Chapter 4

Fast Fourier Transforms

This chapter describes the C* interface to the CMSSL Fast Fourier Transforms (FFTs). The CMSSL currently provides three types of FFTs:

- complex-to-complex FFT (CCFFT)
- real-to-complex FFT (RCFFT)
- complex-to-real FFT (CRFFT)

In addition, CMSSL provides type conversion utilities that convert a real parallel variable to a complex parallel variable suitable for input to the RCFFT, and convert the complex parallel variable produced by the CRFFT back to a real parallel variable.

One section of this chapter is devoted to each of the following topics:

- introduction
- complex-to-complex FFT
- real-to-complex and complex-to-real FFTs
- type conversion utilities
- references

This chapter assumes a basic understanding of Fourier Transforms.

Copyright © 1993 Thinking Machines Corporation

4.1 Introduction

Each FFT provided by the CMSSL can be used for forward and inverse transforms, and can perform a transform on any selected set of axes. The following table compares the three types of FFTs.

Table 5. Comparison of the CCFFT, RCFFT, and CRFFT

	CCFFT	RCFFT	CRFFT
Number of axes	Any	Any	Any
Number of transform axes	Any	Any*	Any*
Length of transform axes	Power of 2	Power of 2	Power of 2
Length of non-transform axes	Any	Any	Any
Input data type	Complex	Real embedded in complex	Complex
Output data type	Complex	Complex	Real embedded in complex
Input ordering	Any	Any	Any
Output ordering	Any	Any	Any
Forward	Twiddles with negative exponent	Twiddles with negative exponent	Twiddles with negative exponent
Inverse	Twiddles with positive exponent	Twiddles with positive exponent	Twiddles with positive exponent
Layout	Any	Any	Any

* The RCFFT and CRFFT are performed along one axis only. The CCFFT can be performed along up to $n-1$ axes, if desired.

Copyright © 1993 Thinking Machines Corporation

4.1.1 The CMSSL FFT Library Calls

CMSSL provides two FFT user interfaces:

- The *Simple FFT* (available for the CCFFT only). This interface is used to transform a parallel variable in the same direction along all axes. You can use the simple FFT to perform multidimensional FFTs, but not to perform multiple independent FFTs concurrently. For multiple instances, you must be able to specify instance axes (axes along which no transform is performed).

- The *Detailed FFT* (available for all FFT types). This interface provides a great deal of flexibility, including support for multiple instances. For each axis of a multidimensional parallel variable, the Detailed FFT allows you to choose the following:

  - Whether a transform is performed along the axis, and if so, in which direction (forward or inverse).

  - The address ordering (normal or bit-reversed) used to store input and output data values along the axis.

  - The scaling factor (none, axis length, or square root of axis length) applied to the transform results. At the end of the FFT, each position of the transformed parallel variable is divided by the product of the scaling factors. (The axes for which you specify "none" as the scaling factor do not contribute to this product.)

These features are described in the man pages later in this chapter.

To perform an FFT, you must follow these steps:

1. Call the **CMSSL_fft_setup** routine.

   This routine computes internal values (including twiddle factors), allocates an internal FFT setup object, and computes an *FFT setup ID*, which you must then use as an argument in all subsequent Simple FFT, Detailed FFT, and deallocation calls associated with this setup call.

2. Call the **CMSSL_fft** or **CMSSL_fft_detailed** routine.

   These routines execute the Simple and Detailed FFTs, respectively. You may follow one call to the setup routine with multiple calls to these routines. However, the result of a call to the setup routine can only be used for parallel variables with the same shape and type as the

Copyright © 1993 Thinking Machines Corporation

parallel variable supplied in the setup routine. If any of these parameters change, you must call **CMSSL_fft_setup** again to create another setup structure.

3. Call the **CMSSL_deallocate_fft_setup** routine.

   This routine deallocates a specified FFT setup strucure and frees the processing node storage space it required.

You may have more than one setup ID active at a time; that is, you may call the setup routine more than once before deallocating any setup IDs. When you call **CMSSL_fft**, **CMSSL_fft_detailed**, or **CMSSL_deallocate_fft_setup**, you must be sure to provide the correct setup ID.

Copyright © 1993 Thinking Machines Corporation

4.2 Complex-to-Complex FFT

The CMSSL complex-to-complex FFT is a Discrete Fourier Transform in which the same call is used to specify a one-dimensional or multidimensional transform on any legal parallel variable of type **`CMSSL_complex_t`** or **`CMSSL_dcomplex_t`**. For example, in Fourier Analysis Cyclic Reduction (FACR) for a three-dimensional problem (see references 1, 6, 16, and 18 in Section 4.5), the parallel variable has at least three axes. You may choose to perform Fourier transformation on planes represented by two of the axes, while tridiagonal systems of equations are solved along the third axis.

4.2.1 Butterfly Computations, Twiddle Factors, and the CCFFT Setup Phase

The CMSSL FFT is of the Cooley-Tukey variety (see references 3 and 4). The Discrete Fourier Transform relies on a set of coefficients known as *twiddle factors*. The twiddle factors are roots of unity that depend only upon the size of the data set and the radix. In the CMSSL CCFFT, the twiddle computations are performed during a setup phase and used during the subsequent evaluation phase. This organization allows you to amortize the expense of twiddle factor computation over several transformations.

The sign of the twiddle factor exponents determines the direction of an FFT. In the CMSSL FFT, the *forward* transform is defined as one that uses twiddle factors with a *negative* exponent, and the *inverse* transform is defined as one that uses twiddle factors with a *positive* exponent.

4.2.2 Bit Ordering and Bit Reversal

The Cooley-Tukey FFT bit-reverses the addresses in which it stores results. For example, if an 8-point FFT is performed, the indices of the output are

0, 4, 2, 6, 1, 5, 3, 7

In binary, this sequence is

000, 100, 010, 110, 001, 101, 011, 111

that is, the normal sequence with bits reversed.

Copyright © 1993 Thinking Machines Corporation

The Detailed FFT allows you to specify whether the address ordering of the input is normal or bit-reversed, and whether the desired address ordering of the result is normal or bit-reversed.

Optimizing CCFFT Address Ordering

For an axis that is being transformed in a CCFFT,

- If you specify *opposite* address orderings for input and output, the Detailed FFT performs the single bit-reversal that is inherent in the Cooley-Tukey algorithm, and produces output that is bit-reversed relative to the input.
- If you specify the *same* address ordering for input and output, the Detailed FFT performs an *extra* bit reversal, and produces output with the same ordering as the input. The extra bit reversal exacts a performance cost.

For an axis that is not being transformed,

- If you specify *opposite* address orderings for input and output, the Detailed FFT performs a single bit-reversal and produces output that is bit-reversed relative to the input.
- If you specify the *same* address ordering for input and output, the Detailed FFT performs *no* bit reversal.

Thus, for optimal CCFFT performance with the Detailed FFT, specify

- Opposite address orderings for input and output, for each axis that is being transformed.
- The same address ordering for input and output, for each axis that is not being transformed.

Note that violating these recommendations along any axis decreases performance.

The Simple FFT does not provide bit-reversal options; it always produces output data with the same bit ordering as the input data. It does this by automatically performing a bit-reversal on the transformed data, thus annihilating the bit-reversal inherent in the FFT.

Copyright © 1993 Thinking Machines Corporation

To bit-reverse the address ordering of a dataset without transforming it, call the Detailed FFT specifying no transform for all axes and specifying *opposite* values for the input and output address ordering.

These points are independent of the directions of the transforms.

4.2.3 Multidimensional and Multiple-Instance FFTs

The CMSSL FFT allows you to specify whether you want to perform a forward transform, an inverse transform, or no transform along each axis of a multidimensional parallel variable. As described later in this chapter, it is also possible to perform bit-reversal on an axis without transforming it. The axes along which you perform a transform or a bit-reversal define the dataset; the axes along which you perform neither a transform nor bit-reversal are the instance axes, representing multiple independent data sets.

The CMSSL FFT routines perform multidimensional and multiple-instance FFTs by combining the following principles:

- If you supply a multidimensional data set, the routines perform a multidimensional FFT by computing a one-dimensional FFT along each axis of the cell in turn.
- If you specify one or more instance axes (possible only with the Detailed FFT, as described below), the routine computes multiple concurrent, independent one-dimensional FFTs along each cell axis in turn.

For example, if your parallel variable has rank 3 and dimensions [*p*][*q*][*r*], you can perform any of the FFTs indicated in Table 6.

Copyright © 1993 Thinking Machines Corporation

Table 6. Possible FFTs with a three-dimensional input parallel variable.

Data Axis Numbers	Instance Axis Numbers	Resulting FFTs
0	1, 2	*qr* concurrent one-dimensional FFTs along axis 0.
1	0, 2	*pr* concurrent one-dimensional FFTs along axis 1.
2	0, 1	*pq* concurrent one-dimensional FFTs along axis 2.
0, 1	2	*r* concurrent two-dimensional FFTs. The dataset for each two-dimensional FFT is defined by axes 0 and 1. Axis 2 defines the instances.
0, 2	1	*q* concurrent two-dimensional FFTs. The dataset for each two-dimensional FFT is defined by axes 0 and 2. Axis 1 defines the instances.
1, 2	0	*p* concurrent two-dimensional FFTs. The dataset for each two-dimensional FFT is defined by axes 1 and 2. Axis 0 defines the instances.
0, 1, 2	none	One three-dimensional FFT on the entire parallel variable.

4.2.4 Current Restrictions

The current FFT algorithm requires that the length of each axis subject to transformation or bit-reversal is equal to some power of 2. All other axes can have any length.

4.2.5 Implementation and Performance

This section discusses the current implementation of the CCFFT and provides performance guidelines.

The CMSSL CCFFT accepts any layout of the parallel variable. Axes may be local or non-local, and may have any weight. (An axis is *local* if it resides on

Copyright © 1993 Thinking Machines Corporation

a single processing element. In terms of the **allocate_detailed_shape** function, which is described in the *C* Programming Guide*, a local axis has a physical mask of 0. A *non-local* axis spans more than one processing element.) However, performance is strongly influenced by the data layout.

The performance characteristics of the CCFFT implementation on the CM are quite different from those observed on many other computers. The differences are a result of the CM memory system. The subsections that follow provide some insight into the implementations and their performance behavior. Several different algorithms are used for the current version of the CCFFT. The combination of algorithms used in a given call depends upon hardware configuration and data layout.

Data Allocation

The default layout used in the CMSSL FFT is block layout, also called the consecutive data allocation scheme. In this scheme, the axis of a shape is divided into groups of successive positions, and the positions in each group are assigned to the same processing element. (This scheme contrasts with the cyclic data allocation scheme, in which successive positions along an axis are assigned to different processing elements.)

The address field of each position of a shape consists of two groups of bits: those that identify the processing element address, and those that identify the local memory address. For example, in the case of a one-dimensional shape of length $N = 2^n$ distributed across $P = 2^p$ processing elements, each address field can be viewed as follows:

$$(i_{n-1} \quad \ldots \quad i_{n-p} \quad | \quad i_{n-p-1} \quad \ldots \quad i_0)$$

The higher-order bits on the left side specify the processing element; the lower-order bits on the right side specify the local memory address. (For multi-dimensional shapes, each axis is encoded separately.)

Communication Requirements for the FFT

In a radix-2 FFT, the data interaction for each stage is between positions of the parallel variable that differ in only one bit in the encoding of their address indices. This bit is the most significant bit in the first stage, and progresses toward the least significant bit in the last stage. A complete FFT has n stages.

Copyright © 1993 Thinking Machines Corporation

An FFT along an axis that is local to a processing element requires no communication between processing elements. An FFT along an axis that spans multiple processing elements requires communication for the butterfly operations whose data elements reside on different processing elements. With block layout, communication is required in the first p FFT stages.

