VisualDSP++™ 3.5
C/C++ Compiler and Library
Manual for ADSP-219x Processors

Revision 4.1, October 2003

Part Number:
82-000390-03

Analog Devices, Inc.
One Technology Way
Norwood, Mass. 02062-9106

Copyright Information

Printed in the USA.

Disclaimer

Analog Devices, Inc. reserves the right to change this product without prior notice. Information furnished by Analog Devices is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices for its use; nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under the patent rights of Analog Devices, Inc.

Trademark and Service Mark Notice

CONTENTS

PREFACE

ACHIEVING OPTIMAL PERFORMANCE FROM C/C++ SOURCE CODE

C/C++ RUN-TIME LIBRARY

DSP RUN-TIME LIBRARY

COMPILER LEGACY SUPPORT

INDEX

PREFACE

Thank you for purchasing Analog Devices development software for digital signal processors (DSPs).

Purpose

The *VisualDSP++ 3.5 C/C++ Compiler and Library Manual for ADSP-219x DSPs* contains information about the C/C++ compiler and run-time library program for ADSP-219x DSPs. It includes syntax for command lines, switches, and language extensions. It leads you through the process of using library routines and writing mixed C/C++/assembly code.

Intended Audience

The primary audience for this manual is DSP programmers who are familiar with Analog Devices DSPs. This manual assumes that the audience has a working knowledge of the ADSP-219x DSPs architecture and instruction set and the C/C++ programming languages.

Programmers who are unfamiliar with ADSP-219x DSPs can use this manual, but they should supplement it with other texts (such as the appropriate hardware reference and instruction set reference) that provide information about your ADSP-219x DSP architecture and instructions).

Manual Contents Description

This manual contains:

- Chapter 1, "Compiler"

  Provides information on compiler options, language extensions and C/C++/assembly interfacing

- Chapter 2, "Achieving Optimal Performance from C/C++ Source Code"
  Shows how to optimize compiler operation.

- Chapter 3, "C/C++ Run-Time Library"

  Shows how to use library functions and provides a complete C/C++ library function reference

- Chapter 4, "DSP Run-Time Library"

  Shows how to use DSP library functions and provides a complete DSP library function reference

- Appendix A, "Compiler Legacy Support"

  Describes support for legacy code that was developed with previous releases of the development tools.

What's New in this Manual

This edition of the *VisualDSP++ 3.5 C/C++Compiler and Library Manual for ADSP-219x DSPs* documents support for all ADSP-219x processors. Refer to *VisualDSP++ 3.5 Product Bulletin for 16-Bit Processors* for information on all new and updated features and other release information.

Technical or Customer Support

You can reach DSP Tools Support in the following ways:

- Visit the DSP Development Tools website at
  `www.analog.com/technology/dsp/developmentTools/index.html`
- Email questions to
  `dsptools.support@analog.com`
- Phone questions to **1-800-ANALOGD**
- Contact your ADI local sales office or authorized distributor
- Send questions by mail to:

```
Analog Devices, Inc.
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106|
USA
```

Supported Processors

The name "ADSP-219x" refers to a family of Analog Devices 16-bit, fixed-point processors. VisualDSP++ currently supports the following ADSP-219x processors:

ADSP-2191 DSP, ADSP-2192-12 DSP, ADSP-2195 DSP,

ADSP-2196 DSP, ADSP-21990 DSP, ADSP-21991 DSP, and

ADSP-21992 DSP

Product Information

You can obtain product information from the Analog Devices website, from the product CD-ROM, or from the printed publications (manuals).

Analog Devices is online at `www.analog.com`. Our website provides information about a broad range of products—analog integrated circuits, amplifiers, converters, and digital signal processors.

MyAnalog.com

MyAnalog.com is a free feature of the Analog Devices website that allows customization of a webpage to display only the latest information on products you are interested in. You can also choose to receive weekly email notification containing updates to the webpages that meet your interests. MyAnalog.com provides access to books, application notes, data sheets, code examples, and more.

Registration:

Visit `www.myanalog.com` to sign up. Click **Register** to use MyAnalog.com. Registration takes about five minutes and serves as means for you to select the information you want to receive.

If you are already a registered user, just log on. Your user name is your email address.

DSP Product Information

For information on digital signal processors, visit our website at `www.analog.com/dsp`, which provides access to technical publications, datasheets, application notes, product overviews, and product announcements.

You may also obtain additional information about Analog Devices and its products in any of the following ways.

- Email questions or requests for information to
  `dsp.support@analog.com`
- Fax questions or requests for information to
  1-781-461-3010 (North America)
  089/76 903-557 (Europe)
- Access the Digital Signal Processing Division's FTP website at
  ftp ftp.analog.com or **ftp 137.71.23.21**
  `ftp://ftp.analog.com`

Related Documents

For information on product related development software, see the following publications:

VisualDSP++ 3.5 Getting Started Guide for 16-Bit Processors

VisualDSP++ 3.5 User's Guide for 16-Bit Processors

VisualDSP++ 3.5 Assembler and Preprocessor Manual for ADSP-218x and ADSP-219x DSPs

VisualDSP++ 3.5 Linker and Utilities Manual for 16-Bit Processors

VisualDSP++ 3.5 Loader Manual for 16-Bit Processors

VisualDSP++ 3.5 Product Bulletin for 16-Bit Processors

VisualDSP++ Kernel (VDK) User's Guide

VisualDSP++ Component Software Engineering User's Guide

Quick Installation Reference Card

For hardware information, refer to the processor's *Hardware Reference Manual* and *Instruction Set Reference Manual*.

Online Technical Documentation

Online documentation comprises VisualDSP++ Help system and tools manuals, Dinkum Abridged C++ library and FlexLM network license manager software documentation. You can easily search across the entire VisualDSP++ documentation set for any topic of interest. For easy printing, supplementary .PDF files for the tools manuals are also provided.

A description of each documentation file type is as follows.

File	Description
.CHM	Help system files and VisualDSP++ tools manuals.
.HTM or .HTML	Dinkum Abridged C++ library and FlexLM network license manager software documentation. Viewing and printing the .HTML files require a browser, such as Internet Explorer 4.0 (or higher).
.PDF	VisualDSP++ tools manuals in Portable Documentation Format, one .PDF file for each manual. Viewing and printing the .PDF files require a PDF reader, such as Adobe Acrobat Reader 4.0 (or higher).

If documentation is not installed on your system as part of the software installation, you can add it from the VisualDSP++ CD-ROM at any time.

Access the online documentation from the VisualDSP++ environment, Windows Explorer, or Analog Devices website.

From VisualDSP++

Access VisualDSP++ online Help from the Help menu's **Contents**, **Search**, and **Index** commands.

Open online Help from context-sensitive user interface items (toolbar buttons, menu commands, and windows).

From Windows

In addition to any shortcuts you may have constructed, there are many ways to open VisualDSP++ online Help or the supplementary documentation from Windows.

Help system files (`.CHM` files) are located in the Help folder, and `.PDF` files are located in the **Docs** folder of your VisualDSP++ installation. The **Docs** folder also contains the Dinkum Abridged C++ library and FlexLM network license manager software documentation.

Using Windows Explorer

- Double-click any file that is part of the VisualDSP++ documentation set.
- Double-click the `vdsp-help.chm` file, which is the master Help system, to access all the other `.CHM` files.

Using the Windows Start Button

- Access the VisualDSP++ online Help by clicking the **Start** button and choosing **Programs**, **Analog Devices**, **VisualDSP++**, and **VisualDSP++ Documentation**.
- Access the `.PDF` files by clicking the **Start** button and choosing **Programs**, **Analog Devices**, **VisualDSP++**, **Documentation for Printing**, and the name of the book.

From the Web

To download the tools manuals, point your browser at
`http://www.analog.com/technology/dsp/developmentTools/gen_purpose.html`

Select a processor family and book title. Download archive (`.ZIP`) files, one for each manual. Use any archive management software, such as WinZip, to decompress downloaded files.

Printed Manuals

For general questions regarding literature ordering, call the Literature Center at **1-800-ANALOGD** (**1-800-262-5643**) and follow the prompts.

VisualDSP++ Documentation Set

VisualDSP++ manuals may be purchased through Analog Devices Customer Service at **1-781-329-4700**; ask for a Customer Service representative. The manuals can be purchased only as a kit. For additional information, call **1-603-883-2430**.

If you do not have an account with Analog Devices, you will be referred to Analog Devices distributors. To get information on our distributors, log onto `http://www.analog.com/salesdir/continent.asp`.

Hardware Manuals

Hardware reference and instruction set reference manuals can be ordered through the Literature Center or downloaded from the Analog Devices website. The phone number is **1-800-ANALOGD** (**1-800-262-5643**). The manuals can be ordered by a title or by product number located on the back cover of each manual.

Datasheets

All datasheets can be downloaded from the Analog Devices website. As a general rule, any datasheet with a letter suffix (L, M, N) can be obtained from the Literature Center at **1-800-ANALOGD** (**1-800-262-5643**) or downloaded from the website. Datasheets without the suffix can be downloaded from the website only—no hard copies are available. You can ask for the datasheet by a part name or by product number.

If you want to have a datasheet faxed to you, call the phone number for that service **1-800-446-6212**. Follow the prompts and a list of datasheet code numbers will be faxed to you. Call the Literature Center first to find out if requested datasheets are available.

Contacting DSP Publications

Please send your comments and recommendation on how to improve our manuals and online Help. You can contact us by sending an E-mail to `dsp.techpubs@analog.com`.

Notation Conventions

The following table identifies and describes text conventions used in this manual. Additional conventions, which apply only to specific chapters, may appear throughout this document.

Example	Description
Close command (**File** menu)	Text in **bold** style indicates the location of an item within the VisualDSP++ environment's menu system. For example, the **Close** command appears on the **File** menu.
`{this | that}`	Alternative required items in syntax descriptions appear within curly brackets and separated by vertical bars; read the example as `this` or `that`.
`[this | that]`	Optional items in syntax descriptions appear within brackets and separated by vertical bars; read the example as an optional `this` or `that`.
`[this,…]`	Optional item lists in syntax descriptions appear within brackets delimited by commas and terminated with an ellipsis; read the example as an optional comma-separated list of `this`.
`.SECTION`	Commands, directives, keywords, and feature names are in text with letter gothic font.
`filename`	Non-keyword placeholders appear in text with italic style format.

Example	Description
	A note, providing information of special interest or identifying a related topic. In the online version of this book, the word **Note** appears instead of this symbol.
	A caution, providing information about critical design or programming issues that influence operation of a product. In the online version of this book, the word **Caution** appears instead of this symbol.

1 COMPILER

The C/C++ compiler (`cc219x`) is part of Analog Devices development software for ADSP-219x DSPs.

This chapter contains:

C/C++ Compiler Overview

The C/C++ compiler (`cc219x`) is designed to aid your DSP project development efforts by:

- Processing C and C++ source files, producing machine level versions of the source code and object files
- Providing relocatable code and debugging information within the object files
- Providing relocatable data and program memory segments for placement by the linker in the processors' memory

Using C/C++, developers can significantly decrease time-to-market since it gives them the ability to efficiently work with complex signal processing data types. It also allows them to take advantage of specialized DSP operations without having to understand the underlying DSP architecture.

The C/C++ compiler (`cc219x`) compiles ISO/ANSI standard C and C++ code for the ADSP-219x DSPs. Additionally, Analog Devices includes within the compiler a number of C language extensions designed to assist in DSP development. The `cc219x` compiler runs from the VisualDSP++ environment or from an operating system command line.

The C/C++ compiler (`cc219x`) processes your C and C++ language source files and produces ADSP-219x DSP's assembler source files. The assembler source files are assembled by the ADSP-219x assembler (`easm219x`). The assembler creates Executable and Linkable Format (ELF) object files that can either be linked (using the linker) to create an ADSP-219x executable file or included in an archive library (using `elfar`). The way in which the compiler controls the assemble, link, and archive phases of the process depends on the source input files and the compiler options used.

Source files contain the C/C++ program to be processed by the compiler. The `cc219x` compiler supports the ANSI/ISO standard definitions of the C and C++ languages. For information on the C language standard, see any of the many reference texts on the C language. Analog Devices recommends the Bjarne Stroustrup text "The C++ Programming Language" from Addison Wesley Longman Publishing Co (ISBN: 0201889544) (1997) as a reference text for the C++ programming language.

The `cc219x` compiler supports the proposed Embedded C++ Standard. The Embedded C++ Standard defines a subset of the full ISO/IEC 14882:1998 C++ language standard. The proposal excludes features that can detract from compiler performance in embedded systems, such as exception handling and run-time type identification. In addition to the embedded C++ standard features, `cc219x` supports the proposed standard definition, templates, and all other features of the full C++ standard except for the exception handling and run-time type identifications. The additional supported features provide extra functionality without degrading the compiler performance.

The `cc219x` compiler supports a set of C/C++ language extensions. These extensions support hardware features of the ADSP-219x DSPs. For information on these extensions, see "C/C++ Compiler Language Extensions" on page 1-59.

Compiler options are set in the VisualDSP++ Integrated Development and Debug Environment (IDDE) from the **Compile** page of the **Project Options** dialog box (see "Specifying Compiler Options in VisualDSP++" on page 1-11). The selections control how the compiler processes your source files, letting you select features that include the language dialect, error reporting, and debugger output.

For more information on the VisualDSP++ environment, see the *VisualDSP++ 3.5 User's Guide for ADSP-21xx DSPs* and online Help.

Standard Conformance

Analog C compilers conform to the ISO/IEC 998:1990 C standard and the ISO/IEC 14882:1998 C++ standard (in C++ mode) with a small number of currently unsupported features or areas of divergence.

Unsupported features are:

- Runtime-type information (RTTI) for C++ is not supported.
- Exceptions for C++ is not supported.
- ANSI features that require operating-system support are generally not supported. This includes `time.h` functionality in C.
- The `cc219x` compiler does not provide comprehensive support of NaN's, overflow and underflow conditions in their compiler support floating-point routines.

Areas of divergence from Standard:

- The `double` type is defined in terms of a single precision 32-bit `floats`, not double precision 64-bit `floats`.
- The `cc219x` compiler makes use of the DSP's double word (long) MAC instruction results to avoid having to explicitly promote integer operand multiplication to long. If the integer multiplication result overflows the integer type, then the result is not truncated as would be the case in strict ANSI terms. This behavior is disabled using the compiler's"-no-widen-muls" switch (on page 1-36).
- Normal ANSI C external linkage does not specifically require standard include files to be used, although it is recommended. In most cases, Analog C compilers do require standard include files to be used. This is because build configurations and optimization levels are used to select the correct and optimal implementation of the functions. For example, the include files may redefine standard C functions to use optimal compiler built-in implementations.

The compilers also support a number of language extensions that are essentially aids to DSP programmers and would not be defined in strict ANSI conforming implementations. These extensions are usually enabled by default and in some cases can be disabled using a command-line switch, if required.

These extensions include:

- Inline (function) which directs the compiler to integrate the function code into the code of the callers.
- Dual memory support keywords (`pm/dm`). Disabled using the `-no-extra-keywords` compiler switch (see on page 1-34).
- Placement support keyword (`section`). Disabled using the `-no-extra-keywords` compiler switch (see on page 1-34).
- Boolean type support keywords in C (`bool`, `true`, `false`). Disabled using the `-no-extra-keywords` compiler switch (see on page 1-34).
- Variable length array support
- Non-constant aggregate initializer support
- Indexed initializer support
- Preprocessor generated warnings
- Support for C++-style comments in C programs

For more information on these extensions, see the "C/C++ Compiler Language Extensions" on page 1-59.

Compiler Command-Line Interface

This section describes how the cc219x compiler is invoked from the command line, the various types of files used by and generated from the compiler, and the switches used to tailor the compiler's operation.

This section contains:

- "Running the Compiler" on page 1-7
- "Specifying Compiler Options in VisualDSP++" on page 1-11
- "Compiler Command-Line Switches" on page 1-12
- "Data Type Sizes" on page 1-54
- "Optimization Control" on page 1-56

By default, the compiler runs with Analog Extensions for C code enabled. This means that the compiler processes source files written in ANSI/ISO standard C language supplemented with Analog Devices extensions. Table 1-1 on page 1-8 lists valid extensions. By default, the compiler processes the input file through the listed stages to produce a .DXE file (see file names in Table 1-2 on page 1-10). Table 1-3 on page 1-13 lists the switches that select the language dialect.

Although many switches are generic between C and C++, some of them are valid in C++ mode only. A summary of the generic C/C++ compiler switches appears in Table 1-4 on page 1-13. A summary of the C++-specific compiler switches appears in Table 1-5 on page 1-20. The summaries are followed by descriptions of each switch.

When developing a DSP project, you may find it useful to modify the compiler's default options settings. The way you set the compiler's options depends on the environment used to run the DSP development software. See "Specifying Compiler Options in VisualDSP++" on page 1-11 for more information.

Running the Compiler

Use the following general syntax for the `cc219x` command line:

```
cc219x [-switch [-switch …]] sourcefile [sourcefile …]
```

where:

- `-switch` is the name of the switch to be processed. The compiler has many switches. These switches select the operations and modes for the compiler and other tools. Command-line switches are case sensitive. For example, `-O` is not the same as `-o`.

- `sourcefile` is the name of the file to be preprocessed, compiled, assembled, and/or linked.

  A file name can include the drive, directory, file name, and file extension. The compiler supports both Win32- and POSIX-style paths, using either forward or back slashes as the directory delimiter. It also supports UNC path names (starting with two slashes and a network name).

  The `cc219x` compiler uses the file extension to determine what the file contains and what operations to perform upon it. Table 1-2 on page 1-10 lists the allowed extensions.

For example, the following command line

```
cc219x -proc ADSP-2191 -O -Wremarks -o program.dxe source.c
```

runs `cc219x` with the following switches:

`-proc ADSP-2191`	Specifies the processor
`-O`	Specifies optimization for the compiler
`-Wremarks`	Selects extra diagnostic remarks in addition to warning and error messages
`-o program.dxe`	Selects a name for the compiled, linked output

`source.c`	Specifies the C language source file to be compiled

The following example command line

```
cc219x -proc ADSP-2191 -c++ source.cpp
```

runs `cc219x` with:

`-c++`	Specifies that all of the source files are written in C++
`source.cpp`	Specifies the C++ language source file for your program

The normal function of `cc219x` is to invoke the compiler, assembler, and linker as required to produce an executable object file. The precise operation is determined by the extensions of the input filenames, and by various switches.

In normal operation the compiler uses the following extension files to perform a specified action:

Table 1-1. File Extensions

Extension	Action
`.C .c .cc .cpp .cxx`	Source file is compiled, assembled, and linked
`.asm`, `.dsp`, or `.s`	Assembly language source file is assembled and linked
`.doj`	Object file (from previous assembly) is linked

If multiple files are specified, each is first processed to produce an object file; then all object files are presented to the linker.

You can stop this sequence at various points by using appropriate compiler switches, or by selecting options with the VisualDSP++ environment. These switches are `-E`, `-P`,`-M`,`-H`,`-S`, and `-c`.

Many of the compiler's switches take a file name as an optional parameter. If you do not use the optional output name switch, `cc219x` names the output for you. Table 1-2 on page 1-10 lists the type of files, names, and extensions `cc219x` appends to output files.

File extensions vary by command-line switch and file type. These extensions are influenced by the program that is processing the file, any search directories that you select, and any path information that you include in the file name. Table 1-2 indicates the searches that the preprocessor, compiler, assembler, and linker support. The compiler supports relative and absolute directory names to define file search paths. For information on additional search directories, see the `-I directory` switch (on page 1-29) and `-L directory` switch (on page 1-31).

When you provide an input or output file name as an optional parameter, use the following guidelines:

- Use a file name (include the file extension) with either an unambiguous relative path or an absolute path. A file name with an absolute path includes the drive, directory, file name, and file extension.

  Enclose long file names within straight quotes; for example, `"long file name.c"`. The `cc219x` compiler uses the file extension conventions listed in Table 1-2 to determine the input file type.

- Verify that the compiler is using the correct file. If you do not provide the complete file path as part of the parameter or add additional search directories, `cc219x` looks for input in the current directory.

ⓘ Using the verbose output switches for the preprocessor, compiler, assembler, and linker causes each of these tools to echo the name of each file as it is processed.

Table 1-2. Input and Output Files

Input File Extension	File Extension Description
.c	C/C++ source file.
.cc .cpp .cxx	C++ source file.
.h	Header file (referenced by a #include statement).
.pch	C++ pre-compiled header file.
.i	Preprocessed C/C++ source, created when preprocess only (-E compiler switch) is specified.
.ipa, opa	Interprocedural analysis files — used internally by the compiler when performing interprocedural analysis.
.s,.dsp,.asm	Assembler source file.
.idl	Interface definition language files for VCSE.
.ii, ti	Template instantiation files -- used internally by the compiler when instantiating C++ templates
.is	Preprocessed assembly source (retained when -save-temps is specified).
.ldf	Linker Description File.
.doj	Object file to be linked.
.dlb	Library of object files to be linked as needed.
.xml	DSP system memory map file output.
.sym	DSP system symbol map file output.

Specifying Compiler Options in VisualDSP++

When using the VisualDSP++ Integrated Development and Debug Environment (IDDE), use the **Compile** tab from the **Project Options** dialog box to get compiler functional options as shown on Figure 1-1.

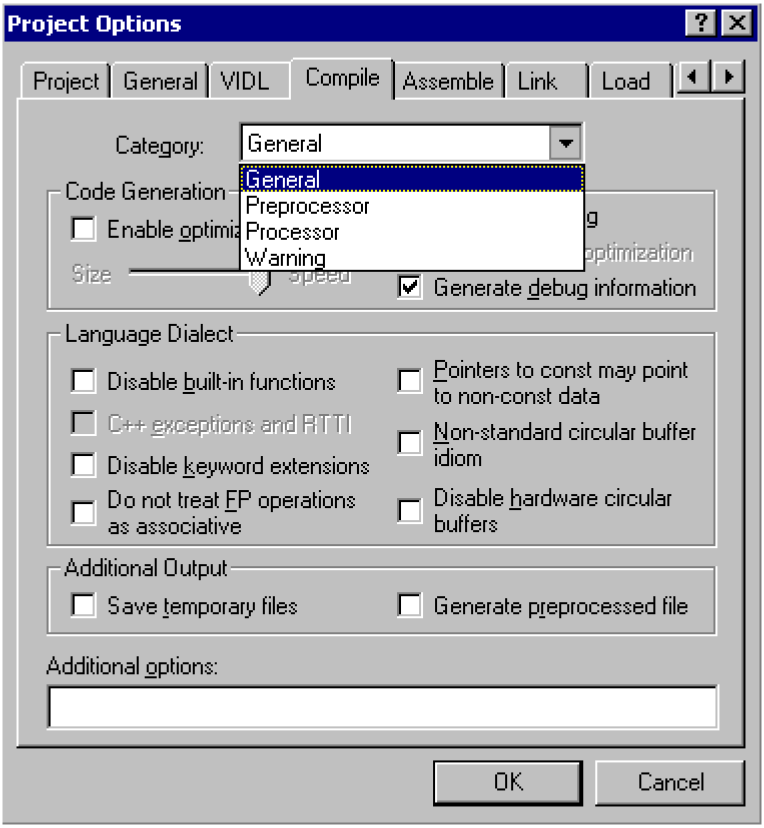

Figure 1-1. Project Options – Compile Property Page

There are four sub-pages you can access—**General**, **Preprocessor**, **Processor**, and **Warning**. Most project options have a corresponding compiler command-line switch described in "Compiler Command-Line Switches"

on page 1-12. The **Additional options** field in each sub-page is used to enter the appropriate file names and options that do not have corresponding controls on the **Compile** tab but are available as compiler switches.

Use the VisualDSP++ context-sensitive online Help to get information on compiler options you can specify in VisualDSP++.

Compiler Command-Line Switches

This section describes the command-line switches you can use when compiling. It contains a set of tables that provide a brief description of each switch. These tables are organized by type of a switch. Following these tables are sections that provide fuller descriptions of each switch.

C/C++ Compiler Switch Summaries

This section contains a set of tables that summarize generic and specific switches (options).

- "C or C++ Mode Selection Switches" in Table 1-3 on page 1-13
- "C/C++ Compiler Common Switches" in Table 1-4 on page 1-13
- "C++ Mode Compiler Switches" in Table 1-5 on page 1-20.

A brief description of each switch follows the tables, beginning on on page 1-21.

Table 1-3. C or C++ Mode Selection Switches

Switch Name	Description
`-c89` (on page 1-21)	Supports programs that conform to the ISO/IEC 9899:1990 standard.
`-c++` (on page 1-21)	Supports ANSI/ISO standard C++ with Analog Devices extensions.

Table 1-4. C/C++ Compiler Common Switches

Switch Name	Description
`sourcefile` (on page 1-21)	Specifies file to be compiled.
`-@` *`filename`* (on page 1-22)	Reads command-line input from the file.
-A *`name(tokens)`* (on page 1-22)	Asserts the specified name as a predicate.
`-alttok` (on page 1-23)	Allows alternative keywords and sequences in sources.
`-bss` (on page 1-24)	Causes the compiler to put global zero-initialized data into a separate BSS-style section.
`-build-lib` (on page 1-24)	Directs the librarian to build a library file.
`-C` (on page 1-24)	Retains preprocessor comments in the output file; active only with the `-E` or `-P` switch).
`-c` (on page 1-24)	Compiles and/or assembles only; does not link.
`-const-read-write` (on page 1-24)	Specifies that data accessed via a pointer to `const` data may be modified elsewhere.
`-D`*`macro`*`[=`*`definition`*`]` (on page 1-24)	Defines a macro.
`-debug-types` (on page 1-25)	Supports building a `*.h` file directly and writing a complete set of debugging information for the header file.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
-default-linkage-{asm|C|C++} (on page 1-25)	Sets the default linkage type (C, C++, asm).
-dry (on page 1-25)	Displays, but does not perform, main driver actions (verbose dry-run).
-dryrun (on page 1-26)	Displays, but does not perform, top-level driver actions (terse dry-run).
-E (on page 1-26)	Preprocesses, but does not compile, the source file.
-ED (on page 1-26)	Produce preprocessed file and compile source.
-EE (on page 1-26)	Preprocesses and compiles the source file.
-extra-keywords (on page 1-26)	Recognizes Analog Devices extensions to ISO/ANSI standards for C and C++. Default mode.
-flags-{tools} <arg1> [,arg2...] (on page 1-27)	Passes command-line switches through the compiler to other build tools.
-fp-associative (on page 1-27)	Treats floating-point multiplication and addition as associative.
-force-cirbuf (on page 1-27)	Treats array references of the form array[i%n] as circular buffer operations.
-full-version (on page 1-27)	Displays version information for build tools.
-g (on page 1-28)	Generates DWARF-2 debug information.
-H ((on page 1-28))	Outputs a list of header files, but does not compile the source file.
-HH (on page 1-28)	Outputs a list of included header files and compiles.
-h[elp] (on page 1-29)	Outputs a list of command-line switches with brief syntax descriptions.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-I-` (on page 1-29)	Establishes the point in the `include` directory list at which the search for header files enclosed in angle brackets should begin.
`-I` *`directory`* (on page 1-29)	Appends *directory* to the standard search path.
`-i` (on page 1-30)	Outputs only header details or makefile dependencies for `include` files specified in double quotes.
`-include` *`filename`* (on page 1-30)	Includes named file prior to preprocessing each source file.
`-ipa` (on page 1-30)	Specifies that interprocedural analysis should be performed for optimization between translation units.
`-jump-{pm|dm|same}` (on page 1-31)	Specifies where the compiler should place jump tables in memory
`-L` *`directory`* (on page 1-31)	Appends the specified directory to the standard library search path when linking.
`-l` *`library`* (on page 1-31)	Searches the specified library for functions when linking.
`-M` (on page 1-32)	Generates make rules only; does not compile.
`-MD` (on page 1-32)	Generates make rule and compiles.
`-MM` (on page 1-32)	Generates make rules and compiles.
`-Mo` *`filename`* (on page 1-32)	Makes dependencies other than `stdout`. This switch is used in conjunction with the `-ED` or `-MD` options.
`-Mt` *`filename`* (on page 1-32)	Makes dependencies, where the target is renamed as `filename`.
`-MQ` (on page 1-33)	Generates make rules only; does not compile. No notification when input files are missing.
`-map` *`filename`* (on page 1-33)	Directs the linker to generate a memory map of all symbols.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-mem` (on page 1-33)	Causes the compiler to invoke the Memory Initializer after linking the executable file.
`-no-alttok` (on page 1-33)	Does not allow alternative keywords and sequences in sources.
`-no-bss` (on page 1-33)	Causes the compiler to group global zero-initialized data into the same section as global data with non-zero initializers.
`-no-builtin` (on page 1-34)	Disables recognition of `__builtin` functions.
`-no-circbuf` (on page 1-34)	Disables the automatic generation of circular buffer code by the compiler.
`-no-defs` (on page 1-34)	Does not define any default preprocessor macros, include directories, library directories, libraries, run-time headers, or keyword extensions.
`-no-extra-keywords` (on page 1-34)	Does not define language extension keywords that could be valid C or C++ identifiers.
`-no-fp-associative` (on page 1-34)	Does not treat floating-point multiply and addition as an associative.
`-no_hardware_pc_stack` (on page 1-35)	Uses software stack instead of default hardware stack for return PC.
`-no-mem` (on page 1-35)	Causes the compiler to not invoke the Memory Initializer after linking. Set by default.
`-no-std-ass` (on page 1-35)	Prevents the compiler from defining standard assertions.
`-no-std-def` (on page 1-35)	Disables normal macro definitions; also disables Analog Devices keyword extensions that do not have leading underscores (__).
`-no-std-inc` (on page 1-35)	Searches for preprocessor header files only in the current directory and in directories specified with the `-I` switch.
`-no-std-lib` (on page 1-36)	Searches for only those linker libraries specified with the `-l` switch when linking.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-nothreads` (on page 1-36)	Specifies that compiled code does not need to be thread-safe.
`-no-widen-muls` (on page 1-36)	Disable widening multiplications optimization.
`-O` (on page 1-37)	Enables optimizations.
`-Oa` (on page 1-37)	Enables automatic function inlining.
`-Os` (on page 1-37)	Optimizes for code size.
`-Ov` *num* (on page 1-38)	Controls speed vs. size optimizations.
`-o` *filename* (on page 1-38)	Specifies the output file name.
`-oldasmcall-{csp|8x}` (on page 1-38)	Switches the operation of the `OldAsmCall` linkage specifier.
`-P` (on page 1-38)	Preprocesses, but does not compile, the source file. Omits line numbers in the preprocessor output.
`-PP` (on page 1-38)	Similar to `-P`, but does not halt compilation after preprocessing.
`-path-{asm|compiler|def|lib|link|mem}` *filename* (on page 1-39)	Uses the specified directory as the location of the specified compilation tool (assembler, compiler, library builder, or linker).
`-path-install` *directory* (on page 1-39)	Uses the specified directory as the location for all compilation tool.
`-path-output` *directory* (on page 1-39)	Specifies the location of non-temporary files.
`-path-temp` *directory* (on page 1-39)	Specifies the location of temporary files.
`-pch` (on page 1-40)	Generates and uses precompiled header files (*.`pch`)

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-pchdir` *directory* (on page 1-40)	Specifies the location of PCHRepository.
`-pedantic` (on page 1-40)	Issues compiler warnings for any constructs that are not ISO/ANSI standard C/C++-compliant.
`-pedantic-errors` (on page 1-40)	Issues compiler errors for any constructs that are not ISO/ANSI standard C/C++-compliant.
`-pplist` *filename* (on page 1-41)	Outputs a raw preprocessed listing to the specified file.
`-proc identifier` (on page 1-41)	Specifies that the compiler should produce code suitable for the specified DSP.
`-R directory` (on page 1-42)	Appends directory to the standard search path for source files.
`-R-` (on page 1-43)	Removes all directories from the standard search path for source files.
`-reserve <reg1>[,reg2...]` (on page 1-43)	Reserves the `I2`, `I3`, and `M0` registers from compiler use. Note: Reserving registers can have a detrimental effect on the compiler's optimization capabilities.
`-S` (on page 1-43)	Stops compilation before running the assembler.
`-s` (on page 1-43)	Removes debugging information from the output executable file when linking.
`-save-temps` (on page 1-44)	Saves intermediate compiler temporary files.
`-show` (on page 1-44)	Displays the driver command-line information.
`-si-revision` *version* (on page 1-44)	Specifies a silicon revision of the specified processor.
`-signed-bitfield` (on page 1-45)	Makes the default type for `int bitfields` signed.
`-signed-char` (on page 1-45)	Makes the default type for `char` signed.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-syntax-only` (on page 1-45)	Checks the source code for compiler syntax errors, but does not write any output.
`-sysdefs` (on page 1-45)	Defines the system definition macros.
`-T filename` (on page 1-46)	Uses the specified the Linker Description File as control input for linking.
`-threads` (on page 1-46)	Specifies that the build and link should be thread-safe.
`-time` (on page 1-47)	Displays the elapsed time as part of the output information on each part of the compilation process.
`-Umacro` (on page 1-47)	Undefines macro (s).
`-unsigned-bitfield` (on page 1-47)	Makes the default type for `bitfield` unsigned.
`-unsigned-char` (on page 1-48)	Makes the default type for `char` unsigned.
`-v` (on page 1-48)	Displays both the version and command-line information for all compilation tools as they process each file (version & verbose).
`-val-global <name-list>` (on page 1-48)	Adds global names.
`-verbose` (on page 1-48)	Displays command-line information for all compilation tools as they process each file.
`-version` (on page 1-49)	Displays version information for all compilation tools as they process each file.
`-W{error|remark|suppress|warn}` *`number`* (on page 1-49)	Overrides the default severity of the specified diagnostic messages (errors, remarks, or warnings).
`-Werror-limit` *`number`* (on page 1-49)	Stops compiling after reaching the specified number of errors.

Table 1-4. C/C++ Compiler Common Switches (Cont'd)

Switch Name	Description
`-Wremarks` (on page 1-49)	Indicates that the compiler may issue remarks, which are diagnostic messages even milder than warnings.
`-Wterse` (on page 1-49)	Issues only the briefest form of compiler warnings, errors, and remarks.
`-w` (on page 1-50)	Does not display compiler warning messages.
`-warn-protos` (on page 1-50)	Produces a warning when a function is called without a prototype.
`-workaround <workaround>[,<workaround>]*` (on page 1-50)	Enables code generator workaround for specific hardware defects.
`-write-files` (on page 1-50)	Enables compiler I/O redirection.
`-write-opts` (on page 1-50)	Passes the user options (but not input filenames) via a temporary file.
`-xref` *`filename`* (on page 1-51)	Outputs cross-reference information to the specified file.

Table 1-5. C++ Mode Compiler Switches

Switch Name	Description
`-anach` (on page 1-52)	Enables C++ anachronisms. The default mode.
`-no-anach` (on page 1-53)	Disables C++ anachronisms.
`-no-demangle` (on page 1-53)	Prevents filtering of any linker errors through the demangler.

C/C++ Mode Selection Switch Descriptions

The following command-line switches provide C/C++ mode selection.

-c89

The `-c89` switch directs the compiler to support programs that conform to the ISO/IEC 9899:1990 standard. For greater conformance to the standard, the following switches should be used: `-alttok`, `-const-read-write`, `no-extra-keywords`, and `-pedantic` (see in Table 1-4 on page 1-13).

-c++

The `-c++` (C++ mode) switch directs the compiler to assume that the source file(s) are written in ANSI/ISO standard C++ with Analog Devices language extensions.

All the standard features of C++ are accepted in the default mode except exception handling and run-time type identification because these impose a run-time overhead that is not desirable for all embedded programs.

C/C++ Compiler Common Switch Descriptions

The following command-line switches apply in C and C++ modes.

sourcefile

The *`sourcefile`* parameter (or parameters) switch specifies the name of the file (or files) to be preprocessed, compiled, assembled, and/or linked. A file name can include the drive, directory, file name, and file extension. The `cc219x` compiler uses the file extension to determine the operations to perform. Table 1-2 on page 1-10 lists the permitted extensions and matching compiler operations.

-@ filename

The `@ filename` (command file) switch directs the compiler to read command-line input from *`filename`*. The specified `filename` must contain driver options but may also contain source *`filenames`* and environment variables. It can be used to store frequently used options as well as to read from a file list.

-A name(tokens)

The `-A` (assert) switch directs the compiler to assert `name` as a predicate with the specified `tokens`. This has the same effect as the `#assert` preprocessor directive. The following assertions are predefined:

system	`embedded`
machine	`adsp219x`
cpu	`adsp219x`
compiler	`cc219x`

The `-A` *`name(value)`* switch is equivalent to including

```
#assert name(value)
```

in your source file, and both may be tested in a preprocessor condition in the following manner:

```
#if #name(value)
   // do something
#else
   // do something else
#endif
```

For example, the default assertions may be tested as:

```
#if #machine(adsp219x)
    // do something
#endif
```

> ⓘ The parentheses in the assertion should be quoted when using the `-A` switch, to prevent misinterpretation. No quotes are needed for a `#assert` directive in a source file.

-alttok

The `-alttok` alternative tokens) switch directs the compiler to allow digraph sequences in C and C++ source files. Additionally, the switch enables the recognition of these alternative operator keywords in C++ source files.

Keyword	Equivalent
`and`	`&&`
`and_eq`	`&=`
`bitand`	`&`
`bitor`	`|`
`compl`	`~`
`not`	`!`
`not_eq`	`!=`
`or`	`||`
`or_eq`	`|=`
`xor`	`^`
`xor_eq`	`^=`

> ⓘ To use them in C, you should use `#include <iso646.h>`.

-bss

The `-bss` switch causes the compiler to place global zero-initialized data into a BSS-style section (called "`bsz`"), rather than into the normal global data section. This is default mode. See also the `-no-bss` switch (on page 1-33).

-build-lib

The `-build-lib` (build library) switch directs the compiler to use `elfar` (librarian) to produce a library file (`.dlb`) as the output instead of using the linker to produce an executable file (`.dxe`). The `-o` option (on page 1-38) must be used to specify the name of the resulting library.

-C

The `-C` (comments) switch, which is only active in combination with the `-E`, `-EE`, `-ED`, `-P` or `-PP` switches, directs the preprocessor to retain comments in its output.

-c

The `-c` (compile only) switch directs the compiler to compile and/or assemble the source files, but stop before linking. The output is an object file (`.doj`) for each source file.

-const-read-write

The `-const-read-write` switch directs the compiler to specify that constants may be accessed as read-write data (as in ANSI C). The compiler's default behavior assumes that data referenced through `const` pointers will never change.

The `-const-read-write` switch changes the compiler's behavior to match the ANSI C assumption, which is that other `non-const` pointers may be used to change the data at some point.

-Dmacro[=definition]

The `-D` (define macro) switch directs the compiler to define a macro. If you do not include the optional definition string, the compiler defines the macro as the string '`1`'. If definition is required to be a character string constant, it must be surrounded by escaped double quotes. Note that the compiler processes `-D` switches on the command line before any `-U` (undefine macro) switches.

> This switch can be invoked with the **Definitions:** dialog field located in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Preprocessor** category.

-debug-types <file.h>

The `-debug-types` switch provides for building an `*.h` file directly and writing a complete set of debugging information for the header file. The `-g` option (on page 1-28) need not be specified with the `-debug-types` switch because it is implied. For example,

```
cc219x -debug-types anyHeader.h
```

The implicit `-g` option writes debugging information for only those `typedefs` that are referenced in the program. The `-debug-types` option provides complete debugging information for all `typedefs` and `structs`.

-default-linkage-{asm|C|C++}

The `-default-linkage-asm` (assembler linkage)/`-default-linkage-C` (C linkage)/`-default-linkage-C++` (C++ linkage) switch directs the compiler to set the default linkage type. C linkage is the default type in C mode, and C++ linkage is the default type in C++ mode.

> This switch can be specified in the **Additional Options** box located in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-dry

The `-dry` (verbose dry-run) switch directs the compiler to display main driver actions, but not to perform them.

-dryrun

The `-dryrun` (terse dry-run) switch directs the compiler to display top-level driver actions, but not to perform them.

-E

The `-E` (stop after preprocessing) switch directs the compiler to stop after the C/C++ preprocessor runs (without compiling). The output (preprocessed source code) prints to the standard output stream (`<stdout>`) unless the output file is specified with the `-o` switch. Note that the `-C` switch can be used in combination with the `-E` switch.

-ED

The `-ED` (run after preprocessing to file) switch directs the compiler to write the output of the C/C++ preprocessor to a file named `original_filename.i`. After preprocessing, compilation proceeds normally.

-EE

The `-EE` (run after preprocessing) switch directs the compiler to write the output of the C/C++ preprocessor to standard output. After preprocessing, compilation proceeds normally.

-extra-keywords

The `-extra-keywords` (enable short-form keywords) switch directs the compiler to recognize the Analog Devices keyword extensions to ISO/ANSI standard C and C++. This recognition includes keywords such as `pm` and `dm` without leading underscores which could affect conforming ISO/ANSI C and C++ programs. This is the default mode. The `-no-extra-keywords` switch (on page 1-34) can be used to disallow support for the additional keywords. Table 1-7 on page 1-60 provides a list and a brief description of keyword extensions.

-flags{-asm | -compiler | -lib | -link | -mem} switch [,switch2 [,...]]

The `-flags` (command-line input) switch directs the compiler to pass command-line switches to the other build tools. These tools are:

Option	Tool
`-flags-asm`	Assembler
`-flags-compiler`	Compiler
`-flags-lib`	Library Builder
`-flags-link`	Linker
`-flags-mem`	Memory Initializer

-force-circbuf

The `-force-circbuf` (circular buffer) switch instructs the compiler to make use of circular buffer facilities, even if the compiler cannot verify that the circular index or pointer is always within the range of the buffer. Without this switch, the compiler's default behavior is to be conservative, and not use circular buffers unless it can verify that the circular index or pointer is always within the circular buffer range. See "Circular Buffer Built-In Functions" on page 1-103.

-fp-associative

The `-fp-associative` switch directs the compiler to treat floating-point multiplication and addition as associative.

-full-version

The `-full-version` (display versions) switch directs the compiler to display version information for build tools used in a compilation.

-g

The `-g` (generate debug information) switch directs the compiler to output symbols and other information used by the debugger.

When the `-g` switch is used in conjunction with the `-O` (enable optimization) switch, the compiler performs standard optimizations. The compiler also outputs symbols and other information to provide limited source level debugging through VisualDSP++. This combination of options provides line debugging and global variable debugging.

You can invoke this switch by selecting the **Generate debug information** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

When `-g` and `-O` are specified, no debug information is available for local variables and the standard optimizations can sometimes re-arrange program code in a way that inaccurate line number information may be produced. For full debugging capabilities, use the `-g` switch without the `-O` switch.

-H

The `-H` (list headers) switch directs the compiler to output only a list of the files included by the preprocessor via the `#include` directive, without compiling. The `-o` switch (on page 1-38) may be used to specify the redirection of the list to a file.

-HH

The -`HH` (list headers and compile) switch directs the compiler to print to the standard output file stream a list of the files included by the preprocessor via the `#include` directive. After preprocessing, compilation proceeds normally.

-h[elp]

The `-help` (command-line help) switch directs the compiler to output a list of command-line switches with a brief syntax description.

-I-

The `-I-` (start include directory list) switch establishes the point in the `include` directory list at which the search for header files enclosed in angle brackets should begin. Normally, for header files enclosed in double quotes, the compiler searches in the directory containing the current input file; then the compiler reverts back to looking in the directories specified with the `-I` switch and then in the standard include directory.

For header files in angle brackets the compiler performs the latter two searches only.

It is possible to replace the initial search (within the directory containing the current input file) by placing the `-I-` switch at the point on the command line where the search for all types of header file should begin. All `include` directories on the command line specified before the `-I-` switch will only be used in the search for header files that are enclosed in double quotes.

This switch removes the directory containing the current input file from the `include` directory list.

-I directory [{,|;} **directory...**]

The `-I` (include search directory) switch directs the C/C++ preprocessor to append the directory (directories) to the search path for `include` files. This option may be specified more than once; all specified directories are added to the search path.

Include files, whose names are not absolute path names and that are enclosed in "..." when included, will be searched for in the following directories in this order:

1. The directory containing the current input file (the primary source file or the file containing the `#include`)
2. Any directories specified with the `-I` switch in the order they are listed on the command line
3. Any directories on the standard list:
   `<VisualDSP++ install dir>/.../include`

> If a file is included using the <...> form, this file will only be searched for by using directories defined in items 2 and 3 above.

-i

The `-i` (less includes) switch may be used with the `-H`, `-HH`, `-M`, or `-MM` switches to direct the compiler to only output header details (`-H`, `-HH`) or makefile dependencies (`-M`, `-MM`) for `include` files specified in double quotes.

-include filename

The `-include` (include file) switch directs the preprocessor to process the specified file before processing the regular input file. Any `-D` and `-U` options on the command line are always processed before an `-include` file.

-ipa

The `-ipa` (interprocedural analysis) switch turns on Interprocedural Analysis (IPA) in the compiler. This option enables optimization across the entire program, including between source files that were compiled separately. The `-ipa` option should be applied to all C and C++ files in the program. For more information, see "Interprocedural Analysis" on page 1-57. Specifying `-ipa` also implies setting the `-O` switch (on page 1-37).

> You can invoke this switch by selecting the **Interprocedural optimization** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-jump-{dm|pm|same}

The `-jump` (select jump table memory type) switch directs the compiler to place jump tables in data memory (`-jump-dm`), program memory (`-jump-pm`), or the same memory section as the function to which it applies (`-jump-same`). Jump tables are storage that might be required to hold in memory target addresses for branch instruction used in complex IF-THEN-ELSE statements or switch statements. The default storage memory for jump tables is data memory (`-jump-dm`).

-L directory [{,|;} directory…]

The `-L` (library search directory) switch directs the linker to append the directory to the search path for library files.

-l library

The `-l` (link library) switch directs the linker to search the library for functions when linking. The library name is the portion of the file name between the `lib` prefix and `.dlb` extension. For example, the compiler command-line switch `-lc` directs the linker to search in the library named c for functions. This library resides in a file named `libc.dlb`.

Normally, you should list all object files on the command line before the `-l` switch. This ensures that the functions and global variables the object files refer to are loaded in the given order. This option may be specified more than once; libraries are searched as encountered during the left-to-right processing of the command line.

-M

The `-M` (generate make rules only) switch directs the compiler not to compile the source file but to output a rule, which is suitable for the make utility, describing the dependencies of the main program file. The format of the make rule output by the preprocessor is:

```
object-file: include-file ...
```

-MD

The `-MM` (generate make rules and compile) switch directs the preprocessor to print to a file called `original_filename.d` a rule describing the dependencies of the main program file. After preprocessing, compilation proceeds normally. See also the `-Mo` switch.

-MM

The `-MM` (generate make rules and compile) switch directs the preprocessor to print to `stdout` a rule describing the dependencies of the main program file. After preprocessing, compilation proceeds normally.

-Mo filename

The `-Mo` *`filename`* (preprocessor output file) switch directs the compiler to use *`filename`* for the output of `-MD` or `-ED` switches.

-Mt filename

The `-Mt` *`filename`* (output make rule for the named source) switch specifies the name of the source file for which the compiler generates the make rule when you use the `-M` or `-MM` switch. If the named file is not in the current directory, you must provide the path name in double quotation marks (“”). The new file name will override the default `base.doj`. The `-Mt` option supports the `.IMPORT` extension.

-MQ

The `-MQ` switch directs the compiler not to compile the source file but to output a rule. In addition, the `-MQ` switch does not produce any notification when input files are missing.

-map filename

The `-map` (generate a memory map) switch directs the linker to output a memory map of all symbols. The map file name corresponds to the *`filename`* argument. For example, if the *`filename`* argument is `test`, the map file name is `test.xml`. The `.xml` extension is added where necessary.

-mem

The `-mem` (invoke memory initializer) switch causes the compiler to invoke the Memory Initializer tool after linking the executable. The MemInit tool can be controlled through the `-flags-mem` switch (on page 1-27).

-no-alttok

The `-no-alttok` (disable alternative tokens) switch directs the compiler to not accept alternative operator keywords and digraph sequences in the source files. This is the default mode. For more information, see the `-alttok` switch on page 1-23.

-no-bss

The `-no-bss` switch causes the compiler to keep zero-initialized and non-zero-initialized data in the same data section, rather than separating zero-initialized data into a different, BSS-style section. See also the `-bss` switch (on page 1-23).

-no-builtin

The `-no-builtin` (no builtin functions) switch directs the compiler to recognize only built-in functions that begin with two underscores (__). Note that this switch influences many functions. This switch also predefines the `__NO_BUILTIN` preprocessor macro. For more information on built-in functions, see "Compiler Built-In Functions" on page 1-94.

-no-circbuf

The `-no-circbuf` (no circular buffer) switch directs the compiler not to define any default preprocessor macros, include directories, library directories, libraries, or run-time headers. It also disables the Analog Devices C/C++ keyword extensions.

-no-defs

The `-no-defs` (disable defaults) switch directs the preprocessor not to define any default preprocessor macros, include directories, library directories, libraries, or run-time headers.

-no-extra-keywords

The `-no-extra-keywords` (disable short-form keywords) switch directs the compiler not to recognize the Analog Devices keyword extensions that might affect conformance to ISO/ANSI programs. These extensions include keywords, such as `pm` and `dm`, which may be used as identifiers in conforming programs. The alternate keywords that are prefixed with two leading underscores, such as `__pm` and `__dm`, continue to work.

The "-extra-keywords" switch (on page 1-26) can be used to explicitly request support for the additional keywords.

-no-fp-associative

The `-no-fp-associative` switch directs the compiler not to treat floating-point multiplication and addition as associative.

-no_hardware_pc_stack

The `-no_hardware_pc_stack` switch directs the compiler to produce code to avoid the possibility of overflowing the hardware call stack. This is done by popping, from the hardware PC stack, the return address at the start of each function and storing this value on the software data stack to be pushed back on the hardware stack just before returning from the function.

-no-mem

The `-no-mem` switch causes the compiler to not invoke the Memory Initializer tool after linking the executable. This is the default setting. See also "-mem" on page 1-33.

-no-std-ass

The `-no-std-ass` (disable standard assertions) switch prevents the compiler from defining the standard assertions. See the `-A` switch (on page 1-22) for the list of standard assertions.

-no-std-def

The `-no-std-def` (disable standard macro definitions) prevents the compiler from defining any default preprocessor macro definitions.

-no-std-inc

The `-no-std-inc` (disable standard include search) switch directs the C preprocessor to limit its search for header files to the current directory and directories specified with `-I` switch.

ⓘ You can invoke this switch by selecting the **Ignore standard include paths** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Preprocessor** category.

-no-std-lib

The `-no-std-lib` (disable standard library search) switch directs the linker to limit its search for libraries to directories specified with the `-L` switch (on page 1-31).

-no-widen-muls

The `-no-widen-muls` (disable widening multiplications) switch disables the compiler optimization which it performs on multiplication operations.

By default, the compiler, attempts to optimize integer multiplication operations which are stored in a long result to utilize the double-word MAC result registers of the ADSP-219x processors. The code produced this way is better suited to the processor and therefore more efficient.

However, this optimization can generate overflow results which are not consistent in some cases and may differ from expected results depending on the optimizations enabled and the way that the source is written. The inconsistency and differences are seen if an overflow and truncation of the integer operands would normally occur.

When the optimization is applied, there is no truncation. When the optimization is disabled, the result of overflow will be truncated to integer size before being stored in the long result.

-nothreads

The `-nothreads` (disable thread-safe build) switch specifies that all compiled code and libraries used in the build need not be thread-safe. This is the default setting when the `-threads` (enable thread-safe build) switch is not used.

-O

The `-O` (enable optimizations) switch directs the compiler to produce code that is optimized for performance. Optimizations are not enabled by default for the `cc219x` compiler.

> You can invoke this switch by selecting the **Enable optimization** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category.

-Oa

The `-Oa` (automatic function inlining) switch enables the inline expansion of C/C++ functions which are not necessarily declared inline in the source code. The amount of auto-inlining the compiler performs is controlled using the `-Ov` (optimize for speed versus size) switch (on page 1-38). Therefore, use of `-Ov100` indicates that as many functions as possible will be auto-inlined whereas `-Ov0` prevents any function from being auto-inlined. Specifying `-Oa` also implies the use of `-O`.

> When remarks are enabled, the compiler will produce a remark to indicate each function that is inlined.

-Os

The `-Os` (optimize for size) switch directs the compiler to produce code that is optimized for size. This is achieved by performing all optimizations except those that increase code size. The optimizations not performed include loop unrolling, some delay slot filling, and jump avoidance. The compiler also uses a function to save and restore preserved registers for a function instead of generating the more cycle-efficient inline default versions.

-Ov num

The `-Ov num` (optimize for speed versus size) switch directs the compiler to produce code that is optimized for speed versus size. The '`num`' should be an integer between 0 (purely size) and 100 (purely speed).

-o filename

The `-o` (output file) switch directs the compiler to use *filename* for the name of the final output file.

-oldasmcall-{csp|8x}

The `-oldasmcall-{csp|8x}` switch changes the operation of the `OldAsmCall` linkage specifier between compatibility call for the ADSP-21csp01 DSP and legacy ADSP-218x DSPs (`-oldasmcall-csp` is default). Therefore, `-oldasmcall-csp` makes code compatible with (legacy) ADSP-21csp01 code, while `-oldasmcall-8x` makes code compatible with (legacy) ADSP-218x code.

-P

The `-P` (omit line numbers) switch directs the compiler to stop after the C preprocessor runs (without compiling) and to omit the `#line` preprocessor directives (with line number information) in the output from the preprocessor. The `-C` switch may be used in combination with the `-P` switch.

-PP

The `-PP` (omit line numbers and compile) switch directs the compiler to omit the `#line` preprocessor directives with line number information from the preprocessor output. After preprocessing, compilation proceeds normally.

-path {-asm | -compiler | -def | -lib | -link | -mem} filename

The `-path` (tool location) switch directs the compiler to use the specified component in place of the default-installed version of the compilation tool. The component should comprise a relative or absolute path to its location. Respectively, the tools are the assembler, compiler, driver definitions file, librarian, linker or memory initializer. Use this switch when you wish to override the normal version of one or more of the tools. The `-path` switch also overrides the directory specified by the `-path-install` switch (on page 1-39).

-path-install directory

The `-path-install` (installation location) switch directs the compiler to use the specified directory as the location for all compilation tools instead of the default path. This is useful when working with multiple versions of the tool set.

You can selectively override this switch with the `-path-tool` switch.

-path-output directory

The `-path-output` (non-temporary files location) switch directs the compiler to place output files in the specified directory.

-path-temp directory

The `-path-temp` (temporary files location) switch directs the compiler to place temporary files in the specified directory.

-pch

The `-pch` (precompiled header) switch directs the compiler to automatically generate and use precompiled header files. A precompiled output header has a `.pch` extension attached to the source file name. By default, all precompiled headers are stored in a directory called `PCHRepository`.

> Precompiled header files can significantly speed compilation; precompiled headers tend to occupy more disk space.

-pchdir directory

The `-pchdir` (locate PCHRepository) switch specifies the location of an alternative `PCHRepository` directory for storing and invocation of precompiled header files. If the directory does not exist, the compiler creates it. Note that `-o` (output) does not influence the `-pchdir` option.

-pedantic

The `-pedantic` (ANSI standard warnings) switch causes the compiler to issue warnings for any constructs found in your program that do not strictly conform to the ISO/ANSI standard for C or C++.

> The compiler may not detect all noconforming constructs. In particular, the `-pedantic` switch does not cause the compiler to issue errors when Analog Devices keyword extensions are used.

-pedantic-errors

The `-pedantic-errors` (ANSI C errors) switch causes the compiler to issue errors instead of warnings for cases described in the `-pedantic` switch.

-pplist filename

The `-pplist` (preprocessor listing) switch directs the preprocessor to output a listing to the named file. When more than one source file has been preprocessed, the listing file contains information about the last file processed. The generated file contains raw source lines, information on transitions into and out of include files, and diagnostics generated by the compiler.

Each listing line begins with a key character that identifies its type as:

Character	Meaning
`N`	Normal line of source
`X`	Expanded line of source
`S`	Line of source skipped by `#if` or `#ifdef`
`L`	Change in source position
`R`	Diagnostic message (remark)
`W`	Diagnostic message (warning)
`E`	Diagnostic message (error)
`C`	Diagnostic message (catastrophic error)

-proc processor

The `-proc` (target processor) switch specifies that the compiler should produce code suitable for the specified processor. The *`processor`* identifiers directly supported in VisualDSP++ 3.5 are:

ADSP-2191, ADSP-21990, ADSP-21991, ADSP-21992, ADSP-2195, ADSP-2196, ADSP-2192-12 and ADSP-219x

For example,

```
cc219x -proc ADSP-2191 -o bin\p1.doj p1.asm
```

> If no target is specified with the `-proc` switch, the system uses the `ADSP-219x` setting as a default.

If the processor identifier is unknown to the compiler, it attempts to read required switches for code generation from the file `<processor>.ini`. The assembler searches for the `.ini` file in the VisualDSP ++ System folder. For custom processors, the compiler searches the section "`proc`" in the `<processor>.ini` for key '`architecture`'. The custom processor must be based on an architecture key that is one of the known processors. For example, `-proc Customxxx` searches the `Customxxx.ini` file.

