Cray C and C++ Reference Manual

S–2179–54

CRAY

New Features

This manual was revised to describe these new features of the Cray C++ 5.4 and Cray C 8.4 releases:

Inlining compiler command options	The `-h ipa`*n* option specifies automatic inlining. Automatic inlining allows the compiler to automatically select, depending on the inlining level *n*, which functions to inline. See Section 2.13.1, page 33.
	The `-h ipafrom=`*source* [`:`*source*] option specifies explicit inlining. The *source* arguments identify the files or directories that contain the functions to consider for inlining. See Section 2.13.2, page 34.
Inlining directives	The `#pragma _CRI inline_enable` directive tells the compiler to attempt inlining at call sites. The `#pragma _CRI inline_disable` directive tells the compiler to not attempt inlining at call sites. See Section 3.10.1, page 97.
	The `#pragma _CRI inline_reset` directive returns the inlining state to the default state. See Section 3.10.2, page 98.
	The `#pragma _CRI inline_always` *name* [`,`*name*] `...` directive specifies functions that the compiler should always attempt to inline. The `#pragma _CRI inline_never` *name* [`,`*name*] `...` directive specifies functions that are never to be inlined. See Section 3.10.3, page 99.

Record of Revision

Version	*Description*
2.0	January 1996 Original Printing. This manual supports the C and C++ compilers contained in the Cray C++ Programming Environment release 2.0. On all Cray systems, the C++ compiler is Cray C++ 2.0. On Cray systems with IEEE floating-point hardware, the C compiler is Cray Standard C 5.0. On Cray systems without IEEE floating-point hardware, the C compiler is Cray Standard C 4.0.
3.0	May 1997 This rewrite supports the C and C++ compilers contained in the Cray C++ Programming Environment release 3.0, which is supported on all systems except the Cray T3D system. On all supported Cray systems, the C++ compiler is Cray C++ 3.0 and the C compiler is Cray C 6.0.
3.0.2	March 1998 This revision supports the C and C++ compilers contained in the Cray C++ Programming Environment release 3.0.2, which is supported on all systems except the Cray T3D system. On all supported Cray systems, the C++ compiler is Cray C++ 3.0.2 and the C compiler is Cray C 6.0.2.
3.1	August 1998 This revision supports the C and C++ compilers contained in the Cray C++ Programming Environment release 3.1, which is supported on all systems except the Cray T3D system. On all supported Cray systems, the C++ compiler is Cray C++ 3.1 and the C compiler is Cray C 6.1.
3.2	January 1999 This revision supports the C and C++ compilers contained in the Cray C++ Programming Environment release 3.2, which is supported on all systems except the Cray T3D system. On all supported Cray systems, the C++ compiler is Cray C++ 3.2 and the C compiler is Cray C 6.2.
3.3	July 1999 This revision supports the C and C++ compilers contained in the Cray C++ Programming Environment release 3.3, which is supported on the Cray SV1, Cray C90, Cray J90, and Cray T90 systems running UNICOS 10.0.0.5 and later, and Cray T3E systems running UNICOS/mk 2.0.4 and later. On all supported Cray systems, the C++ compiler is Cray C++ 3.3 and the C compiler is Cray C 6.3.

| 3.4 | August 2000 |
| | This revision supports the Cray C 6.4 and Cray C++ 3.4 releases running on UNICOS and UNICOS/mk operating systems. It includes updates to revision 3.3. |

| 3.4 | October 2000 |
| | This revision supports the Cray C 6.4 and Cray C++ 3.4 releases running on UNICOS and UNICOS/mk operating systems. This revision supports a new inlining level, inline4. |

| 3.6 | June 2002 |
| | This revision supports the Cray Standard C 6.6 and Cray Standard C++ 3.6 releases running on UNICOS and UNICOS/mk operating systems. |

| 4.1 | August 20, 2002 |
| | Draft version to support Cray C 7.1 and Cray C++ 4.1 releases running on UNICOS/mp operating systems. |

| 4.2 | December 20, 2002 |
| | Draft version to support Cray C 7.2 and Cray C++ 4.2 releases running on UNICOS/mp operating systems. |

| 4.3 | March 31, 2003 |
| | Draft version to support Cray C 7.3 and Cray C++ 4.3 releases running on UNICOS/mp operating systems. |

| 5.0 | June 2003 |
| | Supports Cray C++ 5.0 and Cray C 8.0 releases running on UNICOS/mp 2.1 or later operating systems. |

| 5.1 | October 2003 |
| | Supports Cray C++ 5.1 and Cray C 8.1 releases running on UNICOS/mp 2.2 or later operating systems. |

| 5.2 | April 2004 |
| | Supports Cray C++ 5.2 and Cray C 8.2 releases running on UNICOS/mp 2.3 or later operating systems. |

| 5.3 | November 2004 |
| | Supports Cray C++ 5.3 and Cray C 8.3 releases running on UNICOS/mp 2.5 or later operating systems. |

5.4   March 2005
                    Supports Cray C++ 5.4 and Cray C 8.4 releases running on UNICOS/mp 3.0 or
                    later operating systems.

Contents

Preface

The information in this preface is common to Cray documentation provided with this software release.

Accessing Product Documentation

With each software release, Cray provides books and man pages, and in some cases, third-party documentation. These documents are provided in the following ways:

- CrayDoc, the Cray documentation delivery system that allows you to quickly access and search Cray books, man pages, and in some cases, third-party documentation—Access this HTML and PDF documentation via CrayDoc at the following URLs:

  - The local network location defined by your system administrator

  - The CrayDoc public website: `docs.cray.com`

- Man pages—Access man pages by entering the `man` command followed by the name of the man page. For more information about man pages, see the `man`(1) man page by entering:

  ```
  % man man
  ```

- Third-party documentation not provided through CrayDoc—Access this documentation, if any, according to the information provided with that product.

Conventions

These conventions are used throughout Cray documentation:

Convention	Meaning
command	This fixed-space font denotes literal items, such as file names, pathnames, man page names, command names, and programming language elements.
variable	Italic typeface indicates an element that you will replace with a specific value. For instance, you may replace *filename* with the name datafile in your program. It also denotes a word or concept being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[ ]	Brackets enclose optional portions of a syntax representation for a command, library routine, system call, and so on.
...	Ellipses indicate that a preceding element can be repeated.
name(N)	Denotes man pages that provide system and programming reference information. Each man page is referred to by its name followed by a section number in parentheses.

Enter:

```
% man man
```

to see the meaning of each section number for your particular system.

Reader Comments

Contact us with any comments that will help us to improve the accuracy and usability of this document. Be sure to include the title and number of the document with your comments. We value your comments and will respond to them promptly. Contact us in any of the following ways:

E-mail:
`docs@cray.com`

Telephone (inside U.S., Canada):
1–800–950–2729 (Cray Customer Support Center)

Telephone (outside U.S., Canada):
+1–715–726–4993 (Cray Customer Support Center)

Mail:
Software Publications
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120–1128
USA

Introduction [1]

The Cray C++ Programming Environment contains both the Cray C and C++ compilers. The Cray C compiler conforms to the International Organization of Standards (ISO) standard ISO/IEC 9899:1999 (C99). The Cray C++ compiler conforms to the ISO/IEC 14882:1998 standard, with some exceptions. The exceptions are noted in Appendix D, page 203.

Throughout this manual, the differences between the Cray C and C++ compilers are noted when appropriate. When there is no difference, the phrase *the compiler* refers to both compilers.

The information is presented as follows:

* Chapter 1, page 1 contains introductory information.

* Chapter 2, page 7 contains information on the commands used to invoke the compilers (`CC`, `cc`, `c89`, and `c99`) and the precompiler (`cpp`).

* Chapter 3, page 67 contains information on the `#pragma` directives supported by the Cray C and C++ compilers.

* Chapter 4, page 101 describes Cray Streaming directives (CSDs)

* Chapter 5, page 113 contains information about the OpenMP C and C++ API

* Chapter 6, page 119 contains information about Cray Unified Parallel C (UPC).

* Chapter 7, page 141 contains information about supported and unsupported standard C++ features and about the Dinkum C++ library.

* Chapter 8, page 143 contains information on Cray C++ template instantiation.

* Chapter 9, page 153 contains information on the extensions to the C and C++ languages.

* Chapter 10, page 157 contains information on predefined macros.

* Chapter 11, page 161 describes methods for launching applications.

* Chapter 12, page 163 contains information on debugging Cray C and C++ code.

* Chapter 13, page 165 contains information on interlanguage communication.

- Chapter 14, page 183 contains information on implementation-defined behavior.

- Appendix A, page 191 contains information on requirements for non-C99 code.

- Appendix B, page 193 contains information on the libraries and the loader.

- Appendix C, page 195 contains information on using C++ code developed under Cray C++ Programming Environment 3.5 release or earlier.

- Appendix D, page 203 contains information on the Cray C and C++ dialects.

- Appendix E, page 219 contains information on how to extract information on compiler messages and how to use the message system.

- Appendix F, page 225 contains information on intrinsic functions.

1.1 The Trigger Environment

The user on the Cray X1 series system interacts with the system as if all elements of the Programming Environment are hosted on the Cray X1 series mainframe, including Programming Environment commands hosted on the Cray Programming Environment Server (CPES). CPES-hosted commands have corresponding commands on the Cray X1 series mainframe that have the same names. These are called *triggers.* Triggers are required only for the Programming Environment.

Understanding the trigger environment will aid administrators and end users in identifying what part of the system a problem occurs when using the trigger environment.

When a user enters the name of a CPES-hosted command on the command line of the Cray X1 series mainframe, the corresponding trigger executes, which sets up an environment for the CPES-hosted command. This environment duplicates the portion of the current working environment on the Cray X1 series mainframe that relates to the Programming Environment. This allows the CPES-hosted commands to function properly.

To replicate the current working environment, the trigger captures the current working environment on the Cray X1 series system and copies the standard I/O as follows:

- Copies the standard input of the current working environment to the standard input of the CPES-hosted command

- Copies the standard output of the CPES-hosted command to standard output of the current working environment

- Copies the standard error of the CPES-hosted command to the standard error of the current working environment

All catchable interrupts, quit signals, and terminate signals propagate through the trigger to reach the CPES-hosted command. Upon termination of the CPES-hosted command, the trigger terminates and returns with the CPES-hosted commands return code.

Uncatchable signals have a short processing delay before the signal is passed to the CPES-hosted command. If you execute its trigger again before the CPES-hosted command has time to process the signal, an indeterministic behavior may occur.

Because the trigger has the same name, inputs, and outputs as the CPES-hosted command, user scripts, makefiles, and batch files can function without modification. That is, running a command in the trigger environment is very similar to running the command hosted on the Cray X1 series system.

The commands that have triggers include:

- `ar`

- `as`

- `c++filt`

- `c89`

- `c99`

- `cc`

- `ccp`

- `CC`

- `ftn`

- `ftnlx`

- `ftnsplit`

- `ld`

- `nm`

- `pat_build`

- `pat_help`

- `pat_report`

- `pat_remps`

- `remps`

1.1.1 Working in the Programming Environment

To use the Programming Environment, you must work on a file system that is cross-mounted to the CPES. If you attempt to use the Programming Environment from a directory that is not cross-mounted to the CPES, you will receive this message:

```
trigexecd: trigger command cannot access current directory.
[directory] is not properly cross-mounted on host [CPES]
```

The default files used by the Programming Environment are installed in the `/opt/ctl` file system. The default include file directory is `/opt/ctl/include`. All Programming Environment products are found in the `/opt/ctl` file system.

1.1.2 Preparing the Trigger Environment

To prepare the trigger environment for use, you must use the `module` command to load the PrgEnv module. This module loads all Programming Environment products and sets up the environment variables necessary to find the include files, libraries, and product paths on the CPES and the Cray X1 series system.

Enter the following command on the command line to load the Programming Environment:

```
% module load PrgEnv
```

Loading the PrgEnv module causes all Programming Environment products to be loaded and available to the user. A user may swap an individual product in the product set, but should not unload any one product.

To see the list of products loaded by the PrgEnv module, enter the following on the command line:

```
% module list
```

If you have questions on setting up the programming environment, contact your system support staff.

1.2 General Compiler Description

Both the Cray C and C++ compilers are contained within the same Programming Environment. If you are compiling code written in C, use the cc(1), c89(1), or c99 command to compile source files. If you are compiling code written in C++, use the CC(1) command.

1.2.1 Cray C++ Compiler

The Cray C++ compiler consists of a preprocessor, a language parser, a prelinker, an optimizer, and a code generator. The Cray C++ compiler is invoked by a command called CC(1) in this manual, but it may be renamed at individual sites. The CC(1) command is described in Section 2.1, page 8, and on the CC(1) man page. Command line examples are shown in Section 2.22, page 57.

1.2.2 Cray C Compiler

The Cray C compiler consists of a preprocessor, a language parser, an optimizer, and a code generator. The Cray C compiler is invoked by a command called cc(1), c89(1), or c99(1) in this manual, but it may be renamed at individual sites. The cc(1) and c99(1) commands are discussed in Section 2.2, page 8, the c89(1) command is described in Section 2.3, page 9. All are also discussed in the CC(1) man page. Command line examples are shown in Section 2.22, page 57.

Note: C code developed under other C compilers of the Cray Programming Environments that do not conform to the C99 standard may require modification to successfully compile with the c99 command. Refer to Appendix A, page 191.

1.3 Related Publications

The following documents contain additional information that may be helpful:

- CC(1), ld(1), and pat(1) man pages

- *Optimizing Applications on Cray X1 Series Systems*

- *Cray C++ Tools Library Reference Manual*, Rogue Wave document, *Tools.h++ Introduction and Reference Manual*, publication TPD-0005

- *Cray C++ Mathpack Class Library Reference Manual* by Thomas Keefer and Allan Vermeulen, publication TPD-0006

- *LAPACK.h++ Introduction and Reference Manual, Version 1*, by Allan Vermeulen, publication TPD-0010

Compiler Commands [2]

This chapter describes the compiler commands and the environment variables necessary to execute the Cray C and C++ compilers. These are the commands for the compilers:

- `CC`, which invokes the Cray C++ compiler.

- `cc` and `c99`(1), which invoke the Cray C compiler.

- `c89`, which invokes the Cray C compiler. This command is a subset of the `cc` command. It conforms with POSIX standard (P1003.2, Draft 12).

- `cpp`, which invokes the C language preprocessor. By default, the `CC`, `cc`, `c89`, and `c99`(1) commands invoke the preprocessor automatically. The `cpp` command provides a way for you to invoke only the preprocessor component of the Cray C compiler.

A successful compilation creates an absolute binary file, named `a.out` by default, that reflects the contents of the source code and any referenced library functions. This binary file, `a.out`, can then be executed on the target system. For example, the following sequence compiles file `mysource.c` and executes the resulting executable program:

```
% cc mysource.c
% ./a.out
```

With the use of appropriate options, compilation can be terminated to produce one of several intermediate translations, including relocatable object files (`-c` option), assembly source expansions (`-S` option), or the output of the preprocessor phase of the compiler (`-P` or `-E` option). In general, the intermediate files can be saved and later resubmitted to the `CC`, `cc`, `c89`, or `c99`(1) command, with other files or libraries included as necessary.

By default, the `CC`, `cc`, `c89`, and `c99`(1) commands automatically call the loader, which creates an executable file. If only one source file is specified, the object file is deleted. If more than one source file is specified, the object files are retained. The following command creates object files `file1.o`, `file2.o`, and `file3.o`, and the executable file `a.out`:

```
% cc file1.c file2.c file3.c
```

The following command creates the executable file `a.out` only:

```
% cc file.c
```

2.1 CC **Command**

The CC command invokes the Cray C++ compiler. The CC command accepts C++ source files that have the following suffixes:

.c

.C

.i

.c++

.C++

.cc

.cxx

.Cxx

.CXX

.CC

.cpp

The .i files are created when the preprocessing compiler command option (-P) is used. The CC command also accepts object files with the .o suffix; library files with the .a suffix; and assembler source files with the .s suffix.

The CC command format is as follows:

```
CC [-c] [-C] [-d string] [-D macro[=def]] [-E] [-g] [-G level]
   [-h arg] [-I incldir] [-l libfile] [-L libdir] [-M] [-nostdinc]
   [-o outfile] [-O level] [-P] [-s] [-S] [-U macro] [-V]
   [-Wphase,"opt..."] [-Xnpes] [-Yphase,dirname] [-#] [-##] [-###]
   files ...
```

See Section 2.5, page 10 for an explanation of the command line options.

2.2 cc **and** c99 **Commands**

The cc command invokes the Cray C compiler. The cc and c99 commands accept C source files that have the .c and .i suffixes; object files with the .o suffix; library files with the .a suffix; and assembler source files with the .s suffix.

The `cc` and `c99` commands format are as follows:

```
cc or c99 [-c] [-C] [-d string] [-D macro[=def]] [-E] [-g] [-G level]
   [-h arg] [-I incldir] [-l libfile] [-L libdir] [-M] [-nostdinc]
   [-o outfile] [-O level] [-P] [-s] [-S] [-U macro] [-V]
   [-Wphase,"opt..."] [-Xnpes] [-Yphase,dirname] [-#] [-##] [-###]
   files ...
```

See Section 2.5, page 10 for an explanation of the command line options.

2.3 `c89` Command

The `c89` command invokes the Cray C compiler. This command is a subset of the `cc` command and conforms with the POSIX standard (P1003.2, Draft 12). The `c89` command accepts C source files that have a `.c` or `.i` suffix; object files with the `.o` suffix; library files with the `.a` suffix; and assembler source files with the `.s` suffix.

The `c89` command format is as follows:

```
c89 [-c] [-D macro[=def]] [-E] [-g] [-I incldir] [-l libfile] [-L libdir]
   [-o outfile] [-O level] [-s] [-U macro] [-Yphase,dirname] files ...
```

See Section 2.5, page 10 for an explanation of the command line options.

2.4 `cpp` Command

The `cpp` command explicitly invokes the preprocessor component of the Cray C compiler. Most `cpp` options are also available from the `CC`, `cc`, `c89`, and `c99` commands.

The `cpp` command format is as follows:

```
cpp [-C] [-D macro[=def]] [-E] [-I incldir] [-M] [-N] [-nostdinc] [-P]
   [-U macro] [-V] [-Yphase,dirname] [-#] [-##] [-###] [infile] [outfile]
```

The *infile* and *outfile* files are, respectively, the input and output for the preprocessor. If you do not specify these arguments, input is defaulted to

standard input (`stdin`) and output to standard output (`stdout`). Specifying a minus sign (–) for *infile* also indicates standard input.

See Section 2.5, page 10 for an explanation of the command line options.

2.5 Command Line Options

The following subsections describe options for the `CC`, `cc`, `c89`, `c99`, and `cpp` commands. These options are grouped according to function, as follows:

- Language options:

    – The standard conformance options (Section 2.6, page 12):

Section	Option
Section 2.6.1, page 12	`-h [no]c99`
Section 2.6.2, page 13	`-h [no]conform` and `-h [no]stdc`
Section 2.6.3, page 13	`-h cfront`
Section 2.6.4, page 13	`-h [no]parse_templates`
Section 2.6.5, page 14	`-h [no]dep_name`
Section 2.6.6, page 14	`-h [no]exceptions`
Section 2.6.7, page 14	`-h [no]anachronisms`
Section 2.6.8, page 14	`-h new_for_init`
Section 2.6.9, page 15	`-h [no]tolerant`
Section 2.6.10, page 15	`-h [no] const_string_literals` (`CC`)
Section 2.6.11, page 16	`-h [no]gnu` (`CC`, `cc`)

    – The template options (Section 2.7, page 19):

Section	Option
Section 2.7.1, page 19	`-h simple_templates`
Section 2.7.2, page 19	`-h [no]autoinstantiate`
Section 2.7.3, page 19	`-h one_instantiation_per_object`
Section 2.7.4, page 20	`-h instantiation_dir = ` *dirname*
Section 2.7.5, page 20	`-h instantiate=`*mode*
Section 2.7.6, page 21	`-h [no]implicitinclude`
Section 2.7.7, page 21	`-h remove_instantiation_flags`

Section 2.7.8, page 21   -h prelink_local_copy

Section 2.7.9, page 21   -h prelink_copy_if_nonlocal

- The virtual function options (Section 2.8, page 21): -h forcevtbl and -h suppressvtbl.

- General language options (Section 2.9, page 22):

<u>Section</u>   <u>Options</u>

Section 2.9.1, page 22   -h keep=*file*

Section 2.9.2, page 22   -h restrict=*args*

Section 2.9.3, page 23   -h [no]calchars

Section 2.9.4, page 24   -h [no]signedshifts

- Optimization options:

  - General optimization options (Section 2.10, page 24)

  - Multistreaming Processor (MSP) options (Section 2.11, page 29)

  - Vectorization options (Section 2.12, page 30)

  - Inlining options (Section 2.13, page 32)

  - Scalar optimization options (Section 2.14, page 35)

- Math options (Section 2.15, page 36)

- Debugging options (Section 2.16, page 39)

- Message control options (Section 2.17, page 40)

- Compilation phase control options (Section 2.18, page 42)

- Preprocessing options (Section 2.19, page 44)

- Loader options (Section 2.20, page 47)

- Miscellaneous options (Section 2.21, page 49)

- Command line examples (Section 2.22, page 57)

- Compile-time environment variables (Section 2.23, page 58)

- Run time environment variables (Section 2.24, page 59)

- OpenMP environment variables (Section 2.25, page 62)

Options other than those described in this manual are passed to the loader. For more information on the loader, see the `ld`(1) man page.

There are many options that start with `-h`. Multiple `-h` options can be specified using commas to separate the arguments. For example, the `-h parse_templates` and `-h fp0` command line options can be specified as `-h parse_templates,fp0`.

If conflicting options are specified, the option specified last on the command line overrides the previously specified option. Exceptions to this rule are noted in the individual descriptions of the options.

