configuration according to specified signal sizes and data type. This function is the same as `cufftPlan2d()` except that it takes a third size parameter `nz`.

If `cufftXtSetGPUs()` was called prior to this call with multiple GPUs, then `workSize` will contain multiple sizes. See sections on multiple GPUs for more details.

Input

plan	<code>cufftHandle</code> returned by <code>cufftCreate</code>
nx	The transform size in the x dimension. This is slowest changing dimension of a transform (strided in memory). For multiple GPUs, this must be factorable into primes less than or equal to 127.

ny	The transform size in the y dimension. For multiple GPUs, this must be factorable into primes less than or equal to 127.
nz	The transform size in the z dimension. This is fastest changing dimension of a transform (contiguous in memory). For multiple GPUs, this must be factorable into primes less than or equal to 127.
type	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex). For multiple GPUs this must be a complex to complex transform.
*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

*workSize	Pointer to the size(s) of the work area(s).
------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the <code>nx</code> , <code>ny</code> , or <code>nz</code> parameters is not a supported size.

3.3.5. Function `cufftMakePlanMany()`

```

cufftResult
cufftMakePlanMany(cufftHandle plan, int rank, int *n, int *inembed,
                  int istride, int idist, int *onembed, int ostride,
                  int odist, cufftType type, int batch, size_t *workSize);

```

Following a call to `cufftCreate()` makes a FFT plan configuration of dimension `rank`, with sizes specified in the array `n`. The `batch` input parameter tells cuFFT how many transforms to configure. With this function, batched plans of 1, 2, or 3 dimensions may be created.

The `cufftPlanMany()` API supports more complicated input and output data layouts via the advanced data layout parameters: `inembed`, `istride`, `idist`, `onembed`, `ostride`, and `odist`.

If `inembed` and `onembed` are set to `NULL`, all other stride information is ignored, and default strides are used. The default assumes contiguous data arrays.

If `cufftXtSetGPUs()` was called prior to this call with multiple GPUs, then `workSize` will contain multiple sizes. See sections on multiple GPUs for more details.

All arrays are assumed to be in CPU memory.

Input

plan	cuFFTHandle returned by <code>cuFFTCreate</code>
rank	Dimensionality of the transform (1, 2, or 3)
n	Array of size <code>rank</code> , describing the size of each dimension, <code>n[0]</code> being the size of the outermost and <code>n[rank-1]</code> innermost (contiguous) dimension of a transform. For multiple GPUs and rank equal to 1, the sizes must be a power of 2. For multiple GPUs and rank equal to 2 or 3, the sizes must be factorable into primes less than or equal to 127.
inembed	Pointer of size <code>rank</code> that indicates the storage dimensions of the input data in memory, <code>inembed[0]</code> being the storage dimension of the innermost dimension. If set to NULL all other advanced data layout parameters are ignored.
istride	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
idist	Indicates the distance between the first element of two consecutive signals in a batch of the input data
onembed	Pointer of size <code>rank</code> that indicates the storage dimensions of the output data in memory, <code>onembed[0]</code> being the storage dimension of the innermost dimension. If set to NULL all other advanced data layout parameters are ignored.
ostride	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
odist	Indicates the distance between the first element of two consecutive signals in a batch of the output data
type	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex). For 2 GPUs this must be a complex to complex transform.
batch	Batch size for this transform
*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

*workSize	Pointer to the size(s) of the work areas.
------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle. Handle is not valid when multi-GPU restrictions are not met.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

<code>CUFFT_INVALID_SIZE</code>	One or more of the parameters is not a supported size.
---------------------------------	--

3.3.6. Function `cufftMakePlanMany64()`

```
cufftResult
cufftMakePlanMany64(cufftHandle plan, int rank,
    long long int *n,
    long long int *inembed, long long int istride, long long int idist,
    long long int *onembed, long long int ostride, long long int odist,
    cufftType type,
    long long int batch, size_t *workSize);
```

Following a call to `cufftCreate()` makes a FFT plan configuration of dimension **rank**, with sizes specified in the array **n**. The **batch** input parameter tells cuFFT how many transforms to configure. With this function, batched plans of 1, 2, or 3 dimensions may be created.

This API is identical to `cufftMakePlanMany` except that the arguments specifying sizes and strides are 64 bit integers. This API makes very large transforms possible. cuFFT includes kernels that use 32 bit indexes, and kernels that use 64 bit indexes. cuFFT planning selects 32 bit kernels whenever possible to avoid any overhead due to 64 bit arithmetic.

All sizes and types of transform are supported by this interface, with two exceptions. For transforms whose total size exceeds 4G elements, the dimensions specified in the array **n** must be factorable into primes that are less than or equal to 127. For real to complex and complex to real transforms whose total size exceeds 2G elements, the fastest changing dimension must be even.

The `cufftPlanMany64()` API supports more complicated input and output data layouts via the advanced data layout parameters: **inembed**, **istride**, **idist**, **onembed**, **ostride**, and **odist**.

If **inembed** and **onembed** are set to **NULL**, all other stride information is ignored, and default strides are used. The default assumes contiguous data arrays.

If `cufftXtSetGPUs()` was called prior to this call with multiple GPUs, then **workSize** will contain multiple sizes. See sections on multiple GPUs for more details.

All arrays are assumed to be in CPU memory.

Input

plan	<code>cufftHandle</code> returned by <code>cufftCreate</code>
rank	Dimensionality of the transform (1, 2, or 3)
n	Array of size rank , describing the size of each dimension. For multiple GPUs and rank equal to 1, the sizes must be a power of 2. For multiple GPUs and rank equal to 2 or 3, the sizes must be factorable into primes less than or equal to 127.
inembed	Pointer of size rank that indicates the storage dimensions of the input data in memory. If set to NULL all other advanced data layout parameters are ignored.
istride	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension

idist	Indicates the distance between the first element of two consecutive signals in a batch of the input data
onembed	Pointer of size rank that indicates the storage dimensions of the output data in memory. If set to NULL all other advanced data layout parameters are ignored.
ostride	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
odist	Indicates the distance between the first element of two consecutive signals in a batch of the output data
type	The transform data type (e.g., CUFFT_R2C for single precision real to complex). For 2 GPUs this must be a complex to complex transform.
batch	Batch size for this transform
*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

*workSize	Pointer to the size(s) of the work areas.
------------------	---

Return Values

CUFFT_SUCCESS	cuFFT successfully created the FFT plan.
CUFFT_INVALID_PLAN	The plan parameter is not a valid handle. Handle is not valid when multi-GPU restrictions are not met.
CUFFT_ALLOC_FAILED	The allocation of GPU resources for the plan failed.
CUFFT_INVALID_VALUE	One or more invalid parameters were passed to the API.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_SIZE	One or more of the parameters is not a supported size.

3.3.7. Function `cufftXtMakePlanMany()`

```

cufftResult
cufftXtMakePlanMany(cufftHandle plan, int rank, long long int *n, long long
int *inembed,
long long int istride, long long int idist, cudaDataType inputtype,
long long int *onembed, long long int ostride, long long int odist,
cudaDataType outputtype, long long int batch, size_t *workSize,
cudaDataType executiontype);

```

Following a call to **cufftCreate()** makes an FFT plan configuration of dimension **rank**, with sizes specified in the array **n**. The **batch** input parameter tells cuFFT how many transforms to configure. With this function, batched plans of 1, 2, or 3 dimensions may be created.

Type specifiers **inputtype**, **outputtype** and **executiontype** dictate type and precision of transform to be performed. Not all combinations of parameters are

supported. Currently all three parameters need to match precision. Parameters **inputtype** and **outputtype** need to match transform type complex-to-complex, real-to-complex or complex-to-real. Parameter **executiontype** needs to match precision and be of a complex type. Example: for a 16 bit real-to-complex transform parameters **inputtype**, **outputtype** and **executiontype** would have values of **CUDA_R_16F**, **CUDA_C_16F** and **CUDA_C_16F** respectively.

The **cufftXtMakePlanMany()** API supports more complicated input and output data layouts via the advanced data layout parameters: **inembed**, **istride**, **idist**, **onembed**, **ostride**, and **odist**.

If **inembed** and **onembed** are set to **NULL**, all other stride information is ignored, and default strides are used. The default assumes contiguous data arrays.

If **cufftXtSetGPUs()** was called prior to this call with multiple GPUs, then **workSize** will contain multiple sizes. See sections on multiple GPUs for more details.

All arrays are assumed to be in CPU memory.

Input

plan	cufftHandle returned by cufftCreate
rank	Dimensionality of the transform (1, 2, or 3)
n	Array of size rank , describing the size of each dimension, n[0] being the size of the innermost deminsion. For multiple GPUs and rank equal to 1, the sizes must be a power of 2. For multiple GPUs and rank equal to 2 or 3, the sizes must be factorable into primes less than or equal to 127.
inembed	Pointer of size rank that indicates the storage dimensions of the input data in memory, inembed[0] being the storage dimension of the innermost dimension. If set to NULL all other advanced data layout parameters are ignored.
istride	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
idist	Indicates the distance between the first element of two consecutive signals in a batch of the input data
inputtype	Type of input data.
onembed	Pointer of size rank that indicates the storage dimensions of the output data in memory, inembed[0] being the storage dimension of the innermost dimension. If set to NULL all other advanced data layout parameters are ignored.
ostride	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
odist	Indicates the distance between the first element of two consecutive signals in a batch of the output data
outputtype	Type of output data.
batch	Batch size for this transform
*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

<code>executiontype</code>	Type of data to be used for computations.
----------------------------	---

Output

<code>*workSize</code>	Pointer to the size(s) of the work areas.
------------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully created the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle. Handle is not valid when multi-GPU restrictions are not met.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the parameters is not a supported size.

3.4. cuFFT Estimated Size of Work Area

During plan execution, cuFFT requires a work area for temporary storage of intermediate results. The `cufftEstimate*()` calls return an estimate for the size of the work area required, given the specified parameters, and assuming default plan settings. Some problem sizes require much more storage than others. In particular powers of 2 are very efficient in terms of temporary storage. Large prime numbers, however, use different algorithms and may need up to the eight times that of a similarly sized power of 2. These routines return estimated `workSize` values which may still be smaller than the actual values needed especially for values of `n` that are not multiples of powers of 2, 3, 5 and 7. More refined values are given by the `cufftGetSize*()` routines, but these values may still be conservative.

3.4.1. Function `cufftEstimate1d()`

```
cufftResult
cufftEstimate1d(int nx, cufftType type, int batch, size_t *workSize);
```

During plan execution, cuFFT requires a work area for temporary storage of intermediate results. This call returns an estimate for the size of the work area required, given the specified parameters, and assuming default plan settings.

Input

<code>nx</code>	The transform size (e.g. 256 for a 256-point FFT)
<code>type</code>	The transform data type (e.g., <code>CUFFT_C2C</code> for single precision complex to complex)
<code>batch</code>	Number of transforms of size <code>nx</code> . Please consider using <code>cufftEstimateMany</code> for multiple transforms.
<code>*workSize</code>	Pointer to the size, in bytes, of the work space.

Output

*workSize	Pointer to the size of the work space
------------------	---------------------------------------

Return Values

CUFFT_SUCCESS	cuFFT successfully returned the size of the work space.
CUFFT_ALLOC_FAILED	The allocation of GPU resources for the plan failed.
CUFFT_INVALID_VALUE	One or more invalid parameters were passed to the API.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_SIZE	The nx parameter is not a supported size.