When compiling with the `-proc` switch, the appropriate processor macro as well as `__ADSP21XX__` and `__ADSP219X__` preprocessor macros are defined as 1. For example, `__ADSP2191__` and `__ADSP219X__` are 1.

> See also "-si-revision version" on page 1-44 for more information on silicon revision of the specified processor.

-R directory [{;|,}*directory* …]

The `-R` (add source directory) switch directs the compiler to add the specified directory to the list of directories searched for source files.

On Windows™ platforms, multiple source directories are given as a comma or semicolon separated list. The compiler searches for the source files in the order specified on the command line. The compiler searches the specified directories before reverting to the current project directory. This option is position-dependent on the command line; that is, it affects only source files that follow the option.

> Source files whose file names begin with `/`, `./`, or `../` (or Windows equivalent) and contain drive specifiers (on Windows platforms) are not affected by this option.

-R-

The `-R-` (disable source path) switch removes all directories from the standard search path for source files, effectively disabling this feature.

This option is position-dependent on the command line; it only affects files following it.

-reserve ***register***[, ***register*** ...]

The `-reserve` (reserve register) switch directs the compiler **not** to use the specified register(s). This guarantees that a known register or set of registers is available for auto buffering.

You can reserve the `I2`, `I3`, and `M0` registers. Separate register names with commas on the compiler command line. Reserving registers seriously reduces the effectiveness of compiler optimizations and should only be done when essential.

-S

The `-S` (stop after compilation) switch directs the compiler to stop compilation before running the assembler. The compiler outputs an assembler file with a `.s` extension.

You can invoke this switch by selecting the **Stop after: Compiler** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **General** category selection.

-s

The `-s` (strip debugging information) switch directs the compiler to remove debugging information (symbol table and other items) from the output executable file during linking.

-save-temps

The `-save-temps` (save intermediate files) switch directs the compiler to retain intermediate files generated and normally removed as part of the various compilation stages. These intermediate files are placed in the `-path-output` specified output directory or the build directory if the `-path-output` switch (on page 1-39) is not used. See Table 1-2 on page 1-10 for a list of intermediate files.

-show

The `-show` (display command line) switch directs the compiler to echo all command-line arguments, expanded option files switches, and environment variables used by the compiler

-si-revision version

The `-si-revision` *`version`* (silicon revision) switch sets the version of the hardware which is the required target for the build. It is used to enable inherent behavior relating to any errata in specific silicon revisions. The revision can be specified as "none" or a number of the form described by regular expression `[0-9]+\.[0-9]{1,3}` (for example, 1.123). The compiler defines a macro `__SILICON_REVISION__` to a value specific to each silicon revision. For unknown revisions, the compiler will generate a warning and default to the latest known revision.

The parameter "*`version`*" represents a silicon revision of the processor specified by the `-proc` switch (on page 1-41). The "`none`" revision disables support for silicon errata. For example,

```
cc219x -proc ADSP-2191 -si-revision 0.1
```

ⓘ In the absence of silicon revision, the compiler selects the greatest silicon revision it "knows" about, if any.

A compiler will "pass along" the appropriate `-si-revision` switch setting when invoking another VisualDSP++ tool, for example, when the compiler driver invokes assembler and linker.

-signed-bitfield

The `-signed-bitfield` (make plain bitfields signed) switch directs the compiler to make bitfields which have not been declared with an explicit signed or unsigned keyword to be signed. This switch does not effect plain one-bit bitfields which are always unsigned. This is the default mode. See also the `-unsigned-bitfield` switch (on page 1-47).

-signed-char

The `-signed-char` (make char signed) switch directs the compiler to make the default type for `char signed`. The compiler also defines the `__SIGNED_CHARS__` macro. This is the default mode when the `-unsigned-char` (make char unsigned) switch is not used.

-syntax-only

The `-syntax-only` (only check syntax) switch directs the compiler to check the source code for syntax errors and warnings. No output files will be generated with this switch.

-sysdefs

The `-sysdefs` (system definitions) switch directs the compiler to define several preprocessor macros describing the current user and user's system. The macros are defined as character string constants and are used in functions with null-terminated string arguments.

The following macros are defined if the system returns information for them:

Macro	Description
`__HOSTNAME__`	The name of the host machine
`__MACHINE__`	The machine type of the host machine
`__SYSTEM__`	The OS name of the host machine
`__USERNAME__`	The current user's login name
`__GROUPNAME__`	The current user's group name
`__REALNAME__`	The current user's real name

`__MACHINE__`, `__GROUPNAME__`, and `__REALNAME__` are not available on Windows platforms.

-T filename

The `-T` (Linker Description File) switch directs that the linker, when invoked, will use the specified Linker Description File (LDF). If `-T` is not specified, a default `.LDF` file is selected based on the processor variant.

-threads

The `-threads` (enable thread-safe build) specifies that the build and link should be thread-safe. The macro `_ADI_THREADS` is defined to one (1). It is used for conditional compilation by the preprocessor and by the default Linker Description Files to link with thread-safe libraries.

This switch is only likely to be used by applications involving the VisualDSP++ Kernel (VDK).

-time

The `-time` (time the compiler) switch directs the compiler to display the elapsed time as part of the output information on each part of the compilation process.

-Umacro

The `-U` (undefine macro) switch lets you undefine macros. If you specify a macro name, it will be undefined. The compiler processes all `-D` (define macro) switches on the command line before any `-U` (undefine macro) switches.

You can invoke this switch by selecting the **Undefines** field in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Preprocessor** category.

-unsigned-bitfield

The `-unsigned-bitfield` (make plain bitfields unsigned) switch directs the compiler to make bitfields which have not been declared with an explicit signed or unsigned keyword to be unsigned. This switch does not effect plain one-bit bitfields which are always unsigned.

For example, given the declaration

```
struct {
   int a:2;
   int b:1;
   signed int c:2;
   unsigned int d:2;
} x;
```

the `bitfield` values are:

Field	-signed-bitfield	-unsigned-bitfield	Why
x.a	-2..1	0..3	Plain field
x.b	0..1	0..1	One bit
x.c	-2..1	-2..1	Explicit signed
x.d	0..3	0..3	Explicit unsigned

See also the `-signed-bitfields` switch (on page 1-45).

-unsigned-char

The `-unsigned-char` (make char unsigned) switch directs the compiler to make the default type for `char` unsigned. The compiler also undefines the `__SIGNED_CHARS__` preprocessor macro

-v

The `-v` (version and verbose) switch directs the compiler to display both the version and command-line information for all the compilation tools as they process each file.

-val-global <name-list>

The `-val-global` (add global names) switch directs the compiler that the names given by <**name-list**> are present in all globally defined variables. The list is separated by double colons(::). In C++, these names are encoded as enclosing namespace or classes. In C, the names are prefixed and separated by underscores (_). The compiler will issue an error on any globally defined variable in the current source module(s) not using <**name-list**>. This switch is used to define VCSE components.

-verbose

The `-verbose` (display command line) switch directs the compiler to display command-line information for all the compilation tools as they process each file.

-version

The `-version` (display compiler version) switch directs the compiler to display its version information.

-W {error|remark|suppress|warn} [.number...]

The `-W {...}` *`number`* (override error message) switch directs the compiler to override the specified diagnostic messages (errors, remarks, or warnings). The `num` argument specifies the message to override.

At compilation time, the compiler produces a number for each specific compiler diagnostic message. The `{D}` (discretionary) string after the diagnostic message number indicates that the diagnostic may have its severity overridden. Each diagnostic message is identified by a number that is used across all compiler software releases.

-Werror-limit number

The `-Werror-limit` (maximum compiler errors) switch lets you set a maximum number of errors for the compiler.

-Wremarks

The `-Wremarks` (enable diagnostic warnings) switch directs the compiler to issue remarks, which are diagnostic messages that are even milder than warnings.

You can invoke this switch by selecting the **Enable remarks** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Warning** selection.

-Wterse

The `-Wterse` (enable terse warnings) switch directs the compiler to issue the briefest form of warnings. This also applies to errors and remarks.

-w

The `-w` (disable all warnings) switch directs the compiler not to issue warnings.

You can invoke this switch by selecting the **Disable all warnings and remarks** check box in the VisualDSP++ **Project Options** dialog box, **Compile** tab, **Warning** selection.

-warn-protos

The `-warn-protos` (prototypes warning) switch directs the compiler to produce a warning message when a function is called without a full prototype being supplied.

-workaround <workaround>[,**<workaround>**]*

The -`workaround` switch enables code generator workaround for specific hardware defects. Example of a valid workaround: `type32a-anomaly`.

-write-files

The `-write-files` (enable driver I/O redirection) switch directs the compiler driver to redirect the file name portions of its command line through a temporary file. This technique helps with handling long file names, which can make the compiler driver's command line too long for some operating systems.

-write-opts

The `-write-opts` switch directs the compiler to pass the user-specified options (but **not** the input file names) to the main driver via a temporary file. This can be helpful if the resulting main driver command line would otherwise be too long.

-xref <filename>

The `-xref` (cross-reference list) switch directs the compiler to write cross-reference listing information to the specified file. When more than one source file has been compiled, the listing contains information about the last file processed. For each reference to a symbol in the source program, a line of the form

```
symbol-id name ref-code filename line-number column-number
```

is written to the named file. The `symbol-id` represents a unique decimal number for the symbol, and `ref-code` is one of the following characters:

Character	Meaning
`D`	Definition
d	Declaration
M	Modification
A	Address taken
U	Used
C	Changed (used and modified)
R	Any other type of reference
E	Error (unknown type of reference

C++ Mode Compiler Switch Descriptions

The following switches apply only to C++.

-anach

The `-anach` (enable C++ anachronisms) directs the compiler to accept some language features that are prohibited by the C++ standard but still in common use. This is the default mode. Use the `-no-anach` switch for greater standard compliance.

The following anachronisms are accepted in the default C++ mode:

- Overload is allowed in function declarations. It is accepted and ignored.
- Definitions are not required for static data members that can be initialized using default initialization. The anachronism does not apply to static data members of template classes; they must always be defined.
- The number of elements in an array may be specified in an array delete operation. The value is ignored.
- A single `operator++()` and `operator--()` function can be used to overload both prefix and postfix operations.
- The base class name may be omitted in a base class initializer if there is only one immediate base class.
- Assignment to this in constructors and destructors is allowed. This is allowed only if anachronisms are enabled and the assignment to this configuration parameter is enabled.
- A bound function pointer (a pointer to a member function for a given object) can be cast to a pointer to a function.

- A nested class name may be used as a un-nested class name provided no other class of that name has been declared. The anachronism is not applied to template classes.
- A reference to a `non-const` type may be initialized from a value of a different type. A temporary is created, it is initialized from the (converted) initial value, and the reference is set to the temporary.
- A reference to a `non-const` class type may be initialized from an `rvalue` of the `class` type or a derived class thereof. No (additional) temporary is used.
- A function with old-style parameter declarations is allowed and may participate in function overloading as though it were prototyped. Default argument promotion is not applied to parameter types of such functions when the check for compatibility is done, so that the following statements declare the overload of two functions named `f`.

```
int f(int);
int f(x) char x; { return x; }
```

-no-anach

The `-no-anach` (disable C++ anachronisms) switch directs the compiler to disallow some old C++ language features that are prohibited by the C++ standard. See the `-anach` switch (on page 1-52) for a full description of these features.

-no-demangle

The `-no-demangle` (disable demangler) switch directs the compiler to prevent the driver from filtering any linker errors through the demangler. The demangler's primary role is to convert the encoded name of a function into a more understandable version of the name.

Data Type Sizes

The sizes of intrinsic C/C++ data types are selected by Analog Devices so that normal C/C++ programs execute with hardware-native data types and therefore at high speed.

Table 1-6 shows the size used for each of the intrinsic C/C++ data types

Table 1-6. Data Type Sizes for ADSP-219x Processors

Type	**Bit Size**	`sizeof` **returns**
`char`	8 bits signed	1
`unsigned char`	8 bits unsigned	1
`int`	16 bits signed	1
`unsigned int`	16 bits unsigned	1
`short`	16 bits signed	1
`unsigned short`	16 bits unsigned	1
`long`	32 bits signed	2
`unsigned long`	32 bits unsigned	2
`float`	32 bits float	2
`double`	32 bits float	2
pointer	16 bits	1
function pointer	32 bits	2
`fract16`	16 bits fractional	1
`fract32`	32 bits fractional	2

On any platform the basic type `int` will be the native word size. The data type `long` is 32 bits, as is `float`. A pointer is the same size as an `int`.

In the ADSP-219x processor architecture, the `long long int`, unsigned `long long int`, and `long double` data types are not implemented (they will not be redefined to other types). In general, `double` word data types should be expected to run more slowly, relying largely on software-emulated arithmetic.

Analog Devices does not support data sizes smaller than a single word location for the ADSP-219x processors. For the current processors, this means that both `short` and `char` have the same size as `int`. Although 16-bit `chars` are unusual, they do conform to the standard.

Type `double` poses a special problem. The C language tends to default to `double` for constants and in many floating-point calculations. Without some special handling, many programs would inadvertently end up using slow-speed emulated 64-bit floating-point arithmetic, even when variables are declared consistently as `float`.

In order to avoid this problem and provide the best performance, the size of `double` on the ADSP-219x processors is always 32 bits. This should be acceptable for most DSP programming. It is not, however, fully standard conforming.

The standard `include` files automatically redefine the math library interfaces such that functions like `sin` can be directly called with the proper size operands. Therefore,

```
float sinf (float);      /* 32-bit */
double sin (double);     /* 32-bit */
```

For full descriptions of these functions and their implementation, see Chapter 4, “DSP Run-Time Library”.

Optimization Control

The general aim of compiler optimizations is to generate correct code that executes fast and is small in size. Not all optimizations are suitable for every application or possible all the time so the compiler optimizer has a number of configurations, or optimization levels, which can be applied when suitable. Each of these levels are enabled by one or more compiler switches (and VisualDSP++ project options) or pragmas.

> ⓘ Refer Chapter 2, "Achieving Optimal Performance from C/C++ Source Code" for information on how to obtain maximal code performance from the compiler.

The following list identifies several optimization levels. The levels are notionally ordered with least optimization listed first and most optimization listed last. The descriptions for each level outline the optimizations performed by the compiler and identifies any switches or pragmas required or that have direct influence on the optimization levels performed.

- **Debug**
  The compiler produces debug information to ensure that the object code matches the appropriate source code line. See "-g" on page 1-28 for more information.

- **Default**
  The compiler does not perform any optimizations by default when none of the compiler optimizations switches are used (or enabled in VisualDSP++ project options). Default optimizations level can be enabled using the `optimize_off` pragma (on page 1-125).

- **Procedural Optimizations**
  The compiler performs advanced, aggressive optimization on each procedure in the file being compiled. The optimizations can be directed to favor optimizations for speed (`-O`) or space (`-Os`) or a factor between speed and space (`-Ov`). If debugging is also

requested, the optimization is given priority so the debugging functionality may be limited. See "-O" on page 1-37, "-Os" on page 1-37, and "-Ov num" on page 1-38.

Procedural optimizations for speed and space (`-O` and `-Os`) can be enabled in C/C++ source using the pragma `optimize_{for_speed|for_space}` (see on page 1-125 more information on optimization pragmas).

- **Automatic Inlining**
  The compiler automatically inlines C/C++ functions which are not necessarily declared as inline in the source code. It does this when it has determined that doing so will reduce execution time. How aggressively the compiler performs automatic inlining is controlled using the `-Ov` switch. Automatic inlining is enabled using the `-Oa` switch and additionally enables Procedural Optimizations (`-O`). Refer to "-Oa" on page 1-37, "-Ov num" on page 1-38, and "-O" on page 1-37 for more information.

- **Interprocedural optimizations (IPA)**
  The compiler performs advanced, aggressive optimization over the whole program, in addition to the per-file optimizations in procedural optimization. IPA is enabled using the `-ipa` switch and additionally enables Procedural Optimizations (`-O`).
  See "-ipa" on page 1-30 and "-O" on page 1-37 for more information.

Interprocedural Analysis

The compiler has an optimization capability called *Interprocedural Analysis* (IPA) that allows the compiler to optimize across translation units instead of within individual translation units. This capability allows the compiler to see all of the source files used in a final link at compilation time and to use that information while optimizing.

Interprocedural analysis is enabled by selecting the **Interprocedural analysis** option on the **Compiler** tab (accessed via the VisualDSP++ **Project Options** dialog box), or by specifying the `-ipa` command-line switch (on page 1-30). The `-ipa` switch automatically enables the `-O` switch to turn on optimization.

Use of the `-ipa` switch causes additional files to be generated along with the object file produced by the compiler. These files have `.ipa` and `.opa` filename extensions and should not be deleted manually unless the associated object file is also deleted.

All of the `-ipa` optimizations are invoked after the initial link; when a special program has called, the prelinker re-invokes the compiler to perform the new optimizations.

Because a file may be recompiled by the prelinker, you cannot use the `-S` option to see the final optimized assembler file when `-ipa` is enabled. Instead, you must use the `-save-temps` switch, so that the full compile/link cycle can be performed first.

Because IPA operates only during the final link, the `-ipa` switch has no benefit when compiling the source files to object format for inclusion in a library. Although IPA will generate usage information for potential additional optimizations at the final link stage, neither the usage information nor the module's source file are available when the linker includes a module from a library. Therefore, each library module is compiled to the normal `-O` optimization level.

The prelinker inspects object modules included from libraries and other object files which were not compiled with the `-ipa` switch to see whether there are hidden references to the functions and variables defined in those objects which were compiled with the `-ipa` switch, and optimizes those variables and functions accordingly.

C/C++ Compiler Language Extensions

The compiler supports a set of extensions to the ANSI standard for the C and C++ languages. These extensions add support for DSP hardware and allow some C++ programming features when compiling in C mode. The extensions are also available when compiling in C++ mode.

This section contains:

- "Inline Function Support Keyword (inline)" on page 1-62
- "Inline Assembly Language Support Keyword (asm)" on page 1-63
- "Dual Memory Support Keywords (pm dm)" on page 1-78
- "Placement Support Keyword (section)" on page 1-83
- "Boolean Type Support Keywords (bool, true, false)" on page 1-84
- "Pointer Class Support Keyword (restrict)" on page 1-84
- "Variable Length Array Support" on page 1-85
- "Non-Constant Aggregate Initializer Support" on page 1-87
- "Indexed Initializer Support" on page 1-87
- "Aggregate Constructor Expression Support" on page 1-89
- "Fractional Type Support" on page 1-90
- "Preprocessor Generated Warnings" on page 1-93
- "C++ Style Comments" on page 1-94
- "Compiler Built-In Functions" on page 1-94
- "ETSI Support" on page 1-106
- "Pragmas" on page 1-119
- "GCC Compatibility Extensions" on page 1-138

The additional keywords that are part of these C/C++ extensions do not conflict with any ISO/ANSI C/C++ keywords. The formal definitions of these extension keywords are prefixed with a leading double underscore (__). Unless the `-no-extra-keywords` command-line switch is used, the compiler defines the shorter forms of the keyword extension that omits the leading underscores. See "-extra-keywords" on page 1-26 for more information.

This section describes only the shorter forms of the keyword extensions, but in most cases you can use either form in your code. For example, all references to the `inline` keyword in this text appear without the leading double underscores, but you can use `inline` or `__inline` interchangeably in your code.

You might need to use the longer forms (such as `__inline`) exclusively if you are porting a program that uses the extra Analog Devices keywords as identifiers. For example, a program might declare local variables such as `pm` or `dm`. In this case, you should use the `-no-extra-keywords` switch, and if you need to declare a function as inline, or allocate variables to memory spaces, you can use `__inline` or `__pm`/`__dm` respectively.

Table 1-7 provides a list and a brief description of keyword extensions. Table 1-8 provides a list and a brief description of operational extensions. Both tables direct you to sections of this chapter that document each extension in more detail.

Table 1-7. Keyword Extensions

Keyword extensions	Description
`inline (function)`	Directs the compiler to integrate the function code into the code of the callers.For more information, see "Inline Function Support Keyword (inline)" on page 1-62.
`dm`	Specifies the location of a static or global variable or qualifies a pointer declaration "*" as referring to Data Memory. For more information, see "Dual Memory Support Keywords (pm dm)" on page 1-78.

Table 1-7. Keyword Extensions (Cont'd)

Keyword extensions	Description
`pm`	Specifies the location of a static or global variable or qualifies a pointer declaration "*" as referring to Program Memory. For more information, see "Dual Memory Support Keywords (pm dm)" on page 1-78.
`section("string")`	Specifies the section in which an object or function is placed. For more information, see "Placement Support Keyword (section)" on page 1-83.
`bool, true, false`	A Boolean type. For more information, see "Boolean Type Support Keywords (bool, true, false)" on page 1-84.
`restrict` keyword	Specifies restricted pointer features. For more information, see "Pointer Class Support Keyword (restrict)" on page 1-84.

Table 1-8. Operational Extensions

Operation extensions	Description
Variable length arrays	Support for variable length arrays lets you use automatic arrays whose length is not known until runtime. For more information, see "Variable Length Array Support" on page 1-85.
Non-constant initializers	Support for non-constant initializers lets you use non-constants as elements of aggregate initializers for automatic variables. For more information, see "Non-Constant Aggregate Initializer Support" on page 1-87.
Indexed initializers	Support for indexed initializers lets you specify elements of an aggregate initializer in an arbitrary order. For more information, see "Indexed Initializer Support" on page 1-87.
Preprocessor generated warnings	Support for generating warning messages from the preprocessor. For more information, see "Preprocessor Generated Warnings" on page 1-93.
C++-style comments	Support for C++-style comments in C programs. For more information, see "C++ Style Comments" on page 1-94.

Inline Function Support Keyword (inline)

The `inline` keyword directs `cc219x` to integrate the code for the function you declare as `inline` into the code of its callers. Inline function support and the `inline` keyword is a standard feature of C++; the compiler provides it as a C extension. Using this keyword eliminates the function-call overhead and therefore can increase the speed of your program's execution. Argument values that are constant and that have known values may permit simplifications at compile time.

The following example shows a function definition that uses the `inline` keyword.

```
inline int max3 (int a, int b int c) {
   return max (a, max(b, c));
}
```

A function declared `inline` must be defined (its body must be included) in every file in which the function is used. The normal way to do this is to place the inline definition in a header file. Usually, it will also be declared static.

In some cases, the compiler does not output object code for the function; for example, the address is not needed for an `inline` function called only from within the defining program. However, recursive calls, and functions whose addresses are explicitly referred to by the program, are compiled to assembly code.

ⓘ The compiler only inlines functions, even those declared using the inline keyword, when optimizations are enabled (using the `-O` switches, as described on page 1-37).

Inline Assembly Language Support Keyword (asm)

The `cc219x asm()` construct allows you to code ADSP-219x assembly language instructions within a C or C++ function and to pass declarations and directives through to the assembler. The `asm()` construct is useful for expressing assembly language statements that cannot be expressed easily or efficiently with C constructs.

The `asm()` keyword allows you code complete assembly language instructions or you can specify the operands of the instruction using C expressions. When specifying operands with a C expression, you do not need to know which registers or memory locations contain C variables.

The C compiler *does not analyze* code defined with the `asm()` construct; it passes this code directly to the assembler. The compiler *does* perform substitutions for operands of the formats `%0` through `%9`. However, it passes *everything else* through to the assembler without reading or analyzing it.

The `asm()` constructs are executable statements, and as such, may not appear before declarations within C/C++ functions.

A simplified `asm()` construct without operands takes the form of:

```
asm(" ENA INT;");
```

The complete assembly language instruction, enclosed in quotes, is the argument to `asm()`.

The compiler generates a label before and after inline assembly instructions when generating debug code (`-g` switch). These labels are used to generate the debug line information used by the debugger. If the inline assembler inserts conditionally assembled code, an undefined symbol error is likely to occur at link time. If the inline assembler changes the current section and thereby causes the compiler labels to be placed in another section, such as a data section (instead of the default code section), then the debug line information will be incorrect for these lines.

Using `asm()` constructs with operands requires some additional syntax. The construct syntax is described in:

- "Assembly Construct Template" on page 1-64
- "Assembly Construct Operand Description" on page 1-68
- "Assembly Constructs with Multiple Instructions" on page 1-74
- "Assembly Construct Reordering and Optimization" on page 1-74
- "Assembly Constructs with Input and Output Operands" on page 1-75
- "Assembly Constructs and Macros" on page 1-77

Assembly Construct Template

Using `asm()` constructs, you can specify the operands of the assembly instruction using C expressions. You do not need to know which registers or memory locations contain C variables.

ASM() Construct Syntax:

Use the following general syntax for your `asm()` constructs.

```
   asm(
template
   [:[constraint(output operand)[,constraint(output operand)…]]
      [:[constraint(input operand)[,constraint(input operand)…]]
          [:clobber]]]
);
```

The syntax elements are defined as:

template

- The template is a string containing the assembly instruction(s) with `%number` indicating where the compiler should substitute the operands. Operands are numbered in order of occurrence from left to right, starting at 0. Separate multiple instructions with a semicolon; then enclose the entire string within double quotes.

  For more information on templates containing multiple instructions, see "Assembly Constructs with Multiple Instructions" on page 1-74.

constraint

The constraint is a string that directs the compiler to use certain groups of registers for the input and output operands. Enclose the constraint string within double quotes. For more information on operand constraints, see "Assembly Construct Operand Description" on page 1-68.

output operand

The output operand is the name of a C or C++ variable that receives output from a corresponding operand in the assembly instruction.

input operand

The input operand is a C/C++ expression that provides an input to a corresponding operand in the assembly instruction.

clobber

The clobber notifies the compiler that a list of registers are overwritten by the assembly instructions. Use lowercase characters to name clobbered registers. Enclose each name within double quotes, and separate each quoted register name with a comma. The input and output operands are guaranteed not to use any of the clobbered registers, so you can read and write the clobbered registers as often as you like. See Table 1-10 on page 1-73.

ASM() Construct Syntax Rules

These rules apply to assembly construct template syntax:

- The template is the only mandatory argument to `asm()`. All other arguments are optional.
- An operand constraint string followed by a C/C++ expression in parentheses describes each operand. For output operands, it must be possible to assign to the expression—that is, the expression must be legal on the left side of an assignment statement.
- A colon separates:
  - The template from the first output operand
  - The last output operand from the first input operand
  - The last input operand from the clobbered registers
- A space must be added between adjacent colon field delimiters in order to avoid a clash with the C++ "::" reserved global resolution operator.
- A comma separates operands and registers within arguments.
- The number of operands in arguments must match the number of operands in your template.
- The maximum permissible number of operands is ten (`%0`, `%1`, `%2`, `%3`, `%4`, `%5`, `%6`, `%7`, `%8`, and `%9`).

The compiler cannot check whether the operands have data types that are reasonable for the instruction being executed. The compiler does not parse the assembler instruction template, does not interpret the template, and does not verify whether the template contains valid input for the assembler.

ASM() Construct Template Example

The following example shows how to apply the `asm()` construct template to the ADSP-219x assembly language `abs` instruction:

```
{
int x, result;

   asm ("%0=abs %1;" :
      “=c" (result) :
      “c" (x));
}
```

In the previous example, note the following points:

- The template is “`%0=abs %1;`”. The `%0` is replaced with operand zero `(result)`, the first operand. The `%1` is replaced with operand one `(x)`.
- The output operand is the C/C++ variable `result`. The letter `c` is the *operand constraint* for the variable. This constrains the output to an `ALU` result register. The compiler generates code to copy the output from the register to the variable result, if necessary. The “`=`” in `=c` indicates that the operand is an output.
- The input operand is the C/C++ variable `x`. The letter `c` is the *operand constraint* for the variable. This constrains `x` to an `ALU` register. If `x` is stored in different kinds of registers or in memory, the compiler generates code to copy the values into an register before the `asm()` construct uses them.

Assembly Construct Operand Description

The second and third arguments to the `asm()` construct describe the operands in the assembly language template. There are several pieces of information that need to be conveyed for the compiler to know how to assign registers to operands. This information is conveyed with an operand constraint. The compiler needs to know what kind of registers the assembly instructions can operate on, so it can allocate the correct register type.

You convey this information with a letter in the operand constraint string which describes the class of allowable registers. Table 1-9 on page 1-72 describes the correspondence between constraint letters and register classes.

> The use of any letter not listed in Table 1-9 results in unspecified behavior. The compiler does not check the validity of the code by using the constraint letter.
>
> For example, if your assembly template contains "`ax1 = dm(%0 += m3);`" and the address you want to load from is in the variable `p`, the compiler needs to know that it should put `p` in a `DAG1 I` register (`I0-I3`) before it generates your instruction. You convey this information to `cc219x` by specifying the operand "`w`" `(p)` where "`w`" is the constraint letter for `DAG1 I` registers.

To assign registers to the operands, the compiler must also be told which operands in an assembly language instruction are inputs, which are outputs, and which outputs may not overlap inputs. The compiler is told this in three ways.

- The output operand list appears as the first argument after the assembly language template. The list is separated from the assembly language template with a colon. The input operands are separated from the output operands with a colon and always follow the output operands.

- The operand constraints describe which registers are modified by an assembly language instruction. The = in `=constraint` indicates that the operand is an output; all output operand constraints must use =.
- The compiler may allocate an output operand in the same register as an unrelated input operand, unless the output operand has the `&=` constraint modifier. This situation can occur because the compiler assumes that the inputs are consumed before the outputs are produced.

This assumption may be false if the assembler code actually consists of more than one instruction. In such a case, use `&=` for each output operand that must not overlap an input or supply an "&" for the input operand.

Operand constraints indicate what kind of operand they describe by means of preceding symbols. The possible preceding symbols are: no symbol, `=`, `+`, `&`, `?`, and `#`.

- (no symbol)

  The operand is an input. It must appear as part of the third argument to the `asm()` construct. The allocated register will be loaded with the value of the C/C++ expression before the `asm()` template is executed. Its C/C++ expression will not be modified by the `asm()`, and its value may be a constant or literal. Example: `d`

- **= symbol**

  The operand is an output. It must appear as part of the second argument to the `asm()` construct. Once the `asm()` template has been executed, the value in the allocated register is stored into the location indicated by its C/C++ expression; therefore, the

expression must be one that would be valid as the left-hand side of an assignment.
Example: `=d`

- **+ symbol**

  The operand is both an input and an output. It must appear as part of the second argument to the `asm()` construct. The allocated register is loaded with the C/C++ expression value, the `asm()` template is executed, and then the allocated register's new value is stored back into the C/C++ expression. Therefore, as with pure outputs, the C/C++ expression must be one that is valid on the left-hand side of an assignment.
  Example: `+d`

- **? symbol**

  The operand is temporary. It must appear as part of the third argument to the `asm()` construct. A register is allocated as working space for the duration of the `asm()` template execution. The register's initial value is undefined, and the register's final value is discarded. The corresponding C/C++ expression is not loaded into the register, but must be present. This expression is normally specified using a literal zero.
  Example: `?d`

- **& symbol**

  This operand constraint may be applied to inputs and outputs. It indicates that the register allocated to the input (or output) may not be one of the registers that are allocated to the outputs (or inputs). This operand constraint is used when one or more output registers are set while one or more inputs are still to

be referenced. (This situation sometimes occurs if the `asm()` template contains more than one instruction.)
Example: `&d`

- **# symbol**

  The operand is an input, but the register's value is clobbered by the `asm()` template execution. The compiler may make no assumptions about the register's final value. The operand must appear as part of the second argument to the `asm()` construct.
  Example: `#d`

Table 1-9 on page 1-72 lists the registers that may be allocated for each register constraint letter. The use of any letter not listed in the "Constraint" column of this table results in unspecified behavior. The compiler does not check the validity of the code by using the constraint letter. Table 1-10 on page 1-73 lists the registers that may be named as part of the clobber list.

It is also possible to claim registers directly, instead of requesting a register from a certain class using the constraint letters. You can claim the registers directly by simply naming the register in the location where the class letter would be. The register names are the same as those used to specify the clobber list; see Table 1-10.

For example,

```
asm("%0 = %1 + %2;"
:"=ar"(sum)     /* output */
:"g"(x),"G"(y)  /* input  */
);
```

would load `x` into `ALU-X` register, `y` into `ALU-Y` register, and `sum` will be calculated in register `AR`.

Table 1-9. ASM() Operand Constraints

Constraint[1]	Description	Registers
b	input xregs to MAC	MX1, MX0, SR1, SR0, MR1, MR0, AR
B	input yregs to MAC	MY1, MY0
c	results from ALU	AR
C	result from int multiplies	MR0
cc	Used in the clobber list to tell the compiler that condition codes have been clobbered	ASTAT
d	input xregs to SHIFTER	SI, SR1, SR0, MR1, MR0, AX0, AY0, AX1, AY1, MX0, MX1, MY0, MY1, AR
D	result from shift	SR1
e	data registers, size 16	SI, AX1, AX0, MX1, MX0, MY0, MY1, AY1, AY0, MR1, MR0, SR1, SR0, AR
f	shift amount	SE
g	ALU X registers	AX1 AX0 AR SR1 SR0 MR1 MR0
G	ALU Y registers	AY1 AY0
memory	Used in the clobber list to tell the compiler that the asm() statement writes to memory	
r	all registers	SR1, SR0, SI, MY1, MX1, AY1, AX1, MY0, MX0, AY0, AX0, MR1, MR0, AR, I0-I7, M0-M7, L0-L7
u	DAG1 L registers	L0-L3
v	DAG2 L registers	L4-L7
w	DAG1 I registers	I0-I3
x	DAG1 M registers	M0-M3

Table 1-9. ASM() Operand Constraints (Cont'd)

Constraint[1]	Description	Registers
y	DAG2 I registers	I4-I7
z	DAG2 M registers	M4-M7
=&constraint	Indicates that the constraint is applied to an output operand that may not overlap an input operand	
=constraint	Indicates that the constraint is applied to an output operand	
&constraint	Indicates the constraint is applied to an input operand that may not be overlapped with an output operand	
=&constraint	Indicates the constraint is applied to an output operand that may not overlap an input operand	
?constraint	Indicates the constraint is temporary	
+constraint	Indicates the constraint is both an input and output operand	

1 The use of any letter not listed in the table results in unspecified behavior. The compiler does not check the validity of the code by using the constraint letter.

Table 1-10. Register Names for asm() Constructs

Clobber String	Meaning
"AX1", "AX0", "AY1", "AY0", "AR", "AF"	ALU registers
"MX1", "MX0", "MY1", "MY0", "MR1", "MR0", "MR2"	MAC registers
"SI", "SE", "SR1", "SR0", "SB", "SR2"	Shifter registers
"I0", "I1", "I2", "I3", "I6", "I7"	DAG addressing registers
"M0", "M1", "M2", "M3", "M4", "M6", "M7"	Modifier registers
"B0", "B1", "B2", "B3", "B4", "B5", "B6", "B7"	Base registers
"DMPG1", "DMPG2"	Page registers
"L0", "L1", "L2", "L3", "L5", "L6", "L7"	Length register
"PX"	PMD-DMD bus exchange register

Table 1-10. Register Names for asm() Constructs (Cont'd)

Clobber String	Meaning
`"astat"`	ALU status register
`"MSTAT", "MMODE", "SSTAT"`	mode control register
`"IMASK", "ICNTL", "IFC"`	Interrupt registers
`"CNTR", "STACKA", "STACKP"`	Program sequencer register
`"IJPG", "IOPG",`	Paging/overlay register; ADSP-2190 DSP only
`"cc"`	Condition code register
`"memory"`	Unspecified memory location(s)

Assembly Constructs with Multiple Instructions

There can be many assembly instructions in one template. If the `asm()` string is longer than one line, you may continue it on the next line by placing a backslash (\) at the end of the line or by quoting each line separately.

This is an example of multiple instructions in a template:

```
asm ("se=exp %1 (hi); \
   "sr=norm %1 (hi);  \
   "%0=sr0;"
   : "=e" (normalized)             // output
   : "e" (inval) ;                 // input
   : "se", "sr2", "sr1", "sr0") ;  // clobbers
```

Assembly Construct Reordering and Optimization

For the purpose of optimization, the compiler assumes that the side effects of an `asm()` construct are limited to changes in the output operands or the items specified using the clobber specifiers. This does not mean that you cannot use instructions with side effects, but you must be careful to notify the compiler that you are using them by using the clobber specifiers (see Table 1-10 on page 1-73).

The compiler may eliminate supplied assembly instructions if the output operands are not used, move them out of loops, or replace two with one if they constitute a common subexpression. Also, if the instruction has a side effect on a variable that otherwise appears not to change, the old value of the variable may be reused later if it happens to be found in a register.

Use the keyword `volatile` to prevent an `asm()` instruction from being moved, combined, or deleted. For example:

```
#define IOwrite(val,addr) \
asm volatile ("si="#val";IO("#addr")=si;": : :"si");
```

A sequence of `asm volatile()` constructs is not guaranteed to be completely consecutive; it may be moved across jump instructions or in other ways that are not significant to the compiler. To force the compiler to keep the output consecutive, use only one `asm volatile()` construct, or use the output of the `asm()` construct in a C/C++ statement.

Assembly Constructs with Input and Output Operands

The assembly constructs' output operands must be write only; cc219x assumes that the values in these operands do not need to be preserved. When the assembler instruction has an operand that is both read from and written to, you must logically split its function into two separate operands: one input operand and one write-only output operand. The connection between them is expressed by constraints that say they need to be in the same location when the instruction executes.

You can use the same C expression for both operands, or different expressions. For example, in the following statement, the `modify` instruction uses `sock` as its read only source operand and `shoe` as its read-write destination:

```
/* (pseudo code) modify (shoe += sock); */
asm ("modify (%0 += %2);":"=w"(shoe):"0"(shoe),"x"(sock));
```

The constraint `"0"` for operand 1 says that it must occupy the same location as operand `0`. A digit in an operand constraint is allowed only in an input operand, and it must refer to an output operand.

Only a digit in the constraint can guarantee that one operand is in the same place as another operand. Just because a variable (for example `shoe` in the code that follows) is used for more than one operand does not guarantee that the operands are in the same place in the generated assembler code.

```
/* Do NOT try to control placement with operand names; use
the % digit. The following code might NOT work.*/
asm ("modify (%0 += %2);":"=w"(shoe):"w"(shoe),"x"(sock));
```

In some cases, operands `0` and `1` could be stored in different registers due to reloading or optimizations.

Be aware that `asm()` does not support input operands that are used as both read operands and write operands. The example below shows a *dangerous* use of such an operand. In this example, `my_variable` is modified during the `asm()` operation. The compiler only knows that the output, `result_asm`, has changed. Subsequent use of `my_variable` after the `asm()` instruction may yield incorrect results since those values may have been modified during the `asm()` instruction and may not have been restored.

```
int result_asm;
int *my_variable;
/* NOT recommended */
/* (pseudo code) result_asm = dm(*my_variable += 3); */
/* asm() operation changes value of my_variable */

asm("%0=DM(%1 += 3);":"=e"(result_asm):"w"(my_variable));
```

Assembly Constructs and Macros

One way to use asm() constructs is to encapsulate them in macros that look like functions. For example, the following shows macros that contain asm() constructs. This code defines a macro, abs_macro(), which uses the inline asm() instruction to perform an assembly-language abs operation of variable x_var, putting the result in result_var:

```
#define abs_macro(result,x) \
asm("%0=abs %1;":"=c" (result):"c"(x))

/* (pseudo code) result = abs x */

main(){
   int result_var=0;
   int x_var=10;

   abs_macro(result_var, 10);
   /* or */
   abs_macro(result_var, x_var);
}
```

Assembly Constructs and Flow Control

It is inadvisable to place flow control operations within an asm() construct that "leaves" the asm() construct, such as calling a procedure or performing a jump, to another piece of code that is not within the asm() construct itself. Such operations are invisible to the compiler and may violate assumptions made by the compiler.

For example, the compiler is careful to adhere to the calling conventions for preserved registers when making a procedure call. If an asm() construct calls a procedure, the asm() construct must also ensure that all conventions are obeyed, or the called procedure may corrupt the state used by the function containing the asm() construct.

Dual Memory Support Keywords (pm dm)

This section describes `cc219x` keyword extensions to C and C++ that support the dual-memory space, modified Harvard architecture of the ADSP-219x processors. There are two keywords used to designate memory space: `dm` and `pm`. These keywords can be used to specify the location of a static or global variable or to qualify a pointer declaration.

These keywords allow you to control placement of data in primary (`dm`) or secondary (`pm`) data memory. No data is placed in the memory unit that holds programs. The following rules apply to dual memory support keywords:

- A memory space keyword (`dm` or `pm`) refers to the expression to its right.
- You can specify a memory space for each level of pointer. This corresponds to one memory space for each `*` in the declaration.
- The compiler uses Data Memory as the default memory space for all variables. All undeclared spaces for data are Data Memory spaces.
- The compiler always uses Program Memory as the memory space for functions. Function pointers always point to Program Memory.
- You cannot assign memory spaces to automatic variables. All automatic variables reside on the stack, which is always in Data Memory.
- Literal character strings always reside in Data Memory.
- Although program memory on the ADSP-219x DSPs consists of 24-bit words, only 16 bits of each word are used when C or C++ data is stored in `pm`. (This is normally the case for assembly lan-

guage programming as well.) If you need special access to all 24 bits, you should use an assembly language subroutine and work with the PX register.

The following listing shows examples of dual memory keyword syntax.

```
int pm abc[100];
/* declares an array abc with 100 elements in Program Memory */
int dm def[100];
/* declares an array def with 100 elements in Data Memory    */
int ghi[100];
/* declares an array ghi with 100 elements in Data Memory    */
int pm * pm pp;
/* declares pp to be a pointer which resides in Program Memory
   and points to a Program Memory integer                    */
int dm * dm dd;
/* declares dd to be a pointer which resides in primary Data
   Memory and points to a Data Memory integer                */
int *dd;
/* declares dd to be a pointer which resides in Data Memory
   and points to a Data Memory integer                        */
int pm * dm dp;
/* declares dp to be a pointer which resides in Data Memory
   and points to a Program Memory integer                     */
int pm * dp;
/* declares dp to be a pointer which resides in Data Memory
   and points to a Program Memory integer                     */
int dm * pm pd;
/* declares pd to be a pointer which resides in pm (secondary
   Data Memory) and points to a Data Memory integer          */
int * pm pd;
/* declares pd to be a pointer which resides in Program memory
   and points to a Data Memory integer                        */
float pm * dm * pm fp;
/* the first pm means that *fp is in Program Memory,
   the following dm puts *fp in Data Memory, and fp
   itself is in Program Memory                               */
```

Memory space specification keywords cannot qualify type names and structure tags, but you can use them in pointer declarations. The following listing shows examples of memory space specification keywords in typedef and struct statements.

```
/* Dual Memory Support Keyword typedef & struct Examples */
typedef float pm * PFLOATP;
/* PFLOATP defines a type which is a pointer to a        */
/* float which resides in pm.                            */

struct s {int x; int y; int z;};
static pm struct s mystruct={10,9,8};
/* Note that the pm specification is not used in        */
/* the structure definition. The pm specification        */
/* is used when defining the variable mystruct           */
```

Memory Keywords and Assignments/Type Conversions

Memory space specifications limit the kinds of assignments your program can make:

- You may make assignments between variables allocated in different memory spaces.
- Pointers to program memory must always point to `PM`. Pointers to data memory must always point to `DM`. You may not mix addresses from different memory spaces within one expression. Do not attempt to explicitly cast one type of pointer to another.

The following listings show a code segment with variables in different memory spaces being assigned and a code segment with illegal mixing of memory space assignments.

```
/* Legal Dual Memory Space Variable Assignment Example */
int pm x;
int dm y;
x = y;      /* Legal code */

/* Illegal Dual Memory Space Type Cast Example */
```

```
int pm *x;
int dm *y;
int dm a;
x = y;      /* Compiler will flag error */
x = &a;     /* Compiler will flag error */
```

Memory Keywords and Function Declarations/Pointers

Functions always reside in program memory. Pointers to functions always point to `PM`. The following listing shows some example function declarations with pointers.

```
/* Dual Memory Support Keyword Function Declaration (With Point-
ers) Syntax Examples */

int * y();        /* function y resides in */
                  /* pm and returns a      */
                  /* pointer to an integer */
                  /* which resides in dm   */

int pm * y();     /* function y resides in */
                  /* pm and returns a      */
                  /* pointer to an integer */
                  /* which resides in pm   */

int dm * y();     /* function y resides in */
                  /* pm and returns a      */
                  /* pointer to an integer */
                  /* which resides in dm   */

int * pm * y();   /* function y resides in */
                  /* pm and returns a      */
                  /* pointer to a pointer  */
                  /* residing in pm that   */
                  /* points to an integer  */
                  /* which resides in dm   */
```

Memory Keywords and Function Arguments

The compiler checks calls to prototyped functions for memory space specifications consistent with the function prototype. The following example shows sample code that `cc219x` flags as inconsistent use of memory spaces between a function prototype and a call to the function.

```
/* Illegal Dual Memory Support Keywords & Calls To Prototyped
Functions */

extern int foo(int pm*);
/* declare function foo() which expects a pointer to an int
residing in pm as its argument and which returns an int     */

int x;   /* define int x in dm                                */

foo(&x); /* call function foo()                               */
         /* using pm pointer (location of x) as the           */
         /* argument. cc219x FLAGS AS AN ERROR; this is an    */
         /* inconsistency between the function's              */
         /* declared memory space argument and function       */
         /* call memory space argument                        */
```

Memory Keywords and Macros

Using macros when making memory space specification for variables or pointers can make your code easier to maintain. If you must change the definition of a variable or pointer (moving it to another memory space), declarations that depend on the definition may need to be changed to ensure consistency between different declarations of the same variable or pointer.

To make changes of this kind easier, you can use C preprocessor macros to define common memory spaces that must be coordinated. The following listing shows two code segments that are equivalent after preprocessing.

The following code segment demonstrates how you can redefine the memory space specifications by redefining the macros SPACE1 and SPACE2.

```
/* Dual Memory Support Keywords & Macros */
   #define SPACE1 pm
   #define SPACE2 dm

   char pm * foo (char dm *)  char SPACE1 * foo (char SPACE2 *)
   char pm *x;   char SPACE1 *x;
   char dm y;    char SPACE2 y;

   x = foo(&y);  x = foo(&y);
```

PM and DM Compiler Support for Standard C Library Functions

There are a number of functions defined in the standard C library that take pointer input parameter types. These functions, which include for example `strlen()`, are implemented differently when the pointer input is to program memory (PM) or data memory (DM). The different implementations are called automatically by the compiler because it has specific in-built knowledge about the standard C functions that require pointer parameters. The support requires that the normal standard header file, for example `string.h`, is included prior to use of the function requiring PM and DM variants. The default library function variants are DM should the include file not be used.

Placement Support Keyword (section)

The `section` keyword directs the compiler to place an object or function in an assembly `.SECTION`, in the compiler's intermediate assembly output file. You name the assembly `.SECTION` with `section()`'s string literal parameter. If you do not specify a `section()` for an object or function declaration, the compiler uses a default `section`. The `.LDF` file supplied to the linker must also be updated to support the additional named sections.

Applying `section()` is only meaningful when the data item is something that the compiler can place in the named section.

Apply `section()` only to top-level, named objects that have static duration (they are explicitly `static`) or are given as external-object definitions. The example shows the declaration of a `static` variable that is placed in the section called `bingo`.

```
static section("bingo") int x;
```

Boolean Type Support Keywords (bool, true, false)

The `bool`, `true`, and `false` keywords are extensions to ANSI C that support the C++ Boolean type. The `bool` keyword is a unique signed integral type. There are two built-in constants of this type: `true` and `false`. When converting a numeric or pointer value to `bool`, a zero value becomes `false`; a nonzero value becomes `true`. A `bool` value may be converted to `int` by promotion, taking `true` to one and `false` to zero. A numeric or pointer value is automatically converted to `bool` when needed.

These keywords behave more or less as if the declaration that follows had appeared at the beginning of the file, except that assigning a nonzero integer to a `bool` type always causes it to take on the value `true`.

```
typedef enum { false, true } bool;
```

Pointer Class Support Keyword (restrict)

The `restrict` operator keyword is an extension that supports restricted pointer features. The use of `restrict` is limited to the declaration of a pointer and specifies that the pointer provides exclusive initial access to the object to which it points. More simply, `restrict` is a way that you can identify that a pointer does not create an alias. Also, two different restricted pointers can not designate the same object and therefore are not aliases. The compiler is free to use the information about restricted pointers and aliasing in order to better optimize C or C++ code that uses pointers.

The restrict keyword is most useful when applied to function parameters about which the compiler would otherwise have little information.

```
void fir (short *in,short *c,short *restrict out,int n)
```

The behavior of a program is undefined if it contains an assignment between two restricted pointers except for the following cases:

- A function with a restricted pointer parameter may be called with an argument that is a restricted pointer.
- A function may return the value of a restricted pointer that is local to the function, and the return value may then be assigned to another restricted pointer.

If your program uses a restricted pointer in a way that it does not uniquely refer to storage, then the behavior of the program is undefined.

Variable Length Array Support

The compiler supports variable-length automatic arrays. Unlike other automatic arrays, variable-length ones are declared with a non-constant length. This means that the space is allocated when the array is declared, and deallocated when the brace-level is exited.

The compiler does not allow jumping into the brace-level of the array and produces a compile time error message if this is attempted. The compiler does allow breaking or jumping out of the brace-level, and it deallocates the array when this occurs.

You can use variable-length arrays as function arguments, as shown in the following example.

```
struct entry
   var_array (int array_len, char data[array_len][array_len])
   {
   ...
}
```

The compiler calculates the length of an array at the time of allocation. It then remembers the array length until the brace-level is exited and can return it as the result of the `sizeof()` function performed on the array.

Because variable length arrays must be stored on the stack, it is impossible to have variable length arrays in Program Memory (`pm`). The compiler issues an error if an attempt is made to use a variable length array in `pm`.

As an example, if you were to implement a routine for computation of a product of three matrices, you need to allocate a temporary matrix of the same size as input matrices. Declaring automatic variable size matrix is much easier then explicitly allocating it in a heap.

The expression declares an array with a size that is computed at run time. The length of the array is computed on entry to the block and saved in case `sizeof()` is applied to the array. For multidimensional arrays, the boundaries are also saved for address computation. After leaving the block all the space allocated for the array and size information will be deallocated.

For example, the following program prints 40, not 50:

```
#include <stdio.h>
void foo(int);

main ()
{
   foo(40);
}
```

```
void foo (int n)
()
{
   char c[n];
   n = 50;
      printf("%d", sizeof(c));
}
```

Non-Constant Aggregate Initializer Support

The compiler includes extended support for aggregate initializers. The compiler does not require the elements of an aggregate initializer for an automatic variable to be constant expressions. The following example shows an initializer with elements that vary at run time:

```
void initializer (float a, float b)
{
   float the_array[2] = { a-b, a+b };
}
```

All automatic structures can be initialized by arbitrary expressions involving literals, previously declared variables and functions.

Indexed Initializer Support

ISO/ANSI Standard C and C++ requires the elements of an initializer to appear in a fixed order, the same as the order of the elements in the array or structure being initialized. The `cc219x` compiler, by comparison, supports labeling elements for array initializers. This feature lets you specify array or structure elements in any order by specifying the array indices or structure field names to which they apply. All index values must be constant expressions, even in automatic arrays.

For an array initializer, the syntax `[INDEX]` appearing before an initializer element value specifies the index to be initialized by that value. Subsequent initializer elements are then applied to sequentially following elements of the array, unless another use of the [INDEX] syntax appears. The index values must be constant expressions, even if the array being initialized is automatic.

The following example shows equivalent array initializers—the first initializer is in ISO/ANSI standard C/C++; the second initializer uses the cc219x compiler.

The [index] precedes the value being assigned to that element.

```
/* Example 1 Standard & cc219x C/C++ Array Initializer */
/* Standard Array Initializer */

int a[6] = { 0, 0, 115, 0, 29, 0 };

/* equivalent cc219x C/C++ array initializer */

int a[6] = { [2] 115, [4] 29 };
```

You can combine this technique of naming elements with standard C/C++ initialization of successive elements. The standard and cc219x instructions below are equivalent. Note that any unlabeled initial value is assigned to the next consecutive element of the structure or array.

```
/* Example 2 Standard & cc219x C/C++ Array Initializer */
/* Standard Array Initializer */

int a[6] = { 0, v1, v2, 0, v4, 0 };

/* equivalent cc219x C/C++ array initializer that uses
indexed elements */

int a[6] = { [1] v1, v2, [4] v4 };
```

The following example shows how to label the array initializer elements when the indices are characters or an enum type.

```
/* Example 3 Array Initializer With enum Type Indices */
/* cc219x C/C++ array initializer */

int whitespace[256] =
{
   [' '] 1, ['\t'] 1, ['\v'] 1, ['\f'] 1, ['\n'] 1, ['\r'] 1
};
```

In a structure initializer, specify the name of a field to initialize with *field name* before the element value. The standard C/C++ and `cc219x` C/C++ struct initializers in the example below are equivalent.

```
/* Example 4 Standard & cc219x C/C++ struct Initializer */
/* Standard struct Initializer */

struct point {int x, y;};
struct point p = {xvalue, yvalue};

/* Equivalent cc219x C/C++ struct Initializer With
Labeled Elements */

struct point {int x, y;};
struct point p = {y: yvalue, x: xvalue};
```

Aggregate Constructor Expression Support

Extended initializer support in `cc219x` C/C++ includes support for aggregate constructor expressions. These expressions enable you to assign values to large structure types without requiring each element's value to be individually assigned.

The following example shows an ISO/ANSI standard `struct` usage followed by equivalent `cc219x` code that has been simplified using an constructor expression.

```
/* Standard struct & cc219x C/C++ Constructor struct */
/* Standard struct */

struct foo {int a; char b[2];};
struct foo make_foo(int x, char *s)
{
   struct foo temp;
   temp.a = x;
   temp.b[0] = s[0];
   if (s[0] !=  '\0')
      temp.b[1] = s[1];
   else
      temp.b[1] = '\0';
   return temp;
}

/* Equivalent cc219x C/C++ constructor struct */

struct foo make_foo(int x, char *s)
{
   return((struct foo) {x, {s[0], s[0] ? s[1] : '\0'}});
}
```

Fractional Type Support

While in C++ mode, the `cc219x` compiler supports fractional (fixed-point) arithmetic that provides a way of computing with non-integral values within the confines of the fixed-point representation. The representation on which the fractional support is based is that of a 16-bit integral type. In this release of the compiler, there is no underlying C++ support for 32-bit fractional arithmetic. The ADSP-219x processors provide hardware support for the 16-bit fractional arithmetic.

Fractional values are declared with the `fract` data type. Ensure that your program includes the `fract` header file. `fract` is a C++ class that supports a set of standard arithmetic operators used in arithmetic expressions. Fractional values are represented as signed values in a range of [`-1` … `1`] with a binary point immediately after the sign bit. Other value ranges are obtained by scaling or shifting. In addition to the arithmetic, assignment, and shift operations, `fract` provides several type-conversion operations.

For more information about supported fractional arithmetic operators, see "Fractional Arithmetic Operations" on page 1-92. For sample programs demonstrating the use of the `fract` type, see Listing 1-1 on page 1-177, Listing 1-2 on page 1-178, and Listing 1-3 on page 1-178.

> The current release of the software does not provide for automatic scaling of fractional values.

Format of Fractional Literals

Fractional literals use the floating-point representation with an "`r`" suffix to distinguish them from floating-point literals, for example, `0.5r`. The cc219x compiler validates fractional literal values at run time to ensure they reside within the valid range of values.

Fractional literals are written with the "`r`" suffix to avoid certain precision loss. Literals without an "`r`" are of the type `double`, and are implicitly converted to `fract` as needed.

Conversions Involving Fractional Values

The following notes apply to type-conversion operations:

- Conversion between a fractional value and a floating value is supported. The conversion to the floating-point type may result in some precision loss.
- Conversion between a fractional value and an integer value is supported. The conversion is not recommended because the only common values are `0` and `-1`.

Fractional Arithmetic Operations

The following notes summarize information about fractional arithmetic operators supported by the cc219x compiler:

- Standard arithmetic operations on two fract items include addition, subtraction, and multiplication.
- Assignment operations include +=, -=, and *=.
- Shift operations include left and right shifts. A left shift is implemented as a logical shift and a right shift is an arithmetic shift. Shifting left by a negative amount is not recommended.
- Comparison operations are supported between two fract items.
- Mixed-mode arithmetic has a preference for fract. For more information about the mixed-mode arithmetic, see on page 1-92.
- Multiplication of a fractional and an integer produces an integer result or a fractional result. The program context determines which variant is generated following the conversion algorithm of C++. When the compiler does not have enough context, it generates an ambiguous operator message. For example,

  ```
  error:more than one operator "*" matches these operands:
  ```

  If this error occurs, cast the result of the multiply to the desired type.

Mixed Mode Operations

Most operations supported for fractional values, are supported for mixed fractional/float or fractional/double arithmetic expressions. At run time, a floating-point value is converted to a fractional value, and the operation is completed using fractional arithmetic.

The assignment operations, such as +=, are the exception to the rule. The logic of an assignment operation is defined by the type of a variable positioned on the left side of the expression.

Floating-point operations require an explicit cast of a fractional value to the desired floating type.

Saturated Arithmetic

The `cc219x` compiler supports saturated arithmetic for fractional data in the saturated arithmetic mode.

Whenever a calculation results in a bigger value than the `fract` data type represents, the result is truncated (wrapped around). An overflow flag is set to warn the program that the value has exceeded its limits. To prevent the overflow and to get the result as the maximum representable value when processing signal data, use saturated arithmetic. Saturated arithmetic forces an overflowed value to become the maximum representable value.