The following examples illustrate the use of conflicting options:

- In this example, `-h fp0` overrides `-h fp1`:

  ```
  % CC -h fp1,fp0 myfile.C
  ```

- In this example, `-h vector2` overrides the earlier vector optimization level 3 implied by the `-O3` option:

  ```
  % CC -O3 -h vector2 myfile.C
  ```

Most `#pragma` directives override corresponding command line options. For example, `#pragma _CRI novsearch` overrides the `-h vsearch` option. `#pragma _CRI novsearch` also overrides the `-h vsearch` option implied by the `-h vector2` or `-O2` option. Exceptions to this rule are noted in descriptions of options or `#pragma` directives.

2.6 Standard Language Conformance Options

This section describes standard conformance language options. Each subsection heading shows in parentheses the compiler with which the option can be used.

2.6.1 `-h [no]c99 (cc, c99)`

Default options: `-h noc99 (cc)`

                               `-h c99 (c99)`

This option enables/disables language features new to the C99 standard and Cray C compiler, while providing support for features that were previously defined as Cray extensions. If the previous implementation of the Cray extension differed from the C99 standard, both implementations will be available when the

-h c99 option is enabled. The -h c99 option is also required for C99 features not previously supported as extensions.

When -h noc99 is used, c99 language features such as VLAs and restricted pointers that were available as extensions previously to adoption of the c99 standard remain available to the user.

2.6.2 -h [no]conform (CC, cc, c99), -h [no]stdc (cc, c99)

Default option: -h [no]conform, -h nostdc

The -h conform and -h stdc options specify strict conformance to the ISO C standard or the ISO C++ standard. The -h noconform and -h nostdc options specify partial conformance to the standard. The -h exceptions, -h dep_name, -h parse_templates, and -h const_string_literals options are enabled by the -h conform option in Cray C++.

> **Note:** The c89 command does not accept the -h conform or -h stdc option. It is enabled by default when the command is issued.

2.6.3 -h cfront (CC)

The -h cfront option causes the Cray C++ compiler to accept or reject constructs that were accepted by previous cfront-based compilers (such as Cray C++ 1.0), but which are not accepted in the C++ standard. The -h anachronisms option is implied when -h cfront is specified.

2.6.4 -h [no]parse_templates (CC)

Default option: -h noparse_templates

This option allows existing code that defines templates using previous versions of the Cray STL (before Programming Environment 3.6) to compile successfully with the -h conform option. Consequently, this allows you to compile existing code without having to use the Cray C++ STL. To do this, use the noparse_templates option. Also, the compiler defaults to this mode when the -h dep_name option is used. To have the compiler verify that your code uses the Cray C++ STL properly, use the parse_templates option.

2.6.5 -h [no]dep_name (CC)

Default option: -h nodep_name

This option enables or disables dependent name processing (that is, the separate lookup of names in templates when the template is parsed and when it is instantiated). The -h dep_name option cannot be used with the -h noparse_templates option.

2.6.6 -h [no]exceptions (CC)

Default option: The default is -h exceptions; however, if the CRAYOLDCPPLIB environment variable is set to a nonzero value, the default is -h noexceptions.

The -h exceptions option enables support for exception handling. The -h noexceptions option issues an error whenever an exception construct, a try block, a throw expression, or a throw specification on a function declaration is encountered. -h exceptions is enabled by -h conform.

2.6.7 -h [no]anachronisms (CC)

Default option: -h noanachronisms

The -h [no]anachronisms option enables/disables anachronisms in Cray C++. This option is overridden by -h conform.

2.6.8 -h new_for_init (CC)

The -h new_for_init option enables the new scoping rules for a declaration in a for-init statement. This means that the new (standard-conforming) rules are in effect, which means that the entire for statement is *wrapped* in its own implicitly generated scope. -h new_for_init is implied by the -h conform option.

This is the result of the scoping rule:

```
{
  .
  .
  .
  for (int i = 0; i < n; i++)  {
    .
    .
```

```
                 .
        } // scope of i ends here for -h new_for_init
                 .
                 .
                 .
        }  // scope of i ends here by default
```

2.6.9 -h [no]tolerant (cc, c99)

Default option: -h notolerant

The -h tolerant option allows older, less standard C constructs to facilitate
porting of code written for previous C compilers. Errors involving comparisons
or assignments of pointers and integers become warnings. The compiler
generates casts so that the types agree. With -h notolerant, the compiler
is intolerant of the older constructs.

The use of the -h tolerant option causes the compiler to tolerate accessing an
object with one type through a pointer to an entirely different type. For example,
a pointer to long might be used to access an object declared with type double.
Such references violate the C standard and should be eliminated if possible. They
can reduce the effectiveness of alias analysis and inhibit optimization.

2.6.10 -h [no]const_string_literals (CC)

Default option: -h noconst_string_literals

The -h [no]const_string_literals options controls whether string literals
are const (as required by the standard) or non-const (as was true in earlier
versions of the C++ language).

2.6.11 -h [no]gnu (CC, cc)

Default option: -h nognu

The -h gnu option enables the compiler to recognize the subset of the GCC version 3.3.2 extensions to C listed in Table 1. Table 2 lists the extensions that apply only to C++.

See http://gcc.gnu.org/onlinedocs/ for detailed descriptions of the GCC C and C++ language extensions.

Table 1. GCC C Language Extensions

GCC C Language Extension	Description
Typeof	typeof: referring to the type of an expression
Lvalues	Using ?:, and casts in lvalues
Conditionals	Omitting the middle operand of a ?: expression
Long Long	Double-word integers –long long int
Complex	Data types for complex numbers
Statement Exprs	Putting statements and declarations inside expressions
Zero Length	Zero-length arrays
Variable Length	Arrays whose length is computed at run time
Empty Structures	Structures with no members; applies to C but not C++
Variadic Macros	Macros with a variable number of arguments
Escaped Newlines	Slightly looser rules for escaped newlines
Multiline strings	String literals with embedded newlines
Initializers	Non-constant initializers
Compound Literals	Compound literals give structures, unions or arrays as values
Designated Inits	Labeling elements of initializers
Cast to Union	Casting to union type from any member of the union
Case Ranges	'case 1 ... 9' and such
Mixed Declarations	Mixing declarations and code
Attribute Syntax	Formal syntax for attributes

GCC C Language Extension	Description
Function Prototypes	Prototype declarations and old-style definitions; applies to C but not C++
C++ Comments	C++ comments are recognized
Dollar Signs	Dollar sign is allowed in identifiers
Character Escapes	\e stands for the character `<ESC>`
Alignment	Inquiring about the alignment of a type or variable
Inline	Defining inline functions (as fast as macros)
Alternate Keywords	`__const__`, `__asm__`, etc., for header files
Incomplete Enums	`enum foo;`, with details to follow
Function Names	Printable strings which are the name of the current function
Return Address	Getting the return or frame address of a function
Unnamed Fields	Unnamed struct/union fields within structs/unions
Function Attributes: - nothrow - format, format_arg - deprecated - used - unused - alias - weak	Declaring that functions have no side effects, or that they can never return
Variable Attributes: - alias - deprecated - unused - used - transparent_union - weak	Specifying attributes of variables

GCC C Language Extension	Description
Type Attributes:	Specifying attributes of types
• deprecated	
• unused	
• used	
• transparent_union	
Asm Labels	Specifying the assembler name to use for a C symbol
Other Builtins:	Other built-in functions
• `__builtin_types_compatible_p`	
• `__builtin_choose_expr`	
• `__builtin_constant_p`	
• `__builtin_huge_val`	
• `__builtin_huge_valf`	
• `__builtin_huge_vall`	
• `__builtin_inf`	
• `__builtin_inff`	
• `__builtin_infl`	
• `__builtin_nan`	
• `__builtin_nanf`	
• `__builtin_nanl`	
• `__builtin_nans`	
• `__builtin_nansf`	
• `__builtin_nansl`	

Special files such as `/dev/null` may be used as source files.

The supported subset of the GCC version 3.3.2 extensions to C++ are listed in Table 2.

Table 2. GCC C++ Language Extensions

GCC C++ Extensions	Description
Min and Max	C++ minimum and maximum operators
Restricted Pointers	C99 restricted pointers and references
Backwards Compatibility	Compatibilities with earlier definitions of C++
Strong Using	A using-directive with `__attribute ((strong))`
Explicit template specializations	Attributes may be used on explicit template specializations

2.7 Template Language Options

This section describes template language options. See Chapter 8, page 143 for more information on template instantiation. Each subsection heading shows in parentheses the compiler with which the option can be used.

2.7.1 `-h simple_templates` (CC)

The `-h simple_templates` option enables simple template instantiation by the Cray C++ compiler. For more information on template instantiation, see Chapter 8, page 143. The default is `autoinstantiate`.

2.7.2 `-h [no]autoinstantiate` (CC)

Default option: `-h autoinstantiate`

The `-h [no]autoinstantiate` option enables or disables prelinker (automatic) instantiation of templates by the Cray C++ compiler. For more information on template instantiation, see Chapter 8, page 143.

2.7.3 `-h one_instantiation_per_object` (CC)

The `-h one_instantiation_per_object` option puts each template instantiation used in a compilation into a separate object file that has a `.int.o` extension. The primary object file will contain everything else that is not an instantiation. See the `-h instantiation_dir` option for the location of the object files.

2.7.4 `-h instantiation_dir=`*dirname* (`CC`)

The `-h instantiation_dir = ` *dirname* option specifies the instantiation directory that the `-h one_instantiation_per_object` option should use. If directory *dirname* does not exist, it will be created. The default directory is `./Template.dir`.

2.7.5 `-h instantiate=`*mode* (`CC`)

Default option: `-h instantiate=none`

Normally, during compilation of a source file, no template entities are instantiated (except those assigned to the file by automatic instantiation). However, the overall instantiation mode can be changed by using the `-h instantiate=`*mode* option. *mode* is specified as `none` (the default), `used`, `all`, or `local`. The default is `instantiate=none`. To change the overall instantiation mode, specify one of the following for *mode*:

none   Default. Does not automatically create instantiations of any template entities. This is the most appropriate mode when prelinker (automatic) instantiation is enabled.

used   Instantiates only those template entities that were used in the compilation. This includes all static data members that have template definitions.

all   Instantiates all template functions declared or referenced in the compilation unit. For each fully instantiated template class, all of its member functions and static data members are instantiated regardless of whether they were used. Nonmember template functions are instantiated even if the only reference was a declaration.

local   Similar to `instantiate=used` except that the functions are given internal linkage. This mode provides a simple mechanism for those who are not familiar with templates. The compiler instantiates the functions used in each compilation unit as local functions, and the program links and runs correctly (barring problems due to multiple copies of local static variables). This mode may generate multiple copies of the instantiated functions and is not suitable for production use. This mode cannot be used in conjunction with prelinker (automatic) template instantiation. Automatic template instantiation is disabled by this mode.

If CC is given a single source file to compile and link, all instantiations are done in the single source file and, by default, the instantiate=used mode is used to suppress prelinker instantiation.

2.7.6 -h [no]implicitinclude (CC)

Default option: -h implicitinclude

The -h [no]implicitinclude option enables or disables implicit inclusion of source files as a method of finding definitions of template entities to be instantiated.

2.7.7 -h remove_instantiation_flags (CC)

The -h remove_instantiation_flags option causes the prelinker to recompile all the source files to remove all instantiation flags.

2.7.8 -h prelink_local_copy (CC)

The -h prelink_local_copy indicates that only local files (for example, files in the current directory) are candidates for assignment of instantiations.

2.7.9 -h prelink_copy_if_nonlocal (CC)

The -h prelink_copy_if_nonlocal option specifies that assignment of an instantiation to a nonlocal object file will result in the object file being recompiled in the current directory.

2.8 Virtual Function Options (-h forcevtbl, -h suppressvtbl (CC))

The -h forcevtbl option forces the definition of virtual function tables in cases where the heuristic methods used by the compiler to decide on definition of virtual function tables provide no guidance. The -h suppressvtbl option suppresses the definition of virtual function tables in these cases.

The virtual function table for a class is defined in a compilation if the compilation contains a definition of the first noninline, nonpure virtual function of the class. For classes that contain no such function, the default behavior is to define the virtual function table (but to define it as a local static entity).

The -h forcevtbl option differs from the default behavior in that it does not force the definition to be local.

2.9 General Language Options

This section describes general language options. Each subsection heading shows in parentheses the compiler with which the option can be used.

2.9.1 `-h keep=`*file* (`CC`)

When the `-h keep=`*file* option is specified, the static constructor/destructor object (`.o`) file is retained as *file*. This option is useful when linking `.o` files on a system that does not have a C++ compiler. The use of this option requires that the `main` function must be compiled by C++ and the static constructor/destructor function must be included in the link. With these precautions, mixed object files (files with `.o` suffixes) from C and C++ compilations can be linked into executables by using the loader command instead of the `CC` command.

2.9.2 `-h restrict=`*args* (`CC`, `cc`, `c99`)

The `-h restrict=`*args* option globally tells the compiler to treat certain classes of pointers as restricted pointers. You can use this option to enhance optimizations (this includes vectorization).

Classes of affected pointers are determined by the value contained in *args*, as follows:

args	Description
a	All pointers to object and incomplete types are to be considered restricted pointers, regardless of where they appear in the source code. This includes pointers in `class`, `struct`, and `union` declarations, type casts, function prototypes, and so on.

⚠ **Caution:** Do not specify `restrict=a` if, during execution of any function, an object is modified and that object is referenced through either two different pointers or through the declared name of the object and a pointer. Undefined behavior may result.

f	All function parameters that are pointers to objects or incomplete types can be treated as restricted pointers.

⚠ **Caution:** Do not specify `restrict=f` if, during execution of any function, an object is modified and that object is referenced through either two different pointer function parameters or through the declared name of the object and a pointer function parameter. Undefined behavior may result.

t   All parameters that are `this` pointers can be treated as restricted pointers (Cray C++ only).

    ⚠ **Caution:** Do not specify `restrict=t` if, during execution of any function, an object is modified and that object is referenced through the declared name of the object and a `this` pointer. Undefined behavior may result.

The *args* arguments tell the compiler to assume that, in the current compilation unit, each pointer (=a), or each pointer that is a function parameter (=f), or each `this` pointer (=t) points to a unique object. This assumption eliminates those pointers as sources of potential aliasing, and may allow additional vectorization or other optimizations. These options cause only data dependencies from pointer aliasing to be ignored, rather than all data dependencies, so they can be used safely for more programs than the `-h ivdep` option.

⚠ **Caution:** Like `-h ivdep`, the arguments make assertions about your program that, if incorrect, can introduce undefined behavior. You should not use `-h restrict=a` if, during the execution of any function, an object is modified and that object is referenced through either of the following:

- Two different pointers

- The declared name of the object and a pointer

The `-h restrict=f` and `-h restrict=t` options are subject to the analogous restriction, with "function parameter pointer" replacing "pointer."

2.9.3 `-h [no]calchars` (`CC`, `cc`, `c99`)

Default option: `-h nocalchars`

The `-h calchars` option allows the use of the @ and $ characters in identifier names. This option is useful for porting codes in which identifiers include these characters. With `-h nocalchars`, these characters are not allowed in identifier names.

⚠ **Caution:** Use this option with extreme care, because identifiers with these characters are within UNICOS/mp name space and are included in many library identifiers, internal compiler labels, objects, and functions. You must prevent conflicts between any of these uses, current or future, and identifier declarations or references in your code; any such conflict is an error.

2.9.4 `-h [no]signedshifts` (CC, cc, c99)

Default option: -h signedshifts

The `-h [no]signedshifts` option affects the result of the right shift
operator. For the expression `e1 >> e2` where `e1` has a signed type, when
`-h signedshifts` is in effect, the vacated bits are filled with the sign bit of `e1`.
When `-h nosignedshifts` is in effect, the vacated bits are filled with zeros,
identical to the behavior when `e1` has an unsigned type.

Also refer to Section 14.2.5, page 187 about the effects of this option when
shifting integers.

2.10 General Optimization Options

This section describes general optimization options. Each subsection heading
shows in parentheses the compiler with which the option can be used.

2.10.1 `-h gen_private_callee` (CC, cc, c99)

The `-h gen_private_callee` option is used when compiling source
files containing routines that will be called from streamed regions, whether
those streamed regions are created by CSD directives or by the use of the
`ssp_private` or `concurrent` directives to cause autostreaming. For more
information about the `ssp_private` directive, see Section 3.8.1, page 88. For
more information about CSDs, see Chapter 4, page 101 .

2.10.2 `-h [no]aggress` (CC, cc, c99)

Default option: -h noaggress

The `-h aggress` option provides greater opportunity to optimize loops that
would otherwise by inhibited from optimization due to an internal compiler size
limitation. `-h noaggress` leaves this size limitation in effect.

With `-h aggress`, internal compiler tables are expanded to accommodate larger
loop bodies. This option can increase the compilation's time and memory size.

2.10.3 `-h display_opt`

The `-h display_opt` option displays the current optimization settings for
this compilation.

2.10.4 -h [no]fusion (CC, cc, c99)

Default option: -h fusion

The −h [no]fusion option globally allows or disallows loop fusion. By default, the compiler attempts to fuse all loops, unless the −h nofusion option is specified. Fusing loops generally increases single processor performance by reducing memory traffic and loop overhead. On rare occasions loop fusing may degrade performance.

> **Note:** Loop fusion is disabled when the vectorization level is set to 0 or 1.

See *Optimizing Applications on Cray X1 Series Systems* for more information about loop fusion.

2.10.5 -h [no]intrinsics (CC, cc, c99)

Default option: -h intrinsics

The -h intrinsics option allows the use of intrinsic hardware functions, which allow direct access to some hardware instructions or generate inline code for some functions. This option has no effect on specially-handled library functions.

Intrinsic functions are described in Appendix F, page 225.

2.10.6 -h list=*opt* (CC, cc, c99)

The -h list=*opt* option allows the creation of a loopmark listing and controls its format. The listings are written to *source_file_name_without_suffix*.lst.

For additional information on loopmark listings, see *Optimizing Applications on Cray X1 Series Systems.*

The values for *opt* are:

a   Use all list options; *source_file_name_without_suffix*.lst includes summary report, options report, and source listing

b   Add page breaks to listing

e   Expand include files

> **Note:** Using this option may result in a very large listing file. All system include files are also expanded.

i   Intersperse optimization messages within the source listing rather than at the end

m	Create loopmark listing; *source_file_name_without_suffix*.lst includes summary report and source listing
p	Causes the compiler to insert carriage control characters into column one of each line in the listing. Use this option for line printers which require the carriage control characters to control the vertical position of each printed line.

Table 3 shows the carriage control characters used.

Table 3. Carriage Control Characters

Control character	Action
1	New page
Blank	Single spacing

s	Create a complete source listing (include files not expanded)
w	Create a wide listing rather than the default of 80 characters

Using `-h list=m` creates a loopmark listing. The b, e, i, s, and w options provide additional listing features. Using `-h list=a` combines all options.

2.10.7 `-h msp` (`CC`, `cc`, `c99`)

Default option: `-h msp`

The `-h msp` option causes the compiler to generate code and to select the appropriate libraries to create an executable that runs on one or more multistreaming processors (MSP mode). Any code, including code using Cray-supported distributed memory models, can use MSP mode.

Executables compiled for MSP mode can contain object files compiled with MSP or SSP mode. That is, MSP and SSP object files can be specified during the load step as follows:

```
cc -h msp -c ... /* Produce MSP object files */
cc -h ssp -c ... /* Produce SSP object files */
/* Link MSP and SSP object files */
/* to create an executable to run on MSPs */
cc sspA.o sspB.o msp.o ...
```

For more information about MSP mode, refer to *Optimizing Applications on Cray X1 Series Systems.* For information on SSP mode, see Section 2.10.10, page 27.

2.10.8 `-h` [`no`]`pattern` (`CC`, `cc`, `c99`)

Default option: `-h pattern`

The `-h` [no]pattern option globally enables or disables pattern matching. Pattern matching is on by default. For details on pattern matching, see *Optimizing Applications on Cray X1 Series Systems.*

2.10.9 `-h` [`no`]`overindex` (`CC`, `cc`, `c99`)

Default option: `-h nooverindex`

The `-h overindex` option declares that there are array subscripts that index a dimension of an array that is outside the declared bounds of that array. The `-h nooverindex` option declares that there are no array subscripts that index a dimension of an array that is outside the declared bounds of that array.

2.10.10 `-h ssp` (`CC`, `cc`, `c99`)

Default option: `-h msp`

The `-h ssp` option causes the compiler to compile the code and select the appropriate libraries to create an executable that runs on one single-streaming processor (SSP mode). Any code, including code using Cray-supported distributed memory models, can use SSP mode.

Executables compiled for SSP mode can contain only object files compiled in SSP mode. When loading object files separately from the compile step, the SSP mode must be specified during the load step as this example shows:

```
/* Produce SSP object files */
cc -h ssp -c ...

/* Link SSP object files */
/* to create an executable to run on a single SSP */
cc -h ssp sspA.o sspB.o ...
```

Since SSP mode does not use multistreaming, the `-h ssp` option also changes the compiler's behavior in the same way as the `-h stream0` option. This option then causes the compiler to ignore CSDs.

> **Note:** Code explicitly compiled with the `-h stream0` option can be linked with object files compiled with MSP or SSP mode. You can use this option to create a universal library that can be used in MSP or SSP mode.

For more information about SSP mode, refer to *Optimizing Applications on Cray X1 Series Systems*. For information about MSP mode, see Section 2.10.7, page 26.

> **Note:** The `-h ssp` and `-h command` options both create executables that run on an SSP. The executable created via the `-h ssp` option executes on an application node. The executable created via the `-h command` option executes on the support node.

2.10.11 `-h [no]unroll` (CC, cc, c99)

Default option: `-h unroll`

The `-h [no]unroll` option globally allows or disallows unrolling of loops. By default, the compiler attempts to unroll all loops, unless the `-h nounroll` option is specified, or the `unroll0` or `unroll1` pragma (Section 3.9.5, page 94) is specified for a loop. Loop unrolling generally increases single processor performance at the cost of increased compile time and code size.

See *Optimizing Applications on Cray X1 Series Systems* for more information about loop unrolling.

2.10.12 `-O` *level* (CC, cc, c89, c99)

Default option: Equivalent to the appropriate `-h` option

The `-O` *level* option specifies the optimization level for a group of compiler features. Specifying `-O` with no argument is the same as not specifying the `-O` option; this syntax is supported for compatibility with other vendors.