3.4.2. Function `cufftEstimate2d()`

```

cufftResult
cufftEstimate2d(int nx, int ny, cufftType type, size_t *workSize);

```

During plan execution, cuFFT requires a work area for temporary storage of intermediate results. This call returns an estimate for the size of the work area required, given the specified parameters, and assuming default plan settings.

Input

nx	The transform size in the x dimension (number of rows)
ny	The transform size in the y dimension (number of columns)
type	The transform data type (e.g., CUFFT_C2R for single precision complex to real)
*workSize	Pointer to the size, in bytes, of the work space.

Output

*workSize	Pointer to the size, of the work space
------------------	--

Return Values

CUFFT_SUCCESS	cuFFT successfully returned the size of the work space.
CUFFT_ALLOC_FAILED	The allocation of GPU resources for the plan failed.
CUFFT_INVALID_VALUE	One or more invalid parameters were passed to the API.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_SIZE	Either or both of the nx or ny parameters is not a supported size.

3.4.3. Function `cufftEstimate3d()`

```

cufftResult
cufftEstimate3d(int nx, int ny, int nz, cufftType type, size_t *workSize);

```

During plan execution, cuFFT requires a work area for temporary storage of intermediate results. This call returns an estimate for the size of the work area required, given the specified parameters, and assuming default plan settings.

Input

nx	The transform size in the x dimension
ny	The transform size in the y dimension
nz	The transform size in the z dimension
type	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)
*workSize	Pointer to the size, in bytes, of the work space.

Output

*workSize	Pointer to the size of the work space
------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully returned the size of the work space.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the <code>nx</code> , <code>ny</code> , or <code>nz</code> parameters is not a supported size.

3.4.4. Function `cufftEstimateMany()`

```

cufftResult
cufftEstimateMany(int rank, int *n, int *inembed,
                  int istride, int idist, int *onembed, int ostride,
                  int odist, cufftType type, int batch, size_t *workSize);

```

During plan execution, cuFFT requires a work area for temporary storage of intermediate results. This call returns an estimate for the size of the work area required, given the specified parameters, and assuming default plan settings.

The `cufftEstimateMany()` API supports more complicated input and output data layouts via the advanced data layout parameters: `inembed`, `istride`, `idist`, `onembed`, `ostride`, and `odist`.

All arrays are assumed to be in CPU memory.

Input

rank	Dimensionality of the transform (1, 2, or 3)
n	Array of size <code>rank</code> , describing the size of each dimension
inembed	Pointer of size <code>rank</code> that indicates the storage dimensions of the input data in memory. If set to NULL all other advanced data layout parameters are ignored.

<code>istride</code>	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
<code>idist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the input data
<code>onembed</code>	Pointer of size <code>rank</code> that indicates the storage dimensions of the output data in memory. If set to NULL all other advanced data layout parameters are ignored.
<code>ostride</code>	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
<code>odist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the output data
<code>type</code>	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)
<code>batch</code>	Batch size for this transform
<code>*workSize</code>	Pointer to the size, in bytes, of the work space.

Output

<code>*workSize</code>	Pointer to the size of the work space
------------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully returned the size of the work space.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the parameters is not a supported size.

3.5. cuFFT Refined Estimated Size of Work Area

The `cufftGetSize*()` routines give a more accurate estimate of the work area size required for a plan than the `cufftEstimate*()` routines as they take into account any plan settings that may have been made. As discussed in the section [cuFFT Estimated Size of Work Area](#), the `workSize` value(s) returned may be conservative especially for values of `n` that are not multiples of powers of 2, 3, 5 and 7.

3.5.1. Function `cufftGetSize1d()`

```
cufftResult
cufftGetSize1d(cufftHandle plan, int nx, cufftType type, int batch,
              size_t *workSize);
```

This call gives a more accurate estimate of the work area size required for a plan than `cuFFTEstimate1d()`, given the specified parameters, and taking into account any plan settings that may have been made.

Input

<code>plan</code>	<code>cuFFTHandle</code> returned by <code>cuFFTCreate</code>
<code>nx</code>	The transform size (e.g. 256 for a 256-point FFT)
<code>type</code>	The transform data type (e.g., <code>CUFFT_C2C</code> for single precision complex to complex)
<code>batch</code>	Number of transforms of size <code>nx</code> . Please consider using <code>cuFFTGetSizeMany</code> for multiple transforms.
<code>*workSize</code>	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

<code>*workSize</code>	Pointer to the size of the work space
------------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully returned the size of the work space.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	The <code>nx</code> parameter is not a supported size.

3.5.2. Function `cuFFTGetSize2d()`

```
cuFFTResult
cuFFTGetSize2d(cuFFTHandle plan, int nx, int ny, cuFFTType type,
               size_t *workSize);
```

This call gives a more accurate estimate of the work area size required for a plan than `cuFFTEstimate2d()`, given the specified parameters, and taking into account any plan settings that may have been made.

Input

<code>plan</code>	<code>cuFFTHandle</code> returned by <code>cuFFTCreate</code>
<code>nx</code>	The transform size in the x dimension (number of rows)
<code>ny</code>	The transform size in the y dimension (number of columns)
<code>type</code>	The transform data type (e.g., <code>CUFFT_C2R</code> for single precision complex to real)

*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.
------------------	---

Output

*workSize	Pointer to the size of the work space
------------------	---------------------------------------

Return Values

CUFFT_SUCCESS	cuFFT successfully returned the size of the work space.
CUFFT_INVALID_PLAN	The <code>plan</code> parameter is not a valid handle.
CUFFT_ALLOC_FAILED	The allocation of GPU resources for the plan failed.
CUFFT_INVALID_VALUE	One or more invalid parameters were passed to the API.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_SIZE	Either or both of the <code>nx</code> or <code>ny</code> parameters is not a supported size.

3.5.3. Function `cufftGetSize3d()`

```

cufftResult
cufftGetSize3d(cufftHandle plan, int nx, int ny, int nz, cufftType type,
size_t *workSize);

```

This call gives a more accurate estimate of the work area size required for a plan than `cufftEstimate3d()`, given the specified parameters, and taking into account any plan settings that may have been made.

Input

plan	<code>cufftHandle</code> returned by <code>cufftCreate</code>
nx	The transform size in the x dimension
ny	The transform size in the y dimension
nz	The transform size in the z dimension
type	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)
*workSize	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

*workSize	Pointer to the size of the work space.
------------------	--

Return Values

CUFFT_SUCCESS	cuFFT successfully returned the size of the work space.
CUFFT_INVALID_PLAN	The <code>plan</code> parameter is not a valid handle.

<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the <code>nx</code> , <code>ny</code> , or <code>nz</code> parameters is not a supported size.

3.5.4. Function `cufftGetSizeMany()`

```
cufftResult
cufftGetSizeMany(cufftHandle plan, int rank, int *n, int *inembed,
                 int istride, int idist, int *onembed, int ostride,
                 int odist, cufftType type, int batch, size_t *workSize);
```

This call gives a more accurate estimate of the work area size required for a plan than `cufftEstimateSizeMany()`, given the specified parameters, and taking into account any plan settings that may have been made.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>rank</code>	Dimensionality of the transform (1, 2, or 3)
<code>n</code>	Array of size <code>rank</code> , describing the size of each dimension
<code>inembed</code>	Pointer of size <code>rank</code> that indicates the storage dimensions of the input data in memory. If set to NULL all other advanced data layout parameters are ignored.
<code>istride</code>	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
<code>idist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the input data
<code>onembed</code>	Pointer of size <code>rank</code> that indicates the storage dimensions of the output data in memory. If set to NULL all other advanced data layout parameters are ignored.
<code>ostride</code>	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
<code>odist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the output data
<code>type</code>	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)
<code>batch</code>	Batch size for this transform
<code>*workSize</code>	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

<code>*workSize</code>	Pointer to the size of the work area
------------------------	--------------------------------------

Return Values

CUFFT_SUCCESS	cuFFT successfully returned the size of the work space.
CUFFT_INVALID_PLAN	The <code>plan</code> parameter is not a valid handle.
CUFFT_ALLOC_FAILED	The allocation of GPU resources for the plan failed.
CUFFT_INVALID_VALUE	One or more invalid parameters were passed to the API.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_SIZE	One or more of the parameters is not a supported size.

3.5.5. Function `cufftGetSizeMany64()`

```

cufftResult
cufftGetSizeMany64(cufftHandle plan, int rank,
    long long int *n,
    long long int *inembed, long long int istride, long long int idist,
    long long int *onembed, long long int ostride, long long int odist,
    cufftType type,
    long long int batch, size_t *workSize);

```

This call gives a more accurate estimate of the work area size required for a plan than `cufftEstimateSizeMany()`, given the specified parameters, and taking into account any plan settings that may have been made.

This API is identical to `cufftMakePlanMany` except that the arguments specifying sizes and strides are 64 bit integers. This API makes very large transforms possible. cuFFT includes kernels that use 32 bit indexes, and kernels that use 64 bit indexes. cuFFT planning selects 32 bit kernels whenever possible to avoid any overhead due to 64 bit arithmetic.

All sizes and types of transform are supported by this interface, with two exceptions. For transforms whose total size exceeds 4G elements, the dimensions specified in the array `n` must be factorable into primes that are less than or equal to 127. For real to complex and complex to real transforms whose total size exceeds 2G elements, the fastest changing dimension must be even.

Input

plan	<code>cufftHandle</code> returned by <code>cufftCreate</code>
rank	Dimensionality of the transform (1, 2, or 3)
n	Array of size <code>rank</code> , describing the size of each dimension
inembed	Pointer of size <code>rank</code> that indicates the storage dimensions of the input data in memory. If set to NULL all other advanced data layout parameters are ignored.
istride	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
idist	Indicates the distance between the first element of two consecutive signals in a batch of the input data

<code>onembed</code>	Pointer of size <code>rank</code> that indicates the storage dimensions of the output data in memory. If set to NULL all other advanced data layout parameters are ignored.
<code>ostride</code>	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
<code>odist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the output data
<code>type</code>	The transform data type (e.g., <code>CUFFT_R2C</code> for single precision real to complex)
<code>batch</code>	Batch size for this transform
<code>*workSize</code>	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

<code>*workSize</code>	Pointer to the size of the work area
------------------------	--------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully returned the size of the work space.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the parameters is not a supported size.

3.5.6. Function `cufftXtGetSizeMany()`

```

cufftResult
cufftXtGetSizeMany(cufftHandle plan, int rank, long long int *n, long long
int *inembed,
    long long int istride, long long int idist, cudaDataType inputtype,
    long long int *onembed, long long int ostride, long long int odist,
    cudaDataType outputtype, long long int batch, size_t *workSize,
    cudaDataType executiontype);

```

This call gives a more accurate estimate of the work area size required for a plan than `cufftEstimateSizeMany()`, given the specified parameters that match signature of `cufftXtMakePlanMany` function, and taking into account any plan settings that may have been made.