The mode is set to be saturated or default with the `set_saturate_mode()` and `reset_saturate_mode()` functions. Each arithmetic operator has its corresponding variant effected in the saturated mode. For example, `add_sat`, `sub_sat`, `neg_sat`, ….

Preprocessor Generated Warnings

The preprocessor directive `#warning` causes the preprocessor to generate a warning and continue preprocessing. The text on the remainder of the line that follows `#warning` is used as the warning message.

C++ Style Comments

The compiler accepts C++ style comments, beginning with `//` and ending at the end of the line, in C programs. This is essentially compatible with standard C, except for the following case.

```
a = b
//* highly unusual */ c
```

which a standard C compiler processes as:

```
a = b/c;
```

Compiler Built-In Functions

The compiler supports intrinsic functions that enable efficient use of hardware resources. Knowledge of these functions is built into the `cc219x` compiler. Your program uses them via normal function call syntax. The compiler notices the invocation and generates one or more machine instructions, just as it does for normal operators, such as `+` and `*`.

Built-in functions have names which begin with `__builtin_`. Note that identifiers beginning with double underlines (__) are reserved by the C standard, so these names will not conflict with user program identifiers.

The header files also define more readable names for the built-in functions without the `__builtin_` prefix. These additional names are disabled if the `-no-builtin` switch is used (on page 1-34).

The `cc219x` compiler provides built-in versions of some of the C library functions as described in "Using the Compiler's Built-In C Library Functions" on page 3-5.

The `sysreg.h` header file defines a set of functions that provide efficient system access to registers, modes and addresses not normally accessible from C source. These functions are specific to individual architectures and this section lists the built-in functions supported at this time on ADSP-219x DSPs.

The compiler supports:

- "Access to System Registers" on page 1-95
- "I/O Space Read or Write" on page 1-97
- "Interrupt Control" on page 1-98
- "Mode Control" on page 1-99
- "Near and Far Type Qualifiers" on page 1-99
- "Circular Buffer Built-In Functions" on page 1-103

Access to System Registers

The inclusion of `sysreg.h` allows the use of functions that will generate efficient inline instructions to implement read and write of values from and to general register set and system control set system registers.

General Register set:

```
ASTAT SSTAT MSTAT ICNTL IMASK IRPTL DMPG1 DMPG2 IOPG
```

System Control Register set:

```
B0 B1 B2 B3 B4 B5 B6 B7 SYSCTL CACTL

DBGCTRL DBGSTAT CNT0 CNT1 CNT2 CNT3
```

Also any 8-bit value used as a register identifier.

The prototypes for these functions are, as defined in `sysreg.h`:

```
void sysreg_write(const int sysreg, const int value);
int sysreg_read(const int sysreg);
```

The `sysreg` parameter for these functions can be a member of the SysReg enumeration defined in `sysreg.h`. This enumerations is used to map the actual registers to a small constant defined as a user friendly name.

The SysReg enumeration has the following definitions:

```
/* General Register set */
sysreg_ASTAT  = 0x0,  // ASTAT register - arithmetic status
sysreg_SSTAT  = 0x1,  // SSTAT register - shifter status
sysreg_MSTAT  = 0x2,  // MSTAT register - multiplier status
sysreg_ICNTL  = 0x3,  // ICNTL register - interrupt control
sysreg_IMASK  = 0x4,  // IMASK register - interrupts enabled mask
sysreg_IRPTL  = 0x5,  // Interrupt Latch register
sysreg_DMPG1  = 0x6,  // DMPG1 high address register
sysreg_DMPG2  = 0x7,  // DMPG2 high address register
sysreg_IOPG   = 0x8,  // IOPG I/O page register
/* System Control Register set */
sysreg_B0     = 0x9,  // B0 base register
sysreg_B1     = 0xa,  // B1 base register
sysreg_B2     = 0xb,  // B2 base register
sysreg_B3     = 0xc,  // B3 base register
sysreg_B4     = 0xd,  // B4 base register
sysreg_B5     = 0xe,  // B5 base register
sysreg_B6     = 0xf,  // B6 base register
sysreg_B7     = 0x10, // B7 base register
sysreg_SYSCTL = 0x11, // SYSCTL register
sysreg_CACTL  = 0x12, // Cache Control register
```

The `sysreg` parameter can also be any 8-bit value used to represent a system control register, since each variant of the ADSP-219x DSPs may have a differing System Control Register set. The compiler does not validate the `sysreg` parameter, instead it relies on the assembler to fault erroneous values.

An example use of `sysreg_read` to get the value of `IMASK` might be:

```
#include <sysreg.h>
int read_imask(){
   int value = sysreg_read(sysreg_IMASK);
   return value;
}
```

An example use of `sysreg_write` to set the value of `IMASK` might be:

```
#include <sysreg.h>
void write_imask(int val8bit) {
   sysreg_write(sysreg_IMASK, val8bit);
}
```

I/O Space Read or Write

The inclusion of `sysreg.h` allows the use of functions that will generate efficient inline instructions to implement read and write of values from and to I/O space addresses.

The prototypes for these functions are, as defined in `sysreg.h`:

```
void io_space_write(const unsigned int, const unsigned int);
int io_space_read(const unsigned int addr);
```

These functions are described in "io_space_read" on page 3-67 and "io_space_write" on page 3-69, respectively.

An example use of `io_space_read` to read from address zero might be:

```
#include <sysreg.h>
int read_io_zero(){
   int value = io_space_read(0);
   return value;
}
```

An example use of `io_space_write` to write to address zero might be:

```
#include <sysreg.h>
void write_io_zero(int val) {
   io_space_write(0, val);
}
```

Interrupt Control

The inclusion of `sysreg.h` allows the use of functions that generate the instructions to enable and disable interrupts.

The prototypes for these functions are, as defined in `sysreg.h`:

```
void enable_interrupts(void);
void disable_interrupts(void);
```

The following code provides an example of the use of `enable_interrupts` and `disable_interrupts` to disable and enable interrupts around a call to `printf`:

```
#include <sysreg.h>
#include <stdio.h>
void interrupt_safe_iprint(int val) {
     disable_interrupts();
     printf("%d\n",val);
     enable_interrupts();
}
```

Mode Control

The inclusion of `sysreg.h` allows the use of a function that generates the instructions to enable and disable a series of modes using the zero latency mode control instructions.

The prototype for this function is, as defined in `sysreg.h`:

```
void mode_change(const int _mode_spec);
```

The `_mode_spec` parameter is a bitmask of mode definitions defined in `sysreg.h`. These definitions are:

```
__MODE_ENA_AV_LATCH   =0x1,
__MODE_ENA_AR_SAT     =0x2,
__MODE_ENA_M_MODE     =0x4,
__MODE_ENA_TIMER      =0x8,
__MODE_DIS_AV_LATCH   =0x100,
__MODE_DIS_AR_SAT     =0x200,
__MODE_DIS_M_MODE     =0x400,
__MODE_DIS_INT        =0x1000,
```

Near and Far Type Qualifiers

The ADSP-219x processors can have external memory which will not by default, for reasons of efficiency, be addressable in 16-bits from C/C++ source. The compiler provides an extension to support access to external memory which allows use of external memory in C applications without degrading performance when accessing internal memory. This extension is enabled using a C type qualifiers, "`far`" and "`near`".

Declarations

This extension is enabled using a "`far`" and "`near`" type qualifier in C/C++ variable declarations. The following are example uses of the "`far`" and "`near`"type qualifiers.

```
int near * near nn1;   /* declares nn1 to be a pointer which
resides in near memory and points to a near memory integer */

int * nn2;             /* declares nn2 to be a pointer which
resides in near memory and points to a near memory integer */

int far * near nf1;   /* declares nf1 to be a pointer which
resides in near memory and points to a far memory integer */

int far * nf2;        /* declares nf2 to be a pointer which
resides in near memory and points to a far memory integer */

section("some_far_section")

int far * far ff;     /* declares ff to be a pointer which
resides in far memory and points to a far memory integer */
```

The “`near`” qualifier is a default and may be omitted when declaring types for variable which will normally reside in internal memory. The compiler will treat accesses to “`near`” qualified variable types exactly as if the “`near`” qualifier had not been used.

The “`far`” qualifier will normally be used in conjunction with the section placement specifier on static or global scope variables. Use of “`far`” with automatic declarations will result in a compiler warning indication that the “`far`” qualifier is unnecessary.

The compiler will issue an error if both “`near`” and “`far`” are used in a conflicting way.

Sizes of Far and Near Qualified Types

The size of a “`far`” pointer is 32 bits (the actual address width required, 24 bits, cannot be stored on the data stack). The size of “`near`” qualified data is as normal given that “`near`” will be ignored. The `sizeof` operator will return 1 for `near` pointers and 2 for `far` pointers. For example,

```
#include <assert.h>
```

```
section("some_far_section")
int far abc[100];

int near def[100];

main() {
   assert(sizeof(&abc)==2 && sizeof(&def)==1);
}
```

Conversions Between Far and Near Pointers

Conversions between `far` and near pointers are valid. The compiler will truncate the Most Significant Word (MSW) of a `far` pointer when converting a `far` pointer to a `near` pointer. The conversion from a near-pointer will be done by making the MSW of the output far-pointer the same as that of the internal data area.

Care should be taken to avoid de-referencing truncated `far` pointers as the data accessed will be in the internal memory area. The compiler will generate a warning when it sees a conversion from a `far` pointer to a `near` pointer.

For example,

```
int near * near_iptr;
int far * far_iptr;

void trunc_far (int far *i) {
near_iptr = i;    /* truncate MSW leaving LSW -
                     will cause a compiler warning */
}

void extend_far (int near *i) {
   far_iptr = i;   /* extend MSW with value of internal memory
                     addresses */
}
```

Addressing “Far” Data

The result of the address operator, “&”, when used to generate the address of data which is declared as being in “far” memory will be a far pointer type.

C++ Function Overloading

The compiler, in C++ mode, will allow function overloading of parameters types for different “far” and “near” qualified parameters and return types.

Library Support for “Far" Pointers

All standard library functions that normally have pointer parameters currently only support near pointer parameters. Passing a far pointer parameter to a standard library function will cause the far pointer to be truncated. The compiler will generate a warning when this occurs assuming the correct standard include file has been included before the call to the library.

Legacy Support

Previous releases of VisualDSP++ only supported external memory through compiler intrinsics which had to be passed the full address to be loaded or stored. These intrinsics are still supported although superseded in terms of functionality and flexibility by use of “far” qualified types.

The inclusion of sysreg.h allows the use of functions that will generate inline instructions to implement read and write of values from and to external memory.

The prototypes of the functions are defined in sysreg.h:

```
int external_memory_read(int DMPG_val, int* addr);
void external_memory_write(int DMPG_val, int* addr, int val);
```

```
int __builtin_external_vol_memory_read(int,
    volatile void* addr);
void __builtin_external_vol_memory_write(int,
     volatile void* addr, int);
```

The `DMPG_val` parameter is the value of the upper 8 bits of the 24-bit external memory address. The `addr` parameter is a pointer to the external memory. The `val` parameter to `external_memory_write` is the value to be written to external memory. The variants where `addr` is a volatile pointer should be used when the accessed memory is changed by interrupts of peripherals in a way that may not be visible to the compiler. For example,

```
#include <sysreg.h>
section("external_memory_section")
static int GlobalTable[256];

int main() {
   int page, read_value, value_to_write = 0;
   asm("%0 = PAGE(GlobalTable); " : "=e"(page): : );
   external_memory_write(page, &GlobalTable[0], value_to_write);
   read_value = external_memory_read(page, &GlobalTable[1]);
   return read_value;
}
```

Circular Buffer Built-In Functions

The C/C++ compiler provides the following built-in functions for using the ADSP-219x processor's circular buffer mechanisms. You should include the `builtins.h` file before using these functions. Failure to do so leads to unresolved symbols at link time.

Automatic Circular Buffer Generation

If optimization is enabled, the compiler will automatically attempt to use circular buffer mechanisms where appropriate. For example,

```
void func(int *array, int n, int incr)
{
```

```
    int i;
    for (i = 0; i < n; i++)
         array[i % 10] += incr;
}
```

The compiler will recognize that the "`array[i % 10]`" expression is a circular reference, and will use a circular buffer if possible.

There are cases where the compiler will not be able to verify that the memory access is always within the bounds of the buffer. The compiler is conservative in such cases, and does not generate circular buffer accesses. The compiler can be instructed to still generate circular buffer accesses even in such cases, by specifying "-force-circbuf" on page 1-27.

The compiler also provides built-in functions which can explicitly generate circular buffer accesses, subject to available hardware resources. The built-in functions provide circular indexing, and circular pointer references. Both built-in functions are defined in the `builtins.h` header file.

Circular Buffer Increment of an Index

The following operation performs a circular buffer increment of an index.

```
int __builtin_circindex(int index, int incr, unsigned int nitems);
```

The operation is equivalent to:

```
index += incr;
if (index < 0)
   index += nitems;
else if (index >= nitems)
   index -= nitems;
```

An example of this built-in function is:

```
void func(int *array, int n, int incr, int len)
{
      int i, idx = 0;
      for (i = 0; i < n; i++) {
```

```
            array[idx] += incr;
            idx = __builtin_circindex(idx, incr, len);
        }
    }
```

Circular Buffer Increment of a Pointer

The following operation performs a circular buffer increment of a pointer.

```
void *__builtin_circptr(const void *ptr,
                        unsigned size_t incr,
                        const void * base,
                          size_t buflen);
```

Both *incr* and *buflen* are specified in bytes, since the operation deals in void pointers.

The operation is equivalent to:

```
ptr += incr;
if (ptr < base)
    ptr += buflen;
else if (ptr >= (base+buflen))
    ptr -= buflen;
```

An example of this built-in function is:

```
void func(int *array, int n, int incr, int len)
{
      int i, idx = 0;
      int *ptr = array;

      // scale increment and length by size
      // of item pointed to.
      incr *= sizeof(*ptr);
      len *= sizeof(*ptr);

      for (i = 0; i < n; i++) {
           *ptr += incr;
           ptr = __builtin_circptr(ptr, incr, array, len);
      }
}
```

ETSI Support

The ETSI (European Telecommunications Standards Institute) support for ADSP-219x processors is a collection of functions that provides high performance implementations for operations commonly required by DSP applications. These operations provided by the ETSI library (`libetsi.dlb`) and compiler built-in functions (defined in `ETSI_fract_arith.h`) include support for fractional, or fixed-point, arithmetic. The results obtained from of use of these operations have well defined overflow and saturation conditions. The ETSI support operations are Analog Devices extensions to ANSI C standard.

The ETSI support contains functions that you can call from your source program. The following topics describe how to use this support.

- "ETSI Support Overview" on page 1-106
- "Calling ETSI Library Functions" on page 1-108
- "Using the ETSI Built-In Functions" on page 1-109
- "Linking ETSI Library Functions" on page 1-109
- "Working with ETSI Library Source Code" on page 1-110
- "ETSI Support for Data Types" on page 1-110
- "ETSI Header File" on page 1-111

ETSI Support Overview

The use of fractional arithmetic is vital for many applications on DSP processors as information can be held more compactly than in floating point. It would take 24 bits in floating-point format to match the precision of 16-bit fractional data. Also, control of normalization and precision is more complex with floating point. Many DSPs do not include hardware support for floating-point arithmetic and these operations are therefore very expensive in both code size and performance terms for such DSPs.

Fractional data has a representation similar to that of integers except that while an integer value is considered to have a decimal point to the right of the least significant bit, a fractional value is considered to have a decimal point to the left of the most significant bit. Fractional values are usually held in 16-bit or 32-bit "containers". In each case, signed values are in the range [-1.0, +1.0).

The bit operations on fractional data are identical to those on integer data, but there are three aspects of the result that are normally treated differently:

1. **MSB extraction.** Multiplication is a widening operation, thus multiplying a 16-bit value by another 16-bit value produces a 32-bit result. If a 16-bit integer result is required then this is taken to be the least significant 16 bits of the result, and the upper 16 bits are regarded as overflow. For a fractional operation the upper 16 bits would represent a 16-bit result, and the lower 16 bits would be regarded as an underflow.

2. **Duplicate sign bit elimination.** Following a multiplication of two 16-bit values the nature of the representation results in two "sign bits" in the result. For normal integer arithmetic this causes no problem, but for fractional arithmetic a shift left by one is required to normalize the result.

3. **Saturation.** If we perform an arithmetic operation that would cause us to overflow, it can be useful to return the maximum (appropriately signed) number that can be represented in the result register. The alternatives which include firing an interrupt, saying the result is undefined and is some other number, usually look less attractive to DSP programmers.

These fractional operations can often be done at no extra cost to normal integer operations on DSPs using special instructions or modes of operation.

The C programming language does not include a basic type for fractional data, and rather than introduce a non-standard type, Analog Devices defines `fract16`and`fract32` in terms of appropriately-sized integer data types and provides sets of basic intrinsic functions which perform the required operations. These look like library function calls, but are specially recognized by the compilers, which generate short sequences or single instructions, exploiting any specialized features, which may be available on the architecture. An important aspect of this is that the compiler optimizer is not inhibited in any way by the use of these intrinsics.

Because of the varying nature of the architectures the basic intrinsic functions just discussed cannot be standardized across all the architectures. However, a set of standard functions for manipulating fractional data has been defined by the ITU (International Telecommunications Union) and ETSI (European Telecommunications Standards Institute).

Referred to as the ETSI Standard Functions, these have been very widely used to implement Telecommunications packages such as GSM, EFR and AMR Vocoders, and have become a de-facto industry standard. These functions have been implemented on ADSP-219x DSPs.

The ETSI standard is aimed at DSP processors with 16-bit inputs, saturated arithmetic and 32-bit accumulators.

Calling ETSI Library Functions

To use an ETSI function, call the function by name and give the appropriate arguments. The names and arguments for each function appear on the function's reference page. The names and arguments for each function appear in the section "ETSI Header File" on page 1-111.

Like other functions you use, ETSI functions should be declared. Declarations are supplied in the header file `ETSI_fract_arith.h`, which must be included in any source files where ETSI functions are called. The function names are C function names. If you call C run-time library functions from an assembly language program, you must use the assembly version of the function name—prefix an underscore on the name. For more information on naming conventions, see the section "C/C++ and Assembly Language Interface" on page 1-169.

Several of the ETSI routines are provided with carry and overflow checking. Where overflow or carry occurs, the global variables `Carry` and `Overflow` will be set. It is your responsibility to reset these variables in between operations. The `Carry` and `Overflow` variables are represented by integers and are prototyped in the `libetsi.h` system header file. With the ETSI functions provided by Analog Devices, this can be switched off by compiling with `__NO_ETSI_FLAGS` defined in the compiler command line.

In fact, this is the default for the ADSP-219x DSP implementation. If the user wishes to keep track of these flags, for debugging purposes, they should compile with `__NO_ETSI_FLAGS` set to zero. This will mean that the user is using the functions in accordance with the ETSI standard, but this will result in a reduced performance.

Using the ETSI Built-In Functions

Some of the ETSI functions have been implemented as part of `cc219x` compiler's set of built-in functions. For information on how to use these functions, refer to "Compiler Built-In Functions" on page 1-94. These built-in implementations will be automatically defined when header file `ETSI_fract_arith.h` is included.

Linking ETSI Library Functions

When your C/C++ code calls an ETSI function that is not implemented using a compiler built-in, the call creates a reference that the linker resolves when linking. This requires the linker to be directed to link with the ETSI library, `libetsi.dlb`, in the `219x\lib` directory, which is a subdirectory of the VisualDSP++ installation directory. This is done automatically when using the default Linker Description File (LDF) for ADSP-219x processor targets, as these specify that `libetsi.dlb` will be on each link line.

If not using default .LDF files, then either add libetsi.dlbto the.LDF file which is being used, or alternatively use the compiler's -letsi switch to specify that libetsi.dlb is to be added to the link line.

Working with ETSI Library Source Code

The source code for functions and macros in the ETSI library is provided with your VisualDSP++ software. By default, the installation program copies the source code to a subdirectory of the directory where the run-time libraries are kept named 219x\lib\src\libetsi_src. Each function is kept in a separate file. The file name is the name of the function with the extension .asm. If you do not intend to modify any of the functions, you can delete this directory and its contents to conserve disk space.

The source code is provided so you can customize specific functions for your own needs.

To modify these files, you need proficiency in ADSP-219x assembly language and an understanding of the run-time environment, as explained in "C/C++ and Assembly Language Interface" on page 1-169.

Before you make any modifications to the source code, copy the source code to a file with a different file name and rename the function itself. Test the function before you use it in your system to verify that it is functionally correct.

Analog Devices only supports the run-time library functions as provided.

ETSI Support for Data Types

ETSI functions support fract16 and fract32 data types as follows:

- fract16 is a 16-bit fractional data type (1.15 format) having a range of [-1.0, +1.0). This is defined in the C/C++ language as

```
typedef short fract16
```

- `fract32` is a 32-bit fractional data type (1.31 format) having a range of [-1.0, +1.0). This is defined in the C/C++ language as

  ```
  typedef long fract32
  ```

ETSI Header File

The following are summary descriptions of the functions provided by the ETSI library, as defined in the header file `ETSI_fract_arith.h`.

Short absolute

```
fract16 abs_s (fract16)
```

This function returns the 16-bit value that is the absolute value of the input parameter. Where the input is `0x8000`, saturation occurs and `0x7fff` is returned.

Short add

```
fract16 add (fract16, fract16)
```

This function returns the 16-bit result of addition of the two `fract16` input parameters. Saturation occurs with the result being set to `0x7fff` for overflow and `0x8000` for underflow.

Short division

```
fract16 div_s (fract16, fract16)
```

This function returns the 16-bit result of the fractional integer division of `f1` by `f2`. `f1` and `f2` must both be positive fractional values with `f2` greater than `f1`.

Long division

```
fract16 div_l (fract32, fract16)
```

This function produces a result which is the fractional integer division of the first parameter by the second. Both inputs must be positive and the least significant word of the second parameter must be greater or equal to the first; the result is positive (leading bit equal to 0) and truncated to 16 bits.

Extract high (most significant 16 bits)

```
fract16 extract_h (fract32)
```

This function returns the 16 most significant bits if the 32-bit `fract` parameter provided.

Extract low (least significant 16 bits)

```
fract16 extract_l (fract32)
```

This function returns the 16 least significant bits of the 32-bit `fract` parameter provided.

Multiply and accumulate with rounding

```
fract16 mac_r (fract32, fract16, fract16)
```

This function performs an `L_mac` operation using the three parameters provided. The result is the rounded 16 most significant bits of the 32-bit results from the `L_mac` operation.

Multiply and subtract with rounding

```
fract16 msu_r (fract32, fract16, fract16)
```

This function performs an `L_msu` operation using the three parameters provided. The result is the rounded 16 most significant bits of the 32-bit result from the `L_msu` operation.

Short multiply

```
fract16 mult (fract16, fract16)
```

This function returns the 16-bit result of the fractional multiplication of the input parameters. The result is saturated.

Multiply with rounding

```
fract16 mult_r (fract16, fract16)
```

This function performs a 16-bit multiply with rounding of the result of the fractional multiplication of the two input parameters.

Short negate

```
fract16 negate (fract16)
```

This function returns the 16-bit result of the negation of the input parameter. If the input is `0x8000`, saturation occurs and `0x7fff` is returned.

Long normalize

```
fract16 norm_l (fract16)
```

This function returns the number of left shifts required to normalize the input variable for positive values on the interval with minimum of `0x40000000` and maximum of `0x7fffffff`, and for negative values on the interval with minimum of `0x80000000` and maximum of `0xc0000000`.

Short normalize

```
fract16 norm_s (fract16)
```

This function returns the number of left shifts required to normalize the input 16 bit variable for positive values on the interval with minimum of `0x4000` and maximum of `0x7fff`, and for negative values on the interval with minimum of `0x8000` and maximum of `0xc000`.

Round

```
fract16 round (fract32)
```

This function rounds the lower 16-bits of the 32-bit input parameter into the most significant 16 bits with saturation. The resulting bits are shifted right by 16.

Saturate

```
fract16 saturate (fract32)
```

This function returns the 16 most significant bits of the input parameter. If the input parameter is greater than `0x7fff`, `0x7fff` is returned. If the input parameter is less than `0x8000`, `0x8000` is returned.

Short shift left

```
fract16 shl (fract16, fract16)
```

This function arithmetically shifts the first parameter left by second parameter bits. The empty bits are zero filled. If second parameter is negative the operation shifts right.

Short shift right

```
fract16 shr (fract16, fract16)
```

This function arithmetically shifts the first parameter right by second parameter bits with sign extension. If second parameter is negative the operation shifts left.

Shift right with rounding

```
fract16 shr_r (fract16, fract16)
```

This function performs a shift to the right as per the `shr()` operation with additional rounding and saturation of the result.

Short subtract

```
fract16 sub (fract16, fract16)
```

This function returns the 16-bit result of the subtraction of the two parameters. Saturation occurs with the result being set to `0x7fff` for overflow and `0x8000` for underflow.

Long absolute

```
fract32 L_abs (fract32)
```

This function returns the 32-bit absolute value of the input parameter. In cases where the input is equal to `0x80000000`, saturation occurs and `0x7fffffff` is returned.

Long add

```
fract32 L_add (fract32, fract32)
```

This function returns the 32-bit saturated result of the addition of the two input parameters.

Long add with carry

```
fract32 L_add_c (fract32, fract32)
```

This function performs 32-bit addition of the two input parameters. Uses the `Carry` flag as additional input when using the ETSI flag variables.

16-bit variable -> most significant bits (least significant bits zeroed)

```
fract32 L_deposit_h (fract16)
```

This function deposits the 16-bit parameter into the 16 most significant bits of the 32-bit result. The least 16 bits are zeroed.

16-bit variable -> least significant bits (sign extended)

```
fract32 L_deposit_l (fract16)
```

This function deposits the 16-bit parameter into the 16 least significant bits of the 32-bit result. The most significant bits are set to sign extension for the input.

Multiply and accumulate

```
fract32 L_mac (fract32, fract16, fract16)
```

This function performs a fractional multiplication of the two 16-bit parameters and returns the saturated sum of the multiplication result with the 32-bit parameter.

Multiply and accumulate without saturation

```
fract32 L_macNs (fract32, fract16, fract16)
```

This function performs a non-saturating version of the `L_mac` operation.

Multiply both the most significant bits and the least significant bits of a long, by the same short

```
fract32 L_mls (fract32, fract16)
```

Multiply and subtract

```
fract32 L_msu (fract32, fract16, fract16)
```

This function performs a fractional multiplication of the two 16-bit parameters and returns the saturated subtraction of the multiplication result with the 32-bit parameter.

Multiply and subtract without saturation

```
fract32 L_msuNs (fract32, fract16, fract16)
```

This function performs a non-saturating version of the `L_msu` operation.

Long multiply

```
fract32 L_mult (fract16, fract16)
```

This function returns the 32-bit result of the fractional multiplication of the two 16-bit parameters.

Long negate

```
fract32 L_negate (fract32)
```

This function returns the 32-bit result of the negation of the parameter. Where the input parameter is `0x80000000` saturation occurs and `0x7fffffff` is returned.

Long saturation

```
fract32 L_sat (fract32)
```

The resultant variable is set to `0x80000000` if `Carry` and `Overflow` ETSI flags are set (underflow condition), else if `Overflow` is set, the resultant is set to `0x7fffffff`. The default revision of the library simply returns as no checking or setting of the `Overflow` and `Carry` flags is performed.

Long shift left

```
fract32 L_shl (fract32, fract16)
```

This function arithmetically shifts the 32-bit first parameter to the left by the value given in the 16-bit second parameter. The empty bits of the 32-bit result are zero filled.

If the second parameter is negative, the shift performed is to the right with sign-extended. The result is saturated in cases of overflow and underflow.

Long shift right

```
fract32 L_shr (fract32, fract16)
```

This function arithmetically shifts the 32-bit first parameter to the right by the value given in the 16-bit second parameter with sign extension. If the shifting value is negative, the source is shifted to the left. The result is saturated in cases of overflow and underflow.

Long shift right with rounding

```
fract32 L_shr_r (fract32, fract16)
```

This function performs the shift-right operation as per `L_shr` but with rounding.

Long subtract

```
fract32 L_sub (fract32, fract32)
```

This function returns the 32-bit saturated result of the subtraction of two 32-bit parameters (first-second).

Long subtract with carry

```
fract32 L_sub_c (fract32, fract32)
```

This function performs 32-bit subtraction of the two input parameters. Uses the `Carry` flag as additional input when using the ETSI flag variables.

Compose long

```
fract32 L_Comp (fract16, fract16)
```

This function composes a `fract32` type value from the given `fract16` high (first parameter) and low (second parameter) components. The sign is provided with the low half, the result is calculated to be:

```
high<<16 + low<<1
```

Multiply two longs

```
fract32 Mpy_32 (fract16, fract16, fract16, fract16)
```

This function performs the multiplication of two `fract32` type variables, provided as high and low half parameters. The result returned is calculated as:

```
Res = L_mult(hi1,hi2);
Res = L_mac(Res, mult(hi1,lo2),1);
Res = L_mac(Res, mult(lo1,hi2),1);
```

Multiply short by a long

```
fract32 Mpy_32_16 (fract16, fract16, fract16)
```

Extract a long from two shorts

```
void L_Extract(fract32 src,fract16 *hi,fract16 *lo)
```

This function extracts low and high halves of `fract32` type value into `fract16` variables pointed to by the parameters `hi` and `lo`. The values calculated are:

```
Hi = bit16 to bit31 of src
Lo = (src - hi<<16)>>1
```

Fract integer division of two longs

```
fract32 Div_32(fract32 L_num,fract16 denom_hi,fract16 denom_lo)
```

This is 32-bit fractional divide operation. The result returned is the `fract32` representation of `L_num` divided by `L_denom` (represented by `demon_hi` and `denom_lo`). `L_num` and `L_denom` must both be positive fractional values and `L_num` must be less that `L_denom` to ensure that the result falls within the fractional range.

Pragmas

The compiler supports a number of pragmas. Pragmas are implementation-specific directives that modify the compiler's behavior. There are two types of pragma usage: pragma directives and pragma operators.

Pragma directives have the following syntax:

```
#pragma pragma-directive pragma-directive-operands new-line
```

Pragma operators have the following syntax:

```
_Pragma ( string-literal )
```

When processing a pragma operator, the compiler effectively turns it into a pragma directive using a non-string version of *`string-literal`*. This means that the following pragma directive

```
#pragma linkage_name mylinkname
```

can also be equivalently be expressed using the following pragma operator

```
_Pragma ("linkage_name mylinkname")
```

The examples in this manual use the directive form.

The C compiler supports pragmas for:

- Arranging alignment of data
- Defining functions that can act as interrupt handlers
- Changing the optimization level, midway through a module
- Changing how an externally visible function is linked
- Header file configurations and properties
- Giving additional information about loop usage to improve optimizations

The following sections describe the pragmas that support these features.

- "Data Alignment Pragmas" on page 1-121
- "Interrupt Handler Pragmas" on page 1-122
- "Loop Optimization Pragmas" on page 1-123
- "General Optimization Pragmas" on page 1-125
- "Linking Control Pragmas" on page 1-126
- "Function Side-Effect Pragmas" on page 1-128
- "Template Instantiation Pragmas" on page 1-134
- "Header File Control Pragmas" on page 1-136

The compiler will issue a warning when it encounters an unrecognized pragma directive or pragma operator. The compiler will not expand any pre-processor macros used within any pragma directive or pragma operator.

Data Alignment Pragmas

The data alignment pragmas include `align` and `pad` pragmas. Alignments specified using these pragmas must be a power of two. The compiler will reject uses of those pragmas that specify alignments that are not powers of two.

#pragma align *num*

The `align` *`num`* pragma may be used before variable and field declarations. It applies to the variable or field declaration that immediately follows the pragma. Use of this pragma causes the compiler to generate the next variable or field declaration aligned on a boundary specified by `num`.

The `align` pragma is useful for declaring arrays that need to be on a circular boundary. Such arrays might be required to make use of a bit-reversal sorting algorithm that is implemented using the ADSP-219x processor's `DAG1` bit reversal mode.

```
#pragma align 256
int arr[128];
```

#pragma pad (*alignopt*)

The `#pragma pad (`*`alignopt`*`)` may be applied to `struct` definitions. It applies to `struct` definitions that follow, until the default alignment is restored by omitting *`alignopt`*, for example, by `#pragma pad()` with empty parentheses.

This pragma is effectively a shorthand for placing `#pragma align` before every field within the `struct` definition.

The following example shows how to use `#pragma pad()`.

```
#pragma pad(4)
struct {
   int i;
   int j;
} s = {1,2};
#pragma pad()
```

Interrupt Handler Pragmas

The interrupt pragma s include `interrupt` and `altregisters` pragmas.

#pragma interrupt

The `interrupt` pragma may be used before a function declaration or definition. It applies to the function declaration or definition that immediately follows the pragma. Use of this pragma causes the compiler to generate the function code so that it may be used as a self dispatching interrupt handler.

The compiler arranges for the function to save its context above and beyond the usual caller-preserved set of registers, and to restore the context upon exit. The function will return using a return from interrupt (RTI) instruction.

```
#pragma interrupt
void field_SIG()
{
/* ISR code */
}
```

#pragma altregisters

The `altregisters` pragma may be used in conjunction to the `interrupt` pragma to indicate that the compiler can optimize the saving and restoring of registers through use of the secondary register sets. Note the use of the `altregisters` pragma is not safe when nested interrupts are enabled.

For example,

```
#pragma interrupt
#pragma altregisters
void field_SIG()
{
/* ISR code */
}
```

Loop Optimization Pragmas

Loop optimization pragmas give the compiler additional information about usage within a particular loop, which allows the compiler to perform more aggressive optimization. The pragmas are placed before the loop statement, and apply to the statement that immediately follows, which must be a `for`, `while` or `do` statement to have effect. In general, it is most effective to apply loop pragmas to inner-most loops, since the compiler can achieve the most savings there.

The optimizer always attempts to vectorize loops when it is safe to do so. The optimizer exploits the information generated by the interprocedural analysis (see "Interprocedural Analysis" on page 1-57) to increase the cases where it knows it is safe to do so. Consider the following code:

```
void copy(short *a, short *b) {
     int i;
     for (i=0; i<100; i++)
         a[i] = b[i];
}
```

If you call `copy` with two calls, say `copy(x,y)` and later `copy(y,z)`, the interprocedural analysis will not be able to tell that "`a`" never aliases "`b`". Therefore, the optimizer cannot be sure that one iteration of the loop is not dependent on the data calculated by the previous iteration of the loop. If it is known that each iteration of the loop is not dependent on the previous iteration, then the `vector_for` pragma can be used to explicitly notify the compiler that this is the case.

#pragma loop_count(*min, max, modulo*)

The `loop_count(min, max, modulo)` pragma appears just before the loop it describes. It asserts that the loop will iterate at least `min` times, no more than `max` times, and a multiple of `modulo` times. This information enables the optimizer to omit loop guards, to decide whether the loop is worth completely unrolling, and whether code need be generated for odd iterations. The last two arguments can be omitted if they are unknown.

For example,

```
int i;
#pragma loop_count(24, 48, 8)
for (i=0; i < n; i++)
```

#pragma vector_for

The `#pragma vector_for` notifies the optimizer that it is safe to execute two iterations of the loop in parallel. The `vector_for` pragma does not force the compiler to vectorize the loop; the optimizer checks various properties of the loop and does not vectorize it if it believes it is unsafe or if it cannot deduce that the various properties necessary for the vectorization transformation are valid.

Strictly speaking, the pragma simply disables checking for loop-carried dependencies.

```
void copy(short *a, short *b) {
int i;
#pragma vector_for
   for (i=0; i<100; i++)
        a[i] = b[i];
}
```

In cases where vectorization is impossible (for example, if array `a` were aligned on a word boundary, but array `b` was not), the information given in the assertion made by `vector_for` may still be put to good use in aiding other optimizations.

#pragma no_alias

Use the #pragma no_alias to tell the compiler the following has no loads or stores that conflict due to references to the same location through different pointers, known as “aliases”. In this example,

```
void vadd(int *a, int *b, int *out, int n) {
   int i;
#pragma no_alias
   for (i=0; i < n; i++)
        out[i] = a[i] + b[i];
}
```

the use of #pragma no_alias just before the loop informs the compiler that the pointers a, b and out point to different arrays, so no load from b or a will be using the same address as any store to out. Therefore, a[i] or b[i] is never an alias for out[i].

Using the no_alias pragma can lead to better code because it allows the loads and stores to be reordered and any number of iterations to be performed concurrently, thus providing better software pipelining by the optimizer.

General Optimization Pragmas

There are three pragmas which can change the optimization level while a given module is being compiled. These pragmas must be used at global scope, immediately prior to a function definition. The pragmas do not just apply to the immediately-following function; they remain in effect until the end of the compilation, or until superceded by a following optimize_ pragma.

The pragmas are:

- **#pragma optimize_off**

  This pragma turns off the optimizer, if it was enabled. This pragma has no effect if IPA is enabled.

- **#pragma optimize_for_space**

  This pragma turns the optimizer back on, if it was disabled, or sets focus to give reduced code size a higher priority than high performance, where these conflict.

- **#pragma optimize_for_speed**

  This pragma turns the optimizer back on, if it was disabled, or sets focus to give high performance a higher priority than reduced code size, where these conflict.

- **#pragma optimize_as_cmd_line**

  This pragma resets the optimization settings to be those specified on the `cc219x` command line when the compiler was invoked.

The following shows example uses of these pragmas.

```
#pragma optimize_off
void non_op() { /* non-optimized code */ }

#pragma optimize_for_space
void op_for_si() { /* code optimized for size */ }

#pragma optimize_for_speed
void op_for_sp() { /* code optimized for speed */ }
/* subsequent functions declarations optimized for speed */
```

Linking Control Pragmas

Linking pragmas change how a given global function or variable is viewed during the linking stage. These pragmas are: `linkage_name`, `retain_name`, and `weak_entry`.

#pragma linkage_name *identifier*

The `#pragma linkage_name` associates the *`identifier`* with the next external function declaration. It ensures that identifier is used as the external reference, instead of following the compiler's usual conventions.

If identifier is not a valid function name, as could be used in normal function definitions, the compiler will generate an error. See also the `asm` keyword (described on page 1-145).

The following shows an example use of this pragma.

```
#pragma linkage_name realfuncname
void funcname ();
void func() {
   funcname();   /* compiler will generate a call to
                  realfuncname */
}
```

#pragma retain_name

The `#pragma retain_name` indicates that the external function or variable declaration that follows the pragma is not removed even though Interprocedural Analysis (IPA) sees that it is not used. Use this pragma for C functions that are only called from assembler routines, such as the startup code sequence invoked before `main()`.

The following example shows how to use this pragma.

```
int delete_me(int x) {
   return x-2;
}

#pragma retain_name
int keep_me(int y) {
   return y+2;
}

int main(void) {
   return 0;
}
```

Since the program has no uses of either `delete_me()` or `keep_me()`, the compiler will remove `delete_me()`, but will keep `keep_me()` because of the pragma. You do not need to specify `retain_name` for `main()`.

For more information on IPA, see "Interprocedural Analysis" on page 1-57.

#pragma weak_entry

The `#pragma weak_entry` may be used before a static variable or function declaration or definition. It applies to the function or variable declaration or definition that immediately follows the pragma. Use of this pragma causes the compiler to generate the function or variable definition with weak linkage.

The following are example uses of the `pragma weak_entry` directive.

```
#pragma weak_entry
int w_var = 0;

#pragma weak_entry
void w_func(){}
```

Function Side-Effect Pragmas

The function side-effect pragmas are used before a function declaration to give the compiler additional information about the function in order to enable it to improve the code surrounding the function call. These pragmas should be placed before a function declaration and apply to that function. For example,

```
#pragma pure
long dot(short*, short*, int);
```

The function side-effect pragmas are: `alloc`, `pure`, `const`, `regs_clobbered`, and `result_alignment`.

#pragma alloc

The `#pragma alloc` tells the compiler that the function behaves like the library function "`malloc`", returning a pointer to a newly allocated object. An important property of these functions is that the pointer returned by the function does not point at any other object in the context of the call.

In the example,

```
#pragma alloc
int *new_buf(void);
int *vmul(int *a, int *b) {
   int *out = new_buf();
   for (i = 0; i < N; ++i)
      out[i] = a[i] * b[i];
   return out;
}
```

the compiler can reorder the iterations of the loop because the `#pragma alloc` tells it that `a` and `b` cannot overlap `out`.

The GNU attribute `malloc` is also be supported with the same meaning.

#pragma pure

The `#pragma pure` tells the compiler that the function does not write to any global variables, and does not read or write any volatile variables. Its result, therefore, is a function of its parameters or of global variables. If any of the parameters are pointers the function may read the data they point at but it may not write it.

As this means the function call will have the same effect every time it is called, between assignments to global variables, the compiler need not generate the code for every call. Therefore, in this example,

```
#pragma pure
long sdot(short *, short *, int);

long tendots(short *a, short *b, int n) {
```

```
    int i;
    long s = 0;
  for (i = 1; i < 10; ++i)
      s += sdot(a, b, n);  // call can get hoisted out of loop
    return s;}
```

the compiler can replace the ten calls to `sdot` with a single call made before the loop.

#pragma const

The `#pragma const` is a more restrictive form of the `pure` pragma. It tells the compiler that the function does not read from global variables as well as not writing to them or reading or writing volatile variables. The result of the function is therefore a function of its parameters. If any of the parameters are pointers, the function may not even read the data they point at.

#pragma regs_clobbered *string*

The `#pragma regs_clobbered` *`string`* may be used with a function declaration or definition to specify which registers are modified (or clobbered) by that function. The *`string`* contains a list of registers and is case-insensitive.

When used with an external function declaration, this pragma acts as an assertion telling the compiler something it would not be able to discover for itself. In the example,

```
#pragma regs_clobbered ar m7
void f(void);
```

the compiler will know that only registers `ar` and `m7` may be modified by the call to `f`, so it may keep local variables in other registers across that call.

The `regs_clobbered` pragma may also be used with a function definition, or a declaration preceding a definition, when it acts as a command to the compiler to generate register saves and restores on entry and exit from the function to ensure it only modifies the registers in string. For example,

```
#pragma regs_clobbered "ar m7"
int g(int a) {
   return a+3;
}
```

The `regs_clobbered` pragma may not be used in conjunction with `#pragma interrupt`. If both are specified, a warning is issued and the `regs_clobbered` pragma is ignored.

To obtain best results with the pragma, it is best to restrict the clobbered set to be a subset of the default scratch registers. The compiler is likely to produce more efficient code this way than if the scratch set is changed to use the same number of registers but which does not make a subset of the default volatile set usually scratch.

When considering when to apply the `regs_clobbered` pragma, it may be useful to look at the output of the compiler to see how many scratch registers were used. Restricting the volatile set to these registers will produce no impact on the code produced for the function but may free up registers for the caller to allocate across the call site.

String Syntax

A `regs_clobbered` string consists of a list of registers, register ranges, or register sets that are clobbered. The list is separated by spaces, commas, or semicolons.

A *register* is a single register name, which is the same as that which may be used in an assembly file.

A *register range* consists of `start` and `end` registers which both reside in the same register class, separated by a hyphen. All registers between the two (inclusive) are clobbered.

A *register set* is a name for a specific set of commonly clobbered registers that is predefined by the compiler. The following register sets are defined,

Set	Registers
`CCset`	ASTAT, condition codes
`MR`	MR0 - MR2
`SR`	SR0 - SR2
`B1set`	DAG1 B-registers
`B2set`	DAG2 B-registers
`Bset`	B1set union B2set
`DAG1scratch`	Members of dag1 I, L, B and M-registers that are scratch by default
`DAG2scratch`	Members of dag2 I, L, B and M-registers that are scratch by default
`DAGscratch`	DAG1scratch union DAG2scratch
`Dscratch`	Members of D-registers that are scratch by default
`ALLscratch`	Entire default volatile set

When the compiler detects an illegal string, a warning is issued and the default volatile set as defined in this compiler manual is used instead.

Unclobberable and Must Clobber Registers
There are certain caveats as to what registers may or must be placed in the clobbered set.

On ADSP-219x processors, the registers `I4`, `I5`, `DMPG2`, `MSTAT`, and `M_MODE` may not be specified in the clobbered set, as the correct operation of the function call requires their value to be preserved. If the user specifies them in the clobbered set, a warning will be issued and they will be removed from the specified clobbered set.

The registers `AR` and `M7` are always clobbered. If the user specifies a function definition with the `regs_clobbered` pragma which does not contain these registers, a warning is issued and these registers are added to the clobbered set.

However, if the compiler sees an external function declaration with a `regs_clobbered` pragma that does not contain the `AR` and `M7` registers, a warning will not be issued because an assembly function may have been written which genuinely does not modify these registers.

Registers from these classes,

```
D, I, B, ASTAT, CNTR, PX, DMPG1, STACKA, STACKP, IJPG, IOPG, SB, SE
```

may be specified in clobbered set and code will be generated to save them as necessary.

The L-registers are required to be zero on entry and exit from a function. A user may specify that a function clobbers the L-registers. If it is a compiler-generated function, then it will in fact leave the L-registers as zero at the end of the function. If it is an assembly function, then it may clobber the L-registers. In that case, the L-registers are re-zeroed after any call to that function. The registers `M1`, `M2`, `M6` and `M7` have their required value set in an analogous manner.

The `IMASK`, `ICNTL`, `IRPTL`, `SSTAT`, `SYSCTL`, `CACTL`, `CCODE`, `LPSTACKA`, and `LPSTACKP` registers are never used by the compiler and are never preserved.

User Reserved Registers

User reserved registers will never be preserved in the function wrappers whether in the clobbered set or not.

Function Results

The registers in which a function returns its result must always be clobbered by the callee and retain their new value in the caller. They may appear in the clobbered set of the callee but it will make no difference to the generated code; the return register will not be saved and restored. Only the return register used by the particular function return type is special. Return registers used by different return types will be treated in the clobbered list in the convention way.

For example,

```
typedef struct { int x, int y } Point;
typedef struct { int x[10] } Big;
int f();   // Result in AX1. SR1, SR0 and I0 may be preserved
              across call.
Point g(); // Result in SR1 and SR0. AX1 and I0 may be
              preserved across call.
Big f();    // Result pointer in I0. AX1, SR1 and SR0 may be
              preserved across call
```

#pragma result_alignment (*n*)

The `#pragma result_alignment (n)` asserts that the pointer or integer returned by the function has a value that is a multiple of `n`.

This pragma is often used in conjunction with the `#pragma alloc` of custom allocation functions that return pointers that are more strictly aligned than be deduced from their type.

Template Instantiation Pragmas

The template instantiation pragmas give fine grain control over where (that is, in which object file) the individual instances of template functions, and member functions and static members of template classes are created. The creation of these instances from a template is known in C++ speak as instantiation. As templates are a feature of C++ these pragmas are only allowed in `-c++` mode.

These pragmas take the name of an instance as a parameter, as shown in Table 1-11

Table 1-11. Instance Names

Name	Parameter
a template class name	`A<int>`
a template class declaration	`class A<int>`
a member function name	`A<int>::f`
a static data member name	`A<int>::I`
a static data declaration	`int A<int>::I`
a member function declaration	`void A<int>::f(int, char)`
a template function declaration	`char* f(int, float)`

If instantiation pragmas are not used, the compiler will chose object files in which to instantiate all required instances automatically during the prelinking process.

#pragma instantiate *instance*

The `#pragma instantiate` *`instance`* requests the compiler to instantiate *`instance`* in the current compilation. For example,

```
#pragma instantiate class Stack<int>
```

will cause all static members and member functions for the int instance of a template class Stack to be instantiated, whether they are required in this compilation or not.

The example,

```
#pragma instantiate void Stack<int>::push(int)
```

will cause only the individual member function `Stack<int>::push(int)` to be instantiated.

#pragma do_not_instantiate *instance*

The `#pragma do_not_instantiate` *`instance`* directs the compiler not to instantiate *`instance`* in the current compilation. For example,

```
#pragma do_not_instantiate int Stack<float>::use_count
```

will prevent the compiler from instantiating the static data member `Stack<float>::use_count` in the current compilation.

#pragma can_instantiate *instance*

The `#pragma can_instantiate` *`instance`* tells the compiler that, if instance is required anywhere in the program, it should be instantiated in this compilation.

Currently, this pragma forces the instantiation even if it is not required anywhere in the program. Therefore, it has the same effect as `#pragma instantiate`.

Header File Control Pragmas

The header file control pragmas help the compiler to handle header files. These pragmas are `hdrstop`, `no_pch`, once, and `system_header`.

#pragma hdrstop

The `#pragma hdrstop` is used in conjunction with the `-pch` (precompiled header) switch (on page 1-40). The switch tells the compiler to look for a precompiled header (`.pch` file), and, if it cannot find one, to generate a file for use on a later compilation. The `.pch` file contains a snapshot of all the code preceding the header stop point.

By default, the header stop point is the first non-preprocessing token in the primary source file. The `#pragma hdrstop` can be used to set the point earlier in the source file.

In the example,

```
#include "standard_defs.h"
#include "common_data.h"
#include "frequently_changing_data.h"

int i;
```

the default header stop point is start of the declaration of `i`. This might not be a good choice, as in this example, “`frequently_changing_data.h`” might change frequently, causing the `.pch` file to be regenerated often, and, therefore, losing the benefit of precompiled headers.

The `hdrstop` pragma can be used to move the header stop to a more appropriate place. In this case,

```
#include "standard_defs.h"
#include "common_data.h"
#pragma hdrstop
#include "frequently_changing_data.h"

int i;
```

the precompiled header file would not include the contents of `frequently_changing_data.h`, as it is included after the `hdrstop` pragma. Therefore, the precompiled header file would not need to be regenerated each time `frequently_changing_data.h` was modified.

#pragma no_pch

The `#pragma no_pch` overrides the `-pch` (precomiled headers) switch (on page 1-40) for a particular source file. It directs the compiler not to look for a `.pch` file and not to generate one for the specified source file.

#pragma once

The `#pragma once`, which should appear at the beginning of a header file, tells the compiler that the header is written in such a way that including it several times has the same effect as including it once. For example,

```
#pragma once
#ifndef FILE_H
#define FILE_H
... contents of header file ...
#endif
```

In this example, the `#pragma once` is actually optional because the compiler recognizes the `#ifndef/#define/#endif` idiom and will not reopen a header that uses it.

#pragma system_header

The `#pragma system_header` identifies an `include` file as the file supplied with VisualDSP++. The pragma tells the compiler that every function and variable declared in the file (but not in files included in the file) is the variable or function with that name from the VDSP++ library.

The compiler will take advantage of any special knowledge it has of the behavior of the library.

GCC Compatibility Extensions

The compiler provides compatibility with the C dialect accepted by version 3.2 of the GNU C Compiler. Many of these features are available in the C99 ANSI Standard. A brief description of the extensions is included in this section. For more information, refer to the following web address:

`http://gcc.gnu.org/onlinedocs/gcc-3.2.1/gcc/C-Extensions.html#C%20Extensions`

Statement Expressions

A statement expression is a compound statement enclosed in parentheses. A compound statement itself is enclosed in braces { }, so this construct is enclosed in parentheses-brace pairs ({ }).

The value computed by a statement expression is the value of the last statement which should be an expression statement. The statement expression may be used where expressions of its result type may be used. But they are not allowed in constant expressions.

Statement expressions are useful in the definition of macros as they allow the declaration of variables local to the macro.

In the following example,

```
#define min(a,b) ({                    \
   short __x=(a),__y=(b),__res;   \
   if (__x > __y)                 \
      __res = __y;                \
   else                           \
   __res = __x;                   \
   __res;                         \
})

int use_min() {
   return  min(foo(), thing()) + 2;
}
```

The `foo()` and `thing()` statements get called once each because they are assigned to the variables `__x` and `__y` which are local to the statement expression that `min` expands to and `min()` can be used freely within a larger expression because it expands to an expression.

Labels local to a statement expression can be declared with the `__label__` keyword. For example,

```
({
   __label__ exit;
```

```
    int i;
    for (i=0; p[i]; ++i) {
       int d = get(p[i]);
       if (!check(d)) goto exit;
       process(d);
    }
exit:
    tot;
})
```

Statement expressions are not supported in C++ mode.

Statement expressions are an extension to C originally implemented in the GCC compiler. Analog Devices support the extension primarily to aid porting code written for that compiler. When writing new code consider using inline functions, which are compatible with ANSI/ISO standard C++ and C99, and are as efficient as macros when optimization is enabled.

Type Reference Support Keyword (Typeof)

The `typeof( expression )` construct can be used as a name for the type of expression without actually knowing what that type is. It is useful for making source code that is interpreted more than once such as macros or include files more generic.

The `typeof` keyword may be used where ever a `typedef` name is permitted such as in declarations and in casts.

This example,

```
#define abs(a) ({                         \
    typeof(a) __a = a;                    \
    if (__a < 0) __a = - __a;             \
    __a;                                  \
})
```

shows `typeof` used in conjunction with a statement expression to define a “generic” macro with a local variable declaration.

The argument to `typeof` may also be a type name. Because `typeof` itself is a type name, it may be used in another `typeof( type-name )` construct. This can be used to restructure the C type declaration syntax, such as:

```
#define pointer(T)   typeof(T *)
#define array(T, N)  typeof(T [N])

array (pointer (char), 4) y;
```

declares `y` to be an array of four pointers to `char`.

The `typeof` keyword is not supported in C++ mode.

The `typeof` keyword is an extension to C originally implemented in the GCC compiler. It should be used with caution because it is not compatible with other dialects of C or C++ and has not been adopted by the more recent C99 standard.

GCC Generalized Lvalues

A cast is an `lvalue` (may appear on the left hand side of an assignment) if its operand is an `lvalue`. This is an extension to C, provided for compatibility with GCC. It is not allowed in C++ mode.

A comma operator is an `lvalue` if its right operand is an `lvalue`. This is an extension to C, provided for compatibility with GCC. It is a standard feature of C++.

A conditional operator is an `lvalue` if its last two operands are `lvalues` of the same type. This is an extension to C, provided for compatibility with GCC. It is a standard feature of C++.

Conditional Expressions with Missing Operands

The middle operand of a conditional operator can be left out. If the condition is non-zero (true), then the condition itself is the result of the expression. This can be used for testing and substituting a different value when a pointer is NULL. The condition is only evaluated once; therefore, repeated side effects can be avoided. For example,

```
printf("name = %s\n", lookup(key)?:"-");
```

calls `lookup()` once, and substitutes the string "-" if it returns NULL. This is an extension to C, provided for compatibility with GCC. It is not allowed in C++ mode.

Hexadecimal Floating-Point Numbers

C99 style hexadecimal floating-point constants are accepted. They have the following syntax.

```
hexadecimal-floating-constant:
  {0x|0X} hex-significand binary-exponent-part [ floating-suffix ]
hex-significand: hex-digits [ . [ hex-digits ]]
binary-exponent-part: {p|P} [+|-] decimal-digits
floating-suffix: { f | l | F | L }
```

The hex-significand is interpreted as a hexadecimal rational number. the digit sequence in the exponent part is interpreted as a decimal integer. The exponent indicates the power of two by which the significand is to be scaled. The floating suffix has the same meaning it does for decimal floating constants: a constant with no suffix is of type `double`, a constant with suffix `F` is of type `float`, and a constant with suffix `L` is of type `long double`.

Hexadecimal floating constants enable the programmer to specify the exact bit pattern required for a floating-point constant. For example, the declaration

```
float f = 0x1p-126f;
```

causes f to be initialized with the value `0x800000`.

Hexadecimal floating constants are not supported in C++ mode.

Zero Length Arrays

Arrays may be declared with zero length. This is an anachronism supported to provide compatibility with GCC. Use variable length array members instead.

Variable Argument Macros

The final parameter in a macro declaration may be followed by ... to indicate the parameter stands for a variable number of arguments.

For example,

```
#define trace(msg, args...)  fprintf (stderr, msg, ## args);
```

can be used with differing numbers of arguments,

```
trace("got here\n");
trace("i = %d\n", i);
trace("x = %f, y = %f\n", x, y);
```

The `##` operator has a special meaning when used in a macro definition before the parameter that expands the variable number of arguments: if the parameter expands to nothing then it removes the preceding comma.

The variable argument macro syntax comes from GCC. It is not compatible with C99 variable argument macros and is not supported in C++ mode.

Line Breaks in String Literals

String literals may span many lines. The line breaks do not need to be escaped in any way. They are replaced by the character `\n` in the generated string. This extension is not supported in C++ mode. The extension is not compatible with many dialects of C including ANSI/ISO C89 and C99. However, it is useful in `asm` statements, which are intrinsically non-portable.

Arithmetic on Pointers to Void and Pointers to Functions

Addition and subtraction is allowed on pointers to `void` and pointers to functions. The result is as if the operands had been cast to pointers to `char`. The `sizeof` operator returns one for `void` and function types.

Cast to Union

A type cast can be used to create a value of a union type, by casting a value of one of the unions member's types.

Ranges in Case Labels

A consecutive range of values can be specified in a single case, by separating the first and last values of the range with .... For example,

```
case 200 ... 300:
```

Declarations mixed with Code

In C mode the compiler will accept declarations in the middle of code as in C99 and C++. This allows the declaration of local variables to be placed at the point where they are required. Therefore, the declaration can be combined with initialization of the variable.