A value of `0`, `1`, `2`, or `3` sets that level of optimization for each of the `-h scalar`*n*, `-h stream`*n*, and `-h vector`*n* options.

For example, `-O2` is equivalent to the following:

`-h scalar2,stream2,vector2`

Optimization features specified by `-O` are equivalent to the `-h` options listed in Table 4.

Table 4. `-h` Option Descriptions

`-h` option	Description location
`-h stream`*n*	Section 2.11.1, page 29
`-h vector`*n*	Section 2.12.3, page 31
`-h scalar`*n*	Section 2.14.2, page 35

2.11 Multistreaming Processor Optimization Options

This section describes the multistreaming processor (MSP) options. For information on MSP `#pragma` directives, see Section 3.8, page 87. For information about streaming intrinsics, see Appendix F, page 225. Each subsection heading shows in parentheses the compiler command with which the option can be used.

These options cannot be used in SSP mode, which is enabled with the `-h ssp` option.

2.11.1 `-h stream`*n* `(CC, cc, c99)`

The `-h stream`*n* option specifies the level of automatic MSP optimizations to be performed. Generally, vectorized applications that execute on a one-processor system can expect to execute up to four times faster on a processor with multistreaming enabled.

The default is `-h stream2`.

These can be used for the *n* argument:

n	Description
0	No automatic multistreaming optimizations are performed.
1	Conservative automatic multistreaming optimizations. Automatic multistreaming optimization is limited to inner vectorized loops and some bit matrix multiplication (BMM) operations. MSP operations performed generate the same results that would be obtained from scalar optimizations; for example, no floating-point reductions are performed. This level is compatible with `-h vector1`, `2`, and `3`.

2	Moderate automatic multistreaming optimizations. Automatic multistreaming optimization is performed on loop nests and appropriate BMM operations.

This option also enables conditional streaming. Conditional streaming allows runtime selection between streamed and nonstreamed versions of a loop based on dependence conditions which cannot be evaluated until runtime. For details, see *Optimizing Applications on Cray X1 Series Systems.*

This level is compatible with `-h vector2` and `3`.

3	Aggressive automatic multistreaming optimizations. Automatic multistreaming optimization is performed as with `stream2`. This level is compatible with `-h vector2` and `3`.

2.12 Vector Optimization Options

This section describes vector optimization options. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.12.1 `-h [no]infinitevl` (`CC`, `cc`, `c99`)

Default option: `-h infinitevl`

The `-h infinitevl` option tells the compiler to assume an infinite safe vector length for all `#pragma _CRI ivdep` directives. The `-h noinfinitevl` option tells the compiler to assume a safe vector length equal to the maximum supported vector length on the machine for all `#pragma _CRI ivdep` directives.

2.12.2 `-h [no]ivdep` (`CC`, `cc`, `c99`)

Default option: `-h noivdep`

The `-h ivdep` option tells the compiler to ignore vector dependencies for all loops. This is useful for vectorizing loops that contain pointers. With `-h noivdep`, loop dependencies inhibit vectorization. To control loops individually, use the `#pragma _CRI ivdep` directive, as discussed in Section 3.7.1, page 80.

This option can also be used with "vectorization-like" optimizations found in Section 3.7, page 80.

⚠️ **Caution:** This option should be used with extreme caution because incorrect results can occur if there is a vector dependency within a loop. Combining this option with inlining is dangerous because inlining can introduce vector dependencies.

This option severely constrains other loop optimizations and should be avoided if possible.

2.12.3 `-h vector`*n* (`CC`, `cc`, `c99`)

Default option: `-h vector2`

The `-h vector`*n* option specifies the level of automatic vectorizing to be performed. Vectorization results in dramatic performance improvements with a small increase in object code size. Vectorization directives are unaffected by this option.

Argument *n* can be one of the following:

n	Description
0	No automatic vectorization. Characteristics include low compile time and small compile size. This option is compatible with all scalar optimization levels.
1	Specifies conservative vectorization. Characteristics include moderate compile time and size. No loop nests are restructured; only inner loops are vectorized. Not all vector reductions are performed, so results do not differ from results obtained when the `-h vector0` option is specified. No vectorizations that might create false exceptions are performed.
	The `-h vector1` option is compatible with `-h scalar1`, `-h scalar2`, `-h scalar3`, or `-h stream1`.
2	Specifies moderate vectorization. Characteristics include moderate compile time and size. Loop nests are restructured. Results can differ slightly from results obtained when `-h vector1` is specified because of vector reductions.
	The `-h vector2` option is compatible with `-h scalar2` or `-h scalar3` and with `-h stream0`, `-h stream1`, and `-h stream2`.

| 3 | Specifies aggressive vectorization. Characteristics include potentially high compile time and size. Loop nests are restructured. Results can differ slightly from results obtained when `-h vector1` is specified because of vector reductions. Vectorizations that might create false exceptions in rare cases may be performed. |

Vectorization directives are described in Section 3.7, page 80.

2.12.4 `-h [no]vsearch` (CC, cc, c99)

Default option: `-h vsearch`

The `-h vsearch` option enables vectorization of all search loops. With `-h novsearch`, the default vectorization level applies. The `novsearch` directive is discussed in Section 3.7.5, page 83. This option is affected by the `-h vector`*n* option (see Section 2.12.3, page 31).

2.13 Inlining Optimization Options

Inlining is the process of replacing a user function call with the function definition itself. This saves call overhead and may allow better optimization of the inlined code. If all calls within a loop are inlined, the loop becomes a candidate for vectorization or streaming. Inlining may increase object code size.

Inlining is inhibited if:

- Arguments declared in a function differ in type from arguments in a function call.

- The number of arguments declared in a function differ from the number of arguments in a function call.

- A call site is within the range of a `#pragma inline_disable` directive. For a description of the `inline_disable` directive, see Section 3.10.1, page 97.

- A function being called is specified on a `#pragma inline_never` directive. For a description of the `inline_never` directive, see Section 3.10.3, page 99.

- The compiler determines that the routine is too big to inline. This is determined by an internal limit of the text (that is, the instruction segment in the executable) size of the routine. You can override this limit by inserting a `#pragma inline_always` directive. For a description of the `inline_always` directive, see Section 3.10.3, page 99.

The compiler supports the following inlining modes:

- Automatic inlining (Section 2.13.1).

- Explicit inlining (see Section 2.13.2, page 34).

- Combined inlining (see Section 2.13.3, page 35)

2.13.1 `-h ipan` (`CC`, `cc`, `c89`, `c99`)

Default option: `-h ipa3`

The `-h ipan` option specifies automatic inlining. Automatic inlining allows
the compiler to automatically select, depending on the inlining level *n*, which
functions to inline. Each *n* is a different set of heuristics. The candidates for
expansion are all those functions that are present in the input file to the compile
step. Table 5 explains what is inlined at each level.

Table 5. Automatic Inlining Specifications

Inlining level	Description
0	All inlining is disabled. All inlining compiler directives are ignored. See Section 3.10, page 96 for more information about inlining directives.
1	Directive inlining. Inlining is attempted for call sites and routines that are under the control of an inlining pragma directive.
2	Loop inlining. Inlining is attempted at level 1 plus inlining is attempted for call sites that exist within `for`, `while`, and `do-while` loops.
3	Inlining to a depth of three. Inlining is attempted on call sites containing scalar constant arguments. The inlining must result in "flat-lined" code (i.e., contain no calls); otherwise, the inlining will not be performed.
	For example, assume `main` calls function `p1`, `p1` calls `p2`, `p2` calls `p3`, `p3` calls `p4`, and `p4` calls leaf routine `p5`. Using `-h ipa3` would result in the inlining of functions `p3`, `p4`, and `p5`.
	Leaf node functions do not call any external functions. An external function is any routine that must be branched to.
	`-h ipa3` is the default inlining level.

Inlining level	Description
4	All heuristics inlining. This level is the combination of the heuristics at levels 1, 2, and 3.
5	Aggressive inlining. Inlining is attempted for every call site encountered.

2.13.2 `-h ipafrom=`*source* [`:`*source*] (`CC`, `cc`, `c89`, `c99`)

The `-h ipafrom=`*source* [`:`*source*] option specifies explicit inlining. The *source* arguments identify the files or directories that contain the functions to consider for inlining. Only those functions present in *source* are candidates for inlining. When a call is encountered to a function that resides in *source*, an attempt will be made to expand the function in place at that call site.

Note that blanks are not allowed on either side of the equal sign.

All inlining directives are recognized with explicit inlining. For information on inlining directives, see Section 3.10, page 96.

The functions in *source* are not actually loaded with the final program. They are simply templates for the inliner. To have a function contained in *source* loaded with the program, you must include it in an input file to the compilation.

Use one or more of the following objects in the *source* argument.

Source	Description
C or C++ source files	The functions in C or C++ source files are candidates for inline expansion and must contain error-free code.
	C files that are acceptable for inlining are files with the `.c` extension.
	C++ files that are acceptable for inlining are files that have one of the following extensions: `.C`, `.c++`, `.C++`, `.cc`, `.cxx`, `.Cxx`, `.CXX`, `.CC,`, or `.cpp`.
dir	A directory that contains any of the file types described in this table.

2.13.3 Combined Inlining

Combined inlining is a combination of automatic inlining and explicit inlining. It allows you to specify targets for inline expansion, while applying the selected level of inlining heuristics.

You invoke combined inlining by including both the `-h ipa`*n* and the `-h ipafrom=`*source* [`:`*source*] options on the command line. The only candidates for expansion are those functions that reside in *source*. The rules that apply to deciding whether to inline are defined by the `-h ipa`*n* setting.

2.14 Scalar Optimization Options

This section describes scalar optimization options. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.14.1 `-h [no]interchange` (`CC`, `cc`, `c99`)

Default option: `-h interchange`

The `-h interchange` option allows the compiler to attempt to interchange all loops, a technique that is used to gain performance by having the compiler swap an inner loop with an outer loop. The compiler attempts the interchange only if the interchange will increase performance. Loop interchange is performed only at scalar optimization level 2 or higher.

The `-h nointerchange` option prevents the compiler from attempting to interchange any loops. To disable interchange of loops individually, use the `#pragma _CRI nointerchange` directive.

2.14.2 `-h scalar`*n* (`CC`, `cc`, `c99`)

Default option: `-h scalar1`

The `-h scalar`*n* option specifies the level of automatic scalar optimization to be performed. Scalar optimization directives are unaffected by this option (see Section 3.9, page 91).

Use one of these values for *n*:

n	Description
0	No automatic scalar optimization. The `-h matherror=errno` and `-h zeroinc` options are implied by `-h scalar0`.

1	Conservative automatic scalar optimization. This level implies `-h matherror=abort` and `-h nozeroinc`.
2	Moderate automatic scalar optimization. The scalar optimizations specified by `scalar1` are performed.
3	Aggressive automatic scalar optimization.

2.14.3 `-h [no]reduction` (`CC`, `cc`, `c99`)

Default option: `-h reduction`

The `-h reduction` option tells the compiler to enable vectorization of all reduction loops. The `-h noreduction` option disables vectorization of any loop that contains a reduction in which the order of evaluation is numerically significant. The specific reductions that are disabled are floating point, double precision, complex summation and product reductions, and alternating value computations.

This option is affected by the `-h scalar`*n* option (see Section 2.14.2, page 35). Reduction loops and the `noreduction` directive are discussed in Section 3.9.3, page 93.

2.14.4 `-h [no]zeroinc` (`CC`, `cc`, `c99`)

Default option: `-h nozeroinc`

The `-h nozeroinc` option improves run time performance by causing the compiler to assume that constant increment variables (CIVs) in loops are not incremented by expressions with a value of 0.

The `-h zeroinc` option causes the compiler to assume that some CIVs in loops might be incremented by 0 for each pass through the loop, preventing generation of optimized code. For example, in a loop with index i, the expression *expr* in the statement i += *expr* can evaluate to 0. This rarely happens in actual code. `-h zeroinc` is the safer and slower option. This option is affected by the `-h scalar`*n* option (see Section 2.14.2, page 35).

2.15 Math Options

This section describes compiler options pertaining to math functions. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.15.1 -h fp*n* (CC, cc, c99)

Default option: -h fp2

The -h fp option allows you to control the level of floating-point optimizations. The *n* argument controls the level of allowable optimization; 0 gives the compiler minimum freedom to optimize floating-point operations, while 3 gives it maximum freedom. The higher the level, the lesser the floating-point operations conform to the IEEE standard.

This option is useful for code that use unstable algorithms, but which are optimizable. It is also useful for applications that want aggressive floating-point optimizations that go beyond what the Fortran standard allows.

Generally, this is the behavior and usage for each -h fp level:

- -h fp0—causes your program's executable code to conform more closely to the IEEE floating-point standard than the default mode (-h fp2). When this level is specified, many identity optimizations are disabled, executable code is slower than higher floating-point optimization levels, and a scaled complex divide mechanism is enabled that increases the range of complex values that can be handled without producing an underflow.

  The -h fp0 option should never be used, except when your code pushes the limits of IEEE accuracy, or require strong IEEE standard conformance.

- -h fp1—performs various, generally safe, non-conforming IEEE optimizations, such as folding a == a to true, where a is a floating point object. At this level, floating-point reassociation[1] is greatly limited, which may affect the performance of your code.

  The -h fp1 options should never be used, except when your code pushes the limits of IEEE accuracy, or requires strong IEEE standard conformance.

- -h fp2—includes optimizations of -h fp1.

- -h fp3—includes optimizations of .

  The -h fp3 option should be used when performance is more critical than the level of IEEE standard conformance provided by -h fp2.

Table 6 compares the various optimization levels of the -h fp option (levels 2 and 3 are usually the same). The table lists some of the optimizations performed; the compiler may perform other optimizations not listed.

---

[1]   For example, a+b+c is rearranged to b+a+c, where a, b, and c are floating point variables.

Table 6. Floating-point Optimization Levels

Optimization Type	0	1	2 (default)	3
Inline selected mathematical library functions	N/A	N/A	N/A	Accuracy is slightly reduced.
Complex divisions	Accurate and slower	Accurate and slower	Less accurate (less precision) and faster.	Less accurate (less precision) and faster.
Exponentiation rewrite	None	None	Maximum performance[2]	Maximum performance[2, 3]
Strength reduction	Fast	Fast	Aggressive	Aggressive
Rewrite division as reciprocal equivalent [4]	None	None	Yes	Yes
Safety	Maximum	Moderate	Moderate	Low

If multiple -h fp options are used, the compiler will use only the rightmost option and will issue a message indicating such.

2.15.2 -h matherror=*method* (CC, cc, c99)

Default option: -h matherror=abort

The -h matherror=*method* option specifies the method of error processing used if a standard math function encounters an error. The *method* argument can have one of the following values:

method	Description
abort	If an error is detected, errno is not set. Instead a message is issued and the program aborts. An exception may be raised.

---

[2] Rewriting values raised to a constant power into an algebraically equivalent series of multiplications and/or square roots.

[3] Rewriting exponentiations (a^b) not previously optimized into the algebraically equivalent form exp(b * ln(a)).

[4] For example, x/y is transformed to x * 1.0/y.

errno   If an error is detected, `errno` is set and the math function returns to the caller. This method is implied by the `-h conform`, `-h scalar0`, `-O0`, `-Gn`, and `-g` options.

2.16 Debugging Options

This section describes compiler options used for debugging. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.16.1 `-G` *level* (`CC`, `cc`, `c99`) and `-g` (`CC`, `cc`, `c89`, `c99`)

The `-G` *level* and `-g` options enable the generation of debugging information that is used by symbolic debuggers such as TotalView. These options allow debugging with breakpoints. Table 7 describes the values for the `-G` option.

Table 7. `-G` *level* Definitions

level	Optimization	Breakpoints allowed on
f	Full	Function entry and exit
p	Partial	Block boundaries
n	None	Every executable statement

More extensive debugging (such as full) permits greater optimization opportunities for the compiler. Debugging at any level may inhibit some optimization techniques, such as inlining.

The `-g` option is equivalent to `-Gn`. The `-g` option is included for compatibility with earlier versions of the compiler and many other UNIX systems; the `-G` option is the preferred specification. The `-Gn` and `-g` options disable all optimizations and imply `-O0`.

The debugging options take precedence over any conflicting options that appear on the command line. If more than one debugging option appears, the last one specified overrides the others.

Debugging is described in more detail in Chapter 12, page 163.

2.16.2 -h [no]bounds (cc, c99)

Default option: -h nobounds

The -h bounds option provides checking of pointer and array references to ensure that they are within acceptable boundaries. -h nobounds disables these checks.

The pointer check verifies that the pointer is greater than 0 and less than the machine memory limit. The array check verifies that the subscript is greater than or equal to 0 and is less than the array size, if declared.

2.16.3 -h zero (CC, cc, c99)

The -h zero option causes stack-allocated memory to be initialized to all zeros.

2.17 Compiler Message Options

This section describes compiler options that affect messages. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.17.1 -h msglevel_*n* (CC, cc, c99)

Default option: -h msglevel_3

The -h msglevel_*n* option specifies the lowest level of severity of messages to be issued. Messages at the specified level and above are issued. Argument *n* can be 0 (comment), 1 (note), 2 (caution), 3 (warning), or 4 (error).

2.17.2 -h [no]message=*n*[:*n*...] (CC, cc, c99)

Default option: Determined by -h msglevel_*n*

The -h [no]message=*n*[:*n*...] option enables or disables specified compiler messages. *n* is the number of a message to be enabled or disabled. You can specify more than one message number; multiple numbers must be separated by a colon with no intervening spaces. For example, to disable messages CC-174 and CC-9, specify:

-h nomessage=174:9

The -h [no]message=*n* option overrides -h msglevel_*n* for the specified messages. If *n* is not a valid message number, it is ignored. Any compiler

message except ERROR, INTERNAL, and LIMIT messages can be disabled; attempts to disable these messages by using the -h nomessage=*n* option are ignored.

2.17.3 -h report=*args* (CC, cc, c99)

The -h report=*args* option generates report messages specified in *args* and lets you direct the specified messages to a file. Use any combination of these for *args*:

args	Description
i	Generates inlining optimization messages
m	Generates multistream optimization messages
s	Generates scalar optimization messages
v	Generates vector optimization messages
f	Writes specified messages to file *file*.V where *file* is the source file specified on the command line. If the f option is not specified, messages are written to stderr.

No spaces are allowed around the equal sign (=) or any of the *args* codes. For example, the following example prints inlining and scalar optimization messages for myfile.c:

```
% cc -h report=is myfile.c
```

2.17.4 -h [no]abort (CC, cc, c99)

Default option: -h noabort

The -h [no]abort option controls whether a compilation aborts if an error is detected.

2.17.5 -h errorlimit[=*n*] (CC, cc, c99)

Default option: -h errorlimit=100

The -h errorlimit[=*n*] option specifies the maximum number of error messages the compiler prints before it exits. *n* is a positive integer. Specifying -h errorlimit=0 disables exiting on the basis of the number of errors. Specifying -h errorlimit with no qualifier is the same as setting *n* to 1.

2.18 Compilation Phase Options

This section describes compiler options that affect compilation phases. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.18.1 `-E` (`CC`, `cc`, `c89`, `c99`, `cpp`)

If the `-E` option is specified on the command line (except for `cpp`), it executes only the preprocessor phase of the compiler. The `-E` and `-P` options are equivalent, except that `-E` directs output to `stdout` and inserts appropriate `#line` preprocessing directives. The `-E` option takes precedence over the `-h feonly`, `-S`, and `-c` options.

If the `-E` option is specified on the `cpp` command line, it inserts the appropriate `#line` directives in the preprocessed output. When both the `-P` and `-E` options are specified, the last one specified takes precedence.

2.18.2 `-P` (`CC`, `cc`, `c99`, `cpp`)

When the `-P` option is specified on the command line (except for `cpp`), it executes only the preprocessor phase of the compiler for each source file specified. The preprocessed output for each source file is written to a file with a name that corresponds to the name of the source file and has `.i` suffix substituted for the suffix of the source file. The `-P` option is similar to the `-E` option, except that `#line` directives are suppressed, and the preprocessed source does not go to `stdout`. This option takes precedence over `-h feonly`, `-S`, and `-c`.

When both the `-P` and `-E` options are specified, the last one specified takes precedence.

When the `-P` option is specified on the `cpp` command line, it is ignored.

2.18.3 `-h feonly` (`CC`, `cc`, `c99`)

The `-h feonly` option limits the Cray C and C++ compilers to syntax checking. The optimizer and code generator are not executed. This option takes precedence over `-S` and `-c`.

2.18.4 `-S` (CC, cc, c99)

The `-S` option compiles the named C or C++ source files and leaves their assembly language output in the corresponding files suffixed with a `.s`. If this option is used with `-G` or `-g`, debugging information is not generated. This option takes precedence over `-c`.

2.18.5 `-c` (CC, cc, c89, c99)

The `-c` option creates a relocatable object file for each named source file but does not link the object files. The relocatable object file name corresponds to the name of the source file. The `.o` suffix is substituted for the suffix of the source file.

2.18.6 `-#`, `-##`, and `-###` (CC, cc, c99, cpp)

The `-#` option produces output indicating each phase of the compilation as it is executed. Each succeeding output line overwrites the previous line.

The `-##` option produces output indicating each phase of the compilation as it is executed.

The `-###` option is the same as `-##`, except the compilation phases are not executed.

2.18.7 `-Wphase,"opt..."` (CC, cc, c99)

The `-Wphase` option passes arguments directly to a phase of the compiling system. Table 8 shows the system phases that *phase* can indicate.

Table 8. `-Wphase` Definitions

phase	System phase	Command
p	Preprocessor	cpp
0	Compiler	CC, cc, and c99
a	Assembler	as(1)
l	Loader	ld

Arguments to be passed to system phases can be entered in either of two styles. If spaces appear within a string to be passed, the string is enclosed in double quotes. When double quotes are not used, spaces cannot appear in the string.