Transpositions

In the current FFT implementation on the CM-5, when data elements on different processing elements are to be combined, the FFT routine transposes the parallel variable so that the butterfly operations for these data elements can be performed locally. This procedure is illustrated in the following two examples.

Example 1: One-Dimensional Parallel Variable. Assume that a one-dimensional parallel variable with 16 positions is allocated on four processing elements as shown on the left side of Figure 3. The address field for this parallel variable is

(3 2 | 1 0)

where address bits 2 and 3 specify the processing element and address bits 0 and 1 specify the local memory address.

The first two FFT stages operate on data elements whose addresses differ in bits 2 and 3, respectively. These stages therefore need to combine data elements residing on different processing elements. The FFT routine first makes these data elements local to a processing element by transposing the parallel variable as shown on the right side of Figure 3. This procedure effectively corresponds to exchanging processing element address bits with local memory address bits. After performing the first two FFT stages, the FFT routine transposes the parallel variable back to the original layout and carries out the remaining two local FFT stages.

Copyright © 1993 Thinking Machines Corporation

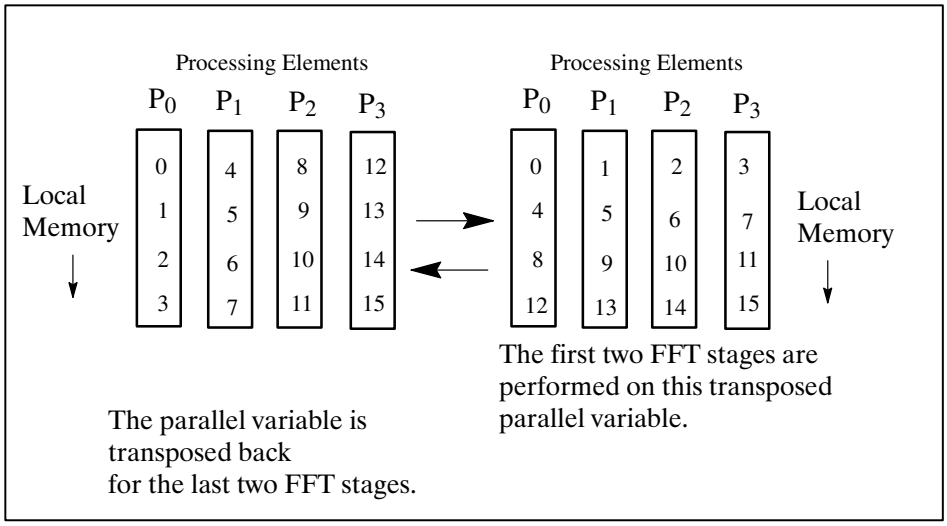

Figure 3. The original parallel variable is transposed for the first two FFT stages and then transposed back for the last two stages.

Example 2: Two-Dimensional Parallel Variable. Assume we want to perform an FFT on a parallel variable of dimensions [*N*][*N*] on *P* processing elements, where *N* is a multiple of *P*. If we lay out this parallel variable so that the first axis is local (in terms of **`allocate_detailed_shape`**, it has a physical mask of 0) and the second is non-local, then the FFT along the first axis can be executed purely locally. For the FFT along the second axis, a transposition is needed. If we view the first axis (axis 0) as an instance axis, the transpose operation places *N*/*P* instances of axis 1 on each processing element, where they can be combined locally. Subsequently, an inverse transposition restores the results into the original geometry.

Performance Considerations

An FFT of a parallel variable of size *N* involves $O(N\log N)$ arithmetic operations. For a non-local FFT, *N* data elements need to be transferred between processing elements during a transposition. Given the fact that on the CM-5 the communication time for transposes dominates the time for FFT arithmetic operations, you can optimize performance by minimizing the number of transpose operations. The following strategies are suggested:

Copyright © 1993 Thinking Machines Corporation

- For a multidimensional parallel variable, specify a layout that makes as many of the FFT axes as possible local. From the discussion in Example 2 above, it follows that for an FFT along a non-local axis, the best performance is obtained if both the length of the FFT axis and the number of instances (the product of the other axis lengths) are multiples of the number of processing elements along the FFT axis.

- Performing an FFT along a non-local axis is equivalent to transposing the data explicitly using the **CMSSL_gen_matrix_transpose** routine, then calling the FFT routine to perform the local FFT, and then calling **CMSSL_gen_matrix_transpose** again to perform the inverse transposition, restoring the data to its original layout. Calling **CMSSL_gen_matrix_transpose** explicitly in this way can be advantageous if your application does not require the inverse transposition. In this case you can reduce the communication time significantly by omitting the second call to **CMSSL_gen_matrix_transpose**. The **CMSSL_gen_matrix_transpose** routine is described in Chapter 6. An illustration of this principle follows.

4.2.6 Efficient Computation of Convolutions and Correlations

One of the main applications of the FFT is to compute convolutions and correlations. To compute the convolution R^*S of two parallel variables R and S, one can proceed as follows:

1. Use the FFT to compute the discrete Fourier transforms $F(R)$ and $F(S)$.

2. Multiply the two transforms together position by position to form the product $F(R)F(S)$.

3. Take the inverse Fourier transform F^{-1} of the product.

(A similar procedure can be used to compute correlations.)

As discussed in Section 4.2.5, an FFT along a non-local axis requires transpositions. The convolution procedure as outlined above would generally require six transpositions. To reduce the number of required transpositions, follow these steps:

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

1. Use **CMSSL_gen_matrix_transpose** to transpose the parallel variables R and S into R^T and S^T so that the FFT axis is local to a processing element.

2. Perform a forward FFT on R^T and S^T. These FFTs will only involve operations on data elements local to a processing element.

3. Form the product $F(R^T)F(S^T)$.

4. Perform an inverse FFT on the product. This FFT will also only involve operations on data elements local to a processing element.

5. Optionally, use **CMSSL_gen_matrix_transpose** to transpose the convolution back to a geometry that conforms with R or S.

This procedure contains only two or three transposes, depending on whether step 5 is included. Because the non-local FFT time is dominated by transpositions, the explicit use of **CMSSL_gen_matrix_transpose** in the above procedure can significantly enhance overall performance.

Copyright © 1993 Thinking Machines Corporation

Complex-to-Complex Fast Fourier Transform

The routines described below use a complex-to-complex Fast Fourier Transform (CCFFT) algorithm to calculate the Discrete Fourier Transform of an *n*-dimensional parallel variable. The Simple FFT call, **CMSSL_fft**, performs the same operation (either a forward or an inverse FFT) along all axes of the parallel variable. With the Detailed FFT call, **CMSSL_fft_detailed**, you can specify different transform operations, bit orderings, and scaling factors for each axis of the parallel variable.

SYNTAX

```
#include <cm/cmssl.h>

int  CMSSL_fft_setup (
     cz:void *A, CMSSL_fft_type_t type, int conj_sym_axis,
     int *setup_id);

int  CMSSL_fft (
     cz:void *A, CMSSL_fft_type_t type, CMSSL_fft_op_t op,
     int setup_id);

int  CMSSL_fft_detailed (
     cz:void *A, CMSSL_fft_type_t type,
     CMSSL_fft_control_t *controls, int setup_id);

void CMSSL_deallocate_fft_setup (
     int setup_id);
```

OVERLOADED DATA TYPES

cz **CMSSL_complex_t** or **CMSSL_dcomplex_t**

The **CMSSL_complex_t** and **CMSSL_dcomplex_t** data types are defined in the header file **<cm/cmssl.h>**.

ARGUMENTS

A Pointer to a complex parallel variable. The **A** you supply in the **CMSSL_fft_setup** call must have the same shape and type as the **A** arguments you supply in all subsequent **CMSSL_fft** or **CMSSL_fft_detailed** calls associated with the setup call. All axes that are to be transformed must have power-of-2 lengths.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

When you call **CMSSL_fft_setup**, you may supply arbitrary values in **A**; the setup routine neither examines nor modifies the contents of this parallel variable, but rather uses its shape and type to create the setup object.

When you call **CMSSL_fft** or **CMSSL_fft_detailed**, **A** is the parallel variable to be transformed. Upon completion, **A** is overwritten with the result. The floating-point precision of the result always matches that of the input.

type Specifies the type of FFT to be performed, as defined below and in the header file **<cm/cmssl.h>**. Supply the value **CMSSL_cmpx_to_cmpx** for a complex-to-complex transform.

```
typedef enum CMSSL_fft_type {
  CMSSL_cmpx_to_cmpx = 0,
  CMSSL_real_to_cmpx = 1,
  CMSSL_cmpx_to_real = 2
} CMSSL_fft_type_t;
```

conj_sym_axis Ignored in the complex-to-complex FFT. Used in the real-to-complex and complex-to-real FFTs; see the man page for the the RCFFT and CRFFT.

op Indicates the direction in which the **A** axes are to be transformed in a Simple FFT, as defined below and in the header file **<cm/cmssl.h>**. Supply the value **CMSSL_f_xform** for a forward transform, **CMSSL_i_xform** for an inverse transform, or **CMSSL_nop** for no transform.

```
typedef enum CMSSL_fft_op {
  CMSSL_nop          = 0,
  CMSSL_f_xform      = 1,
  CMSSL_i_xform      = 2
} CMSSL_fft_op_t;
```

setup_id Internal variable (pointer to an internal variable, when you call the setup routine). When you call **CMSSL_fft**, **CMSSL_fft_detailed**, or **CMSSL_deallocate_fft_setup**, you must supply the value returned by the associated **CMSSL_fft_setup** call.

Copyright © 1993 Thinking Machines Corporation

controls Pointer to an array of structures of type **CMSSL_fft_control_t** (see below). This array has one structure for each axis of **A**. Each element of **controls** specifies the direction of transformation (if any), input address ordering, and output address ordering for the corresponding axis of **A** in a Detailed FFT:

```
typedef struct CMSSL_fft_control {
  CMSSL_fft_op_t      op;
  CMSSL_fft_order_t   in_order;
  CMSSL_fft_order_t   out_order;
  CMSSL_fft_scale_t   scale;
} CMSSL_fft_control_t;
```

The **op** slot specifies the direction (if any) in which the corresponding **A** axis is to be transformed. Supply the value **CMSSL_f_xform** for a forward transform, **CMSSL_i_xform** for an inverse transform, or **CMSSL_nop** for no transform:

```
typedef enum CMSSL_fft_op {
  CMSSL_nop          = 0,
  CMSSL_f_xform      = 1,
  CMSSL_i_xform      = 2
} CMSSL_fft_op_t;
```

The **in_order** and **out_order** slots specify the initial address ordering and desired output address ordering (respectively) of the corresponding **A** axis:

```
typedef enum CMSSL_fft_order {
  CMSSL_normal        = 0,
  CMSSL_bit_reversed = 1
} CMSSL_fft_order_t;
```

The **scale** slot specifies the scaling factor (if any) for the corresponding **A** axis:

```
typedef enum CMSSL_fft_scale {
  CMSSL_noscale       = 0,
  CMSSL_scale_sqrt    = 1,
  CMSSL_scale_n       = 2
} CMSSL_fft_scale_t;
```


Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

The value **CMSSL_scale_n** scales by the axis length; the value **CMSSL_scale_sqrt** scales by the square root of the axis length; and the value **CMSSL_noscale** specifies no scaling. At the end of the **CMSSL_fft_detailed** call, each position of **A** is divided by the product of the scaling factors for all axes. (The axes for which you specified **CMSSL_noscale** do not contribute to this product.)

RETURNED VALUE

The **CMSSL_fft_setup**, **CMSSL_fft**, and **CMSSL_fft_detailed** routines return 0 upon successful return. The **CMSSL_fft** routine may return the following error:

−1 You are attempting to transform axes whose lengths are not a power of 2.

DESCRIPTION

To perform a CCFFT, you must follow these steps:

1. Call **CMSSL_fft_setup**.

   This routine computes internal values (including the twiddle factors), allocates an internal FFT setup object, and computes an FFT setup ID, which you must supply in all subsequent Simple FFT, Detailed FFT, and deallocation calls associated with this setup call.