For example, in the following function

```
void func(Key k) {
Node *p = list;
while (p && p->key != k)
   p = p->next;
if (!p)
   return;
Data *d = p->data;
while (*d)
   process(*d++);
}
```

the declaration of `d` is delayed until its initial value is available, so that no variable is uninitialized at any point in the function.

Escape Character Constant

The character escape '`\e`' may be used in character and string literals and maps to the ASCII Escape code, 27.

Alignment Inquiry Keyword (__alignof__)

The `__alignof__ (type-name)` construct evaluates to the alignment required for an object of a type. The `__alignof__` *`expression`* construct can also be used to give the alignment required for an object of the *`expression`* type.

If expression is an `lvalue` (may appear on the left hand side of an assignment), the alignment returned takes into account alignment requested by pragmas and the default variable allocation rules.

Keyword for Specifying Names in Generated Assembler (asm)

The `asm` keyword can be used to direct the compiler to use a different name for a global variable or function (see also "#pragma linkage_name identifier" on page 1-126).

For example,

```
int N asm("C11045");
```

tells the compiler to use the label C11045 in the assembly code it generates wherever it needs to access the source level variable `N`. By default the compiler would use the label `_N`.

The `asm` keyword can also be used in function declarations but not function definition. However, a definition preceded by a declaration has the desired effect.

For example,

```
extern int f(int, int) asm("func");

int f(int a, int b) {
. . .
}
```

Function, Variable and Type Attribute Keyword (__attribute__)

The `__attribute__` keyword can be used to specify attributes of functions, variables and types, as in these examples.

```
void func(void) __attribute__ ((section("fred")));

int a __attribute__ ((aligned (8)));

typedef struct {int a[4];} __attribute__((aligned (4))) Q;
```

The `__attribute__` keyword is supported, and therefore code, written for GCC, can be ported. All attributes accepted by GCC on `ix86` are accepted. The ones that are actually interpreted by the compiler are described in the sections of this manual describing the corresponding pragmas (see “Pragmas” on page 1-119).

Preprocessor Features

The cc219x compiler provides standard preprocessor functionality, as described in any C text. The following extensions to standard C are also supported:

```
// end of line (C++-style) comments
#warning directive
```

For more information about these extensions, refer to "Preprocessor Generated Warnings" on page 1-93 and "C++ Style Comments" on page 1-94.

This section contains:

- "Predefined Preprocessor Macros"
- "Header Files" on page 1-149
- "Writing Preprocessor Macros" on page 1-149
- "Preprocessing of .IDL Files" on page 1-151

Predefined Preprocessor Macros

The cc219x compiler defines a number of macros to produce information about the compiler, source file, and options specified. These macros can be tested, using the #ifdef and related directives, to support your program's needs. Similar tailoring is done in the system header files.

Macros such as __DATE__ can be useful to incorporate in text strings. The "#" operator with a macro body is useful in converting such symbols into text constructs.

Table 1-12 describes the predefined preprocessor macros.

Table 1-12. Predefined Macro Listing

Macro	Function
`__ADSP21XX__`	Always defines `__ADSP21XX__` and `__ADSP219X__` as 1
`__ADSP2191__`	Defines `__ADSP2191__` as 1 when you compile with the `-proc ADSP-2191` command-line switch
`__ADSP2192_12__`	Defines `__ADSP2192_12__` as 1 when you compile with the `-proc ADSP-2192-12` command-line switch
`__ANALOG_EXTENSIONS__`	Defines `__ANALOG_EXTENSIONS__` as 1, unless you compile with `-pedantic` or `-pedantic-errors`
`__cplusplus`	Defines `__cplusplus` as 199711L when compiling in C++ mode
`__DATE__`	The preprocessor expands this macro into the current date as a string constant. The date string constant takes the form `mm dd yyyy` (ANSI standard).
`__DOUBLES_ARE_FLOATS__`	Always defines `__DOUBLES_ARE_FLOATS__` as 1
`__ECC__`	Always defines `__ECC__` as 1
`__EDG__`	Always defines `__EDG__` as 1. This signifies that an Edison Design Group front end is being used
`__EDG_VERSION__`	Always defines `__EDG_VERSION__` as an integral value representing the version of the compiler's front end
`__FILE__`	The preprocessor expands this macro into the current input file name as a string constant. The string matches the name of the file specified on the `cc219x` command line or in a preprocessor `#include` command (ANSI standard).
`_LANGUAGE_C`	Always defines `_LANGUAGE_C` as 1 when compiling C or C++ source
`__LINE__`	The preprocessor expands the `__LINE__` macro into the current input line number as a decimal integer constant (ANSI standard)
`__NO_BUILTIN`	Defines `__NO_BUILTIN` as 1 when you compile with the `-no-builtin` command-line switch
`_NO_LONG_LONG`	Always defines `_NO_LONG_LONG` as 1 for C and C++ source. This definition signifies no support is present for the `long long int` type.

Table 1-12. Predefined Macro Listing (Cont'd)

Macro	Function
`__SIGNED_CHARS__`	Defines `__SIGNED_CHARS__` as 1 unless you compile with the `-unsigned-char` command-line switch
`__STDC__`	Defines `__STDC__` as 1
`__STDC_VERSION__`	Defines `__STDC_VERSION__` as 199409L
`__TIME__`	The preprocessor expands this macro into the current time as a string constant. The time string constant takes the form `hh:mm:ss` (ANSI standard).
`__VERSION__`	Defines `__VERSION__` as a string constant giving the version number of the compiler used to compile this module.

Header Files

A header file contains C or C++ declarations and macro definitions. Use the `#include` preprocessor directive to access header files for your program. Header file names have an `.h` extension. There are two main categories of header files:

- System header files declare the interfaces to the libraries supplied with the ADSP-219x product. Include them in your program for the definitions and declarations you need to access. Use angle brackets to indicate a system header file: `#include <file>`
- User header files contain declarations for interfaces between the source files of your program. Use double quotes to indicate a user header file: `#include "file"`

Writing Preprocessor Macros

A macro is a name of a block of text that the preprocessor substitutes. Use the `#define` preprocessor command to create a macro definition. When a macro definition has arguments, the block of text the preprocessor substitutes can vary with each new set of arguments.

Multi-statement Macros

Whenever possible, use inline functions rather than multi-statement macros. If multi-statement macros are necessary, define such macros to allow invocation like function calls. This will make your source code easier to read and maintain.

The following two code segments define two versions of the macro `SKIP_SPACES`.

```
/* SKIP_SPACES, regular macro */
#define SKIP_SPACES (p, limit)   \{
   char *lim = (limit); \
      while ((p) != lim)           { \
         if (*(p)++ != ' ')      { \
            (p)--; \
            break; \
         } \
      } \
}
/* SKIP_SPACES, enclosed macro */
#define SKIP_SPACES (p, limit)  \
   do { \
      char *lim = (limit);         \
      while ((p) != lim)         { \
         if (*(p)++ != ' ')      { \
            (p)--;   \
            break;   \
         }   \
      }   \
} while (0)
```

The second definition is the same as the first, except that it is enclosed in a `do {...} while (0)` construct. Enclosing the definition within the `do {…} while (0)` pair means that the macro can be invoked more like a function. With the first definition, sometimes you would have to follow the macro with a semi-colon, and sometimes you would not. Whereas with the second definition, the `do {...} while (0)` pair means that the macro is followed by a semi-colon.

With the `do {…} while (0)` construct, you can treat the macro as a function and put the semicolon after it.

For example,

```
/* SKIP_SPACES, enclosed macro, ends without ';' */
if (*p != 0)
   SKIP_SPACES (p, lim);
else …
```

This expands to

```
if (*p != 0)
   do {
      ...
   } while (0);
else …
```

Without the `do {…} while (0)` construct, the expansion would be:

```
if (*p != 0)
   {
      ,,,
   };     /* Probably not intended syntax */
  else
```

Preprocessing of .IDL Files

Every VisualDSP++ Interface Definition Language (VIDL) specification is analyzed by the C++ language preprocessor prior to syntax analysis. For more information, refer to the *VisualDSP++ Component Software Engineering User's Guide*.

The `#include` directive is used to control the inclusion of additional VIDL source text from a secondary input file that is named in the directive. Two available forms of `#include` are shown in Figure 1-2.

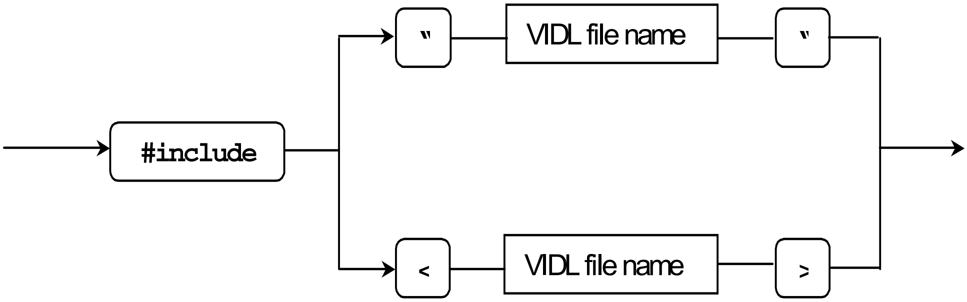

Figure 1-2. #INCLUDE Syntax Diagram

The file identified by the file name is located by searching a list of directories. When the name is delimited by quote characters, the search begins in the directory containing the primary input file, then proceeds with the list of directories specified by the `-I` command-line switch. When the name is delimited by angle-bracket characters, the search proceeds directly with the directories specified by `-I`. If the file is not located within any directory on the search list, the search may be continued in one or more platform dependent system directories.

C/C++ Run-Time Model and Environment

This section provides a full description of the ADSP-219x DSP run-time model and run-time environment. The run-time model, which applies to compiler-generated code, includes descriptions of the layout of the stack, data access, and call/entry sequence. The C/C++ run-time environment includes the conventions that C/C++ routines must follow to run on ADSP-219x DSPs. Assembly routines linked to C/C++ routines must follow these conventions.

ADI recommends that assembly programmers maintain stack conventions.

This section contains:

Using the Run-Time Header

The run-time header is an assembly language procedure that initializes the processor and sets up processor features to support the C run-time environment. The default run-time header source code for the ADSP-219x processors is in the `219x_hdr.asm` file. This run-time header performs the following operations:

- Initializes the C run-time environment
- Calls your `main` routine
- Calls `exit` routine, defined in the C run-time library (`libc.dlb`), if `main` returns.
- Defines system halt instruction called from `exit` routine.

Interrupt Table and Interface

The interrupt table is an assembly language set of functions defined in named sections. These sections get placed appropriately in the Linker Description File (LDF) to be executed at interrupt vector addresses. The default code for the ADSP-219x DSP's interrupt table is defined in `219x_int_tab.asm`.

The default interrupt table uses the following external symbols:

`_lib_int_table`	Static table holding interrupt information defined in the C run-time library
`__lib_int_determiner`	An interrupt dispatcher defined in the C run-time library
`_____system_start`	C run-time initialization defined in the run-time header

The `219x_int_tab` file contains a section of code for each hardware interrupt. The `.LDF` file places these code sections in the correct interrupt vector slots for each interrupt.

If an interrupt occurs, program execution begins at the interrupt vector addresses. Program execution causes a jump to `__lib_int_determiner` in the default vector code. If `__lib_int_determiner` finds (by inspecting `__lib_int_table`) a handler set for the interrupt, it will call the handler. `__lib_int_determiner` saves and restores all scratch registers around the handler call. The function `__lib_int_determiner` terminates by executing a return from interrupt (RTI) instruction, which restores program execution to the point at which the interrupt was raised.

A handler for an interrupt or signal is set using the `interrupt` or `signal` C run-time library functions. These functions pass the signal name and a handler function pointer as parameters. The signal macro names are defined in `signal.h`.

The default interrupt vector code may be replaced with custom code by modifying or creating a new piece of code to be placed at the vector addresses. This is usually done by copying the default `219x_int_tab.asm` file and `.LDF` file into your project and modifying them as required.

An `interrupt` pragma defined function can be placed in the interrupt vector code directly or be jumped to from the vector if it does not fit in the interrupt vector space (see "Interrupt Handler Pragmas" on page 1-122).

> The reset vector code, which is placed at address zero (0) and does a jump to `_____system_start`, should not be replaced.

Stack Frame

The stack frame (or activation record) provides for the following activities:

- space for local variables for the current procedure. For the compiler, this includes temporary storage as well as that required for explicitly declared user automatic variables.
- place to save linkage information such as return addresses, location information for the previous caller's stack frame, and to allow this procedure to return to its caller

- space to save information that must be preserved and restored
- arguments passed to the current procedure

In addition, if this is not a leaf procedure (if it is going to call other procedures), its stack frame also contains outgoing linkage and parameter space:

- space for the arguments to the called procedure.
- space for the callee to save basic linkage information.

Figure 1-3 provides a general overview of the stack. Note that the stack grows downward on the page. Because the stacks grow towards smaller addresses, higher addresses are found in the upwards direction. "Stack Frame Description" on page 1-157

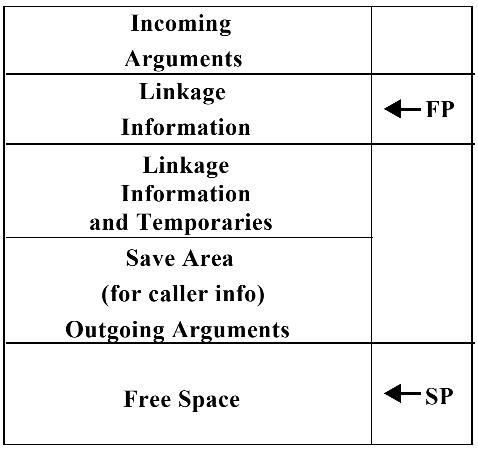

Figure 1-3. ADSP-219x DSP Stack

The stack resides in primary Data Memory (`DM`). It is controlled by a pair of pointers: a Stack Pointer (`SP`), which identifies the boundary of the in-use portion of the stack space, and a Frame Pointer (`FP`), which provides stable addressing to the current frame.

For increased efficiency the ADSP-219x DSP run-time model assumes 32 words past the stack pointer are available for use within a leaf function (functions that make no calls) and are *protected* from changes by an interrupt handler. This provides easy access to temporary space for use within simple leaf functions without the overhead of establishing a stack frame.

Stack Frame Description

This section describes the stack, as shown in Figure 1-3 on page 1-156.

Incoming Arguments

The memory area for incoming arguments begins at the `FP` value `+1`. Argument words are mapped by ascending addresses, so the second argument word is mapped at `FP+2`.

Linkage Information

The return address is saved on the hardware stack by the `CALL` instruction. In the called function, the address can then be popped from the hardware stack and saved as part of the stack frame. This information is used by the debugger to generate call stack debug information for source level debugging. Saving the return address on the stack frame is also useful in avoiding overflowing the finite hardware call stack, when using recursion for example. The value stored on the stack gets pushed back on the hardware call stack before the return of the function. The compiler detects recursive routines and also offers a switch to avoid overflowing the hardware call stack, called `-no_hardware_pc_stack`, which might be required for highly nested software.

Local Variables and Temporaries

Space for a register save area and local variables/temporaries is allocated on the stack by the function prologue. Local variables and temporaries are typically placed first in this area so they can be addressed with the smallest offsets from FP. The register save area is located at the farthest end of this area and can be accessed by SP-relative addressing.

Outgoing Arguments

Outgoing arguments are located at the top of the stack prior to the call. Space may be pre-allocated or claimed at the time of each call.

Free Space

Space below SP is normally considered free and unprotected; it is available for use (in growing the stack) at any time, synchronously or asynchronously (the latter for interrupt handling). However, on the ADSP-219x DSPs, the 32 words past the SP are reserved as a protected temporary space for use within a procedure.

General System-Wide Specifications

Some general specifications that apply to the stacks are:

- The stacks grow down in memory from higher to lower addresses.
- The current frames' "base" is addressed by the FP register.
- The first free word in each stack is addressed by the SP register.

Data can be pushed onto the stack by executing an instruction like the following one for the ADSP-219x DSPs:

```
DM (I4 += M5) = rej.
```

- The return address of the caller is stored at offsets -1 and -2 from the address carried by the current FP, if it is stored on the stack.
- The linkage back to the previous stack frame is stored at offset 0 from the current FP.

At a procedure call, the following must be true:

- For the ADSP-219x DSPs, no space beyond the `SP` must be in use.
- There must be at least one free slot on the `PC` stack to hold the return address.

At an interrupt, the following must be true:

- Space beyond the `SP` must be available.
- For the ADSP-219x DSPs, the first 32 words beyond the `SP` must be protected; the interrupt routine should decrement the `SP` by 32 and then restore the original value of `SP` prior to return.

Return Values

Return values always use registers. Single word return values come back in register `AX1`. Double word return values are stored in `SR1:0`, with the most significant part in `SR1`.

If the return value is larger than two words, then the caller must allocate space and pass the address in, as a "hidden argument". The register `I0` is used for this purpose.

Procedure Call and Return

To call a procedure:

1. Evaluate the arguments and push them onto the stack.
2. Call the procedure.
3. On return, remove arguments if desired.

On Entry:

1. Save the old `FP` and set `FP` to current `SP`.
2. If `-no_hardware_pc_stack` is specified (for debugging), pop the PC stack and store it on the main stack.
3. Move the `SP` to create space for the new frame.
4. If `-g` is specified but `-no_hardware_pc_stack` is not specified, push the return address back onto the hardware stack.

After this step, the new frame is officially in place.

5. Continue saving registers, and then executing the procedure.

A leaf procedure which does not require much stack space might choose to omit steps (1) and (2), operating without its own stack frame. On the ADSP-219x DSPs, the 32 words of protected space beyond the `SP` can be used for temporary storage.

To Return from a Procedure:

1. If the hardware PC stack is not used, the return address must be loaded from the stack and restored.
2. Restore miscellaneous saved registers.
3. Place the return value in the correct register (if not there already).
4. Restore `FP` for the previous frame.
5. Reset `SP` to remove the frame for procedure.
6. Return to the caller.

Miscellaneous Information

This section contains a number of miscellaneous aspects of the design that may be helpful in understanding stack functionality.

- Procedures without prototypes can be called successfully, provided the argument types correspond properly. Prototypes are always good programming practice. Programs that call library subroutines should always include header files.
- There is no special interface for calling system library functions. They use the standard calling convention.

Register Classification

This section describes all of the ADSP-219x processor registers. Registers are listed in order of preferred allocation by the compiler

Callee Preserved Registers ("Preserved")

Registers `I2`, `I3`, `I7`, `M0, M2 and M4` are preserved. A subroutine which uses any of these registers must save (preserve) it and restore it.

Dedicated Registers

Certain registers have dedicated purposes and will not be used for other things. Compiled code and libraries expect the dedicated registers to be correct.

Caller Save Registers ("Scratch")

All registers not preserved or dedicated are *scratch*. A subroutine may use a scratch register without having to save it.

Circular Buffer Length Registers

Registers L0 through L7 are the circular buffer length registers. The compiler assumes that these registers contain zero, which disables circular buffering. They **must** be set to zero when compiled code is executing to avoid incorrect behavior. There is no restriction on the value of an L register when the corresponding I register has been reserved from compiler use. See "-reserve register[, register …]" on page 1-43 for information on reserving registers.

Mode Status (MSTAT) Register

The C runtime initializes the MSTAT register as part of the run-time header code. The compiler and run-time libraries assume to be running in these preset modes. If you change any of the modes listed in Table 1-13, ensure that they are reverted before calling C/C++ compiled functions or functions from the C run-time library. Failure to revert to the default modes may cause applications to fail when running.

Table 1-13. MSTAT Register Modes

Mode	Description	State
SEC_REG	Secondary Data Registers	disabled
BIT_REV	Bit-reversed address output	disabled
AR_SAT	ALU saturation mode	disabled
M_MODE	MAC result mode	Integer Mode, 16.0 format
SEC_DAG	Secondary DAG registers	disabled

Complete List of Registers

The following tables describe all of the registers for the ADSP-219x DSPs.

- Table 1-14 lists the data register's file registers
- Table 1-15 lists the DAG1 registers
- Table 1-16 lists the DAG2 registers
- Table 1-17 lists miscellaneous registers

You must **always** specify **all parts** of the SR and MR register groups. For example, either MR1 or SR0 is valid, but either SR or MR by itself is **not** valid.

Table 1-14. List of Data Register File Registers

Register	Description	Notes
AX0	scratch	
AX1	scratch; single-word return	
AY0	scratch	
AY1	scratch	
AR	scratch;	
AF	scratch	
MX0	scratch	
MX1	scratch	
MY0	scratch	
MY1	scratch	
MR1:0	scratch	
MR2	scratch	

Table 1-14. List of Data Register File Registers (Cont'd)

Register	Description	Notes
PX	scratch	
SB	scratch	
SE	scratch	
SI	scratch	
SR1:0	scratch; double-word return	
SR2	scratch	

Table 1-15. List of DAG1 Registers

Register	Description
I0	scratch
I1	scratch
I2	preserved
I3	preserved
M0	preserved
M1	scratch
M2	preserved
M3	scratch
L0-3	not used, must be zero
B0-3	not used
DMPG1	preserved

Table 1-16. List of DAG2 Registers

Register	Description
I4	Dedicated: SP
I5	Dedicated: FP
I6	scratch
I7	preserved
M4	preserved
M5	dedicated: -1
M6	scratch
M7	scratch
L4-7	not used, must be zero
B4-7	not used
DMPG2	dedicated, must not change

Table 1-17. Miscellaneous Registers

Register	Description
ASTAT	scratch
CCODE	preserved; (0x8 ALU result sign default in C/C++ code)
CNTR	scratch
ICNTL	scratch
IJPG	scratch
IMASK	scratch
IRPTL	scratch
LPSTACKA	scratch

Table 1-17. Miscellaneous Registers (Cont'd)

LPSTACKP	scratch
MSTAT	scratch
STACKA	scratch
STACKP	scratch
STACKA	scratch

File I/O Support

The VisualDSP++ environment provides access to files on a host system, using `stdio` functions. File I/O support is provided through a set of low-level primitives that implement the `open`, `close`, `read`, `write`, and `seek` operations.The `stdio` functions make use of these primitives to provide conventional C input and output facilities. The source files for the I/O primitives are available under the VisualDSP++ installation in the sub-directory `219x\lib\src\libio_src`.

Refer to "stdio.h" on page 3-13 for more information.

Extending I/O Support To New Devices

The I/O primitives are implemented using an extensible device driver mechanism. The default start-up code includes a device driver that can perform I/O through the VisualDSP++ simulator and EZ-KIT Lite. Other device drivers may be registered and then used through the normal `stdio` functions.

A device driver is a set of primitive functions, grouped together into a `DevEntry` structure. This structure is defined in `device.h`:

```
struct DevEntry {
   int   DeviceID;
   void  *data;
```

```
        int   (*init)(struct DevEntry *entry);
        int   (*open)(const char *name, int mode);
        int   (*close)(int fd);
        int   (*write)(int fd, unsigned char *buf, int size);
        int   (*read)(int fd, unsigned char *buf, int size);
        long  (*seek)(long fd, long offset, int whence);
    }

    typedef struct DevEntry DevEntry;
    typedef struct DevEntry *DevEntry_t;
```

The `DeviceID` field is a unique identifier for the device, known to the user. Device IDs are used globally across an application. The data field is a pointer for any private data the device may need; it is not used by the run-time libraries. The function pointed to by the `init` field is invoked by the run-time library when the device is first registered. It returns a negative value for failure, positive value for success.

The functions pointed to by the `open`, `close`, `write` and `read` fields are the functions that provide the same functionality used in the default I/O device. `Seek` is another function at the same level, for those devices which support such functionality. If a device does not support an operation (such as seeking, writing on read-only devices or reading write-only devices), then a function pointer must still be provided; the function must arrange to always return failure codes when the operation is attempted.

A new device can be registered with the following call:

```
int add_devtab_entry(DevEntry_t entry);
```

If the device is successfully registered, the `init()` routine of the device is called, with `entry` as its parameter. `add_devtab_entry()` returns the DeviceID of the device registered.

If the device is not successfully registered, a negative value is returned. Reasons for failure include, but are not limited to:

- The DeviceID is the same as another device, already registered
- There are no more spaces left in the device registry table
- The DeviceID is less than zero
- Some of the function pointers are NULL
- The device's `init()` routine returned a failure result

Once a device is registered, it can be made the default device, using the following function:

```
void set_default_io_device(int);
```

The user passes the `DeviceID`. There is a corresponding function for retrieving the current default device:

```
int get_default_io_device(void);
```

The default device is used by `fopen()` when a file is first opened. The `fopen()` function passes the open request to the `open()` function of the device indicated by `get_default_io_device()`. The device file identifier (`dfid`) returned by the `open()` function is private to the device; other devices may simultaneously have other files open which use the same identifier. An open file is uniquely identified by the combination of `DeviceID` and `dfid`.

The `fopen()` function records the `DeviceID` and `dfid` in the global open file table, and allocates its own internal `fid` to this combination. All future operations on the file reads, writes, seeks and close —use this `fid` to retrieve the `DeviceID`—and thus direct the request to the appropriate device's primitive functions, passing the `dfid` along with other parameters. Once a file has been opened by `fopen()`, the current value of `get_default_io_device()` is irrelevant to that file.

C/C++ and Assembly Language Interface

This section describes how to call assembly language subroutines from within C or C++ programs and C or C++ functions from within assembly language programs.

Before attempting to perform either of these calls, be sure to familiarize yourself with the information about the C/C++ run-time model (including details about the stack, data types, and how arguments are handled) contained in "C/C++ Run-Time Model and Environment" on page 1-153.

This section contains:

Calling Assembly Subroutines from C/C++ Programs

Before calling an assembly language subroutine from a C/C++ program, create a prototype to define the arguments for the assembly language subroutine and the interface from the C/C++ program to the assembly language subroutine. Even though it is legal to use a function without a prototype in C/C++, prototypes are a strongly recommended practice for good software engineering. When the prototype is omitted, the compiler cannot perform argument type checking and assumes that the return value is of type integer and uses K&R promotion rules instead of ANSI promotion rules.

The run-time model defines some registers as *scratch* registers and others as *preserved* or *dedicated* registers. Scratch registers can be used within the assembly language program without worrying about their previous contents. If more room is needed (or an existing code is used) and you wish to use the preserved registers, you *must save* their contents and then *restore* those contents before returning.

Do *not* use the dedicated or stack registers for other than their intended purpose; the compiler, libraries, debugger, and interrupt routines depend on having a stack available as defined by those registers.

The compiler also assumes the machine state does not change during execution of the assembly language subroutine.

Do not change any machine modes (for example, certain registers may be used to indicate circular buffering when those register values are nonzero).

The compiler prefaces the name of any external entry point with an underscore. Therefore, declare your assembly language subroutine's name with a leading underscore. If you're using the function from assembly programs as well, you might want your function's name to be just as you write it. Then you'll also need to tell the C/C++ compiler that it's an `asm` function, by placing `'extern "asm" {}'` around the prototype.

The C/C++ runtime determines that all function parameters are passed on the stack. A good way to observe and understand how arguments are passed is to write a dummy function in C or C++ and compile it using the `-save-temps` command-line switch (see on page 1-44). The resulting compiler generated assembly file (.s) can then be viewed.

The following example includes the global volatile variable assignments to indicate where the arguments can be found upon entry to `asmfunc`.

```
// Sample file for exploring compiler interface...
// global variables ... assign arguments there just so
// we can track which registers were used
// (type of each variable corresponds to one of arguments)

int global_a;
float global_b;
int *  global_p;

// the function itself

int asmfunc(int a, float b, int * p, int d, int e) {
   //do some assignments so .s file will show where args are
   global_a = a;
   global_b = b;
   global_p = p;
   //value gets loaded into the return register
   return 12345;
}
```

When compiled with the `-save-temps` option set, this produces the following code.

```
_asmfunc:
      AX1 = DM(I4 +  2);
      SI = DM(I4 +  1);
      AX0 = DM(I4 +  3);
      DM(ADDRESS(_global_b)) = AX1;
      AY0 = DM(I4 +  4);
      DM(ADDRESS(_global_a)) = SI;
      DM(ADDRESS(_global_b+1)) = AX0;
      RTS (DB);
      DM(ADDRESS(_global_p)) = AY0;
      AX1 = 12345;

_asmfunc.end:
   .global _asmfunc;
   .type _asmfunc,STT_FUNC;
```

For a more complicated function, you might find it useful to follow the general run-time model, and use the run-time stack for local storage, etc. A simple C program, passed through the compiler, will provide a good template to build on. Alternatively, you may find it just as convenient to use local static storage for temporaries.

Calling C/C++ Functions from Assembly Programs

You may want to call a C/C++ callable library and other functions from within an assembly language program. As discussed in "Calling Assembly Subroutines from C/C++ Programs" on page 1-170, you may want to create a test function to do this in C/C++, and then use the code generated by the compiler as a reference when creating your assembly language program and the argument setup. Using volatile global variables may help clarify the essential code in your test function.

The run-time model defines some registers as *scratch* registers and others as *preserved* or *dedicated*. The contents of the scratch registers may be changed without warning by the called C/C++ function. If the assembly language program needs the contents of any of those registers, you *must save* their contents before the call to the C/C++ function and then *restore* those contents after returning from the call.

Do *not* use the dedicated registers for other than their intended purpose; the compiler, libraries, debugger, and interrupt routines all depend on having a stack available as defined by those registers.

Preserved registers can be used; their contents will not be changed by calling a C/C++ function. The function will always save and restore the contents of preserved registers if they are going to change.

If arguments are on the stack, they are addressed via an offset from the stack pointer or frame pointer. Explore how arguments are passed between an assembly language program and a function by writing a dummy function in C/C++ and compiling it with the `save temporary files` option (the "-save-temps" command-line switch). By examining the contents of volatile global variables in `*.s` file, you can determine how the C/C++ function passes arguments, and then duplicate that argument setup process in the assembly language program.

The stack must be set up correctly before calling a C/C++ callable function. If you call other functions, maintaining the basic stack model also facilitates the use of the debugger. The easiest way to do this is to define a C/C++ main program to initialize the run-time system; maintain the stack until it is needed by the C/C++ function being called from the assembly language program; and then continue to maintain that stack until it is needed to call back into C/C++. However, make sure the dedicated registers are correct. You do not need to set the `FP` prior to the call; the caller's `FP` is never used by the recipient.

Using Mixed C/C++ and Assembly Naming Conventions

It is necessary to be able to use C/C++ symbols (function names or variable names) in assembly routines and use assembly symbols in C routines. This section describes how to name C/C++ and assembly symbols and shows how to use C/C++ and assembly symbols.

To name an assembly symbol that corresponds to a C/C++ symbol, add an underscore prefix to the C/C++ symbol name when declaring the symbol in assembly. For example, the C/C++ symbol `main` becomes the assembly symbol `_main`.

To use a C/C++ function or variable in your assembly routine, declare it as global in the C/C++ program and import the symbol into the assembly routine by declaring the symbol with the `.EXTERN` assembler directive.

The C++ language performs name mangling on function names it defines according to the output and input parameter types of the function. If calling into a C++ defined function from assembly code, the `.EXTERN` symbol will need to be the mangled C++ output name. This is best retrieved by looking at the compiler's assembly output for the C++ source that defines the required function.

To use an assembly function or variable in your C/C++ program, declare the symbol with the `.GLOBAL` assembler directive in the assembly routine and import the symbol by declaring the symbol as `extern` in the C/C++ program.

Alternatively, the `cc219x` compiler provides an "`asm`" linkage specifier (used similarly to the "C" linkage specifier of C++), which when used, removes the need to add an underscore prefix to the symbol that is defined in assembly.

Table 1-18 shows several examples of the C/C++ and assembly interface naming conventions.

Table 1-18. C/C++ Naming Conventions for Symbols

In C/C++ Program	In Assembly Subroutine
`int c_var; /*declared global*/`	`.extern _c_var;`
`void c_func(); /* in C code */`	`.extern _c_func;`
`void cpp_func(void); /* in C++ source */`	`.extern _cpp_func__Fv;`
`extern int asm_var;`	`.global _asm_var;`
`extern void asm_func();`	`.global _asm_func;` `_asm_func:`
`extern "asm" void asm_func();`	`.global asm_func;` `asm_func:`

C++ Programming Examples

This section shows examples of the features specific to C++. These examples are:

- "Using Fract Type Support" on page 1-176
- "Using Complex Number Support" on page 1-177

By default, the `cc219x` compiler runs in C mode. To run the compiler in C++ mode, use the corresponding option on the command line, or select the option in the **Project Options** dialog box in the VisualDSP++ environment.

For example, the following command line

```
cc219x -c++ source.cpp -proc ADSP-2191
```

runs `cc219x` with:

`-c++`
Specifies that the following source file is written in ANSI/ISO standard C++ extended with the Analog Devices keywords.

`source.cpp`
Specifies the source file for your program.

`-proc ADSP-2191`
Specifies that the compiler should produce code suitable for the ADSP-2191 DSP.

Using Fract Type Support

Listing 1-1 on page 1-177 demonstrates the compiler support for the `fract` type and associated arithmetic operators, such as `+` and `*`. The dot product algorithm is expressed using the standard arithmetic operators. The code demonstrates how two variable-length arrays are initialized with fractional literals.

For more information abut the fractional data type and arithmetic, see "Fractional Type Support" on page 1-90.

Listing 1-1. Example Code: Using Fract Data Type — C++ code

```
#include <fract>
#define N 20
fract x[N] = {.5r,.5r,.5r,.5r,.5r,.5r,.5r,.5r,.5r,
     .5r,.5r,.5r,.5r,.5r,.5r,.5r,.5r,.5r,.5r,.5r};
fract y[N] = {0,.1r,.2r,.3r,.4r,.5r,.6r,.7r,.8r,.9r,.10r,.1r,
     .2r,.3r,.4r,.5r,.6r,.7r,.8r,.9r};
fract fdot(int N, fract *x, fract *y)
{
   int j;
   fract s;
   s = 0;
   for (j=0; j<N; j++)
   {
   s += x[j] * y[j];
   }
   return s;
}
int main(void)
{
   fdot(N,x,y);
}
```

Using Complex Number Support

The Mandelbrot fractal set is defined by the following iteration on complex numbers:

```
z := z * z + c
```

The `c` values belong to the set for which the above iteration does not diverge to infinity. The canonical set is defined when `z` starts from `zero`.

Listing 1-2 demonstrates the Mandelbrot generator expressed in a simple algorithm using the C++ library `complex` class:

Listing 1-2. Mandelbrot Generator Example — C++ code

```
#include <complex>
int iterate (complex<double> c, complex<double> z, int max)
{
int n;
   for (n = 0; n<max && abs(z)<2.0; n++)
   {
      z = z * z + c;
   }
   return (n == max ? 0 : n);
}
```

Listing 1-3 shows a C version of the inner computational function of the Mandelbrot generator extracts performance and programming penalties (compared with the C++ version).

Listing 1-3. Mandelbrot Generator Example — C Code

```
int iterate (double creal, double cimag,
double zreal, double zimag, int max)
{
double real, imag;
   int n;
   real = zreal * zreal;
   imag = zimag * zimag;
   for (n = 0; n<max && (real+imag)<4.0; n++)
   {
      zimag = 2.0 * zreal * zimag + cimag;
      zreal = real - imag + creal;
      real = zreal * zreal;
      imag = zimag * zimag;
   }
   return (n == max ? 0 : n);
}
```

2 ACHIEVING OPTIMAL PERFORMANCE FROM C/C++ SOURCE CODE

This chapter provides guidance to help you to tune your application to achieve the best possible code from the compiler. Some implementation choices are available when coding an algorithm, and understanding their impact is crucial to attaining optimal performance.

The focus of what follows is on how to obtain maximal code performance from the compiler. Most of these guidelines also apply when optimizing for minimum code size, although some techniques specific to that goal are also discussed. The first section looks at some general principles, and how the compiler can lend the most help to your optimization effort. Optimal coding styles are then considered in detail. Special features such as compiler switches, built-in functions, and pragmas are also discussed. The chapter ends with a short example to demonstrate how the optimizer works.

Small examples are included throughout this chapter to demonstrate points being made. Some show recommended coding styles, others identify styles to be avoided or code that it may be possible to improve. These are marked as "**Good**" and "**Bad**", respectively.

This chapter contains:

- "Pragmas" on page 2-31
- "Useful Optimization Switches" on page 2-37

General Guidelines

It is useful to bear in mind the following basic strategy when writing an application:

1. Try to choose an algorithm that is suited to the architecture being targeted. For example, there may be a trade-off between memory usage and algorithm complexity that may be influenced by the target architecture.
2. Code the algorithm in a simple, high-level generic form. Keep the target in mind, especially with respect to choices of data types.
3. You can then turn your attention towards code tuning. For critical code sections, think more carefully about the strengths of the target platform, and make any non-portable changes where necessary.

Tip: Choose the language as appropriate.

As the programmer your first decision is to choose whether to implement your application in C or C++. This decision may be influenced by performance considerations. C++ code using only C features will have very similar performance to pure C source. Many higher-level C++ features (for example those resolved at compile-time, such as namespaces, overloaded functions and inheritance) have no performance cost. However, use of some other features may degrade performance, and so the performance loss must be weighed against the richness of expression available in C++. Examples of features that may degrade performance are virtual functions or using classes to implement basic data types.

This section contains:

- "How the Compiler Can Help" on page 2-4
- "Data Types" on page 2-7
- "Getting the Most from IPA" on page 2-9

How the Compiler Can Help

The compiler provides many facilities designed to help the programmer.

Using the Compiler Optimizer

There is a vast difference in performance between code compiled optimized and code compiled non-optimized. In some cases optimized code can run ten or twenty times faster. Optimization should always be used when measuring performance or shipping code as product.

The optimizer in the C/C++ compiler is designed to generate efficient code from source that has been written in a straightforward manner. The basic strategy for tuning a program is to present the algorithm in a way that gives the optimizer excellent visibility of the operations and data, and hence the greatest freedom to safely manipulate the code. Future releases of the compiler will continue to enhance the optimizer, and expressing algorithms simply will provide the best chance of benefiting from such enhancements.

Note that the default setting (or "debug" mode within the VisualDSP++ IDDE) is for non-optimized compilation in order to assist programmers in diagnosing problems with their initial coding. The optimizer is enabled in VisualDSP++ by checking the **Enable optimization** checkbox under the **Project Options ->Compile** tab. This adds the `-O` (enable optimization) switch (on page 1-37) to the compiler invocation. A "release" build from within VisualDSP++ will automatically enable optimization.

Using the Statistical Profiler

Tuning an application begins with an understanding of which areas of the application are most frequently executed and therefore where improvements would provide the largest gains. The statistical profiling feature provided in VisualDSP++ is an excellent means for finding these areas. More details about how to use it may be found in the *VisualDSP++ 3.5 User's Guide*.

The particular advantage of statistical profiling is that it is completely unobtrusive. Other forms of profiling insert instrumentation into the code, perturbing the original optimization, code size and register allocation to some degree.

The best methodology is usually to compile with both optimization and debug information generation enabled. In this way, you can obtain a profile of the optimized code while retaining function names and line number information. This will give you accurate results that correspond directly to the C/C++ source. Note that the compiler optimizer may have moved code between lines.

You can obtain a more accurate view of your application if you build it optimized but without debug information generation. You will then obtain statistics that relate directly to the assembly code. The only problem with doing this may be in relating assembly lines to the original source. Do not strip out function names when linking, since keeping function names means you can scroll through the assembly window to instructions of interest.

In very complicated code, you can locate the exact source lines by counting the loops, unless they are unrolled. Looking at the line numbers in the assembly file (use the `-save-temps` switch (on page 1-44) to retain compiler generated assembly files, which will have the `.s` filename extension) may also help. The compiler optimizer may have moved code around so that it does not appear in the same order as in your original source.

Using Interprocedural Optimization

To obtain the best performance, the optimizer often requires information that can only be determined by looking outside the function that it is working on. For example, it helps to know what data can be referenced by pointer parameters, or whether a variable actually has a constant value. The `-ipa` compiler switch (on page 1-30) enables interprocedural analysis (IPA), which can make this available. When this switch is used the compiler will be called again from the link phase to recompile the program using additional information obtained during previous compilations.

Because it only operates at link time, the effects of IPA will not be seen if you compile with the `-S` switch (on page 1-43). To see the assembly file when IPA is enabled, use the `-save-temps` switch (on page 1-44), and look at the `.s` file produced after your program has been built.

As an alternative to IPA, you can achieve many of the same benefits by adding pragma directives and other declarations such as `__builtin_aligned` to provide information to the compiler about how each function interacts with the rest of the program.

These directives are further described "Using Built-In Functions in Code Optimization" on page 2-25 and "Pragmas" on page 2-31.

Data Types

The compiler directly supports the following scalar data types.

Single-Word Fixed-Point Data Types: Native Arithmetic	
`char`	16-bit signed integer
`unsigned char`	16-bit unsigned integer
`short`	16-bit signed integer
`unsigned short`	16-bit unsigned integer
`int`	16-bit signed integer
`unsigned int`	16-bit unsigned integer
Double-Word Fixed-Point Data Types: Emulated Arithmetic	
`long`	32-bit signed integer
`unsigned long`	32-bit unsigned integer
Floating-Point Data Types: Emulated Arithmetic	
`double`	32-bit float
`float`	32-bit float

Fractional data types are represented using the integer types. Manipulation of these is best done by use of built-in functions, described in "System Support Built-In Functions" on page 2-26.

Avoiding Emulated Arithmetic

Arithmetic operations for some types are implemented by library functions because the DSP hardware does not directly support these types. Consequently, operations for these types are far slower than native operations-sometimes by a factor of a hundred-and also produce larger code. These types are marked as “Emulated Arithmetic” in “Data Types” on page 2-7.

The hardware does not provide direct support for division, so division and modulus operations are almost always multi-cycle operations, even on integral type inputs. If the compiler has to issue a full division operation, it will usually need to generate a call to a library function. One notable situation in which a library call is avoided is for integer division when the divisor is a compile-time constant and is a power of two. In that case the compiler generates a shift instruction. Even then, a few fix-up instructions are needed after the shift if the types are signed. If you have a signed division by a power of two, consider whether you can change it to unsigned in order to obtain a single-instruction operation.

When the compiler has to generate a call to a library function for one of these arithmetic operators that are not supported by the hardware, performance will suffer not only because the operation will take multiple cycles, but also because the effectiveness of the compiler optimizer will be reduced.

For example, such an operation in a loop can prevent the compiler from making use of efficient zero-overhead hardware loop instructions. Also, calling the library to perform the required operation can change values held in scratch registers before the call, so the compiler will have to generate more stores and loads from the data stack to keep values required after the call returns. Emulated arithmetic operators should therefore be avoided where possible, especially in loops.

Getting the Most from IPA

Interprocedural analysis (IPA) is designed to try to propagate information about the program to parts of the optimizer that can use it. This section looks at what information is useful, and how to structure your code to make this information easily accessible to the analysis.

Initializing Constants Statically

IPA will identify variables that have only one value and replace them with constants, resulting in a host of benefits for the optimizer's analysis. For this to happen a variable must have a single value throughout the program. If the variable is statically initialized to zero, as all global variables are by default, and is subsequently assigned some other value at another point in the program, then the analysis sees two values and will not consider the variable to have a constant value.

For example,

```
#include <stdio.h>
int val;            // initialized to zero
void init() {
   val = 3;         // re-assigned
}
void func() {
   printf("val %d",val);
}
int main() {
   init();
   func();
}
```

Bad: IPA cannot see that `val` is a constant

is better written as

```
#include <stdio.h>
const int val = 3;    // initialized once
```

```
void init() {
}
void func() {
   printf("val %d",val);
}
int main() {
   init();
   func();
}
```

Good: IPA knows `val` is 3.

Avoiding Aliases

It may seem that the iterations may be performed in any order in the following loop:

```
void fn(char a[], char b[], int n) {
int i;
for (i = 0; i < n; ++i)
   a[i] = b[i];
}
```

Bad: `a` and `b` may alias each other.

but `a` and `b` are both parameters, and, although they are declared with `[]`, they are in fact pointers, which may point to the same array. When the same data may be reachable through two pointers, they are said to alias each other.

If IPA is enabled, the compiler will look at the call sites of `fn` and try to determine whether `a` and `b` can ever point to the same array.

Even with IPA, it is quite easy to create what appear to the compiler as aliases. The analysis works by associating pointers with sets of variables that they may refer to at some point in the program. If the sets for two pointers are found to intersect, then both pointers are assumed to point to the union of the two sets.

If `fn` above were called in two places with global arrays as arguments, then IPA would have the results shown below:

```
fn(glob1, glob2, N);
fn(glob1, glob2, N);
```

Good: sets for `a` and `b` do not intersect: `a` and `b` are not aliases.

```
fn(glob1, glob2, N);
fn(glob3, glob4, N);
```

Good: sets for `a` and `b` do not intersect: `a` and `b` are not aliases.

```
fn(glob1, glob2, N);
fn(glob3, glob1, N);
```

Bad: sets intersect - both `a` and `b` may access `glob1`; `a` and `b` may be aliases.

The third case arises because IPA considers the union of all calls at once, rather than considering each call individually, when determining whether there is a risk of aliasing. If each call were considered individually, IPA would have to take flow control into account and the number of permutations would make compilation time impracticably long.

The lack of control flow analysis can also create problems when a single pointer is used in multiple contexts. For example, it is better to write:

```
int *p = a;
int *q = b;
   // some use of p
   // some use of q
```

Good: `p` and `q` do not alias.

than

```
int *p = a;
   // some use of p
```

```
p = b;
   // some use of p
```

Bad: uses of `p` in different contexts may alias.

because the latter may cause extra apparent aliases between the two uses.

Indexed Arrays vs. Pointers

C allows a program to access data from an array in two ways: either by indexing from an invariant base pointer, or by incrementing a pointer. These two versions of vector addition illustrate the two styles:

Style 1: using indexed arrays

```
void va_ind(const short a[], const short b[], short out[], int n) {
   int i;
   for (i = 0; i < n; ++i)
      out[i] = a[i] + b[i];
}
```

Style 2: using pointers

```
void va_ptr(const short a[], const short b[], short out[], int n) {
   int i;
   short *pout = out;
   const short *pa = a, *pb = b;
   for (i = 0; i < n; ++i)
      *pout++ = *pa++ + *pb++;
}
```

Trying Pointer and Indexed Styles

One might hope that the chosen style would not make any difference to the generated code, but this is not always the case. Sometimes, one version of an algorithm will generate better optimized code than the other, but it is not always the same style that is better.

Tip: Try both pointer and index styles.

The pointer style introduces additional variables that compete with the surrounding code for resources during the compiler optimizer's analysis. Array accesses, on the other hand, must be transformed to pointers by the compiler and sometimes it does not do the job as well as you could do by hand.

The best strategy is to start with array notation. If the generated code looks unsatisfactory, try using pointers. Outside the critical loops, use the indexed style, since it is easier to understand.

Function Inlining

The function inlining may be used in two ways

- By annotating functions in the source code with the `inline` keyword. In this case, function inlining is only performed when optimization is enabled.
- By turning on automatic inlining with the `-Oa` switch (on page 1-37). This switch automatically enables optimization.

Tip: Inline small, frequently executed functions.

You can use the compiler's `inline` keyword to indicate that functions should have code generated inline at the point of call. Doing this avoids various costs such as program flow latencies, function entry and exit instructions and parameter passing overheads. Using an `inline` function also has the advantage that the compiler can optimize through the inline code and does not have to assume that scratch registers and condition states are modified by the call. Prime candidates for inlining are small, frequently used functions because they will cause the least code-size increase while giving most performance benefit.

As an example of the usage of the `inline` keyword, the function below sums two input parameters and returns the result.

```
inline int add(int a, int b) {
   return (a+b);
}
```

Good: use of the `inline` keyword.

Inlining has a code-size to performance trade-off that should be considered when it is used. With `-Oa`, the compiler will automatically inline small functions where possible. If the application has a tight upper code-size limit, the resulting code-size expansion may be too great. It is worth considering using automatic inlining in conjunction with the `-Ov n` switch (on page 1-38) to restrict inlining (and other optimizations with a code-size cost) to parts of the application that are performance-critical. This will be considered in more detail later in this chapter.

Using Inline asm Statements

The compiler allows use of inline `asm` statements to insert small sections of assembly into C code.

Tip: Avoid use of inline `asm` statements where built-in functions may be used instead

The compiler does not intensively optimize code that contains inline `asm` statements because it has little understanding about what the code in the statement does. In particular, use of an `asm` statement in a loop may inhibit useful transformations.

The compiler has been enhanced with a large number of built-in functions. These generate specific hardware instructions and are designed to allow the programmer to more finely tune the code produced by the compiler, or to allow access to system support functions. A complete list of compiler's built-in functions is given in "Compiler Built-In Functions" on page 1-94.

Use of these builtins is much preferred to using inline `asm` statements. Since the compiler knows what each builtin does, it can easily optimize around them. Conversely, since the compiler does not parse asm statements, it does not know what they do, and so is hindered in optimizing code that uses them. Note also that errors in the text string of an `asm` statement will be caught by the assembler and not the compiler.

Examples of efficient use of built-in functions are given in "System Support Built-In Functions" on page 2-26.

Memory Usage

The compiler, in conjunction with the use of the linker description file (`.LDF`), allows the programmer control over where data is placed in memory. This section describes how to best lay out data for maximum performance.

Tip: Try to put arrays into different memory sections.

The DSP hardware can support two memory operations on a single instruction line, combined with a compute instruction. However, two memory operations will only complete in one cycle if the two addresses are situated in different memory blocks; if both access the same block, then a stall will be incurred.

Take as an example the dot product loop below. Because data is loaded from both array `a` and array `b` in every iteration of the loop, it may be useful to ensure that these arrays are located in different blocks.

```
for (i=0; i<100; i++)
    sum += a[i] * b[i];
```

Bad: compiler assumes that two memory accesses together may give a stall.

This is done by using the compiler extension described in "Dual Memory Support Keywords (pm dm)" on page 1-78. Placing a `pm` qualifier before the type definition tells the compiler that the array is located in what is notionally called "Program Memory" (PM). The memory of a ADSP-219x DSP is in one unified address space and there is no restriction concerning in which part of memory program code or data can be placed as on previous generations of DSP architectures like the ADSP-218x DSPs. However, the default `.LDF` files ensure that `pm` qualified data is placed in a different memory block than non-qualified (or `dm` qualified) data, thus allowing two accesses to occur simultaneously without incurring a stall.

The array declaration of one of either `a` or b is modified to

```
pm int a[100];
```

and any pointers to the buffer a become, for example,

```
pm int *p = a;
```

to allow simultaneous accesses to the two buffers.

Note that the explicit placement of data in PM can only be done for global data.

Loop Guidelines

Loops are where an application will ordinarily spend the majority of its time. It is therefore useful to look in detail at how to help the compiler to produce the most efficient code possible for them.

Keeping Loops Short

For best code efficiency, loops should be as small as possible. Large loop bodies are usually more complex and difficult to optimize. Additionally, they may require register data to be stored in memory. This will cause a decrease in code density and execution performance.

Avoiding Unrolling Loops

Tip: Do not unroll loops yourself.

Not only does loop unrolling make the program harder to read but it also prevents optimization by complicating the code for the compiler.

```
void va1(const short a[], const short b[], short c[], int n)
{
   int i;
   for (i = 0; i < n; ++i) {
   c[i] = b[i] + a[i];
   }
}
```

Good: the compiler will unroll if it helps.

```
void va2(const short a[], const short b[], short c[], int n)
{
   short xa, xb, xc, ya, yb, yc;
   int i;
   for (i = 0; i < n; i+=2) {
      xb = b[i]; yb = b[i+1];
      xa = a[i]; ya = a[i+1];
```

```
        xc = xa + xb; yc = ya + yb;
        c[i] = xc; c[i+1] = yc;
    }
}
```

Bad: harder for the compiler to optimize.

Avoiding Loop Rotation by Hand

Tip: Do not rotate loops by hand.

Programmers are often tempted to “rotate” loops in DSP code by “hand” attempting to execute loads and stores from earlier or future iterations at the same time as computation from the current iteration. This technique introduces loop-carried dependencies that prevent the compiler from rearranging the code effectively. However, it is better to give the compiler a “normalized” version, and leave the rotation to the compiler.

```
int ss(short *a, short *b, int n) {
    int sum = 0;
    int i;
    for (i = 0; i < n; i++) {
        sum += a[i] + b[i];
    }
    return sum;
}
```

Good: will be rotated by the compiler.

```
int ss(short *a, short *b, int n) {
    short ta, tb;
    int sum = 0;
    int i = 0;
    ta = a[i]; tb = b[i];
    for (i = 1; i < n; i++) {
        sum += ta + tb;
        ta = a[i]; tb = b[i];
    }
    sum += ta + tb;
```

```
    return sum;
}
```

Bad: rotated by hand—hard for the compiler to optimize.

By rotating the loop, the scalar variables `ta` and `tb` have been added, introducing loop-carried dependencies.

Avoiding Array Writes in Loops

Other dependencies can be caused by writes to array elements. In the following loop, the optimizer cannot determine whether the load from a reads a value defined on a previous iteration or one that will be overwritten in a subsequent iteration.

```
for (i = 0; i < n; ++i)
    a[i] = b[i] * a[c[i]];
```

Bad: has array dependency.

The optimizer can resolve access patterns where the addresses are expressions that vary by a fixed amount on each iteration. These are known as "induction variables".

```
for (i = 0; i < n; ++i)
    a[i+4] = b[i] * a[i];
```

Good: uses induction variables.

Inner Loops vs. Outer Loops

Tip: Inner loops should iterate more than outer loops.

The optimizer focuses on improving the performance of inner loops because this is where most programs spend the majority of their time. It is considered a good trade-off for an optimization to slow down the code before and after a loop if it is going to make the loop body run faster. Therefore, try to make sure that your algorithm also spends most of its time in the inner loop; otherwise it may actually be made to run slower by optimization. If you have nested loops where the outer loop runs many times and the inner loop a small number of times, it may be possible to rewrite the loops so that the outer loop has the fewer iterations.

Avoiding Conditional Code in Loops

If a loop contains conditional code, control-flow latencies may incur large penalties if the compiler has to generate conditional jumps within the loop. In some cases, the compiler will be able to convert `IF-THEN-ELSE` and `?:` constructs into conditional instructions. In other cases, it will be able to relocate the expression evaluation outside of the loop entirely. However, for important loops, linear code should be written where possible.

The compiler will not perform the loop transformation that interchanges conditional code and loop structures. Instead of writing

```
for (i=0; i<100; i++) {
   if (mult_by_b)
      sum1 += a[i] * b[i];
   else
      sum1 += a[i] * c[i];
}
```

Bad: loop contains conditional code.

it is better to write

```
if (mult_by_b) {
   for (i=0; i<100; i++)
      sum1 += a[i] * b[i];
} else {
   for (i=0; i<100; i++)
      sum1 += a[i] * c[i];
}
```

Good: two simple loops can be optimized well.

if this is an important loop.

Avoiding Placing Function Calls in Loops

The compiler will not usually be able to generate a hardware loop if the loop contains a function call due to the expense of saving and restoring the context of a hardware loop. In addition to obvious function calls, such as `printf()`, hardware loop generation can also be prevented by operations such as division, modulus, and some type coercions. These operations may require implicit calls to library functions. For more details, see "Data Types" on page 2-7.

Avoiding Non-Unit Strides

If you write a loop such as:

```
for (i=0; i<n; i+=3) {
      // some code
}
```

Bad: non-unit stride means division may be required.

then in order for the compiler to turn this into a hardware loop, it will need to work out the loop trip count. To do so, it must divide `n` by 3. The compiler will decide that this is worthwhile as it will speed up the loop, but as discussed above, division is an expensive operation. Try to avoid creating loop control variables with strides of non-unit magnitude.

Loop Control

Tip: Use `int` types for loop control variables and array indices.

Tip: Use automatic variables for loop control and loop exit test.

For loop control variables and array indices, it is always better to use signed `ints` rather than any other integral type. The C standard requires various type promotions and standard conversions that complicate the code for the compiler optimizer. Frequently, the compiler is still able to deal with such code and create hardware loops and pointer induction variables. However, it does make it more difficult for the compiler to optimize and may occasionally result in under-optimized code.

The same advice goes for using automatic (local) variables for loop control. It is easy for a compiler to see that an automatic scalar, whose address is not taken, may be held in a register during a loop. But it is not as easy when the variable is a global or a function static. Therefore, code such as

```
for (i=0; i<globvar; i++)
   a[i] = 10;
```

Bad: may need to reload `globvar` on every iteration.

may not create a hardware loop if the compiler cannot be sure that the write into the array `a` does not change the value of the global variable. The `globvar` must be re-loaded each time around the loop before performing the exit test.

In this circumstance, the programmer can make the compiler's job easier by writing:

```
int upper_bound = globvar;
for (i=0; i<upper_bound; i++)
   a[i] = 10;
```

Good: easily becomes hardware loop.

Using the Restrict Qualifier

The `restrict` qualifier provides one way to help the compiler resolve pointer aliasing ambiguities. Accesses from distinct `restricted` pointers do not interfere with each other. The loads and stores in the following loop

```
for (i=0; i<100; i++)
   a[i] = b[i];
```

Bad: possible alias of arrays a and b.

may be disambiguated by writing

```
int * restrict p = a;
int * restrict q = b;
for (i=0; i<100; i++)
   *p++ = *q++;
```

Good: restrict qualifier tells compiler that memory accesses do not alias.

The restrict keyword is particularly useful on function parameters.