Commas can appear wherever spaces normally appear; an option and its argument can be either separated by a comma or not separated. If a comma is part of an argument, it must be preceded by the \ character. For example, any of the following command lines would send -e name and -s to the loader:

```
% cc -Wl,"-e name -s" file.c

% cc -Wl,-e,name,-s file.c

% cc -Wl,"-ename",-s file.c
```

Because the preprocessor is built into the compiler, -Wp and -W0 are equivalent.

2.18.8 -Y*phase*,*dirname* (CC, cc, c89, c99, cpp)

The -Y*phase*,*dirname* option specifies a new directory (*dirname*) from which the designated *phase* should be executed. *phase* can be one or more of the values shown in Table 9.

Table 9. -Y*phase* Definitions

phase	System phase	Command
p	Preprocessor	cpp
0	Compiler	CC,cc,c89,c89,cpp
a	Assembler	as
l	Loader	ld

Because there is no separate preprocessor, -Yp and -Y0 are equivalent. If you are using the -Y option on the cpp command line, p is the only argument for *phase* that is allowed.

2.19 Preprocessing Options

This section describes compiler options that affect preprocessing. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.19.1 −C (CC, cc, c99, cpp)

The −C option retains all comments in the preprocessed source code, except those on preprocessor directive lines. By default, the preprocessor phase strips comments from the source code. This option is useful with cpp or in combination with the −P or −E option on the CC, cc, and c99 commands.

2.19.2 −D *macro*[=*def*] (CC, cc, c89, c99 cpp)

The −D *macro*[=*def*] option defines a macro named *macro* as if it were defined by a #define directive. If no =*def* argument is specified, *macro* is defined as 1.

Predefined macros also exist; these are described in Chapter 10, page 157. Any predefined macro except those required by the standard (see Section 10.1, page 157) can be redefined by the −D option. The −U option overrides the −D option when the same macro name is specified regardless of the order of options on the command line.

2.19.3 −h [no]pragma=*name*[:*name...*] (CC, cc, c99)

Default option: −h pragma

The [no]pragma=*name*[:*name...*] option enables or disables the processing of specified directives in the source code. *name* can be the name of a directive or a word shown in Table 10 to specify a group of directives. More than one name can be specified. Multiple names must be separated by a colon and have no intervening spaces.

Table 10. −h pragma Directive Processing

name	Group	Directives affected
all	All	All directives
allinline	Inlining	inline, nolinline
allscalar	Scalar optimization	concurrent, nointerchange, noreduction, suppress, unroll
allvector	Vectorization	ivdep, novector, novsearch, prefervector, shortloop

When using this option to enable or disable individual directives, note that some directives must occur in pairs. For these directives, you must disable both directives if you want to disable either; otherwise, the disabling of one of the directives may cause errors when the other directive is (or is not) present in the compilation unit.

2.19.4 −I *incldir* (CC, cc, c89, c99, cpp)

The −I *incldir* option specifies a directory for files named in #include directives when the #include file names do not have a specified path. Each directory specified must be specified by a separate −I option.

The order in which directories are searched for files named on #include directives is determined by enclosing the file name in either quotation marks ("") or angle brackets (< and >).

Directories for #include "*file*" are searched in the following order:

1. Directory of the input file.

2. Directories named in −I options, in command line order.

3. Site-specific and compiler release-specific include files directories.

4. Directory /usr/include.

Directories for #include <*file*> are searched in the following order:

1. Directories named in −I options, in command line order.

2. Site-specific and compiler release-specific include files directories.

3. Directory /usr/include.

If the −I option specifies a directory name that does not begin with a slash (/), the directory is interpreted as relative to the current working directory and not relative to the directory of the input file (if different from the current working directory). For example:

```
% cc -I. -I yourdir mydir/b.c
```

The preceding command line produces the following search order:

1. mydir (#include "*file*" only).

2. Current working directory, specified by −I.

3. `yourdir` (relative to the current working directory), specified by `-I yourdir`.

4. Site-specific and compiler release-specific include files directories.

5. Directory `/usr/include`.

2.19.5 `-M` (`CC`, `cc`, `c99`, `cpp`)

The `-M` option provides information about recompilation dependencies that the source file invokes on `#include` files and other source files. This information is printed in the form expected by `make`. Such dependencies are introduced by the `#include` directive. The output is directed to `stdout`.

2.19.6 `-N` (`cpp`)

The `-N` option specified on the `cpp` command line enables the old style (referred to as K & R) preprocessing. If you have problems with preprocessing (especially non-C source code), use this option.

2.19.7 `-nostdinc` (`CC`, `cc`, `c89`, `c99`, `cpp`)

The `-nostdinc` option stops the preprocessor from searching for include files in the standard directories (`/usr/include/CC` and `/usr/include`).

2.19.8 `-U` *macro* (`CC`, `cc`, `c89`, `c99`, `cpp`)

The `-U` option removes any initial definition of *macro*. Any predefined macro except those required by the standard (see Section 10.1, page 157) can be undefined by the `-U` option. The `-U` option overrides the `-D` option when the same macro name is specified, regardless of the order of options on the command line.

Predefined macros are described in Chapter 10, page 157. Macros defined in the system headers are not predefined macros and are not affected by the `-U` option.

2.20 Loader Options

This section describes compiler options that affect loader tasks. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.20.1 -l *libfile* **(CC, cc, c89, c99)**

The -l *libfile* option identifies a library file. To request more than one library file, specify multiple -l options.

The loader searches for libraries by prepending ldir/lib on the front of *libfile* and appending .a on the end of it, for each ldir that has been specified by using the -L option. It uses the first file it finds. See also the -L option (Section 2.20.2, page 48).

There is no search order dependency for libraries. Default libraries are shown in the following list:

libC.a (Cray C++ only)

libu.a

libm.a

libc.a

libsma.a

libf.a

libfi.a

libsci.a

If you specify personal libraries by using the -l command line option, as in the following example, those libraries are added to the top of the preceding list. (The -l option is passed to the loader.)

```
cc -l mylib target.c
```

When the previous command line is issued, the loader looks for a library named libmylib.a (following the naming convention) and adds it to the top of the list of default libraries.

2.20.2 -L *libdir* **(CC, cc, c89, c99)**

The -L *libdir* option changes the -l option search algorithm to look for library files in directory ldir. To request more than one library directory, specify multiple -L options.

The loader searches for library files in the compiler release-specific directories.

Note: Multiple −L options are treated cumulatively as if all *libdir* arguments appeared on one −L option preceding all −l options. Therefore, do not attempt to load functions of the same name from different libraries through the use of alternating −L and −l options.

2.20.3 −o *outfile* (CC, cc, c89, c99)

The −o *outfile* option produces an absolute binary file named *outfile*. A file named a.out is produced by default. When this option is used in conjunction with the −c option and a single C or C++ source file, a relocatable object file named *outfile* is produced.

2.20.4 −s (CC, cc, c89, c99)

(Deferred implementation) The −s option produces executable files from which symbolic and other information not required for proper execution has been removed. If both the −s and −g (or −G) options are present, −s is ignored.

2.21 Miscellaneous Options

This section describes compiler options that affect general tasks. Each subsection heading shows in parentheses the compiler command with which the option can be used.

2.21.1 −h command (cc, c99)

The command mode option (−h command) allows you to create commands for Cray X1 series systems to supplement commands developed by Cray. Such commands run serially on a single-streaming processor (SSP) within a support node; they execute immediately without assistance from aprun or psched.

The commands created with the command mode option cannot multistream. If you want to disable vectorization, add the −h vector0 option to the compiler command line. The compiled commands will have less debugging information, unless you specify a debugging option. The debugging information does not slow execution time, but it does result in a larger executable that may take longer to load.

For simplicity, you should use the C compiler to load your programs built with the command mode option, because the required options and libraries are automatically specified and loaded for you.

If you decide to load the libraries manually, you must use the loader command (`ld`) and specify on its command line the `-command` and `-ssp` options and the `-L` option with the path to the command mode libraries. The command mode libraries are found in the `cmdlibs` directory under the path defined by the `CRAYLIBS_SV2` environment variable. These must also be linked:

- `Start0.o`
- `libc` library
- `libm` library
- `libu` library

The following sample command line illustrates compiling the code for a command named `fierce`:

```
% cc -h command -h vector0 -o fierce fierce.c
```

Note: The `-h ssp` and `-h command` options both create executables that run on an SSP. The executable created via the `-h ssp` option runs on an application node. The executable created via the `-h command` option runs on the support node.

2.21.2 `-h cpu=`*target_system* (`CC`, `cc`, `c99`)

The `-h cpu=`*target_system* option specifies the Cray X1 series system on which the absolute binary file is to be executed.

Default   `-h cpu=cray-x1`

Use one of these values for *target_system*:

target_system	Description
`cray-x1`	Use this option (default) if the absolute binary file will be executed on a Cray X1 system
`cray-x1e`	Use this option if the absolute binary file will be executed on a Cray X1E system

Note: Currently, there are no differences in the code produced for the `cray-x1` and `cray-x1e` targets. This option was created to allow us to support future changes in optimization or code generation based on our experience with the Cray X1E hardware. It is possible that compilations with the `-hcpu=cray-x1e` option will not be compatible with Cray X1 machines in the future.

2.21.3 -h decomp (CC, cc, c99)

The -h decomp option decompiles (translates) the intermediate representation of the compiler into listings that resemble the format of the source code. This is performed twice, resulting in two output files, at different points during the optimization process. You can use these files to examine the restructuring and optimization changes made by the compiler, which can lead to insights about changes you can make to your C or C++ source to improve its performance. See *Optimizing Applications on Cray X1 Series Systems* for further information.

The compiler produces two decompilation listing files, with these extensions, per source file specified on the command line: .opt and .cg. The compiler generates the .opt file after applying most high level loop nest transformations to the code. The code structure of this listing most resembles your source code and is readable by most users. In some cases, because of optimizations, the structure of the loops and conditionals will be significantly different than the structure in your source file.

The .cg file contains a much lower level of decompilation. It is still displayed in a C or C++ like format, but is quite close to what will be produced as assembly output. This version displays the intermediate text after all multistreaming translation, vector translation, and other optimizations have been performed. An intimate knowledge of the hardware architecture of the system is helpful to understanding this listing.

The .opt and .cg files are intended as a tool for performance analysis and are not valid C or C++ functions. The format and contents of the files can be expected to change from release to release.

The following examples show the listings generated when the -h decomp is applied to this example:

```
/* Source code, in file example.c */
void
example( double a[restrict], double b[restrict],
             double c[restrict], const int n )
{
    int i;
    for ( i = 0; i < n; i++ ) {
        a[i] = b[i] * c[i];
    }
}
```

This is the listing of the `example.opt` file after loop optimizations are performed:

```
2. void
2. example( a, b, c, n )
2. {
6.     if ( 0 < n ) {
6.         @Induc01_N2 = 0;
6. #pragma ivdep
6. #pragma stream
6.         do {
7.             a[@Induc01_N2] = c[@Induc01_N2] * b[@Induc01_N2];
6.             @Induc01_N2 = 1 + @Induc01_N2;
6.         } while ( @Induc01_N2 < n );
6.     }
9.     return;
9. }
```

This is the listing of the `example.cg` file after other optimizations are performed:

```
2. void
2. example( a, b, c, n )
2. {
2.     _park( 1 );
2.     $MR_a_0 = a;
2.     $MR_c_1 = c;
2.     $MR_b_2 = b;
2.     r_suppress ( $SC_n_I3, $ssp_0, $strind_1, $MR_c_1, $MR_b_2 );
2.     d_suppress ( $ssp_0, $strind_1, $MR_a_0 );
2.     @SB_a_8 = $MR_a_0;
2.     @SB_c_9 = $MR_c_1;
2.     @SB_b_10 = $MR_b_2;
2.     _unpark();
2.     $MR_b_2 = @SB_b_10;
2.     $MR_c_1 = @SB_c_9;
2.     $MR_a_0 = @SB_a_8;
6.     $MR_n_3 = n;
6.     if ( 0 < $MR_n_3 ) {
6.         $ssp_0 = _mype();
6.         $InvSSP_5 = 3 & ~$ssp_0;
6.         $start_2 = $InvSSP_5 * $MR_n_3 >> 2;
6.         $psize_4 = ( ( 1 + $InvSSP_5 ) * $MR_n_3 >> 2 ) -
6.                     $start_2;
```

```
6.        if ( $psize_4 > 0 ) {
6.            $VL_1 = _cvl( $psize_4 );
6.            $LC_2 = $psize_4;
6.            $SI_4 = 0;
7.            $LCS_0 = $start_2 << 3;
7.            @LIS_E2 = $LCS_0 + (long) $MR_b_2;
7.            @LIS_E1 = $LCS_0 + (long) $MR_c_1;
7.            @LIS_E0 = $LCS_0 + (long) $MR_a_0;
7.            do {
7.                $LCS_1 = $SI_4 << 3;
7.                0[@LIS_E0 + $LCS_1:$VL_1:1].L = 0[@LIS_E1 +
7.                    $LCS_1:$VL_1:1].L * 0[@LIS_E2 +
7.                    $LCS_1:$VL_1:1].L;
6.                $SI_4 = $VL_1 + $SI_4;
6.                $LC_2 = $LC_2 - $VL_1;
6.                $VL_1 = _cvl( $LC_2 );
6.            } while ( $LC_2 > 0 );
6.        }
6.    }
6.    r_suppress ( $SC_n_I3, $ssp_0, $strind_1, $MR_c_1, $MR_b_2 );
6.    d_suppress ( $ssp_0, $strind_1, $MR_a_0 );
9.    return;
9. }
```

2.21.4 -h ident=*name* (CC, cc, c99)

Default option: File name specified on the command line

The -h ident=*name* option changes the ident name to *name*. This *name* is used as the module name in the object file (.o suffix) and assembler file (.s suffix). Regardless of whether the ident name is specified or the default name is used, the following transformations are performed on the ident name:

- All . characters in the ident name are changed to $.

- If the ident name starts with a number, a $ is added to the beginning of the ident name.

2.21.5 `-h keepfiles` (`CC`, `cc`, `c89`, `c99`)

The `-h keepfiles` option prevents the removal of the object (`.o`) files after an executable is created. Normally, the compiler automatically removes these files after linking them to create an executable. Since the original object files are required in order to instrument a program for performance analysis, if you plan to use CrayPat to conduct performance analysis experiments, you can use this option to preserve the object files.

2.21.6 `-h [no]mpmd` (`CC`, `cc`)

Default: `-h nompmd`

Used when compiling co-array Fortran (CAF) programs for multiple program, multiple data (MPMD) launch. For details, see Section 11.3, page 162.

2.21.7 `-h [no]omp` (`CC`, `cc`)

Default: `-h noomp`

Enables or disables the C or C++ compiler recognition of OpenMP directives. For details, see Chapter 5, page 113.

2.21.8 `-h prototype_intrinsics` (`CC`, `cc`, `c99`, `cpp`)

Simulates the effect of including `intrinsics.h` at the beginning of a compilation. Use this option if the source code does not include the `intrinsics.h` statement and you cannot modify the code. This option is off by default. See Appendix F, page 225 for details.

2.21.9 `-h task`*n* (`CC`, `cc`)

This option enables tasking in C or C++ applications that contain OpenMP directives.

Default: `-h task0`

n	Description
0	Disables tasking. OpenMP directives are ignored. Using this option can reduce compile time and the size of the executable. The `-h task0` option is compatible with all vectorization and scalar optimization levels.

1   The -h task1 option specifies user tasking, so OpenMP directives are recognized. No level for scalar optimization is enabled automatically. The -h task1 option is compatible with all vectorization and scalar optimization levels.

2.21.10 -h [no]threadsafe (CC)

Default: -h threadsafe

This option enables or disables the generation of threadsafe code. Code that is threadsafe can be used with pthreads and OpenMP. This option is not binary-compatible with code generated by Cray C++ 5.1 and earlier compilers. Users who need binary compatibility with previously compiled code can use -h nothreadsafe, which causes the compiler to be compatible with Cray C++ 5.1 and earlier compilers at the expense of not being threadsafe.

C++ code compiled with -h threadsafe (the default) cannot be linked with C++ code compiled with -h nothreadsafe or with code compiled with a Cray C++ 5.1 or earlier compiler.

2.21.11 -h upc (cc)

The -h upc option enables compilation of Unified Parallel C (UPC) code. UPC is a C language extension for parallel program development that allows you to explicitly specify parallel programming through language syntax rather than through library functions such as are used in MPI or SHMEM.

The Cray X1 series implementation of UPC is discussed in greater detail in Chapter 6, page 119.

2.21.12 -V (CC, cc, c99, cpp)

The -V option displays compiler version information. If the command line specifies no source file, no compilation occurs. Version information consists of the product name, the version number, and the current date and time, as shown in the following example:

```
% CC -V
Cray Standard C: Version 5.3.0.0.35  Thu Oct 21, 2004  14:06:15
```

2.21.13 -X *npes* **(CC, cc, c99)**

The `-X` *npes* option specifies the number of processing elements to use during execution. The value for *npes* ranges from 1 through 4096 inclusive.

Once set, the number of processing elements to use cannot be changed at load or run time. You must recompile the program with a different value for *npes* to change the number of processing elements.

If you use the `ld` command to manually load a program compiled with the `-X` option, you must specify the same value to the loader as was specified at compile time.

You can execute the compiled program without using the `aprun` command just by entering the name of the output file. If you use the command and specify the number of processing elements on the `aprun` command line, you must specify the same number on the `aprun` command as was specified at compile time.

The `_num_pes` intrinsic function can be used when programming UNICOS/mp systems. The value returned by `_num_pes` is equal to the number processing elements available to your program. The number of the first processing element is always 0, and the number of the last processing element is `_num_pes() - 1`. When the `-X` *npes* option is specified at compile time, the `_num_pes` intrinsic function returns the value specified by the *npes* argument.

On the Cray X1 series system, the `_num_pes` intrinsic can be used only in either of these situations:

- When the `-X` *npes* option is specified on the command line

- When the value of the expression containing the `_num_pes` intrinsic function is not known until run time (that is, it can only be used in run time expressions)

One of the many uses for the `_num_pes` intrinsic is illustrated in the following example, which declares a variable length array of size equal to the number of processing elements:

```
int a[_num_pes()];
```

Using the `_num_pes` intrinsic in conjunction with the `-X` *npes* option allows the user to program the number of processing elements into code in places that do not accept run time values. Specifying the number of processing elements at compile time can also enhance compiler optimization.

2.22 Command Line Examples

These examples illustrate a variety of command lines for the C and C++ compiler commands:

- This example compiles `myprog.C`, fixes the number of processing elements to 8, and instantiates all template entities declared or referenced in the compilation unit. Because the program is compiled in default MSP mode, each processing element is an MSP.

  ```
  % CC -X8 -h instantiate=all myprog.C
  ```

- This example compiles `myprog.C`. The `-h conform` option specifies strict conformance to the ISO C++ standard. No automatic instantiation of templates is performed.

  ```
  % CC -h conform -h noautoinstantiate myprog.C
  ```

- This example compiles input files `myprog.C` and `subprog.C`. The `-c` option tells the compiler to create object files `myprog.o` and `subprog.o` but not call the loader. Option `-h ipa1` tells the compiler to inline function calls marked with the `inline_always` pragma.

  ```
  % CC -c -h ipa1 myprog.C subprog.C
  ```

- This example specifies that the compiler search the current working directory, represented by a period (`.`), for `#include` files before searching the default `#include` file locations.

  ```
  % CC -I. disc.C vend.C
  ```

- This example specifies that source file `newprog.c` be preprocessed only. Compilation and linking are suppressed. In addition, the macro `DEBUG` is defined.

  ```
  % cc -P -D DEBUG newprog.c
  ```

- This example compiles `mydata1.C`, creates object file `mydata1.o`, and produces a scalar optimization report to `stdout`.

  ```
  % CC -c -h report=s mydata1.C
  ```

- This example compiles `mydata3.c` and produces the executable file `a.out`. A 132-column pseudo assembly listing file is also produced in file `mydata3.L`.

  ```
  % cc -h listing mydata3.c
  ```

- This example compiles myfile.C and tells the compiler to attempt to aggressively inline calls to functions defined within myfile.C. An inlining report is directed to myfile.V.

```
% CC -h ipa5,report=if myfile.C
```

2.23 Compile Time Environment Variables

These environment variables are used during compilation.

Variable	Description
CRAYOLDCPPLIB	Enables, when set to a nonzero value, C++ code to use these nonstandard Cray C++ headers files:

- common.h

- complex.h

- fstream.h

- generic.h

- iomanip.h

- iostream.h

- stdiostream.h

- stream.h

- strstream.h

- vector.h

If you want to use the standard header files, your code may require modification to compile successfully. Refer to Appendix C, page 195.

Note: Setting the CRAYOLDCPPLIB environment variable disables exception handling, unless you compile with the -h exceptions option.

CRI_CC_OPTIONS, CRI_cc_OPTIONS, CRI_c89_OPTIONS, CRI_cpp_OPTIONS	Specifies command line options that are applied to all compilations. Options specified by this environment variable are added following the options specified directly on the command line. This is especially useful for adding options to compilations done with build tools.
LANG	Identifies your requirements for native language, local customs, and coded character set with regard to compiler messages.
MSG_FORMAT	Controls the format in which you receive compiler messages.
NLSPATH	Specifies the message system catalogs that should be used.
NPROC	Specifies the number of processes used for simultaneous compilations. The default is 1. When more than one source file is specified on the command line, compilations may be multiprocessed by setting the environment variable NPROC to a value greater than 1. You can set NPROC to any value; however, large values can overload the system.

2.24 Run Time Environment Variables

These environment variables are used during run time.

Variable	Description
CRAY_AUTO_APRUN_OPTIONS	

The CRAY_AUTO_APRUN_OPTIONS environment variable specifies options for the aprun command when the command is called automatically (auto aprun). Calling the aprun command automatically occurs when only the name of the program and, where applicable, associated program options are entered on the command line; this will cause the system to automatically call aprun to run the program.

The CRAY_AUTO_APRUN_OPTIONS environment variable does not specify options for the aprun command when you explicitly specify the command on the command line, nor does it specify options for your program.