2. Call **CMSSL_fft** or **CMSSL_fft_detailed**.

   These routines execute the Simple and Detailed FFTs, respectively. You may follow one call to **CMSSL_fft_setup** with multiple calls to these routines. However, the result of a call to **CMSSL_fft_setup** can only be used for parallel variables with the same shape and type as the parallel variable supplied in the setup routine. If these parameters change, you must call **CMSSL_fft_setup** again to create another setup structure.

3. Call **CMSSL_deallocate_fft_setup**.

   This routine removes an FFT setup object that is no longer needed. Because an FFT setup object occupies both partition manager and processing element memory, you should free this space by deallocating the setup object after the completion of all FFT invocations that use it.

The time required to compute the contents of an FFT setup structure is substantially longer than the time required to actually perform an FFT. For this reason, and because it is common to perform FFTs on many parallel variables with the same shape and type, the setup phase is separated from the transform phase.

Copyright © 1993 Thinking Machines Corporation

You may have more than one setup ID active at a time; that is, you may call the setup routine more than once before deallocating any setup IDs. When you call **CMSSL_fft**, **CMSSL_fft_detailed**, or **CMSSL_deallocate_fft_setup**, you must be sure to provide the correct setup ID.

Twiddle Factor Exponents. The sign of the twiddle factor exponents determines the direction of an FFT. In the CMSSL FFT, the *forward* transform is defined as one that uses twiddle factors with a *negative* exponent, and the *inverse* transform is defined as one that uses twiddle factors with a *positive* exponent.

Multidimensional and Multiple-Instance FFTs. The **CMSSL_fft** routine computes a multidimensional transform by performing a one-dimensional transform along each axis in turn. The **CMSSL_fft_detailed** routine supports multiple instances, and computes one or more concurrent, independent one-dimensional FFTs along each data axis in turn.

Scalar Operation Value for Simple FFT. In the **CMSSL_fft** call, the **op** argument determines whether the routine performs a forward or an inverse FFT transform along each axis. The transform direction is the same for all axes. The simple FFT uses the following conventions:

- Along each axis, input and output data are stored using the same bit ordering.
- For a forward FFT, no scaling factors are used. For an inverse FFT, the axis length is used as the scaling factor for each axis.

Scalar Control Array for Detailed FFT. The **CMSSL_fft_detailed** call requires a pointer to an array of structures, **CMSSL_fft_control**, of type **CMSSL_fft_control_t** (one structure for each axis of **A**). Each element of **controls** specifies the direction of transformation (if any), input address ordering, and output address ordering for the corresponding axis of **A** in a Detailed FFT: The slots of **CMSSL_fft_control** are as follows:

- The **op** slot specifies the direction in which the corresponding **A** axis is to be transformed, or specifies that the axis is not to be transformed. In a multiple-instance FFT, the axes to be transformed are the data axes, and the axes that are not to be transformed are the instance axes.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

- The **in_order** and **out_order** slots specify the order in which the original and resulting **A** values are stored along each axis. In normal address ordering (**CMSSL_normal**), data values are stored consecutively. Consecutive order is the only compiler-supported data allocation scheme. In bit-reversed ordering, consecutive data values are stored according to sequential addresses whose high- and low-order bits have been swapped; that is, the consecutive allocation scheme is applied to the bit-reversed indices.

  Although **CMSSL_fft_detailed** can handle any combination of bit orderings, performance is best when the axes of an untransformed parallel variable and a transformed parallel variable have opposite bit-orderings — that is, input normal and output bit-reversed, or input bit-reversed and output normal.

- The **scale** slot specifies a scaling factor (none, square root of axis length, or axis length) for each axis. At the end of the **CMSSL_fft_detailed** call, each position of **A** is divided by the product of the scaling factors. (The axes for which you specified **CMSSL_noscale** do not contribute to this product.)

  For example, suppose **A** has dimensions [m][n][p] and you supply **CMSSL_scale_n** for the first axis, **CMSSL_scale_sqrt** for the second, and **CMSSL_noscale** for the third. At the end of the **CMSSL_fft_detailed** call, each position of **A** is divided by m * sqrt(n).

  One artifact of the Fourier Transform algorithm is that if no scaling is performed, then a forward FFT followed immediately by an inverse FFT results in the original parallel variable with all values multiplied by the FFT size — that is, the product of the lengths of the transformed axes. Scaling is typically used to correct this effect and prevent arithmetic overflow. Typically, a scaling factor of **CMSSL_noscale** is used for all axes in a forward transform, and a scaling factor of **CMSSL_scale_n** is used for all axes being transformed in an inverse transform. Alternatively, you may specify a scaling factor of **CMSSL_scale_sqrt** for all transformed axes in both the forward and the inverse FFTs. This alternative is less efficient, but can reduce the risk of arithmetic overflow for very large numbers.

  Scaling is independent of the operations performed along the axes. For example, it is possible to specify a scaling factor of **CMSSL_scale_n** for an axis that is not being transformed.

Bit-Reversing the Addressing of a Parallel Variable. You can use the Detailed FFT to bit-reverse the addressing of any n-dimensional parallel variable. To do this, specify **CMSSL_nop** for the **op** slot in all elements of **controls**, and specify opposite values for the input and output orderings; that is, either **CMSSL_normal** for the **in_order** values and **CMSSL_bit_reversed** for the **out_order** values, or

Copyright © 1993 Thinking Machines Corporation

CMSSL_bit_reversed for the **in_order** values and **CMSSL_normal** for the **out_order** values.

NOTES

Setup Uses Shape and Type. The setup ID computed by the **CMSSL_fft_setup** call can be used for all parallel variables with the same shape and type as the **A** argument supplied in the **CMSSL_fft_setup** call — and *only* for such parallel variables. If a transform is to be performed on two parallel variables, **A** and **B**, of the same shape and type, then *one* call to the setup routine suffices, even if transforms are performed on different axes of the two parallel variables. But if **A** and **B** have different shapes or types, then you must make one call for each parallel variable.

Setup is Private. The FFT setup data contains information private to the FFT. The format of this data is not documented and may change between CMSSL releases. For this reason, application code should never access or modify the contents of an FFT setup structure.

Bit Reversal and Performance. For an axis that is being transformed,

- If you specify *opposite* address orderings for input and output, the Detailed FFT performs the single bit-reversal that is inherent in the Cooley-Tukey algorithm, and produces output that is bit-reversed relative to the input.
- If you specify the *same* address ordering for input and output, the Detailed FFT performs an *extra* bit reversal, and produces output with the same ordering as the input. The extra bit reversal exacts a performance cost.

For an axis that is not being transformed,

- If you specify *opposite* address orderings for input and output, the Detailed FFT performs a single bit-reversal and produces output that is bit-reversed relative to the input.
- If you specify the *same* address ordering for input and output, the Detailed FFT performs *no* bit reversal.

Thus, for optimal FFT performance, specify

- Opposite address orderings for input and output, for each axis that is being transformed.
- The same address ordering for input and output, for each axis that is not being transformed.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

To bit-reverse the address ordering of a dataset without transforming it, call the Detailed FFT specifying no transform for all axes and specifying *opposite* values for the input and output address ordering.

These points are independent of the directions of the transforms.

EXAMPLES

Sample C* code that uses the complex-to-complex FFT can be found on-line in the subdirectory

```
fft/cstar/
```

of a CMSSL examples directory whose location is site-specific.

Copyright © 1993 Thinking Machines Corporation

4.3 Real-to-Complex and Complex-to-Real FFTs

The CMSSL provides a real-to-complex FFT (RCFFT) for computing the Fourier transform of real data, and a complex-to-real FFT (CRFFT) for the transformation of *conjugate symmetric* complex sequences. The transform X_k (for $k = 0, 1, \ldots, N-1$) of a real sequence x_n ($n = 0, 1, \ldots, N-1$), is conjugate symmetric, i.e., $X_k = \overline{X_{N-k}}$, and conversely, the transform of such a sequence is real. The Fourier Transform of a real or conjugate symmetric sequence can be computed using half the storage and half the arithmetic of a complex-to-complex FFT (CCFFT). The savings in arithmetic and storage are realized by only computing X_k for $k = 0, 1, \ldots, N/2$, and exploiting the fact that X_0 and $X_{N/2}$ are both real. These numbers are stored as one complex number. As a result, the conjugate symmetric sequence requires precisely the same amount of storage as the real sequence. We refer to the domain in which the real data exists as the *real domain*, and the domain in which the conjugate symmetric sequence exists as the *complex domain*.

The CMSSL RCFFT and CRFFT can be used to perform a transform along any axis of a multidimensional parallel variable. We refer to this axis as the *conjugate symmetric axis*, since after a real-to-complex transform along this axis, it represents half of a conjugate symmetric sequence. The CMSSL RCFFT and CRFFT compute the transform in place and fully realize the savings in storage.

The reduced need for storage also results in a performance gain. If the conjugate symmetric axis extends over several processing elements, then the bit-reversal operation needed to make the output ordering the same as the input ordering only needs to move half as much data as if the real parallel variable of size N had been mapped to a complex parallel variable of size N. Furthermore, for multidimensional parallel variables, subsequent Fourier transforms, if any, operate on a data set half the size, and a factor of two performance gain is realized.

The performance gain in the CMSSL RCFFT or CRFFT itself is best when the conjugate symmetric axis is local to a processing element. Some data reordering is required, and the efficiency in performing the arithmetic is not quite as high as in the CCFFT. Therefore, the performance gain is slightly less than the ideal factor of two. When the conjugate symmetric axis is distributed across several processing elements, then the reordering time is significant, and the total communication time for a real parallel variable of size N is comparable to the communication time for a complex parallel variable of size N. Hence, in this case the performance gain for the RCFFT or

Copyright © 1993 Thinking Machines Corporation

CRFFT proper is marginal. However, the aforementioned performance gains from the reduced data set size can still be realized.

4.3.1 Relationship of CRFFT, RCFFT, and CCFFT

The CMSSL RCFFT is based upon the CMSSL CCFFT followed by postprocessing. Similarly, the CRFFT has a preprocessing stage followed by a CCFFT.

The input parallel variable to the RCFFT must be complex, with a pair of successive data elements along the conjugate symmetric axis forming a complex number. For one-dimensional data, this is a natural ordering that may be performed without any data motion. But for a multidimensional parallel variable, a data reordering is required, in general, as discussed in Section 4.3.2 below. If there are at least two data elements per processing element along the axis to be transformed, then local memory reordering suffices; otherwise, communication is required. Similarly, the output of the CRFFT is a complex parallel variable, with complex numbers forming pairs of successive real elements along the conjugate symmetric axis. The type conversion utilities described in Section 4.4 convert a real parallel variable to a complex parallel variable suitable for transformation by the RCFFT, and convert the complex parallel variable resulting from the CRFFT into a real parallel variable.

In the CMSSL RCFFT and CRFFT, a forward transform uses twiddle factors with a negative exponent, and an inverse transform uses twiddle factors with a positive exponent, just as in the CCFFT. Therefore, to invert an RCFFT, it is necessary to change the operation from forward to inverse as well as changing the FFT type from real-to-complex to complex-to-real.

4.3.2 Data Orderings

For general information about normal and bit-reversed data orderings, refer to Section 4.2.2.

For the RCFFT and CRFFT, there are several important data ordering issues. In the complex domain, the ordering is normal when successive complex data elements along an axis are assigned to successive locations in the complex parallel variable; the the ordering is bit-reversed if the complex data

Copyright © 1993 Thinking Machines Corporation

elements are assigned to locations in the parallel variable based on the bit-reversed values of their indices, as described for the CCFFT. Figure 4 illustrates the data allocation in the complex domain for a one-dimensional parallel variable. For the axis representing the conjugate symmetric sequence, location zero is special. It represents two real elements, namely the real components of frequency component 0 and frequency component $N/2$, both of which have an imaginary component equal to zero.