Using the Const Qualifier

By default, the compiler assumes that the data referenced by a pointer to const type will not change. Therefore, another way to tell the compiler that the two arrays a and b do not overlap is to use the const keyword.

```
void copy(short *a, const short *b) {
   int i;
   for (i=0; i<100; i++)
      a[i] = b[i];
}
```

Good: pointers disambiguated via const qualifier.

The use of const in the above example will have a similar effect on the no_alias pragma (see in "#pragma no_alias" on page 2-36). In fact, the const implementation is better since it also allows the optimizer to use the fact that accesses via a and b are independent in other parts of the code, not just the inner loop.

In C, it is legal, though bad programming practice, to use casts to allow the data pointed to by pointers to const type to change. This should be avoided since, by default, the compiler will generate code that assumes `const` data does not change. If you have a program that modifies `const` data through a pointer, you can generate standard-conforming code by using the compile-time flag `-const-read-write`.

Avoiding Long Latencies

All pipelined machines will introduce stall cycles when you cannot execute the current instruction until a prior instruction has exited the pipeline.

If a stall is seen empirically, but it is not obvious to you exactly why it is occurring, a good way to learn about the cause is the **Pipeline Viewer**. This can be accessed through **Debug Windows -> Pipeline Viewer** in the VisualDSP++ 3.5 IDDE. By single-stepping through the program, you will see where the stall occurs. Note that the **Pipeline Viewer** is only available within a simulator session.

Using Built-In Functions in Code Optimization

Built-in functions, also known as compiler intrinsics, provide a method for the programmer to efficiently use low-level features of the DSP hardware while programming in C. Although this section does not cover all the built-in functions available (for more information, refer to "Compiler Built-In Functions" on page 1-94), it presents some code examples where implementation choices are available to the programmer.

Fractional Data

Fractional data, represented as an integral type, can be manipulated in two ways: one way is the use of long promoted shifts and multiply constructs, and the other is the use of compiler built-in functions. The built-in functions are recommended as they give you the most control over your data. Let's consider the fractional dot product algorithm. This may be written as:

```
long dot_product (short *a, short *b) {
   int i;
   long sum=0;
   for (i=0; i<100; i++) {
   /* this line is performance critical */
      sum += (((long)a[i]*b[i]) << 1);
   }
   return sum;
}
```

Bad: uses shifts to implement fractional multiplication.

This presents some problems to the optimizer. Normally, the code generated here would be a multiply, followed by a shift, followed by an accumulation. However, the DSP hardware has a fractional multiply/accumulate instruction that performs all these tasks in one cycle.

In the example code, the compiler recognizes this idiom and replaces the multiply followed by shift with a fractional multiply. In more complicated cases, where perhaps the multiply is further separated from the shift, the compiler may not detect the possibility of using a fractional multiply.

The recommended coding style is to use built-in functions. In the following example, a builtin is used to multiply fractional 16-bit data.

```
#include <math.h>
fract32 dot_product(fract16 *a, fract16 *b) {
   int i;
   fract32 sum=0;
   for (i=0; i<100; i++) {
      /* this line is performance critical */
      sum += __builtin_mult_fr1x32(a[i],b[i]);
   }
   return sum;
}
```

Good: uses builtins to implement fractional multiplication.

Note that the `fract16` and `fract32` types used in the example above are merely `typedefs` to C integer types used by convention in standard include files. The compiler does not have any in-built knowledge of these types and treats them exactly as the integer types that they are `typedefed` to.

System Support Built-In Functions

Built-in functions are also provided to perform low-level system management, in particular for the manipulation of system registers (defined in `sysreg.h`). It is usually better to use these built-in functions rather than inline `asm` statements. The built-in functions cause the compiler to generate efficient inline instructions and their use often results in better optimization of the surrounding code at the point where they are used. Using builtins will also usually result in improved code-readability. For more information on built-in functions supported by the compiler, refer to “Compiler Built-In Functions” on page 1-94.

Examples of the two styles are:

```
int read_io(int iopg_val) {
   int ret_val;
   ("IOPG=%1 ; %0=IO(0x20);" : "=e"(ret_val) : "e"(iopg_val) : "IOPG");
   return ret_val;
}
```

Bad: uses inline `asm` statement.

```
#include <sysreg.h>
#define ADDR 0x20
int read_io(int iopg_val) {
   sysreg_write(sysreg_IOPG, iopg_val);
   return sysreg_read(reg_CYCLES);
}
```

Good: uses `sysreg.h`.

This example reads and returns the `CYCLES` register.

Using Circular Buffers

Circular buffers are often extremely useful in DSP code. They can be used in several ways. Consider the C code:

```
for (i=0; i<1000; i++) {
   sum += a[i] * b[i%20];
}
```

Good: the compiler knows that `b` is accessed as a circular buffer.

Clearly the access to array `b` is a circular buffer. When optimization is enabled the compiler will produce a hardware circular buffer instruction for this access.

Consider the slightly more complicated example:

```
for (i=0; i<1000; i+=n) {
   sum += a[i] * b[i%20];
}
```

Bad: may not be able to use circular buffer to access `b`.

In this case, the compiler does not know if `n` is positive and less than 20. If it is, then the access may be correctly implemented as a hardware circular buffer. On the other hand, if it is greater than 20, a circular buffer increment may not yield the same results as the C code.

The programmer has two options here. One is to compile with the `-force-circbuf` switch (on page 1-27). This tells the compiler that any access of the form `a[i%n]` should be considered as a circular buffer. Before using this switch, you should check that this assumption is valid for your application.

The preferred option, however, is to use builtins to perform the circular buffering. Two are provided for this purpose. To make it clear to the compiler that a circular buffer should be used, you may write either

```
for (i=0, j=0; i<1000; i+=n) {
   sum += a[i] * b[j];
   j = __builtin_circindex(j, n, 20);
}
```

Good: explicit use of circular buffer via `__builtin_circindex`.

or

```
int *p = b;
for (i=0, j=0; i<1000; i+=n) {
   sum += a[i] * (*p);
   p = __builtin_circptr(p, n, b, 20);
}
```

Good: explicit use of circular buffer via `__builtin_circptr`.

For more information, see “Compiler Built-In Functions” on page 1-94.

Smaller Applications: Optimizing for Code Size

The same ethos for producing fast code also applies to producing small code. You should present the algorithm in a way that gives the optimizer excellent visibility of the operations and data, and hence the greatest freedom to safely manipulate the code to produce small applications.

Once the program is presented in this way, the optimization strategy will depend on the code-size constraint that the program must obey. The first step should be to optimize the application for full performance, using `-O` or `-ipa` switches. If this obeys the code-size constraints, then no more need be done.

The "optimize for space" switch `-Os` (on page 1-37). which may be used in conjunction with IPA, will perform every performance-enhancing transformation except those that increase code-size. In addition, the `-e` linker switch (`-flags-link -e` if used from the compiler command line) may be helpful (on page 1-27). This performs section elimination in the linker to remove unneeded data and code. If the code produced with `-Os` and `-e` does not meet the code-size constraint, some analysis of the source code will be required to try to reduce the code-size further.

Note that loop transformations such as unrolling and software pipelining increase code size. But it is these loop transformations that also give the greatest performance benefit. Therefore, in many cases compiling for minimum code size will produce significantly slower code than optimizing for speed.

The compiler provides a way to balance between the two extremes of `-O` and `-Os`. This is the sliding-scale `-Ov num` switch (adjustable using the optimization slider bar under **Project Options** in the VisualDSP++ IDDE), described on page 1-38. The *`num`* is a value between 0 and 100, where the lower value corresponds to minimum code size and the upper to maximum performance. A value in between will try to optimize the frequently executed regions of code for maximum performance, while keeping the infrequently executed parts as small as possible. The switch is most reliable when using profile-guided optimization, since the execution counts of the various code regions have been measured experimentally. Without PGO, the execution counts are estimated, based on the depth of loop nesting.

Tip: Avoid the use of inline code.

Avoid using the `inline` keyword to inline code for functions that are used a number of times, especially if they not very small functions. The `-Os` switch does not have any effect on the use of the `inline` keyword. It does, however, prevent automatic inlining (using the `-Oa` switch) from increasing the code size. Macro functions can also cause code expansion and should be used with care.

Pragmas

Pragmas can assist optimization by allowing the programmer to make assertions or suggestions to the compiler. This section looks at how they can be used to finely tune source code. Refer to "Pragmas" on page 1-119 for full details of how each pragma works; the emphasis here will be in considering under what circumstances they are useful during the optimization process.

In most cases, the pragmas serve to give the compiler information which it is unable to deduce for itself. It must be emphasized that the programmer is responsible for making sure that the information given by the pragma is valid in the context in which it is used. Use of a pragma to assert that a function or loop has a quality that it does not in fact have is likely to result in incorrect code and hence a malfunctioning application.

An advantage of the use of pragmas is that they allow code to remain portable, since they will normally be ignored by a compiler that does not recognize them.

Function Pragmas

Function pragmas include `#pragma const`, `#pragma pure`, `#pragma alloc`, `#pragma result_alignment`, `#pragma regs_clobbered`, and `#pragma optimize_{off|for_speed|for_space|as_cmd_line}`.

#pragma const

The `pragma const` pragma asserts to the compiler that a function does not have any side effects (such as modifying global variables or data buffers), and the result returned is only a function of the parameter values. The pragma may be applied to a function prototype or definition. It helps the compiler since two calls to the function with identical parameters will always yield the same result. This way, calls to `#pragma const` functions may be hoisted out of loops if their parameters are loop independent.

#pragma pure

Like #pragma const, this pragma asserts to the compiler that a function does not have any side effects (such as modifying global variables or data buffers). However, the result returned may be a function of both the parameter values and any global variables. The pragma may be applied to a function prototype or definition. Two calls to the function with identical parameters will always yield the same result provided that no global variables have been modified between the calls. Hence, calls to #pragma pure functions may be hoisted out of loops if their parameters are loop independent and no global variables are modified in the loop.

#pragma alloc

The pragma alloc pragma asserts that the function behaves like the malloc library function. In particular, it returns a pointer to new memory that cannot alias any pre-existing buffers. In the following code,

```
#pragma alloc
int *new_buf(void);
int *vmul(int *a, int *b) {
   int i;
   int *out = new_buf();
   for (i=0; i<100; i++)
      out[i] = a[i] * b[i];
}
```

Good: uses #pragma alloc to disambiguate out from a and b.

the use of the pragma allows the compiler to be sure that the write into buffer out does not modify either of the two input buffers a or b, and, therefore, the iterations of the loop may be re-ordered.

#pragma regs_clobbered

The `regs_clobbered` pragma is a useful way to improve the performance of code that makes function calls. The best use of the pragma is to increase the number of call-preserved registers available across a function call. There are two complementary ways in which this may be done.

First of all, suppose that you have a function written in assembly that you wish to call from C source code. The `regs_clobbered` pragma may be applied to the function prototype to specify which registers are "clobbered" by the assembly function, that is, which registers may have different values before and after the function call. Consider for example an simple assembly function to add two integers and mask the result to fit into 8 bits:

```
_add_mask:
   AY1 = DM(I4 +  1);
   AX1 = DM(I4 +  2);
   AY0 =  255;
   AR = AX1 + AY1;
   RTS (DB);
   AR = AR AND AY0;
   AX1 = AR;
._add_mask.end
```

Clearly the function does not modify the majority of the scratch registers available and thus these could instead be used as call-preserved registers. In this way fewer spills to the stack would be needed in the caller function. Using the prototype

```
#pragma regs_clobbered "AX1, AY0, AY1, AR,  ASTAT"
int add_mask(int, int);
```

Good: uses `regs_clobbered` to increase call-preserved register set.

the compiler is told which registers are modified by a call to the `add_mask` function. The registers not specified by the pragma are assumed to preserve their values across such a call and the compiler may use these spare registers to its advantage when optimizing the call sites.

The pragma is also powerful when all of the source code is written in C. In the above example, a C implementation might be:

```
int add_mask(int a, int b) {
   return ((a+b)&255);
}
```

Bad: function thought to clobber entire volatile register set.

Since this function will not need many registers when compiled, it can be defined using:

```
#pragma regs_clobbered "AX1, AY0, AY1, AR, M7, CCset"
int add_mask(int a, int b) {
   return ((a+b)&255);
}
```

Good: function compiled to preserve most registers.

to ensure that any other registers aside from `AX1`, `AY0`, `AY1`, `AR`, `M7` and the condition codes will not be modified by the function. If any other registers are used in the compilation of the function, they will be saved and restored during the function prologue and epilogue.

In general, it is not very helpful to specify any of the condition codes as call-preserved as they are difficult to save and restore and are usually clobbered by any function. Moreover, it is usually of limited benefit to be able to keep them live across a function call. Therefore, it is better to use `CCset` (all condition codes) rather than `ASTAT` in the clobbered set above. For more information, refer to "#pragma regs_clobbered string" on page 1-130.

#pragma optimize_{off|for_speed|for_space|as_cmd_line}

The `optimize_` pragma may be used to change the optimization setting on a function-by-function basis. In particular, it may be useful to optimize functions that are rarely called (for example, error handling code) for space (using `#pragma optimize_for_space`), whereas functions critical to performance should be compiled for maximum speed (`#pragma optimize_for_speed`). The `#pragma optimize_off` is useful for debugging specific functions without increasing the size or decreasing the performance of the overall application unnecessarily.

For more information, refer to "General Optimization Pragmas" on page 1-125.

Loop Optimization Pragmas

Many pragmas are targeted towards helping to produce optimal code for inner loops. These are the `loop_count` and `no_alias` pragmas.

#pragma loop_count

The `loop_count` pragma enables the programmer to inform the compiler about a loop's iteration count. The compiler is able to make more reliable decisions about the optimization strategy for a loop if it knows the iteration count range. If you know that the loop count is always a multiple of some constant, this can also be useful as it allows a loop to be partially unrolled or vectorized without the need for conditionally-executed iterations. Knowledge of the minimum trip count may allow the compiler to omit the guards that are usually required after software pipelining. Any of the parameters of the pragma that are unknown may be left blank.

An example of the use of the `loop_count` pragma might be:

```
#pragma loop_count(/*minimum*/ 40, /*maximum*/ 100, /*modulo*/ 4)
for (i=0; i<n; i++)
   a[i] = b[i];
```

Good: the `loop_count` pragma gives compiler helpful information to assist optimization.

For more information, refer to "#pragma loop_count(min, max, modulo)" on page 1-124.

#pragma no_alias

When immediately preceding a loop, the `no_alias` pragma asserts that no load or store in the loop accesses the same memory as any other. This helps to produce shorter loop kernels as it permits instructions in the loop to be rearranged more freely.

See "#pragma no_alias" on page 1-125 for more information.

Useful Optimization Switches

Table 2-1 lists the compiler switches useful during the optimization process.

Table 2-1. C/C++ Compiler Optimization Switches

Switch Name	Description
`-const-read-write` on page 1-24	Specifies that data accessed via a pointer to `const` data may be modified elsewhere.
`-flags-link -e` on page 1-27	Specifies linker section elimination.
`-force-circbuf` on page 1-27	Treats array references of the form `array[i%n]` as circular buffer operations.
`-ipa` on page 1-30	Turns on inter-procedural optimization. Implies use of `-O`. May be used in conjunction with `-Os` or `-Ov`.
`-no-fp-associative` on page 1-34	Does not treat floating-point multiply and addition as an associative.
`-O` on page 1-37	Enables code optimizations and optimizes the file for speed.
`-Os` on page 1-37	Optimizes the file for size.
`-Ov num` on page 1-38	Controls speed vs. size optimizations (sliding scale).
`-save-temps` on page 1-44	Saves intermediate files (for example, `.s`).

3 C/C++ RUN-TIME LIBRARY

The C and C++ run-time libraries are collections of functions, macros, and class templates that you can call from your source programs. The libraries provide a broad range of services including those that are basic to the languages such as memory allocation, character and string conversions, and math calculations. Using the library simplifies your software development by providing code for a variety of common needs.

This chapter contains

- "C and C++ Run-Time Library Guide" on page 3-3
  It provides introductory information about the ANSI/ISO standard C and C++ libraries. It also provides information about the ANSI standard header files and built-in functions that are included with this release of the `cc219x` compiler.

- "Documented Library Functions" on page 3-23
  It tabulates the functions that are defined by ANSI standard header files.

- "C Run-Time Library Reference" on page 3-26"
  It provides reference information about the C run-time library functions included with this release of the `cc219x` compiler.

The `cc219x` compiler provides a broad collection of library functions including those required by the ANSI standard and functions supplied by Analog Devices that are of value in signal processing applications. In addition to the Standard C Library, this release of the compiler software includes the abridged C++ library, a conforming subset of the standard C++ library. The abridged C++ library includes the embedded C++ and embedded standard template libraries.

This chapter describes the standard C/C++ library functions in the current release of the run-time library. Chapter 4, "DSP Run-Time Library" describes a number of signal processing, matrix, and statistical functions that assist DSP code development.

For more information on the C standard library, see The Standard C Library by P.J. Plauger, Prentice Hall, 1992. For more information on the algorithms on which many of the C library's math functions are based, see Cody, W. J. and W. Waite, *Software Manual for the Elementary Functions*, Englewood Cliffs, New Jersey: Prentice Hall, 1980. For more information on the C++ library portion of the ANSI/ISO Standard for C++, see Plauger, P. J. (Preface), *The Draft Standard C++ Library*, Englewood Cliffs, New Jersey: Prentice Hall, 1994, (ISBN: 0131170031).

The C++ library reference information in HTML format is included on the software distribution CD-ROM. To access the reference files from VisualDSP++, use the **Help Topics** command (**Help** menu) and select the **Reference** book icon. From the **Online Manuals** topic, you can open any of the library files. You can also manually access the HTML files using a web browser.

C and C++ Run-Time Library Guide

The C and C++ run-time libraries contain routines that you can call from your source program. This section describes how to use the libraries and provides information on the following topics:

- "Calling Library Functions" on page 3-4
- "Using the Compiler's Built-In C Library Functions" on page 3-5
- "Linking Library Functions" on page 3-6
- "Working With Library Header Files" on page 3-8
- "Abridged C++ Library Support" on page 3-16

For information on the C library's contents, see "Documented Library Functions" on page 3-23. For information on the Abridged C++ library's contents, see "Abridged C++ Library Support" on page 3-16 and on-line Help.

Calling Library Functions

To use a C/C++ library function, call the function by name and give the appropriate arguments. The name and arguments for each function appear on the function's reference page. The reference pages appear in the "Documented Library Functions" on page 3-23 and in the **C++ Run-Time Library** topic of the on-line Help.

Like other functions you use, library functions should be declared. Declarations are supplied in header files. For more information about the header files see "Working With Library Header Files" on page 3-8.

Function names are C/C++ function names. If you call a C or C++ run-time library function from an assembler program, you must use the assembly version of the function name.

- For C functions, this is an underscore at the beginning of the C function name. For example, the C function `main()` is referred to as `_main` from assembler.
- Functions in C++ modules are normally compiled with an encoded function name. Function names in C++ contain abbreviations for the parameters to the function and also the return type. As such, they can become very large. The compiler "mangles" these names to a shorter form. You can instruct the C++ compiler to use the single-underscore convention from C, as shown by the following example.

```
extern "C" {
    int myfunc(int);   // external name is _myfunc
}
```

Alternatively, compile C++ files to assembler, and see how the function has been declared in the assembly file.

It may not be possible to call inline functions from an assembler program as the compiler may have removed the definition of the function if all calls to the function are inlined. Global static variables cannot be referred to in assembly routines as their names are encrypted. For more information on the naming conventions, see "C/C++ and Assembly Language Interface" on page 1-169.

You can use `elfar` (the archiver), described in the *VisualDSP++ 3.5 Linker and Utilities Manual for 16-Bit Processors*, to build library archive files of your own functions.

Using the Compiler's Built-In C Library Functions

The C/C++ compiler's built-in functions are a set of functions that the compiler immediately recognizes and replaces with inline assembly code instead of a function call. Typically, in-line assembly code is faster than an library routine, and it does not incur the calling overhead.

To use built-in functions, your source must include the required standard include file. For the abs functions this would require `stdlib.h` to be included. There are built-in functions used to define some ANSI C `math.h`, `string.h` and `stdlib.h` functions. There are also built-in functions to support various ANALOG extensions to the ANSI standard defined in the include file `math_builtins.h`. Not all built-in functions have a library alternate definition. Therefore, the failure to use the required include files can result in your program build failing to link.

If you want to use the C run-time library functions of the same name, compile with the `-no-builtin` compiler switch (on page 1-34).

Linking Library Functions

The C/C++ run-time library is organized as several libraries which are catalogued in Table 3-1. The libraries and start-up files are installed within the subdirectory ...\219x\lib of your VisualDSP++ installation.

Several library files are built twice: once for single-threaded environments, and once for multi-threaded environments. The multi-threaded versions have a "mt" suffix. Another variant for some of the libraries is also provided that avoids a hardware anomaly involving instruction type 32a; these libraries have a "_type32aworkaround" suffix.

When you call a run-time library function, the call creates a reference that the linker resolves. One way to direct the linker to the library's location is to use the default Linker Description File (ADSP-21<your_target>.ldf).

If you are not using the default .LDF file, then either add the appropriate library/libraries to the .LDF file used for your project, or use the compiler's -l switch ("-l library" on page 1-31) to specify the library to be added to the link line.

For example, the switches -lc -letsi will add the C library libc.dlb and the ETSI support library libetsi.dlb to the list of libraries to be searched by the linker. For more information on the .LDF file, see the *VisualDSP++ 3.5 Linker and Utilities Manual for 16-Bit Processors*.

Table 3-1 briefly describes the ADSP-219x DSP library functions.

Table 3-1. C and C++ Library Files

219x\lib Directory	Description
2192-12_int_tab.doj	Default interrupt vector code for ADSP-2192-12 DSP
219x_int_tab.doj	Default interrupt vector code for ADSP-219x DSPs
219x_hdr.doj	Startup file — set-up routines and call main()

Table 3-1. C and C++ Library Files (Cont'd)

219x\lib Directory	Description
`219x_cpp_hdr.doj` `219x_cpp_mt_hdr.doj`	C++ startup file — set-up routines and call main()
`219x_exit.doj`	Dummy exit object for backwards compatibility.
`219x_ezkit_hdr.doj`	Startup file — set-up routines, call `main()`, and enable use of the EZ-kit monitor program
`219x_cpp_ezkit_hdr.doj` `219x_cpp_mt_ezkit_hdr.doj`	C++ start-up file — set-up routines, call `main()`, and enable use of the EZ-kit monitor program

`libc.dlb`
`libcmt.dlb`
`libc_type32aworkaround.dlb`
`libcmt_type32aworkaround.dlb`
	C run-time library

`libcpp.dlb`
`libcppmt.dlb`
`libcpp_type32aworkaround.dlb`
`libcppmt_type32aworkaround.dlb`
	C++ run-time library

`libcpprt.dlb`
`libcpprtmt.dlb`
`libcpprt_type32aworkaround.dlb`
`libcpprtmt_type32aworkaround.dlb`
	C++ run-time support library

`libdsp.dlb` `libdsp_type32aworkaround.dlb`	DSP library
`libetsi.dlb`	ETSI run-time library

`libio.dlb`
`ibiomt.dlb`
`libio_type32aworkaround.dlb`
`libiomt_type32aworkaround.dlb`
	I/O library

`libsim.dlb`	Simulator library support

Ensure that the C++ library of run-time routines, `libcpprt.dlb`, is the last library specified on the link line.

If all the objects supplied to the driver have been built as C, but are referencing a C++ object which is in a library, the standard C++ libraries are not searched and the linker may issue an error concerning unresolved symbol(s). This can be avoided by using the `flags-link` switch (see on page 1-27), which ensures that that the C++ libraries are linked from the default `.LDF` files.

For example,

```
flags-link -MD__cplusplus=1
```

Note that this problem will only occur if the C++ object is in a library. If it is in an object file, the compiler will recognize it as a C++ object and link with the C++ libraries.

Working With Library Header Files

When you use a library function in your program, you should also include the function's header file with the `#include` preprocessor command. The header file for each function is identified in the **Synopsis** section of the function's reference page. Header files contain function prototypes. The compiler uses these prototypes to check that each function is called with the correct arguments.

A list of the header files that are supplied with this release of the `cc219x` compiler appears in Table 3-2. You should use a C standard text to augment the information supplied in this chapter.

Table 3-2. C Run-Time Library Header Files

Header	Purpose	Standard
assert.h	Diagnostics	ANSI
ctype.h	Character Handling	ANSI
def2191.h	Memory Map Register and System Definitions for ADSP-2191 DSPs	C Extension

Table 3-2. C Run-Time Library Header Files (Cont'd)

Header	Purpose	Standard
def2192-12.h	Memory Map Register and System Definitions for ADSP-2192-12 DSPs	C Extension
def219x.h	Memory Map Definitions	C Extension
errno.h	Error Handling	ANSI
float.h	Floating Point	ANSI
iso646.h	Boolean Operators	ANSI
limits.h	Limits	ANSI
locale.h	Localization	ANSI
math.h	Mathematics	ANSI
setjmp.h	Non-Local Jumps	ANSI
signal.h	Signal Handling	ANSI
stdarg.h	Variable Arguments	ANSI
stddef.h	Standard Definitions	ANSI
stdio.h	Input/Output	ANSI
stdlib.h	Standard Library	ANSI
string.h	String Handling	ANSI
sysreg.h	Efficient system access	C Extension

The following sections provide descriptions of the header files contained in the C library. The header files are listed in alphabetical order.

assert.h

The `assert.h` header file contains the `assert` macro.

ctype.h

The `ctype.h` header file contains functions for character handling, such as `isalpha`, `tolower`, etc.

def2191.h – Memory Map Definitions

The `def2191.h` header file contains macro definitions for the ADSP-2191 processor's system addresses, and system register bits. These symbolic names can be used in programs to access specific system registers and system register bits in the ADSP-2191 DSP.

def2192-12.h – Memory Map Definitions

The `def2192-12.h` header file contains macro definitions for the ADSP-2191-12 processor's system addresses, and system register bits. These symbolic names can be used in programs to access specific system registers and system register bits in the ADSP-2191-12 DSP.

def219x.h– Memory Map Definitions

The `def219x.h` header file contains macro definitions for a ADSP-219x processor's system addresses, and system register bits. These symbolic names can be used in programs to access specific system registers and system register bits in ADSP-219x DSPs.

errno.h

The `errno.h` header file provides access to `errno`. This facility is not, in general, supported by the rest of the library.

float.h

The `float.h` file defines the format of floating-point data types. The `FLT_ROUNDS` macro, defined in the header file, is set to the C run-time environment definition of the rounding mode for `float` variables, which is *round-towards-nearest*.

iso646.h

The `iso646.h` header file defines symbolic names for certain C operators; the symbolic names and their associated value are shown in Table 3-3.

Table 3-3. Symbolic Names Defined in iso646.h

<table>
<tr><th>Symbolic Name</th><th>Equivalent</th></tr>
<tr><td><code>and</code></td><td>&&</td></tr>
<tr><td><code>and_eq</code></td><td>&=</td></tr>
<tr><td><code>bitand</code></td><td>&</td></tr>
<tr><td><code>bitor</code></td><td>|</td></tr>
<tr><td><code>compl</code></td><td>~</td></tr>
<tr><td><code>not</code></td><td>!</td></tr>
<tr><td><code>not_eq</code></td><td>!=</td></tr>
<tr><td><code>or</code></td><td>||</td></tr>
<tr><td><code>or_eq</code></td><td>|=</td></tr>
<tr><td><code>xor</code></td><td>^</td></tr>
<tr><td><code>xor_eq</code></td><td>^=</td></tr>
</table>

The symbolic names have the same name as the C++ keywords that are accepted by the compiler when the `-alttok` switch (see on page 1-23) is specified.

limits.h

The `limits.h` header file contains definitions of maximum and minimum values for each C data type other than floating-point.

locale.h

The `locale.h` header file contains definitions for expressing numeric, monetary, time, and other data.

math.h

This category includes the floating-point mathematical functions of the C run-time library. The mathematical functions are ANSI standard. The `math.h` header file contains prototypes for functions used to calculate mathematical properties of single-precision floating type variables. On the ADSP-219x processors, `double` and `float` are both single-precision floating point types. Additionally, some functions support a 16-bit fractional data type.

The `math.h` file also defines the macro `HUGE_VAL`. `HUGE_VAL` evaluates to the maximum positive value that the type `double` can support. The macros `EDOM` and `ERANGE`, defined in `errno.h`, are used by `math.h` functions to indicate domain and range errors.

Some of the functions in this header file exist as both integer and floating point. The floating-point functions typically have an `f` prefix. Make sure you are using the correct one.

ⓘ The C language provides for implicit type conversion, so the following sequence produces surprising results with no warnings:

```
float x,y;
y = abs(x);
```

The value in `x` is truncated to an integer prior to calculating the absolute value, then reconverted to floating point for the assignment to `y` .

setjmp.h

The `setjmp.h` header file contains `setjmp` and `longjmp` for non-local jumps.

signal.h

The `signal.h` header file provides function prototypes for the standard ANSI `signal.h` routines and also for several ADSP-219x DSP's extensions such as `interrupt()` and `clear_interrupt()`.

The signal handling functions process conditions (hardware signals) that can occur during program execution. They determine the way that your C program responds to these signals. The functions are designed to process such signals as external interrupts and timer interrupts.

stdarg.h

The `stdarg.h` header file contains definitions needed for functions that accept a variable number of arguments. Programs that call such functions must include a prototype for the functions referenced.

stddef.h

The `stddef.h` header file contains a few common definitions useful for portable programs, such as `size_t`.

stdio.h

The `stdio.h` header file defines a set of functions, macros, and data types for performing input and output. Applications that use the facilities of this header file should link with the I/O library `libio.dlb` in the same way as linking with the C run-time library `libc.dlb`. The library is thread-safe but it is not interrupt-safe and should not therefore be called either directly or indirectly from an interrupt service routine.

The implementation of the stdio.h routines is based on a simple interface with a device driver that provides a set of low-level primitives for open, close, read, write, and seek operations. By default, these operations are provided by the VisualDSP++ simulator and EZ-KIT Lite systems. However, alternative device drivers may be registered (see "Extending I/O Support To New Devices" on page 1-166) that can then be used through the stdio.h functions.

The following restrictions apply to this software release:

- Functions tmpfile() and tmpnam() are not available,
- Functions rename() and remove() are only supported under the default device driver supplied by the VisualDSP++ simulator and EZ-kits, and they only operate on the host file system,
- Positioning within a file that has been opened as a text stream is only supported if the lines within the file are terminated by the character sequence \r\n.

When using the default device driver, all I/O operations are channeled through the C function _primIO(). The assembly label has two underscores, __primIO. The _primIO() function accepts no arguments. Instead, it examines the I/O control block at label _primIOCB. Without external intervention by a host environment, the _primIO routine simply returns, which indicates failure of the request.

When the host environment is providing I/O support, the host places a breakpoint at the start of _primIO(). Upon entry to _primIO(), the data for the request will reside in a control block at the label _primIOCB. The host arranges to intercept control when it enters the _primIO() routine, and, after servicing the request, returns control to the calling routine. See "File I/O Support" on page 1-166 for more information.

At program termination, the host environment will close down any physical connection between the application and an opened file. However, the I/O library will not implicitly close any opened streams to avoid an unnecessary overheads (particularly with respect to memory occupancy).

Therefore, unless explicit action is taken by an application any unflushed output may be lost. Any output generated by `printf` is always flushed but output generated by other library functions, such as `putchar`, `fwrite`, `fprintf`, will not be automatically flushed. Applications should therefore arrange to close down any streams that they open. Note that the function reference `fflush (NULL);` will flush the buffers of all opened streams.

stdlib.h

The `stdlib.h` header file contains general utilities specified by the C standard. These include some integer math functions, such as `abs`, `div`, and `rand`; general string-to-numeric conversions; memory allocation functions, such as `malloc` and `free`, and termination functions, such as `exit`. This header file also contains prototypes for miscellaneous functions such as `bsearch` and `qsort`.

string.h

The `string.h` header file contains string handling functions, including `strcpy` and `memcpy`.

sysreg.h

The `sysreg.h` header file defines a set of functions that provide efficient system access to registers, modes and addresses not normally accessible from C source. See “Compiler Built-In Functions” on page 1-94 for more information on these functions.

Abridged C++ Library Support

When in C++ mode, the `cc219x` compiler can call a large number of functions from the Abridged Library, a conforming subset of the C++ library.

The Abridged Library has two major components—embedded C++ library (EC++) and embedded standard template library (ESTL). The embedded C++ library is a conforming implementation of the Embedded C++ library, as specified by the Embedded C++ Technical Committee.

This section lists and briefly describes the following components of the Abridged Library:

- "Embedded C++ Library Header Files" on page 3-16
- "C++ Header Files for C Library Facilities" on page 3-19
- "Embedded Standard Template Library Header Files" on page 3-20

For more information on the Abridged Library, see online Help.

Embedded C++ Library Header Files

The following sections provide a brief description of the header files in the embedded C++ library

complex

The `complex` header file defines a template class complex and a set of associated arithmetic operators. Predefined types include `complex_float` and `complex_long_double`.

This implementation does not support the full set of complex operations as specified by the C++ standard. In particular, it does not support either the transcendental functions or the I/O operators << and >>.

The `complex` header and the C library header file `complex.h` refer to two different and incompatible implementations of the complex data type.

exception

The `exception` header defines the `exception` and `bad_exception` classes and several functions for exception handling.

fract

The `fract` header defines the `fract` data type, which supports fractional arithmetic, assignment, and type-conversion operations. The header file is fully described under "Fractional Type Support" on page 1-90, and an example that demonstrates its use appears under "C++ Programming Examples" on page 1-176.

fstream

The `fstream` header defines the `filebuf`, `ifstream`, and `ofstream` classes for external file manipulations.

iomanip

The `iomanip` header declares several `iostream` manipulators. Each manipulator accepts a single argument.

ios

The `ios` header defines several classes and functions for basic `iostream` manipulations.

Most of the `iostream` header files include `ios`.

iosfwd

The `iosfwd` header declares forward references to various `iostream` template classes defined in other standard headers.

iostream

The `iostream` header declares most of the `iostream` objects used for the standard stream manipulations.

istream

The `istream` header defines the `istream` class for `iostream` extractions.

Most of the iostream header files include `istream`.

new

The `new` header declares several classes and functions for memory allocations and deallocations.

ostream

The `ostream` header defines the `ostream` class for `iostream` insertions.

sstream

The `sstream` header defines the `stringbuf`, `istringstream`, and `ostringstream` classes for various `string` object manipulations.

stdexcept

The `stdexcept` header defines a variety of classes for `exception` reporting.

streambuf

The `streambuf` header defines the `streambuf` classes for basic operations of the `iostream` classes.

Most of the iostream header files include `streambuf`.

string

The `string` header defines the `string` template and various supporting classes and functions for `string` manipulations.

> Objects of the `string` type should not be confused with the null-terminated C strings.

strstream

The `strstream` header defines the `strstreambuf`, `istrstream`, and `ostream` classes for `iostream` manipulations on allocated, extended, and freed character sequences.

C++ Header Files for C Library Facilities

For each C standard library header there is a corresponding standard C++ header. If the name of a C standard library header file is `foo.h`, then the name of the equivalent C++ header file will be `cfoo`. For example, the C++ header file `cstdio` provides the same facilities as the C header file `stdio.h`.

Normally, the C standard headers files may be used to define names in the C++ global namespace while the equivalent C++ header files define names in the standard namespace. However, the standard namespace is not supported in this release of the compiler, and the effect of including one of the C++ header files listed in Table 3-4 is the same as including the equivalent C standard library header file.

Table 3-4 lists the C++ header files that provide access to the C library facilities.

Table 3-4. C++ Header Files for C Library Facilities

Header	Description
`cassert`	Enforces assertions during function executions.
`cctype`	Classifies characters.

Table 3-4. C++ Header Files for C Library Facilities (Cont'd)

Header	Description
cerrno	Tests error codes reported by library functions.
cfloat	Tests floating-point type properties.
climits	Tests integer type properties.
clocale	Adapts to different cultural conventions.
cmath	Provides common mathematical operations.
csetjmp	Executes non-local goto statements.
csignal	Controls various exceptional conditions.
cstdarg	Accesses a variable number of arguments.
cstddef	Defines several useful data types and macros.
cstdio	Performs input and output.
cstdlib	Performs a variety of operations
cstring	Manipulates several kinds of strings

Embedded Standard Template Library Header Files

Templates and the associated header files are not part of the embedded C++ standard, but they are supported by the `cc219x` compiler in C++ mode.

The embedded standard template library headers are:

algorithm

The `algorithm` header defines numerous common operations on sequences.

deque

The `deque` header defines a deque template container.

functional

The `functional` header defines numerous function objects.

hash_map

The `hash_map` header defines two hashed map template containers.

hash_set

The `hash_set` header defines two hashed set template containers.

iterator

The `iterator` header defines common iterators and operations on iterators.

list

The `list` header defines a list template container.

map

The `map` header defines two map template containers.

memory

The `memory` header defines facilities for managing memory.

numeric

The `numeric` header defines several numeric operations on sequences.

queue

The `queue` header defines two queue template container adapters.

set

The `set` header defines two set template containers.

stack

The `stack` header defines a stack template container adapter.

utility

The `utility` header defines an assortment of utility templates.

vector

The `vector` header defines a vector template container.

The embedded C++ library also includes several headers for compatibility with traditional C++ libraries, such as:

fstream.h

The `fstream.h` header defines several `iostreams` template classes that manipulate external files.

iomanip.h

The `iomanip.h` header declares several `iostreams` manipulators that take a single argument.

iostream.h

The `iostream.h` header declares the `iostreams` objects that manipulate the standard streams.

new.h

The `new.h` header declares several functions that allocate and free storage.

Documented Library Functions

The following tables list the library functions documented in this chapter.

These tables list the functions by the header file in which they are located, whereas the "C Run-Time Library Reference" on page 3-26 presents the functions in alphabetic order.

Table 3-5. Documented Library Functions in the `ctype.h` Header File

isalnum	isalpha	iscntrl
isdigit	isgraph	islower
isprint	ispunct	isspace
isupper	isxdigit	tolower
toupper		

Table 3-6. Documented Library Functions in the `math.h` Header File

acos	asin	atan
atan2	ceil	cos
cosh	exp	fabs
floor	fmod	frexp
isinf	isnan	ldexp
log	log10	modf
pow	sin	sinh
sqrt	tan	tanh

Table 3-7. Documented Library Functions in the `setjmp.h` Header File

longjmp	setjmp

Table 3-8. Documented Library Functions in the `signal.h` Header File

clear_interrupt	interrupt	raise
signal		

Table 3-9. Documented Library Functions in the `stdarg.h` Header File

va_arg	va_end	va_start

Table 3-10. Supported Library Functions in the `stdio.h` Header File

clearerr	fclose	feof
ferror	fflush	fgetc
fgetpos	fgets	fprintf
fputc	fputs	fopen
fread	freopen	fscanf
fseek	fsetpos	ftell
fwrite	getc	getchar
gets	perror	putc
putchar	puts	remove
rename	rewind	scanf
setbuf	setvbuf	sprintf
sscanf	ungetc	vfprintf
vprintf	vsprintf	

Table 3-11. Documented Library Functions in `stdlib.h` Header File

abort	abs	atexit
atof	atoi	atol
bsearch	calloc	div
exit	free	labs

Table 3-11. Documented Library Functions in `stdlib.h` Header File

ldiv	malloc	qsort
rand	realloc	srand
strtod	strtodf	strtol
strtoul		

Table 3-12. Documented Library Functions in `string.h` Header File

memchr	memcmp	memcpy
memcpy_from_shared	memcpy_to_shared	memmove
memset	strncat	strchr
strcmp	strcoll	strcpy
strcspn	strerror	strlen
strncat	strncmp	strncpy
strpbrk	strrchr	strspn
strstr	strtok	strxfrm

Table 3-13. Documented Library Functions in `sysreg.h` Header File

disable_interrupts	enable_interrupts	external_memory_read
external_memory_write	io_space_read	io_space_write
mode_change	sysreg_read	sysreg_write

C Run-Time Library Reference

The C run-time library is a collection of functions that you can call from your C programs.

The information that follows applies to all of the functions in the library.

Notation Conventions

An interval of numbers is indicated by the minimum and maximum, separated by a comma, and enclosed in two square brackets, two parentheses, or one of each. A square bracket indicates that the endpoint is included in the set of numbers; a parenthesis indicates that the endpoint is not included.

Reference Format

Each function in the library has a reference page. These pages have the following format:

Name and Purpose of the function

Synopsis—Required header file and functional prototype

Description—Function specification

Error Conditions—How the function indicates an error

Example—Typical function usage

See Also—Related functions

abort

abnormal program end

Synopsis

```
#include <stdlib.h>
void abort(void);
```

Description

The `abort` function causes an abnormal program termination by raising the `SIGABRT` signal. If the `SIGABRT` handler returns, `abort()` calls `exit()` to terminate the program with a failure condition.

Error Conditions

The `abort` function does not return.

Example

```
#include <stdlib.h>
extern int errors;

if(errors)       /* terminate program if */
   abort();      /* errors are present   */
```

See Also

atexit, exit

abs

absolute value

Synopsis

```
#include <stdlib.h>
int abs(int j);
```

Description

The `abs` function returns the absolute value of its `int` input. The `abs` function is implemented through a built-in. The built-in causes the compiler to emit an inline instruction to perform the required function at the point where `abs` is called.

`abs(INT_MIN)` returns `INT_MIN`.

Error Conditions

The `abs` function does not return an error condition.

Example

```
#include <stdlib.h>
int i;
i = abs(-5);       /* i == 5 */
```

See Also

fabs, labs

acos

arc cosine

Synopsis

```
#include <math.h>
double acos(double x);
float acosf (float x);
fract16 acos_fr16 (fract16 x);
```

Description

The `acos` function returns the arc cosine of the argument. The input must be in the range [-1, 1]. The output, in radians, is in the range [0, π].

The `acos_fr16` function is only defined for input values between 0 and 0.9 (=`0x7333`). The input argument is in radians. Output values range from `acos(0)`*2/π (=`0x7FFF`) to `acos(0.9)`*2/π (=`0x24C1`).

Error Conditions

The `acos` function returns a zero if the input is not in the defined range.

Example

```
#include <math.h>
double y;
y = acos(0.0);       /* y = π/2 */
```

See Also

cos

asin

arc sine

Synopsis

```
#include <math.h>
double asin(double x);
float asinf (float x);
fract16 asin_fr16(fract16 x);
```

Description

The `asin` function returns the arc sine of the argument. The input must be in the range [-1, 1]. The output, in radians, is in the range -π/2 to π/2.

The `asin_fr16` function is only defined for input values between -0.9 (=`0x8ccd`) and 0.9 (=`0x7333`). The input argument is in radians. Output values range from `asin(-0.9)`*2/π (=`0xa4C1`) to `asin(0.9)`*2/π (=`0x5B3F`).

Error Conditions

The `asin` function returns a zero if the input is not in the defined range.

Example

```
#include <math.h>
double y;
y = asin(1.0);      /* y = π/2 */
```

See Also

sin

atan

arc tangent

Synopsis

```
#include <math.h>
double atan(double x);
float atanf (float x);
fract16 atan_fr16 (fract16 x);
```

Description

The `atan` function returns the arc tangent of the argument. The output, in radians, is in the range -π/2 to π/2.

The `atan_fr16` function covers the output range from -π/4 (input value `0x8000`, output value `0x9B78`) to π/4 (input value `0x7FFF`, output value `0x6488`). The input argument is in radians.

Error Conditions

The `atan` function does not return an error condition.

Example

```
#include <math.h>
double y;
y = atan(0.0);       /* y = 0.0 */
```

See Also

atan2, tan

atan2

arc tangent of quotient

Synopsis

```
#include <math.h>
double atan2 (double x, double y);
float atan2f (float x, float y);
fract16 atan2 (fract16 x, fract16 y);
```

Description

The `atan2` function computes the arc tangent of the input value `x` divided by input value `y`. The output, in radians, is in the range [-π, π].

The `atan2_fr16` function uses the full range from -π/4 to π/4 (`0x8000` to `0x7FFF`) for both input and output arguments. This corresponds to a scaling by π compared to the floating-point function. The input argument is in radians.

Error Conditions

The `atan2` function returns a zero if x = 0 and y <> 0.

Example

```
#include <math.h>
double a;
float b;

a = atan2 (0.0, 0.5);      /* the error condition: a = 0.0 */
b = atan2f (1.0, 0.0);     /* b = π/2                      */
```

See Also

atan, tan

atexit

register a function to call at program termination

Synopsis

```
#include <stdlib.h>
int atexit(void (*func)(void));
```

Description

The `atexit` function registers a function to be called at program termination. Functions are called once for each time they are registered, in the reverse order of registration. Up to 32 functions can be registered using `atexit`.

Error Conditions

The `atexit` function returns a non-zero value if the function cannot be registered.

Example

```
#include <stdlib.h>

extern void goodbye(void);

if (atexit(goodbye))
    exit(1);
```

See Also

abort, exit

atof

convert string to a double

Synopsis

```
#include <stdlib.h>
double atof(const char *nptr);
```

Description

The `atof` function converts a character string into a floating-point value of type `double`, and returns its value. The character string is pointed to by the argument `nptr` and may contain any number of leading whitespace characters (as determined by the function `isspace`) followed by a floating-point number. The floating-point number may either be of the form of a decimal floating-point number or a hexadecimal floating-point number.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and `digits` are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (-) followed by the hexadecimal prefix `0x` or `0X` . This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter `p` or `P` , an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens `[hexdigs] [.hexdigs]`.

The first character that does not fit either form of number will stop the scan.

Error Conditions

The `atof` function returns a zero if no conversion could be made. If the correct value results in an overflow, a positive or negative (as appropriate) `HUGE_VAL` is returned. If the correct value results in an underflow, 0.0 is returned. The `ERANGE` value is stored in `errno` in the case of either an overflow or underflow.

Notes

The function reference `atof (pdata)` is functionally equivalent to:

```
strtod (pdata, (char *) NULL);
```

and therefore, if the function returns zero, it is not possible to determine whether the character string contained a (valid) representation of 0.0 or some invalid numerical string.

Example

```
#include <stdlib.h>
double x;

x = atof("5.5");       /* x == 5.5 */
```

See Also

atoi, atol, strtod

atoi

convert string to integer

Synopsis

```
#include <stdlib.h>
int atoi(const char *nptr);
```

Description

The `atoi` function converts a character string to an integer value. The character string to be converted is pointed to by the input pointer, `nptr`. The function clears any leading characters for which `isspace` would return true. Conversion begins at the first digit (with an optional preceding sign) and terminates at the first non-digit.

Error Conditions

The `atoi` function returns a zero if no conversion can be made.

Example

```
#include <stdlib.h>
int i;

i = atoi("5");      /* i == 5 */
```

See Also

atol, atof, strtod, strtol, strtoul

atol

convert string to long integer

Synopsis

```
#include <stdlib.h>
long atol(const char *nptr);
```

Description

The `atol` function converts a character string to a long integer value. The character string to be converted is pointed to by the input pointer, `nptr`. The function clears any leading characters for which `isspace` would return true. Conversion begins at the first digit (with an optional preceding sign) and terminates at the first non-digit.

There is no way to determine if a zero is a valid result or an indicator of an invalid string.

Error Conditions

The `atol` function returns a zero if no conversion can be made.

Example

```
#include <stdlib.h>
long int i;

i = atol("5");       /* i == 5 */
```

See Also

atoi, atof, strtod, strtol, strtoul

bsearch

perform binary search in a sorted array

Synopsis

```
#include <stdlib.h>
void *bsearch(const void *key, const void *base,
      size_t nelem, size_t size,
      int (*compare)(const void *, const void *));
```

Description

The `bsearch` function executes a binary search operation on a pre-sorted array, where:

- `key` is a pointer to the element to search for
- `base` points to the start of the array
- `nelem` is the number of elements in the array
- `size` is the size of each element of the array
- `*compare` points to the function used to compare two elements. It takes as parameters a pointer to the `key` and a pointer to an array element and should return a value less than, equal to, or greater than zero, according to whether the first parameter is less than, equal to, or greater than the second.

The `bsearch` function returns a pointer to the first occurrence of `key` in the array.

Error Conditions

The `bsearch` function returns a null pointer if the key is not found in the array.

Example

```
#include <stdlib.h>
char *answer;
char base[50][3];

answer = bsearch("g", base, 50, 3, strcmp);
```

See Also

qsort

calloc

allocate and initialize memory

Synopsis

```
#include <stdlib.h>
void *calloc(size_t nmemb, size_t size);
```

Description

The `calloc` function dynamically allocates a range of memory and initializes all locations to zero. The number of elements (the first argument) multiplied by the size of each element (the second argument) is the total memory allocated. The memory may be deallocated with the `free` function.

Error Conditions

The `calloc` function returns a null pointer if unable to allocate the requested memory.

Example

```
#include <stdlib.h>
int *ptr;

ptr = (int *) calloc(10, sizeof(int));
      /* ptr points to a zeroed array of length 10 */
```

See Also

free, malloc, realloc

ceil

ceiling

Synopsis

```
#include <math.h>
double ceil(double);
float ceilf(float);
```

Description

The `ceil` functions return the smallest integral value, that is not less than its input.

Error Conditions

The `ceil` functions do not return an error condition.

Example

```
#include <math.h>
double y;
y = ceil (1.05);          /* y = 2.0  */
y = ceilf (-1.05);        /* y = -1.0 */
```

See Also

floor

clear_interrupt

clear a pending signal

Synopsis

```
#include <signal.h>
int clear_interrupt(int sig);
```

Description

The `clear_interrupt` function clears the signal `sig` in the `IRPTL` register. The `sig` argument must be one of the processor signals shown below for the ADSP-219x DSPs.

Table 3-14. ADSP-219x Signals

Sig Value	Definition
SIG_PWRDWN	power down interrupt
SIG_STACKINT	`PC`, `LOOP`, or `COUNTER` overflow on `push`, or on `pop` when empty
SIG_KERNEL	kernel interrupt
SIG_INT4	user-assignable
SIG_INT5	user-assignable
SIG_INT6	user-assignable
SIG_INT7	user-assignable
SIG_INT8	user-assignable
SIG_INT9	user-assignable
SIG_INT10	user-assignable
SIG_INT11	user-assignable
SIG_INT12	user-assignable
SIG_INT13	user-assignable

Error Conditions

The `clear_interrupt` function returns a `1` if the interrupt was pending, a `-1` if the parameter is not a valid signal, or `0` is returned otherwise.

Example

```
#include <signal.h>
clear_interrupt(SIG_PWRDWN);
/* clear the interrupt 2 latch */
```

See Also

interrupt, raise, signal

cos

cosine

Synopsis

```
#include <math.h>
double cos(double);
float cosf (float);
fract16 cos_fr16 (fract16);
```

Description

The `cos` function returns the cosine of the argument. The input is interpreted as radians; the output is in the range [-1, 1].

The `cos_fr16` function inputs a fractional value in the range [-1.0, 1.0) corresponding to [-π/2, π/2]. The domain represents half a cycle which can be used to derive a full cycle if required (see Notes below). The result, in radians, is in the range [-1.0, 1.0).

Error Conditions

The `cos` function does not return an error condition.

Example

```
#include <math.h>
double y;

y = cos(3.14159);        /* y = -1.0 */
```

Notes

The domain of the `cos_fr16` function is restricted to the fractional range [`0x8000`, `0x7fff`] which corresponds to half a period from –(π /2) to π/2. It is possible however to derive the full period using the following properties of the function.

```
cosine [0, π/2] = -cosine [π, 3/2 π]
cosine [-π/2, 0] = -cosine [π/2, π]
```

The function below uses these properties to calculate the full period (from 0 to 2π) of the `cosine` function using an input domain of `[0, 0x7fff]`.

```
#include <math.h>

fract16 cos2pi_fr16 (fract16 x)
{
   if (x < 0x2000) {                    /* <0.25 */
      /* first quadrant [0..π/2):              */
      /* cos_fr16([0x0..0x7fff]) = [0..0x7fff)  */
      return cos_fr16(x * 4);

   } else if (x < 0x6000) {             /* < 0.75 */
      /* if (x < 0x4000)                       */
      /* second quadrant [π/2..π):              */
      /* -cos_fr16([0x8000..0x0)) = [0x7fff..0) */
      /*                                       */
      /* if (x < 0x6000)                       */
      /* third quadrant [π..3/2π):              */
      /* -cos_fr16([0x0..0x7fff]) = [0..0x8000) */
      return -cos_fr16((0xc000 + x) * 4);

   } else {
      /* fourth quadrant [3/2π..π):             */
      /* cos_fr16([0x8000..0x0)) = [0x8000..0)  */
      return cos_fr16((0x8000 + x) * 4);
   }
}
```

See Also

acos, sin

cosh

hyperbolic cosine

Synopsis

```
#include <math.h>
double cosh(double);
float coshf (float);
```

Description

The `cosh` function returns the hyperbolic cosine of its argument.

Error Conditions

The `cosh` function returns the IEEE constant `+Inf` if the argument is outside the domain.

Example

```
#include <math.h>
double x,y;

y = cosh(x);
```

See Also

sinh

disable_interrupts

disable interrupts

Synopsis

```
#include <sysreg.h>
void disable_interrupts(void)
```

Description

The `disable_interrupts` function causes the compiler to emit an instruction to disable hardware interrupts.

This function is implemented as a compiler built-in. The emitted instruction will be inlined at the point of its use. The inclusion of the `sysreg.h` include file is mandatory when using `disable_interrupts`.

The `disable_interrupts` function does not return a value.

Error Conditions

The `disable_interrupts` function does not return, raise, or set any error conditions.

Example

```
#include <sysreg.h>
main(){
   disable_interrupts();      // emits "DIS INTS;"
                              // instruction inline
}
```

See Also

enable_interrupts, io_space_read, io_space_write, mode_change, sysreg_read, sysreg_write

div

division

Synopsis

```
#include <stdlib.h>
div_t div(int numer, int denom);
```

Description

The `div` function divides `numer` by `denom`, both of type `int`, and returns a structure of type `div_t`. The type `div_t` is defined as

```
typedef struct {
   int quot;
   int rem;
} div_t
```

where `quot` is the quotient of the division and `rem` is the remainder, such that if `result` is of type `div_t`,

```
result.quot * denom + result.rem == numer
```

Error Conditions

If `denom` is zero, the behavior of the `div` function is undefined.

Example

```
#include <stdlib.h>
div_t result;

result = div(5, 2);       /* result.quot=2, result.rem=1 */
```

See Also

ldiv, fmod, modf

enable_interrupts

enable interrupts

Synopsis

```
#include <sysreg.h>
void enable_interrupts(void)
```

Description

The `enable_interrupts` function causes the compiler to emit an instruction to enable hardware interrupts.

The `enable_interrupts` function is implemented as a compiler built-in. The emitted instruction will be inlined at the point of `enable_interrupts` use. The inclusion of the `sysreg.h` include file is mandatory when using `enable_interrupts`.

The `enable_interrupts` function does not return a value.

Error Conditions

The `enable_interrupts` function does not return, raise or set any error conditions.

Example

```
#include <sysreg.h>
main(){
   enable_interrupts();       // emits "ENA INTS;"
                              // instruction inline
}
```

See Also

disable_interrupts, io_space_read, io_space_write, mode_change, sysreg_read, sysreg_write

exit

normal program termination

Synopsis

```
#include <stdlib.h>
void exit(int status);
```

Description

The `exit` function causes normal program termination. The functions registered by the `atexit` function are called in reverse order of their registration and the microprocessor is put into the `IDLE` state. The `status` argument is stored in register `AX1`, and control is passed to the label `___lib_prog_term`, which is defined in the run-time startup file.

Error Conditions

The `exit` function does not return an error condition.

Example

```
#include <stdlib.h>

exit(EXIT_SUCCESS);
```

See Also

abort, atexit

exp

exponential

Synopsis

```
#include <math.h>
double exp(double);
float expf (float);
```

Description

The `exp` function computes the exponential value `e` to the power of its argument. The argument must be in the range [-87.9 , 88.6].

Error Conditions

The `exp` function returns the value `HUGE_VAL` and stores the value `ERANGE` in `errno` when there is an overflow error. In the case of underflow, the `exp` function returns a zero.

Example

```
#include <math.h>
double y;
y = exp(1.0);       /* y = 2.71828...*/
```

See Also

alog, log, pow

external_memory_read

read from external memory

Synopsis

```
#include <sysreg.h>
void external_memory_read(int, void*)
```

Description

The `external_memory_read` function causes the compiler to emit instructions to read from external memory at the address passed as two parameters and set the value read from that address as a return value. The first parameter is the value of the top eight bits of the 24-bit external memory address to read. The second parameter is the lower 16 bits of the address of the external memory to read.

This function is implemented as a compiler built-in. The emitted instructions will be inlined at the point of `external_memory_read` use. The inclusion of the `sysreg.h` include file is mandatory when using `external_memory_read`.

Error Conditions

The `external_memory_read` function does not return, raise, or set any error conditions.

Example

```
#include <sysreg.h>
section("external_memory_section")
static int GlobalTable[256];

int main() {
   int page, read_value;
   asm("%0 = PAGE(GlobalTable);" : "=e"(page): : );
```

```
    read_value = external_memory_read(page, &GlobalTable[1]);
    return read_value;
}
```

See Also

external_memory_write, disable_interrupts, enable_interrupts, io_space_read, io_space_write, mode_change, sysreg_read, sysreg_write

external_memory_write

write to external memory

Synopsis

```
#include <sysreg.h>
void external_memory_write(int, void*, int)
```

Description

The `external_memory_write` function causes the compiler to emit instructions to write to external memory at the address passed in the first two parameters with the value passed in the last parameter. The first parameter is the value of the top eight bits of the 24-bit external memory address to write. The second parameter is the lower 16-bits of the address of the external memory to write.