When setting options for the `aprun` command in the `CRAY_AUTO_APRUN_OPTIONS` environment variable, surround the options within double quotes and separate each option with a space. Do not use spaces between an option and its associated value. For example,

```
setenv CRAY_AUTO_APRUN_OPTIONS "-n10 -m16G"
```

If you execute a program compiled with a fixed number of processing elements (that is, the `-X` compiler option was specified at compile time) and the `CRAY_AUTO_APRUN_OPTIONS` also specifies the `-n` option, you must ensure that the values used for both options are the same. To do otherwise is an error.

X1_DYNAMIC_COMMON_SIZE

The `X1_DYNAMIC_COMMON_SIZE` sets the size of the dynamic `COMMON` block defined by the loader. Refer to the `-LD_LAYOUT:dynamic=` option in the `ld(1)` man page. Also refer to *Optimizing Applications on Cray X1 Series Systems* for more information about dynamic `COMMON` blocks.

X1_COMMON_STACK_SIZE
X1_PRIVATE_STACK_SIZE
X1_STACK_SIZE
X1_LOCAL_HEAP_SIZE
X1_SYMMETRIC_HEAP_SIZE
X1_HEAP_SIZE
X1_PRIVATE_STACK_GAP

These environment variables allow you to change the default size of the application stacks or heaps, or consolidate the private stacks:

- `X1_COMMON_STACK_SIZE`, change the common stack size to the specified value.

- `X1_PRIVATE_STACK_SIZE`, change the private stack size to the specified value.

- `X1_STACK_SIZE`, set the size of the common and private stack to the specified value.

- `X1_LOCAL_HEAP_SIZE`, change the local heap size to the specified value.

- `X1_SYMMETRIC_HEAP_SIZE`, change the symmetric heap size to the specified value.

- `X1_HEAP_SIZE`, change the local and symmetric heap size to the specified value.

- `X1_PRIVATE_STACK_GAP`, consolidate, when used with `X1_PRIVATE_STACK_SIZE`, the four private stacks within an MSP into one segment, which frees up nontext pages for application use. The specified value, in bytes, indicates the gap to separate each stack. This gap serves as a guard region in case any of the stacks overflow.

The default size of each application stack or heap is 1 GB.

The `X1_STACK_SIZE` and `X1_HEAP_SIZE` are termed general environment variables in that they set the values for multiple stacks or heaps, respectively. The other variables in this section are termed specific because they set the value for a particular stack or heap. A specific variable overrides a general variable if both are specified as follows:

- The `X1_COMMON_STACK_SIZE` variable overrides the `X1_STACK_SIZE` variable if both are specified.

- The `X1_PRIVATE_STACK_SIZE` variable overrides the `X1_STACK_SIZE` if both are specified.

- The `X1_LOCAL_HEAP_SIZE` variable overrides the `X1_HEAP_SIZE` variable if both are specified.

- The `X1_SYMMETRIC_HEAP_SIZE` overrides the `X1_HEAP_SIZE` variable if both are specified.

The value you specify for a variable sets the size of a stack or heap in bytes. This number can be expressed as a decimal number, an octal number with a leading zero, or a hexadecimal number with a leading "0x".

If you specify a number smaller than the page size you gave to the `aprun` or `mpirun` command, the system will silently enforce a single-page minimum size. If you do not use the `aprun` command or do not specify a page size for `aprun`, the minimum page size is set to 64 KB. Refer to the `-p text:other` option of the `aprun`(1) man page for more information about page sizes.

Using the `X1_PRIVATE_STACK_GAP` and `X1_PRIVATE_STACK_SIZE` environment variables together to consolidate the private stacks may help applications that have problems obtaining a sufficient number of large nontext pages via the `aprun` or `mpirun` commands. When the private stacks are consolidated, the pages that would have been used by the other private stacks are freed so they can be used by the application.

Each MSP used by an application uses four private stacks where each private stack occupies an integral number of pages, but if the application actually needs a private stack that is much smaller than the integral number of pages, space is wasted. In some of these cases, consolidating all four private stacks into one segment will free up the wasted space so it can be used by the application. For example, an application uses 256 MB pages, which means the size of each private stack is a multiple of 256 MB. If the application only needs 60 MB for each private stack, we can consolidate all four private stacks into a 256 MB page by setting `X1_PRIVATE_STACK_SIZE` to 0x3c00000 (60 MB) and `X1_PRIVATE_STACK_GAP` to 0x400000 (4 Mb). This packs the four private stacks into one 256 MB page with a 4 MB guard region between the stacks. This saves three 256 MB physical pages on each MSP.

Warning: You should be aware that there is no protection against overflowing the private stacks; one private stack may corrupt another with unpredictable results if stack overflow occurs.

2.25 OpenMP Environment Variables

This section describes the OpenMP C and C++ API environment variables that control the execution of parallel code. The names of environment variables must be uppercase. The values assigned to them are case insensitive and may have leading and trailing white space. Modifications to the values after the program has started are ignored.

The environment variables are as follows:

- `OMP_SCHEDULE` sets the run time schedule type and chunk size

- `OMP_NUM_THREADS` sets the number of threads to use during execution

- `OMP_DYNAMIC` enables or disables dynamic adjustment of the number of threads

- `OMP_NESTED` enables or disables nested parallelism

- `OMP_THREAD_STACK_SIZE` is a Cray-specific, nonstandard variable used to change the size of the thread stack from the default size of 16 MB to the specified size.

The examples in this section only demonstrate how these variables might be set in UNIX C shell (`csh`) environments:

```
setenv OMP_SCHEDULE "dynamic"
```

In Korn shell environments, the actions are similar, as follows:

```
export OMP_SCHEDULE="dynamic"
```

2.25.1 `OMP_SCHEDULE`

`OMP_SCHEDULE` applies only to `for` and `parallel for` directives that have the schedule type `runtime`. The schedule type and chunk size for all such loops can be set at run time by setting this environment variable to any of the recognized schedule types and to an optional *chunk_size*.

For `for` and `parallel for` directives that have a schedule type other than `runtime`, `OMP_SCHEDULE` is ignored. The default value for this environment variable is implementation-defined. If the optional *chunk_size* is set, the value must be positive. If *chunk_size* is not set, a value of 1 is assumed, except in the case of a `static` schedule. For a `static` schedule, the default chunk size is set to the loop iteration space divided by the number of threads applied to the loop.

Example:

```
setenv OMP_SCHEDULE "guided,4"
setenv OMP_SCHEDULE "dynamic"
```

2.25.2 `OMP_NUM_THREADS`

The `OMP_NUM_THREADS` environment variable sets the default number of threads to use during execution, unless that number is explicitly changed by calling the `omp_set_num_threads` library routine (see the `omp_threads`(3) man page) or by an explicit `num_threads` clause on a `parallel` directive.

The value of the OMP_NUM_THREADS environment variable must be a positive
integer. Its effect depends upon whether dynamic adjustment of the number
of threads is enabled. For information about the interaction between the
OMP_NUM_THREADS environment variable and dynamic adjustment of threads,
see Section 5.2, page 114.

If no value is specified for the OMP_NUM_THREADS environment variable, or if
the value specified is not a positive integer, or if the value is greater than the
maximum number of threads the system can support, the number of threads to
use is implementation-defined.

Example:

```
setenv OMP_NUM_THREADS 16
```

2.25.3 OMP_DYNAMIC

The OMP_DYNAMIC environment variable enables or disables dynamic
adjustment of the number of threads available for execution of parallel regions
unless dynamic adjustment is explicitly enabled or disabled by calling the
omp_set_dynamic library routine (see the omp_threads(3) man page). Its
value must be TRUE or FALSE. The default condition is FALSE.

If set to TRUE, the number of threads that are used for executing parallel regions
may be adjusted by the run time environment to best utilize system resources.

If set to FALSE, dynamic adjustment is disabled.

Example:

```
setenv OMP_DYNAMIC TRUE
```

2.25.4 OMP_NESTED

The OMP_NESTED environment variable enables or disables nested parallelism
unless nested parallelism is enabled or disabled by calling the omp_set_nested
library routine (see the omp_nested(3) man page). If set to TRUE, nested
parallelism is enabled; if it is set to FALSE, nested parallelism is disabled. The
default value is FALSE.

Example:

```
setenv OMP_NESTED TRUE
```

2.25.5 `OMP_THREAD_STACK_SIZE`

The `OMP_THREAD_STACK_SIZE` environment variable changes the size of the thread stack from the default size of 16 MB to the specified size. The size of the thread stack should be increased when thread-private variables may utilize more than 16 MB of memory.

The requested thread stack space is allocated from the local heap when the threads are created. The amount of space used by each thread for thread stacks depend on whether you are using MSP or SSP mode. In MSP mode, the memory used is five times the specified thread stack size because each SSP is assigned one thread stack and one thread stack is used as the MSP common stack. For SSP mode, the memory used is one times the specified thread stack size.

This is the format for the `OMP_THREAD_STACK_SIZE` environment variable:

`OMP_THREAD_STACK_SIZE` *n*

where *n* is a decimal number, an octal number with a leading zero, or a hexadecimal number with a leading "0x" specifying the amount of memory, in bytes, to allocate for a thread's stack.

For more information about memory on the Cray X1 series system, see the `memory`(7) man page.

Example:

`setenv OMP_THREAD_STACK_SIZE 18000000`

`#pragma` directives are used within the source program to request certain kinds of special processing. `#pragma` directives are part of the C and C++ languages, but the meaning of any `#pragma` directive is defined by the implementation. `#pragma` directives are expressed in the following form:

> `#pragma [ _CRI]` *identifier* [*arguments*]

The `_CRI` specification is optional and ensures that the compiler will issue a message concerning any directives that it does not recognize. Diagnostics are not generated for directives that do not contain the `_CRI` specification.

These directives are classified according to the following types:

- General (Section 3.5, page 70)

- Instantiation (Cray C++ only) (Section 3.6, page 79)

- Vectorization (Section 3.7, page 80)

- Multistreaming (Section 3.8, page 87)

- Scalar (Section 3.9, page 91)

- Inlining (Section 3.10, page 96)

Macro expansion occurs on the directive line after the directive name. That is, macro expansion is applied only to *arguments*.

> **Note:** OpenMP `#pragma` directives are described in Chapter 5, page 113. UPC `#pragma` directives are described in Chapter 6, page 119.

At the beginning of each section that describes a directive, information is included about the compilers that allow the use of the directive and the scope of the directive. Unless otherwise noted, the following default information applies to each directive:

Compiler: Cray C and Cray C++

Scope: Local and global

The scoping list may also indicate that a directive has a lexical block scope. A lexical block is the scope within which a directive is on or off and is bounded by the opening curly brace just before the directive was declared and the

corresponding closing curly brace. Only applicable executable statements within the lexical block are affected as indicated by the directive. The lexical block does not include the statements contained within a procedure that is called from the lexical block.

This example code fragment shows the lexical block for the upc strict and upc relaxed directives:

```
void Example(void)
{
    #pragma _CRI upc strict  // UPC strict state is on
    ...

    {
        ...                 // UPC strict state is still on
        #pragma _CRI upc relaxed // UPC strict state is now off
        ...
    }

// UPC strict state is back on
    ...
}
```

3.1 Protecting Directives

To ensure that your directives are interpreted only by the Cray C and C++ compilers, use the following coding technique in which *directive* represents the name of the directive:

```
#if _CRAYC
    #pragma _CRI directive
#endif
```

This ensures that other compilers used to compile this code will not interpret the directive. Some compilers diagnose any directives that they do not recognize. The Cray C and C++ compilers diagnose directives that are not recognized only if the _CRI specification is used.

3.2 Directives in Cray C++

C++ prohibits referencing undeclared objects or functions. Objects and functions must be declared prior to using them in a #pragma directive. This is not always the case with C.

Some #pragma directives take function names as arguments (for example: #pragma _CRI weak, #pragma _CRI suppress, and #pragma _CRI inline_always *name* [, *name*] ...). Member functions and qualified names are allowed for these directives.

3.3 Loop Directives

Many directives apply to groups. Unless otherwise noted, these directives must appear before a for, while, or do while loop. These directives may also appear before a label for if...goto loops. If a loop directive appears before a label that is not the top of an if...goto loop, it is ignored.

3.4 Alternative Directive form: **_Pragma**

Compiler directives can also be specified in the following form, which has the advantage in that it can appear inside macro definitions:

_Pragma("_CRI *identifier*");

This form has the same effect as using the #pragma form, except that everything that appeared on the line following the #pragma must now appear inside the double quotation marks and parentheses. The expression inside the parentheses must be a single string literal; it cannot be a macro that expands into a string literal. _Pragma is an extension to the C and C++ standards.

The following is an example using the #pragma form:

#pragma _CRI ivdep

The following is the same example using the alternative form:

_Pragma("_CRI ivdep");

In the following example, the loop automatically vectorizes wherever the macro is used:

```
#define SEARCH(A, B, KEY, SIZE, RES)
{
   int i;
   _Pragma("_CRI ivdep");
   for (i = 0; i < (SIZE); i++)
      if ( (A)[ (B)[i] ] == (KEY)) break;
   (RES)=i;
}
```

Macros are expanded in the string literal argument for `_Pragma` in an identical fashion to the general specification of a `#pragma` directive.

3.5 General Directives

General directives specify compiler actions that are specific to the directive and have no similarities to the other types of directives. The following sections describe general directives.

3.5.1 `[no]bounds` Directive (Cray C Compiler)

The `bounds` directive specifies that pointer and array references are to be checked. The `nobounds` directive specifies that this checking is to be disabled.

When bounds checking is in effect, pointer references are checked to ensure that they are not 0 or are not greater than the machine memory limit. Array references are checked to ensure that the array subscript is not less than 0 or greater than or equal to the declared size of the array.

Both directives may be used only within function bodies. They apply until the end of the function body or until another `bounds`/`nobounds` directive appears. They ignore block boundaries.

These directives have the following format:

```
#pragma _CRI bounds
```

```
#pragma _CRI nobounds
```

The following example illustrates the use of the `bounds` directive:

```
int a[30];
#pragma _CRI bounds
void f(void)
{
   int x;
   x = a[30];
   .
   .
   .
}
```

3.5.2 `duplicate` Directive (Cray C Compiler)

Scope: Global

The `duplicate` directive lets you provide additional, externally visible names for specified functions. You can specify duplicate names for functions by using a directive with one of the following forms:

```
#pragma _CRI duplicate actual as dupname...

#pragma _CRI duplicate actual as (dupname...)
```

The *actual* argument is the name of the actual function to which duplicate names will be assigned. The *dupname* list contains the duplicate names that will be assigned to the actual function. The *dupname* list may be optionally parenthesized. The word `as` must appear as shown between the *actual* argument and the comma-separated list of *dupname* arguments.

The `duplicate` directive can appear anywhere in the source file and it must appear in global scope. The actual name specified on the directive line must be defined somewhere in the source as an externally accessible function; the actual function cannot have a static storage class.

The following example illustrates the use of the `duplicate` directive:

```
#include <complex.h>

extern void maxhits(void);

#pragma _CRI duplicate maxhits as count, quantity      /* OK */

void maxhits(void)
{
   #pragma _CRI duplicate maxhits as tempcount
   /* Error: #pragma _CRI duplicate can't appear in local scope */
}

double _Complex minhits;

#pragma _CRI duplicate minhits as lower_limit
/* Error: minhits is not declared as a function */

extern void derivspeed(void);
```

```
#pragma _CRI duplicate derivspeed as accel
/* Error: derivspeed is not defined */

static void endtime(void)
{
}

#pragma _CRI duplicate endtime as limit
/* Error: endtime is defined as a static function */
```

Because duplicate names are simply additional names for functions and are not functions themselves, they cannot be declared or defined anywhere in the compilation unit. To avoid aliasing problems, duplicate names may not be referenced anywhere within the source file, including appearances on other directives. In other words, duplicate names may only be referenced from outside the compilation unit in which they are defined.

The following example references duplicate names:

```
void converter(void)
{
   structured(void);
}

#pragma _CRI duplicate converter as factor, multiplier /* OK */

void remainder(void)
{
}

#pragma _CRI duplicate remainder as factor, structured
/* Error: factor and structured are referenced in this file */
```

Duplicate names can be used to provide alternate external names for functions, as shown in the following examples.

main.c:

```
extern void fctn(void), FCTN(void);

main()
{
  fctn();
  FCTN();
}
```

fctn.c:

```
#include <stdio.h>

void fctn(void)
{
  printf("Hello world\n");
}

#pragma _CRI duplicate fctn as FCTN
```

Files `main.c` and `fctn.c` are compiled and linked using the following command line:

```
% cc main.c fctn.c
```

When the executable file `a.out` is run, the program generates the following output:

```
Hello world
Hello world
```

3.5.3 `message` Directive

The `message` directive directs the compiler to write the message defined by *text* to `stderr` as a warning message. Unlike the `error` directive, the compiler continues after processing a `message` directive. The format of this directive is as follows:

```
#pragma _CRI message "text"
```

The following example illustrates the use of the `message` compiler directive:

```
#define FLAG 1

#ifdef FLAG
#pragma _CRI message "FLAG is Set"
#else
#pragma _CRI message "FLAG is NOT Set"
#endif
```

3.5.4 `no_cache_alloc` Directive

The `no_cache_alloc` directive is an advisory directive that specifies objects that should not be placed into the cache. Advisory directives are directives the compiler will honor if conditions permit it to. When this directive is honored, the performance of your code may be improved because the cache is not occupied by objects that have a lower cache hit rate. Theoretically, this makes room for objects that have a higher cache hit rate.

Here are some guidelines that will help you determine when to use this directive. This directive works only on objects that are vectorized. That is, other objects with low cache hit rates can still be placed into the cache. Also, you should use this directive for objects you feel should not be placed into the cache.

To use the directive, you must place it only in the specification part, before any executable statement.

The format of the `no_cache_alloc` directive is:

```
#pragma _CRI no_cache_alloc base_name [,base_name] ...
```

base_name   The base name of the object that should not be placed into the cache. This can be the base name of any object such as an array, scalar structure, etc., without member references like `C[10]`. If you specify a pointer in the list, only the references, not the pointer itself, have the no cache allocate property.

3.5.5 `cache_shared` Directive

Scope: Declaration

This directive asserts that all vector loads with the specified symbols as the base are to be made using cache-shared instructions. This an advisory directive; if the compiler honors it, vector load misses cause the cache line to be allocated in a shared state, in anticipation of a subsequent load by a different MSP. For vector store operations, this directive is not meaningful and will be ignored. Scalar loads and stores also are unaffected. The compiler may override the directive if it determines the directive is not beneficial. The scope of this directive is the scope of the declaration of the specified symbol.

The format of the `cache_shared` directive is:

```
#pragma _CRI cache_shared symbol [,symbol...]
```

symbol   A base symbol (an array or scalar structure, but not a member reference or array element).

Examples of valid cache_shared symbols are A, B, C. Expressions such as B.E or C[10] cannot be used as cache_shared symbols.

3.5.6 [`no`]`opt` Directive

Scope: Global

The `noopt` directive disables all automatic optimizations and causes optimization directives to be ignored in the source code that follows the directive. Disabling optimization removes various sources of potential confusion in debugging. The `opt` directive restores the state specified on the command line for automatic optimization and directive recognition. These directives have global scope and override related command line options.

The format of these directives is as follows:

```
#pragma _CRI opt
```

```
#pragma _CRI noopt
```

The following example illustrates the use of the `opt` and `noopt` compiler directives:

```
#include <stdio.h>

void sub1(void)
{
        printf("In sub1, default optimization\n");
}

#pragma _CRI noopt
void sub2(void)
{
        printf("In sub2, optimization disabled\n");
}
#pragma _CRI opt

void sub3(void)
{
        printf("In sub3, optimization enabled\n");
}

main()
{
        printf("Start main\n");
        sub1();
        sub2();
        sub3();
}
```

3.5.7 `weak` Directive

Scope: Global

The `weak` directive specifies an external identifier that may remain unresolved throughout the compilation. A `weak` external reference can be to a function or to a data object. A `weak` external does not increase the total memory requirements of your program.

Declaring an object as a weak external directs the loader to do one of these tasks:

• Link the object only if it is already linked (that is, if a strong reference exists); otherwise, leave it is as an unsatisfied external. The loader does not display an unsatisfied external message if weak references are not resolved.

- If a strong reference is specified in the weak directive, resolve all weak references to it.

  Note: The loader treats weak externals as unsatisfied externals, so they remain silently unresolved if no strong reference occurs during compilation. Thus, it is your responsibility to ensure that run time references to weak external names do not occur unless the loader (using some "strong" reference elsewhere) has actually loaded the entry point in question.

These are the forms of the weak directive:

```
#pragma _CRI weak var
```

```
#pragma
   _CRI weak sym1 = sym2
```

var	The name of an external
sym1	Defines an externally visible weak symbol
sym2	Defines an externally visible strong symbol defined in the current compilation.

The first form allows you to declare one or more weak references on one line. The second form allows you to assign a strong reference to a weak reference.

The weak directive must appear at global scope.

The attributes that weak externals must have depend on the form of the weak directive that you use:

- First form, weak externals must be declared, but not defined or initialized, in the source file.

- Second form, weak externals may be declared, but not defined or initialized, in the source file.

- Either form, weak externals cannot be declared with a static storage class.

The following example illustrates these restrictions:

```
extern long x;
#pragma _CRI weak x  /* x is a weak external data object */
extern void f(void);
#pragma _CRI weak f  /* f is a weak external function */

extern void g(void);
#pragma _CRI weak g=fun;  /* g is a weak external function
                                 with a strong reference to fun */

long y = 4;
#pragma _CRI weak y  /* ERROR - y is actually defined */

static long z;
#pragma _CRI weak z  /* ERROR - z is declared static */

void fctn(void)
{
#pragma _CRI weak a  /* ERROR - directive must be at global scope */
}
```

3.5.8 `vfunction` Directive

Scope: Global

The `vfunction` directive lists external functions that use the call-by-register calling sequence. Such functions can be vectorized but must be written in Cray Assembly Language. See *Cray Assembly Language (CAL) for Cray X1 Systems Reference Manual.*

The format of this directive is as follows:

```
#pragma _CRI vfunction func
```

The *func* variable specifies the name of the external function.