Order	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**
Normal	0_r	8_r	1_r	1_c	2_r	2_c	3_r	3_c	4_r	4_c	5_r	5_c	6_r	6_c	7_r	7_c
Bit-reversed	0_r	8_r	4_r	4_c	2_r	2_c	6_r	6_c	1_r	1_c	5_r	5_c	3_r	3_c	7_r	7_c

Figure 4. Normal and bit-reversed order of a conjugate symmetric sequence allocated to a complex one-dimensional parallel variable.

In the real domain, the order is normal if the lower-ordered element in a pair of successive elements along the conjugate symmetric axis is assigned to the real component, and the higher-ordered element is assigned to the imaginary component, of a complex number. Furthermore, successive pairs must be assigned to successive complex locations along the axes to which the RCFFT is applied. The order is bit-reversed in the real domain if successive pairs of elements are assigned to complex locations along the conjugate symmetric axis based on the bit-reversed value of the index of the pair. Figure 5 illustrates the normal and bit-reversed ordering in the real domain.

Order	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**
Normal	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Bit-reversed	0	1	8	9	4	5	12	13	2	3	10	11	6	7	14	15

Figure 5. Normal and bit-reversed order of a real sequence allocated to a complex one-dimensional parallel variable.

The mapping of real data into a complex parallel variable of higher dimension depends upon which axis is to be transformed. Figure 6 shows an 8 x 8 real parallel variable. For an RCFFT along the vertical axis, you must map this parallel variable into a complex parallel variable as shown in Figure 7. For an RCFFT along the horizontal axis, you must map it into a complex parallel variable as shown in Figure 8. The **CMSSL_complex_from_real** routine, described in Section 4.4, performs these mappings.

Copyright © 1993 Thinking Machines Corporation

[0][0]x	[0][1]x	[0][2]x	[0][3]x	[0][4]x	[0][5]x	[0][6]x	[0][7]x
[1][0]x	[1][1]x	[1][2]x	[1][3]x	[1][4]x	[1][5]x	[1][6]x	[1][7]x
[2][0]x	[2][1]x	[2][2]x	[2][3]x	[2][4]x	[2][5]x	[2][6]x	[2][7]x
[3][0]x	[3][1]x	[3][2]x	[3][3]x	[3][4]x	[3][5]x	[3][6]x	[3][7]x
[4][0]x	[4][1]x	[4][2]x	[4][3]x	[4][4]x	[4][5]x	[4][6]x	[4][7]x
[5][0]x	[5][1]x	[5][2]x	[5][3]x	[5][4]x	[5][5]x	[5][6]x	[5][7]x
[6][0]x	[6][1]x	[6][2]x	[6][3]x	[6][4]x	[6][5]x	[6][6]x	[6][7]x
[7][0]x	[7][1]x	[7][2]x	[7][3]x	[7][4]x	[7][5]x	[7][6]x	[7][7]x

Figure 6. A real 8 x 8 parallel variable.

[0][0]w_r	[0][1]w_r	[0][2]w_r	[0][3]w_r	[0][4]w_r	[0][5]w_r	[0][6]w_r	[0][7]w_r
[0][0]w_i	[0][1]w_i	[0][2]w_i	[0][3]w_i	[0][4]w_i	[0][5]w_i	[0][6]w_i	[0][7]w_i
[1][0]w_r	[1][1]w_r	[1][2]w_r	[1][3]w_r	[1][4]w_r	[1][5]w_r	[1][6]w_r	[1][7]w_r
[1][0]w_i	[1][1]w_i	[1][2]w_i	[1][3]w_i	[1][4]w_i	[1][5]w_i	[1][6]w_i	[1][7]w_i
[2][0]w_r	[2][1]w_r	[2][2]w_r	[2][3]w_r	[2][4]w_r	[2][5]w_r	[2][6]w_r	[2][7]w_r
[2][0]w_i	[2][1]w_i	[2][2]w_i	[2][3]w_i	[2][4]w_i	[2][5]w_i	[2][6]w_i	[2][7]w_i
[3][0]w_r	[3][1]w_r	[3][2]w_r	[3][3]w_r	[3][4]w_r	[3][5]w_r	[3][6]w_r	[3][7]w_r
[3][0]w_i	[3][1]w_i	[3][2]w_i	[3][3]w_i	[3][4]w_i	[3][5]w_i	[3][6]w_i	[3][7]w_i

Figure 7. The mapping of an 8 x 8 real parallel variable into a 4 x 8 complex parallel variable for transformation on the vertical axis.

[0][0]w_r	[0][0]w_i	[0][1]w_r	[0][1]w_i	[0][2]w_r	[0][2]w_i	[0][3]w_r	[0][3]w_i
[1][0]w_r	[1][0]w_i	[1][1]w_r	[1][1]w_i	[1][2]w_r	[1][2]w_i	[1][3]w_r	[1][3]w_i
[2][0]w_r	[2][0]w_i	[2][1]w_r	[2][1]w_i	[2][2]w_r	[2][2]w_i	[2][3]w_r	[2][3]w_i
[3][0]w_r	[3][0]w_i	[3][1]w_r	[3][1]w_i	[3][2]w_r	[3][2]w_i	[3][3]w_r	[3][3]w_i
[4][0]w_r	[4][0]w_i	[4][1]w_r	[4][1]w_i	[4][2]w_r	[4][2]w_i	[4][3]w_r	[4][3]w_i
[5][0]w_r	[5][0]w_i	[5][1]w_r	[5][1]w_i	[5][2]w_r	[5][2]w_i	[5][3]w_r	[5][3]w_i
[6][0]w_r	[6][0]w_i	[6][1]w_r	[6][1]w_i	[6][2]w_r	[6][2]w_i	[6][3]w_r	[6][3]w_i
[7][0]w_r	[7][0]w_i	[7][1]w_r	[7][1]w_i	[7][2]w_r	[7][2]w_i	[7][3]w_r	[7][3]w_i

Figure 8. The mapping of an 8 x 8 real parallel variable into an 8 x 4 complex parallel variable for transformation on the horizontal axis.

Copyright © 1993 Thinking Machines Corporation

4.3.3 Multidimensional and Multiple-Instance Real-to-Complex and Complex-to-Real FFTs

For general information about multidimensional and multiple-instance FFTs, refer to Section 4.2.3.

The CMSSL RCFFT and CRFFT support multiple instances, but do not properly support multidimensional transforms. The discussion below explains how to use these routines to perform multidimensional transforms.

The fact that the real components of frequency components 0 and $N/2$ for the transformation of a real sequence of length N are stored in complex location 0 along the transformed axis implies that a ''complex'' parallel variable of conjugate symmetric sequences does not contain complex numbers in a strict sense. This may be important if a CCFFT is desired along other axes of a multidimensional parallel variable. For instance, if the RCFFT is performed along the vertical axis of Figure 7, then the result is returned in the form shown in Figure 9. The first row of complex elements contains two real sequences: one in the real component, one in the imaginary component. Applying a CCFFT to the horizontal axis yields the correct Discrete Fourier Transform of all sequences along the horizontal axis except for the two real sequences stored in complex row location zero. These two real sequences are transformed as a pair, as described in Section 4.3.4. The transform of each sequence can be obtained by applying the necessary postprocessing *only to row zero*. Such postprocessing is currently not supported in CMSSL.

Storage location								
[0][:]re	$[0][0]X_r$	$[0][1]X_r$	$[0][2]X_r$	$[0][3]X_r$	$[0][4]X_r$	$[0][5]X_r$	$[0][6]X_r$	$[0][7]X_r$
[0][:]im	$[4][0]X_r$	$[4][1]X_r$	$[4][2]X_r$	$[4][3]X_r$	$[4][4]X_r$	$[4][5]X_r$	$[4][6]X_r$	$[4][7]X_r$
[1][:]re	$[1][0]X_r$	$[1][1]X_r$	$[1][2]X_r$	$[1][3]X_r$	$[1][4]X_r$	$[1][5]X_r$	$[1][6]X_r$	$[1][7]X_r$
[1][:]im	$[1][0]X_i$	$[1][1]X_i$	$[1][2]X_i$	$[1][3]X_i$	$[1][4]X_i$	$[1][5]X_i$	$[1][6]X_i$	$[1][7]X_i$
[2][:]re	$[2][0]X_r$	$[2][1]X_r$	$[2][2]X_r$	$[2][3]X_r$	$[2][4]X_r$	$[2][5]X_r$	$[2][6]X_r$	$[2][7]X_r$
[2][:]im	$[2][0]X_i$	$[2][1]X_i$	$[2][2]X_i$	$[2][3]X_i$	$[2][4]X_i$	$[2][5]X_i$	$[2][6]X_i$	$[2][7]X_i$
[3][:]re	$[3][0]X_r$	$[3][1]X_r$	$[3][2]X_r$	$[3][3]X_r$	$[3][4]X_r$	$[3][5]X_r$	$[3][6]X_r$	$[3][7]X_r$
[3][:]im	$[3][0]X_i$	$[3][1]X_i$	$[3][2]X_i$	$[3][3]X_i$	$[3][4]X_i$	$[3][5]X_i$	$[3][6]X_i$	$[3][7]X_i$

Figure 9. The conjugate symmetric sequences in normal order, resulting from the transformation along the vertical axis of the 4 x 8 complex parallel variable of real numbers. The subscripts *r* and *i* indicate the real and imaginary components of the data elements.

Copyright © 1993 Thinking Machines Corporation

The postprocessing of row zero is different from the postprocessing in the CMSSL RCFFT. The postprocessing of row zero should uncover *two* conjugate symmetric sequences, and a reordering is required. Let the transform of X in Figure 9 along the horizontal axis be χ, that is,

$$X(km) = \sum_{n=0}^{N-1} x(nm) e^{-ikn(2\pi/N)}$$

and

$$\chi(kl) = \sum_{m=0}^{M-1} X(km) e^{-ilm(2\pi/M)}$$

Then the transformation on row zero, [0][:] X+ i*[4][:]X in Figure 10, by a CCFFT represents the transformation of a pair of real sequences embedded in a one-dimensional complex parallel variable. The postprocessing of a pair of real sequences embedded in a one-dimensional complex parallel variable requires postprocessing as described in Section 4.3.4. This postprocessing produces two conjugate symmetric sequences, each of half length. The elements of these two sequences can be stored as shown in Figure 10, in which the result is in normal order. The conjugate symmetric sequence resulting from the transformation of the real sequence [0][:]X is stored in the first half of complex row zero, and the conjugate symmetric sequence resulting from the transformation of real sequence [4][:]X is stored in the second half of complex row zero. Within each half, the storage is consistent with the storage for the RCFFT along the vertical axis. Whereas the convention is to store the two real components of the transformation of a real sequence — i.e., [0]X_r and [N/2]X_r — in complex location zero, there is no conventional storage scheme for the transform of a pair of real sequences. Since the postprocessing is required only for a parallel variable with rank one less than the original parallel variable, the current implementation of FFTs in the CMSSL leaves the postprocessing of multidimensional transforms to the user.