For more information about valid combinations of `inputtype`, `outputtype` and `executiontype` parameters please refer to documentation of `cufftXtMakePlanMany` function.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>rank</code>	Dimensionality of the transform (1, 2, or 3)
<code>n</code>	Array of size <code>rank</code> , describing the size of each dimension
<code>inembed</code>	Pointer of size <code>rank</code> that indicates the storage dimensions of the input data in memory. If set to NULL all other advanced data layout parameters are ignored.
<code>istride</code>	Indicates the distance between two successive input elements in the least significant (i.e., innermost) dimension
<code>idist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the input data
<code>cudaDataType inputtype</code>	Type of input data.
<code>onembed</code>	Pointer of size <code>rank</code> that indicates the storage dimensions of the output data in memory. If set to NULL all other advanced data layout parameters are ignored.
<code>ostride</code>	Indicates the distance between two successive output elements in the output array in the least significant (i.e., innermost) dimension
<code>odist</code>	Indicates the distance between the first element of two consecutive signals in a batch of the output data
<code>cudaDataType outputtype</code>	Type of output data.
<code>batch</code>	Batch size for this transform
<code>*workSize</code>	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.
<code>cudaDataType executiontype</code>	Type of data to be used for computations.

Output

<code>*workSize</code>	Pointer to the size of the work area
------------------------	--------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully returned the size of the work space.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INVALID_VALUE</code>	One or more invalid parameters were passed to the API.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_SIZE</code>	One or more of the parameters is not a supported size.

3.6. Function `cufftGetSize()`

```
cufftResult
cufftGetSize(cufftHandle plan, size_t *workSize);
```

Once plan generation has been done, either with the original API or the extensible API, this call returns the actual size of the work area required to support the plan. Callers who choose to manage work area allocation within their application must use this call after plan generation, and after any `cufftSet*` () calls subsequent to plan generation, if those calls might alter the required work space size.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>*workSize</code>	Pointer to the size(s), in bytes, of the work areas. For example for two GPUs worksize must be declared to have two elements.

Output

<code>*workSize</code>	Pointer to the size of the work space
------------------------	---------------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully returned the size of the work space.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.

3.7. cuFFT Caller Allocated Work Area Support

3.7.1. Function `cufftSetAutoAllocation()`

```
cufftResult
cufftSetAutoAllocation(cufftHandle plan, int autoAllocate);
```

`cufftSetAutoAllocation()` indicates that the caller intends to allocate and manage work areas for plans that have been generated. cuFFT default behavior is to allocate the work area at plan generation time. If `cufftSetAutoAllocation()` has been called with `autoAllocate` set to 0 ("false") prior to one of the `cufftMakePlan*` () calls, cuFFT does not allocate the work area. This is the preferred sequence for callers wishing to manage work area allocation.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code> .
<code>autoAllocate</code>	Indicates whether to allocate work area.

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully allows user to manage work area.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.

3.7.2. Function `cufftSetWorkArea()`

```
cufftResult
    cufftSetWorkArea(cufftHandle plan, void *workArea);
```

cufftSetWorkArea() overrides the work area pointer associated with a plan. If the work area was auto-allocated, cuFFT frees the auto-allocated space. The **cufftExecute***() calls assume that the work area pointer is valid and that it points to a contiguous region in device memory that does not overlap with any other work area. If this is not the case, results are indeterminate.

Input

plan	cufftHandle returned by cufftCreate
workArea	Pointer to workArea. For multiple GPUs, multiple work area pointers must be given.

Return Values

CUFFT_SUCCESS	cuFFT successfully allows user to override workArea pointer.
CUFFT_INVALID_PLAN	The plan parameter is not a valid handle.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.

3.7.3. Function `cufftXtSetWorkAreaPolicy()`

```
cufftResult
    cufftXtSetWorkAreaPolicy(cufftHandle plan, cufftXtWorkAreaPolicy policy,
    size_t *workSize);
```

cufftXtSetWorkAreaPolicy() indicates that the caller intends to change work area size for a given plan handle. cuFFT's default behavior is to allocate the work area at plan generation time with a default size that depends on the plan type and other parameters. If **cufftXtSetWorkAreaPolicy()** has been called with the **policy** parameter set to **CUFFT_WORKAREA_MINIMAL**, cuFFT will attempt to re-plan the handle to use zero bytes of work area memory. If the **cufftXtSetWorkAreaPolicy()** call is successful the auto-allocated work area memory is released.

Currently the policies **CUFFT_WORKAREA_PERFORMANCE**, **CUFFT_WORKAREA_USER** and the **workSize** parameter are not supported and reserved for use in future cuFFT releases.

This function can be called once per lifetime of a plan handle.

Input

plan	cufftHandle returned by cufftCreate and followed by one of the planning calls
policy	Type of work area policy to apply.
workSize	Reserved for future use.

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully applied new policy.
<code>CUFFT_INVALID_SIZE</code>	FFT size does not allow use of the selected policy.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle, or <code>cufftXtSetWorkAreaPolicy()</code> has been called already.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

3.8. Function `cufftDestroy()`

```
cufftResult
cufftDestroy(cufftHandle plan);
```

Frees all GPU resources associated with a cuFFT plan and destroys the internal plan data structure. This function should be called once a plan is no longer needed, to avoid wasting GPU memory.

Input

<code>plan</code>	The <code>cufftHandle</code> object of the plan to be destroyed.
-------------------	--

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully destroyed the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.

3.9. cuFFT Execution

3.9.1. Functions `cufftExecC2C()` and `cufftExecZ2Z()`

```
cufftResult
cufftExecC2C(cufftHandle plan, cufftComplex *idata,
             cufftComplex *odata, int direction);
cufftResult
cufftExecZ2Z(cufftHandle plan, cufftDoubleComplex *idata,
             cufftDoubleComplex *odata, int direction);
```

`cufftExecC2C()` (**`cufftExecZ2Z()`**) executes a single-precision (double-precision) complex-to-complex transform plan in the transform direction as specified by **`direction`** parameter. cuFFT uses the GPU memory pointed to by the **`idata`** parameter as input data. This function stores the Fourier coefficients in the **`odata`** array. If **`idata`** and **`odata`** are the same, this method does an in-place transform.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>idata</code>	Pointer to the complex input data (in GPU memory) to transform

<code>odata</code>	Pointer to the complex output data (in GPU memory)
<code>direction</code>	The transform direction: <code>CUFFT_FORWARD</code> or <code>CUFFT_INVERSE</code>

Output

<code>odata</code>	Contains the complex Fourier coefficients
--------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully executed the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INVALID_VALUE</code>	At least one of the parameters <code>idata</code> , <code>odata</code> , and <code>direction</code> is not valid.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_EXEC_FAILED</code>	cuFFT failed to execute the transform on the GPU.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

3.9.2. Functions `cufftExecR2C()` and `cufftExecD2Z()`

```

cufftResult
    cufftExecR2C(cufftHandle plan, cufftReal *idata, cufftComplex *odata);
cufftResult
    cufftExecD2Z(cufftHandle plan, cufftDoubleReal *idata, cufftDoubleComplex
        *odata);

```

`cufftExecR2C()` (**`cufftExecD2Z()`**) executes a single-precision (double-precision) real-to-complex, implicitly forward, cuFFT transform plan. cuFFT uses as input data the GPU memory pointed to by the `idata` parameter. This function stores the nonredundant Fourier coefficients in the `odata` array. Pointers to `idata` and `odata` are both required to be aligned to **`cufftComplex`** data type in single-precision transforms and **`cufftDoubleComplex`** data type in double-precision transforms. If `idata` and `odata` are the same, this method does an in-place transform. Note the data layout differences between in-place and out-of-place transforms as described in [Parameter `cufftType`](#).

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>idata</code>	Pointer to the real input data (in GPU memory) to transform
<code>odata</code>	Pointer to the complex output data (in GPU memory)

Output

<code>odata</code>	Contains the complex Fourier coefficients
--------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully executed the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.

<code>CUFFT_INVALID_VALUE</code>	At least one of the parameters <code>idata</code> and <code>odata</code> is not valid.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_EXEC_FAILED</code>	cuFFT failed to execute the transform on the GPU.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

3.9.3. Functions `cufftExecC2R()` and `cufftExecZ2D()`

```

cufftResult
    cufftExecC2R(cufftHandle plan, cufftComplex *idata, cufftReal *odata);
cufftResult
    cufftExecZ2D(cufftHandle plan, cufftComplex *idata, cufftReal *odata);

```

`cufftExecC2R()` (**`cufftExecZ2D()`**) executes a single-precision (double-precision) complex-to-real, implicitly inverse, cuFFT transform plan. cuFFT uses as input data the GPU memory pointed to by the `idata` parameter. The input array holds only the nonredundant complex Fourier coefficients. This function stores the real output values in the `odata` array. and pointers are both required to be aligned to **`cufftComplex`** data type in single-precision transforms and **`cufftDoubleComplex`** type in double-precision transforms. If `idata` and `odata` are the same, this method does an in-place transform.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>idata</code>	Pointer to the complex input data (in GPU memory) to transform
<code>odata</code>	Pointer to the real output data (in GPU memory)

Output

<code>odata</code>	Contains the real output data
--------------------	-------------------------------

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully executed the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INVALID_VALUE</code>	At least one of the parameters <code>idata</code> and <code>odata</code> is not valid.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_EXEC_FAILED</code>	cuFFT failed to execute the transform on the GPU.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

3.9.4. Function `cufftXtExec()`

```

cufftResult
    cufftXtExec(cufftHandle plan, void *input,
                void *output, int direction);

```

Function **`cufftXtExec`** executes any cuFFT transform regardless of precision and type. In case of complex-to-real and real-to-complex transforms **`direction`** parameter is ignored. cuFFT uses the GPU memory pointed to by the **`input`** parameter as input data.

This function stores the Fourier coefficients in the **output** array. If **input** and **output** are the same, this method does an in-place transform.

Input

plan	cufftHandle returned by cufftCreate
input	Pointer to the input data (in GPU memory) to transform
output	Pointer to the output data (in GPU memory)
direction	The transform direction: CUFFT_FORWARD or CUFFT_INVERSE . Ignored for complex-to-real and real-to-complex transforms.

Output

output	Contains the complex Fourier coefficients
---------------	---

Return Values

CUFFT_SUCCESS	cuFFT successfully executed the FFT plan.
CUFFT_INVALID_PLAN	The plan parameter is not a valid handle.
CUFFT_INVALID_VALUE	At least one of the parameters idata , odata , and direction is not valid.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_EXEC_FAILED	cuFFT failed to execute the transform on the GPU.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.

3.9.5. Function **cufftXtExecDescriptor()**

```
cufftResult
cufftXtExecDescriptor(cufftHandle plan, cudaLibXtDesc *input,
                     cudaLibXtDesc *output, int direction);
```

Function **cufftXtExecDescriptor()** executes any cuFFT transform regardless of precision and type. In case of complex-to-real and real-to-complex transforms **direction** parameter is ignored. cuFFT uses the GPU memory pointed to by **cudaLibXtDesc *input** descriptor as input data and **cudaLibXtDesc *output** as output data.

Input

plan	cufftHandle returned by cufftCreate
input	Pointer to the complex input data (in GPU memory) to transform
output	Pointer to the complex output data (in GPU memory)
direction	The transform direction: CUFFT_FORWARD or CUFFT_INVERSE . Ignored for complex-to-real and real-to-complex transforms.