The following example illustrates the use of the `vfunction` compiler directive:

```
extern double vf(double);
#pragma _CRI vfunction vf


void f3(int n) {
    int i;
    for (i = 0; i < n; i++) {     /* Vectorized */
        b[i] = vf(c[i]);
    }
}
```

3.5.9 `ident` Directive

The `ident` directive directs the compiler to store the string indicated by *text* into the object (`.o`) file. This can be used to place a source identification string into an object file.

The format of this directive is as follows:

```
#pragma _CRI ident
    text
```

3.6 Instantiation Directives

The Cray C++ compiler recognizes three instantiation directives. Instantiation directives can be used to control the instantiation of specific template entities or sets of template entities. The following directives are described in detail in Section 8.5, page 149:

- `#pragma _CRI instantiate`

- `#pragma _CRI do_not_instantiate`

- `#pragma _CRI can_instantiate`

- The `#pragma _CRI instantiate` directive causes a specified entity to be instantiated.

- The `#pragma _CRI do_not_instantiate` directive suppresses the instantiation of a specified entity. It is typically used to suppress the instantiation of an entity for which a specific definition is supplied.

- The `#pragma _CRI can_instantiate` directive indicates that a specified entity can be instantiated in the current compilation, but need not be. It is used in conjunction with automatic instantiation to indicate potential sites for instantiation if the template entity is deemed to be required by the compiler.

  See Chapter 8, page 143 for more information on template instantiation.

3.7 Vectorization Directives

Because vector operations cannot be expressed directly in Cray C and C++, the compilers must be capable of vectorization, which means transforming scalar operations into equivalent vector operations. The candidates for vectorization are operations in loops and assignments of structures. For more information, see *Optimizing Applications on Cray X1 Series Systems.*

The subsections that follow describe the compiler directives used to control vectorization.

3.7.1 `ivdep` Directive

Scope: Local

The `ivdep` directive tells the compiler to ignore vector dependencies for the loop immediately following the directive. Conditions other than vector dependencies can inhibit vectorization. If these conditions are satisfactory, the loop vectorizes. This directive is useful for some loops that contain pointers and indirect addressing. The format of this directive is as follows:

```
#pragma _CRI ivdep safevl=vlen|infinitevl
```

vlen   Specifies a vector length in which no dependency will occur. *vlen* must be an integer between 1 and 1024 inclusive.

infinitevl   Specifies an infinite safe vector length. This option asserts that no data dependency will occur at any vector length.

The following example illustrates the use of the `ivdep` compiler directive:

```
p = a; q = b;
#pragma _CRI ivdep
for (i = 0; i < n; i++) {          /* Vectorized */
        *p++ = *q++;
}
```

On the Cray X1 series system, the compiler by default assumes an infinite safe vector length; that is, any vector length can safely be used to vectorize the loop. You can use the `-h noinfinitevl` compiler option to change this behavior for all loops in the compilation unit.

Caution: Use the `ivdep` pragma with caution. Asserting a safe vector length that proves to be not safe can produce incorrect results. Refer to *Optimizing Applications on Cray X1 Series Systems* for further information.

3.7.2 `loop_info` Directive

Scope: Local

The `loop_info` directive provides loop trip count information that the optimizer uses to produce faster code sequences.

The format of this directive is as follows:

```
#pragma _CRI loop_info [min_trips(c)] [est_trips(c)]
   [max_trips(c)]
```

c	An expression that evaluates to an integer constant at compilation time
`min_trips`	Keyword for the guaranteed minimum number of trips
`est_trips`	Keyword for the estimated/average number of trips
`max_trips`	Keyword for the guaranteed maximum number of trips

In the following example, the estimated trip count is 100 and the maximum trip count is 10000:

```
#pragma _CRI loop_info min_trips(100) max_trips(10000)
for (i = 0; i < n; i++) {
  a[i] = b[i];
}
```

3.7.3 `nopattern` Directive

Scope: Local

The `nopattern` directive disables pattern matching for the loop immediately following the directive.

The format of this directive is as follows:

```
#pragma _CRI nopattern
```

By default, the compiler detects coding patterns in source code sequences and replaces these sequences with calls to optimized library functions. In most cases, this replacement improves performance. There are cases, however, in which this substitution degrades performance. This can occur, for example, in loops with very low trip counts. In such a case, you can use the `nopattern` directive to disable pattern matching and cause the compiler to generate inline code.

In the following example, placing the `nopattern` directive in front of the outer loop of a nested loop turns off pattern matching for the matrix multiply that takes place inside the inner loop:

```
double a[100][100], b[100][100], c[100][100];

void nopat(int n)
{
  int i, j, k;

#pragma _CRI nopattern
  for (i=0; i < n; ++i) {
    for (j = 0; j < n; ++j) {
      for (k = 0; k < n; ++k) {
```

```
                c[i][j] += a[i][k] * b[k][j];
              }
          }
      }
}
```

3.7.4 `novector` Directive

Scope: Local

The `novector` directive directs the compiler to not vectorize the loop that immediately follows the directive. It overrides any other vectorization-related directives, as well as the `-h vector` and `-h ivdep` command line options. The format of this directive is as follows:

```
#pragma _CRI novector
```

The following example illustrates the use of the `novector` compiler directive:

```
#pragma _CRI novector
for (i = 0; i < h; i++) {    /* Loop not vectorized */
    a[i] = b[i] + c[i];
}
```

3.7.5 `novsearch` Directive

Scope: Local

The `novsearch` directive directs the compiler to not vectorize the search loop that immediately follows the directive. A search loop is a loop with one or more early exit statements. It overrides any other vectorization-related directives as well as the `-h vector` and `-h ivdep` command line options. The format of this directive is as follows:

```
#pragma _CRI novsearch
```

The following example illustrates the use of the `novsearch` compiler directive:

```
#pragma _CRI novsearch
for (i = 0; i < h; i++) {    /* Loop not vectorized */
    if (a[i] < b[i]) break;
    a[i] = b[i];
}
```

3.7.6 [`no`]`pipeline` Directive

Software-based vector pipelining (software vector pipelining) provides additional optimization beyond the normal hardware-based vector pipelining. In software vector pipelining, the compiler analyzes all vector loops and automatically attempts to pipeline a loop if doing so can be expected to produce a significant performance gain. This optimization also performs any necessary loop unrolling.

In some cases the compiler either does not pipeline a loop that could be pipelined or pipelines a loop without producing performance gains. In these situations, you can use the `pipeline` or `nopipeline` directive to advise the compiler to pipeline or not pipeline the loop immediately following the directive.

Software vector pipelining is valid only for the innermost loop of a loop nest.

The `pipeline` and `nopipeline` directives are advisory only. While you can use the `nopipeline` directive to inhibit automatic pipelining, and you can use the `pipeline` directive to attempt to override the compiler's decision not to pipeline a loop, you cannot force the compiler to pipeline a loop that cannot be pipelined.

Loops that have been pipelined are so noted in loopmark listing messages.

The formats of the pipelining directives are as follows:

```
#pragma _CRI pipeline

#pragma _CRI nopipeline
```

For more information about software vector pipelining, see *Optimizing Applications on Cray X1 Series Systems.*

3.7.7 **`prefervector` Directive**

Scope: Local

The `prefervector` directive tells the compiler to vectorize the loop that immediately follows the directive if the loop contains more than one loop in the nest that can be vectorized. The directive states a vectorization preference and does not guarantee that the loop has no memory dependence hazard.

The format of this directive is as follows:

```
#pragma _CRI prefervector
```

The following example illustrates the use of the `prefervector` directive:

```
#pragma _CRI prefervector
for (i = 0; i < n; i++) {
   #pragma _CRI ivdep
   for (j = 0; j < m; j++)
      a[i] += b[j][i];
}
```

In the preceding example, both loops can be vectorized, but the directive directs the compiler to vectorize the outer `for` loop. Without the directive and without any knowledge of `n` and `m`, the compiler vectorizes the inner `for` loop. In this example, the outer `for` loop is vectorized even though the inner `for` loop had an `ivdep` directive.

3.7.8 **`safe_address` Directive**

Scope: Local

The format of this directive is as follows:

```
#pragma _CRI safe_address
```

The `safe_address` directive allows you to tell the compiler that it is safe to speculatively execute memory references within all conditional branches of a loop. In other words, you know that these memory references can be safely executed in each iteration of the loop.

For most code, the `safe_address` directive can improve performance significantly by preloading vector expressions. However, most loops do not

require this directive to have preloading performed. The directive is required only when the safety of the operation cannot be determined or index expressions are very complicated.

The `safe_address` directive is an advisory directive. That is, the compiler may override the directive if it determines the directive is not beneficial.

If you do not use the directive on a loop and the compiler determines that it would benefit from the directive, it issues a message indicating such. The message is similar to this:

```
CC-6375 cc: VECTOR File = ctest.c, Line = 6
  A loop would benefit from "#pragma safe_address".
```

If you use the directive on a loop and the compiler determines that it does not benefit from the directive, it issues a message that states the directive is superfluous and can be removed.

To see the messages, you must use the `-hreport=v` option.

⚠️ **Caution:** Incorrect use of the directive can result in segmentation faults, bus errors, or excessive page faulting. However, it should not result in incorrect answers. Incorrect usage can result in very severe performance degradations or program aborts.

In the example below, the compiler will not preload vector expressions, because the value of `j` is unknown. However, if you know that references to `b[i][j]` is safe to evaluate for all iterations of the loop, regardless of the condition, we can use the `SAFE_ADDRESS` directive for this loop as shown below:

```
void x3( double a[restrict 1000], int j )
{
    int i;
#pragma _CRI safe_address
    for ( i = 0; i < 1000; i++ ) {
        if ( a[i] != 0.0 ) {
            b[j][i] = 0.0;
        }
    }
}
```

With the directive, the compiler can load `b[i][j]` with a full vector mask, merge `0.0` where the condition is true, and store the resulting vector using a full mask.

3.7.9 `shortloop` and `shortloop128` Directives

Scope: Local

The `shortloop` and `shortloop128` directives improve performance of a vectorized loop by allowing the compiler to omit the run time test to determine whether it has been completed. The `shortloop` compiler directive identifies vector loops that execute with a maximum iteration count of 64 and a minimum iteration count of 1. The `shortloop128` compiler directive identifies vector loops that execute with a maximum iteration count of 128 and a minimum iteration count of 1. If the iteration count is outside the range for the directive, results are unpredictable.

These directives are ignored if the loop trip count is known at compile time and is greater than the target machine's vector length. The maximum hardware vector length is 64.

The syntax of these directives are as follows:

```
#pragma _CRI shortloop
#pragma _CRI shortloop128
```

The following examples illustrate the use of the `shortloop` and `shortloop128` directives:

```
#pragma _CRI shortloop
for (i = 0; i < n; i++) {  /* 0 < = n < = 63 */
    a[i] = b[i] + c[i];
}


#pragma _CRI shortloop128
for (i = 0; i < n; i++) { /* 0 < = n < = 127 */
    a[i] = b[i] + c[i];
}
```

3.8 Multistreaming Processor (MSP) Directives

This section describes the multistreaming processor (MSP) optimization directives. MSPs are advisory directives; the compiler is not obligated to honor them. For information about MSP compiler options, refer to Section 2.11, page 29 and for streaming intrinsics, refer to Appendix F, page 225. For details on Cray Streaming Directives, see Chapter 4, page 101.

The MSP directives work with the `-h stream`*n* command line option to determine whether parts of your program are optimized for the MSP. The level of multistreaming must be greater than 0 in order for these directives to be recognized. For more information on the `-h stream`*n* command line option, see Section 2.11.1, page 29.

The MSP `#pragma` directives are as follows:

- `#pragma _CRI ssp_private` (see Section 3.8.1, page 88)

- `#pragma _CRI nostream` (see Section 3.8.2, page 90)

- `#pragma _CRI preferstream` (see Section 3.8.3, page 91)

3.8.1 `ssp_private` Directive (`cc`, `c99`)

The `ssp_private` directive allows the compiler to multistream loops that contain function calls. By default, the compiler does not multistream loops containing function calls, because the function may cause side effects that interfere with correct parallel execution. The `ssp_private` directive asserts that the specified function is free of side effects that inhibit parallelism and that the specified function, and all functions it calls, will run on an SSP.

An implied condition for multistreaming a loop containing a call to a function specified with the `ssp_private` directive is that the loop body must not contain any data reference patterns that prevent parallelism. The compiler can disregard an `ssp_private` directive if it detects possible loop-carried dependencies that are not directly related to a call inside the loop.

Note: The `ssp_private` directive affects only whether or not loops are multistreamed. It has no effect on loops within CSD parallel regions.

When using the `ssp_private` directive, you must ensure that the function called within the body of the loop follows these criteria:

- The function does not modify an object in one iteration and reference this same data in another iteration of the multistreamed loop.

- The function does not reference data in one iteration that is defined in another iteration.

- If the function modifies data, the iterations cannot modify data at the same storage location, unless these variables are scoped as `PRIVATE`. Following the multistreamed loop, the content of private variables are undefined.

The ssp_private directive does not force the master thread to execute the last iteration of the multistreamed loop.

- If the function uses shared data that can be written to and read, you must protect it with a guard (such as the CSD critical directive or the lock command) or have the SSPs access the data disjointedly (where access does not overlap).

- The function calls only other routines that are capable of being called privately.

- The function calls I/O properly.

    Note: The preceding list assumes that you have a working knowledge of race conditions.

To use the ssp_private directive, it must placed in the specification part, before any executable statements. This is the syntax of the ssp_private directive:

```
#pragma _CRI ssp_private PROC_NAME[,PROC_NAME] ...
```

PROC_NAME	The name of a function. Any number of ssp_private directives may be specified in a function. If a function is specified with the ssp_private directive, the function retains this attribute throughout the entire program unit. Also, the ssp_private directive is considered a declarative directive and must be specified before the start of any executable statements.

The following example demonstrates use of the ssp_private pragma:

```
/* Code in example.c */
extern void poly_eval( float *y, float x, int m, float p[m] );
#pragma _CRI ssp_private poly_eval

void example(int n, int m, float x[n], float y[n], float p[])
{
  int i;

  for (i = 0; i < n; ++i) {
    poly_eval( &y[i], x[i], m, p );
  }
}

/* Code in example poly_eval.c */
```

```
void poly_eval( float *y, float x, int m, float p[] )
{
  float result = p[m];
  int    i;

  for (i = m-1; m >= 0; --m) {
    result = x * result + p[i];
  }
  *y = result;
}
```

This example compiles the code:

```
% cc -c example.c
% cc -c -h gen_private_callee poly_eval.c
% cc -o example example.o poly_eval.o
```

Now run the code:

```
% aprun -L1 ./example
```

SSP private routines are appropriate for user-specified math support functions. Intrinsic math functions, like COS are effectively SSP private routines.

3.8.2 `nostream` Directive

Scope: Local

The #pragma _CRI nostream directive directs the compiler to not perform MSP optimizations on the loop that immediately follows the directive. It overrides any other MSP-related directives as well as the -h stream*n* command line option.

The format of this directive is as follows:

```
#pragma _CRI nostream
```

The following example illustrates the use of the nostream directive:

```
#pragma _CRI nostream
   for ( i = 0; i < n1; i++ ) {
     x[i] = y[i] + z[i];
     }
```

3.8.3 `preferstream` Directive

Scope: Local

The `preferstream` directive tells the compiler to multistream the following loop. It can be used when one of these conditions apply:

- The compiler issues a message saying there are too few iterations in the loop to make multistreaming worthwhile.

- The compiler multistreams a loop in a loop nest, and you want it to multistream a different eligible loop in the same nest.

The format of this directive is as follows:

```
#pragma _CRI preferstream
```

The following example illustrates the use of the `preferstream` directive:

```
   for (j = 0; j< n2; j++ ) {
#pragma _CRI preferstream
      for (i = 0; i < n1; i++ ) {
         a[j][i] = b[j][i] + c[j][i];
      }
   }
```

3.9 Scalar Directives

This section describes the scalar optimization directives, which control aspects of code generation, register storage, and other scalar operations.

3.9.1 `concurrent` Directive

Scope: Local

The `concurrent` directive indicates that no data dependence exists between array references in different iterations of the loop that follows the directive. This can be useful for vectorization and multistreaming optimizations.

The format of the `concurrent` directive is as follows:

```
#pragma _CRI concurrent [safe_distance=n]
```

> *n*   An integer that represents the number of additional consecutive loop iterations that can be executed in parallel without danger of data conflict. *n* must be an integral constant > 0.
>
> The `concurrent` directive is ignored if the `safe_distance` clause is used and MSP optimizations, multistreaming, or vectorization is requested on the command line.

In the following example, the `concurrent` directive indicates that the relationship `k>3` is true. The compiler will safely load all the array references `x[i-k]`, `x[i-k+1]`, `x[i-k+2]`, and `x[i-k+3]` during `i`-th loop iteration.

```
#pragma _CRI concurrent safe_distance=3

for (i = k + 1; i < n;i++) {
  x[i] = a[i] + x[i-k];
}
```

3.9.2 `nointerchange` Directive

> Scope: Local

The `nointerchange` directive inhibits the compiler's ability to interchange the loop that follows the directive with another inner or outer loop.

The format of this directive is as follows:

```
#pragma _CRI nointerchange
```

In the following example, the `nointerchange` directive prevents the `iv` loop from being interchanged by the compiler with either the `jv` loop or the `kv` loop:

```
    for (jv = 0; jv < 128; jv++) {
#pragma _CRI nointerchange
        for (iv = 0; iv < m; iv++) {
            for (kv = 0; kv < n; kv++) {
                p1[iv][jv][kv] = pw[iv][jv][kv] * s;
            }
        }
    }
```

3.9.3 `noreduction` Directive

Scope: Local

The `noreduction` compiler directive tells the compiler to not optimize the loop that immediately follows the directive. If the loop is not a reduction loop, the directive is ignored.

A *reduction loop* is a loop that contains at least one statement that reduces an array to a scalar value by doing a cumulative operation on many of the array elements. This involves including the result of the previous iteration in the expression of the current iteration.

You may choose to use this directive when the loop iteration count is small or when the order of evaluation is numerically significant. It overrides any vectorization-related directives as well as the `-h vector` and `-h ivdep` command line options.

The `noreduction` directive disables vectorization of any loop that contains a reduction. The specific reductions that are disabled are summation and product reductions, and alternating value computations. The directive also prevents the compiler from rewriting loops involving multiplication or exponentiation by an induction variable to be a series of additions or multiplications of a value.

Regardless of platform, however, the format of this directive is as follows:

```
#pragma _CRI noreduction
```

The following example illustrates the use of the `noreduction` compiler directive:

```
sum = 0;
#pragma _CRI noreduction
for (i = 0; i < n; i++) {
    sum += a[i];
}
```

3.9.4 `suppress` Directive

The `suppress` directive suppresses optimization in two ways, determined by its use with either global or local scope.

The global scope `suppress` directive specifies that all associated local variables are to be written to memory before a call to the specified function. This ensures

that the value of the variables will always be current. The global `suppress` directive takes the following form:

```
#pragma _CRI suppress func...
```

The local scope `suppress` directive stores current values of the specified variables in memory. If the directive lists no variables, all variables are stored to memory. This directive causes the values of these variables to be reloaded from memory at the first reference following the directive. The local `suppress` directive has the following format:

```
#pragma _CRI suppress [var] ...
```

The net effect of the local `suppress` directive is similar to declaring the affected variables to be `volatile` except that the `volatile` qualifier affects the entire program whereas the local `suppress` directive affects only the block of code in which it resides.

3.9.5 [`no`]`unroll` Directive

Scope: Local

The `unroll` directive allows the user to control unrolling for individual loops or to specify no unrolling of a loop. Loop unrolling can improve program performance by revealing cross-iteration memory optimization opportunities such as read-after-write and read-after-read. The effects of loop unrolling also include:

- Improved loop scheduling by increasing basic block size

- Reduced loop overhead

- Improved chances for cache hits

The format for this compiler directive is as follows:

```
#pragma _CRI [no]unroll[n]
```

The `nounroll` directive disables loop unrolling for the next loop and does not accept the integer argument *n*. The `nounroll` directive is equivalent to the `unroll0` and `unroll1` directives.

The *n* argument applies only to the unroll directive and specifies no loop unrolling (*n* = 0 or 1) or the total number of loop body copies to be generated (2 ≤ *n* ≤ 63).

If you do not specify a value for *n*, the compiler will determine the number of copies to generate based on the number of statements in the loop nest.

Note: The compiler cannot always safely unroll non-innermost loops due to data dependencies. In these cases, the directive is ignored (see Example 2, page 96).

The unroll directive can be used only on loops with iteration counts that can be calculated before entering the loop. If unroll is specified on a loop that is not the innermost loop in a loop nest, the inner loops must be nested perfectly. That is, all loops in the nest can contain only one loop, and the innermost loop can contain work.

Example 1: Unrolling Outer Loops

In the following example, assume that the outer loop of the following nest will be unrolled by 2:

```
#pragma _CRI  unroll2
for (i = 0; i < 10; i++) {
      for (j = 0; j < 100; j++) {
        a[i][j] = b[i][j] + 1;
      }
}
```

With outer loop unrolling, the compiler produces the following nest, in which the two bodies of the inner loop are adjacent:

```
for (i = 0; i < 10; i += 2) {
      for (j = 0; j < 100; j++) {
        a[i][j] = b[i][j] + 1;
      }
      for (j = 0; j < 100; j++) {
        a[i+1][j] = b[i+1][j] + 1;
      }
}
```

The compiler then *jams*, or *fuses*, the inner two loop bodies, producing the following nest:

```
for (i = 0; i < 10; i += 2) {
      for (j = 0; j < 100; j++) {
```

```
                        a[i][j] = b[i][j] + 1;
                        a[i+1][j] = b[i+1][j] + 1;
                    }
}
```

Example 2: Illegal Unrolling of Outer Loops

Outer loop unrolling is not always legal because the transformation can change
the semantics of the original program. For example, unrolling the following loop
nest on the outer loop would change the program semantics because of the
dependency between `a[i][...]` and `a[i+1][...]`:

```
/* directive will cause incorrect code due to dependencies! */
#pragma _CRI  unroll2
for (i = 0; i < 10; i++) {
  for (j = 1; j < 100; j++) {
    a[i][j] = a[i+1][j-1] + 1;
  }
}
```

3.10 Inlining Directives

Inlining replaces calls to user-defined functions with the code in the calling
process that represents the function. This can improve performance by saving
the expense of the function call overhead. It also enhances the possibility of
additional code optimization and vectorization, especially if the function call was
an inhibiting factor. Inlining may increase object code size.