This function is implemented as a compiler built-in. The emitted instructions will be inlined at the point of `external_memory_write` use. The inclusion of the `sysreg.h` include file is mandatory when using `external_memory_write`.

Error Conditions

The `external_memory_write` function does not return, raise, or set any error conditions.

Example

```
#include <sysreg.h>
section("external_memory_section")
static int GlobalTable[256];

int main() {
   int page, value_to_write = 0;
   asm("%0 = PAGE(GlobalTable);" : "=e"(page): : );
```

```
    external_memory_write(page, &GlobalTable[1], value_to_write);
}
```

See Also

external_memory_read, disable_interrupts, enable_interrupts, io_space_read, io_space_write, sysreg_read, sysreg_write

fabs

float absolute value

Synopsis

```
#include <math.h>
double fabs(double f);
float fabsf(float f);
```

Description

The `fabs` function returns the absolute value of the argument.

Error Conditions

The `fabs` function does not return an error condition.

Example

```
#include <math.h>
double y;

y = fabs(-2.3);      /* y = 2.3 */
y = fabs(2.3);       /* y = 2.3 */
```

See Also

abs, labs

floor

floor

Synopsis

```
#include <math.h>
double floor(double);
float floorf (float);
```

Description

The `floor` function returns the largest integral value that is not greater than its input.

Error Conditions

The `floor` function does not return an error condition.

Example

```
#include <math.h>
double y;

y = floor(1.25);        /* y = 1.0  */
y = floor(-1.25);       /* y = -2.0 */
```

See Also

ceil

fmod

floating-point modulus

Synopsis

```
#include <math.h>
double fmod(double numer, double denom);
float fmodf(float numer, float denom);
```

Description

The `fmod` function computes the floating-point remainder that results from dividing the first argument into the second argument. This value is less than the second argument and has the same sign as the first argument. If the second argument is equal to zero, `fmod` returns a zero.

Error Conditions

The `fmod` function does not return an error condition.

Example

```
#include <math.h>
double y;

y = fmod(5.0, 2.0);      /* y = 1.0 */
```

See Also

div, ldiv, modf

free

deallocate memory

Synopsis

```
#include <stdlib.h>
void free(void *ptr);
```

Description

The `free` function deallocates a pointer previously allocated to a range of memory (by `calloc` or `malloc`) to the free memory heap. If the pointer was not previously allocated by `calloc`, `malloc` or `realloc`, the behavior is undefined.

The `free` function returns the allocated memory to the heap from which it was allocated.

Error Conditions

The `free` function does not return an error condition.

Example

```
#include <stdlib.h>
char *ptr;

ptr = malloc(10);      /* Allocate 10 words from heap */
free(ptr);             /* Return space to free heap   */
```

See Also

calloc, malloc, realloc

frexp

separate fraction and exponent

Synopsis

```
#include <math.h>
double frexp(double f, int *expptr);
float frexpf(float f, int *expptr);
```

Description

The `frexp` function separates a floating-point input into a normalized fraction and a (base 2) exponent. The function returns the first argument as a fraction in the interval [½, 1), and stores a power of 2 in the integer pointed to by the second argument. If the input is zero, then the fraction and exponent will both be set to zero.

Error Conditions

The `frexp` function does not return an error condition.

Example

```
#include <math.h>
double y;
int exponent;

y = frexp(2.0, &exponent);      /* y=0.5, exponent=2 */
```

See Also

modf

interrupt

define interrupt handling

Synopsis

```
#include <signal.h>
void (*interrupt (int sig, void(*func)(int))) (int);
void (*interruptf(int sig, void(*func)(int))) (int);
void (*interrupts(int sig, void(*func)(int))) (int);
```

Description

These functions are Analog Devices extensions to the ANSI standard.

The `interrupt` function determines how a signal received during program execution is handled. The `interrupt` function executes the function pointed to by `func` at every interrupt `sig`; the `signal` function executes the function only once.

The different variants of the `interrupt` functions differentiate between handler dispatching functions. The variants will be appropriate for some applications and provide improved efficiency. The default `interrupt` function dispatcher saves and restores all scratch registers and modes on the data stack around a call to the handler (`func`) when servicing an interrupt. This dispatcher will pass the interrupt ID (for example, `SIG_PWRDWN`) to the handler as its parameter.

The `interruptf` interrupt dispatcher is similar to `interrupt`, except that it switches between primary and secondary register sets to save and restore registers instead of using the data stack. The `interruptf` function cannot be used in applications where nested interrupts are enabled. This interrupt dispatcher will pass the interrupt ID to the handler as its parameter.

The `interrupts` interrupt dispatcher saves and restores only the smallest number of registers and modes required to determine if a handler has been registered and to call that handler. The handler passed as input to `interrupts` must be declared using the `#pragma interrupt` directive (see on page 1-122). The `#pragma altregisters` directive (see on page 1-122) may be used in conjunction with the `interrupt pragma` in the definition of the handler. This dispatcher will *not* pass the interrupt ID to the handler.

The `sig` argument must be one of the signals listed in priority order in Table 3-15.

Table 3-15. Interrupt Function Signals - Values and Meanings

Sig Value	Definition
SIG_PWRDWN	power down interrupt
SIG_STACKINT	`PC`, `LOOP`, or `COUNTER` overflow on `push`, or on `pop` when empty
SIG_KERNEL	kernel interrupt
SIG_INT4	user-assignable
SIG_INT5	user-assignable
SIG_INT6	user-assignable
SIG_INT7	user-assignable
SIG_INT8	user-assignable
SIG_INT9	user-assignable
SIG_INT10	user-assignable
SIG_INT11	user-assignable
SIG_INT12	user-assignable
SIG_INT13	user-assignable
SIGABRT	software interrupt
SIGILL	software interrupt
SIGINT	software interrupt
SIGSEGV	software interrupt

Table 3-15. Interrupt Function Signals - Values and Meanings (Cont'd)

Sig Value	Definition
SIGTERM	software interrupt
SIGFPE	software interrupt

The `interrupt` functions cause the receipt of the signal number `sig` to be handled in one of the following ways:

- `SIG_DFL`—The default action is taken.
- `SIG_IGN`—The signal is ignored.
- Function Address—The function pointed to by `func` is executed.

The function pointed to by `func` is executed each time the interrupt is received. The `interrupt` function must be called with the `SIG_IGN` argument to disable interrupt handling.

Interrupts are *not* nested by the default start-up file.

Error Conditions

The `interrupt` functions return `SIG_ERR` and set `errno` equal to `SIG_ERR` if the requested interrupt is not recognized.

Example

```
include <signal.h>

void handler (int sig) {  /* Interrupt Service Routine (ISR) */
}

main () {
   /* enable power down interrupt and register ISR */
   interrupt(SIG_PWRDWN, handler);

   /* disable power down interrupt */
```

```
    interrupt(SIG_PWRDWN, SIG_IGN);

    /* enable power down interrupt and register ISR */
    interruptf(SIG_PWRDWN, handler);

    /* disable power down interrupt */
    interruptf(SIG_PWRDWN, SIG_IGN);
}
```

See Also

signal, raise

io_space_read

read I/O space

Synopsis

```
#include <sysreg.h>
int io_space_read(const int)
```

Description

The `io_space_read` function returns the value read from I/O memory space at the address specified by the parameter.

The function is implemented as a compiler built-in. If the input argument is a constant literal value the compiler will emit a Type 34 instruction that will be inlined at the point of `io_space_read` use. For non-literal inputs the compiler will call a library compiler support routine to perform the required read.

Error Conditions

The `io_space_read` function does not return, raise, or set any error conditions.

Example

```
#include <sysreg.h>
int addr = 0xA;

main(){
   int v1 = io_space_read(0xA);    /* inline instruction
                                      will be generated       */
   int v2 = io_space_read(addr);   /* library support routine
                                      will be called          */
}
```

See Also

disable_interrupts, enable_interrupts, io_space_write, mode_change, sysreg_read,sysreg_write

io_space_write

write I/O space

Synopsis

```
#include <sysreg.h>
void io_space_write(const int address, const unsigned int value)
```

Description

The `io_space_write` function stores the value passed as the second parameter to I/O memory space at the address passed as the first parameter.

This function is implemented as a compiler built-in. If the address parameter is a constant literal value the compiler will emit a Type 34 instruction that will be inlined at the point of `io_space_write` use. For non-literal addresses, the compiler will call a library compiler support routine to perform the required write.

The inclusion of the `sysreg.h` include file is mandatory when using `io_space_write`.

Error Conditions

The `io_space_write` function does not return, raise or set any error conditions.

Example

```
#include <sysreg.h>
int addr = 0xA;
int val = 0xA;

main(){
  int v1 = io_space_write(0xA, val);  /* inline instruction
                                          will be generated  */
```

```
    int v2 = io_space_write(addr, 0xFF);  /* support routine
                                              will be called    */
  }
```

See Also

disable_interrupts, enable_interrupts, io_space_read, mode_change, sysreg_read, sysreg_write

isalnum

detect alphanumeric character

Synopsis

```
#include <ctype.h>
int isalnum(int c);
```

Description

The `isalnum` function determines if the argument is an alphanumeric character (`A-Z`, `a-z`, or `0-9`). If the argument is not alphanumeric, `isalnum` returns a zero. If the argument is alphanumeric, `isalnum` returns a non-zero value.

Error Conditions

The `isalnum` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%3s", isalnum(ch) ? "alphanumeric" : "");
    putchar('\n');
}
```

See Also

isalpha, isdigit

isalpha

detect alphabetic character

Synopsis

```
#include <ctype.h>
int isalpha(int c);
```

Description

The `isalpha` function determines if the input is an alphabetic character (`A-Z` or `a-z`). If the input is not alphabetic, `isalpha` returns a zero. If the input is alphabetic, `isalpha` returns a non-zero value.

Error Conditions

The `isalpha` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", isalpha(ch) ? "alphabetic" : "");
    putchar('\n');
}
```

See Also

isdigit, isalnum

iscntrl

detect control character

Synopsis

```
#include <ctype.h>
int iscntrl(int c);
```

Description

The `iscntrl` function determines if the argument is a control character (`0x00-0x1F` or `0x7F`). If the argument is not a control character, `iscntrl` returns a zero. If the argument is a control character, `iscntrl` returns a non-zero value.

Error Conditions

The `iscntrl` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", iscntrl(ch) ? "control" : "");
    putchar('\n');
}
```

See Also

isalnum, isgraph

isdigit

detect decimal digit

Synopsis

```
#include <ctype.h>
int isdigit(int c);
```

Description

The `isdigit` function determines if the input character is a decimal digit (`0-9`). If the input is not a digit, `isdigit` returns a zero. If the input is a digit, `isdigit` returns a non-zero value.

Error Conditions

The `isdigit` function does not return an error condition.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", isdigit(ch) ? "digit" : "");
    putchar('\n');
}
```

See Also

isalpha, isalnum, isxdigit

isgraph

detect printable character, not including white space

Synopsis

```
#include <ctype.h>
int isgraph(int c);
```

Description

The `isgraph` function determines if the argument is a printable character, not including a white space (`0x21-0x7e`). If the argument is not a printable character, `isgraph` returns a zero. If the argument is a printable character, `isgraph` returns a non-zero value.

Error Conditions

The `isgraph` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", isgraph(ch) ? "graph" : "");
    putchar('\n');
}
```

See Also

isalnum, iscntrl, isprint

isinf

test for infinity

Synopsis

```
#include <math.h>
int isinff(float x);
int isinf(double x);
```

Description

The `isinf` function returns a zero if the argument is not set to the IEEE constant for `+Infnity` or `-Infinity`; otherwise, the function will return a non-zero value.

Error Conditions

The `isinf` function does not return or set any error conditions.

Example

```
#include <stdio.h>
#include <math.h>

static int fail=0;

main(){
   /* test int isinf(double) */
   union {
            double d; float f; unsigned long l;
   } u;

   #ifdef __DOUBLES_ARE_FLOATS__
      u.l=0xFF800000L; if ( isinf(u.d)==0 ) fail++;
      u.l=0xFF800001L; if ( isinf(u.d)!=0 ) fail++;
```

```
      u.l=0x7F800000L; if ( isinf(u.d)==0 ) fail++;
      u.l=0x7F800001L; if ( isinf(u.d)!=0 ) fail++;
   #endif

   /* test int isinff(float) */
      u.l=0xFF800000L; if ( isinff(u.f)==0 ) fail++;
      u.l=0xFF800001L; if ( isinff(u.f)!=0 ) fail++;
      u.l=0x7F800000L; if ( isinff(u.f)==0 ) fail++;
      u.l=0x7F800001L; if ( isinff(u.f)!=0 ) fail++;

   /* print pass/fail message */
   if ( fail==0 )
      printf("Test passed\n");
   else
      printf("Test failed: %d\n", fail);
}
```

See Also

isnan

islower

detect lowercase character

Synopsis

```
#include <ctype.h>
int islower(int c);
```

Description

The `islower` function determines if the argument is a lowercase character (`a-z`). If the argument is not lowercase, `islower` returns a zero. If the argument is lowercase, `islower` returns a non-zero value.

Error Conditions

The `islower` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", islower(ch) ? "lowercase" : "");
    putchar('\n');
}
```

See Also

isalpha, isupper

isnan

test for not-a-number (NAN)

Synopsis

```
#include <math.h>
int isnanf(float x);
int isnan(double x);
```

Description

The `isnan` function returns a zero if the argument is not set to an IEEE `NaN` (Not a Number); otherwise, the function will return a non-zero value.

Error Conditions

The `isnan` function does not return or set any error conditions.

Example

```
#include <stdio.h>
#include <math.h>

static int fail=0;

main(){
   /* test int isnan(double) */
   union {
      double d; float f; unsigned long l;
   } u;

   #ifdef __DOUBLES_ARE_FLOATS__
      u.l=0xFF800000L; if ( isnan(u.d)!=0 ) fail++;
      u.l=0xFF800001L; if ( isnan(u.d)==0 ) fail++;
      u.l=0x7F800000L; if ( isnan(u.d)!=0 ) fail++;
```

```
        u.l=0x7F800001L; if ( isnan(u.d)==0 ) fail++;
    #endif

    /* test int isnanf(float) */
        u.l=0xFF800000L; if ( isnanf(u.f)!=0 ) fail++;
        u.l=0xFF800001L; if ( isnanf(u.f)==0 ) fail++;
        u.l=0x7F800000L; if ( isnanf(u.f)!=0 ) fail++;
        u.l=0x7F800001L; if ( isnanf(u.f)==0 ) fail++;

    /* print pass/fail message */
    if ( fail==0 )
        printf("Test passed\n");
    else
        printf("Test failed: %d\n", fail);
}
```

See Also

isinf

isprint

detect printable character

Synopsis

```
#include <ctype.h>
int isprint(int c);
```

Description

The `isprint` function determines if the argument is a printable character (`0x20-0x7E`). If the argument is not a printable character, `isprint` returns a zero. If the argument is a printable character, `isprint` returns a non-zero value.

Error Conditions

The `isprint` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%3s", isprint(ch) ? "printable" : "");
    putchar('\n');
}
```

See Also

isgraph, isspace

ispunct

detect punctuation character

Synopsis

```
#include <ctype.h>
int ispunct(int c);
```

Description

The `ispunct` function determines if the argument is a punctuation character. If the argument is not a punctuation character, `ispunct` returns a zero. If the argument is a punctuation character, `ispunct` returns a non-zero value.

Error Conditions

The `ispunct` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%3s", ispunct(ch) ? "punctuation" : "");
    putchar('\n');
}
```

See Also

isalnum

isspace

detect whitespace character

Synopsis

```
#include <ctype.h>
int isspace(int c);
```

Description

The `isspace` function determines if the argument is a blank whitespace character (`0x09-0x0D` or `0x20`). This includes space, form feed (`\f`), new line (`\n`), carriage return (`\r`), horizontal tab (`\t`) and vertical tab (`\v`).

If the argument is not a blank space character, `isspace` returns a zero. If the argument is a blank space character, `isspace` returns a non-zero value.

Error Conditions

The `isspace` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", isspace(ch) ? "space" : "");
    putchar('\n');
}
```

See Also

iscntrl, isgraph

isupper

detect uppercase character

Synopsis

```
#include <ctype.h>
int isupper(int c);
```

Description

The `isupper` function determines if the argument is an uppercase character (`A-Z`). If the argument is not an uppercase character, `isupper` returns a zero. If the argument is an uppercase character, `isupper` returns a non-zero value.

Error Conditions

The `isupper` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", isupper(ch) ? "uppercase" : "");
    putchar('\n');
}
```

See Also

isalpha, islower

isxdigit

detect hexadecimal digit

Synopsis

```
#include <ctype.h>
int isxdigit(int c);
```

Description

The `isxdigit` function determines if the argument is a hexadecimal digit character (`A-F`, `a-f`, or `0-9`). If the argument is not a hexadecimal digit, `isxdigit` returns a zero. If the argument is a hexadecimal digit, `isxdigit` returns a non-zero value.

Error Conditions

The `isxdigit` function does not return any error conditions.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
    printf("%#04x", ch);
    printf("%2s", isxdigit(ch) ? "hexadecimal" : "");
    putchar('\n');
}
```

See Also

isalnum, isdigit

labs

long integer absolute value

Synopsis

```
#include <stdlib.h>
long int labs(long int j);
```

Description

The `labs` function returns the absolute value of its integer input.

Error Conditions

The `labs` function does not return an error condition.

Example

```
#include <stdlib.h>
long int j;
j = labs(-285128);    /* j = 285128 */
```

See Also

abs, fabs

ldexp

multiply by power of 2

Synopsis

```
#include <math.h>
double ldexp(double x, int n);
float ldexpf(float x, int n);
```

Description

The `ldexp` function returns the value of the floating-point input multiplied by 2 raised to the power of `n`. It adds the value of the second argument `n` to the exponent of the first argument `x`.

Error Conditions

If the result overflows, `ldexp` returns a NaN. If the result underflows, `ldexp` returns a zero.

Example

```
#include <math.h>
double y;

y = ldexp(0.5, 2);     /* y = 2.0 */
```

See Also

exp, pow

ldiv

long division

Synopsis

```
#include <stdlib.h>
ldiv_t ldiv(long int numer, long int denom);
```

Description

The `ldiv` function divides `numer` by `denom`, and returns a structure of type `ldiv_t`. The type `ldiv_t` is defined as:

```
typedef struct {
   long int quot;
   long int rem;
} ldiv_t
```

where `quot` is the quotient of the division and `rem` is the remainder, such that if result is of type `ldiv_t`, then

```
result.quot * denom + result.rem = numer
```

Error Conditions

If `denom` is zero, the behavior of the `ldiv` function is undefined.

Example

```
#include <stdlib.h>
ldiv_t result;

result = ldiv(7, 2);     /* result.quot=3, result.rem=1 */
```

See Also

div, fmod

log

natural logarithm

Synopsis

```
#include <math.h>
double log(double);
float logf(float);
```

Description

The `log` function computes the natural (base e) logarithm of its input.

Error Conditions

The `log` function returns a zero and sets `errno` to `EDOM` if the input value is negative.

Example

```
#include <math.h>
double y;

y = log(1.0);      /* y = 0.0 */
```

See Also

alog, exp, log10

log10

base 10 logarithm

Synopsis

```
#include <math.h>
double log10(double);
float log10f(float);
```

Description

The `log10` function returns the base 10 logarithm of its input.

Error Conditions

The `log10` function indicates a domain error (sets `errno` to `EDOM`) and returns a zero if the input is negative.

Example

```
#include <math.h>
double y;

y = log10(100.0);       /* y = 2.0 */
```

See Also

alog10, log, pow

longjmp

second return from `setjmp`

Synopsis

```
#include <setjmp.h>
void longjmp(jmp_buf env, int return_val);
```

Description

The `longjmp` function causes the program to execute a second return from the place where `setjmp (env)` was called (with the same `jmp_buf` argument).

The `longjmp` function takes as its arguments a jump buffer that contains the context at the time of the original call to `setjmp`. It also takes an integer, `return_val`, which `setjmp` returns if `return_val` is non-zero. Otherwise, `setjmp` returns a 1.

If `env` was not initialized through a previous call to `setjmp` or the function that called `setjmp` has since returned, the behavior is undefined. Also, automatic variables that are local to the original function calling `setjmp`, that do not have `volatile`-qualified type, and that have changed their value prior to the `longjmp` call, have indeterminate value.

Error Conditions

The `longjmp` function does not return an error condition.

Example

```
#include <setjmp.h>
#include <stdio.h>
#include <errno.h>
#include <stdlib.h>
```

```
jmp_buf env;
int res;

if ((res == setjmp(env)) != 0) {
   printf ("Problem %d reported by func ()", res);
   exit (EXIT_FAILURE);
}
func ();

void func (void)
{
   if (errno != 0) {
      longjmp (env, errno);
   }
}
```

See Also

setjmp

malloc

allocate memory

Synopsis

```
#include <stdlib.h>
void *malloc(size_t size);
```

Description

The `malloc` function returns a pointer to a block of memory of length `size`. The block of memory is uninitialized.

Error Conditions

The `malloc` function returns a null pointer if it is unable to allocate the requested memory.

Example

```
#include <stdlib.h>
int *ptr;

ptr = (int *)malloc(10);      /* ptr points to an   */
                              /* array of length 10 */
```

See Also

calloc, free, realloc

memchr

find first occurrence of character

Synopsis

```
#include <string.h>
void *memchr(const void *s1, int c, size_t n);
```

Description

The `memchr` function compares the range of memory pointed to by `s1` with the input character `c` and returns a pointer to the first occurrence of `c`. A null pointer is returned if `c` does not occur in the first `n` characters.

Error Conditions

The `memchr` function does not return an error condition.

Example

```
#include <string.h>
char *ptr;

ptr= memchr("TESTING", 'E', 7);
     /* ptr points to the E in TESTING  */
```

See Also

strchr, strrchr

memcmp

compare objects

Synopsis

```
#include <string.h>
int memcmp(const void *s1, const void *s2, size_t n);
```

Description

The `memcmp` function compares the first `n` characters of the objects pointed to by `s1` and `s2`. It returns a positive value if the `s1` object is lexically greater than the `s2` object, a negative value if the `s2` object is lexically greater than the `s1` object, and a zero if the objects are the same.

Error Conditions

The `memcmp` function does not return an error condition.

Example

```
#include <string.h>
char string1 = "ABC";
char string2 = "BCD";
int result;
result = memcmp (string1, string2, 3);     /* result < 0 */
```

See Also

strcmp, strcoll, strncmp

memcpy

copy characters from one object to another

Synopsis

```
#include <string.h>
void *memcpy(void *s1, const void *s2, size_t n);
```

Description

The memcpy function copies n characters from the object pointed to by s2 into the object pointed to by s1. The behavior of memcpy is undefined if the two objects overlap.

The memcpy function returns the address of s1.

Error Conditions

The memcpy function does not return an error condition.

Example

```
#include <string.h>
char *a = "SRC";
char *b = "DEST";
memcpy (b, a, 3);     /* *b="SRCT" */
```

See Also

memmove, memcpy_from_shared, memcpy_to_shared, strcpy, strncpy

memcpy_from_shared

copy characters from an object in the ADSP-2192-12 processor's shared memory (`0x20000-0x20FFF`) to non-shared memory in default data area.

Synopsis

```
#include <string.h>
void *memcpy_from_shared(void *s1, void *s2, size_t n);
```

Description

The `memcpy_from_shared` function copies `n` characters from the object pointed to by `s2` into the object pointed to by `s1`.

The `s2` parameter must be an object in the ADSP-2192-12 processor's shared memory (`0x20000-0x20FFF`).

Error Conditions

The `memcpy_from_shared` function does not return, raise, or set any error conditions.

See Also

memcpy, memcpy_to_shared

memcpy_to_shared

copy characters from object in default data area to an object in the ADSP-2192-12 processor's shared memory (0x20000-0x20FFF).

Synopsis

```
#include <string.h>
void *memcpy_to_shared(void *s1, void *s2, size_t n);
```

Description

The memcpy_to_shared function copies n characters from the object pointed to by s2 into the object pointed to by s1.

The s1 parameter must be an object in the ADSP-2192-12 processor's shared memory (0x20000-0x20FFF).

Error Conditions

The memcpy_to_shared function does not return, raise, or set any error conditions.

See Also

memcpy, memcpy_from_shared

memmove

copy characters between overlapping objects

Synopsis

```
#include <string.h>
void *memmove(void *s1, const void *s2, size_t n);
```

Description

The `memmove` function copies `n` characters from the object pointed to by `s2` into the object pointed to by `s1`. The entire object is copied correctly even if the objects overlap.

The `memmove` function returns a pointer to `s1`.

Error Conditions

The `memmove` function does not return an error condition.

Example

```
#include <string.h>
char *ptr, *str = "ABCDE";

ptr = str + 2;
memmove(str, str, 3);     /* *ptr = "ABC",  *str = "ABABC */
```

See Also

memcpy, strcpy, strncpy

memset

set range of memory to a character

Synopsis

```
#include <string.h>
void *memset(void *s1, int c, size_t n);
```

Description

The `memset` function sets a range of memory to the input character `c`. The first `n` characters of `s1` are set to `c`.

The `memset` function returns a pointer to `s1`.

Error Conditions

The `memset` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];
memset(string1, '\0', 50);      /* set string1 to 0 */
```

See Also

memcpy

mode_change

change selected system modes

Synopsis

```
#include <sysreg.h>
void mode_change(const int);
```

Description

The `mode_change` function causes the compiler to emit instructions to enable and disable a series of modes using the zero latency mode control instructions.

The `mode_change` function takes as a parameter a constant integer bitmask that the compiler converts into a series of enable and disable mode settings.

The `sysreg.h` include file defines each mode bit setting value as a symbolic variable that can be used as a parameter to `mode_change`. These definitions are:

```
__MODE_ENA_AV_LATCH =0x1,
__MODE_ENA_AR_SAT   =0x2,
__MODE_ENA_M_MODE   =0x4,
__MODE_ENA_TIMER    =0x8,
__MODE_ENA_INT      =0x10,
__MODE_DIS_AV_LATCH =0x100,
__MODE_DIS_AR_SAT   =0x200,
__MODE_DIS_M_MODE   =0x400,
__MODE_DIS_TIMER    =0x800,
__MODE_DIS_INT      =0x1000,
```

The `mode_change` function is implemented as a compiler built-in and the emitted instructions will be inlined at the point of `mode_change` use.

The inclusion of the `sysreg.h` include file is mandatory when using `mode_change`.

Error Conditions

The `mode_change` function does not return, raise, or set any error conditions.

Example

```
#include <sysreg.h>
main(){
   /* enable TIMER and disable AR saturation */
   mode_change( __MODE_ENA_TIMER |  __MODE_DIS_AR_SAT );
}
```

See Also

disable_interrupts, enable_interrupts, io_space_read, io_space_write, sysreg_read, sysreg_write

modf

separate integral and fractional parts

Synopsis

```
#include <math.h>
double modf(double f, double *fraction);
float modff (float f, float *fraction);
```

Description

The `modf` function separates the first argument into integral and fractional portions. The fractional portion is returned and the integral portion is stored in the object pointed to by the second argument. The integral and fractional portions have the same sign as the input.

Error Conditions

The `modf` function does not return an error condition.

Example

```
#include <math.h>
double y, n;

y = modf(-12.345, &n);      /* y = -0.345, n = -12.0 */
```

See Also

frexp

pow

raise to a power

Synopsis

```
#include <math.h>
double pow(double, double);
float powf(float, float);
```

Description

The `pow` function computes the value of the first argument raised to the power of the second argument.

Error Conditions

A domain error occurs if the first argument is negative and the second argument cannot be represented as an integer. If the first argument is zero, the second argument is less than or equal to zero, and the result cannot be represented, `EDOM` is stored in `errno`.

Example

```
#include <math.h>
double z;
z = pow(4.0, 2.0);      /* z = 16.0 */
```

See Also

exp, ldexp

qsort

quicksort

Synopsis

```
#include <stdlib.h>
void qsort(void *base, size_t nelem, size_t size,
     int (*compare) (const void *, const void *));
```

Description

The `qsort` function sorts an array of `nelem` objects, pointed to by `base`. The size of each object is specified by `size`.

The contents of the array are sorted into ascending order according to a comparison function pointed to by `compare`, which is called with two arguments that point to the objects being compared. The function shall return an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second.

If two elements compare as equal, their order in the sorted array is unspecified. The `qsort` function executes a binary search operation on a pre-sorted array. Note that:

- `base` points to the start of the array
- `nelem` is the number of elements in the array
- `size` is the size of each element of the array
- `compare` is a pointer to a function that is called by `qsort` to compare two elements of the array. The function should return a value less than, equal to, or greater than zero, according to whether the first argument is less than, equal to, or greater than the second.

Return Value

The `qsort` function returns no value.

Example

```
#include <stdlib.h>
float a[10];

int compare_float (const void *a, const void *b)
{
   float aval = *(float *)a;
   float bval = *(float *)b;
   if (aval < bval)
      return -1;
   else if (aval == bval)
      return 0;
   else
      return 1;
}

qsort (a, sizeof (a)/sizeof (a[0]), sizeof (a[0]), compare_float);
```

See Also

bsearch

raise

force a signal

Synopsis

```
#include <signal.h>
int raise(int sig);
```

Description

The `raise` function sends the signal `sig` to the executing program. The `raise` function forces interrupts wherever possible and simulates an interrupt otherwise. The `sig` argument must be one of the signals listed in priority order in Table 3-16

Table 3-16. Raise Function Signals - Values and Meanings

Sig Value	Definition
SIG_PWRDWN	power down interrupt
SIG_STACKINT	`PC`, `LOOP`, or `COUNTER` overflow on `push`, or on `pop` when empty
SIG_KERNEL	kernel interrupt
SIG_INT4	user-assignable
SIG_INT5	user-assignable
SIG_INT6	user-assignable
SIG_INT7	user-assignable
SIG_INT8	user-assignable
SIG_INT9	user-assignable
SIG_INT10	user-assignable
SIG_INT11	user-assignable
SIG_INT12	user-assignable
SIG_INT13	user-assignable

Table 3-16. Raise Function Signals - Values and Meanings (Cont'd)

Sig Value	Definition
SIGABRT	software interrupt
SIGILL	software interrupt
SIGINT	software interrupt
SIGSEGV	software interrupt
SIGTERM	software interrupt
SIGFPE	software interrupt

Interrupts are *not* nested by the default start-up file.

Error Conditions

The `raise` function returns a zero if successful, a non-zero value if it fails.

Example

```
#include <signal.h>
raise(SIGABRT);
```

See Also

interrupt, signal

rand

random number generator

Synopsis

```
#include <stdlib.h>
int rand(void);
```

Description

The `rand` function returns a pseudo-random integer value in the range $[0, 2^{15} - 1]$.

For this function, the measure of randomness is its periodicity, the number of values it is likely to generate before repeating a pattern. The output of the pseudo-random number generator has a period in the order of $2^{15} - 1$.

Error Conditions

The `rand` function does not return an error condition.

Example

```
#include <stdlib.h>
int i;

i = rand();
```

See Also

srand

realloc

change memory allocation

Synopsis

```
#include <stdlib.h>
void *realloc(void *ptr, size_t size);
```

Description

The `realloc` function changes the memory allocation of the object pointed to by `ptr` to `size`. Initial values for the new object are taken from those in the object pointed to by `ptr`. If the size of the new object is greater than the size of the object pointed to by `ptr`, then the values in the newly allocated section are undefined.

If `ptr` is a non-null pointer that was not allocated with `malloc` or `calloc`, the behavior is undefined. If `ptr` is a null pointer, `realloc` imitates `malloc`. If `size` is zero and `ptr` is not a null pointer, `realloc` imitates `free`.

Error Conditions

If memory cannot be allocated, `ptr` remains unchanged and `realloc` returns a null pointer.

Example

```
#include <stdlib.h>
int *ptr;

ptr = (int *)malloc(10);           /* intervening code    */
ptr = (int *)realloc(ptr, 20);     /* the size is now 20  */
```

See Also

calloc, free, malloc

setjmp

define a run-time label

Synopsis

```
#include <setjmp.h>
int setjmp(jmp_buf env);
```

Description

The `setjmp` function saves the calling environment in the `jmp_buf` argument. The effect of the call is to declare a run-time label that can be jumped to via a subsequent call to `longjmp`.

When `setjmp` is called, it immediately returns with a result of zero to indicate that the environment has been saved in the `jmp_buf` argument. If, at some later point, `longjmp` is called with the same `jmp_buf` argument, `longjmp` will restore the environment from the argument. The execution will then resume at the statement immediately following the corresponding call to `setjmp`. The effect is as if the call to `setjmp` has returned for a second time but this time the function will return a non-zero result.

The effect of calling `longjmp` will be undefined if the function that called `setjmp` has returned in the interim.

Error Conditions

The `setjmp` function does not return an error condition.

Example

See code example for "longjmp" on page 3-91.

See Also

longjmp

signal

define signal handling

Synopsis

```
#include <signal.h>
void (*signal(int sig, void (*func)(int))) (int);
void (*signalf(int sig, void (*func)(int))) (int);
void (*signals(int sig, void (*func)(int))) (int);
```

Description

These functions are Analog Devices extensions to the ANSI standard.

The `signal` function determines how a signal received during program execution is handled. The `signal` functions cause a single execution the function pointed to by `func`; the `interrupt` functions cause the function to be executed for every interrupt.

The different variants of the `signal` functions differentiate between handler dispatching functions. The variants will be appropriate for some applications and provide improved efficiency. The default `signal` function dispatcher saves and restores all scratch registers and modes on the data stack around a call to the handler (`func`) when servicing an interrupt. This dispatcher will pass the interrupt ID (for example, `SIG_PWRDWN`) to the handler as its parameter.

The `signalf` interrupt dispatcher is similar to `interrupt`, except that it switches between primary and secondary register sets to save and restore registers instead of using the data stack. The `signalf` function cannot be used in applications where nested interrupts are enabled. This interrupt dispatcher will pass the interrupt ID to the handler as its parameter.

The `signals` interrupt dispatcher saves and restores only the smallest number of registers and modes required to determine if a handler has been registered and to call that handler. The handler passed as input to `signals` must be declared using the `#pragma interrupt` directive (see on page 1-122). The `altregisters` directive (see on page 1-122) may be used in conjunction with the `interrupt` pragma in the definition of the handler. This dispatcher will *not* pass the interrupt ID to the handler.

The `sig` argument must be one of the signals listed in highest to lowest priority of interrupts in Table 3-17.

Table 3-17. Signal Function Signals - Values and Meanings

Sig Value	Definition
SIG_PWRDWN	power down interrupt
SIG_STACKINT	`PC`, `LOOP`, or `COUNTER` overflow on `push`, or on `pop` when empty
SIG_KERNEL	kernel interrupt
SIG_INT4	user-assignable
SIG_INT5	user-assignable
SIG_INT6	user-assignable
SIG_INT7	user-assignable
SIG_INT8	user-assignable
SIG_INT9	user-assignable
SIG_INT10	user-assignable
SIG_INT11	user-assignable
SIG_INT12	user-assignable
SIG_INT13	user-assignable
SIGABRT	software interrupt
SIGILL	software interrupt
SIGINT	software interrupt
SIGSEGV	software interrupt
SIGTERM	software interrupt
SIGFPE	software interrupt

The `signal` function causes the receipt of the signal number `sig` to be handled in one of the following ways:

- `SIG_DFL`—The default action is taken.
- `SIG_IGN`—The signal is ignored.
- Function address—The function pointed to by `func` is executed. The function pointed to by `func` is executed once when the signal is received. Handling is then returned to the default state.

Interrupts are *not* nested by the default start-up file.

Error Conditions

The `signal` function returns `SIG_ERR` and sets `errno` to `SIG_ERR` if it does not recognize the requested signal.

Example

```
#include <signal.h>

void handler (int sig) {   /* Interrupt Service Routine (ISR) */
}

main () {

    /* enable power down interrupt and register ISR */
    signal(SIG_PWRDWN, handler);

    /* disable power down interrupt */
    signal(SIG_PWRDWN, SIG_IGN);

    /* enable power down interrupt and register ISR */
    signalf(SIG_PWRDWN, handler);

    /* disable power down interrupt */
    signalf(SIG_PWRDWN, SIG_IGN);
}
```

See Also

interrupt, raise

sin

sine

Synopsis

```
#include <math.h>
double sin(double x);
float sinf (float x);
fract16 sin_fr16 (fract16 x);
```

Description

The `sin` function returns the sine of the argument `x`. The input is interpreted as a radian; the output is in the range [-1, 1].

The `sin_fr16` function inputs a fractional value in the range [-1.0, 1.0) corresponding to [-π/2, π/2]. The domain represents half a cycle which can be used to derive a full cycle if required (see Notes below). The result, in radians, is in the range [-1.0, 1.0).

Error Conditions

The `sin` function does not return an error condition.

Example

```
#include <math.h>
double y;
y = sin(3.14159);       /* y = 0.0 */
```

Notes

The domain of the `sin_fr16` function is restricted to the fractional range [`0x8000`, `0x7fff`] which corresponds to half a period from –(π /2) to π/2. It is possible however to derive the full period using the following properties of the function.

```
sine [0, π/2] = -sine [π, 3/2 π]

sine [-π/2, 0] = -sine [π/2, π]
```

The function below uses these properties to calculate the full period (from 0 to 2π) of the `sine` function using an input domain of `[0, 0x7fff]`.

```
#include <math.h>

fract16 sin2pi_fr16 (fract16 x)
{
   if (x < 0x2000) {                  /* <0.25 */
      /* first quadrant [0..π/2):            */
      /* sin_fr16([0x0..0x7fff]) = [0..0x7fff)  */
      return sin_fr16(x * 4);

   } else if (x < 0x6000) {           /* < 0.75 */
      /* if (x < 0x4000)                     */
      /* second quadrant [π/2..π):           */
      /* -sin_fr16([0x8000..0x0)) = [0x7fff..0) */
      /*                                     */
      /* if (x < 0x6000)                     */
      /* third quadrant [π..3/2π):           */
      /* -sin_fr16([0x0..0x7fff]) = [0..0x8000) */
      return -sin_fr16((0xc000 + x) * 4);

   } else {
      /* fourth quadrant [3/2π..π):          */
      /* sin_fr16([0x8000..0x0)) = [0x8000..0)  */
      return sin_fr16((0x8000 + x) * 4);
   }
}
```

See Also

asin, cos

sinh

hyperbolic sine

Synopsis

```
#include <math.h>
double sinh(double x);
float sinhf (float x);
```

Description

The `sinh` function returns the hyperbolic sine of the argument `x`.

Error Conditions

The `sinh` function returns the IEEE constant `+Inf` if the argument is outside the domain

Example

```
#include <math.h>
double x,y;
y = sinh(x);
```

See Also

cosh

sqrt

square root

Synopsis

```
#include <math.h>
double sqrt(double x);
float sqrtf (float x);
fract16 sqrt_fr16 (fract16 x);
```

Description

The `sqrt` function returns the positive square root of the argument.

Error Conditions

The `sqrt` function returns a zero for a negative input.

Example

```
#include <math.h>
double y;
y = sqrt(2.0);      /* y = 1.414..... */
```

See Also

rsqrt

srand

random number seed

Synopsis

```
#include <stdlib.h>
void srand(unsigned int seed);
```

Description

The `srand` function is used to set the seed value for the `rand` function. A particular seed value always produces the same sequence of pseudo-random numbers.

Error Conditions

The `srand` function does not return an error condition.

Example

```
#include <stdlib.h>

srand(22);
```

See Also

rand

strcat

concatenate strings

Synopsis

```
#include <string.h>
char *strcat(char *s1, const char *s2);
```

Description

The `strcat` function appends a copy of the null-terminated string pointed to by `s2` to the end of the null-terminated string pointed to by `s1`. It returns a pointer to the new `s1` string, which is null-terminated. The behavior of `strcat` is undefined if the two strings overlap.

Error Conditions

The `strcat` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];

string1[0] = 'A';
string1[1] = 'B';
string1[2] = '\0';
strcat(string1, "CD");      /* new string is "ABCD" */
```

See Also

strncat

strchr

find first occurrence of character in string

Synopsis

```
#include <string.h>
char *strchr(const char *s1, int c);
```

Description

The `strchr` function returns a null pointer to the first location in `s1`, a null-terminated string, that contains the character `c`.

Error Conditions

The `strchr` function returns a null pointer if `c` is not part of the string.

Example

```
#include <string.h>
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = strchr(ptr1, 'E');
       /* ptr2 points to the E in TESTING */
```

See Also

memchr, strrchr

strcmp

compare strings

Synopsis

```
#include <string.h>
int strcmp(const char *s1, const char *s2);
```

Description

The `strcmp` function lexicographically compares the null-terminated strings pointed to by `s1` and `s2`. It returns a positive value if the `s1` string is greater than the `s2` string, a negative value if the `s2` string is greater than the `s1` string, and a zero if the strings are the same.

Error Conditions

The `strcmp` function does not return an error condition.

Example

```
#include <string.h>
char string1[50], string2[50];

if (strcmp(string1, string2))
    printf("%s is different than %s \n", string1, string2);
```

See Also

memcmp, strncmp

strcoll

compare strings

Synopsis

```
#include <string.h>
int strcoll(const char *s1, const char *s2);
```

Description

The `strcoll` function compares the string pointed to by `s1` with the string pointed to by `s2`. The comparison is based on the locale macro, `LC_COLLATE`. Because only the C locale is defined in the ADSP-219x DSP environment, the `strcoll` function is identical to the `strcmp` function. The function returns a positive value if the `s1` string is greater than the `s2` string, a negative value if the `s2` string is greater than the `s1` string, and a zero if the strings are the same.

Error Conditions

The `strcoll` function does not return an error condition.

Example

```
#include <string.h>
char string1[50], string2[50];

if (strcoll(string1, string2))
    printf("%s is different than %s \n", string1, string2);
```

See Also

strrchr, strncmp

strcpy

copy from one string to another

Synopsis

```
#include <string.h>
void *strcpy(char *s1, const char *s2);
```

Description

The `strcpy` function copies the null-terminated string pointed to by `s2` into the space pointed to by `s1`. Memory allocated for `s1` must be large enough to hold `s2`, plus one space for the null character (‘`\0`’). The behavior of `strcpy` is undefined if the two objects overlap or if `s1` is not large enough. The `strcpy` function returns the new `s1`.

Error Conditions

The `strcpy` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];

strcpy(string1, "SOMEFUN");
      /* SOMEFUN is copied into string1 */
```

See Also

memcpy, memmove, strncpy

strcspn

length of character segment in one string but not the other

Synopsis

```
#include <string.h>
size_t strcspn(const char *s1, const char *s2);
```

Description

The `strcspn` function returns the length of the initial segment of `s1` which consists entirely of characters not in the string pointed to by `s2`. The string pointed to by `s2` is treated as a set of characters. The order of the characters in the string is not significant.

Error Conditions

The `strcspn` function does not return an error condition.

Example

```
#include <string.h>
char *ptr1, *ptr2;
size_t len;

ptr1 = "Tried and Tested";
ptr2 = "aeiou";
len = strcspn (ptr1, ptr2);      /* len = 2 */
```

See Also

strlen, strspn

strerror

get string containing error message

Synopsis

```
#include <string.h>
char *strerror(int errnum);
```

Description

The `strerror` function returns a pointer to a string containing an error message by mapping the number in `errnum` to that string.

Error Conditions

The `strerror` function does not return an error condition.

Example

```
#include <string.h>
char *ptr1;

ptr1 = strerror(1);
```

See Also

No references to this function.

strlen

string length

Synopsis

```
#include <string.h>
size_t strlen(const char *s1);
```

Description

The `strlen` function returns the length of the null-terminated string pointed to by `s1` (not including the terminating null character).

Error Conditions

The `strlen` function does not return an error condition.

Example

```
#include <string.h>
size_t len;

len = strlen("SOMEFUN");      /* len = 7 */
```

See Also

No references to this function.

strncat

concatenate characters from one string to another

Synopsis

```
#include <string.h>
char *strncat(char *s1, const char *s2, size_t n);
```

Description

The `strncat` function appends a copy of up to `n` characters in the null-terminated string pointed to by `s2` to the end of the null-terminated string pointed to by `s1`. It returns a pointer to the new `s1` string.

The behavior of `strncat` is undefined if the two strings overlap. The new `s1` string is terminated with a null (‘`\0`’).

Error Conditions

The `strncat` function does not return an error condition.

Example

```
#include <string.h>
char string1[50], *ptr;

string1[0]='\0';
strncat(string1, "MOREFUN", 4);
       /* string1 equals "MORE" */
```

See Also

strcat

strncmp

compare characters in strings

Synopsis

```
#include <string.h>
int strncmp(const char *s1, const char *s2, size_t n);
```

Description

The `strncmp` function lexicographically compares up to `n` characters of the null-terminated strings pointed to by `s1` and `s2`. It returns a positive value if the `s1` string is greater than the `s2` string, a negative value if the `s2` string is greater than the `s1` string, and a zero if the strings are the same.

Error Conditions

The `strncmp` function does not return an error condition.

Example

```
#include <string.h>
char *ptr1;
ptr1 = "TEST1";

if (strncmp(ptr1, "TEST", 4) == 0)
    printf("%s starts with TEST \n", ptr1);
```

See Also

memcmp, strcmp

strncpy

copy characters from one string to another

Synopsis

```
#include <string.h>
char *strncpy(char *s1, const char *s2, size_t n);
```

Description

The `strncpy` function copies up to `n` characters of the null-terminated string pointed to by `s2` into the space pointed to by `s1`. If the last character copied from `s2` is not a null, the result does not end with a null. The behavior of `strncpy` is undefined when the two objects overlap. The `strncpy` function returns the new `s1`.

If the `s2` string contains fewer than `n` characters, the `s1` string is padded with the null character until all `n` characters have been written.

Error Conditions

The `strncpy` function does not return an error condition.

Example

```
#include <string.h>
char string1[50];

strncpy(string1, "MOREFUN", 4);
                               /* MORE is copied into string1 */
string1[4] = '\0';             /* must null-terminate string1 */
```

See Also

memcpy, memmove, strcpy

strpbrk

find character match in two strings

Synopsis

```
#include <string.h>
char *strpbrk(const char *s1, const char *s2);
```

Description

The `strpbrk` function returns a pointer to the first character in `s1` that is also found in `s2`. The string pointed to by `s2` is treated as a set of characters. The order of the characters in the string is not significant.

Error Conditions

In the event that no character in `s1` matches any in `s2`, a null pointer is returned.

Example

```
#include <string.h>
char *ptr1, *ptr2, *ptr3;

ptr1 = "TESTING";
ptr2 = "SHOP"
ptr3 = strpbrk(ptr1, ptr2);
       /* ptr3 points to the S in TESTING */
```

See Also

strspn

strrchr

find last occurrence of character in string

Synopsis

```
#include <string.h>
char *strrchr(const char *s1, int c);
```

Description

The `strrchr` function returns a pointer to the last occurrence of character `c` in the null-terminated input string `s1`.

Error Conditions

The `strrchr` function returns a null pointer if `c` is not found.

Example

```
#include <string.h>
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = strrchr(ptr1, 'T');
      /* ptr2 points to the second T of TESTING  */
```

See Also

memchr, strchr

strspn

length of segment of characters in both strings

Synopsis

```
#include <string.h>
size_t strspn(const char *s1, const char *s2);
```

Description

The `strspn` function returns the length of the initial segment of `s1` which consists entirely of characters in the string pointed to by `s2`. The string pointed to by `s2` is treated as a set of characters. The order of the characters in the string is not significant.

Error Conditions

The `strspn` function does not return an error condition.

Example

```
#include <string.h>
size_t len;
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = "ERST";
len = strspn(ptr1, ptr2);      /* len = 4 */
```

See Also

strcspn, strlen

strstr

find string within string

Synopsis

```
#include <string.h>
char *strstr(const char *s1, const char *s2);
```

Description

The `strstr` function returns a pointer to the first occurrence in the string pointed to by s1 of the characters in the string pointed to by s2. This excludes the terminating null character in `s1`.

Error Conditions

If the string is not found, `strstr` returns a null pointer. If `s2` points to a string of zero length, `s1` is returned.

Example

```
#include <string.h>
char *ptr1, *ptr2;

ptr1 = "TESTING";
ptr2 = strstr(ptr1, ‘E’);
      /* ptr2 points to the E in TESTING  */
```

See Also

strchr

strtod

convert string to double

Synopsis

```
#include <stdlib.h>
double strtod(const char *nptr, char **endptr)
```

Description

The `strtod` function extracts a value from the string pointed to by `nptr`, and returns the value as a `double`. The `strtod` function expects `nptr` to point to a string that represents either a decimal floating-point number or a hexadecimal floating-point number. Either form of number may be preceded by a sequence of whitespace characters (as determined by the `isspace` function) that the function ignores.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and digits are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (–) followed by the hexadecimal prefix `0x` or `0X`. This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter p or P, an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens [hexdigs] [.hexdigs].

The first character that does not fit either form of number will stop the scan. If endptr is not NULL, a pointer to the character that stopped the scan is stored at the location pointed to by endptr. If no conversion can be performed, the value of nptr is stored at the location pointed to by endptr.

Error Conditions

The strtod function returns a zero if no conversion can be made and a pointer to the invalid string is stored in the object pointed to by endptr. If the correct value results in an overflow, a positive or negative (as appropriate) HUGE_VAL is returned. If the correct value results in an underflow, 0 is returned. The ERANGE value is stored in errno in the case of either an overflow or underflow.

Example

```
#include <stdlib.h>
char *rem;
double dd;

dd = strtod ("2345.5E4 abc",&rem);
     /* dd = 2.3455E+7, rem = "abc" */

dd = strtod ("-0x1.800p+9,123",&rem);
     /* dd = -768.0, rem = ",123"   */
```

See Also

atof, strtol, strtoul

strtodf

convert string to float

Synopsis

```
#include <stdlib.h>
float strtodf(const char *nptr, char **endptr)
```

Description

The `strtodf` function extracts a value from the string pointed to by `nptr`, and returns the value as a `float`. The `strtodf` function expects `nptr` to point to a string that represents either a decimal floating-point number or a hexadecimal floating-point number. Either form of number may be preceded by a sequence of whitespace characters (as determined by the `isspace` function) that the function ignores.

A decimal floating-point number has the form:

```
[sign] [digits] [.digits] [{e|E} [sign] [digits]]
```

The `sign` token is optional and is either plus (+) or minus (–); and digits are one or more decimal digits. The sequence of digits may contain a decimal point (.).

The decimal digits can be followed by an exponent, which consists of an introductory letter (`e` or `E`) and an optionally signed integer. If neither an exponent part nor a decimal point appears, a decimal point is assumed to follow the last digit in the string.

The form of a hexadecimal floating-point number is:

```
[sign] [{0x}|{0X}] [hexdigs] [.hexdigs] [{p|P} [sign] [digits]]
```

A hexadecimal floating-point number may start with an optional plus (+) or minus (–) followed by the hexadecimal prefix `0x` or `0X`. This character sequence must be followed by one or more hexadecimal characters that optionally contain a decimal point (.).

The hexadecimal digits are followed by a binary exponent that consists of the letter p or P, an optional sign, and a non-empty sequence of decimal digits. The exponent is interpreted as a power of two that is used to scale the fraction represented by the tokens [hexdigs] [.hexdigs].

The first character that does not fit either form of number will stop the scan. If endptr is not NULL, a pointer to the character that stopped the scan is stored at the location pointed to by endptr. If no conversion can be performed, the value of nptr is stored at the location pointed to by endptr.

Error Conditions

The strtodf function returns a zero if no conversion can be made and a pointer to the invalid string is stored in the object pointed to by endptr. If the correct value results in an overflow, a positive or negative (as appropriate) HUGE_VAL is returned. If the correct value results in an underflow, 0 is returned. The ERANGE value is stored in errno in the case of either an overflow or underflow.

Example

```
#include <stdlib.h>
char *rem;
float f;

f = strtodf ("2345.5E4 abc",&rem);
       /* f = 2.3455E+7, rem = "abc" */
f = strtodf ("-0x1.800p+9,123",&rem);
       /* f = -768.0, rem = ",123    */
```

See Also

atof, strtol, strtoul

strtok

convert string to tokens

Synopsis

```
#include <string.h>
char *strtok(char *s1, const char *s2);
```

Description

The `strtok` function returns successive tokens from the string `s1`, where each token is delimited by characters from `s2`.

A call to `strtok`, with `s1` not NULL, returns a pointer to the first token in `s1`, where a token is a consecutive sequence of characters not in `s2`. `s1` is modified in place to insert a null character at the end of the token returned. If `s1` consists entirely of characters from `s2`, NULL is returned.

Subsequent calls to `strtok` with `s1` equal to NULL will return successive tokens from the same string. When the string contains no further tokens, NULL is returned. Each new call to `strtok` may use a new delimiter string, even if `s1` is NULL, in which case the remainder of the string is tokenized using the new delimiter characters.

Error Conditions

The `strtok` function returns a null pointer if there are no tokens remaining in the string.

Example

```
#include <string.h>
static char str[] = "a phrase to be tested, today";
char *t;

t = strtok(str, " ");       /* t points to "a"            */
```

```
t = strtok(NULL, " ");      /* t points to "phrase"       */
t = strtok(NULL, ",");      /* t points to "to be tested" */
t = strtok(NULL, ".");      /* t points to " today"       */
t = strtok(NULL, ".");      /* t = NULL                   */
```

See Also

No references to this function.

strtol

convert string to long integer

Synopsis

```
#include <stdlib.h>
long int strtol(const char *nptr, char **endptr, int base);
```

Description

The `strtol` function returns as a `long int` the value that was represented by the string `nptr`. If `endptr` is not a null pointer, `strtol` stores a pointer to the unconverted remainder in `*endptr`.

The `strtol` function breaks down the input into three sections: white space (as determined by `isspace`), initial characters, and unrecognized characters, including a terminating null character. The initial characters may comprise an optional sign character, `0x` or `0X`, when `base` is 16, and those letters and digits which represent an integer with a radix of `base`. The letters (`a-z` or `A-Z`) are assigned the values 10 to 35 and are permitted only when those values are less than the value of `base`.

If `base` is zero, the base is taken from the initial characters. A leading `0x` indicates base 16; a leading 0 indicates base 8. For any other leading characters, base 10 is used. If `base` is between 2 and 36, it is used as a base for conversion.

Error Conditions

The `strtol` function returns a zero if no conversion can be made and the invalid string is stored in the object pointed to by `endptr`. If the correct value results in an overflow, positive or negative (as appropriate) `LONG_MAX` is returned. If the correct value results in an underflow, `LONG_MIN` is returned. `ERANGE` is stored in `errno` in the case of either overflow or underflow.

Example

```
#include <stdlib.h>
#define base 10
char *rem;
long int i;

i = strtol("2345.5", &rem, base);
    /* i=2345, rem=".5" */
```

See Also

atoi, atol, strtoul

strtoul

convert string to unsigned long integer

Synopsis

```
#include <stdlib.h>
unsigned long int strtoul(const char *nptr,
         char **endptr, int base);
```

Description

The `strtoul` function returns as an `unsigned long int` the value represented by the string `nptr`. If `endptr` is not a null pointer, `strtoul` stores a pointer to the unconverted remainder in `*endptr`.

The `strtoul` function breaks down the input into three sections:

- white space (as determined by `isspace`)
- initial characters
- unrecognized characters including a terminating null character

The initial characters may comprise an optional sign character, `0x` or `0X`, when `base` is 16, and those letters and digits which represent an integer with a radix of `base`. The letters (`a-z` or `A-Z`) are assigned the values 10 to 35, and are permitted only when those values are less than the value of `base`.

If `base` is zero, the base is taken from the initial characters. A leading `0x` indicates base 16; a leading 0 indicates base 8. For any other leading characters, base 10 is used. If `base` is between 2 and 36, it is used as a base for conversion.

Error Conditions

The `strtoul` function returns a zero if no conversion can be made and a pointer to the invalid string is stored in the object pointed to by `endptr`. If the correct value results in an overflow, `ULONG_MAX` is returned. `ERANGE` is stored in `errno` in the case of overflow.

Example

```
#include <stdlib.h>
#define base 10

char *rem;
unsigned long int i;

i = strtoul("2345.5", &rem, base);
    /* i = 2345, rem = ".5" */
```

See Also

atoi, atol, strtol

strxfrm

transform string using LC_COLLATE

Synopsis

```
#include <string.h>
size_t strxfrm(char *s1, const char *s2, size_t n);
```

Description

The strxfrm function transforms the string pointed to by s2 using the locale specific category LC_COLLATE, and copies no more than n characters of the transformed string into the string pointed to by s1. The resultant string will include the terminating null character.

The function returns the length of the transformed string (not including the terminating null character). If n is zero and s1 is set to the null pointer, then strxfrm will return the number of characters required for the transformed string. Overlapping strings are not supported

The transformation is such that strcmp will return the same result for two transformed strings as strcoll would for the same original strings. However, because only the C locale is defined in the ADSP-219x DSP environment, the strxfrm function is similar to the strncpy function except that the null character is always appended at the end of the output string.

Error Conditions

The strxfrm function does not return an error condition.