Output

<code>idata</code>	Contains the complex Fourier coefficients
--------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully executed the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INVALID_VALUE</code>	At least one of the parameters <code>idata</code> and <code>direction</code> is not valid.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_EXEC_FAILED</code>	cuFFT failed to execute the transform on the GPU.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_DEVICE</code>	An invalid GPU index was specified in a descriptor.

3.10. cuFFT and Multiple GPUs

3.10.1. Function `cufftXtSetGPUs()`

```

cufftResult
cufftXtSetGPUs(cufftHandle plan, int nGPUs, int *whichGPUs);

```

`cufftXtSetGPUs()` identifies which GPUs are to be used with the plan. As in the single GPU case `cufftCreate()` creates a plan and `cufftMakePlan*` does the plan generation. This call will return an error if a non-default stream has been associated with the plan.

Note that the call to `cufftXtSetGPUs()` must occur after the call to `cufftCreate()` and prior to the call to `cufftMakePlan*`. Parameter `whichGPUs` of `cufftXtSetGPUs()` function determines ordering of the GPUs with respect to data decomposition (first data chunk is placed on GPU denoted by first element of `whichGPUs`).

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>nGPUs</code>	Number of GPUs to use
<code>whichGPUs</code>	The GPUs to use

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully set the GPUs to use.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle, or a non-default stream has been associated with the plan.
<code>CUFFT_ALLOC_FAILED</code>	The allocation of GPU resources for the plan failed.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

<code>CUFFT_INVALID_VALUE</code>	The requested number of GPUs was less than 2 or more than 8.
<code>CUFFT_INVALID_DEVICE</code>	An invalid GPU index was specified.
<code>CUFFT_INVALID_SIZE</code>	Transform size that <code>plan</code> was created for does not meet minimum size criteria.

3.10.2. Function `cufftXtSetWorkArea()`

```
cufftResult
cufftXtSetWorkArea(cufftHandle plan, void **workArea);
```

`cufftXtSetWorkArea()` overrides the work areas associated with a plan. If the work area was auto-allocated, cuFFT frees the auto-allocated space. The **`cufftXtExec*()`** calls assume that the work area is valid and that it points to a contiguous region in each device memory that does not overlap with any other work area. If this is not the case, results are indeterminate.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>workArea</code>	Pointer to the pointers to workArea

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully allows user to override workArea pointer.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_DEVICE</code>	A GPU associated with the plan could not be selected.

3.10.3. cuFFT Multiple GPU Execution

3.10.3.1. Functions `cufftXtExecDescriptorC2C()` and `cufftXtExecDescriptorZ2Z()`

```
cufftResult
cufftXtExecDescriptorC2C(cufftHandle plan, cudaLibXtDesc *input,
                        cudaLibXtDesc *output, int direction);
cufftResult
cufftXtExecDescriptorZ2Z(cufftHandle plan, cudaLibXtDesc *input,
                        cudaLibXtDesc *output, int direction);
```

`cufftXtExecDescriptorC2C()` (**`cufftXtExecDescriptorZ2Z()`**) executes a single-precision (double-precision) complex-to-complex transform plan in the transform direction as specified by `direction` parameter. cuFFT uses the GPU memory pointed to by `cudaLibXtDesc *input` as input data. Since only in-place multiple GPU functionality is support, this function also stores the result in the `cudaLibXtDesc *input` arrays.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>input</code>	Pointer to the complex input data (in GPU memory) to transform
<code>output</code>	Pointer to the complex output data (in GPU memory)
<code>direction</code>	The transform direction: <code>CUFFT_FORWARD</code> or <code>CUFFT_INVERSE</code>

Output

<code>input</code>	Contains the complex Fourier coefficients
--------------------	---

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully executed the FFT plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle.
<code>CUFFT_INVALID_VALUE</code>	At least one of the parameters <code>input</code> and <code>direction</code> is not valid.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_EXEC_FAILED</code>	cuFFT failed to execute the transform on the GPU.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.
<code>CUFFT_INVALID_DEVICE</code>	An invalid GPU index was specified in a descriptor.

3.10.4. Memory Allocation and Data Movement Functions

Multiple GPU cuFFT execution functions assume a certain data layout in terms of what input data has been copied to which GPUs prior to execution, and what output data resides in which GPUs post execution. The following functions assist in allocation, setup and retrieval of the data. They must be called after the call to `cufftMakePlan*()`.

3.10.4.1. Function `cufftXtMalloc()`

```
cufftResult
cufftXtMalloc(cufftHandle plan, cudaLibXtDesc **descriptor,
             cufftXtSubFormat format);
```

`cufftXtMalloc()` allocates a descriptor, and all memory for data in GPUs associated with the plan, and returns a pointer to the descriptor. Note the descriptor contains an array of device pointers so that the application may preprocess or postprocess the data on the GPUs. The enumerated parameter `cufftXtSubFormat_t` indicates if the buffer will be used for input or output.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>descriptor</code>	Pointer to a pointer to a <code>cudaLibXtDesc</code> object
<code>format</code>	<code>cufftXtSubFormat</code> value

Output

descriptor	Pointer to a pointer to a <code>cudaLibXtDesc</code> object
-------------------	---

Return Values

CUFFT_SUCCESS	cuFFT successfully allows user to allocate descriptor and GPU memory.
CUFFT_INVALID_PLAN	The <code>plan</code> parameter is not a valid handle or it is not a multiple GPU <code>plan</code> .
CUFFT_ALLOC_FAILED	The allocation of GPU resources for the plan failed.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_DEVICE	An invalid GPU index was specified in the descriptor.

3.10.4.1.1. Parameter `cufftXtSubFormat`

`cufftXtSubFormat_t` is an enumerated type that indicates if the buffer will be used for input or output and the ordering of the data.

```
typedef enum cufftXtSubFormat_t {
    CUFFT_XT_FORMAT_INPUT,           //by default input is in linear order
    across GPUs
    CUFFT_XT_FORMAT_OUTPUT,         //by default output is in scrambled
    order depending on transform
    CUFFT_XT_FORMAT_INPLACE,        //by default inplace is input order,
    which is linear across GPUs
    CUFFT_XT_FORMAT_INPLACE_SHUFFLED, //shuffled output order after execution
    of the transform
    CUFFT_FORMAT_UNDEFINED
} cufftXtSubFormat;
```

3.10.4.2. Function `cufftXtFree()`

```
cufftResult
cufftXtFree(cudaLibXtDesc *descriptor);
```

`cufftXtFree()` frees the descriptor and all memory associated with it. The descriptor and memory must have been returned by a previous call to **`cufftXtMalloc()`**.

Input

descriptor	Pointer to a <code>cudaLibXtDesc</code> object
-------------------	--

Return Values

CUFFT_SUCCESS	cuFFT successfully allows user to free descriptor and associated GPU memory.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.

3.10.4.3. Function `cufftXtMemcpy()`

```
cufftResult
cufftXtMemcpy(cufftHandle plan, void *dstPointer, void *srcPointer,
cufftXtCopyType type);
```

cuFFT_XtMemcpy () copies data between buffers on the host and GPUs or between GPUs. The enumerated parameter **cuFFT_XtCopyType_t** indicates the type and direction of transfer.

Input

plan	cuFFTHandle returned by cuFFTCreate
dstPointer	Pointer to the destination address(es)
srcPointer	Pointer to the source address(es)
type	cuFFT_XtCopyType value

Return Values

CUFFT_SUCCESS	cuFFT successfully allows user to copy memory between host and GPUs or between GPUs.
CUFFT_INVALID_PLAN	The plan parameter is not a valid handle.
CUFFT_INVALID_VALUE	One or more invalid parameters were passed to the API.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_SETUP_FAILED	The cuFFT library failed to initialize.
CUFFT_INVALID_DEVICE	An invalid GPU index was specified in a descriptor.

3.10.4.3.1. Parameter **cuFFT_XtCopyType**

cuFFT_XtCopyType_t is an enumerated type for multiple GPU functions that specifies the type of copy for **cuFFT_XtMemcpy ()**.

CUFFT_COPY_HOST_TO_DEVICE copies data from a contiguous host buffer to multiple device buffers, in the layout cuFFT requires for input data. **dstPointer** must point to a **cudaLibXtDesc** structure, and **srcPointer** must point to a host memory buffer.

CUFFT_COPY_DEVICE_TO_HOST copies data from multiple device buffers, in the layout cuFFT produces for output data, to a contiguous host buffer. **dstPointer** must point to a host memory buffer, and **srcPointer** must point to a **cudaLibXtDesc** structure.

CUFFT_COPY_DEVICE_TO_DEVICE copies data from multiple device buffers, in the layout cuFFT produces for output data, to multiple device buffers, in the layout cuFFT requires for input data. **dstPointer** and **srcPointer** must point to different **cudaLibXtDesc** structures (and therefore memory locations). That is, the copy cannot be in-place.

```
typedef enum cuFFT_XtCopyType_t {
    CUFFT_COPY_HOST_TO_DEVICE,
    CUFFT_COPY_DEVICE_TO_HOST,
    CUFFT_COPY_DEVICE_TO_DEVICE
} cuFFT_XtCopyType;
```

3.10.5. General Multiple GPU Descriptor Types

3.10.5.1. cudaXtDesc

A descriptor type used in multiple GPU routines that contains information about the GPUs and their memory locations.

```

struct cudaXtDesc_t{
    int version;                //descriptor version
    int nGPUs;                  //number of GPUs
    int GPUs[MAX_CUDA_DESCRIPTOR_GPUS]; //array of device IDs
    void *data[MAX_CUDA_DESCRIPTOR_GPUS]; //array of pointers to data, one
per GPU
    size_t size[MAX_CUDA_DESCRIPTOR_GPUS]; //array of data sizes, one per GPU
    void *cudaXtState;         //opaque CUDA utility structure
};
typedef struct cudaXtDesc_t cudaXtDesc;

```

3.10.5.2. cudaLibXtDesc

A descriptor type used in multiple GPU routines that contains information about the library used.

```

struct cudaLibXtDesc_t{
    int version;                //descriptor version
    cudaXtDesc *descriptor;     //multi-GPU memory descriptor
    libFormat library;          //which library recognizes the format
    int subFormat;              //library specific enumerator of sub formats
    void *libDescriptor;        //library specific descriptor e.g. FFT transform
plan object
};
typedef struct cudaLibXtDesc_t cudaLibXtDesc;

```

3.11. cuFFT Callbacks

3.11.1. Function cufftXtSetCallback()

```

cufftResult
    cufftXtSetCallback(cufftHandle plan, void **callbackRoutine,
    cufftXtCallbackType type, void **callerInfo)

```

cufftXtSetCallback() specifies a load or store callback to be used with the plan. This call is valid only after a call to **cufftMakePlan*()**, which does the plan generation. If there was already a callback of this type associated with the plan, this new callback routine replaces it. If the new callback requires shared memory, you must call **cufftXtSetCallbackSharedSize** with the amount of shared memory it needs. cuFFT will not retain the amount of shared memory associated with the previous callback.

Input

plan	cufftHandle returned by cufftCreate
callbackRoutine	Array of callback routine pointers, one per GPU
type	type of callback routine
callerInfo	optional array of device pointers to caller specific information, one per GPU

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully associated the callback function with the plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle, or a non-default stream has been associated with the plan.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.
<code>CUFFT_SETUP_FAILED</code>	The cuFFT library failed to initialize.