Inlining is invoked in the following ways:

- Automatic inlining of functions in an entire compilation unit is enabled by the
  -h ipa*n* option, as described in Section 2.13, page 32. Automatic inlining of
  portions of code is enabled by the `inline_enable` directive and disabled by
  the `inline_disable` directive (Section 3.10.1, page 97).

- Explicit inlining of particular function calls in portions of code is enabled by
  the #pragma _CRI inline_always *name* [, *name*] directive and disabled
  by the #pragma _CRI inline_never *name* [, *name*] directive (Section
  3.10.3, page 99). Explicit inlining of particular function calls within an entire
  compilation unit is enabled by the -h ipafrom=*source* [:*source*] option
  described in Section 2.13.2, page 34.

- Combined inlining by the use of both the -h ipa*n* and -h ipafrom=*source*
  [:*source*] options (see Section 2.13.3, page 35).

Inlining directives can appear in global scope (that is, not inside a function definition). Global inlining directives specify whether all calls to the specified functions should be inlined.

Inlining directives can also appear in local scope; that is, inside a function definition. A local inlining directive applies only to the next call to the function specified on the directive. Although the function specified on an inlining directive does not need to appear in the next statement, a call to the function must occur before the end of the function definition.

Inlining directives always take precedence over the automatic inlining requested on the command line. This means that function calls that are associated with inlining directives are inlined before any function calls selected to be inlined by automatic inlining.

Note: A function that contains a variable length array argument is not currently inlined.

The -h report=i option writes messages identifying where functions are inlined or briefly explains why functions are not inlined.

3.10.1 inline_enable and inline_disable Directives

The inline_enable directive tells the compiler to attempt to inline functions at call sites. It has the following format:

```
#pragma _CRI inline_enable
```

The inline_disable directive tells the compiler to not inline functions at call sites. It has the following format:

```
#pragma _CRI inline_disable
```

The following example illustrates the use of the inline_enable directive.

Example 3: Using the inline_enable Directive

Source code:

```
#include <stdio.h>

int f(int a)
{
```

```
   return a*a;
}

main()
{

  #pragma _CRI inline_enable /* Direct the compiler to inline
    calls. */

  int b = 5;
  printf("%d\n", f(b)); /* f is inlined here */
}
```

Compiler command:

`% cc -h report=i -o ipaex3 ipaex3.c`

Report:

```
CC-3001 cc: INLINE File = ipaex3.c, Line = 14
  The call to f was textually inlined.
```

Program execution:

```
% ./ipaex3
25
```

3.10.2 `inline_reset` Directive

The `inline_reset` directive returns the inlining state to the default state. The format of the `inline_reset` directive is as follows:

```
#pragma _CRI inline_reset
```

Example 4: Using the `inline_reset` Directive

The following code fragment shows how the `#pragma _CRI inline_reset` directive would affect code compiled with the `-h ipa3` option:

```
...
#pragma _CRI inline_disable  /* From here to inline_reset,
no inlining will be done */
```

```
void f1()
{
  ...
}

#pragma _CRI inline_reset  /* From here forward, the
inlining state is -h ipa3 */
void f2()
{
  ...
}
```

3.10.3 `inline_always` and `inline_never` Directives

The `inline_always` directive specifies functions that the compiler should always attempt to inline. The format of the `inline_always` directive is as follows:

```
#pragma _CRI inline_always name [, name ] ...
```

The `inline_never` directive specifies functions that are never to be inlined. The format of the `inline_never` directive is as follows:

```
#pragma _CRI inline_never name [, name ] ...
```

The *name* argument is the name of a function.

Cray Streaming Directives (CSDs) [4]

The Cray streaming directives (CSDs) consist of six non-advisory directives which allow you to more closely control multistreaming for key loops in C and C++ programs. Non-advisory means that the compiler must honor these directives. The intention of these directives is not to create an additional parallel programming style or demand large effort in code development. They are meant to assist the compiler in multistreaming your program. On its own, the compiler should perform multistreaming correctly in most cases. However, if you feel that multistreaming for key loops is not occurring as you desire, then use the CSDs to override the compiler.

CSDs are modeled after the OpenMP directives and are compatible with Pthreads and all distributed-memory parallel programming models on Cray X1 series systems. Multistreaming advisory directives (MSP directives) and CSDs cannot be mixed within the same block of code. For information on MSPs, see Section 3.8, page 87.

Before explaining guidelines and other issues, you need an understanding of these CSD items:

*   CSD parallel regions

*   CSD `parallel` (defines a CSD parallel region)

*   CSD `for` (multistreams a `for` loop)

*   CSD `parallel for` (combines the CSD `parallel` and `for` directives into one directive)

*   CSD `sync` (synchronizes all SSPs within an MSP)

*   CSD `critical` (defines a critical section of code)

*   CSD `ordered` (specifies that SSPs execute in order)

When you are familiar with the directives, these topics will be beneficial to you:

*   Using CSDs with Cray programming models

*   CSD Placement

*   Protection of shared data

*   Dynamic memory allocation for CSD parallel regions

*   Compiler options affecting CSDs

> **Note:** Refer to *Optimizing Applications on Cray X1 Series Systems* for information about how to use the CSDs to optimize your code.

4.1 CSD Parallel Regions

CSDs are applied to a block of code (for example a loop), which will be referred to as the CSD parallel region. All CSDs must be used within this region. You must not branch into or out of the region.

Multiple CSD parallel regions can exist within a program; however, only one parallel region will be active at any given time. For example, if a parallel region calls a function containing a parallel region, the function will execute as if it did not contain a parallel region.

The CSD parallel region can contain loops and nonloop constructs, but only loops preceded by a `for` directive are partitioned. Parallel execution of other loops and nonloop constructs, such as initializing variables for the targeted loop, are performed redundantly on all SSPs. Functions called from the region will be executed redundantly, and loops within them can be partitioned with the `for` directive. Parallel execution of the function is independent on all SSPs, except for code blocks containing standalone CSDs. Refer to Section 4.9, page 110.

4.2 `parallel` Directive

The `parallel` directive defines the CSD parallel region, tells the compiler to multistream the region, and specifies private data objects. All other CSDs must be used within the region. You cannot place the `parallel` directive in the middle of a construct.

This is the form of the parallel directives:

```
#pragma _CRI csd parallel [private(list)] [ordered]
{
  structured_block
}   /* End of CSD parallel region */
```

The `private` clause allows you to specify data objects that are private to each SSP within the CSD parallel region; that is, each SSP has its own copy of that object and is not shared with other SSPs. The main reason for having private objects is because updating them within the CSD parallel region could cause incorrect updates because of race conditions on their addresses. The *list* argument specifies a comma separated list of objects to make private.

By default, the variables used for loop indexing are assumed to be private. Variables declared in the inner scope of a parallel region are implicitly private. Other variables, unless specified in the private clause, are assumed to be shared.

You may need to take special steps when using private variables. If a data object existed before the parallel region is entered and the object is made private, the object may not have the same contents inside of the region as it did outside the region. The same is true when exiting the parallel region. This same object may not have the same content outside the region as it did within the region. Therefore, if you desire that a private object keep the same value when transitioning in and out of the parallel region, copy its value to a protected shared object so you can copy it back into the private object later.

The ordered clause is needed if there is within the parallel region, but outside the loops within the region, any call to a function containing a CSD ordered directive. That is, if only the loops contain calls to functions that contain the CSD ordered directive, the clause is not needed. If the clause is used and there are no called functions containing a CSD ordered directive, the results produced by the code encapsulated by the directive will be correct, but performance of that code will be slightly degraded. If the ordered clause is missing and there is a called function containing a CSD ordered directive, your results will be incorrect. The following example shows when the ordered clause is needed:

```
#pragma _CRI csd parallel ordered
{
  fun();  /* fun contains ordered directive */

  for_loop_block
    . . .
}
```

The end of the CSD parallel region has an implicit barrier synchronization. The implicit barrier protects an SSP from prematurely accessing shared data.

Note: At the point of the parallel directive, all SSPs are enabled and are disabled at the end of the CSD parallel region.

This example shows how to use the parallel directive:

```
#pragma _CRI csd parallel private(jx)
{
   x = 2 * PI;   /* This line is computed on all SSPs */
   for(i=1; i<n; i++)
   {
     jx = y[i] * z[i] * x;  /* jx is private to each SSP */
```

```
         ...
      }
} /* End of CSD parallel region */
```

4.3 `for` Directive

The compiler distributes among the SSPs the iteration of `for` loops modified by the CSD `for` directive. Iterations of `for` loops not modified by the CSD `for` directives are not distributed among the SSPs, but are all redundantly executed on all SSPs.

Refer to Section 4.9, page 110 for placement restrictions of the CSD `for` directive.

This is the syntax of the CSD `for` directive:

```
#pragma _CRI csd for [schedule(static [, chunk_size])]
[nowait] [ordered]
for_statement {
        ...
} /* End of for loop and CSD for region */
```

The `schedule` clause specifies how the loop iterations are distributed among the SSPs. This iteration distribution is fixed (`static`) at compile time and cannot be changed by run time events.

The iteration distribution is calculated by you or the compiler. You or the compiler will divide the number of iterations into groups or *chunks*. The compiler will then statically assign the chunks to the 4 SSPs in a round-robin fashion according to iteration order (in other words, from the first iteration to the last iteration). Therefore, an SSP could have one or more chunks. The number of iterations in each chunk is called the *chunk size* which is specified by the *chunk_size* argument.

You can use these tips to calculate the chunk size:

- Balance the parallel work load across all 4 SSPs (the number of SSPs in an MSP) by dividing the number of iterations by 4. If you have a remainder, add one to the chunk size. Using 4 chunks gives you the best performance, because less overhead is incurred when using fewer chunks per SSP.

- The work load distribution among the SSPs will be imbalanced if the chunk size is greater than $1/4^{\text{th}}$ of the total number of iterations.

- If the chunk size is greater than the total number of iterations, the first SSP (SSP0) will do all the work.

The compiler calculates the iteration distribution (*chunk_size*) if the schedule clause or *chunk_size* argument is not specified. The value used is dependent on the conditions shown in Table 11.

Table 11. Compiler-calculated Chunk Size

Calculated chunk size	Condition
1	When a sync, critical, or ordered CSD directive or a function call appears in the loop.
Iterations / 4	The number of iterations are divided as evenly as possible into four chunks if these are not present in the CSD parallel region: sync, critical, or ordered directive or a function call.

An implicit barrier synchronization occurs at the end of the for region, unless the nowait clause is also specified. The implicit barrier protects an SSP from prematurely accessing shared data. The nowait clause assumes that you are guaranteeing that consumption-before-production cannot occur.

The ordered clause is needed if the for loop encapsulated by the CSD for directive calls any function containing a CSD ordered directive. If the clause is used and there are no called functions containing a CSD ordered directive, the results produced by the code encapsulated by the directive will be correct, but performance of that code will be slightly degraded. If the ordered clause is missing and there is a called function containing a CSD ordered directive, the results produced by the code encapsulated by the directive will be incorrect. The following example shows when the ORDERED clause is needed:

```
#pragma _CRI csd parallel
{
   ...

#pragma csd for ordered
  for(i=1, i<n; i++)
    fun(i) /* fun contains ordered directive */
}
```

The following examples illustrate compiler and user calculated chunk sizes. For this example, the compiler calculates the chunk size as 1, because of the function call (a chunk size of 1 causes SSP0 to perform iterations 1, 5, 9, ... , SSP1 to perform iterations 2, 6, 10, ...):

```
#pragma _CRI csd for
for(i=1; i<num_samples; i++)
{
    process_sample(sample[i]);
} /* End of CSD for region */
```

For this example, because there are no `sync`, `critical`, or `ordered` directives, or subprogram calls, the compiler calculates the chunk size as `(arraySize + 3) / 4`:

```
#pragma _CRI csd for
for(i=1; i<arraySize; i++) {
  product[i] = operand[i] * operand[i];
} /* End of CSD for region */
```

Adding 3 to the array size produces an optimal chunk size by grouping the maximum number of iterations into 4 chunks.

This example specifies the `schedule` clause and a chunk size of 128:

```
#pragma _CRI csd for schedule(static,128)
for(i=1; i<array_size; i++) {
  product[i] = operand[i] * operand[i];
} /* End of CSD for region */
```

In the above example, the compiler will use the chunk size based on this statement `min(array_size, 128)`. If the chunk size is larger than the array size, the compiler will use the array as the chunk size. If this is the case, then all the work will be done by SSP0.

4.4 `parallel for` Directive

The `parallel for` directive combines most of the functionality of the CSD `parallel` and `for` directives into one directive. The `parallel for` directive is used on a single `for` loop that contains or does not contain nested loops and is the equivalent to the following statements:

```
#pragma _CRI csd parallel [private(list)]
{
  #pragma _CRI csd for [schedule(static [,chunk])]
  for_loop_block
} /* End of CSD parallel for region */
```

The differences between the `parallel for` and its counter parts include the lack of the `nowait` clause, because it is not needed.

This is the form of the `parallel for` directive:

```
#pragma _CRI csd parallel for [private(list)] [schedule(static
[, chunk_size])]
for_statement {
    loop_block
} /* End of CSD parallel for region */
```

For a description the `parallel for` directive, refer to the `parallel` and `for` directives at Section 4.2, page 102 and Section 4.3, page 104.

4.5 `sync` Directive

The `sync` directive synchronizes all SSPs within a multistreaming processor (MSP) and may under certain conditions synchronize memory with physical storage by calling `msync`. The `sync` directive is normally used where additional intra-MSP synchronization is needed to prevent race conditions caused by forced multistreaming.

The `sync` directive can appear anywhere within the CSD parallel region, even within the CSD `for` and `parallel for` directives. If the `sync` directive appears within a CSD parallel region, but outside of an enclosed CSD `for` directive, then it performs an `msync` on all four SSPs.

This example shows how to use the `sync` directive:

```
#pragma _CRI csd parallel for private(j)
{
  for(i=1; i<4; i++) {
    for(j=1; j<100000; j++) {
      x[j][i] = ... ;         /* Produce x */
    }

      #pragma _CRI csd sync

    for(j=1; j<100000; j++) {
        ... = x[j][5-i]; * ... /* Consume x */
    }
  }
}
```

The two inner loops provide a producer and consumer pair for array `x`. The `sync` directive prevents the use of the array by the second inner loop before it is completely populated.

4.6 `critical` Directive

The `critical` directive specifies a critical region where only one SSP at a time will execute the enclosed region.

This is the form of the `critical` directive:

```
#pragma _CRI csd critical
{
    block_of_code
} /* End of critical region */
```

This example performs a multistreamed sum reduction of array `a` and uses the `critical` directive to calculate the complete sum:

```
sum = 0;   /* Shared variable */

#pragma _CRI csd parallel private(private_sum)
{
  private_sum = 0;

  #pragma _CRI csd for
  for(i=1; i<a_size; i++) {
  private_sum = private_sum + a(i);
  }

  #pragma _CRI csd critical
  {
    sum = sum + private_sum;
  }
}
```

4.7 `ordered` Directive

The `ordered` directive allows you to have loops with particular dependencies on other loops in the parallel region by ensuring the execution order of the SSPs. That is, SSP0 completes execution of its block of code in the ordered region before

SSP1 executes that same block of code; SSP1 completes execution of that block of code before SSP2 can execute it, etc.

If the CSD `ordered` directive is placed in a function that is called from a parallel region, the CSD `parallel`, `parallel for`, or `for` directives that encapsulate the call may also need the `ordered` clause to ensure correct results. See the appropriate CSD directive for more information.

This is the format of the `ordered` directive:

```
#pragma _CRI csd ordered
{
    block_of_code
} /* End of ordered region */
```

In following example, successive iterations of the loop depend upon previous iterations, because of `a[i-1]` and `a[i-2]` on the right side of the first assignment statement. The `ordered` directive ensures that each computation of `a[i]` is complete before the next iteration (which occurs on the next SSP) uses this value as its `a[i-1]` and similarly for `a[i-2]`:

```
#pragma _CRI csd parallel for schedule(static,1)
for(i=3; i<a_size; i++) {
   #pragma _CRI csd ordered
   {
      a[i] = a[i-1] + a[i-2];
   }

   ...    /* other processing */
}
```

If the execution time for the code indicated by the `other processing` comment is larger than the time to compute the assignment within the ordered directive, then the loop will mostly run concurrently on the 4 SSPs, even if the ordered directive is used.

4.8 Nested CSDs Within Cray Parallel Programming Models

CSDs can be mixed with all parallel programming models within the same program on Cray X1 series systems. If you nest them, the CSDs must be at the inner most level. These are the nesting levels:

1. Distributed memory models (MPI, SHMEM, UPC, and CAF)

2. Shared memory models (OpenMP and Pthreads)

3. CSDs

If the shared or distributed memory model is used, then you can nest the CSDs within either one. These models cannot be nested within the CSDs. If both memory models are used, then the CSDs must be nested within the shared memory model, and the shared memory model nested within the distributed memory model.

4.9 CSD Placement

CSDs must be used within the CSD parallel region as defined by the `parallel` directive. Some must be used where the parallel directives are used; that is, used within the same block of code. Other CSDs can be used in the same block of code or be placed in a function and called from the parallel region (in effect, appearing as if they were within the parallel region). These CSDs will be referred to as standalone CSDs.

The CSD `for` directive is the only one that must be used within the same block of code as this example shows:

```
#pragma _CRI csd parallel
{
   ...
  #pragma _CRI csd for
  for_loop_block
  ...
}
```

The standalone CSDs are `sync`, `critical`, and `ordered`. If standalone CSDs are placed in a function and the function is not called from a parallel region, the code will execute as if no CSD code exists.

4.10 Protection of Shared Data

Updates to shared data, both directly in a CSD parallel region and within functions called from a CSD parallel region, must be protected against simultaneous execution by SSPs used for the CSD parallel region. Shared data include statically allocated data objects (such as file scope variables or variables declared with the static storage class), dynamically allocated data objects pointed to by more than one SSP, and function parameters that point to shared objects.

Protecting your shared data includes using the `private` list of the `parallel` and `parallel for` directives, the `critical` directive, or `for` loop indices.

Accesses to shared arrays made within a CSD `for` loop are in effect private and therefore need no synchronization if the accesses use indices that involve the loop control variable of the `for` loop.

This example shows access to the `sum` shared array using loop control variable `i`:

```
#pragma csd parallel for
for(i=0; i<n; i++) {
  initialize ( &a[i] );
}
```

The `critical` directive can protect updates to shared data by ensuring that only one SSP at any one time can execute the enclosed code that accesses the shared data.

4.11 Dynamic Memory Allocation for CSD Parallel Regions

There are certain precautions you should remember as you allocate or free dynamic memory that is accessed through private or shared pointers.

Calls to the `libc` dynamic memory allocation routines (`malloc`, `free`, etc.) within CSD parallel regions must be made by only one SSP at a time. In general, this will require that they be made from within CSD critical regions. This requirement may be relaxed in future releases.

Dynamic memory for pointers specified in the `private` list of the `parallel` directive must be allocated and freed within the CSD parallel region. Dynamic memory cannot be allocated for private objects before entering the CSD parallel region and made private when within the region simply by specifying a pointer to that object in the private list of the parallel directive.

Dynamic memory can be allocated for shared pointers outside or within the CSD parallel region. If memory for the shared pointer is allocated or freed within the CSD parallel region, you must ensure that it is allocated or freed by only one SSP.

This example shows how to ensure that only one SSP deallocates the memory for private variable `a`:

```
...
double *a;
...
```

```
#pragma csd parallel private(a)
{
  #pragma critical
  a = malloc(SIZE * sizeof(double));
  ...
  #pragma critical
  free(a);
}
```

4.12 Compiler Options Affecting CSDs

To enable CSDs, compile your code with the -h stream*n* option with *n* set to 1
or greater. Also, specify the -h gen_private_callee option to compile
procedures called from the CSD parallel region. To disable CSDs, specify the
-h stream0 option.

OpenMP C and C++ API Directives [5]

This chapter describes the OpenMP directives that the Cray C and C++ Compilers support. These directives are based on the *OpenMP C and C++ Application Program Interface Version 2.0 March 2002* standard. Copyright © 1997–2002 OpenMP Architecture Review Board.

In addition to directives, the OpenMP C and C++ API describes several run time library routines and environment variables. For information on the library routines, see the OpenMP man pages. For information on the environment variables, see Section 2.25, page 62.

The sections in this chapter are as follows:

- OpenMP Feature Restriction (Section 5.1, page 113)

- Cray implementation differences (Section 5.2, page 114)

- `OMP_THREAD_STACK_SIZE` environment variable (Section 5.3, page 115)

- Compiler options affecting OpenMP (Section 5.4, page 116)

- OpenMP program execution (Section 5.5, page 116)

5.1 Deferred OpenMP Features

The following OpenMP Fortran features are not yet supported by the Cray C and Cray C++ compilers:

- An object of a class with a nontrivial default constructor, a nontrivial copy constructor, a nontrivial destructor, or a nontrivial copy assignment operator can be used in a scoping clause. Currently, these objects can be used only in a shared scope.

  In C++, an object of a class with a nontrivial constructor, a nontrivial copy constructor, a nontrivial destructor, or a nontrivial copy assignment operator cannot be used in a scoping clause for any scope except shared. This restriction will be lifted in a future release.

- The appearance of array names in the `REDUCTION` clause.

- The reprivatization of variables.

5.2 Cray Implementation Differences

The OpenMP C and C++ Application Program Interface specification defines areas of implementation that have vendor-specific behaviors. Those areas are described in the following list:

Note: The following information was created by copying from the OpenMP specification the text surrounding the phrase "implementation-dependent" and replacing that phrase with the Cray specific implementation information (shown in bold).