Example

```
#include <string.h>
char string1[50];
```

```
strxfrm(string1, "SOMEFUN", 49);
        /* SOMEFUN is copied into string1 */
```

See Also

strcmp, strcoll, strncpy

sysreg_read

read from non-memory-mapped register

Synopsis

```
#include <sysreg.h>
int sysreg_read(const int)
```

Description

The `sysreg_read` function causes the compiler to emit instructions to read the non-memory-mapped register passed as a parameter and set the value read from that register as a return value.

The input parameter for `sysreg_read` can be one of the symbolic variable constants of the enumerated data type `SysReg` defined in `sysreg.h` or any valid 8-bit value used to represent a system control register.

The symbolic definitions in `sysreg.h` are listed below.

General Register set:

`sysreg_ASTAT`	arithmetic status
`sysreg_SSTAT`	shifter status
`sysreg_MSTAT`	multiplier status
`sysreg_ICNTL`	interrupt control
`sysreg_IMASK`	interrupts enabled mask
`sysreg_IRPTL`	`Interrupt Latch` register
`sysreg_DMPG1`	`DMPG1` high address register
`sysreg_DMPG2`	`DMPG2` high address register
`sysreg_IOPG`	`IOPG` I/O page register

System Control Register set:

`sysreg_B0`	`B0` base register
`sysreg_B1`	`B1` base register
`sysreg_B2`	`B2` base register

`sysreg_B3`	`B3` base register
`sysreg_B4`	`B4` base register
`sysreg_B5`	`B5` base register
`sysreg_B6`	`B6` base register
`sysreg_B7`	`B7` base register
`sysreg_SYSCTL`	`SYSCTL` register

Cache Control Register:

`sysreg_CACTL`	`CACTL` register

Emulation Debug Control Register:

`sysreg_DBGCTRL`	`DBGCTRL` register

Emulation Debug Status Register:

`sysreg_DBGSTAT`	`DBGSTAT` register

Cycle Counter Registers

`sysreg_CNT0`	`CNT0` register
`sysreg_CNT1`	`CNT1` register
`sysreg_CNT2`	`CNT2` register
`sysreg_CNT3`	`CNT3` register

The `sysreg_read` function is implemented as a compiler built-in and the emitted instructions will be inlined at the point of `sysreg_read` use.

The inclusion of the `sysreg.h` include file is mandatory when using `sysreg_read`.

Error Conditions

The `sysreg_read` function does not return, raise, or set any error conditions. The compiler will not validate the input, instead it will rely on the assembler to fault erroneous values.

Example

```
#include <sysreg.h>

main(){
   int value = sysreg_read(sysreg_IMASK);
}
```

See Also

disable_interrupts, enable_interrupts, io_space_read, io_space_write, mode_change, sysreg_write

sysreg_write

write to non-memory-mapped register

Synopsis

```
#include <sysreg.h>
void sysreg_write(const int, const unsigned int)
```

Description

The `sysreg_write` function causes the compiler to emit instructions to write the non-memory-mapped register, which is passed as the first parameter, with the value passed as the second parameter.

The first parameter for `sysreg_write` can be one of the symbolic variable constants of the enumerated data type `SysReg` defined in `sysreg.h` or any valid 8-bit value used to represent a system control register. The second parameter is `unsigned`.

The symbolic definitions in `sysreg.h` are listed below.

General Register set:

`sysreg_ASTAT`	arithmetic status
`sysreg_SSTAT`	shifter status
`sysreg_MSTAT`	multiplier status
`sysreg_ICNTL`	interrupt control
`sysreg_IMASK`	interrupts enabled mask
`sysreg_IRPTL`	`Interrupt Latch` register
`sysreg_DMPG1`	`DMPG1` high address register
`sysreg_DMPG2`	`DMPG2` high address register
`sysreg_IOPG`	`IOPG` I/O page register

System Control Register set:

`sysreg_B0`	`B0` base register
`sysreg_B1`	`B1` base register

`sysreg_B2`	`B2` base register
`sysreg_B3`	`B3` base register
`sysreg_B4`	`B4` base register
`sysreg_B5`	`B5` base register
`sysreg_B6`	`B6` base register
`sysreg_B7`	`B7` base register
`sysreg_SYSCTL`	`SYSCTL` register

Cache Control Register:

`sysreg_CACTL`	`CACTL` register

Emulation Debug Control Register:

`sysreg_DBGCTRL`	`DBGCTRL` register

Emulation Debug Status Register:

`sysreg_DBGSTAT`	`DBGSTAT` register

Cycle Counter Registers

`sysreg_CNT0`	`CNT0` register
`sysreg_CNT1`	`CNT1` register
`sysreg_CNT2`	`CNT2` register
`sysreg_CNT3`	`CNT3` register

The `sysreg_write` function is implemented as a compiler built-in. The emitted instructions will be inlined at the point of `sysreg_write` use.

The inclusion of the `sysreg.h` include file is mandatory when using `sysreg_write`.

Error Conditions

The `sysreg_write` function does not return, raise, or set any error conditions.

Example

```
#include <sysreg.h>

main(){
   sysreg_write(sysreg_IMASK,0x1);
}
```

See Also

disable_interrupts, enable_interrupts, io_space_read,io_space_write, mode_change, sysreg_read

tan

tangent

Synopsis

```
#include <math.h>
double tan(double x);
float tanf(float x);
fract16 tan_fr16(fract16 x);
```

Description

The `tan` function returns the tangent of the argument `x`. The input, in radians, must be in the range [-9099, 9099].

The `tan_fr16` function is only defined for input values between -π/4 (=`0x9B78`) and π/4 (=`0x6488`). The input argument is in radians. Output values range from `0x8000` to `0x7FFF`. The library function returns 0 for any input argument that is outside the defined domain.

Error Conditions

The `tan` function returns zero if the input argument is outside the defined domain.

Example

```
#include <math.h>
double y;
y = tan(3.14159/4.0);      /* y = 1.0 */
```

See Also

atan, atan2

tanh

hyperbolic tangent

Synopsis

```
#include <math.h>
double tanh(double x);
float tanhf (float x)
```

Description

The `tanh` function returns the hyperbolic tangent of the argument `x`.

Error Conditions

The `tanh` function does not return an error condition.

Example

```
#include <math.h>
double x,y;
y = tanh(x);
```

See Also

cosh, sinh

tolower

convert from uppercase to lowercase

Synopsis

```
#include <ctype.h>
int tolower(int c);
```

Description

The `tolower` function converts the input character to lowercase if it is uppercase; otherwise, it returns the character.

Error Conditions

The `tolower` function does not return an error condition.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   if(isupper(ch))
      printf("tolower=%#04x", tolower(ch));
   putchar('\n');
}
```

See Also

islower, isupper, toupper

toupper

convert from lowercase to uppercase

Synopsis

```
#include <ctype.h>
int toupper(int c);
```

Description

The `toupper` function converts the input character to uppercase if it is in lowercase; otherwise, it returns the character.

Error Conditions

The `toupper` function does not return an error condition.

Example

```
#include <ctype.h>
int ch;

for (ch=0; ch<=0x7f; ch++) {
   printf("%#04x", ch);
   if(islower(ch))
      printf("toupper=%#04x", toupper(ch));
   putchar('\n');
}
```

See Also

islower, isupper, tolower

va_arg

get next argument in variable-length list of arguments

Synopsis

```
#include <stdarg.h>
void va_arg(va_list ap, type);
```

Description

The `va_arg` macro is used to walk through the variable length list of arguments to a function.

After starting to process a variable-length list of arguments with `va_start`, call `va_arg` with the same `va_list` variable to extract arguments from the list. Each call to `va_arg` returns a new argument from the list.

Substitute a `type` name corresponding to the type of the next argument for the `type` parameter in each call to `va_arg`. After processing the list, call `va_end`.

The header file `stdarg.h` defines a pointer type called `va_list` that is used to access the list of variable arguments.

The function calling `va_arg` is responsible for determining the number and types of arguments in the list. It needs this information to determine how many times to call `va_arg` and what to pass for the type parameter each time. There are several common ways for a function to determine this type of information. The standard C `printf` function reads its first argument looking for %-sequences to determine the number and types of its extra arguments. In the example below, all of the arguments are of the same type (`char*`), and a termination value (NULL) is used to indicate the end of the argument list. Other methods are also possible.

If a call to `va_arg` is made after all arguments have been processed, or if `va_arg` is called with a type parameter that is different from the type of the next argument in the list, the behavior of `va_arg` is undefined.

Error Conditions

The `va_arg` macro does not return an error condition.

Example

```
#include <stdarg.h>
#include <string.h>
#include <stdlib.h>

char *concat(char *s1,...)
{
   int len = 0;
   char *result;
   char *s;
   va_list ap;

   va_start (ap,s1);
   s = s1;
   while (s){
       len += strlen (s);
       s = va_arg (ap,char *);
   }
   va_end (ap);

   result = malloc (len +7);
   if (!result)
       return result;
   *result = '';
   va_start (ap,s1);
   s = s1;
   while (s){
       strcat (result,s);
       s = va_arg (ap,char *);
   }
   va_end (ap);
   return result;
}
```

See Also

va_end, va_start

va_end

finish variable-length argument list processing

Synopsis

```
#include <stdarg.h>
void va_end(va_list ap);
```

Description

The `va_end` macro can only be used after the `va_start` macro has been invoked. A call to `va_end` concludes the processing of a variable-length list of arguments that was begun by `va_start`.

Error Conditions

The `va_end` macro does not return an error condition.

Example

See "va_arg" on page 3-158.

See Also

va_arg, va_start

va_start

initialize the variable-length argument list processing

Synopsis

```
#include <stdarg.h>
void va_start(va_list ap, parmN);
```

Description

The `va_start` macro is used in a function declared to take a variable number of arguments to start processing those variable arguments. The first argument to `va_start` should be a variable of type `va_list`, which is used by `va_arg` to walk through the arguments.

The second argument is the name of the last *named* parameter in the function's parameter list; the list of variable arguments immediately follows this parameter. The `va_start` macro must be invoked before either the `va_arg` or `va_end` macro can be invoked.

Error Conditions

The `va_start` macro does not return an error condition.

Example

See "va_arg" on page 3-158.

See Also

va_arg, va_end

4 DSP RUN-TIME LIBRARY

This chapter describes the DSP run-time library which contains a broad collection of functions that are commonly required by signal processing applications. The services provided by the library include support for general purpose signal processing such as companders, filters, and Fast Fourier Transform (FFT) functions. All these services are Analog Devices extensions to ANSI standard C. These functions are in addition to the C/C++ run-time library functions that are described in Chapter 3, "C/C++ Run-Time Library".

For more information on the algorithms on which many of DSP run-time library's math functions are based, see the Cody and Waite text "*Software Manual for the Elementary Functions*" from Prentice Hall (1980).

This chapter contains:

- "DSP Run-Time Library Guide" on page 4-2
  It contains information about the library and provides a description of the DSP header files that are included with this release of the `cc219x` compiler.

- "DSP Run-Time Library Reference" on page 4-22
  It provides the complete reference for each DSP run-time library function provided with this release of the `cc219x` compiler.

DSP Run-Time Library Guide

The DSP run-time library contains functions that you can call from your source program. This section describes how to use the library and provides information on the following topics:

- "Calling DSP Library Functions"
- "Linking DSP Library Functions" on page 4-3
- "Working with Library Source Code" on page 4-3
- "DSP Header Files" on page 4-4

Calling DSP Library Functions

To use a DSP library function, call the function by name and give the appropriate arguments. The names and arguments for each function are described in the function's reference page in the section "DSP Run-Time Library Reference" on page 4-22.

Like other functions you use, library functions should be declared. Declarations are supplied in header files, as described in the section, "Working With Library Header Files" on page 3-8.

> The function names are C function names. If you call C run-time library functions from an assembly language program, you must use the assembly version of the function name, which is the function name prefixed with an underscore. For more information on naming conventions, see the section, "C/C++ and Assembly Language Interface" on page 1-169.

> You can use the archiver, described in the *VisualDSP++ 3.5 Linker and Utilities Manual for 16-Bit Processors*, to build library archive files of your own functions.

Linking DSP Library Functions

When your C/C++ code calls a DSP run-time library function, the call creates a reference that the linker resolves when linking your program. This requires the linker to be directed to link with the DSP run-time library, `libdsp.dlb`, in the `219x\lib` directory, which is a subdirectory of the VisualDSP++ installation directory. The default Linker Description File (LDF) will do this automatically as it specifies that `libdsp.dlb` will be included in every link. If an application uses a customized `.LDF` file, then either add `libdsp.dlb` to the customized `.LDF` file, or alternatively use the compiler's `-ldsp` switch to specify that `libdsp.dlb` is to be added to the link line.

Working with Library Source Code

The source code for the functions in the C and DSP run-time libraries is provided with your VisualDSP++ software. By default, the installation program copies the source code to a subdirectory of the directory where the run-time libraries are kept, named `219x\lib\src`. The directory contains the source for the C run-time library, the DSP run-time library, the ETSI functions, and for the I/O run-time library, as well as the source for the main program start-up functions. If you do not intend to modify any of the run-time library functions, you can delete this directory and its contents to conserve disk space.

The source code is provided so you can customize specific functions for your own needs. Each function is kept in a separate file. The file name is the name of the function with the appropriate extension for C or assembler source. To modify these files, you need proficiency in ADSP-219x DSP assembly language and an understanding of the run-time environment, as explained in "C/C++ and Assembly Language Interface" on page 1-169. Before you make any modifications to the source code, copy the source code to a file with a different filename and rename the function itself. Test the function before you use it in your system to verify that it is functionally correct.

Analog Devices only supports the run-time library functions as provided.

DSP Header Files

The DSP header files contains prototypes for all the DSP library functions. When the appropriate `#include` preprocessor command is included in your source, the compiler will use the prototypes to check that each function is called with the correct arguments. The DSP header files included in this release of the `cc219x` compiler are:

- "complex.h — Basic Complex Arithmetic Functions" on page 4-4
- "filter.h — DSP Filters and Transformations" on page 4-6
- "math.h — Math Functions" on page 4-10
- "matrix.h — Matrix Functions" on page 4-12
- "stats.h — Statistical Functions" on page 4-16
- "vector.h — Vector Functions" on page 4-17
- "window.h — Window Generators" on page 4-20

complex.h — Basic Complex Arithmetic Functions

The `complex.h` header file contains type definitions and basic arithmetic operations for variables of type `complex_float`, `complex_double`, and `complex_fract16`. The complex functions defined in the `complex.h` header file are listed in Table 4-1 on page 4-5.

The following structures are used to represent complex numbers in rectangular coordinates:

```
typedef struct
{
   float re;
   float im;
} complex_float;

typedef struct
{
   double re;
   double im;
} complex_double;

typedef struct
{
   fract16 re;
   fract16 im;
} complex_fract16
```

Details of the basic complex arithmetic functions are included in "DSP Run-Time Library Reference" on page 4-22.

Table 4-1. Complex Functions

Description	Prototype

Complex Absolute Value	`double cabs (complex_double a);`
	`float cabsf (complex_float a);`
	`fract16 cabs_fr16 (complex_fract16 a);`

Complex Addition	`complex_double cadd (complex_double a, complex_double b);`
	`complex_float caddf (complex_float a, complex_float b);`
	`complex_fract16 cadd_fr16 (complex_fract16 a,`
	`   complex_fract16 b);`

Complex Subtraction	`complex_double csub (complex_double a, complex_double b);`
	`complex_float csubf (complex_float a, complex_float b);`
	`complex_fract16 csub_fr16 (complex_fract16 a,`
	`   complex_fract16 b);`

Table 4-1. Complex Functions (Cont'd)

Description	Prototype

Complex Multiply	`complex_double cmlt (complex_double a, complex_double b);`
	`complex_float cmltf (complex_float a, complex_float b);`
	`complex_fract16 cmlt_fr16 (complex_fract16 a,`
	`   complex_fract16 b);`

Complex Division	`complex_double cdiv (complex_double a, complex_double b);`
	`complex_float cdivf (complex_float a, complex_float b);`
	`complex_fract16 cdiv_fr16 (complex_fract16 a,`
	`   complex_fract16 b);`

Get Phase of Complex Number	`double arg (complex_double a);`
	`float argf (complex_float a);`
	`fract16 arg_fr16 (complex_fr16 a);`

Complex Conjugate	`complex_double conj (complex_double a);`
	`complex_float conjf (complex_float a);`
	`complex_fract16 conj_fr16 (complex_fract16 a);`

Convert Cartesian to Polar Coordinates	`double cartesian (complex_double a, double* phase)`
	`float cartesianf (complex_float a, float* phase)`
	`fract16 cartesian_fr16 (complex_fract16 a, fract16* phase)`

Convert Polar to Cartesian Coordinates	`complex_double polar (double mag, double phase);`
	`complex_float polarf (float mag, float phase);`
	`complex_fract16 polar_fr16 (fract16 mag, fract16 phase);`

Complex Exponential	`complex_double cexp (double a);` `complex_float cexpf (float a);`
Normalization	`complex_double norm (complex_double a);` `complex_float normf (complex_float a);`

filter.h — DSP Filters and Transformations

The `filter.h` header file contains filters used in signal processing. It also includes the A-law and μ-law companders that are used by voice-band compression and expansion applications.

The header file also contains functions that perform key transformations used in DSPs, including FFTs and convolution.

Various forms of the FFT function are provided by the library corresponding to radix-2, radix-4, and two-dimensional FFTs. The number of points is provided as an argument. The twiddle table for the FFT functions is supplied as a separate argument and is normally calculated once during program initialization.

Library functions are provided to initialize a twiddle table. A table can accommodate several FFTs of different sizes by allocating the table at maximum size, and then using the stride argument of the FFT function to specify the step size through the table. If the stride argument is set to 1, the FFT function will use all the table; if the FFT uses only half the number of points of the largest, the stride should be 2.

The functions defined in the `filter.h` header file are listed in Table 4-2 and Table 4-3 and are described in detail in "DSP Run-Time Library Reference" on page 4-22.

Table 4-2. Filter Library

Description	Prototype
Finite Impulse Response Filter	`void fir_fr16 (const fract16 x[], fract16 y[], int n, fir_state_fr16 *s);`
Infinite Impulse Response Filter	`void iir_fr16 (const fract16 x[], fract16 y[], int n, iir_state_fr16 *s);`
FIR Decimation Filter	`void fir_decima_fr16 (const fract16 x[], fract16 y[], int n, fir_state_fr16 *s);`
FIR Interpolation Filter	`void fir_interp_fr16 (const fract16 x[], fract16 y[], int n, fir_state_fr16 *s);`
Complex Finite Impulse Response Filter	`void cfir_fr16 (const complex_fract16 x[], complex_fract16 y[], int n, cfir_state_fr16 *s);`

Table 4-3. Transformation Functions

Description	Prototype
Fast Fourier Transform	
Generate FFT Twiddle Factor	`void twidfft_fr16 (complex_fract16 w[], int n);`
Generate FFT Twiddle Factors for Radix 2 FFT	`void twidfftrad2_fr16 (complex_fract16 w[], int n);`
Generate FFT Twiddle Factors for Radix 4 FFT	`void twidfftrad4_fr16 (complex_fract16 w[], int n);`
Generate FFT Twiddle Factors for 2-D FFT	`void twidfft2d_fr16 (complex_fract16 w[], int n);`
N Point Radix 2 Complex Input FFT	`void cfft_fr16 (const complex_fract16 *in, complex fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int *block_exponent, int scale_method);`
N Point Radix 2 Real Input FFT	`void rfft_fr16 (const fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int *block_exponent, int scale_method);`
N Point Radix 2 Inverse Input FFT	`void ifft_fr16 (const complex_fract *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int *block_exponent, int scale_method);`
N Point Radix 4 Complex Input FFT	`void cfftrad4_fr16 (const complex_fract16 *in, complex fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int *block_exponent, int scale_method);`
N Point Radix 4 Real Input FFT	`void rfftrad4_fr16 (const fract16 *in, complex_fract16 *t, complex_fract16 *out, const complex_fract16 *w, int wst, int n, int *block_exponent, int scale_method);`

Table 4-3. Transformation Functions (Cont'd)

Description	Prototype

N Point Radix 4 Inverse Input FFT	`void ifftrad4_fr16`
	`(const complex_fract *in,`
	`complex_fract16 *t, complex_fract16 *out,`
	`const complex_fract16 *w, int wst, int n,`
	`int *block_exponent, int scale_method);`

Nxn Point 2-D Complex Input FFT	`void cfft2d_fr16`
	`(const complex_fract16 *in,`
	`complex fract16 *t, complex_fract16 *out,`
	`const complex_fract16 *w, int wst, int n,`
	`int *block_exponent, int scale_method);`

Nxn Point 2-D Real Input FFT	`void rfft2d_fr16`
	`(const fract16 *in, complex_fract16 *t,`
	`complex_fract16 *out,`
	`const complex_fract16 *w, int wst, int n,`
	`int *block_exponent, int scale_method);`

Nxn Point 2-D Inverse FFT	`void ifft2d_fr16`
	`(const complex_fract *in,`
	`complex_fract16 *t, complex_fract16 *out,`
	`const complex_fract16 *w, int wst, int n,`
	`int *block_exponent, int scale_method);`

Convolutions	

Convolution	`void convolve_fr16`
	`(const fract16 cin1[], int clen1,`
	`const fract16 cin2[], int clen2,`
	`fract16 cout[]);`

2-D Convolution	`void conv2d_fr16`
	`(const fract16 *cin1, int crow1, int ccol1,`
	`const fract16 *cin2, int crow2, int ccol2,`
	`fract16 *cout);`

2-D Convolution 3x3 Matrix	`void conv2d3x3_fr16`
	`(const fract16 *cin, int crow1, int ccol1,`
	`const fract16 cin2 [3] [3], fract16 *cout);`

Compression/Expansion	
A-law Compression	`void a_compress` `(const short in[], short out[], int n);`

Table 4-3. Transformation Functions (Cont'd)

Description	Prototype
A-law Expansion	`void a_expand (const short in[], short out[], int n);`
µ-law Compression	`void mu_compress (const short in[], short out[], int n);`
µ-law Expansion	`void mu_expand (const char in[], short out[], int n);`

math.h — Math Functions

The standard math functions have been augmented by implementations for the float data type, and in some cases for the fract16 data type.

Table 4-4 provides a summary of the functions defined by the `math.h` header file. Descriptions of these functions are given under the name of the `double` version in the "C Run-Time Library Reference" on page 3-26.

This header also provides prototypes for a number of additional math functions `clip`, `copysign`, `max`, `min`, and an integer function `countones`. These functions are described in "DSP Run-Time Library Reference" on page 4-22. .

Table 4-4. Math Library

Description	Prototype
Anti-log	`double alog (double x);` `float alogf (float x);`
Base 10 Anti-log	`double alog10 (double x);` `float alog10f (float x);`

Arc Cosine	`double acos (double x);`
	`float acosf (float x);`
	`fract16 acos_fr16 (fract16 x);`

Table 4-4. Math Library (Cont'd)

Description	Prototype

Arc Sine	`double asin (double x);`
	`float asinf (float x);`
	`fract16 asin_fr16 (fract16 x);`

Arc Tangent	`double atan (float x);`
	`float atanf (float x);`
	`fract16 atan_fr16 (fract16 x);`

Arc Tangent of Quotient	`double atan2 (double x, double y);`
	`float atan2f (float x, float y);`
	`fract16 atan2_fr16 (fract16 x, fract16 y);`

Ceiling	`double ceil (double x);` `float ceilf (float x);`

Cosine	`double cos (double x);`
	`float cosf (float x);`
	`fract16 cos_fr16 (fract16 x);`

Hyperbolic Cosine	`double cosh (double x);` `float coshf (float x);`
Cotangent	`double cot (double x);` `float cotf (float x);`
Exponential	`double exp (double x);` `float expf (float x);`
Floor	`double floor (double x);` `float floorf (float x);`
Floating Point Remainder	`double fmod (double x, double y);` `float fmodf (float x, float y);`
Get Mantissa and Exponent	`double frexp (double x, int *n);` `float frexpf (float x, int *n);`
Multiply by Power of 2	`double ldexp (double x, int n);` `float ldexpf (float x, int n);`
Natural Logarithm	`double log (double x);` `float logf (float x);`

Table 4-4. Math Library (Cont'd)

Description	Prototype
Logarithm Base 10	`double log10 (double x);` `float log10f (float x);`
Get Fraction and Integer	`double modf (double x, double *i);` `float modff (float x, float *i);`
Power	`double pow (double x, double y);` `float powf (float x, float y);`

Sine	`double sin (double x);`
	`float sinf (float x);`
	`fract16 sin_fr16 (fract16 x);`

Hyperbolic Sine	`double sinh (double x);` `float sinhf (float x);`

Square Root	`double sqrt (double x);`
	`float sqrtf (float x);`
	`fract16 sqrt_fr16 (fract16 x);`

Reciprocal of Square Root	`double rsqrt (double x);` `float rsqrtf (float x);`

Tangent	`double tan (double x);`
	`float tanf (float x);`
	`fract16 tan_fr16 (fract16 x);`

Hyperbolic Tangent	`double tanh (double x);` `float tanhf (float x);`

matrix.h — Matrix Functions

The `matrix.h` header file contains matrix functions for operating on real and complex matrices, both matrix-scalar and matrix-matrix operations. See "complex.h — Basic Complex Arithmetic Functions" on page 4-4 for definition of the complex types.

The matrix functions defined in the `matrix.h` header file are listed in Table 4-5. In most of the function prototypes:

`*a` is a pointer to input matrix `a [] []`

`*b` is a pointer to input matrix `b [] []`

`b` is an input scalar

`n` is the number of rows

`m` is the number of columns

`*c` is a pointer to output matrix `c [] []`

In the `matrix*matrix` functions, `n` and `k` are the dimensions of matrix `a` and `k` and `m` are the dimensions of matrix `b`.

The functions described by this header assume that input array arguments are constant; that is, their contents do not change during the course of the routine. In particular, this means the input arguments do not overlap with any output argument.

Table 4-5. Matrix Functions

Description	Prototype

Real
Matrix + Scalar
Addition
	`void matsadd (const double *a, const double b, int n, int m, double *c);`
	`void matsaddf (const float *a, const float b, int n, int m, float *c);`
	`void matsadd_fr16 (const fract16 *a, const fract16 b, int n, int m, fract16 *c);`

Real
Matrix – Scalar
Subtraction
	`void matssub (const double *a, const double b, int n, int m, double *c);`
	`void matssubf (const float *a, const float b, int n, int m, float *c);`
	`void matssub_fr16 (const fract16 *a, const fract16 b, int n, int m, fract16 *c);`

Table 4-5. Matrix Functions (Cont'd)

Description	Prototype

Real Matrix * Scalar Multiplication	`void matsmlt (const double *a, double b, int n, int m, double *c);`
	`void matsmltf (const float *a, float b, int n, int m, float *c);`
	`void matsmlt_fr16 (const fract16 *a, fract16 b, int n, int m, fract16 *c);`

Real Matrix + Matrix Addition	`void matmadd (const double *a, const double *b, int n, int m, double *c);`
	`void matmaddf (const float *a, const float *b, int n, int m, float *c);`
	`void matmadd_fr16 (const fract16 *a, const fract16 *b, int n, int m, fract16 *c);`

Real Matrix – Matrix Subtraction	`void matmsub (const double *a, const double *b, int n, int m, double *c);`
	`void matmsubf (const float *a, const float *b, int n, int m, float *c);`
	`void matmsub_fr16 (const fract16 *a, const fract16 *b, int n, int m, fract16 *c);`

Real Matrix * Matrix Multiplication	`void matmmlt (const double *a, int n, int k, const double *b, int m, double *c);`
	`void matmmltf (const float *a, int n, int k, const float *b, int m, float *c);`
	`void matmmlt_fr16 (const fract16 *a, int n, int k, const fract16 *b, int m, fract16 *c);`

Complex Matrix + Scalar Addition	`void cmatsadd (const complex_double *a, complex_double b, int n, int m, complex_double *c);`
	`void cmatsaddf (const complex_float *a, complex_float b, int n, int m, complex_float *c);`
	`void cmatsadd_fr16 (const complex_fract16 *a, complex_fract16 b, int n, int m, complex_fract16 *c);`

Complex Matrix – Scalar Subtraction	`void cmatssub (const complex_double *a, complex_double b, int n, int m, complex_double *c);`
	`void cmatssubf (const complex_float *a, complex_float b, int n, int m, complex_float *c);`
	`void cmatssub_fr16 (const complex_fract16 *a, complex_fract16 b, int n, int m, complex_fract16 *c);`

Table 4-5. Matrix Functions (Cont'd)

Description	Prototype

Complex Matrix * Scalar Multiplication	`void cmatsmlt (const complex_double *a, complex_double b, int n, int m, complex_double *c);`
	`void cmatsmltf (const complex_float *a, complex_float b, int n, int m, complex_float *c);`
	`void cmatsmlt_fr16 (const complex_fract16 *a, complex_fract16 b, int n, int m, complex_fract16 *c);`

Complex Matrix + Matrix Addition	`void cmatmadd (const complex_double *a, const complex_double *b, int n, int m, complex_double *c);`
	`void cmatmaddf (const complex_float *a, const complex_float *b, int n, int m, complex_float *c);`
	`void cmatmadd_fr16 (const complex_fract16 *a, const complex_fract16 *b, int n, int m, complex_fract16 *c);`

Complex Matrix – Matrix Subtraction	`void cmatmsub (const complex_double *a, const complex_double *b, int n, int m, complex_double *c);`
	`void cmatmsubf (const complex_float *a, const complex_float *b, int n, int m, complex_float *c);`
	`void cmatmsub_fr16 (const complex_fract16 *a, const complex_fract16 *b, int n, int m, complex_fract16 *c);`

Complex Matrix * Matrix Multiplication	`void cmatmmlt (const complex_double *a, int n, int k, const complex_double *b, int m, complex_double *c);`
	`void cmatmmltf (const complex_float *a, int n, int k, const complex_float *b, int m, complex_float *c);`
	`void cmatmmlt_fr16 (const complex_fract16 *a, int n, int k, const complex_fract16 *b, int m, complex_fract16 *c);`

Transpose	`void transpm (const double *a, int n, int m, double *c);`
	`void transpmf (const float *a, int n, int m, float *c);`
	`void transpm_fr16 (const fract16 *a, int n, int m, fract16 *c);`

stats.h — Statistical Functions

Table 4-6 lists the statistical functions defined in the `stats.h` header file that are described in "DSP Run-Time Library Reference" on page 4-22.

Table 4-6. Stats Functions

Description	Prototype

Autocoherence	`void autocohf (const float a[ ], int n,`
	`int m, float c[ ]);`
	`void autocoh_fr16 (const fract16 a[ ], int n,`
	`int m, fract16 c[ ]);`

Autocorrelation	`void autocorrf (const float a[ ], int n,`
	`int m, float c[ ]);`
	`void autocorr_fr16 (const fract16 a[ ], int n,`
	`int m, fract16 c[ ]);`

Cross-coherence	`void crosscohf (const float a[ ], const float b[ ], int n,`
	`int m, float c[ ]);`
	`void crosscoh_fr16 (const fract16 a[ ], const fract16 b[ ],`
	`int n, int m, fract16 c[ ]);`

Cross-correlation	`void crosscorrf (const float a[ ], const float b[ ], int n,`
	`int m, float c[ ]);`
	`void crosscorr_fr16 (const fract16 a[ ], const fract16 b[ ],`
	`int n, int m, fract16 c[ ]);`

Histogram	`void histogramf (const float a[ ], int c[ ], float max,`
	`float min, int n, int m);`
	`void histogram_fr16 (const fract16 a[ ], int c[ ],`
	`fract16 max, fract16 min, int n, int m);`

Mean	`float meanf (const float a[ ], int n);` `fract16 mean_fr16 (const fract16 a[ ], int n);`
Root Mean Square	`float rmsf (const float a[ ], int n);` `fract16 rms_fr16 (const fract16 a[ ], int n);`
Variance	`float varf (const float a[ ], int n)` `fract16 var_fr16 (const fract16 a[ ], int n);`
Count Zero Crossing	`float zero_crossf (const float a[ ], int n);` `fract16 zero_cross_fr16 (const fract16 a[ ], int n);`

vector.h — Vector Functions

The `vector.h` header file contains functions for operating on real and complex vectors, both vector-scalar and vector-vector operations. See "complex.h — Basic Complex Arithmetic Functions" on page 4-4 for the definition of the complex types.

The functions in the `vector.h` header file are listed in Table 4-7. In the **Prototype** column, `a[]` and `b[]` are input vectors, `b` is an input scalar, `c[]` is an output vector and `n` is the number of elements.

The functions described by this header assume that input array arguments are constant—their contents will not change during the course of the routine. In particular, this means the input arguments do not overlap with any output argument.

Table 4-7. Vector Functions

Description	Prototype

Real
Vector + Scalar
Addition
	`void vecsadd (const double a [ ], double b, double c [ ], int n);`
	`void vecsaddf (const float a[ ], float b, float c [ ], int n);`
	`void vecsadd_fr16 (const fract16 a [ ], fract16 b, fract16 c [ ], int n);`

Real
Vector – Scalar
Subtraction
	`void vecssub (const double a [ ], double b, double c [ ], int n);`
	`void vecssubf (const float a [ ], float b, float c [ ], int n);`
	`void vecssub_fr16 (const fract16 a [ ], fract16 b, fract16 c [ ], int n);`

Real
Vector * Scalar
Multiplication
	`void vecsmlt (const double a [ ], double b, double c [ ], int n);`
	`void vecsmltf (const float a [ ], float b, float c [ ], int n);`
	`void vecsmlt_fr16 (const fract16 a [ ], fract16 b, fract16 c [ ], int n);`

Table 4-7. Vector Functions (Cont'd)

Description	Prototype

Real
Vector + Vector
Addition
	`void vecvadd (const double a [ ], const double b [ ],`
	`double c [ ], int n);`
	`void vecvaddf (const float a [ ], const float b [ ],`
	`float c [ ], int n);`
	`void vecvadd_fr16 (const fract16 a [ ], const fract16 b[ ],`
	`fract16 c [ ], int n);`

Real
Vector – Vector
Subtraction
	`void vecvsub (const double a [ ], const double b [ ],`
	`double c [ ], int n);`
	`void vecvsubf (const float a [ ], const float b [ ],`
	`float c [ ], int n);`
	`void vecvsub_fr16 (const fract16 a [ ], const fract16 b [ ],`
	`fract16 c [ ], int n);`

Real
Vector * Vector
Multiplication
	`void vecvmlt (const double a [ ], const double b [ ],`
	`double c [ ], int n);`
	`void vecvmltf (const float a [ ], const float b [ ],`
	`float c [ ], int n);`
	`void vecvmlt_fr16 (const fract16 a [ ], const fract16 b [ ],`
	`fract16 c [ ], int n);`

Maximum Value of Vector Elements	`double vecmax (const double a [ ], int n);`
	`float vecmaxf (const float a [ ], int n);`
	`fract16 vecmax_fr16 (const fract 16 a [ ], int n);`

Minimum Value of Vector Elements	`double vecmin (const double a [ ], int n);`
	`float vecminf (const float a [ ], int n);`
	`fract16 vecmin_fr16 (const fract16 a [ ], int n);`

Index of
Maximum Value of Vector Elements
	`int vecmaxloc (const double a [ ], int n);`
	`int vecmaxlocf (const float a [ ], int n);`
	`int vecmaxloc_fr16 (const fract16 a [ ], int n);`

Index of Minimum Value of Vector Elements	`int vecminloc (const double a [ ], int n);`
	`int vecminlocf (const float a [ ], int n);`
	`int vecminloc_fr16 (const fract16 a [ ], int n);`

Table 4-7. Vector Functions (Cont'd)

Description	Prototype

Complex Vector + Scalar Addition	`void cvecsadd (const complex_double a [ ], complex_double b, complex_double c [ ], int n);`
	`void cvecsaddf (const complex_float a [ ], complex_float b, complex_float c [ ], int n);`
	`void cvecsadd_fr16 (const complex_fract16 a [ ], complex_fract16 b, complex_fract16 c [ ], int n);`

Complex Vector – Scalar Subtraction	`void cvecssub (const complex_double a [ ], complex_double b, complex_double c [ ], int n);`
	`void cvecssubf (const complex_float a [ ], complex_float b, complex_float c [ ], int n);`
	`void cvecssub_fr16 (const complex_fract16 a [ ], complex_fract16 b, complex_fract16 c [ ], int n);`

Complex Vector * Scalar Multiplication	`void cvecsmlt (const complex_double a [ ], complex_double b, complex_double c [ ], int n);`
	`void cvecsmltf (const complex_float a [ ], complex_float b, complex_float c [ ], int n);`
	`void cvecsmlt_fr16 (const complex_fract16 a [ ], complex_fract16 b, complex_fract16 c [ ], int n);`

Complex Vector + Vector Addition	`void cvecvadd (const complex_double a [ ], const complex_double b [ ], complex_double c [ ], int n);`
	`void cvecvaddf (const complex_float a [ ], const complex_float b [ ], complex_float c [ ], int n);`
	`void cvecvadd_fr16 (const complex_fract16 a [ ], const complex_fract16 b [ ], complex_fract16 c [ ], int n);`

Complex Vector – Vector Subtraction	`void cvecvsub (const complex_double a [ ], const complex_double b [ ], complex_double c [ ], int n);`
	`void cvecvsubf (const complex_float a [ ], const complex_float b [ ], complex_float c [ ], int n);`
	`void cvecvsub_fr16 (const complex_fract16 a [ ], const complex_fract16 b [ ], complex_fract16 c [ ], int n);`

Complex Vector * Vector Multiplication	`void cvecvmlt (const complex_double a [ ], const complex_double b [ ], complex_double c [ ], int n);`
	`void cvecvmltf (const complex_float a [ ], const complex_float b [ ], complex_float c [ ], int n);`
	`void cvecvmlt_fr16 (const complex_fract16 a [ ], const complex_fract16 b [ ], complex_fract16 c [ ], int n);`

Table 4-7. Vector Functions (Cont'd)

Description	Prototype

Real Vector Dot Product	`double vecdot (const double a [ ],`
	`   const double b [ ], int n);`
	`float vecdotf (const float a [ ],`
	`   const float b [ ], int n);`
	`fract16 vecdot_fr16 (const fract16 a [ ],`
	`   const fract16 b [ ], int n);`

Complex Vector Dot Product	`complex_double cvecdot (const complex_double a [ ],`
	`   const complex_double b [ ], int n);`
	`complex_float cvecdotf (const complex_float a [ ],`
	`   const complex_float b [ ], int n);`
	`complex_fract16 cvecdot_fr16 (const complex_fract16 a [ ],`
	`   const complex_fract16 b [ ], complex_fract16 c [ ], int n);`

window.h — Window Generators

The `window.h` header file contains various functions to generate windows based on various methodologies. The functions defined in the `window.h` header file are listed in Table 4-8 and are described in detail in "DSP Run-Time Library Reference" on page 4-22.

For all window functions, a stride parameter `a` can be used to space the window values. The window length parameter `n` equates to the number of elements in the window. Therefore, for a stride `a` of 2 and a length `n` of 10, an array of length 20 is required, where every second entry is untouched.

Table 4-8. Window Generator Functions

Description	Prototype
Generate Bartlett Window	`void gen_bartlett_fr16 (fract16 w[ ], int a, int n);`
Generate Blackman Window	`void gen_blackman_fr16 (fract16 w[ ], int a, int n);`
Generate Gaussian Window	`void gen_gaussian_fr16 (fract16 w[ ], float alpha, int a, int n);`

Table 4-8. Window Generator Functions (Cont'd)

Description	Prototype
Generate Hamming Window	`void gen_hamming_fr16 (fract16 w[ ], int a, int n);`
Generate Hanning Window	`void gen_hanning_fr16 (fract16 w[ ], int a, int n);`
Generate Harris Window	`void gen_harris_fr16 (fract16 w[ ], int a, int n);`
Generate Kaiser Window	`void gen_kaiser_fr16 (fract16 w[ ], float beta, int a, int n);`
Generate Rectangular Window	`void gen_rectangular_fr16 (fract16 w[ ], int a, int n);`
Generate Triangle Window	`void gen_triangle_fr16 (fract16 w[ ], int a, int _n);`
Generate Vonhann Window	`void gen_vonhann_fr16 (fract16 _w[ ], int a, int n);`

DSP Run-Time Library Reference

This section provides descriptions of the DSP run-time library functions.

Notation Conventions

An interval of numbers is indicated by the minimum and maximum, separated by a comma, and enclosed in two square brackets, two parentheses, or one of each. A square bracket indicates that the endpoint is included in the set of numbers; a parenthesis indicates that the endpoint is not included.

The reference pages for the library functions use the following format.

- **Name** and purpose of the function
- **Synopsis** — Required header file and functional prototype. When the functionality is provided for several data types (for example, float, double, or fract16), several prototypes are given.
- **Description** — Function specification
- **Algorithm** — High level mathematical representation of the function
- **Domain** — Range of values supported by the function
- **Notes** — Other miscellaneous notations

For some functions, the interface is presented using the "K&R" style for ease of documentation. An ANSI C prototype is provided in the header file.

a_compress

A-law compression

Synopsis

```
#include <filter.h>

void a_compress(in, out, n)
const short in[];      /* Input array                       */
short out[];           /* Output array                      */
int n;                 /* Number of elements to be compressed */
```

Description

The `a_compress` function takes a vector of linear 13-bit signed speech samples and performs A-law compression according to ITU recommendation G.711. Each sample is compressed to 8 bits and is returned in the vector pointed to by `out`.

Algorithm

```
C(k)=a-law compression of A(k)
```

for k=0 to n-1

Domain

content of input array: -4096 to 4095

a_expand

A-law expansion

Synopsis

```
#include <filter.h>

void a_expand(in, out, n)
const short in[];       /* Input array                      */
short out[];            /* Output array                     */
int n;                  /* Number of elements to be expanded */
```

Description

The `a_expand` function inputs a vector of 8-bit compressed speech samples and expands them according to ITU recommendation G.711. Each input value is expanded to a linear 13-bit signed sample in accordance with the A-law definition and is returned in the vector pointed to by `out`.

Algorithm

```
C(k)=a-law expansion of A(k)
```

for k=0 to n-1

Domain

content of input array: 0 to 255

alog

anti-log

Synopsis

```
#include <math.h>

double alog(double x);
float alogf (float x);
```

Description

The `alog` function calculates the natural (base `e`) anti-log of its argument. An anti-log function performs the reverse of a `log` function and is therefore equivalent to exponentiation.

The input argument `x` must be in the range [-87.9 , 88.6]. The function will return `HUGE_VAL` if `x` is greater than the domain, and it will return 0.0 if `x` is less than the domain.

Algorithm

$c = e^{x}$

Domain

x = [-87.9 , 88.6]   for `alog()`, `alogf()`

Example

```
#include <math.h>

double y;
y = alog(1.0);          /* y = 2.71828... */
```

alog10

base 10 anti-log

Synopsis

```
#include <math.h>

double alog10(double x);
float alog10f (float x);
```

Description

The `alog` function calculates the base 10 anti-log of its argument. An anti-log function performs the reverse of a `log` function and is therefore equivalent to exponentiation. Therefore, `alog10(x)` is equivalent to `exp(x * log(10.0))`.

If the argument `x` is greater than the domain of [-38.2 , 38.5], the function will return `HUGE_VAL`. For input values less than the domain, the function will return 0.0.

Algorithm

$c = e^{(x * \log(10.0))}$

Domain

x = [-38.2 , 38.5]   for `alog10()`, `alog10f()`

Example

```
#include <math.h>

double y;
y = alog10(1.0);          /* y = 10.0 */
```

arg

get phase of a complex number

Synopsis

```
#include <complex.h>

float argf (complex_float a);
double arg (complex_double a);
fract16 arg_fr16 (complex_fract16 a);
```

Description

This function computes the phase associated with a Cartesian number represented by the complex argument `a`, and returns the result.

Refer to the description of the `polar_fr16` function which explains how a phase, represented as a fractional number, is interpreted in polar notation (see "polar" on page 4-91).

Algorithm

$$c = atan\left(\frac{\mathrm{Im}(a)}{\mathrm{Re}(a)}\right)$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `argf( )`, `arg( )`
-1.0 to +1.0	for `arg_fr16( )`

Note

Im (a) /Re (a) < =1	for `arg_fr16( )`

autocoh

autocoherence

Synopsis

```
#include <stats.h>

void autocohf(a,n,m,c)
const float a[];              /* Input vector a   */
int n;                        /* Input samples    */
int m;                        /* Lag count        */
float c[];                    /* Output vector c  */
void autocoh_fr16 (a,n,m,c)
const fract16 a[];            /* Input vector a   */
int n;                        /* Input samples    */
int m;                        /* Lag count        */
fract16 c[];                  /* Output vector c  */
```

Description

This function computes the autocoherence of the input elements contained within input vector `a`, and stores the result to output vector `c`.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} (a_j - \overline{a}) * (a_{j+k} - \overline{a}) \right)$$

where k={0,1,...,m-1} and $\overline{a}$ is the mean value of input vector `a`.

Domain

-3.4 x 10^{38} to +3.4 x 10^{38}	for `autocohf( )`
-1.0 to 1.0	for `autocoh_fr16( )`

autocorr

autocorrelation

Synopsis

```
#include <stats.h>

void autocorrf(a,n,m,c)
const float a[];                  /* Input vector a          */
int n;                            /* Number of input samples */
int m;                            /* Lag count               */
float c[];                        /* Output vector c         */
void autocorr_fr16 (a,n,m,c)
const fract16 a[];                /* Input vector a          */
int n;                            /* Number of input samples */
int m;                            /* Lag count               */
fract16 c[];                      /* Output vector c         */
```

Description

This function computes the autocorrelation of the input elements contained within input vector `a`, and stores the result to output vector `c`. The autocorr function is used in digital signal processing applications such as speech analysis.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} a_j * a_{j+k} \right)$$

where k={0,1,...,m-1}

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `autocorrf( )`
-1.0 to + 1.0	for `autocorr_fr16( )`

cabs

complex absolute value

Synopsis

```
#include <complex.h>

float cabsf (complex_float a)
double cabs (complex_double a)
fract16 cabs_fr16 (complex_fract16 a)
```

Description

This function computes the complex absolute value of a complex input and returns the result.

Algorithm

$$c = \sqrt{\mathrm{Re}^2(a) + \mathrm{Im}^2(a)}$$

Domain

Re^2 (a) + Im^2 (a) <= 3.4×10^{38} for `cabsf( )`, `cabs( )`

Re^2 (a) + Im^2 (a) <= 1.0 for `cabs_fr16( )`

Note

The minimum input value for both real and imaginary parts can be less than 1/256 for `cabs_fr16` but the result may have bit error of 2 to 3 bits.

cadd

complex addition

Synopsis

```
#include <complex.h>

complex_float caddf (complex_float a, complex_float b)
complex_double cadd (complex_double a, complex_double b)
complex_fract16 cadd_fr16 (complex_fract16 a, complex_fract16 b)
```

Description

This function computes the complex addition of two complex inputs: `a`, and `b`, and returns the result.

Algorithm

```
Re(c) = Re(a) + Re(b)
Im(c) = Im(a) + Im(b)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `caddf( ), cadd( )`
-1.0 to +1.0	for `cadd_fr16( )`

cartesian

convert Cartesian to polar notation

Synopsis

```
#include <complex.h>

float cartesianf (complex_float a, float *phase);
double cartesian (complex_double a, double *phase);
fract16 cartesian_fr16 (complex_fract16 a, fract16 *phase);
```

Description

This function transforms a complex number from Cartesian notation to polar notation. The Cartesian number is represented by the argument `a` that the function converts into a corresponding magnitude, which it returns as the function's result, and a phase that is returned via the second argument `phase` .

Refer to the description of the `polar_fr16` function which explains how a phase, represented as a fractional number, is interpreted in polar notation (see "polar" on page 4-91).

Algorithm

```
magnitude = cabs(a)
phase = arg(a)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `cartesianf ( )`, `cartesian ( )`
-1.0 to +1.0	for `cartesian_fr16 ( )`

Example

```
#include <complex.h>

complex_float point = {-2.0 , 0.0};
float phase;
float mag;
mag = cartesianf (point,&phase);     /* mag = 2.0, phase = π */
```

cdiv

complex division

Synopsis

```
#include <complex.h>

complex_float cdivf (complex_float a, complex_float b)
complex_double cdiv (complex_double a, complex_double b)
complex_fract16 cdiv_fr16 (complex_fract16 a, complex_fract16 b)
```

Description

This function computes the complex division of two complex inputs: `a` and `b`, and returns the result.

Algorithm

$$\mathrm{Re}(c) = \frac{\mathrm{Re}(a) * \mathrm{Re}(b) + \mathrm{Im}(a) * \mathrm{Im}(b)}{\mathrm{Re}^2(b) + \mathrm{Im}^2(b)}$$

$$\mathrm{Im}(c) = \frac{\mathrm{Re}(b) * \mathrm{Im}(a) - \mathrm{Im}(b) * \mathrm{Re}(a)}{\mathrm{Re}^2(b) + \mathrm{Im}^2(b)}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `cdivf( )`, `cdiv( )`
-1.0 to 1.0	for `cdiv_fr16( )`

cexp

complex exponential

Synopsis

```
#include <complex.h>

complex_float cexpf (float a)
complex_double cexp (double a)
```

Description

This function computes the complex exponential of real input `a` and returns the result.

Algorithm

```
Re(c) = cos(a)
Im(c) = sin(a)
```

Domain

a = [-9099 ... 9099]   for `cexpf( )`, `cexp( )`

cfft

N point complex input FFT

Synopsis

```
#include <filter.h>

void cfft_fr16(in[], t[], out[], w[], wst, n, block_exponent,
     scale_method)
const complex_fract16 in[]   /* input sequence               */
complex_fract16 t[];         /* temporary working buffer     */
complex_fract16 out[];       /* output sequence              */
const complex_fract16 w[];   /* twiddle sequence             */
int wst;                     /* twiddle factor stride        */
int n;                       /* number of FFT points         */
int *block_exponent;         /* block exponent of output data */
int scale method;            /* scaling method desired 0-none,
                                1-static, 2-dynamic          */
```

Description

This function transforms the time domain complex input signal sequence to the frequency domain by using the radix-2 Fast Fourier Transform (FFT).

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is a multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as either the output array or as the temporary array. Specifying the input array as the temporary array will also result in increased run-time performance.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle factors. The function `twidfftrad2_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `cfft_fr16`, then the stride factor has to be set appropriately; otherwise it should be 1.

The argument `scale_method` controls how the function should scale the output to avoid overflow. If no scaling is selected by setting `scale_method` to zero, then the input signal should be sufficiently conditioned to avoid overflow. The `block_exponent` argument will be set to zero.

The function will perform static scaling if `scale_method` is set to 1. For static scaling, the function will scale intermediate results to prevent overflow. The final output will be scaled by 1/`n`, and `block_exponent` will be set to `log2(n)`.

If `scale_method` is set to 2, then the function will select dynamic scaling. Under dynamic scaling, the function will inspect the intermediate results and will only scale to avoid overflow. Dynamic scaling therefore minimizes loss of precision but at the possible cost of slightly reduced performance. The `block_exponent` argument will be set to a value between 0 (which indicates that no scaling was performed) and `log2(n)` (as if static scaling was performed).

Algorithm

$$X(k) = \sum_{n=0}^{N-1} x(n) W_N^{nk}$$

When the sequence length, `n`, equals power of four, the `cfftrad4` algorithm is also available.

Domain

Input sequence length `n` must be a power of two and at least 16.

cfftrad4

N point complex input FFT

Synopsis

```
#include <filter.h>

void cfftrad4_fr16 (in[], t[], out[], w[], wst, n,
                     block_exponent, scale_method)
const complex_fract16 in[];   /* input sequence              */
complex_fract16 t[];          /* temporary working buffer    */
complex_fract16 out[];        /* output sequence             */
const complex_fract16 w[];    /* twiddle sequence            */
int wst;                      /* twiddle factor stride       */
int n;                        /* number of FFT points        */
int *block_exponent;          /* block exponent of output data */
int scale method;             /* scaling method desired
                                 0-none, 1-static, 2-dynamic  */
```

Description

This function transforms the time domain complex input signal sequence to the frequency domain by using the radix-4 Fast Fourier Transform. The `cfftrad4_fr16` function will decimate in frequency by the radix-4 FFT algorithm.

The size of the input array `in`, the output array `out`, and the temporary working buffer t is `n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is a multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as either the output array or as the temporary array. Specifying the input array as the temporary array will also result in increased run-time performance.

The twiddle table is passed in the argument w, which must contain at least ¾n twiddle coefficients. The function twidfftrad4_fr16 may be used to initialize the array. If the twiddle table contains more coefficients than needed for a particular call on cfftrad4_fr16, then the stride factor has to be set appropriately; otherwise it should be one.

The argument scale_method controls how the function should scale the output to avoid overflow. If no scaling is selected by setting scale_method to zero, then the input signal should be sufficiently conditioned to avoid overflow. The block_exponent argument will be set to zero.

The function will perform static scaling if scale_method is set to 1. For static scaling, the function will scale intermediate results to prevent overflow. The final output will be scaled by 1/n, and block_exponent argument will be set to log2(n).

If scale_method is set to 2, then the function will select dynamic scaling. Under dynamic scaling, the function will inspect the intermediate results and will only scale to avoid overflow. Dynamic scaling therefore minimizes loss of precision but at the possible cost of slightly reduced performance. The block_exponent argument will be set to a value between 0 (which indicates that no scaling was performed) and log2(n) (as if static scaling was performed).

Algorithm

$$X(k) = \sum_{n=0}^{N-1} x(n) W_N^{nk}$$

When the sequence length, n, is not a power of four, the radix2 method must be used. See "cfft" on page 4-36 for more information.

Domain

Input sequence length n must be a power of four, and at least 16.

cfft2d

NxN point 2-D complex input FFT

Synopsis

```
#include <filter.h>

void cfft2d_fr16(*in, *t, *out, w[], wst, n, block_exponent,
     scale_method)
const complex_fract16 *in;  /* pointer to input matrix a[n][n]  */
complex_fract16 *t;         /* pointer to working buffer t[n][n] */
complex_fract16 *out;       /* pointer to output matrix c[n][n]  */
const complex_fract16 w[];  /* twiddle sequence                  */
int wst;                    /* twiddle factor stride             */
int n;                      /* number of FFT points              */
int *block_exponent;        /* block exponent of output data     */
int scale method;           /* scaling method desired 0-none,
                               1-static, 2-dynamic              */
```

Description

This function computes the two dimensional Fast Fourier Transform of the complex input matrix `a[n][n]`, and stores the result to the complex output matrix `c[n][n]`.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n*n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is a multiple of twice the FFT size.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle factors. The function `twidfft2d_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `cfft2d_fr16`, then the stride factor has to be set appropriately; otherwise it should be 1.

The arguments block_exponent and scale_method have been added for future expansion. However, the current version of the function ignores the argument and always scales the output by n*n; this is equivalent to static scaling. The function will also set block_exponent to log2(n).

Algorithm

$$c(i,j) = \sum_{k=0}^{n-1}\sum_{l=0}^{n-1} a(k,l) * e^{-2\pi j(i*k+j*l)/n}$$

where i={0,1,...,n -1}, j={0,1,2,...,n-1}

Domain

Input sequence length n must be a power of two and at least 16.

cfir

complex finite impulse response filter

Synopsis

```
#include <filter.h>

void cfir_fr16(x,y,n,s)
const complex_fract16 x[];    /* Input sample vector x    */
complex_fract16 y[];          /* Output sample vector y   */
int n;                        /* Number of input samples  */
cfir_state_fr16 *s;           /* Pointer to filter state
                                 structure                */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   int k;                     /* Number of coefficients   */
   complex_fract16 *h;        /* Filter coefficients      */
   complex_fract16 *d;        /* Start of delay line      */
   complex_fract16 *p;        /* Read/write pointer       */
}  cfir_state_fr16;
```

Description

The `cfir_fr16` function implements a complex finite impulse response (CFIR) filter. It generates the filtered response of the complex input data `x` and stores the result in the complex output vector `y`.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `cfir_init`, in the `filter.h` header file, is available to initialize the structure and is defined as:

```
#define cfir_init(state, coeffs, delay, ncoeffs)  \
   (state).h = (coeffs);  \
   (state).d = (delay);   \
   (state).p = (delay);   \
   (state).k = (ncoeffs)
```

The characteristics of the filter (passband, stopband, etc.) are dependent upon the number of complex filter coefficients and their values. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients.

Each filter should have its own delay line which is a vector of type `complex_fract16` and whose length is equal to the number of coefficients. The vector should be cleared to zero before calling the function for the first time and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

Algorithm

$$y(k) = \sum_{i=0}^{p-1} h(i) * x(k-i) \quad \text{for } k = 0,1,\ldots n$$

Domain

-1.0 to +1.0

clip

clip

Synopsis

```
#include <math.h>

int clip (int parm1, int parm2)
long int lclip (long int parm1, long int parm2)
float fclipf (float parm1, float parm2)
double fclip (double parm1, double parm2)
fract16 clip_fr16 (fract16 parm1, fract16 parm2)
```

Description

This function clips a value if it is too large.

Algorithm

```
if (|parm1| < |parm2|)
   return(parm1)
else
   return(|parm2| * signof(parm1))
```

Domain

Full range for various input parameter types.

cmlt

complex multiply

Synopsis

```
#include <complex.h>

complex_float cmltf (complex_float a, complex_float b)
complex_double cmlt (complex_double a, complex_double b)
complex_fract16 cmlt_fr16 (complex_fract16 a, complex_fract16 b)
```

Description

This function computes the complex multiplication of two complex inputs `a` and `b`, and returns the result.

Algorithm

```
Re(c) = Re(a) * Re(b) - Im(a) * Im(b)
Im(c) = Re(a) * Im(b) + Im(a) * Re(b)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `cmltf( )`, `cmlt()`
-1.0 to 1.0	for `cmlt_fr16( )`

conj

complex conjugate

Synopsis

```
#include <complex.h>

complex_float conjf (complex_float a, complex_float b)
complex_double conj (complex_double a, complex_double b)
complex_fract16 conj_fr16 (complex_fract16 a, complex_fract16 b)
```

Description

This function conjugates the complex input `a` and returns the result.