3.11.2. Function `cufftXtClearCallback()`

```
cufftResult
cufftXtClearCallback(cufftHandle plan, cufftXtCallbackType type)
```

`cufftXtClearCallback()` instructs cuFFT to stop invoking the specified callback type when executing the plan. Only the specified callback is cleared. If no callback of this type had been specified, the return code is `CUFFT_SUCCESS`.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>type</code>	type of callback routine

Return Values

<code>CUFFT_SUCCESS</code>	cuFFT successfully disassociated the callback function with the plan.
<code>CUFFT_INVALID_PLAN</code>	The <code>plan</code> parameter is not a valid handle, or a non-default stream has been associated with the plan.
<code>CUFFT_INTERNAL_ERROR</code>	An internal driver error was detected.

3.11.3. Function `cufftXtSetCallbackSharedSize()`

```
cufftResult
cufftXtSetCallbackSharedSize(cufftHandle plan, cufftXtCallbackType type,
size_t sharedSize)
```

`cufftXtSetCallbackSharedSize()` instructs cuFFT to dynamically allocate shared memory at launch time, for use by the callback. The maximum allowable amount of shared memory is 16K bytes. cuFFT passes a pointer to this shared memory to the callback routine at execution time. This shared memory is only valid for the life of the load or store callback operation. During execution, cuFFT may overwrite shared memory for its own purposes.

Input

<code>plan</code>	<code>cufftHandle</code> returned by <code>cufftCreate</code>
<code>type</code>	type of callback routine
<code>sharedSize</code>	amount of shared memory requested

Return Values

CUFFT_SUCCESS	cuFFT will invoke the callback routine with a pointer to the requested amount of shared memory.
CUFFT_INVALID_PLAN	The <code>plan</code> parameter is not a valid handle, or a non-default stream has been associated with the plan.
CUFFT_INTERNAL_ERROR	An internal driver error was detected.
CUFFT_ALLOC_FAILED	cuFFT will not be able to allocate the requested amount of shared memory.

3.12. Function `cufftSetStream()`

```
cufftResult
cufftSetStream(cufftHandle plan, cudaStream_t stream);
```

Associates a CUDA stream with a cuFFT plan. All kernel launches made during plan execution are now done through the associated stream, enabling overlap with activity in other streams (e.g. data copying). The association remains until the plan is destroyed or the stream is changed with another call to `cufftSetStream()`. This call will return an error for multiple GPU plans.

Input

plan	The <code>cufftHandle</code> object to associate with the stream
stream	A valid CUDA stream created with <code>cudaStreamCreate()</code> ; 0 for the default stream

Status Returned

CUFFT_SUCCESS	The stream was associated with the plan.
CUFFT_INVALID_PLAN	The <code>plan</code> parameter is not a valid handle, or it is a multiple GPU <code>plan</code> .

3.13. Function `cufftGetVersion()`

```
cufftResult
cufftGetVersion(int *version);
```

Returns the version number of cuFFT.

Input

version	Pointer to the version number
----------------	-------------------------------

Output

version	Contains the version number
----------------	-----------------------------

Return Values

CUFFT_SUCCESS	cuFFT successfully returned the version number.
----------------------	---

3.14. Function cufftGetProperty()

```
cufftResult
    cufftGetProperty(libraryPropertyType type, int *value);
```

Return in ***value** the number for the property described by **type** of the dynamically linked CUFFT library.

Input

type	CUDA library property
-------------	-----------------------

Output

value	Contains the integer value for the requested property
--------------	---

Return Values

CUFFT_SUCCESS	the property value was successfully returned.
CUFFT_INVALID_TYPE	the property type is not recognized
CUFFT_INVALID_VALUE	value is NULL

3.15. cuFFT Types

3.15.1. Parameter cufftType

The cuFFT library supports complex- and real-data transforms. The **cufftType** data type is an enumeration of the types of transform data supported by cuFFT.

```
typedef enum cufftType_t {
    CUFFT_R2C = 0x2a, // Real to complex (interleaved)
    CUFFT_C2R = 0x2c, // Complex (interleaved) to real
    CUFFT_C2C = 0x29, // Complex to complex (interleaved)
    CUFFT_D2Z = 0x6a, // Double to double-complex (interleaved)
    CUFFT_Z2D = 0x6c, // Double-complex (interleaved) to double
    CUFFT_Z2Z = 0x69  // Double-complex to double-complex (interleaved)
} cufftType;
```

3.15.2. Parameters for Transform Direction

The cuFFT library defines forward and inverse Fast Fourier Transforms according to the sign of the complex exponential term.

```
#define cuFFTFORWARD -1
#define cuFFTINVERSE 1
```

cuFFT performs un-normalized FFTs; that is, performing a forward FFT on an input data set followed by an inverse FFT on the resulting set yields data that is equal to the input,

scaled by the number of elements. Scaling either transform by the reciprocal of the size of the data set is left for the user to perform as seen fit.

3.15.3. Type definitions for callbacks

The cuFFT library supports callback functions for all combinations of single or double precision, real or complex data, load or store. These are enumerated in the parameter **cuFFT_Xt_CallbackType**.

```
typedef enum cuFFT_Xt_CallbackType_t {
    CUFFT_CB_LD_COMPLEX = 0x0,
    CUFFT_CB_LD_COMPLEX_DOUBLE = 0x1,
    CUFFT_CB_LD_REAL = 0x2,
    CUFFT_CB_LD_REAL_DOUBLE = 0x3,
    CUFFT_CB_ST_COMPLEX = 0x4,
    CUFFT_CB_ST_COMPLEX_DOUBLE = 0x5,
    CUFFT_CB_ST_REAL = 0x6,
    CUFFT_CB_ST_REAL_DOUBLE = 0x7,
    CUFFT_CB_UNDEFINED = 0x8
} cuFFT_Xt_CallbackType;
```

The corresponding function prototypes and pointer type definitions are as follows:

```
typedef cuFFT_Complex (*cuFFT_CallbackLoadC)(void *dataIn, size_t offset, void
    *callerInfo, void *sharedPointer);

typedef cuFFT_DoubleComplex (*cuFFT_CallbackLoadZ)(void *dataIn, size_t offset,
    void *callerInfo, void *sharedPointer);

typedef cuFFT_Real (*cuFFT_CallbackLoadR)(void *dataIn, size_t offset, void
    *callerInfo, void *sharedPointer);

typedef cuFFT_DoubleReal (*cuFFT_CallbackLoadD)(void *dataIn, size_t offset, void
    *callerInfo, void *sharedPointer);

typedef void (*cuFFT_CallbackStoreC)(void *dataOut, size_t offset, cuFFT_Complex
    element, void *callerInfo, void *sharedPointer);

typedef void (*cuFFT_CallbackStoreZ)(void *dataOut, size_t offset,
    cuFFT_DoubleComplex element, void *callerInfo, void *sharedPointer);

typedef void (*cuFFT_CallbackStoreR)(void *dataOut, size_t offset, cuFFT_Real
    element, void *callerInfo, void *sharedPointer);

typedef void (*cuFFT_CallbackStoreD)(void *dataOut, size_t offset,
    cuFFT_DoubleReal element, void *callerInfo, void *sharedPointer);
```

3.15.4. Other cuFFT Types

3.15.4.1. cuFFT_Handle

A handle type used to store and access cuFFT plans. The user receives a handle after creating a cuFFT plan and uses this handle to execute the plan.

```
typedef unsigned int cuFFT_Handle;
```

3.15.4.2. cufftReal

A single-precision, floating-point real data type.

```
typedef float cufftReal;
```

3.15.4.3. cufftDoubleReal

A double-precision, floating-point real data type.

```
typedef double cufftDoubleReal;
```

3.15.4.4. cufftComplex

A single-precision, floating-point complex data type that consists of interleaved real and imaginary components.

```
typedef cuComplex cufftComplex;
```

3.15.4.5. cufftDoubleComplex

A double-precision, floating-point complex data type that consists of interleaved real and imaginary components.

```
typedef cuDoubleComplex cufftDoubleComplex;
```

3.16. Common types

3.16.1. cudaDataType

The **cudaDataType** data type is an enumeration of the types supported by CUDA libraries.

```
typedef enum cudaDataType_t
{
    CUDA_R_16F= 2, // 16 bit real
    CUDA_C_16F= 6, // 16 bit complex
    CUDA_R_32F= 0, // 32 bit real
    CUDA_C_32F= 4, // 32 bit complex
    CUDA_R_64F= 1, // 64 bit real
    CUDA_C_64F= 5, // 64 bit complex
    CUDA_R_8I= 3, // 8 bit real as a signed integer
    CUDA_C_8I= 7, // 8 bit complex as a pair of signed integers
    CUDA_R_8U= 8, // 8 bit real as a signed integer
    CUDA_C_8U= 9 // 8 bit complex as a pair of signed integers
} cudaDataType;
```

3.16.2. libraryPropertyType

The **libraryPropertyType** data type is an enumeration of library property types. (ie. CUDA version X.Y.Z would yield **MAJOR_VERSION=X**, **MINOR_VERSION=Y**, **PATCH_LEVEL=Z**)

```
typedef enum libraryPropertyType_t
{
    MAJOR_VERSION,
    MINOR_VERSION,
    PATCH_LEVEL
} libraryPropertyType;
```

Chapter 4.

CUFFT CODE EXAMPLES

This chapter provides multiple simple examples of complex and real 1D, 2D, and 3D transforms that use cuFFT to perform forward and inverse FFTs.

4.1. 1D Complex-to-Complex Transforms

In this example a one-dimensional complex-to-complex transform is applied to the input data. Afterwards an inverse transform is performed on the computed frequency domain representation.

```

#define NX 256
#define BATCH 1

cufftHandle plan;
cufftComplex *data;
cudaMalloc((void**)&data, sizeof(cufftComplex)*NX*BATCH);
if (cudaGetLastError() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to allocate\n");
    return;
}

if (cufftPlan1d(&plan, NX, CUFFT_C2C, BATCH) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: Plan creation failed");
    return;
}

...

/* Note:
 * Identical pointers to input and output arrays implies in-place
 * transformation
 */

if (cufftExecC2C(plan, data, data, CUFFT_FORWARD) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: ExecC2C Forward failed");
    return;
}

if (cufftExecC2C(plan, data, data, CUFFT_INVERSE) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: ExecC2C Inverse failed");
    return;
}

/*
 * Results may not be immediately available so block device until all
 * tasks have completed
 */

if (cudaDeviceSynchronize() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to synchronize\n");
    return;
}

/*
 * Divide by number of elements in data set to get back original data
 */

...

cufftDestroy(plan);
cudaFree(data);

```

4.2. 1D Real-to-Complex Transforms

In this example a one-dimensional real-to-complex transform is applied to the input data.

```
#define NX 256
#define BATCH 1

cufftHandle plan;
cufftComplex *data;
cudaMalloc((void**)&data, sizeof(cufftComplex)*(NX/2+1)*BATCH);
if (cudaGetLastError() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to allocate\n");
    return;
}

if (cufftPlan1d(&plan, NX, CUFFT_R2C, BATCH) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: Plan creation failed");
    return;
}

...

/* Use the CUFFT plan to transform the signal in place. */
if (cufftExecR2C(plan, (cufftReal*)data, data) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: ExecC2C Forward failed");
    return;
}

if (cudaDeviceSynchronize() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to synchronize\n");
    return;
}

...

cufftDestroy(plan);
cudaFree(data);
```