Copyright © 1993 Thinking Machines Corporation

Storage location								
[0][:]re	$[0][0]\chi_r$	$[0][1]\chi_r$	$[0][2]\chi_r$	$[0][3]\chi_r$	$[0][4]\chi_r$	$[0][5]\chi_r$	$[0][6]\chi_r$	$[0][7]\chi_r$
[0][:]im	$[0][0]\chi_i$	$[0][1]\chi_i$	$[0][2]\chi_i$	$[0][3]\chi_i$	$[0][4]\chi_r$	$[0][5]\chi_i$	$[0][6]\chi_i$	$[0][7]\chi_i$
[1][:]re	$[1][0]\chi_r$	$[1][1]\chi_r$	$[1][2]\chi_r$	$[1][3]\chi_r$	$[1][4]\chi_r$	$[1][5]\chi_r$	$[1][6]\chi_r$	$[1][7]\chi_r$
[1][:]im	$[1][0]\chi_i$	$[1][1]\chi_i$	$[1][2]\chi_i$	$[1][3]\chi_i$	$[1][4]\chi_i$	$[1][5]\chi_i$	$[1][6]\chi_i$	$[1][7]\chi_i$
[2][:]re	$[2][0]\chi_r$	$[2][1]\chi_r$	$[2][2]\chi_r$	$[2][3]\chi_r$	$[2][4]\chi_r$	$[2][5]\chi_r$	$[2][6]\chi_r$	$[2][7]\chi_r$
[2][:]im	$[2][0]\chi_i$	$[2][1]\chi_i$	$[2][2]\chi_i$	$[2][3]\chi_i$	$[2][4]\chi_i$	$[2][5]\chi_i$	$[2][6]\chi_i$	$[2][7]\chi_i$
[3][:]re	$[3][0]\chi_r$	$[3][1]\chi_r$	$[3][2]\chi_r$	$[3][3]\chi_r$	$[3][4]\chi_r$	$[3][5]\chi_r$	$[3][6]\chi_r$	$[3][7]\chi_r$
[3][:]im	$[3][0]\chi_i$	$[3][1]\chi_i$	$[3][2]\chi_i$	$[3][3]\chi_i$	$[3][4]\chi_i$	$[3][5]\chi_i$	$[3][6]\chi_i$	$[3][7]\chi_i$

Figure 10. One potential data allocation of complex and conjugate symmetric sequences resulting from a two-dimensional transform of an 8 x 8 real parallel variable initially mapped into a 4 x 8 complex parallel variable. The subscripts *r* and *i* indicate the real and imaginary components of the data elements.

Note that in many applications the values $[0][:]X$ and $[N/2][:]X$ may not be needed, and can be set to zero. Note further that even if the postprocessing is omitted, a real-to-complex multidimensional transform followed by a complex-to-real multidimensional inverse transform always reproduces the original result.

4.3.4 Implementation

The Fourier transform

$$X_k = \sum_{n=0}^{N-1} x_n e^{-ikn(2\pi/N)}$$

and its inverse

Copyright © 1993 Thinking Machines Corporation

$$x_n = 1/N \sum_{k=0}^{N-1} X_k e^{ikn(2\pi/N)}$$

are computed efficiently using the Fast Fourier Transform. But if x_n is a real sequence, $x_n = \overline{x_n}$, then X_k can be computed with half as many arithmetic operations as required for a complex sequence, and only half as much storage is needed. Many algorithms for real sequences are developed by modifying the Cooley-Tukey complex-to-complex FFT that follows from the splitting formula

$$X_k = \sum_{n=0}^{N/2-1} x_{2n} e^{-ik2n(2\pi/N)} + \sum_{n=0}^{N/2-1} x_{2n+1} e^{-ik(2n+1)(2\pi/N)}$$

Ignoring the scale factor $1/N$ and defining

$$Y_k = \sum_{n=0}^{N/2-1} x_{2n} e^{-ikn\frac{2\pi}{N/2}}$$

$$Z_k = \sum_{n=0}^{N/2-1} x_{2n+1} e^{-ikn\frac{2\pi}{N/2}}$$

we have

$$X_k = Y_k + e^{-ik(2\pi/N)} Z_k, \quad k = 0, \ldots, N/2 - 1$$

But since $Y_{k+N/2} = Y_k$, $Z_{k+N/2} = Z_k$, and $e^{-i(k+N/2)\,2\pi/N} = -e^{-ik\,2\pi/N}$, we have

$$X_{k+N/2} = Y_k - e^{-ik\,2\pi/N} Z_k, k = 0, \ldots, N/2 - 1.$$

This splitting algorithm for X_k is the basis for the Cooley-Tukey complex-to-complex FFT, which consists of repeated application of the splitting algorithm. In order to evaluate, say, only the first half, X_k ($k = 0, 1, \ldots, N/2$), for a conjugate symmetric sequence, it is necessary to modify the splitting formulas so that only the first half of Y_k and Z_k are referenced.

Copyright © 1993 Thinking Machines Corporation

Transformation of Pairs of Sequences

A *pair* of real sequences, y_n and z_n, can be transformed using a single complex transform by defining the complex sequence $x_n = y_n + iz_n$ and computing its complex transform X_k. Since the Fourier transform is linear, we have

$$X_k = Y_k + i\,Z_k$$

where Y_k and Z_k are the Fourier transforms of y_n and z_n, respectively. Since Y_k and Z_k are complex, they do not correspond to the real and imaginary parts of X_k. But since y_n and z_n are real, Y_k and Z_k are conjugate symmetric. Therefore

$$X_{N-k} = Y_{N-k} + i\,Z_{N-k} = \overline{Y_k} + i\,\overline{Z_k}$$

or

$$\overline{X_{N-k}} = Y_k - i\,Z_k$$

and

$$Y_k = 1/2\ (X_k + \overline{X_{N-k}}), \quad k = 0, 1, \ldots, N/2$$

$$Z_k = 1/2i\ (X_k - \overline{X_{N-k}}), \quad k = 0, 1, \ldots, N/2$$

Therefore, the transform of a pair of real sequences of length N can be computed from a single complex transform of length N: first, form $x_n = y_n + iz_n$; next, compute X_k; and finally, compute Y_k and Z_k. This technique is particularly useful for multidimensional transforms on real data, but is not used in CMSSL.

Note that since the Fourier Transform is a linear operation, if the same operation is applied to both transformed sequences, followed by an inverse transform, then explicit computation of Y_k and Z_k is not necessary. There is a large class of applications, including the solution of certain partial differential equations, to which this case applies. For example, an approximate derivative of the sequences y_n and z_n can be computed by multiplying X_k by ik and performing the inverse FFT.

Transformation of a Single Sequence

For the transformation of a *single* real sequence of length N by pre- and postprocessing around a complex transform of length $N/2$, let x_n be a real

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

sequence of length N, where N is an even integer. Define $y_n = x_{2n}$, $z_n = x_{2n+1}$, and $w_n = y_n + iz_n$. The sequence w_n is of length $N/2$. Next, compute the Fourier transform W_k of w_n. From the previous derivations, we have

$$Y_k = 1/2\ (W_k + \overline{W_{N/2-k}}), \quad k = 0, 1, \ldots, N/4$$

$$Z_k = 1/2i\ (W_k - \overline{W_{N/2-k}}), \quad k = 0, 1, \ldots, N/4$$

where Y_k and Z_k are the Discrete Fourier transforms of the sequences $y_n = x_{2n}$ and $z_n = x_{2n+1}$. We wish to compute X_k, $k = 0, 1, \ldots, N/2$, where

$$X_k = Y_k + e^{-ik(2\pi/N)}\ Z_k, \quad k = 0, \ldots, N/2 - 1$$

The sequence X_k for $k = 0, 1, \ldots, N/4$ can be computed from

$$X_k = Y_k + e^{-ik(2\pi/N)}\ Z_k, \quad k = 0, \ldots, N/4$$

For the indices $k = N/4 + 1, \ldots, N/2$, we have

$$X_{N/2-k} = Y_{N/2-k} + e^{-i(N/2-k)(2\pi/N)}\ Z_{N/2-k}, \quad k = 0, \ldots, N/4 - 1$$

But since the sequence x_n is real, both Y_k and Z_k are conjugate symmetric, and

$$\overline{X_{N/2-k}} = Y_k - e^{-ik(2\pi/N)}\ Z_k, \quad k = 0, \ldots, N/4 - 1$$

In summary, for the transformation of a single real sequence of length N, form the complex sequence $w_n = x_{2n} + ix_{2n+1}$ of length $N/2$, compute its transform using a complex-to-complex transform, and then obtain the transforms Y_k of x_{2n} and Z_k of x_{2n+1} through one postprocessing step, followed by another postprocessing step to obtain X_k. The two postprocessing steps can be combined into

$$X_k = 1/2\ [(W_k + \overline{W_{N/2-k}}) - ie^{-ik(2\pi/N)}\ (W_k - \overline{W_{N/2-k}})], \quad k = 0, 1, \ldots, N/4$$

$$\overline{X_{N/2-k}} = 1/2\ [(W_k + \overline{W_{N/2-k}}) + ie^{-ik(2\pi/N)}\ (W_k - \overline{W_{N/2-k}})], \quad k = 0, 1, \ldots, N/4-1$$

Copyright © 1993 Thinking Machines Corporation

Summary

The single real-sequence transform as outlined above is the algorithm used in the current CMSSL. The arithmetic requirements are the same as for the transform of pairs of sequences. The postprocessing requires data motion in both algorithms. The communication is incurred because of the *index* reversal:

$$k \longleftrightarrow N/2 - k$$

For data motion between processing elements, the communication time between processing elements for a CCFFT on $N/2$ points distributed over P processing elements is proportional to $N/2P$. The communication time for the postprocessing is also proportional to $N/2P$; i.e., the total communication time is proportional to N/P. Hence, in the current release, the communication time for the RCFFT on a real sequence of length N is approximately the same as the communication time for a CCFFT on a complex sequence of length N.

4.3.5 Steps in Performing the Real-to-Complex and Complex-to-Real FFTs

To perform an RCFFT, follow these steps:

1. Use the **CMSSL_complex_from_real** routine (described in Section 4.4) to place the real data into a complex parallel variable, with pairs of real data elements along the axis to be transformed (the conjugate symmetric axis) forming complex numbers. Note that the length of the conjugate symmetric axis must be a power of 2.

2. Call **CMSSL_fft_setup** and **CMSSL_fft_detailed** specifying an RCFFT. The **CMSSL_fft_detailed** routine computes the first half of the conjugate symmetric sequence in place. The real components of elements X_0 and $X_{N/2}$ are stored in location zero along the transformed axis. This routine

    a. Internally calls the CCFFT with the proper twiddle factors.

    b. Performs the postprocessing by

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

- Calling index-reversal routines if needed, performing the data communication in a close to optimal manner.
- Performing the butterfly-like computations for the postprocessing.

To perform a CRFFT, follow these steps:

1. Place the conjugate symmetric data into a complex parallel variable with the real components of elements X_0 and $X_{N/2}$ in the first location of the axis to be transformed (that is, the conjugate symmetric axis).
2. Call **CMSSL_fft_setup** and **CMSSL_fft_detailed** specifying a CRFFT. The **CMSSL_fft_detailed** routine computes in place the real sequence resulting from the Fourier Transform of the conjugate symmetric sequence. Pairs of successive real elements are allocated into the real and imaginary components of complex numbers along the axis representing the conjugate symmetric sequence. This routine
   a. Internally performs the preprocessing by
      - Calling index-reversal routines if needed, performing the data communication in a close to optimal manner.
      - Performing the butterfly-like computations for the preprocessing.
   b. Internally calls the CCFFT with the proper twiddle factors.
3. Call the **CMSSL_real_from_complex** routine to extract the resulting real parallel variable.

Copyright © 1993 Thinking Machines Corporation

Real-to-Complex and Complex-to-Real Fast Fourier Transform

The Fourier Transform routines described below use a real-to-complex Fast Fourier Transform (RCFFT) algorithm to compute the Discrete Fourier Transform of real data stored in a complex *n*-dimensional parallel variable. The real components of elements 0 and *N*/2 of each resulting conjugate symmetric sequence, together with components 1 through *N*/2 − 1, are stored in the complex parallel variable. The real components of elements 0 and *N*/2 are stored in location zero along the axis representing each conjugate symmetric sequence.

These routines also use a complex-to-real Fast Fourier Transform (CRFFT) to compute the Discrete Fourier Transform of an *n*-dimensional parallel variable containing conjugate symmetric sequences stored as described above.