Algorithm

```
Re(c) = Re(a)
Im(c) = -Im(a)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `conjf ( )`, `conj ( )`
-1.0 to 1.0	for `conj_fr16 ( )`

convolve

convolution

Synopsis

```
#include <filter.h>

void convolve_fr16(cin1, clen1, cin2, clen2, cout)
const fract16 cin1[];       /* pointer to input sequence 1     */
int clen1;                  /* length of the input sequence 1  */
const fract16 cin2[];       /* pointer to input sequence 2     */
int clen2;                  /* length of the input sequence 2  */
float cout[];               /* pointer to output sequence      */
```

Description

This function convolves two sequences pointed to by `cin1` and `cin2`. If `cin1` points to the sequence whose length is `clen1` and `cin2` points to the sequence whose length is `clen2`, then resulting sequence pointed to by `cout` has length `clen1 + clen2 - 1`

Algorithm

Convolution between two sequences `cin1` and `cin2` is described as:

$$cout[n] = \sum_{k=0}^{clen2-1} cin1[n + k - (clen2 - 1)] \bullet cin2[(clen2 - 1) - k]$$

for `n = 0` to `clen1 + clen2-2`.
(Values for `cin1[j]` are considered to be zero for `j < 0` or `j > clen1-1`).

Example

Here is an example of a convolution where `cin1` is of length 4 and `cin2` is of length 3. If we represent `cin1` as “A” and `cin2` as “B”, the elements of the output vector are:

```
{A[0]*B[0],
  A[1]*B[0] + A[0]*B[1],
  A[2]*B[0] + A[1]*B[1] + A[0]*B[2],
  A[3]*B[0] + A[2]*B[1] + A[1]*B[2],
              A[3]*B[1] + A[2]*B[2],
                          A[3]*B[2]}
```

Domain

-1.0 to +1.0

conv2d

2-D convolution

Synopsis

```
#include <filter.h>

void conv2d_fr16(min1, mrow1, mcol1, min2, mrow2, mcol2, mout )
const fract16 *min1;          /* pointer to input matrix 1      */
int mrow1;                    /* number of rows in matrix 1     */
int mcol1;                    /* number of columns in matrix 1  */
const fract16 *min2;          /* pointer to input matrix 2      */
fract16 *mrow2;               /* number of rows in matrix 1     */
int mcol2;                    /* number of columns in matrix 2  */
fract16 *mout;                /* pointer to output matrix       */
```

Description

This function computes the two-dimensional convolution of input matrix `min1` of size `mrow1` x `mcol1` and `min2` of size `mrow2` x `mcol2` and stores the result in matrix `mout` of dimension (`mrow1 +mrow2-1`) x (`mcol1 + mcol2-1`).

Algorithm

Two dimensional input matrix `min1` is convolved with input matrix `min2`, placing the result in a matrix pointed to by `mout`.

$$mout[c,r] = \sum_{i=0}^{mcol2-1} \sum_{j=0}^{mrow2-1} min1[c+i,\ r+j] \bullet min2[(mcol2-1)-i,\ (mrow2-1)-j]$$

for c = 0 to `mcol1+mcol2-1` and r = 0 to `mrow2-1`

Domain

-1.0 to +1.0

conv2d3x3

2-D convolution

Synopsis

```
#include <filter.h>

void conv2d3x3_fr16(min1, mrow1, mcol1, min2, mout )
const fract16 *min1[];       /* pointer to input matrix 1      */
int mrow1;                   /* number of rows in matrix 1     */
int mcol1;                   /* number of columns in matrix 1  */
const fract16 *min2[];       /* pointer to input matrix 2      */
fract16 *mout[];             /* pointer to output matrix       */
```

Description

This function computes the two-dimensional convolution of matrix `min1` (size `mrow1` x `mcol1`) with matrix `min2` (size `3x3`).

Algorithm

Two dimensional input matrix `min1` is convolved with input matrix `min2`, placing the result in a matrix pointed to by `mout`.

$$mout\ [c, r] = \sum_{i=0}^{2} \sum_{j=0}^{2} min1\ [c+i,\ r+j] \bullet min2\ [2-i, 2-j]$$

for c = 0 to `mcol1+2` and r = 0 to `mrow1+2`

Domain

-1.0 to +1.0

copysign

copy the sign

Synopsis

```
#include <math.h>

float copysignf (float parm1, float parm2)
double copysign (double parm1, double parm2)
fract16 copysign_fr16 (fract16 parm1, fract16 parm2)
```

Description

This function copies the sign of the second argument to the first argument.

Algorithm

```
return (|parm1| * copysignof(parm2))
```

Domain

Full range for type of parameters used.

cot

cotangent

Synopsis

```
#include <math.h>

float cotf (float a)
double cot (double a)
```

Description

This function calculates the cotangent of its argument `a`, which is measured in radians. If `a` is outside of the domain, the function returns 0.

Algorithm

```
c = cot(a)
```

Domain

x = [–9099 ... 9099] for `cotf( )`, `cot( )`

countones

count one bits in word

Synopsis

```
#include <math.h>

int countones(int parm)
int lcountones(long parm)
```

Description

This function counts the number of one bits in `parm`.

Algorithm

$$\text{return} = \sum_{j=0}^{j=N-1} \text{bit[j] of parm}$$

where `N` is the number of bits in `parm`

crosscoh

cross-coherence

Synopsis

```
#include <stats.h>

void crosscohf(a,b,n,m,c)
const float a[];                  /* Input vector a           */
const float b[];                  /* Input vector b           */
int n;                            /* Number of input samples  */
int m;                            /* Lag count                */
float c[];                        /* Output vector c          */
void crosscoh_fr16(a,n,m,c)
const fract16 a[];                /* Input vector a           */
const fract16 b[];                /* Input vector b           */
int n;                            /* Number of input samples  */
int m;                            /* Lag count                */
fract16 c[];                      /* Output vector c          */
```

Description

This function computes the cross-coherence of the input elements contained within input vector `a` and input vector `b`, and stores the result to output vector `c`.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} (a_j - \overline{a}) * (b_{j+k} - \overline{b}) \right)$$

where k={0,1,...,m-1}, $\overline{a}$ is the mean value of input vector `a` and $\overline{b}$ is the mean value of input vector `b`.

Domain

-3.4 x 10^{38} to +3.4 x 10^{38}	for `crosscohf( )`
-1.0 to +1.0	for `crosscoh_fr16( )`

crosscorr

cross-correlation

Synopsis

```
#include <stats.h>

void crosscorrf(a,b,n,m,c)
const float a[];                  /* Input vector a            */
const float b[];                  /* Input vector b            */
int n;                            /* Number of input samples   */
int m;                            /* Lag count                 */
float c[];                        /* Pointer to output vector c */
void crosscorr_fr16(a,b,n,m,c)
const fract16 a[];                /* Input vector a            */
const fract16 b[];                /* Input vector b            */
int n;                            /* Number of input samples   */
int m;                            /* Lag count                 */
fract16 c[];                      /* Pointer to output vector c */
```

Description

This function computes the cross-correlation of the input elements contained within input vector `a` and input vector `b`, and stores the result to output vector `c`.

Algorithm

$$c_k = \frac{1}{n} * \left(\sum_{j=0}^{n-k-1} a_j * b_{j+k} \right)$$

where k={0,1,...,m-1}

Domain

-3.4 x 10^{38} to +3.4 x 10^{38}	for `crosscorrf( )`
-1.0 to +1.0	for `crosscorr_fr16( )`

csub

complex subtraction

Synopsis

```
#include <complex.h>

complex_float csubf (complex_float a, complex_float b)
complex_double csub (complex_double a, complex_double b)
complex_fract16 csub_fr16 (complex_fract16 a, complex_fract16 b)
```

Description

This function computes the complex subtraction of two complex inputs `a` and `b`, and returns the result.

Algorithm

```
Re(c) = Re(a) – Re(b)
Im(c) = Im(a) – Im(b)
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `csubf( )`, `csub( )`
-1.0 to 1.0	for `csub_fr16( )`

fir

finite impulse response filter

Synopsis

```
#include <filter.h>

void fir_fr16(x,y,n,s)
   const fract16 x[];     /* Input sample vector x            */
   fract16 y[];           /* Output sample vector y           */
   int n;                 /* Number of input samples          */
   fir_state_fr16 *s;     /* Pointer to filter state structure */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *h;            /* Filter coefficients              */
   fract16 *d;            /* Start of delay line              */
   fract16 *p;            /* Read/Write pointer               */
   int k;                 /* Number of coefficients           */
   int l;                 /* Interpolation/decimation index   */
} fir_state_fr16;
```

Description

The `fir_fr16` function implements a finite impulse response (FIR) filter. The function generates the filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples and the length of the output vector is specified by the argument `n`.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `fir_init`, in the `filter.h` header file, is available to initialize the structure and is defined as:

```
#define fir_init(state, coeffs, delay, ncoeffs. index)  \
   (state).h = (coeffs);  \
   (state).d = (delay);   \
   (state).p = (delay);   \
   (state).k = (ncoeffs); \
   (state).l = (index)
```

The characteristics of the filter (passband, stopband, and so on) are dependent upon the number of filter coefficients and their values. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to the number of coefficients. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

The structure member `s->l` is not used by `fir_fr16`. This field is normally set to an interpolation/decimation index before calling either the `fir_interp_fr16` or `fir_decima_fr16` functions.

Algorithm

$$y(i) = \sum_{j=0}^{k-1} h(j) * x(i - j) \quad \text{for } i = 0,1,..n-1$$

Domain

-1.0 to +1.0

fir_decima

FIR decimation filter

Synopsis

```
#include <filter.h>

void fir_decima_fr16(x,y,n,s)
const fract16 x[];      /* Input sample vector x           */
fract16 y[];            /* Output sample vector y          */
int n;                  /* Number of input samples         */
fir_state_fr16 *s;      /* Pointer to filter state structure */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *h;          /* Filter coefficients             */
   fract16 *d;          /* Start of delay line             */
   fract16 *p;          /* Read/Write pointer              */
   int k;               /* Number of coefficients          */
   int l;               /* Interpolation/decimation index  */
} fir_state_fr16;
```

Description

The `fir_decima_fr16` function performs an FIR-based decimation filter. It generates the filtered decimated response of the input data `x` and stores the result in the output vector `y`. The number of input samples is specified by the argument `n`, and the size of the output vector should be `n/l` where l is the decimation index.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `fir_init`, in the `filter.h` header file, is available to initialize the structure and is defined as:

```
#define fir_init(state, coeffs, delay, ncoeffs, index) \
   (state).h = (coeffs);  \
   (state).d = (delay);   \
   (state).p = (delay);   \
   (state).k = (ncoeffs); \
   (state).l = (index)
```

The characteristics of the filter are dependent upon the number of filter coefficients and their values, and on the decimation index supplied by the calling program. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients. The decimation index is supplied to the function in `s->l`.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to the number of coefficients. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

Algorithm

$$y(i) = \sum_{j=0}^{k-1} x(i*l - j) * h(k - 1 + j)$$

where i = 0,1,...,(n/l) - 1

Domain

-1.0 to + 1.0

fir_interp

FIR interpolation filter

Synopsis

```
#include <filter.h>

void fir_interp_fr16(x,y,n,s)
const fract16 x[];          /* Input sample vector x           */
fract16 y[];                /* Output sample vector y          */
int n;                      /* Number of input samples         */
fir_state_fr16 *s;          /* Pointer to filter state structure */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *h;              /* Filter coefficients             */
   fract16 *d;              /* Start of delay line             */
   fract16 *p;              /* Read/Write pointer              */
   int k;                   /* Number of coefficients          */
   int l;                   /* Interpolation/decimation index  */
} fir_state_fr16;
```

Description

The `fir_interp_fr16` function performs an FIR-based interpolation filter. It generates the interpolated filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples is specified by the argument `n`, and the size of the output vector should be `n*l` where l is the interpolation index.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `fir_init`, in the `filter.h` header file, is available to initialize the structure and is defined as:

```
#define fir_init(state, coeffs, delay, ncoeffs, index) \
   (state).h = (coeffs);  \
   (state).d = (delay);   \
   (state).p = (delay);   \
   (state).k = (ncoeffs); \
   (state).l = (index)
```

The filter characteristics are dependent upon the number of polyphase filter coefficients and their values, and on the interpolation index supplied by the calling program. A pointer to the coefficients should be stored in `s->h`, and `s->k` should be set to the number of coefficients. The interpolation index is supplied to the function in `s->l`.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to the number of coefficients. The vector should be cleared to zero before calling the function for the first time and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line, and the function uses `s->p` to keep track of its current position within the vector.

Algorithm

$$\mathrm{y}(\mathrm{l}^* i + m) = \sum_{j=0}^{k-1} x(i-j) * h((k\text{-}1+j)+m*k) \quad \text{for m} = 0,1,\ldots,\text{l-1}$$

where i = 0,1,...,n-1

Domain

-1.0 to +1.0

gen_bartlett

generate bartlett window

Synopsis

```
#include <window.h>

void gen_bartlett_fr16(w,a,N)
fract16 w[];     /* Window vector                              */
int a;           /* Address stride in samples for window vector */
int N;           /* Length of window vector                     */
```

Description

This function generates a vector containing the Bartlett window. The length of the window required is specified by the parameter N, and the stride parameter a is used to space the window values within the output vector w. The length of the output vector should therefore be N*a.

The Bartlett window is similar to the Triangle window (see "gen_triangle" on page 4-72) but has the following different properties:

- The Bartlett window always returns a window with two zeros on either end of the sequence, so that for odd n, the center section of an N+2 Bartlett window equals an N Triangle window.
- For even n, the Bartlett window is still the convolution of two rectangular sequences. There is no standard definition for the Triangle window for even n; the slopes of the Triangle window are slightly steeper than those of the Bartlett window.

Algorithm

$$w[n] = 1 - \left| \frac{n - \frac{N-1}{2}}{\frac{N-1}{2}} \right|$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_blackman

generate blackman window

Synopsis

```
#include <window.h>

void gen_blackman_fr16(w,a,N)
fract16 w[];    /* Window vector                              */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                     */
```

Description

This function generates a vector containing the Blackman window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$$w[n] = 0.42 - 0.5\cos\left(\frac{2\pi n}{N-1}\right) + 0.08\cos\left(\frac{4\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_gaussian

generate gaussian window

Synopsis

```
#include <window.h>

void gen_gaussian_fr16(w,alpha,a,N)
fract16 w[];     /* Window vector                            */
float alpha;     /* Gaussian alpha parameter                 */
int a;           /* Address stride in samples for window vector */
int N;           /* Length of window vector                  */
```

Description

This function generates a vector containing the Gaussian window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$$w(n) = \exp\left[-\frac{1}{2}\left(\alpha\frac{n - N/2 - 1/2}{N/2}\right)^2\right]$$

where n = {0, 1, 2, ..., N-1} and α is an input parameter

Domain

a > 0; N > 0; α > 0.0

gen_hamming

generate hamming window

Synopsis

```
#include <window.h>

void gen_hamming_fr16(w,a,N)
fract16 w[];      /* Window vector                              */
int a;            /* Address stride in samples for window vector */
int N;            /* Length of window vector                     */
```

Description

This function generates a vector containing the Hamming window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

$$w[n] = 0.54 - 0.46\cos\left(\frac{2\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_hanning

generate hanning window

Synopsis

```
#include <window.h>

void gen_hanning_fr16(w,a,N)
fract16 w[];    /* Window vector                              */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                     */
```

Description

This function generates a vector containing the Hanning window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`. This window is also known as the Cosine window.

Algorithm

$$w[n] = 0.5 - 0.5\cos\left(\frac{2\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_harris

generate harris window

Synopsis

```
#include <window.h>

void gen_harris_fr16(w,a,N)
fract16 w[];    /* Window vector                              */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                     */
```

Description

This function generates a vector containing the Harris window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`. This window is also known as the Blackman-Harris window.

Algorithm

$$w[n] = 0.35875 - 0.48829 * \cos\left(\frac{2\pi n}{N-1}\right) + 0.14128 * \cos\left(\frac{4\pi n}{N-1}\right) - 0.01168 * \cos\left(\frac{6\pi n}{N-1}\right)$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_kaiser

generate kaiser window

Synopsis

```
#include <window.h>

void gen_kaiser_fr16(w,beta,a,N)
fract16 w[];    /* Window vector                               */
float beta;     /* Kaiser beta parameter                       */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                     */
```

Description

This function generates a vector containing the Kaiser window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`. The β value is specified by parameter `beta`.

Algorithm

$$w[n] = \frac{I_0\left[\beta\left(1-\left[\frac{n-\alpha}{\alpha}\right]^2\right)^{1/2}\right]}{I_0(\beta)}$$

where n = {0, 1, 2, ..., N-1}, α = (N - 1) / 2, and $I_0(\beta)$ represents the zeroth-order modified Bessel function of the first kind.

Domain

a > 0; N > 0; β > 0.0

gen_rectangular

generate rectangular window

Synopsis

```
#include <window.h>

void gen_rectangular_fr16(w,a,N)
fract16 w[];    /* Window vector                               */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                     */
```

Description

This function generates a vector containing the Rectangular window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`. The length of the output vector should therefore be `N*a`.

Algorithm

```
w[n] = 1
```

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_triangle

generate triangle window

Synopsis

```
#include <window.h>

void gen_triangle_fr16(w,a,N)
fract16 w[];     /* Window vector                           */
int a;           /* Address stride in samples for window vector */
int N;           /* Length of window vector                 */
```

Description

This function generates a vector containing the Triangle window. The length of the window required is specified by the parameter `N`, and the stride parameter `a` is used to space the window values within the output vector `w`.

Refer to the Bartlett window (on page 4-63) regarding the relationship between it and the Triangle window.

Algorithm

For even `n`, the following equation applies:

$$w[n] = \begin{cases} \dfrac{2n+1}{N} & n < N/2 \\ \dfrac{2N-2n-1}{N} & n > N/2 \end{cases}$$

where n = {0, 1, 2, ..., N-1}

For odd n, the following equation applies:

$$w[n] = \begin{cases} \dfrac{2n+2}{N+1} & n < N/2 \\ \dfrac{2N-2n}{N+1} & n > N/2 \end{cases}$$

where n = {0, 1, 2, ..., N-1}

Domain

a > 0; N > 0

gen_vonhann

generate von hann window

Synopsis

```
#include <window.h>

void gen_vonhann_fr16(w,a,N)
fract16 w[];    /* Window vector                             */
int a;          /* Address stride in samples for window vector */
int N;          /* Length of window vector                   */
```

Description

This function is identical to `gen_hanning` window (on page 4-68).

Domain

a > 0; N > 0

histogram

histogram

Synopsis

```
#include <stats.h>

void histogramf(a,c,max,min,n,m)
const float a[];            /* Pointer to input vector a    */
int c[];                    /* Pointer to output vector c   */
float max;                  /* Maximum value of the bin     */
float min;                  /* Minimum value of the bin     */
int n;                      /* Number of input samples      */
int m;                      /* Number of bins               */

void histogram_fr16(a,c,max,min,n,m)
const fract16 a[];          /* Pointer to input vector a    */
int c[];                    /* Pointer to output vector c   */
fract16 max;                /* Maximum value of the bin     */
fract16 min;                /* Minimum value of the bin     */
int n;                      /* Number of input samples      */
int m;                      /* Number of bins               */
```

Description

The `histogram` function computes a histogram of the input vector `a` that contains `n` samples, and stores the result in the output vector `c`

The minimum and maximum value of any input sample is specified by `min` and `max`, respectively. These values are used by the function to calculate the size of each bin as `(max - min) / m`, where `m` is the size of the output vector `c`.

Any input value that is outside the range `[ min, max )` will exceed the boundaries of the output vector and will be discarded.

To preserve maximum performance while performing out of bounds checking, the `histogram_fr16` function allocates a temporary work area on the stack. The work area is allocated with `(m + 2)` elements and the stack may therefore overflow if the number of bins is sufficiently large. The size of the stack may be adjusted by making appropriate changes to the `.LDF` file.

Algorithm

Each input value is adjusted by `min`, multiplied by `1/binsize` and is rounded. The appropriate bin in the output vector is then incremented.

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `histogramf()`
-1.0 to +1.0	for `histogram_fr16()`

ifft

N point inverse FFT

Synopsis

```
#include <filter.h>

void ifft_fr16(in[], t[], out[], w[], wst, n, block_exponent,
               scale_method)
const complex_fract16 in[];   /* input sequence              */
complex_fract16 t[];          /* temporary working buffer    */
complex_fract16 out[];        /* output sequence             */
const complex_fract16 w[];    /* twiddle sequence            */
int wst;                      /* twiddle factor stride       */
int n;                        /* number of FFT points        */
int *block_exponent;          /* block exponent of output data */
int scale method;             /* scaling method desired 0-none,
                                 1-static,  2-dynamic        */
```

Description

This function transforms the frequency domain complex input signal sequence to the time domain by using the radix-2 Fast Fourier Transform.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle coefficients. The function `twidfftrad2_fr16` may be used to initialize the array. If the twiddle table contains more coefficients than needed for a particular call on `ifft_fr16`, then the stride factor has to be set appropriately; otherwise it should be 1.

The argument `scale_method` controls how the function should scale the output to avoid overflow. If no scaling is selected by setting `scale_method` to zero, then the input signal should be sufficiently conditioned to avoid overflow. The `block_exponent` argument will be set to zero.

The function will perform static scaling if `scale_method` is set to 1. For static scaling, the function will scale intermediate results to prevent overflow. The final output will be scaled by `1/n`, and `block_exponent` will be set to `log2(n)`.

If `scale_method` is set to 2, then the function will select dynamic scaling. Under dynamic scaling, the function will inspect the intermediate results and will only scale to avoid overflow. Dynamic scaling therefore minimizes loss of precision but at the possible cost of slightly reduced performance. The `block_exponent` argument will be set to a value between 0 (which indicates that no scaling was performed) and `log2(n)` (as if static scaling was performed).

Algorithm

$$x(n) = \frac{1}{N}\sum_{k=0}^{N-1} X(k) W_N^{-nk}$$

The implementation uses core FFT functions. To get the inverse effect, it first swaps the real and imaginary parts of the input, performs the direct radix-2 transformation and finally swaps the real and imaginary parts of the output.

Domain

Input sequence length `n` must be a power of two and at least 16.

ifftrad4

N point inverse FFT

Synopsis

```
#include <filter.h>

void ifftrad4_fr16 (in[], t[], out[], w[], wst, n, block_exponent,
                   scale_method)
const complex_fract16 in[];   /* input sequence                */
complex_fract16 t[];          /* temporary working buffer      */
complex_fract16 out[];        /* output sequence               */
const complex_fract16 w[];    /* twiddle sequence              */
int wst;                      /* twiddle factor stride         */
int n;                        /* number of FFT points          */
int *block_exponent;          /* block exponent of output data  */
int scale method;             /* scaling method desired 0-none,
                                 1-static,  2-dynamic          */
```

Description

This function transforms the frequency domain complex input signal sequence to the time domain by using the radix-4 Inverse Fast Fourier Transform.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `¾n` twiddle factors. The function `twidfftrad4_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `ifftrad4_fr16`, then the stride factor has to be set appropriately; otherwise it should be 1.

The argument `scale_method` controls how the function should scale the output to avoid overflow. If no scaling is selected by setting `scale_method` to zero, then the input signal should be sufficiently conditioned to avoid overflow. The `block_exponent` argument will be set to zero.

The function will perform static scaling if `scale_method` is set to 1. For static scaling, the function will scale intermediate results to prevent overflow. The final output will be scaled by `1/n`, and `block_exponent` will be set to `log2(n)`.

If `scale_method` is set to 2, then the function will select dynamic scaling. Under dynamic scaling, the function will inspect the intermediate results and will only scale to avoid overflow. Dynamic scaling therefore minimizes loss of precision but at the cost of slightly reduced performance. The `block_exponent` will be set to a value between 0 (which indicates that no scaling was performed) and `log2(n)` (if static scaling was performed).

Algorithm

$$x(n) = \frac{1}{N}\sum_{k=0}^{N-1} X(k)W_N^{-nk}$$

The implementation uses core FFT functions implemented as direct radix-4 algorithm. To get the inverse effect, it first swaps the real and imaginary parts of the input, performs the direct radix-4 transformation and finally swaps the real and imaginary parts of the output.

Domain

Input sequence length `n` must be a power of four and at least 16.

ifft2d

NxN point 2-D inverse input FFT

Synopsis

```
#include <filter.h>

void ifft2d_fr16(*in, *t, *out, w[], wst, n, block_exponent,
                 scale_method)
const complex_float *in;   /* pointer to input matrix a[n][n]   */
complex_fract16 *t;        /* pointer to working buffer t[n][n] */
complex_fract16 *out;      /* pointer to output matrix c[n][n]  */
const complex_fract16 w[]; /* twiddle sequence                  */
int wst;                   /* twiddle factor stride             */
int n;                     /* number of FFT points              */
int *block_exponent;       /* block exponent of output data     */
int scale method;          /* scaling method desired 0-none,
                              1-static, 2-dynamic               */
```

Description

This function computes a two-dimensional Inverse Fast Fourier Transform of the complex input matrix `a[n][n]` and stores the result to the complex output matrix `c[n][n]`.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n*n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as the output array.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle factors. The function `twidfft2d_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `ifft2d_fr16`, then the stride factor has to be set appropriately; otherwise it should be 1.

The arguments `block_exponent` and `scale_method` have been added for future expansion. However, the current version of the function ignores the argument and always scales the output by `n*n`; this is equivalent to static scaling. The function will also set `block_exponent` to `log2(n)`.

Algorithm

$$c(i,j) = \frac{1}{n^2}\sum_{k=0}^{n-1}\sum_{l=0}^{n-1} a(k,l) * e^{2\pi j(i*k+j*l)/n}$$

where i={0,1,...,n -1}, j={0,1,2,...,n-1}

Domain

Input sequence length `n` must be a power of two and at least 16.

iir

infinite impulse response filter

Synopsis

```
#include <filter.h>

void iir_fr16(x,y,n,s)
const fract16 x[];        /* Input sample vector x          */
fract16 y[];              /* Output sample vector y         */
int n;                    /* Number of input samples        */
iir_state_fr16 *s         /* Pointer to filter state structure */
```

The function uses the following structure to maintain the state of the filter.

```
typedef struct
{
   fract16 *c;             /* coefficients                  */
   fract16 *d;             /* start of delay line           */
   int k;                  /* number of bi-quad stages      */
} iir_state_fr16;
```

Description

The `iir_fr16` function implements a bi-quad, canonical form, infinite impulse response (IIR) filter. It generates the filtered response of the input data `x` and stores the result in the output vector `y`. The number of input samples and the length of the output vector is specified by the argument `n`.

The function maintains the filter state in the structured variable `s`, which must be declared and initialized before calling the function. The macro `iir_init`, in the `filter.h` header file, is available to initialize the structure and is defined as:

```
#define iir_init(state, coeffs, delay, stages)  \
   (state).c = (coeffs);  \
   (state).d = (delay);   \
   (state).k = (stages)
```

The characteristics of the filter are dependent upon filter coefficients and the number of stages. Each stage has five coefficients which must be stored in the order `B2`, `B1`, `B0`, `A2`, `A1`. The value of `A0` is implied to be 1.0 and `A1` and `A2` should be scaled accordingly. This requires that the value of the A0 coefficient is greater than both `A1` and `A2` for all the stages. A pointer to the coefficients should be stored in `s->c`, and `s->k` should be set to the number of stages.

> The `iir_fr16` function is implemented using a direct form II algorithm. When importing coefficients from a filter design tool that employs a transposed direct form II, the `A1` and `A2` coefficients have to be negated. For example, if a filter design tool returns `A = [1.0, 0.2, -0.9]`, then the A coefficients have to be modified to `A = [1.0, -0.2, 0.9]`.

Each filter should have its own delay line which is a vector of type `fract16` and whose length is equal to twice the number of stages. The vector should be initially cleared to zero and should not otherwise be modified by the user program. The structure member `s->d` should be set to the start of the delay line.

Algorithm

$$H(z) = \frac{B_0 + B_1 z^{-1} + B_2 z^{-2}}{1 - A_1 z^{-1} - A_2 z^{-2}}$$

where

$$D_m = A_2 * D_{m-2} + A_1 * D_{m-1} + x_m$$
$$Y_m = B_2 * D_{m-2} + B_1 * D_{m-1} + B_0 * D_m$$

where m={0,1,2,...,n-1}

Domain

-1.0 to +1.0

max

maximum

Synopsis

```
#include <math.h>

int max (int parm1, int parm2)
long int lmax (long int parm1, long int parm2)
float fmaxf (float parm1, float parm2)
double fmax (double parm1, double parm2)
fract16 max_fr16 (fract16 parm1, fract16 parm2)
```

Description

This function returns the larger of its two arguments.

Algorithm

```
if (parm1 > parm2)
   return(parm1)
else
   return(parm2)
```

Domain

Full range for type of parameters.

mean

mean

Synopsis

```
#include <stats.h>

fract16 mean_fr16(a,n)
const fract16 a[];           /* Input vector a          */
int n;                       /* Number of input samples  */
float meanf(a,n)
const float a[];             /* Input vector a          */
int n;                       /* Number of input samples  */
```

Description

This function computes the mean of the input elements contained within input vector `a` and returns the result.

Algorithm

$$c = \frac{1}{n} * \left(\sum_{i=0}^{n-1} a_i\right)$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `meanf( )`
-1.0 to +1.0	for `mean_fr16( )`

min

minimum

Synopsis

```
#include <math.h>

int min (int parm1, int parm2)
long int lmin (long int parm1, long int parm2)
float fminf (float parm1, float parm2)
double fmin (double parm1, double parm2)
fract16 min_fr16 (fract16 parm1, fract16 parm2)
```

Description

This function returns the smaller of its two arguments.

Algorithm

```
if (parm1 < parm2)
   return(parm1)
else
   return(parm2)
```

Domain

Full range for type of parameters used.

mu_compress

μ-law compression

Synopsis

```
#include <filter.h>

void mu_compress(in, out, n)
const short in[];     /* Input array                          */
short out[];          /* Output array                         */
int n;                /* Number of elements to be compressed */
```

Description

The `mu_compress` function takes a vector of linear 14-bit signed speech samples and performs μ-law compression according to ITU recommendation G.711. Each sample is compressed to 8 bits and is returned in the vector pointed to by `out`.

Algorithm

```
C(k)= mu_law compression of A(k)
```

for k=0 to n-1

Domain

Content of input array: -8192 to 8191

mu_expand

µ-law expansion

Synopsis

```
#include <filter.h>

void mu_expand(in, out, n)
const short in[];      /* Input array                    */
short out[];           /* Output array                   */
int n;                 /* Number of elements to be expanded */
```

Description

The `mu_expand` function inputs a vector of 8-bit compressed speech samples and expands them according to ITU recommendation G.711. Each input value is expanded to a linear 14-bit signed sample in accordance with the µ-law definition and is returned in the vector pointed to `out`.

Algorithm

```
C(k)= mu_law expansion of A(k)
```

for k=0 to n-1

Domain

Content of input array: 0 to 255

norm

normalization

Synopsis

```
#include <complex.h>

complex_float normf (complex_float a)
complex_double norm (complex_double a)
```

Description

This function normalizes the complex input `a` and returns the result.

Algorithm

$$\mathrm{Re}(c) = \frac{\mathrm{Re}(a)}{\sqrt{\mathrm{Re}^2(a) + \mathrm{Im}^2(a)}}$$

$$\mathrm{Im}(c) = \frac{\mathrm{Im}(a)}{\sqrt{\mathrm{Re}^2(a) + \mathrm{Im}^2(a)}}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `normf ()`

polar

construct from polar coordinates

Synopsis

```
#include <complex.h>

complex_float polarf (float magnitude, float phase)
complex_double polar (double magnitude, double phase)
complex_fract16 polar_fr16 (fract16 magnitude, fract16 phase)
```

Description

This function transforms the polar coordinate, specified by the arguments `magnitude` and `phase`, into a Cartesian coordinate and returns the result as a complex number in which the x-axis is represented by the real part, and the y-axis by the imaginary part. The phase argument is interpreted as radians.

For the `polar_fr16` function, the phase must be scaled by 2π and must be in the range `[0x8000, 0x7ff0]`. The value of the phase may be either positive or negative. Positive values are interpreted as an anti-clockwise motion around a circle with a radius equal to the magnitude as shown in Table 4-9.

Table 4-9. Positive Phases

Phase	Radians
0.0	0
0.25 (0x2000)	π/2
0.50 (0x4000)	π
0.75 (0x6000)	3/2π
0.999 (0x7ff0)	<2π

Table 4-10 shows how negative values for the phase argument are interpreted as a clockwise movement around a circle.

Table 4-10. Negative Phases

Phase	Radians
-0.25 (0xe000)	$3/2\pi$
-0.50 (0xc000)	π
-0.75 (0xa000)	$\pi/2$
-1.00 (0x8000)	2π

Algorithm

Re(c) = r*cos(θ)

Im(c) = r*sin(θ)

where θ is the phase, and `r` is the magnitude

Domain

phase = [-9099 ... 9099] for `polarf( )`, `polar( )`

magnitude = $-3.4 \times 10^{38} \ldots +3.4 \times 10^{38}$ for `polarf( )`, `polar( )`

phase = [-1.0 ...+0.999969] for `polar_fr16( )`

magnitude = [-1.0 ... 1.0) for `polar_fr16( )`

Example

```
#include <complex.h>

#define PI 3.14159265

complex_fract16 point;
float phase_float;
```

```
fract16 phase_fr16;
fract16 mag_fr16;

phase_float = PI;
phase_fr16 = (phase_float / (2*PI)) * 32768.0;
mag_fr16 = 0x0200;

point = polar_fr16 (mag_fr16,phase_fr16);
        /* point.re = 0xfe00 */
        /* point.im = 0x0000 */
```

rfft

N point real input FFT

Synopsis

```
#include <filter.h>

void rfft_fr16(in[], t[], out[], w[], wst, n, block_exponent,
               scale_method)
const fract16 in[];            /* input/output sequence         */
complex_fract16 t[];           /* temporary working buffer      */
complex_fract16 out[];         /* working buffer                */
const complex_fract16 w[];     /* twiddle sequence              */
int wst;                       /* twiddle factor stride         */
int n;                         /* number of FFT points          */
int block_exponent;            /* block exponent of output data */
int scale method;              /* scaling method desired 0-none,
                                  1-static,  2-dynamic           */
```

Description

This function transforms the time domain real input signal sequence to the frequency domain by using the radix-2 FFT. The function takes advantage of the fact that the imaginary part of the input equals zero, which in turn eliminates half of the multiplications in the butterfly.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as either the output array or as the temporary array provided that the memory size of the input array is at least `2*n`. Specifying the input array as the temporary array will also result in increased run-time performance.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle factors. The function `twidfftrad2_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `rfft_fr16`, then the stride factor has to be set appropriately; otherwise it should be 1.

The argument `scale_method` controls how the function should scale the output to avoid overflow. If no scaling is selected by setting `scale_method` to zero, then the input signal should be sufficiently conditioned to avoid overflow. The `block_exponent` argument will be set to zero.

The function will perform static scaling if `scale_method` is set to 1. For static scaling, the function will scale intermediate results to prevent overflow. The final output will be scaled by `1/n`, and `block_exponent` will be set to `log2(n)`.

If `scale_method` is set to 2, then the function will select dynamic scaling. Under dynamic scaling, the function will inspect the intermediate results and will only scale to avoid overflow. Dynamic scaling therefore minimizes loss of precision but at the possible cost of slightly reduced performance. The `block_exponent` argument will be set to a value between 0 (which indicates that no scaling was performed) and `log2(n)` (as if static scaling was performed).

Algorithm

See `cfft` on page 4-36.

Domain

Input sequence length `n` must be a power of two and at least 16.

rfftrad4

N point real input FFT

Synopsis

```
#include <filter.h>

void rfftrad4_fr16(in[], t[], out[], w[], wst, n, block_exponent,
                   scale_method)
const fract16 in[];          /* input/output sequence        */
complex_fract16 t[];         /* temporary working buffer     */
complex_fract16 out[];       /* working buffer               */
const complex_fract16 w[];   /* twiddle sequence             */
int wst;                     /* twiddle factor stride        */
int n;                       /* number of FFT points         */
int *block_exponent;         /* block exponent of output data */
int scale method;            /* scaling method desired 0-none,
                                1-static,  2-dynamic          */
```

Description

This function transforms the time domain real input signal sequence to the frequency domain by using the radix-4 Fast Fourier Transform. The `rfftrad4_fr16` function takes advantage of the fact that the imaginary part of the input equals zero, which in turn eliminates half of the multiplications in the butterfly.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is multiple of twice the FFT size. If the input data can be overwritten, then optimum memory usage can be achieved by specifying the input array as either the output array or as the temporary array provided that the memory size of the input array is at least `2*n`. Specifying the input array as the temporary array will also result in increased run-time performance.

The twiddle table is passed in the argument `w`, which must contain at least ¾`n` twiddle factors. The function `twidfftrad4_fr16` may be used to initialize the array. If the twiddle table contains more factors than needed for a particular call on `rfftrad4_fr16`, then the stride factor has to be set appropriately; otherwise it should be one.

The argument `scale_method` controls how the function should scale the output to avoid overflow. If no scaling is selected by setting `scale_method` to zero, then the input signal should be sufficiently conditioned to avoid overflow. The `block_exponent` argument will be set to zero.

The function will perform static scaling if `scale_method` is set to 1. For static scaling, the function will scale intermediate results to prevent overflow. The final output will be scaled by `1/n`, and `block_exponent` will be set to `log2(n)`.

If `scale_method` is set to 2, then the function will select dynamic scaling. Under dynamic scaling, the function will inspect the intermediate results and will only scale to avoid overflow. Dynamic scaling therefore minimizes loss of precision but at the possible cost of slightly reduced performance. The `block_exponent` argument will be set to a value between 0 (which indicates that no scaling was performed) and `log2(n)` (as if static scaling was performed).

Algorithm

See `cfftrad4_fr16` on page 4-38.

Domain

Input sequence length `n` must be a power of four and at least 16.

rfft2d

NxN point 2-D real input FFT

Synopsis

```
#include <filter.h>

void rfft2d_fr16(*in, *t, *out, w[], wst, n, block_exponent,
                 scale_method)
const fract16 *in;          /* pointer to input matrix a[n][n]  */
complex_fract16 *t;         /* pointer to working buffer t[n][n] */
complex_fract16 *out;       /* pointer to output matrix c[n][n] */
const complex_fract16 w[]; /* twiddle sequence                 */
int wst;                    /* twiddle factor stride            */
int n;                      /* number of FFT points             */
int *block_exponent;        /* block exponent of output data    */
int scale method;           /* scaling method desired 0-none,
                               1-static,  2-dynamic             */
```

Description

This function computes a two-dimensional Fast Fourier Transform of the real input matrix `a[n][n]`, and stores the result to the complex output matrix `c[n][n]`.

The size of the input array `in`, the output array `out`, and the temporary working buffer `t` is `n*n`, where `n` represents the number of points in the FFT. The function does not impose any special memory alignment requirements on the arrays. However, benefits in run-time performance will be realized if the output array is allocated on an address boundary that is a multiple of twice the FFT size.

The twiddle table is passed in the argument `w`, which must contain at least `n/2` twiddle coefficients. The function `twidfft2d_fr16` may be used to initialize the array. If the twiddle table contains more coefficients than needed for a particular call on `rfft2d_fr16`, then the stride factor has to be set appropriately; otherwise it should be one.

The arguments `block_exponent` and `scale_method` have been added for future expansion. However, the current version of the function ignores the argument and always scales the output by `n*n`; this is equivalent to static scaling. The function will also set `block_exponent` to `log2(n)`.

Algorithm

$$c(i,j) = \sum_{k=0}^{n-1}\sum_{l=0}^{n-1} a(k,l) * e^{-2\pi j(i*k+j*l)/n}$$

where i={0,1,...,n -1}, j={0,1,2,...,n-1}

Domain

Input sequence length `n` must be a power of two and at least 16.

rms

root mean square

Synopsis

```
#include <stats.h>

float rmsf(a,n)
const float a[];                /* Pointer to input vector a  */
int n;                          /* Number of input samples    */
fract16 rms_fr16(a,n)
const fract16 a[];              /* Pointer to input vector a  */
int n;                          /* Number of input samples    */
```

Description

This function computes the root mean square of the input elements contained within input vector `a` and returns the result.

Algorithm

$$c = \sqrt{\frac{\sum_{i=0}^{n-1} a_i^{\,2}}{n}}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `rmsf( )`
-1.0 to +1.0	for `rms_fr16( )`

rsqrt

reciprocal square root

Synopsis

```
#include <math.h>

float rsqrtf (float a);
double rsqrt (double a);
```

Description

This function calculates the reciprocal of the square root of the number `a`. If `a` is negative, the function returns 0.

Algorithm

$c = 1/\sqrt{a}$

Domain

[0.0 to $+3.4 \times 10^{38}$ for `rsqrtf()`

twidfftrad2

generate FFT twiddle factors for radix-2 FFT

Synopsis

```
#include <filter.h>

void twidfftrad2_fr16 (complex_fract16 w[], int n)
```

Description

This function calculates complex twiddle coefficients for a radix-2 FFT with `n` points and returns the coefficients in the vector `w`. The vector `w`, known as the twiddle table, is normally calculated once and is then passed to an FFT function as a separate argument. The size of the table must be at least ½ of `n`, the number of points in the FFT.

FFTs of different sizes can be accommodated with the same twiddle table. Simply allocate the table at the maximum size. Each FFT has an additional parameter, the "stride" of the twiddle table. To use the whole table, specify a stride of 1. If the FFT uses only half the points of the largest FFT, the stride should be 2 (this takes only every other element).

Algorithm

This function takes FFT length `n` as an input parameter and generates the lookup table of complex twiddle coefficients. The samples are:

$$twid\_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid\_im(k) = -\sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., n/2 - 1}

Domain

The FFT length `n` must be a power of two and at least 16.

twidfftrad4

generate FFT twiddle factors for radix-4 FFT

Synopsis

```
#include <filter.h>

void twidfftrad4_fr16 (complex_fract16 w[], int n)
void twidfft_fr16(complex_fract16 w[], int n)
```

Description

The `twidfftrad4_fr16` function initializes a table with complex twiddle factors for a radix-4 FFT. The number of points in the FFT are defined by `n`, and the coefficients are returned in the twiddle table `w`.

The size of the twiddle table must be at least ¾`n`, the length of the FFT input sequence. A table can accommodate several FFTs of different sizes by allocating the table at maximum size, and then using the stride argument of the FFT function to specify the step size through the table. If the stride is set to 1, the FFT function uses all the table; if your FFT has only a quarter of the number of points of the largest FFT, the stride should be 4.

For efficiency, the twiddle table is normally generated once during program initialization and is then supplied to the FFT routine as a separate argument.

The `twidfft_fr16` routine is provided as an alternative to the `twidfftrad4_fr16` routine and performs the same function.

Algorithm

This function takes FFT length `n` as an input parameter and generates the lookup table of complex twiddle coefficients.

The samples generated are:

$$twid\_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid\_im(k) = -\sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., ¾n - 1}

Domain

The FFT length n must be a power of four and at least 16.

twidfft2d

generate FFT twiddle factors for 2-D FFT

Synopsis

```
#include <filter.h>

void twidfft2d_fr16 (complex_fract16 w[], int n)
```

Description

The `twidfft2d_fr16` function generates complex twiddle factors for a 2-D FFT. The size of the FFT input sequence is given by the argument `n` and the function writes the twiddle factors to the vector `w`, known as the twiddle table.

The size of the twiddle table must be `n/2`, the number of points in the FFT. Normally, the table is only calculated once and is then passed to an FFT function as an argument. A twiddle table may be used to generate several FFTs of different sizes by initializing the table for the largest FFT and then using the stride argument of the FFT function to specify the step size through the table. For example, to generate the largest FFT, the stride would be set to 1; and to generate an FFT of half this size, the stride would be set to 2.

Algorithm

This function takes FFT length `n` as an input parameter and generates the lookup table of complex twiddle coefficients.

The samples generated are:

$$twid\_re(k) = \cos\left(\frac{2\pi}{n}k\right)$$

$$twid\_im(k) = -\sin\left(\frac{2\pi}{n}k\right)$$

where k = {0, 1, 2, ..., n/2 - 1}

Domain

The FFT length `n` must be a power of two and at least 16.

var

variance

Synopsis

```
#include <stats.h>

float varf(a,n)
const float a[];              /* Pointer to input vector a  */
int n;                        /* Number of input samples    */
fract16 var_fr16(a, n)
const fract16 a[];            /* Pointer to input vector a  */
int n;                        /* Number of input samples    */
```

Description

This function computes the variance of the input elements contained within input vector a and returns the result.

Algorithm

$$c = \frac{n * \sum_{i=0}^{n-1} a_i^{\,2} - \left(\sum_{i=0}^{n-1} a_i\right)^2}{n(n-1)}$$

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$ for `varf( )`

-1.0 to +1.0 for `var_fr16( )`

zero_cross

count zero crossings

Synopsis

```
#include <stats.h>

int zero_crossf(a,n)
const float a[];              /* Pointer to input vector a  */
int n;                        /* Number of input samples    */
zero_cross_fr16 (a, n)
const fract16 a[];            /* Pointer to input vector a  */
int n;                        /* Number of input samples    */
```

Description

This function computes the number of times that a signal crosses over the zero line and returns the result. If all the input values are zero, the function returns a zero.

Algorithm

The actual algorithm is different from the one shown below because the algorithm needs to handle the case where an element of the array is zero. However, this example should give you a basic understanding.

```
if ( a(i) > 0 && a(i+1) < 0 )|| (a(i) < 0) && a(i+1) > 0 )

   the number of zeros is increased by one
```

Domain

-3.4×10^{38} to $+3.4 \times 10^{38}$	for `zero_crossf( )`
-1.0 to +1.0	for `zero_cross_fr16( )`

A COMPILER LEGACY SUPPORT

The VisualDSP++ environment and tools provide several types of support for legacy code that was developed with previous releases of the development tools. For more information on legacy code support, see the *VisualDSP++3.5 Linker and Utilities Manual for 16-Bit Processors* and *VisualDSP++ 3.5 Assembler and Preprocessor Manual for ADSP-218x and ADSP-219x DSPs*.

Tools Differences

VisualDSP++ 3.5 includes an updated C/C++ compiler, linker, and debugger, and a binary file format, ELF. Due to use of the VisualDSP++ Integrated Development and Debugging Environment (IDDE) and other enhancements, VisualDSP++ 3.5 has significant differences from Release 6.1 that you will need to be aware of. Most of these software differences originated in VisualDSP++ 3.0 release. In some cases, you will need to modify your sources to use the new tools.

Of the new features and enhancements, the following have the most impact on your existing projects:

- Some tools switches have changed. If you use any of the modified or obsolete switches, you must revise your command line scripts or batch files in order to rebuild your project.
- The code generation tools no longer support AEXE-format DSP executables (`.EXE`). They now generate ELF-format DSP executables (`.DXE`), and the debugger requires DSP executables to be in the

ELF/DWARF-2 format. As a result, AEXE-formatted files must be recompiled or reassembled in order to be debugged under VisualDSP++ 3.5. An ELF/DWARF-to-AEXE conversion utility is available in VisualDSP++ 3.5 that will perform back-conversion. An AEXE-to-ELF conversion utility performs forward conversion.

- Some assembly instructions and directives have changed from the VisualDSP 6.1 syntax, but a `-legacy` assembler switch has been provided to assemble files in the old syntax. You may need to review diagnostic messages and revise your source code in order to reassemble your source. Legacy syntax and the new syntax under VisualDSP++ 3.5 cannot be used together in the same source file. They can be mixed together within the same project, as long as they are assembled in different source files.

- Some C compiler extensions to the ISO/ANSI standard have changed. If you use any of the modified or removed extensions, you must revise your code in order to rebuild your project.

- The run-time model has changed. If you call a VisualDSP Release 6.1 assembly language subroutine from your C/C++ program, you must revise the assembly code to comply with the new rules for the C/C++ run-time environment.

- The Architecture File (`.ACH`) is no longer supported. If you re-link using your Release 6.1 object files or object libraries, you must create a Linker Description File for each object or object library before using the new Linker.

The remainder of this section describes these and other known differences between VisualDSP Releases 6.1 and VisualDSP++ 3.5. It also provides assistance when possible for making these required changes.

C/C++ Compiler and Run-Time Library

The new cc219x compiler provided in VisualDSP++ 3.5 does not support some switches and extensions that were available in the g21 compiler. As a result, the compiler supports a set of new rules for the run-time environment. This section lists the extensions and switches that have been removed, replaced, or whose function works differently than in Release 6.1.

For further details about the cc219x compiler, see Chapter 1, “Compiler”.

Segment Placement Support Keyword Changed to Section

The segment() placement keyword has changed to section(). The section() construct now precedes the variable declaration, and its argument is a string. For example,

```
section("my_sec”) int myvar;
```

For more information about the section() construct, see “Placement Support Keyword (section)” on page 1-83.

G21 Compatibility Call

The cc219x compiler provides a special G21 compatibility call that enables use of existing libraries with the new compiler. The extern OldAsmCall declaration can be added to the prototype(s) of the functions developed under Release 6.1. Your programs will be faster, smaller, and more reliable after the C code is upgraded to use the new compiler.

This convention is similar to the C++/C linkage specification.

Support for G21-Based Options And Extensions

The `cc219x` compiler supports most of the switches and extensions of the previous GNU-based compiler release. For a list of absolute or modified options, see "Compiler Switch Modifications" on page A-5.

ANSI C Extensions

The following extensions are no longer supported, or their functions have been modified:

- `typeof` — This extension was used to define the type of an expression.
- Complex types: `complex`, `creal`, `cimag`, and `conj` — These extensions were used to define complex numbers. Although you cannot write complex number literals, you can have a complex type defined with real and imaginary components. These types need to be managed by the programmer. Support for complex types using such an approach is used in the `libdsp` definitions and use of various complex types.
- Compound statements within expressions — This extension was used to declare variables within an expression. You can achieve these results using `inline` functions.
- Iterator types: `iter` and `sum` — These extensions created loop expressions that were used as a shorthand for working with arrays.
- Assigning variables to specific registers: `asm` — This extension was used to declare a variable and specify a machine register in which to store it.

If you use any of these extensions in your C source code, revise that source.

Compiler Switch Modifications

The switches listed in Table A-1 have been removed or their actions have been modified. If you use any of these switches to compile your C/C++ code, remove or replace the switch before recompiling the code with the new `cc219x` compiler.

Table A-1. C/C++ Compiler — Obsolete and Replaced Switches

Release 6.1Switch	Operation under Release 6.1	Change for VisualDSP++ 3.5
`-a`	Specify Architecture File.	Replaced with `-T <Linker Description File>` in linker command line. Subsumed in VisualDSP++ build process.
`-dD, -dM, and -dN`	Output the results of preprocessing and/or list of `#defines`.	Removed.
`-deps`	Instruct driver to rebuild only components of a target that have changed.	Removed.
`-Fno-<high|med|low>`	Disable specified optimization level.	Removed. Optimization controlled with `-O` switches.
`-fcond-mismatch`	Allow conditional expression mismatch.	Removed.
`-finline-functions`	Force function inlining.	Removed.
`-fkeep-inline-funtions`	Force keeping of inlined functions.	Removed.
`-fno-asm`	Don't recognize asm as a keyword.	Replaced with `-no-extra-keywords` (see on page 1-34)
`-fno-builtin`	Don't recognize builtin functions.	Replaced with `-no-builtin` (see on page 1-34)

Table A-1. C/C++ Compiler — Obsolete and Replaced Switches (Cont'd)

Release 6.1Switch	Operation under Release 6.1	Change for VisualDSP++ 3.5
`-fsigned-bitfields` `-funsigned-bitfields`	Control whether bit field is signed or unsigned.	Removed. Bit field is signed or unsigned based on the sign of the type definition declaring the bitfield.
`-fno-signed-bitfields` `-fno-unsigned-bit-fields`	Negative form of the `-fsigned-bitfield` and `-funsigned-bitfield`.	Removed. Bit field is signed or unsigned based on the sign of the type definition declaring the bitfield.
`-fsigned-char` `-funsigned-char`	Specify whether to default to signed or unsigned char type.	Replaced with `-signed-char` and `-unsigned-char`.
`-fsyntax-only`	Check syntax only; no output.	Replaced with `-syntax-only` (see on page 1-45)
`-fwritable-strings`	Store string constants in the writable data segment.	Removed. String constants are placed in the `seg_data1` section. Definition in the `.LDF` file can place this section in RAM or ROM.
`-imacros`	Process macro file.	Removed.
`-MD and -MMD`	Output rules for the make utility; used with -E.	Removed.
`-mboot-page=`	Specify boot page.	Removed. The ADSP-219x processors do not support paging.

`-mdmdata=`
`-mpmdata=`
`-mdcode=`
	Specify target architecture file segments.	Removed.
		You can control placement of object file segments using the `SECTIONS` command in the `.LDF` file..

`-mlistm`	Merge C code with assembler-generated code.	Removed.

Table A-1. C/C++ Compiler — Obsolete and Replaced Switches (Cont'd)

Release 6.1Switch	Operation under Release 6.1	Change for VisualDSP++ 3.5
-mno-doloops	Do not generate loop structures in assembled code.	Removed. The compiler only generates do loop control structures when it is safe to do so and the compiler is generating optimized code (-O, -Os). Interrupt handling routines should save and restore the loop stack if they are to use the same construct to avoid overflowing the loop stacks.
-mno-inits	Do not initialize variables in assembled code.	Removed.
-mpjump	Place the jump table in pm memory.	Removed.
-mreserved=	Instructs the compiler not to use specified registers.	Replaced with -reserve (see on page 1-43)
-mrom	Make the module a ROM module.	Removed.
-msmall-code	Optimize for size, not for speed.	Replaced with -Os (see on page 1-37)
-mstatic-spill	Use dm memory when all register are used.	Removed.
-nostdinc	Do not search standard system directories for header files.	Replaced with -no-std-inc (see on page 1-35)
-nostdlib	Do not use standard system libraries and startup files when linking.	Replaced with -no-std-lib (see on page 1-36)
-runhdr	Specify a particular runtime header.	Removed.
-traditional-cpp	Support some preprocessing features.	Removed.

New and Obsolete Warnings

The VisualDSP++ 3.5 compiler includes several new warning switches. These switches control the number and type of messages reported during a given compilation. They are described in Table A-2 on page A-8.

Table A-2. C/C++ Compiler — New Warning Switches

VisualDSP ++ 3.5 Warning Switch	Description
`-warn-protos`	Produce a warning when a function is called without a full prototype.
`-Wdriver-limit number`	Set a maximum *number* of driver errors.
`-Werror-limit number`	Set a maximum *number* of compiler errors.
`-Wremarks`	Indicates that the compiler may issue remarks, which are diagnostic messages even milder than warnings.
`-Wterse`	Enable terse warnings.
`-pedantic`	Causes the compiler to generate warnings for any constructs in a C or C++ source file that does not conform to the ANSI standard.
`-W<error|remark|suppress|warn> <num> [, num ...]`	Overrides the severity of specific compilation diagnostic messages, where `num` is the number representing the message to override.

The Release 6.1 warning switches that are no longer supported are listed in Table A-3 on page A-8:

Table A-3. C Compiler — Obsolete Warning Switches

`-Wall`	`-Wformat`	`-Wswitch`
`-Wchar-subsripts`	`-Wimplictit`	`-Wtrigraphs`
`-Wcomment`	`-Wparentheses`	`-Wuninitialized`
`-Werror`	`-Wreturn-type`	`-Wunused`

See “Compiler Command-Line Switches” on page 1-12 for further information on the compiler’s switch set.

Run-Time Model

The `cc219x` compiler in VisualDSP++ 3.5 produces code that is not fully compatible with the Release 6.1 run-time model. VisualDSP++ 3.5 has significant changes in the registers and stack usage. These changes are especially important if you call an assembly language subroutine from a C/C++ program, or a C/C++ function from an assembly language program. For more information about the VisualDSP++ 3.5 run-time model, see “C/C++ Run-Time Model and Environment” on page 1-153.

C/C++ Run-Time Library

This release includes a set of documented ANSI standard routines that you can call from your C/C++ programs. Many of these routines have been modified to provide better support for the improved performance and code compilation of `cc219x`. For complete information on the library contents, see Chapter 3, “C/C++ Run-Time Library”.

The `cc219x` compiler release now includes a complete documented DSP library (`libdsp.dlb`) and associated include files, providing efficient standard functions required by DSP application designers. For details, refer to Chapter 4, “DSP Run-Time Library”.

Future `cc219x` compiler releases may include additional library functions (refer to Chapter 3, “C/C++ Run-Time Library”).

I INDEX

H

O

P

Q

R

S

T

U

V