4.3. 2D Complex-to-Real Transforms

In this example a two-dimensional complex-to-real transform is applied to the input data arranged according to the requirements of the default FFTW padding mode.

```

#define NX 256
#define NY 128
#define NRANK 2
#define BATCH 1

cufftHandle plan;
cufftComplex *data;
int n[NRANK] = {NX, NY};

cudaMalloc((void**)&data, sizeof(cufftComplex)*NX*(NY/2+1));
if (cudaGetLastError() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to allocate\n");
    return;
}

/* Create a 2D FFT plan. */
if (cufftPlanMany(&plan, NRANK, n,
    NULL, 1, 0,
    NULL, 1, 0,
    CUFFT_C2R,BATCH) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT Error: Unable to create plan\n");
    return;
}

...

if (cufftExecC2R(plan, data, data) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT Error: Unable to execute plan\n");
    return;
}

if (cudaDeviceSynchronize() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to synchronize\n");
    return;
}

...

cufftDestroy(plan);
cudaFree(data);

```

4.4. 3D Complex-to-Complex Transforms

In this example a three-dimensional complex-to-complex transform is applied to the input data.

```

#define NX 64
#define NY 128
#define NZ 128
#define BATCH 10
#define NRANK 3

cufftHandle plan;
cufftComplex *data;
int n[NRANK] = {NX, NY, NZ};

cudaMalloc((void**)&data, sizeof(cufftComplex)*NX*NY*NZ*BATCH);
if (cudaGetLastError() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to allocate\n");
    return;
}

/* Create a 3D FFT plan. */
if (cufftPlanMany(&plan, NRANK, n,
    NULL, 1, NX*NY*NZ, // *inembed, istride, idist
    NULL, 1, NX*NY*NZ, // *onembed, ostride, odist
    CUFFT_C2C, BATCH) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: Plan creation failed");
    return;
}

/* Use the CUFFT plan to transform the signal in place. */
if (cufftExecC2C(plan, data, data, CUFFT_FORWARD) != CUFFT_SUCCESS){
    fprintf(stderr, "CUFFT error: ExecC2C Forward failed");
    return;
}

if (cudaDeviceSynchronize() != cudaSuccess){
    fprintf(stderr, "Cuda error: Failed to synchronize\n");
    return;
}

...

cufftDestroy(plan);
cudaFree(data);

```

4.5. 2D Advanced Data Layout Use

In this example a two-dimensional complex-to-complex transform is applied to the input data arranged according to the requirements the advanced layout.

```

#define NX 128
#define NY 256
#define BATCH 10
#define NRANK 2

/* Advanced interface parameters, arbitrary strides */
#define ISTRIDE 2 // distance between successive input elements in innermost
  dimension
#define OSTRIDE 1 // distance between successive output elements in innermost
  dimension
#define IX (NX+2)
#define IY (NY+1)
#define OX (NX+3)
#define OY (NY+4)
#define IDIST (IX*IY*ISTRIDE+3) // distance between first element of two
  consecutive signals in a batch of input data
#define ODIST (OX*OY*OSTRIDE+5) // distance between first element of two
  consecutive signals in a batch of output data

cufftHandle plan;
cufftComplex *idata, *odata;
int isize = IDIST * BATCH;
int osize = ODIST * BATCH;
int n[NRANK] = {NX, NY};
int inembed[NRANK] = {IX, IY}; // pointer that indicates storage dimensions of
  input data
int onembed[NRANK] = {OX, OY}; // pointer that indicates storage dimensions of
  output data

cudaMalloc((void **)&idata, sizeof(cufftComplex)*isize);
cudaMalloc((void **)&odata, sizeof(cufftComplex)*osize);
if (cudaGetLastError() != cudaSuccess){
  fprintf(stderr, "Cuda error: Failed to allocate\n");
  return;
}

/* Create a batched 2D plan */
if (cufftPlanMany(&plan, NRANK, n,
  inembed, ISTRIDE, IDIST,
  onembed, OSTRIDE, ODIST,
  CUFFT_C2C, BATCH) != CUFFT_SUCCESS){
  fprintf(stderr, "CUFFT Error: Unable to create plan\n");
  return;
}

...

/* Execute the transform out-of-place */
if (cufftExecC2C(plan, idata, odata, CUFFT_FORWARD) != CUFFT_SUCCESS){
  fprintf(stderr, "CUFFT Error: Failed to execute plan\n");
  return;
}

if (cudaDeviceSynchronize() != cudaSuccess){
  fprintf(stderr, "Cuda error: Failed to synchronize\n");
  return;
}

...

cufftDestroy(plan);
cudaFree(idata);
cudaFree(odata);

```

4.6. 3D Complex-to-Complex Transforms using Two GPUs

In this example a three-dimensional complex-to-complex transform is applied to the input data using two GPUs.

```
// Demonstrate how to use CUFFT to perform 3-d FFTs using 2 GPUs
//
// cufftCreate() - Create an empty plan
cufftHandle plan_input; cufftResult result;
result = cufftCreate(&plan_input);
if (result != CUFFT_SUCCESS) { printf ("*Create failed\n"); return; }
//
// cufftXtSetGPUs() - Define which GPUs to use
int nGPUs = 2, whichGPUs[2];
whichGPUs[0] = 0; whichGPUs[1] = 1;
result = cufftXtSetGPUs (plan_input, nGPUs, whichGPUs);
if (result != CUFFT_SUCCESS) { printf ("*XtSetGPUs failed\n"); return; }
//
// Initialize FFT input data
size_t worksize[2];
cufftComplex *host_data_input, *host_data_output;
int nx = 64, ny = 128, nz = 32;
int size_of_data = sizeof(cufftComplex) * nx * ny * nz;
host_data_input = malloc(size_of_data);
if (host_data_input == NULL) { printf ("malloc failed\n"); return; }
host_data_output = malloc(size_of_data);
if (host_data_output == NULL) { printf ("malloc failed\n"); return; }
initialize_3d_data (nx, ny, nz, host_data_input, host_data_output);
//
// cufftMakePlan3d() - Create the plan
result = cufftMakePlan3d (plan_input, nz, ny, nx, CUFFT_C2C, worksize);
if (result != CUFFT_SUCCESS) { printf ("*MakePlan* failed\n"); return; }
//
// cufftXtMalloc() - Malloc data on multiple GPUs
cudaLibXtDesc *device_data_input;
result = cufftXtMalloc (plan_input, &device_data_input,
    CUFFT_XT_FORMAT_INPLACE);
if (result != CUFFT_SUCCESS) { printf ("*XtMalloc failed\n"); return; }
//
// cufftXtMemcpy() - Copy data from host to multiple GPUs
result = cufftXtMemcpy (plan_input, device_data_input,
    host_data_input, CUFFT_COPY_HOST_TO_DEVICE);
if (result != CUFFT_SUCCESS) { printf ("*XtMemcpy failed\n"); return; }
//
// cufftXtExecDescriptorC2C() - Execute FFT on multiple GPUs
result = cufftXtExecDescriptorC2C (plan_input, device_data_input,
    device_data_input, CUFFT_FORWARD);
if (result != CUFFT_SUCCESS) { printf ("*XtExec* failed\n"); return; }
//
// cufftXtMemcpy() - Copy data from multiple GPUs to host
result = cufftXtMemcpy (plan_input, host_data_output,
    device_data_input, CUFFT_COPY_DEVICE_TO_HOST);
if (result != CUFFT_SUCCESS) { printf ("*XtMemcpy failed\n"); return; }
//
// Print output and check results
int output_return = output_3d_results (nx, ny, nz,
    host_data_input, host_data_output);
if (output_return != 0) { return; }
//
// cufftXtFree() - Free GPU memory
result = cufftXtFree(device_data_input);
if (result != CUFFT_SUCCESS) { printf ("*XtFree failed\n"); return; }
//
// cufftDestroy() - Destroy FFT plan
result = cufftDestroy(plan_input);
if (result != CUFFT_SUCCESS) { printf ("*Destroy failed: code\n"); return; }
free(host_data_input); free(host_data_output);
```

4.7. 1D Complex-to-Complex Transforms using Two GPUs with Natural Order

In this example a one-dimensional complex-to-complex transform is applied to the input data using two GPUs. The output data is in natural order in GPU memory.

```
// Demonstrate how to use CUFFT to perform 1-d FFTs using 2 GPUs
// Output on the GPUs is in natural output
// Function return codes should be checked for errors in actual code
//
// cufftCreate() - Create an empty plan
cufftHandle plan_input; cufftResult result;
result = cufftCreate(&plan_input);
//
// cufftXtSetGPUs() - Define which GPUs to use
int nGPUs = 2, whichGPUs[2];
whichGPUs[0] = 0; whichGPUs[1] = 1;
result = cufftXtSetGPUs (plan_input, nGPUs, whichGPUs);
//
// Initialize FFT input data
size_t worksize[2];
cufftComplex *host_data_input, *host_data_output;
int nx = 1024, batch = 1, rank = 1, n[1];
int inembed[1], istride, idist, onembed[1], ostride, odist;
n[0] = nx;
int size_of_data = sizeof(cufftComplex) * nx * batch;
host_data_input = malloc(size_of_data);
host_data_output = malloc(size_of_data);
initialize_1d_data (nx, batch, rank, n, inembed, &istride, &idist,
onembed, &ostride, &odist, host_data_input, host_data_output);
//
// cufftMakePlanMany() - Create the plan
result = cufftMakePlanMany (plan_input, rank, n, inembed, istride, idist,
onembed, ostride, odist, CUFFT_C2C, batch, worksize);
//
// cufftXtMalloc() - Malloc data on multiple GPUs
cudaLibXtDesc *device_data_input, *device_data_output;
result = cufftXtMalloc (plan_input, &device_data_input,
CUFFT_XT_FORMAT_INPLACE);
result = cufftXtMalloc (plan_input, &device_data_output,
CUFFT_XT_FORMAT_INPLACE);
//
// cufftXtMemcpy() - Copy data from host to multiple GPUs
result = cufftXtMemcpy (plan_input, device_data_input,
host_data_input, CUFFT_COPY_HOST_TO_DEVICE);
//
// cufftXtExecDescriptorC2C() - Execute FFT on multiple GPUs
result = cufftXtExecDescriptorC2C (plan_input, device_data_input,
device_data_input, CUFFT_FORWARD);
//
// cufftXtMemcpy() - Copy the data to natural order on GPUs
result = cufftXtMemcpy (plan_input, device_data_output,
device_data_input, CUFFT_COPY_DEVICE_TO_DEVICE);
//
// cufftXtMemcpy() - Copy natural order data from multiple GPUs to host
result = cufftXtMemcpy (plan_input, host_data_output,
device_data_output, CUFFT_COPY_DEVICE_TO_HOST);
//
// Print output and check results
int output_return = output_1d_results (nx, batch,
host_data_input, host_data_output);
//
// cufftXtFree() - Free GPU memory
result = cufftXtFree(device_data_input);
result = cufftXtFree(device_data_output);
//
// cufftDestroy() - Destroy FFT plan
result = cufftDestroy(plan_input);
free(host_data_input); free(host_data_output);
```