SYNTAX

```
#include <cm/cmssl.h>

int  CMSSL_fft_setup (
     cz:void *A, CMSSL_fft_type_t type, int conj_sym_axis,
     int *setup_id);

int  CMSSL_fft_detailed (
     cz:void *A, CMSSL_fft_type_t type,
     CMSSL_fft_control_t *controls, int setup_id);

void CMSSL_deallocate_fft_setup (
     int setup_id);
```

OVERLOADED DATA TYPES

cz **CMSSL_complex_t** or **CMSSL_dcomplex_t**

The **CMSSL_complex_t** and **CMSSL_dcomplex_t** data types are defined in the header file **<cm/cmssl.h>**.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

ARGUMENTS

A Pointer to a complex parallel variable. The **A** you supply in the **CMSSL_fft_setup** call must have the same shape and type as the **A** arguments you supply in all subsequent **CMSSL_fft_detailed** calls associated with the setup call. All axes that are to be transformed must have power-of-2 lengths.

When you call **CMSSL_fft_setup**, you may supply arbitrary values in **A**; the setup routine neither examines nor modifies the contents of this parallel variable, but rather uses its shape and type to create the setup object.

When you call **CMSSL_fft_detailed**, **A** is the parallel variable to be transformed. Upon completion, **A** is overwritten with the result. The floating-point precision of the result always matches that of the input.

type Specifies the type of FFT to be performed, as defined below and in the header file **<cm/cmssl.h>**. Supply the value **CMSSL_real_to_cmpx** for a real-to-complex transform, or the value **CMSSL_cmpx_to_real** for a complex-to-real transform. (Note: in the setup call, these two types are interchangeable, as the setup action is the same for real-to-complex and complex-to-real transforms.)

```
typedef enum CMSSL_fft_type {
   CMSSL_cmpx_to_cmpx = 0,
   CMSSL_real_to_cmpx = 1,
   CMSSL_cmpx_to_real = 2
} CMSSL_fft_type_t;
```

conj_sym_axis Identifies the *conjugate symmetric axis*: the axis along which real data is aligned for a real-to-complex transform, or along which each conjugate symmetric sequence is aligned for a complex-to-real transform. The ordering of parallel variable dimensions is such that the left-most axis is axis 1. The length of the conjugate symmetric axis must be a power of 2.

setup_id Internal variable (pointer to an internal variable, when you call the setup routine). When you call **CMSSL_fft_detailed** or **CMSSL_deallocate_fft_setup**, you must supply the value returned by the associated **CMSSL_fft_setup** call.

Copyright © 1993 Thinking Machines Corporation

controls Pointer to an array of structures of type **CMSSL_fft_control_t** (see below). This array has one structure for each axis of **A**. Each element of **controls** specifies the direction of transformation (if any), input address ordering, and output address ordering for the corresponding axis of **A** in a Detailed FFT:

```
typedef struct CMSSL_fft_control {
  CMSSL_fft_op_t      op;
  CMSSL_fft_order_t   in_order;
  CMSSL_fft_order_t   out_order;
  CMSSL_fft_scale_t   scale;
} CMSSL_fft_control_t;
```

The **op** slot specifies the direction (if any) in which the corresponding **A** axis is to be transformed. Supply the value **CMSSL_f_xform** for a forward transform, **CMSSL_i_xform** for an inverse transform, or **CMSSL_nop** for no transform:

```
typedef enum CMSSL_fft_op {
  CMSSL_nop          = 0,
  CMSSL_f_xform      = 1,
  CMSSL_i_xform      = 2
} CMSSL_fft_op_t;
```

The **in_order** and **out_order** slots specify the initial address ordering and desired output address ordering (respectively) of the corresponding **A** axis:

```
typedef enum CMSSL_fft_order {
  CMSSL_normal        = 0,
  CMSSL_bit_reversed = 1
} CMSSL_fft_order_t;
```

The **scale** slot specifies the scaling factor (if any) for the corresponding **A** axis:

```
typedef enum CMSSL_fft_scale {
  CMSSL_noscale       = 0,
  CMSSL_scale_sqrt    = 1,
  CMSSL_scale_n       = 2
} CMSSL_fft_scale_t;
```

The value **CMSSL_scale_n** scales by the axis length; the value **CMSSL_scale_sqrt** scales by the square root of the axis length; and the value **CMSSL_noscale** specifies no scaling.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

Note that on the conjugate symmetric axis, the "axis length" is the declared length of the complex parallel variable axis, just as in the case of the CCFFT.

At the end of the **CMSSL_fft_detailed** call, each position of **A** is divided by the product of the scaling factors for all axes. (The axes for which you specified **CMSSL_noscale** do not contribute to this product.)

RETURNED VALUE

The **CMSSL_fft_setup** and **CMSSL_fft_detailed** routines return 0 upon successful return.

DESCRIPTION

The RCFFT computes the Fourier transform of real data, and the CRFFT transforms *conjugate symmetric* complex sequences. The transform X_k (for $k = 0, 1, \ldots, N-1$) of a real sequence x_n ($n = 0, 1, \ldots, N-1$), is conjugate symmetric, i.e., $X_k = \overline{X_{N-k}}$, and conversely, the transform of such a sequence is real. We refer to the domain in which the real data exists as the *real domain*, and the domain in which the conjugate symmetric sequence exists as the *complex domain*.

You can use the CMSSL RCFFT and CRFFT to perform a transform along any axis of a multidimensional parallel variable. The axis you select for the RCFFT or CRFFT is referred to as the *conjugate symmetric axis*, since after a real-to-complex transform along this axis, it represents half of a conjugate symmetric sequence. The CMSSL RCFFT and CRFFT compute the transform in place.

To perform a RCFFT or CRFFT, you must follow these steps:

1. Call **CMSSL_fft_setup**.

   This routine computes internal values (including twiddle factors), allocates an internal FFT setup object, and computes an FFT setup ID, which you must supply in all subsequent Detailed FFT and deallocation calls associated with this setup call.

2. Call **CMSSL_fft_detailed**.

   This routine executes the FFT. You may follow one call to **CMSSL_fft_setup** with multiple calls to **CMSSL_fft_detailed**. However, the result of a call to **CMSSL_fft_setup** can only be used for **CMSSL_fft_detailed** calls using parallel variables with the same shape and type as the parallel variable supplied in the setup routine. If these parameters change, you must call **CMSSL_fft_setup** again to create another setup structure.

Copyright © 1993 Thinking Machines Corporation

3. Call **CMSSL_deallocate_fft_setup**.

   This routine removes an FFT setup object that is no longer needed. Because an FFT setup object occupies both partition manager and processing element memory, you should free this space by deallocating the setup object after the completion of all FFT invocations that use it.

You may have more than one setup ID active at a time; that is, you may call the setup routine more than once before deallocating any setup IDs. When you call **CMSSL_fft_detailed** or **CMSSL_deallocate_fft_setup**, you must be sure to provide the correct setup ID.

Twiddle Factor Exponents. The sign of the twiddle factor exponents determines the direction of an FFT. In the CMSSL FFT, the *forward* transform is defined as one that uses twiddle factors with a *negative* exponent, and the *inverse* transform is defined as one that uses twiddle factors with a *positive* exponent.

Data Orderings. For the RCFFT and CRFFT, there are several important data ordering issues. In the complex domain, the ordering is normal when successive complex data elements along an axis are assigned to successive locations in the complex parallel variable; the the ordering is bit-reversed if the complex data elements are assigned to locations in the parallel variable based on the bit-reversed values of their indices, as described for the CCFFT. The figure below illustrates the data allocation in the complex domain for a one-dimensional parallel variable. For the axis representing the conjugate symmetric sequence, location zero is special. It represents two real elements, namely the real components of frequency component 0 and frequency component $N/2$, both of which have an imaginary component equal to zero.

Order	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**
Normal	0_r	8_r	1_r	1_c	2_r	2_c	3_r	3_c	4_r	4_c	5_r	5_c	6_r	6_c	7_r	7_c
Bit-reversed	0_r	8_r	4_r	4_c	2_r	2_c	6_r	6_c	1_r	1_c	5_r	5_c	3_r	3_c	7_r	7_c

Normal and bit-reversed order of a conjugate symmetric sequence allocated to a complex one-dimensional parallel variable.

In the real domain, the order is normal if the lower-ordered element in a pair of successive elements along the conjugate symmetric axis is assigned to the real component, and the higher-ordered element is assigned to the imaginary component, of a complex number. Furthermore, successive pairs must be assigned to successive complex locations along the axes to which the RCFFT is applied. The order is bit-reversed in the real domain if successive pairs of elements are assigned to complex locations along the conjugate symmetric axis based on the bit-reversed value

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

of the index of the pair. The next figure illustrates the normal and bit-reversed ordering in the real domain.

Order	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**	**re**	**im**
Normal	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Bit-reversed	0	1	8	9	4	5	12	13	2	3	10	11	6	7	14	15

Normal and bit-reversed order of a real sequence allocated to a complex one-dimensional parallel variable.

The mapping of real data into a complex parallel variable of higher dimension depends upon which axis is to be transformed. For example, consider the following 8 x 8 real parallel variable:

[0][0]*x*	[0][1]*x*	[0][2]*x*	[0][3]*x*	[0][4]*x*	[0][5]*x*	[0][6]*x*	[0][7]*x*
[1][0]*x*	[1][1]*x*	[1][2]*x*	[1][3]*x*	[1][4]*x*	[1][5]*x*	[1][6]*x*	[1][7]*x*
[2][0]*x*	[2][1]*x*	[2][2]*x*	[2][3]*x*	[2][4]*x*	[2][5]*x*	[2][6]*x*	[2][7]*x*
[3][0]*x*	[3][1]*x*	[3][2]*x*	[3][3]*x*	[3][4]*x*	[3][5]*x*	[3][6]*x*	[3][7]*x*
[4][0]*x*	[4][1]*x*	[4][2]*x*	[4][3]*x*	[4][4]*x*	[4][5]*x*	[4][6]*x*	[4][7]*x*
[5][0]*x*	[5][1]*x*	[5][2]*x*	[5][3]*x*	[5][4]*x*	[5][5]*x*	[5][6]*x*	[5][7]*x*
[6][0]*x*	[6][1]*x*	[6][2]*x*	[6][3]*x*	[6][4]*x*	[6][5]*x*	[6][6]*x*	[6][7]*x*
[7][0]*x*	[7][1]*x*	[7][2]*x*	[7][3]*x*	[7][4]*x*	[7][5]*x*	[7][6]*x*	[7][7]*x*

A real 8 x 8 parallel variable.

For an RCFFT along the vertical axis, you must map this parallel variable into a complex parallel variable as follows:

[0][0]w_r	[0][1]w_r	[0][2]w_r	[0][3]w_r	[0][4]w_r	[0][5]w_r	[0][6]w_r	[0][7]w_r
[0][0]w_i	[0][1]w_i	[0][2]w_i	[0][3]w_i	[0][4]w_i	[0][5]w_i	[0][6]w_i	[0][7]w_i
[1][0]w_r	[1][1]w_r	[1][2]w_r	[1][3]w_r	[1][4]w_r	[1][5]w_r	[1][6]w_r	[1][7]w_r
[1][0]w_i	[1][1]w_i	[1][2]w_i	[1][3]w_i	[1][4]w_i	[1][5]w_i	[1][6]w_i	[1][7]w_i
[2][0]w_r	[2][1]w_r	[2][2]w_r	[2][3]w_r	[2][4]w_r	[2][5]w_r	[2][6]w_r	[2][7]w_r
[2][0]w_i	[2][1]w_i	[2][2]w_i	[2][3]w_i	[2][4]w_i	[2][5]w_i	[2][6]w_i	[2][7]w_i
[3][0]w_r	[3][1]w_r	[3][2]w_r	[3][3]w_r	[3][4]w_r	[3][5]w_r	[3][6]w_r	[3][7]w_r
[3][0]w_i	[3][1]w_i	[3][2]w_i	[3][3]w_i	[3][4]w_i	[3][5]w_i	[3][6]w_i	[3][7]w_i

The mapping of an 8 x 8 real parallel variable into a 4 x 8 complex parallel variable for transformation on the vertical axis.