4.8. 1D Complex-to-Complex Convolution using Two GPUs

In this example a one-dimensional convolution is calculated using complex-to-complex transforms.

```
//
// Demonstrate how to use CUFFT to perform a convolution using 1-d FFTs and
// 2 GPUs. The forward FFTs use both GPUs, while the inverse FFT uses one.
// Function return codes should be checked for errors in actual code.
//
// cufftCreate() - Create an empty plan
cufftResult result; cudaError_t cuda_status;
cufftHandle plan_forward_2_gpus, plan_inverse_1_gpu;
result = cufftCreate(&plan_forward_2_gpus);
result = cufftCreate(&plan_inverse_1_gpu);
//
// cufftXtSetGPUs() - Define which GPUs to use
int nGPUs = 2, whichGPUs[2];
whichGPUs[0] = 0; whichGPUs[1] = 1;
result = cufftXtSetGPUs (plan_forward_2_gpus, nGPUs, whichGPUs);
//
// Initialize FFT input data
size_t worksize[2];
cufftComplex *host_data_input, *host_data_output;
int nx = 1048576, batch = 2, rank = 1, n[1];
int inembed[1], istride, idist, onembed[1], ostride, odist;
n[0] = nx;
int size_of_one_set = sizeof(cufftComplex) * nx;
int size_of_data = size_of_one_set * batch;
host_data_input = (cufftComplex*)malloc(size_of_data);
host_data_output = (cufftComplex*)malloc(size_of_one_set);
initialize_1d_data (nx, batch, rank, n, inembed, &istride, &idist,
onembed, &ostride, &odist, host_data_input, host_data_output);
//
// cufftMakePlanMany(), cufftPlan1d - Create the plans
result = cufftMakePlanMany (plan_forward_2_gpus, rank, n, inembed,
istride, idist, onembed, ostride, odist, CUFFT_C2C, batch, worksize);
result = cufftPlan1d (&plan_inverse_1_gpu, nx, CUFFT_C2C, 1);
//
// cufftXtMalloc(), cudaMallocHost - Allocate data for GPUs
cudaLibXtDesc *device_data_input; cufftComplex *GPU0_data_from_GPU1;
result = cufftXtMalloc (plan_forward_2_gpus, &device_data_input,
CUFFT_XT_FORMAT_INPLACE);
int device0 = device_data_input->descriptor->GPUs[0];
cudaSetDevice(device0) ;
cuda_status = cudaMallocHost ((void*)&GPU0_data_from_GPU1,size_of_one_set);
//
// cufftXtMemcpy() - Copy data from host to multiple GPUs
result = cufftXtMemcpy (plan_forward_2_gpus, device_data_input,
host_data_input, CUFFT_COPY_HOST_TO_DEVICE);
//
// cufftXtExecDescriptorC2C() - Execute forward FFTs on multiple GPUs
result = cufftXtExecDescriptorC2C (plan_forward_2_gpus, device_data_input,
device_data_input, CUFFT_FORWARD);
//
// cudaMemcpy result from GPU1 to GPU0
cufftComplex *device_data_on_GPU1;
device_data_on_GPU1 = (cufftComplex*)
(device_data_input->descriptor->data[1]);
cuda_status = cudaMemcpy (GPU0_data_from_GPU1, device_data_on_GPU1,
size_of_one_set, cudaMemcpyDeviceToDevice);
//
// Continued on next page
```

```

//
// Demonstrate how to use CUFFT to perform a convolution using 1-d FFTs and
// 2 GPUs. The forward FFTs use both GPUs, while the inverse FFT uses one.
// Function return codes should be checked for errors in actual code.
//
// Part 2
//
// Multiply results and scale output
cufftComplex *device_data_on_GPU0;
device_data_on_GPU0 = (cufftComplex*)
    (device_data_input->descriptor->data[0]);
cudaSetDevice(device0) ;
ComplexPointwiseMulAndScale<<<32, 256>>>((cufftComplex*)device_data_on_GPU0,
    (cufftComplex*) GPU0_data_from_GPU1, nx);
//
// cufftExecC2C() - Execute inverse FFT on one GPU
result = cufftExecC2C (plan_inverse_1_gpu, GPU0_data_from_GPU1,
    GPU0_data_from_GPU1, CUFFT_INVERSE);
//
// cudaMemcpy() - Copy results from GPU0 to host
cuda_status = cudaMemcpy(host_data_output, GPU0_data_from_GPU1,
    size_of_one_set, cudaMemcpyDeviceToHost);
//
// Print output and check results
int output_return = output_1d_results (nx, batch,
    host_data_input, host_data_output);
//
// cufftDestroy() - Destroy FFT plans
result = cufftDestroy(plan_forward_2_gpus);
result = cufftDestroy(plan_inverse_1_gpu);
//
// cufftXtFree(), cudaFreeHost(), free() - Free GPU and host memory
result = cufftXtFree(device_data_input);
cuda_status = cudaFreeHost (GPU0_data_from_GPU1);
free(host_data_input); free(host_data_output);

```

```

//
// Utility routine to perform complex pointwise multiplication with scaling
__global__ void ComplexPointwiseMulAndScale
    (cufftComplex *a, cufftComplex *b, int size)
{
    const int numThreads = blockDim.x * gridDim.x;
    const int threadID = blockIdx.x * blockDim.x + threadIdx.x;
    float scale = 1.0f / (float)size;
    cufftComplex c;
    for (int i = threadID; i < size; i += numThreads)
    {
        c = cuCmulf(a[i], b[i]);
        b[i] = make_cuFloatComplex(scale*cuCreal(c), scale*cuCimag(c));
    }
    return;
}

```

Chapter 5.

MULTIPLE GPU DATA ORGANIZATION

This chapter explains how data are distributed between the GPUs, before and after a multiple GPU transform. For simplicity, it is assumed in this chapter that the caller has specified GPU 0 and GPU 1 to perform the transform.

5.1. Multiple GPU Data Organization for Batched Transforms

For batches of transforms, each individual transform is executed on a single GPU. If possible the batches are evenly distributed among the GPUs. For a batch of size m performed on n GPUs, where m is not divisible by n , the first $m \% n$ GPUs will perform $\lfloor \frac{m}{n} \rfloor + 1$ transforms. The remaining GPUs will perform $\lfloor \frac{m}{n} \rfloor$ transforms. For example, in a batch of 15 transforms performed on 4 GPUs, the first three GPUs would perform 4 transforms, and the last GPU would perform 3 transforms. This approach removes the need for data exchange between the GPUs, and results in nearly perfect scaling for cases where the batch size is divisible by the number of GPUs.

5.2. Multiple GPU Data Organization for Single 2D and 3D Transforms

Single transforms performed on multiple GPUs require the data to be divided between the GPUs. Then execution takes place in phases. For example with 2 GPUs, for 2D and 3D transforms with even sized dimensions, each GPU does half of the transform in (rank - 1) dimensions. Then data are exchanged between the GPUs so that the final dimension can be processed.

Since 2D and 3D transforms support sizes other than powers of 2, it is possible that the data can not be evenly distributed among the GPUs. In general for the case of n GPUs, a dimension of size m that is not a multiple of n would be distributed such that the first $m \% n$ GPUs would get one extra row for 2D transforms, one extra plane for 3D transforms.

Take for example, a 2D transform on 4 GPUs, using an array declared in C as `data[x][y]`, where `x` is 65 and `y` is 99. The surface is distributed prior to the transform such that GPU 0 receives a surface with dimensions `[17][99]`, and GPUs 1...3 receive surfaces with dimensions `[16][99]`. After the transform, each GPU again has a portion of the surface, but divided in the `y` dimension. GPUs 0...2 have surfaces with dimensions `[65][25]`. GPU 3 has a surface with dimensions `[65][24]`.

For a 3D transform on 4 GPUs consider an array declared in C as `data[x][y][z]`, where `x` is 103, `y` is 122, and `z` is 64. The volume is distributed prior to the transform such that each GPUs 0...2 receive volumes with dimensions `[26][122][64]`, and GPU 3 receives a volume with dimensions `[26][101][64]`. After the transform, each GPU again has a portion of the surface, but divided in the `y` dimension. GPUs 0 and 1 have a volumes with dimensions `[103][31][64]`, and GPUs 2 and 3 have volumes with dimensions `[103][30][64]`.

5.3. Multiple-GPU Data Organization for Single 1D Transforms

By default for 1D transforms, the initial distribution of data to the GPUs is similar to the 2D and 3D cases. For a transform of dimension `x` on two GPUs, GPU 0 receives data ranging from `0...(x/2-1)`. GPU 1 receives data ranging from `(x/2)...(x-1)`. Similarly, with 4 GPUs, the data are evenly distributed among all 4 GPUs.

Before computation can begin, data are redistributed among the GPUs. It is possible to perform this redistribution in the copy from host memory, in cases where the application does not need to pre-process the data prior to the transform. To do this, the application can create the data descriptor with `cuFFTxFtMalloc` using the sub-format `CUFFT_XT_FORMAT_1D_INPUT_SHUFFLED`. This can significantly reduce the time it takes to execute the transform.

cuFFT performs multiple GPU 1D transforms by decomposing the transform size into factors `Factor1` and `Factor2`, and treating the data as a grid of size `Factor1` x `Factor2`. The four steps done to calculate the 1D FFT are: `Factor1` transforms of size `Factor2`, data exchange between the GPUs, a pointwise twiddle multiplication, and `Factor2` transforms of size `Factor1`.

To gain efficiency by overlapping computation with data exchange, cuFFT breaks the whole transform into independent segments or strings, which can be processed while others are in flight. A side effect of this algorithm is that the output of the transform is not in linear order. The output in GPU memory is in strings, each of which is composed of `Factor2` substrings of equal size. Each substring contains contiguous results starting `Factor1` elements subsequent to start of the previous substring. Each string starts

substring size elements after the start of the previous string. The strings appear in order, the first half on GPU 0, and the second half on GPU 1. See the example below:

```
transform size = 1024
number of strings = 8
Factor1 = 64
Factor2 = 16
substrings per string for output layout is Factor2 (16)
string size = 1024/8 = 128
substring size = 128/16 = 8
stride between substrings = 1024/16 = Factor1 (64)

On GPU 0:
string 0 has substrings with indices 0...7 64...71 128...135 ... 960...967
string 1 has substrings with indices 8...15 72...79 136...143 ... 968...975
...
On GPU 1:
string 4 has substrings with indices 32...39 96...103 160...167 ... 992...999
...
string 7 has substrings with indices 56...63 120...127 184...191 ... 1016...1023
```

The `cufftXtQueryPlan` API allows the caller to retrieve a structure containing the number of strings, the decomposition factors, and (in the case of power of 2 size) some useful mask and shift elements. The example below shows how `cufftXtQueryPlan` is invoked. It

also shows how to translate from an index in the host input array to the corresponding index on the device, and vice versa.

```

/*
 * These routines demonstrate the use of cufftXtQueryPlan to get the 1D
 * factorization and convert between permuted and linear indexes.
 */
/*
 * Set up a 1D plan that will execute on GPU 0 and GPU1, and query
 * the decomposition factors
 */
int main(int argc, char **argv){
    cufftHandle plan;
    cufftResult stat;
    int whichGPUs[2] = { 0, 1 };
    cufftXt1dFactors factors;
    stat = cufftCreate( &plan );
    if (stat != CUFFT_SUCCESS) {
        printf("Create error %d\n",stat);
        return 1;
    }
    stat = cufftXtSetGPUs( plan, 2, whichGPUs );
    if (stat != CUFFT_SUCCESS) {
        printf("SetGPU error %d\n",stat);
        return 1;
    }
    stat = cufftMakePlan1d( plan, size, CUFFT_C2C, 1, workSizes );
    if (stat != CUFFT_SUCCESS) {
        printf("MakePlan error %d\n",stat);
        return 1;
    }
    stat = cufftXtQueryPlan( plan, (void *) &factors, CUFFT_QUERY_1D_FACTORS );
    if (stat != CUFFT_SUCCESS) {
        printf("QueryPlan error %d\n",stat);
        return 1;
    }
    printf("Factor 1 %zd, Factor2 %zd\n",factors.factor1,factors.factor2);
    cufftDestroy(plan);
    return 0;
}

```

```

/*
 * Given an index into a permuted array, and the GPU index return the
 * corresponding linear index from the beginning of the input buffer.
 *
 * Parameters:
 *   factors      input:  pointer to cufftXtldFactors as returned by
 *                       cufftXtQueryPlan
 *   permutedIx   input:  index of the desired element in the device output
 *                       array
 *   linearIx     output: index of the corresponding input element in the
 *                       host array
 *   GPUix        input:  index of the GPU containing the desired element
 */
cufftResult permuted2Linear( cufftXtldFactors * factors,
                             size_t permutedIx,
                             size_t *linearIx,
                             int GPUix ) {
    size_t indexInSubstring;
    size_t whichString;
    size_t whichSubstring;
    // the low order bits of the permuted index match those of the linear index
    indexInSubstring = permutedIx & factors->substringMask;
    // the next higher bits are the substring index
    whichSubstring = (permutedIx >> factors->substringShift) &
                     factors->factor2Mask;
    // the next higher bits are the string index on this GPU
    whichString = (permutedIx >> factors->stringShift) & factors->stringMask;
    // now adjust the index for the second GPU
    if (GPUix) {
        whichString += factors->stringCount/2;
    }
    // linear index low order bits are the same
    // next higher linear index bits are the string index
    *linearIx = indexInSubstring + ( whichString << factors->substringShift );
    // next higher bits of linear address are the substring index
    *linearIx += whichSubstring << factors->factor1Shift;
    return CUFFT_SUCCESS;
}

```

```

/*
 * Given a linear index into a 1D array, return the GPU containing the permuted
 * result, and index from the start of the data buffer for that element.
 *
 * Parameters:
 *   factors      input: pointer to cufftXtldFactors as returned by
 *                       cufftXtQueryPlan
 *   linearIx     input: index of the desired element in the host input
 *                       array
 *   permutedIx   output: index of the corresponding result in the device
 *                       output array
 *   GPUix        output: index of the GPU containing the result
 */
cufftResult linear2Permuted( cufftXtldFactors * factors,
                             size_t linearIx,
                             size_t *permutedIx,
                             int *GPUix ) {
    size_t indexInSubstring;
    size_t whichString;
    size_t whichSubstring;
    size_t whichStringMask;
    int whichStringShift;
    if (linearIx >= factors->size) {
        return CUFFT_INVALID_VALUE;
    }
    // get a useful additional mask and shift count
    whichStringMask = factors->stringCount - 1;
    whichStringShift = (factors->factor1Shift + factors->factor2Shift) -
        factors->stringShift ;
    // the low order bits identify the index within the substring
    indexInSubstring = linearIx & factors->substringMask;
    // first determine which string has our linear index.
    // the low order bits identify the index within the substring.
    // the next higher order bits identify which string.
    whichString = (linearIx >> factors->substringShift) & whichStringMask;
    // the first stringCount/2 strings are in the first GPU,
    // the rest are in the second.
    *GPUix = whichString / (factors->stringCount / 2);
    // next determine which substring within the string has our index
    // the substring index is in the next higher order bits of the index
    whichSubstring = (linearIx >> (factors->substringShift + whichStringShift)) &
        factors->factor2Mask;
    // now we can re-assemble the index
    *permutedIx = indexInSubstring;
    *permutedIx += whichSubstring << factors->substringShift;
    if ( !*GPUix ) {
        *permutedIx += whichString << factors->stringShift;
    } else {
        *permutedIx += (whichString - (factors->stringCount / 2) ) <<
            factors->stringShift;
    }
    return CUFFT_SUCCESS;
}

```

Chapter 6.

FFTW CONVERSION GUIDE

cuFFT differs from FFTW in that FFTW has many plans and a single execute function while cuFFT has fewer plans, but multiple execute functions. The cuFFT execute functions determine the precision (single or double) and whether the input is complex or real valued. The following table shows the relationship between the two interfaces.

FFTW function	cuFFT function
<code>fftw_plan_dft_1d()</code> , <code>fftw_plan_dft_r2c_1d()</code> , <code>fftw_plan_dft_c2r_1d()</code>	<code>cufftPlan1d()</code>
<code>fftw_plan_dft_2d()</code> , <code>fftw_plan_dft_r2c_2d()</code> , <code>fftw_plan_dft_c2r_2d()</code>	<code>cufftPlan2d()</code>
<code>fftw_plan_dft_3d()</code> , <code>fftw_plan_dft_r2c_3d()</code> , <code>fftw_plan_dft_c2r_3d()</code>	<code>cufftPlan3d()</code>
<code>fftw_plan_dft()</code> , <code>fftw_plan_dft_r2c()</code> , <code>fftw_plan_dft_c2r()</code>	<code>cufftPlanMany()</code>
<code>fftw_plan_many_dft()</code> , <code>fftw_plan_many_dft_r2c()</code> , <code>fftw_plan_many_dft_c2r()</code>	<code>cufftPlanMany()</code>
<code>fftw_execute()</code>	<code>cufftExecC2C()</code> , <code>cufftExecZ2Z()</code> , <code>cufftExecR2C()</code> , <code>cufftExecD2Z()</code> , <code>cufftExecC2R()</code> , <code>cufftExecZ2D()</code>
<code>fftw_destroy_plan()</code>	<code>cufftDestroy()</code>

Chapter 7.

FFTW INTERFACE TO CUFFT

NVIDIA provides FFTW3 interfaces to the cuFFT library. This allows applications using FFTW to use NVIDIA GPUs with minimal modifications to program source code. To use the interface first do the following two steps

- ▶ It is recommended that you replace the include file `fftw3.h` with `cufftw.h`
- ▶ Instead of linking with the double/single precision libraries such as `fftw3/fftw3f` libraries, link with both the cuFFT and cuFFTW libraries
- ▶ Ensure the search path includes the directory containing `cuda_runtime_api.h`

After an application is working using the FFTW3 interface, users may want to modify their code to move data to and from the GPU and use the routines documented in the [FFTW Conversion Guide](#) for the best performance.

The following tables show which components and functions of FFTW3 are supported in cuFFT.

Section in FFTW manual	Supported	Unsupported
Complex numbers	<code>fftw_complex</code> , <code>fftwf_complex</code> types	
Precision	double <code>fftw3</code> , single <code>fftwf3</code>	long double <code>fftw3l</code> , quad precision <code>fftw3q</code> are not supported since CUDA functions operate on double and single precision floating-point quantities
Memory Allocation		<code>fftw_malloc()</code> , <code>fftw_free()</code> , <code>fftw_alloc_real()</code> , <code>fftw_alloc_complex()</code> , <code>fftwf_alloc_real()</code> , <code>fftwf_alloc_complex()</code>
Multi-threaded FFTW		<code>fftw3_threads</code> , <code>fftw3_omp</code> are not supported
Distributed-memory FFTW with MPI		<code>fftw3_mpi</code> , <code>fftw3f_mpi</code> are not supported

Note that for each of the double precision functions below there is a corresponding single precision version with the letters **fftw** replaced by **fftwf**.

Section in FFTW manual	Supported	Unsupported
Using Plans	<code>fftw_execute()</code> , <code>fftw_destroy_plan()</code> , <code>fftw_cleanup()</code> , <code>fftw_print_plan()</code>	<code>fftw_cost()</code> , <code>fftw_flops()</code> exist but are not functional
Basic Interface		
Complex DFTs	<code>fftw_plan_dft_1d()</code> , <code>fftw_plan_dft_2d()</code> , <code>fftw_plan_dft_3d()</code> , <code>fftw_plan_dft()</code>	
Planner Flags		Planner flags are ignored and the same plan is returned regardless
Real-data DFTs	<code>fftw_plan_dft_r2c_1d()</code> , <code>fftw_plan_dft_r2c_2d()</code> , <code>fftw_plan_dft_r2c_3d()</code> , <code>fftw_plan_dft_r2c()</code> , <code>fftw_plan_dft_c2r_1d()</code> , <code>fftw_plan_dft_c2r_2d()</code> , <code>fftw_plan_dft_c2r_3d()</code> , <code>fftw_plan_dft_c2r()</code>	
Read-data DFT Array Format		Not supported
Read-to-Real Transform		Not supported
Read-to-Real Transform Kinds		Not supported
Advanced Interface		
Advanced Complex DFTs	<code>fftw_plan_many_dft()</code> with multiple 1D, 2D, 3D transforms	<code>fftw_plan_many_dft()</code> with 4D or higher transforms or a 2D or higher batch of embedded transforms
Advanced Real-data DFTs	<code>fftw_plan_many_dft_r2c()</code> , <code>fftw_plan_many_dft_c2r()</code> with multiple 1D, 2D, 3D transforms	<code>fftw_plan_many_dft_r2c()</code> , <code>fftw_plan_many_dft_c2r()</code> with 4D or higher transforms or a 2D or higher batch of embedded transforms
Advanced Real-to-Real Transforms		Not supported
Guru Interface		
Interleaved and split arrays	Interleaved format	Split format
Guru vector and transform sizes	<code>fftw_iodim</code> struct	
Guru Complex DFTs	<code>fftw_plan_guru_dft()</code> , <code>fftw_plan_guru_dft_r2c()</code> , <code>fftw_plan_guru_dft_c2r()</code> with multiple 1D, 2D, 3D transforms	<code>fftw_plan_guru_dft()</code> , <code>fftw_plan_guru_dft_r2c()</code> , <code>fftw_plan_guru_dft_c2r()</code> with

Section in FFTW manual	Supported	Unsupported
		4D or higher transforms or a 2D or higher batch of transforms
Guru Real-data DFTs		Not supported
Guru Real-to-real Transforms		Not supported
64-bit Guru Interface		Not supported
New-array Execute Functions	<code>fftw_execute_dft()</code> , <code>fftw_execute_dft_r2c()</code> , <code>fftw_execute_dft_c2r()</code> with interleaved format	Split format and real-to-real functions
Wisdom		<code>fftw_export_wisdom_to_file()</code> , <code>fftw_import_wisdom_from_file()</code> exist but are not functional. Other wisdom functions do not have entry points in the library.

Chapter 8.

DEPRECATED FUNCTIONALITY

Function `cufftSetCompatibilityMode` was removed in version 9.1.

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2007-2018 NVIDIA Corporation. All rights reserved.