On the other hand, for an RCFFT along the horizontal axis, you must map the parallel variable into a complex parallel variable as follows:

Copyright © 1993 Thinking Machines Corporation

[0][0]w_r	[0][0]w_i	[0][1]w_r	[0][1]w_i	[0][2]w_r	[0][2]w_i	[0][3]w_r	[0][3]w_i
[1][0]w_r	[1][0]w_i	[1][1]w_r	[1][1]w_i	[1][2]w_r	[1][2]w_i	[1][3]w_r	[1][3]w_i
[2][0]w_r	[2][0]w_i	[2][1]w_r	[2][1]w_i	[2][2]w_r	[2][2]w_i	[2][3]w_r	[2][3]w_i
[3][0]w_r	[3][0]w_i	[3][1]w_r	[3][1]w_i	[3][2]w_r	[3][2]w_i	[3][3]w_r	[3][3]w_i
[4][0]w_r	[4][0]w_i	[4][1]w_r	[4][1]w_i	[4][2]w_r	[4][2]w_i	[4][3]w_r	[4][3]w_i
[5][0]w_r	[5][0]w_i	[5][1]w_r	[5][1]w_i	[5][2]w_r	[5][2]w_i	[5][3]w_r	[5][3]w_i
[6][0]w_r	[6][0]w_i	[6][1]w_r	[6][1]w_i	[6][2]w_r	[6][2]w_i	[6][3]w_r	[6][3]w_i
[7][0]w_r	[7][0]w_i	[7][1]w_r	[7][1]w_i	[7][2]w_r	[7][2]w_i	[7][3]w_r	[7][3]w_i

The mapping of an 8 x 8 real parallel variable into an 8 x 4 complex parallel variable for transformation on the horizontal axis.

The **CMSSL_complex_from_real** routine performs these mappings.

Multidimensional and Multiple-Instance FFTs. An RCFFT can only be performed on the conjugate symmetric axis, and must be performed *prior* to any CCFFTs you may perform on the parallel variable. Moreover, if you perform a CCFFT on the other axes of the parallel variable following an RCFFT, the result is a Fourier transform for all components of the parallel variable *except* the portion of rank one less than the parallel variable, formed by excluding the conjugate symmetric axis and defined by location 0 along the conjugate symmetric axis.

Similarly, if the parallel variable is formatted for a CRFFT, and you perform a CCFFT prior to a CRFFT, the result is not a Fourier transform for all components of the parallel variable.

In an RCFFT, axis **conj_sym_axis** is transformed first. The CRFFT operates in the reverse order, and therefore transforms axis **conj_sym_axis** last.

Scalar Control Array for Detailed FFT. The **CMSSL_fft_detailed** call requires a pointer to an array of structures, **CMSSL_fft_control**, of type **CMSSL_fft_control_t** (one structure for each axis of **A**). Each element of **controls** specifies the direction of transformation (if any), input address ordering, and output address ordering for the corresponding axis of **A** in a Detailed FFT: The slots of **CMSSL_fft_control** are as follows:

- The **op** slot specifies the direction in which the corresponding **A** axis is to be transformed, or specifies that the axis is not to be transformed. In a multiple-instance FFT, the axes to be transformed are the data axes, and the axes that are not to be transformed are the instance axes.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

- The **in_order** and **out_order** slots specify the order in which the original and resulting **A** values are stored along each axis. In normal address ordering (**CMSSL_normal**), data values are stored consecutively. Consecutive order is the only compiler-supported data allocation scheme. In bit-reversed ordering, consecutive data values are stored according to sequential addresses whose high- and low-order bits have been swapped; that is, the consecutive allocation scheme is applied to the bit-reversed indices. Note that bit reversal is always done with respect to the *complex* parallel variable, as described above.

  Although **CMSSL_fft_detailed** can handle any combination of bit orderings, performance is best when the axes of an untransformed parallel variable and a transformed parallel variable have opposite bit-orderings — that is, input normal and output bit-reversed, or input bit-reversed and output normal.

- The **scale** slot specifies a scaling factor (none, square root of axis length, or axis length) for each axis. (Note that on the conjugate symmetric axis, the "axis length" is the declared length of the complex parallel variable axis, just as in the case of the CCFFT.) At the end of the **CMSSL_fft_detailed** call, each position of **A** is divided by the product of the scaling factors. (The axes for which you specified **CMSSL_noscale** do not contribute to this product.)

  For example, suppose **A** has dimensions [*m*][*n*][*p*] and you supply **CMSSL_scale_n** for the first axis, **CMSSL_scale_sqrt** for the second, and **CMSSL_noscale** for the third. At the end of the **CMSSL_fft_detailed** call, each position of **A** is divided by *m* * sqrt(*n*).

  One artifact of the Fourier Transform algorithm is that if no scaling is performed, then a forward FFT followed immediately by an inverse FFT results in the original parallel variable with all values multiplied by the FFT size — that is, the product of the lengths of the transformed axes. Scaling is typically used to correct this effect and prevent arithmetic overflow. Typically, a scaling factor of **CMSSL_noscale** is used for all axes in a forward transform, and a scaling factor of **CMSSL_scale_n** is used for all axes being transformed in an inverse transform. Alternatively, you may specify a scaling factor of **CMSSL_scale_sqrt** for all transformed axes in both the forward and the inverse FFTs. This alternative is less efficient, but can reduce the risk of arithmetic overflow for very large numbers.

  Scaling is independent of the operations performed along the axes. For example, it is possible to specify a scaling factor of **CMSSL_scale_n** for an axis that is not being transformed.

Copyright © 1993 Thinking Machines Corporation

Bit-Reversing the Addressing of a Parallel Variable. You can use the Detailed FFT to bit-reverse the addressing of any *n*-dimensional parallel variable. To do this, specify **CMSSL_nop** for the **op** slot in all elements of **controls**, and specify opposite values for the input and output orderings; that is, either **CMSSL_normal** for the **in_order** values and **CMSSL_bit_reversed** for the **out_order** values, or **CMSSL_bit_reversed** for the **in_order** values and **CMSSL_normal** for the **out_order** values.

NOTES

Setup Uses Shape and Type. The setup ID computed by the **CMSSL_fft_setup** call can be used for all parallel variables with the same shape and type as the **A** argument supplied in the **CMSSL_fft_setup** call — and *only* for such parallel variables. If a transform is to be performed on two parallel variables, **A** and **B**, of the same shape and type, then *one* call to the setup routine suffices, even if transforms are performed on different axes of the two parallel variables. But if **A** and **B** have different shapes or types, then you must make one call for each parallel variable.

Setup is Private. The FFT setup data contains information private to the FFT. The format of this data is not documented and may change between CMSSL releases. For this reason, application code should never access or modify the contents of an FFT setup structure.

Bit Reversal and Performance. For an axis that is being transformed and that has a non-local component,

- If you specify *opposite* address orderings for input and output, the Detailed FFT performs the single bit-reversal that is inherent in the Cooley-Tukey algorithm, and produces output that is bit-reversed relative to the input.
- If you specify the *same* address ordering for input and output, the Detailed FFT performs an *extra* bit reversal, and produces output with the same ordering as the input. The extra bit reversal exacts a performance cost.

For an axis that is being transformed and that is local to a processing element,

- If you specify normal input and bit-reversed output on the conjugate symmetric axis for an RCFFT, *or* bit-reversed input and normal output on the conjugate symmetric axis for a CRFFT, the Detailed FFT performs *two extra* bit reversals.
- All other cases behave as in the non-local case, above.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

For an axis that is not being transformed,

- If you specify *opposite* address orderings for input and output, the Detailed FFT performs a single bit-reversal and produces output that is bit-reversed relative to the input.
- If you specify the *same* address ordering for input and output, the Detailed FFT performs *no* bit reversal.

To bit-reverse the address ordering of a dataset without transforming it, call the Detailed FFT specifying no transform for all axes and specifying *opposite* values for the input and output address ordering.

These points are independent of the directions of the transforms.

Axis Transformation Order. In an RCFFT, axis **`conj_sym_axis`** is transformed first. The CRFFT operates in the reverse order, and therefore transforms axis **`conj_sym_axis`** last. See the section called "Multidimensional and Multiple-Instance FFTs" in the description above for cautions related to multidimensional FFTs.

FFT Inversion. To invert an RCFFT, you must change the operation from forward to inverse as well as changing the FFT type from real-to-complex to complex-to-real. The **`type`** values **`CMSSL_real_to_cmpx`** and **`CMSSL_cmpx_to_real`** select the algorithm (FFT type); the **`op`** slot of the **`CMSSL_fft_control`** structure chooses the sign of the exponents of the twiddles, and thus the direction of the transform.

EXAMPLES

Sample C* code that uses the RCFFT and CRFFT can be found on-line in the subdirectory

```
fft/cstar/
```

of a CMSSL examples directory whose location is site-specific.

Copyright © 1993 Thinking Machines Corporation

4.4 Type Conversion Utilities for the Real-to-Complex and Complex-to-Real FFTs

The CMSSL offers two type conversion utilities for use with the FFTs:

- The **CMSSL_complex_from_real** routine converts a real parallel variable to a complex parallel variable suitable for input to the RCFFT.
- The **CMSSL_real_from_complex** routine converts a complex parallel variable to a real parallel variable. The conversion assumes that the elements of the complex parallel variable are ordered as in a parallel variable transformed by the CRFFT.

These routines are intended primarily for use with the RCFFT and CRFFT, but can be used with any real and complex parallel variables whose shapes satisfy the necessary conditions. Details are provided in the man page that follows.

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

Type Conversion Utilities for the Real-to-Complex and Complex-to-Real FFTs

The **CMSSL_complex_from_real** routine converts a real parallel variable to a complex parallel variable suitable for input to the real-to-complex Fast Fourier Transform (RCFFT). The **CMSSL_real_from_complex** routine converts a complex parallel variable, whose elements are ordered as in a parallel variable transformed by the complex-to-real FFT (CRFFT), to a real parallel variable.

SYNTAX

```
#include <cm/cmssl.h>

void CMSSL_complex_from_real (
     cz:void *A, fdfd:void *B, int conj_sym_axis);

void CMSSL_real_from_complex (
     fdfd:void *B, cz:void *A, int conj_sym_axis);
```

OVERLOADED DATA TYPES

cz **CMSSL_complex_t** or **CMSSL_dcomplex_t**
fdfd **float** or **double**; same precision as **A**.

The **CMSSL_complex_t** and **CMSSL_dcomplex_t** data types are defined in the header file **<cm/cmssl.h>**.

ARGUMENTS

A Pointer to a parallel variable with the same rank as **B**. Corresponding axes of **A** and **B** must have equal lengths, with the exception of the axis identified by the **conj_sym_axis** argument. Axis **conj_sym_axis** of **A** must have half the length of axis **conj_sym_axis** of **B**.

When you call **CMSSL_complex_from_real**, **A** is the complex parallel variable resulting from the conversion, and is an output argument.

When you call **CMSSL_real_from_complex**, **A** is the complex parallel variable you wish to convert to a real parallel variable, and is an input argument.

Copyright © 1993 Thinking Machines Corporation

B Pointer to a parallel variable with the same rank as **A**. Corresponding axes of **A** and **B** must have equal lengths, with the exception of the axis identified by the **conj_sym_axis** argument. Axis **conj_sym_axis** of **B** must have twice the length of axis **conj_sym_axis** of **A**.

When you call **CMSSL_complex_from_real**, **B** is the real parallel variable you wish to convert to a complex parallel variable, and is an input argument.

When you call **CMSSL_real_from_complex**, **B** is the real parallel variable resulting from the conversion, and is an output argument.

conj_sym_axis The axis that determines the mapping of real to complex or complex to real data elements, as explained in the Description below. When **CMSSL_complex_from_real** and **CMSSL_real_from_complex** are used in conjunction with the RCFFT and CRFFT, **conj_sym_axis** is the axis along which the transform is performed (the conjugate symmetric axis).

DESCRIPTION

Real-to-Complex Conversion. The **CMSSL_complex_from_real** routine converts **B** to a complex parallel variable and places the results in **A**. Each element of **B** is mapped to the real or imaginary part of an element of **A**. (This storage scheme is more efficient than a scheme that maps real elements exclusively to the real parts of a complex parallel variable.) If your parallel variables are one-dimensional, and if **A** has length n and **B** has length $2n$, then for each i between 0 and $n-1$, inclusive, **CMSSL_complex_from_real** sets the real part of $[i]$**A** equal to `[2*i]`**B** and the imaginary part of $[i]$**A** equal to `[2*i+1]`**B**.

If your parallel variables have rank greater than 1, the mapping of real to complex data elements depends on the axis you specify in the **conj_sym_axis** argument; the mapping defined above occurs along axis **conj_sym_axis**. If you define

$[n_1][n_2] \ldots [2n_{\texttt{conj_sym_axis}}] \ldots [n_p]$**B**
$[n_1][n_2] \ldots [n_{\texttt{conj_sym_axis}}] \ldots [n_p]$**A**

where $2n_{\texttt{conj_sym_axis}}$ is the length of axis **conj_sym_axis** of **B** and $n_{\texttt{conj_sym_axis}}$ is the length of axis **conj_sym_axis** of **A**, then for

Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

$i = 0$ to $n_1 - 1$,
$j = 0$ to $n_2 - 1$,
...,
$m = 0$ to $n_{\texttt{conj_sym_axis}} - 1$,
...,
$t = 0$ to $n_p - 1$

CMSSL_complex_from_real sets the real part of [*i*][*j*]...[*m*]...[*t*]**A** to [*i*][*j*]...[2*m*]...[*t*]**B** and the imaginary part of [*i*][*j*]...[*m*]...[*t*]**A** to [*i*][*j*]...[2*m*+1]...[*t*]**B**.

For example, if **B** is

[0][0]*B*	[0][1]*B*	[0][2]*B*	[0][3]*B*	[0][4]*B*	[0][5]*B*	[0][6]*B*	[0][7]*B*
[1][0]*B*	[1][1]*B*	[1][2]*B*	[1][3]*B*	[1][4]*B*	[1][5]*B*	[1][6]*B*	[1][7]*B*
[2][0]*B*	[2][1]*B*	[2][2]*B*	[2][3]*B*	[2][4]*B*	[2][5]*B*	[2][6]*B*	[2][7]*B*
[3][0]*B*	[3][1]*B*	[3][2]*B*	[3][3]*B*	[3][4]*B*	[3][5]*B*	[3][6]*B*	[3][7]*B*

and **conj_sym_axis** = 1, then each element is mapped to the corresponding location in the following depiction of **A**:

[0][0]A_r	[0][0]A_i	[0][1]A_r	[0][1]A_i	[0][2]A_r	[0][2]A_i	[0][3]A_r	[0][3]A_i
[1][0]A_r	[1][0]A_i	[1][1]A_r	[1][1]A_i	[1][2]A_r	[1][2]A_i	[1][3]A_r	[1][3]A_i
[2][0]A_r	[2][0]A_i	[2][1]A_r	[2][1]A_i	[2][2]A_r	[2][2]A_i	[2][3]A_r	[2][3]A_i
[3][0]A_r	[3][0]A_i	[3][1]A_r	[3][1]A_i	[3][2]A_r	[3][2]A_i	[3][3]A_r	[3][3]A_i

In this depiction, *r* and *i* indicate the real and imaginary parts of a data element, respectively. Thus, elements [0][0]**B** and [0][1]**B** are mapped to the real and imaginary parts, respectively, of element [0][0]**A**.

Complex-to-Real Conversion. The **CMSSL_real_from_complex** routine converts **A** to a real parallel variable and places the results in **B**. Each real or imaginary part of an element of **A** is mapped to an element of **B**. If your parallel variables are one-dimensional, and if **A** has length *n* and **B** has length 2*n*, then **CMSSL_real_from_complex** performs the equivalent of the following C* statements for $i = 0$ to $n-1$, inclusive:

```
[2*i]B = re([i]A);
[2*i+1]B = im([i]A);
```

If your parallel variables have rank greater than 1, the mapping of complex to real data elements depends on the axis you specify in the **conj_sym_axis** argument; the mapping defined above occurs along axis **conj_sym_axis**. If you define

Copyright © 1993 Thinking Machines Corporation

$[n_1][n_2]$ ... $[2n_{\texttt{conj_sym_axis}}]$ ... $[n_p]$**B**
$[n_1][n_2]$ ... $[n_{\texttt{conj_sym_axis}}]$ ... $[n_p]$**A**

where $2n_{\texttt{conj_sym_axis}}$ is the length of axis **conj_sym_axis** of **B** and $n_{\texttt{conj_sym_axis}}$ is the length of axis **conj_sym_axis** of **A**, then for

$i = 0$ to $n_1 - 1$,
$j = 0$ to $n_2 - 1$,
...,
$m = 0$ to $n_{\texttt{conj_sym_axis}} - 1$,
...,
$t = 0$ to $n_p - 1$

CMSSL_complex_from_real sets $[i][j]...[2m]...[t]$**B** to the real part of $[i][j]...[m]...[t]$**A** and $[i][j]...[2m+1]...[t]$**B** to the imaginary part of $[i][j]...[m]...[t]$**A**. This mapping is the inverse of that performed by **CMSSL_complex_from_real**.

Use with FFT. The complex parallel variable produced by the **CMSSL_complex_from_real** routine is set up correctly for input to the RCFFT. The axis identified by the **conj_sym_axis** argument is the axis to be transformed (the *conjugate symmetric axis*). Similarly, the conversion performed by the **CMSSL_real_from_complex** routine assumes that the data elements of the input argument **A** are ordered as in an parallel variable transformed by the CRFFT. A typical program structure is as follows:

```
type = CMSSL_real_to_cmpx;
CMSSL_fft_setup(&A, type, conj_sym_axis, &setup_id);
B = ...
                          .
                          .
                          .
CMSSL_complex_from_real (&A, &B, conj_sym_axis);
CMSSL_fft_detailed (&A, type, conj_sym_axis, setup_id);
                          .
                          .
                          .
type = CMSSL_cmpx_to_real
CMSSL_fft_detailed (&A, type, conj_sym_axis,
                    setup_id);
                          .
                          .
CMSSL_real_from_complex (&B, &A, conj_sym_axis);
CMSSL_deallocate_fft_setup(setup_id);
```


Version 3.2 Beta, December 1993
Copyright © 1993 Thinking Machines Corporation

EXAMPLES

Sample C* code that uses the type conversion utilities can be found on-line in the subdirectory

```
fft/cstar/
```

of a CMSSL examples directory whose location is site-specific.

Copyright © 1993 Thinking Machines Corporation

4.5 References

For further information about the CCFFT, see the following references:

1. Buzbee, B. L., G. H. Golub, and C. W. Nielson. On Direct Methods for Solving Poisson's Equations. *SIAM J. Numer. Anal.* **7**, no. 4 (1970): 627–56.

2. Chu, C. *The Fast Fourier Transform on Hypercube Parallel Computers*. Ph.D. Thesis, Center for Applied Mathematics, Cornell University, Ithaca NY 14585.

3. Cooley, J. C., P. Lewis, and P. D. Welch. The Fast Fourier Transform Algorithm: Programming Considerations in the Calculation of the Sine, Cosine, and Laplace Transforms. *J. Sound Vibrations* **12**, no. 3 (1970): 315–37.

4. Cooley, J. C., and J. W. Tukey. An Algorithm for the Machine Computation of Complex Fourier Series. *Math. Comp.* **19** (1965): 291–301.

5. Edelman, A. *Optimal Matrix Transposition and Bit-Reversal on Hypercubes: Node Address–Memory Address Exchanges.* Thinking Machines Corporation Technical Report, 1989.

6. Hockney, R. W. The Potential Calculation and Some Applications. *Methods Comput. Phys.* **9** (1970): 135–211.

7. Hong, J. W., and H. T. Kung. I/O Complexity: The Red-Blue Pebble Game. In *Proc. of the 13th ACM Symposium on the Theory of Computation*. ACM, 1981. Pp. 326–33.

8. Johnsson, S. L. Communication Efficient Basic Linear Algebra Computations on Hypercube Architectures. *J. Parallel Distributed Comput.***4**, no. 2 (1987): 133–72.

9. Johnsson, S. L. and C-T. Ho. Matrix Transposition on Boolean n-Cube Configured Ensemble Architectures. *SIAM J. Matrix Anal. Appl.* **9**, no. 3 (1988): 419–54.

10. Johnsson, S. L. and C-T. Ho. *Optimal Communication Channel Utilitzation for Matrix Transposition and Related Permutations on Boolean Cubes.* Technical Report TR–03–91, Harvard University, Division of Applied Sciences, January 1991.

Copyright © 1993 Thinking Machines Corporation

11. Johnsson, S. L., C-T. Ho, M. Jacquemin, and A. Ruttenberg. Computing Fast Fourier Transforms on Boolean Cubes and Related Networks. In *Advanced Algorithms and Architectures for Signal Processing II,* **826**. Society of Photo-Optical Instrumentation Engineers, 1987. Pp. 223–31.

12. Johnsson, S. L., Jacquemin, M., and C-T. Ho. *High Radix FFT on Boolean Cube Networks*. Harvard University, Division of Applied Sciences, Technical Report TR–25–91. To appear in the *Journal of Computational Physics*.

13. Johnsson, S. L. and R. L. Krawitz. Cooley-Tukey FFT on the Connection Machine. *Parallel Computing* **18**, no. 11 (1992): 1201–21.

14. Johnsson, S. L., M. Jacquemin, and R. L. Krawitz. Communication Efficient Multi-Processor FFT. *J. Comp. Phys* **102**, no. 2 (1992): 381–97.

15. Johnsson, S. L and C.–T. Ho. Boolean Cube Emulation of Butterfly Networks Encoded by Gray Code. Yale University Department of Computer Science, Technical Report YALEU/DCS/RR–764, 1990. Also Thinking Machines Corporation Technical Report BA90–1, TMC–5. To appear in *Journal of Parallel and Distributed Computing*.

16. Swarztrauber, P. N. The Methods of Cyclic Reduction, Fourier Analysis, and the FACR Algorithm for the Discrete Solution of Poisson's Equation on a Rectangle. *SIAM Review* **19** (1977): 490–501.

17. Swarztrauber, P. N. Multiprocessor FFTs. *Parallel Computing* **5** (1987): 197–210.

18. Temperton, C. On the FACR(1) Algorithm for the Discrete Poisson Equation. *J. of Computational Physics* **34** (1980): 314–29.

19. Van Loan, C. *Computational Frameworks for the Fast Fourier Transform*. SIAM, 1992.

In addition, for information about the CRFFT and RCFFT, see the following references:

20. Bergland, G.D. A fast Fourier transform for real-valued series. *Comm. ACM* **11** (1968): 703–10.

21. Briggs, W.L. Further symmetries of in-place FFTs. *SIAM J. Sci. Stat. Comput*. **8** (1987): 644–54.

Copyright © 1993 Thinking Machines Corporation

22. Pelz, R.B. Parallel compact FFTs for real sequences. Submitted to *SIAM J. Sci. Stat. Comput.*

23. Swarztrauber, P.N. Symmetric FFTs. *Math. Comp*. **47** (1986): 323−46.

24. Swarztrauber, P.N. and S. L. Johnsson. Multiprocessor FFTs for real sequences. Technical Report, National Center for Atmospheric Research, August 1991.

25. Sweet, R.A., V.E. Henson, and J. Porsche. A fast Fourier transform for real data on the hypercube. In preparation.

Copyright © 1993 Thinking Machines Corporation