- Implementation-dependent areas of the `parallel` construct:

    – If none of the methods[1] above were used, then the number of threads requested is defined by the *depth* value you define through the `aprun -d depth` option. If this option is not set, the `aprun` command defaults the depth to one.

    – If a parallel region is encountered while dynamic adjustment of the number of threads is disabled, and the number of threads requested for the parallel region exceeds the number that the runtime system can supply, the program will terminate.

    – The number of physical processors actually hosting the threads at any given time is fixed at program startup and is specified by the `aprun -d depth` option.

    – The number of threads in a team that execute a nested parallel region is one because all nested parallel regions are serialized.

- Implementation-dependent areas of the `for` construct:

    – When `schedule(runtime)` is specified, the decision regarding scheduling is deferred until runtime. The schedule kind and size of the chunks can be chosen at run time by setting the environment variable `OMP_SCHEDULE`. If this environment variable is not set, the schedule type and chunk size default to `GUIDED` and 1, respectively. When `schedule(runtime)` is specified, `chunk_size` must not be specified.

    – In the absence of an explicitly defined schedule clause, the default schedule is `STATIC` and the default chunk size is roughly the number of iterations divided by the number of threads..

---

[1] Methods that explicitly specify the number of threads were not used. In other words, the number of threads was not explicitly specified through the `num_threads` clause of the parallel region and was not specified by a call to `omp_set_num_threads` or by the `OMP_NUM_THREADS` environment variable.

- Implementation-dependent area of the `ATOMIC` construct—The `ATOMIC` directive is replaced with a critical section that encloses the statement.

- Implementation-dependent areas in the OpenMP library functions:

  - `omp_get_nested`—This procedure always returns `0` because nested parallel regions are always serialized.

  - `omp_get_num_threads`—If the number of threads has not been explicitly set by the user, the default is the *depth* value defined through the `aprun -d depth` option. If this option is not set, the `aprun` command defaults `depth` to one, which sets the number of threads to one, which value `omp_get_num_threads` returns.

  - `omp_set_dynamic`—The default for the dynamic adjustment of threads is **on**.

  - `omp_set_nested`—Calls to this function are ignored since nested parallel regions are always serialized and executed by a team of one thread.

- Implementation-dependent areas of the OpenMP environment variables:

  - `OMP_DYNAMIC`—The default value is **TRUE**.

  - `OMP_NUM_THREADS`—If no value is specified for the `OMP_NUM_THREADS` environment variable, it defaults to *depth* as defined by the `aprun -d depth` option or to 1 if the option is not specified.

    If the value specified for the `OMP_NUM_THREADS` environment variable is not a positive integer, a warning message is printed and the program behaves as if no value was specified.

    If the value specified for the `OMP_NUM_THREADS` environment variable is greater than the maximum number of threads the system can support, the behavior of the program depends on the value of the `OMP_DYNAMIC` environment variable. If `OMP_DYNAMIC` is `FALSE`, the program terminates, otherwise it uses as many threads as possible.

  - `OMP_SCHEDULE`—The default values for this environment variable are `GUIDED` for schedule and 1 for chunk size.

5.3 `OMP_THREAD_STACK_SIZE`

`OMP_THREAD_STACK_SIZE` is a Cray specific OpenMP environment variable that changes the size of the thread stack from the default size of 16 MB to

the specified size. The size of the thread stack should be increased when thread-private variables may utilize more than 16 MB of memory.

The requested thread stack space is allocated from the local heap when the threads are created. The amount of space used by each thread for thread stacks depend on whether you are using MSP or SSP mode. In MSP mode, the memory used is five times the specified thread stack size because each SSP is assigned one thread stack and one thread stack is used as the MSP common stack. For SSP mode, the memory used equals the specified thread stack size.

This is the format for the `OMP_THREAD_STACK_SIZE` environment variable:

`OMP_THREAD_STACK_SIZE` *n*

where *n* is a decimal number, an octal number with a leading zero, or a hexadecimal number with a leading "0x" specifying the amount of memory, in bytes, to allocate for a thread's stack.

For more information about memory on the Cray X1 series system, see the `memory`(7) man page.

Example:

`setenv OMP_THREAD_STACK_SIZE 18000000`

5.4 Compiler Options Affecting OpenMP

These Cray C and C++ Compiler options enable or disable the compiler recognition of OpenMP directives:

- Enable OpenMP directive recognition: `-h omp`

- Disable OpenMP directive recognition: `-h noomp`

5.5 OpenMP Program Execution

The `aprun` command can be used to define the default thread count to use for OpenMP parallel regions for programs that do not explicitly define the number of threads to use (that is, you do not use the `OMP_NUM_THREADS` environment variable or the program does not use the `omp_set_num_threads` library procedure). If these programs do not use `aprun` to set the number of threads, all OpenMP directives are ignored.

Use this command line to set the default thread count:

`aprun -d` *depth*

where *depth* is the number of threads to run. For programs that do not specify the number of threads, you must set *depth* to either 4 when using MSP mode or 16 when using SSP mode.

Using the `-d` option causes your OpenMP program to run as if the `OMP_NUM_THREADS` environment variable was set.

These options of the `aprun` command do not affect the number of OpenMP threads: `-n` (define number of processors) and `-N` (define number of processors per node).

Cray Unified Parallel C (UPC) [6]

Unified Parallel C (UPC) is a C language extension for parallel program development. UPC allows you to explicitly specify parallel programming through language syntax rather than library functions such as those used in MPI and SHMEM by allowing you to read and write memory of other processes with simple assignment statements. Program synchronization occurs only when explicitly programmed; there is no implied synchronization. These methods map very well onto the Cray X1 series systems and enable users to achieve high performance.

> **Note:** UPC is a dialect of the C language. It is not available in C++.

UPC allows you to maintain a view of your program as a collection of threads operating in a common global address space without burdening you with details of how parallelism is implemented on the machine (for example, as shared memory or as a collection of physically distributed memories).

UPC data objects are private to a single thread or shared among all threads of execution. Each thread has a unique memory space that holds its private data objects, and access to a globally shared memory space that is distributed across the threads. Thus, every part of a shared data object has an affinity to a single thread.

Cray UPC is compatible with MPI, SHMEM, and CAF.

> **Note:** Currently, the UPC model does not define an I/O model. Therefore, you must supply the controls as needed to remove race conditions. File I/O under UPC is very similar to standard C because one thread opens a file and shares the file handle, and multiple threads may read or write to the same file.

We assume that you are familiar with UPC and understand the differences between the published UPC Introduction and Language Specification paper and the current UPC specification.

If you are not familiar with UPC, refer to the UPC home page at `http://upc.gwu.edu/`. Under the Publications link, select the *Introduction to UPC and Language Specification* paper. This paper is slightly outdated but contains valuable information about understanding and using UPC.

The UPC home page also contains, under the Documentation link, the *UPC Language Specification 1.1.1* paper, which is up to date.

Cray supports the UPC Language Specification 1.1.1 and adds Cray-specific functions as noted in the following sections.

Note: Because of changes made to the Cray UPC libraries to support the UPC 1.1.1 specification, UPC binaries produced by earlier Cray C++ Programming Environment releases are not compatible with UPC binaries produced by the Cray C++ Programming Environment 5.2 release. They must be recompiled by the 5.2 compiler. If you link incompatible binaries, the linker displays an error message explaining the presence of incompatible UPC object files and does not create an executable file.

After familiarizing yourself with UPC, refer to the following sections for details:

- Predefined identifiers (Section 6.1, page 121)

- UPC expressions (Section 6.2, page 121)

- UPC statements (Section 6.3, page 121)

- UPC `#pragma` directives (Section 6.4, page 124)

- Predefined macro names (Section 6.5, page 125)

- Standard headers (Section 6.6, page 125)

- UPC functions (Section 6.7, page 126)

- Cray implementation differences (Section 6.8, page 138)

- Compiling and executing UPC code (Section 6.9, page 138)

For details, see the UPC man pages.

For a description of the `-h upc` command line option see Section 2.21.11, page 55.

> **Note:** Some UPC constructs perform more efficiently than others. For more information about UPC optimization guidelines, refer to *Optimizing Applications on Cray X1 Series Systems.*

6.1 Predefined Identifiers

UPC recognizes the following predefined identifiers:

THREADS   `THREADS` is a value of type `int`; it specifies the number of independent computational units and has the same value on every thread. Under the static `THREADS` translation environment, `THREADS` is an integer constant suitable for use in `#if` preprocessing directives.

MYTHREAD   `MYTHREAD` is a value of type `int`; it specifies the unique thread index. The range of possible values is `0...THREADS-1`.

UPC_MAX_BLOCK_SIZE   `UPC_MAX_BLOCK_SIZE` is a predefined integer constant value. It indicates the maximum value allowed in a layout qualifier for shared data. It is suitable for use in `#if` preprocessing directives.

6.2 UPC Expressions

Cray supports the following expressions:

- `upc_localsizeof` returns the size, in bytes, of the local portion of its operand, which may be a shared object or a shared-qualified type

- `upc_blocksizeof` returns the block size of the operand, which may be a shared object or a shared-qualified type

- `upc_elemsizeof` returns the size, in bytes, of the highest-level (leftmost) type that is not an array; for non-array objects, `upc_elemsizeof` returns the same value as `sizeof`

6.3 UPC Statements

Cray supports UPC barrier statements and the iteration statements.

6.3.1 UPC Barrier Statements

Cray supports the following UPC barrier statements:

- `upc_notify`

  The syntax is:

  `upc_notify` [*exp*]`;`

- `upc_wait`

  The syntax is:

  `upc_wait` [*exp*]`;`

- `upc_barrier`

  The syntax is:

  `upc_barrier` [*exp*]`;`

- `upc_fence`

  The syntax is:

  `upc_fence;`

where *exp* is an integer expression.

Each thread must execute an alternating sequence of `upc notify` and `upc wait` statements, starting with a `upc notify` and ending with a `upc wait` statement. A synchronization phase consists of the execution of all statements between the completion of one `upc wait` and the start of the next.

A `upc wait` statement completes and the thread enters the next synchronization phase only after all threads have completed the `upc notify` statement in the current synchronization phase. `upc wait` and `upc notify` are collective operations.

The `upc fence` statement is equivalent to a null strict reference. This insures that all shared references issued before the fence are complete before any after it are issued.

A null strict reference is implied before a `upc notify` statement and after a `upc wait` statement.

The `upc wait` statement will generate a runtime error if the value of its expression does not equal the value of the expression by the `upc notify`

statement for the current synchronization phase. No error will be generated if either statement does not have an expression.

The upc wait statement will generate a runtime error if the value of its expression differs from any expression in the upc wait and upc notify statements issued by any thread in the current synchronization phase. No error will be generated from a difference involving a statement for which no expression is given.

The barrier operations at thread startup and termination have a value of expression which is not in the range of user expressible values.

6.3.2 UPC Iteration Statements

Cray supports the forall statement.

The syntax is:

upc_forall ([*exp*][;*exp*][;*exp*][*affinity*]) *statement*

The expression for *affinity* has pointer-to-shared type or integer type.

upc_forall is a collective operation in which, for each execution of the loop body, the controlling expression and affinity expression are single-valued.

The *affinity* field specifies the executions of the loop body which are to be performed by a thread.

When *affinity* is of pointer-to-shared type, the loop body of the upc_forall statement is executed for each iteration in which the value of MYTHREAD equals the value of upc_threadof(*affinity*). Each iteration of the loop body is executed by precisely one thread.

When *affinity* is an integer expression, the loop body of the upc_forall statement is executed for each iteration in which the value of MYTHREAD equals the value *affinity* mod THREADS.

When *affinity* is continue or not specified, each loop body of the upc_forall statement is performed by every thread.

If the loop body of a upc_forall statement contains one or more upc_forall statements, either directly or through one or more function calls, the construct is called a nested upc_forall statement. In a nested upc_forall, the outermost upc_forall statement that has an *affinity* expression which is not continue is called the controlling upc_forall statement. All upc_forall statements

which are not controlling in a nested `upc_forall` behave as if their *affinity* expressions were `continue`.

If the execution of any loop body of a `upc_forall` statement produces a side-effect which affects the execution of another loop body of the same `upc_forall` statement which is executed by a different thread, the behavior is undefined.

If any thread terminates or executes a `upc_barrier`, `upc_notify`, or `upc_wait` statement within the dynamic scope of a `upc_forall` statement, the result is undefined. If any thread terminates a `upc_forall` statement using a `break`, `goto`, or `return` statement, the result is undefined. If any thread enters the body of a `upc_forall` statement using a `goto` statement, the result is undefined.

6.4 UPC #pragma Directives

Cray supports the UPC `strict` and `relaxed` pragma directives. The syntax is:

```
#pragma _CRI upc strict
```

```
#pragma _CRI upc relaxed
```

These pragmas affect the strict or relaxed categorization of references to shared objects where the referenced type is neither strict-qualified nor relaxed-qualified. Such references are strict if a `strict` pragma is in effect or relaxed if a `relaxed` pragma is in effect.

Shared references which are not categorized by either referenced type or by these pragmas behave in an implementation-defined manner in which either all such references are strict or all are relaxed. Users wishing portable programs are strongly encouraged to categorize all shared references either by using type qualifiers, these directives, or by including a `upc strict.h` or `upc relaxed.h` header file.

The pragmas must occur either outside external declarations (i.e., not within a function body) or preceding all explicit declarations and statements inside a compound statement (i.e., a block). When `upc strict` or `upc relaxed` is outside a function body, the directive applies to everything from that point on in the source file and any `#include` files, unless another file-scope directive appears and overrides it. When inside a compound statement, the directives follow the scoping rules; their applicability ends at the `}`, which ends the

compound statement. If these pragmas are used in any other context, their behavior is undefined.

6.5 Predefined Macro Names

The following macro names are defined as follows:

- `__UPC__`

    The integer constant 1, indicating a conforming implementation.

- `__UPC_VERSION__`

    The integer constant 200310L.

- `__UPC_DYNAMIC_THREADS__`

    The integer constant 1 in the dynamic `THREADS` translation environment, otherwise undefined.

- `__UPC_STATIC_THREADS__`

    The integer constant 1 in the static `THREADS` translation environment, otherwise undefined.

6.6 Standard Headers

The standard headers are:

```
<upc_strict.h>
<upc_relaxed.h>
<upc.h>
```

`upc_strict.h` must contain at least:

```
#pragma _CRI upc strict
#include <upc.h>
```

`upc relaxed.h` must contain at least:

```
#pragma _CRI upc relaxed
#include <upc.h>
```

6.7 UPC Functions

Cray supports the following UPC functions. Cray-specific functions are noted in the function descriptions. See the UPC man pages for further information.

6.7.1 Termination of All Threads Function

6.7.1.1 upc_global_exit

The synopsis is:

```
upc_global_exit(int status);
```

upc_global_exit flushes all I/O, releases all memory, and terminates the execution for all active threads

6.7.2 Shared Memory Allocation Functions

The following sections describe the shared memory allocation functions.

6.7.2.1 upc_global_alloc

The synopsis is:

```
#include <upc.h>

   shared void *upc_global_alloc(size_t nblocks,
   size_t nbytes);
```

where *nblocks* is the number of blocks and *nbytes* is the block size in bytes.

upc_global_alloc allocates shared space compatible with the declaration:

shared [*nbytes*] char[*nblocks* * *nbytes*]

The upc_global_alloc function is not a collective function. If upc_global_alloc is called by multiple threads, all threads which make the call get different allocations. If nblocks*nbytes is zero, the result is a null pointer-to-shared.

⚠ **Caution:** `upc_global_alloc` must be used with MPT 2.3.0.1 and UNICOS/mp 2.4 to work correctly. If `upc_global_alloc` is executed with a previous version of MPT or UNICOS/mp, the function will issue a descriptive message and abort the application.

6.7.2.2 `upc_all_alloc`

The synopsis is:

```
#include <upc.h>

    shared void *upc_all_alloc(size_t nblocks, size_t nbytes);
```

`upc_all_alloc` is a collective function with single-valued arguments. `upc_all_alloc` allocates shared space compatible with the following declaration:

`shared [`*nbytes*`] char[`*nblocks* `* ` *nbytes*`]`

The `upc_all_alloc` function returns the same pointer value on all threads. If *nblocks**nbytes* is zero, the result is a null pointer-to-shared value.

The dynamic lifetime of an allocated object extends from the time any thread completes the call to `upc_all_alloc` until any thread has deallocated the object.

6.7.2.3 `upc_all_free`

The synopsis is:

```
#include <upc.h>

    void upc_all_free(shared void *ptr);
```

`upc_all_free` deallocates memory allocated by the `upc_all_alloc` function.

Note: This is a Cray-specific function.

6.7.2.4 upc_alloc

The synopsis is:

```
#include <upc.h>

    shared void *upc_alloc(size_t nbytes);
```

where *nbytes* is the total number of bytes to allocate.

The upc_alloc function allocates shared space of at least *nbytes* bytes with affinity to the calling thread.

upc_alloc is similar to malloc() except that it returns a pointer-to-shared value. It is not a collective function. If *nbytes* is zero, the result is a null pointer-to-shared.

6.7.2.5 upc_local_alloc

upc_local_alloc

The synopsis is:

```
#include <upc.h>

    shared void *upc_local_alloc(size_t nblocks, size_t nbytes);
```

The upc_local_alloc function is deprecated and should not be used. UPC programs should use the upc_alloc function instead. Support may be removed in future versions of the UPC specification.

upc_local_alloc allocates shared space of at least *nblocks * nbytes* with affinity to the calling thread. If *nblocks*nbytes* is zero, the result is a null pointer-to-shared value.

upc_local_alloc is similar to malloc() except that it returns a pointer-to shared value. It is not a collective function.

6.7.2.6 `upc_local_free`

The synopsis is:

```
#include <upc.h>

    void upc_local_free(shared void *ptr);
```

The `upc_local_free` function deallocates shared memory allocated by a call to either `upc_alloc` or `upc_local_alloc`. If the *ptr* argument does not point to memory that was allocated by either `upc_alloc` or `upc_local_alloc` or points to memory that was already deallocated, the behavior of the function is undefined. If the *ptr* argument is NULL, no action occurs.

Note that program termination does not imply that shared data allocated dynamically is freed.

Note: This is a Cray-specific function.

6.7.2.7 `upc_free`

The synopsis is:

```
#include <upc.h>

    void upc_free(shared void *ptr;
```

`upc_free` frees the dynamically allocated shared storage pointed to by *ptr*. If *ptr* is a null pointer, no action occurs. Otherwise, if the argument does not match a pointer earlier returned by the `upc_alloc`, `upc_global_alloc`, `upc_all_alloc`, or `upc_local_alloc` function, or if the space has been deallocated by a previous call by any thread to `upc_free`, the behavior is undefined.

⚠ **Caution:** `upc_free` must be used with MPT 2.3.0.1 and UNICOS/mp 2.4 to work correctly. If `upc_free` is executed with a previous version of MPT or UNICOS/mp, the function will issue a descriptive message and abort the application.

6.7.3 Pointer-to-shared Manipulation Functions

The following sections describe the pointer-to-shared manipulation functions.

6.7.3.1 `upc_threadof`

The synopsis is:

```
#include <upc.h>

    size_t upc_threadof(shared void *ptr);
```

`upc_threadof` returns the number of the thread that has affinity to the shared object pointed to by *ptr*

6.7.3.2 `upc_phaseof`

The synopsis is:

```
#include <upc.h>

    size_t upc_phaseof(shared void *ptr);
```

`upc_phaseof` returns the phase component of the pointer-to-shared argument

6.7.3.3 `upc_resetphase`

The synopsis is:

```
#include <upc.h>

    shared void *upc_resetphase(shared void *ptr);
```

`upc_resetphase` returns a pointer-to-shared value which is identical to its input except that it has zero phase.

6.7.3.4 `upc_addrfield`

The synopsis is:

```
#include <upc.h>

   size_t upc_addrfield(shared void *ptr);
```

`upc_addrfield` returns an implementation-defined value reflecting the local address of the object pointed to by the pointer-to-shared argument.

6.7.3.5 `upc_affinitysize`

The synopsis is:

```
#include <upc.h>

   size_t upc_affinitysize(size_t totalsize, size_t nbytes,
      size_t threadid);
```

where *totalsize* is the total size of the allocation in bytes, *nbytes* is the number of bytes in a block, and *threadid* is the thread whose affinity size is to be evaluated

`upc_affinitysize` is a convenience function which calculates the exact size of the local portion of the data in a shared object with affinity to a given thread.

In the case of a dynamically allocated shared object, the *totalsize* argument is *nbytes**nblocks and the *nbytes* argument is *nbytes,* where *nblocks* and *nbytes* are exactly as passed to `upc_global_alloc` or `upc_all_alloc` when the object was allocated.

In the case of a statically allocated shared object with declaration:

```
shared [b] t d[s];
```

the *totalsize* argument is `s * sizeof (t)` and the *nbytes* argument must be `b * sizeof (t)`. If block size is unspecified, *nbytes* must be 1. If the block size is indefinite, *nbytes* must be 0.

threadid must be a value in `0..(THREADS-1)`.

6.7.4 Lock Functions

The following sections describe the lock functions.

6.7.4.1 `upc_lock_t`

`upc_lock_t` is an opaque UPC type. `upc_lock_t` is a shared datatype with incomplete type. Objects of type `upc_lock_t` may therefore be manipulated only through pointers.

6.7.4.2 `upc_global_lock_alloc`

The synopsis is:

```
#include <upc.h>
   upc_lock_t *upc_global_lock_alloc(void);
```

`upc_global_lock_alloc` dynamically allocates a lock and returns a pointer to it. The lock is created in an unlocked state.

The `upc_global_lock_alloc` function is not a collective function. If called by multiple threads, all threads which make the call get different allocations.

⚠ **Caution:** `upc_global_lock_alloc` must be used with MPT 2.3.0.1 and UNICOS/mp 2.4 to work correctly. If `upc_global_lock_alloc` is executed with a previous version of MPT or UNICOS/mp, the function will issue a descriptive message and abort the application.

6.7.4.3 `upc_all_lock_alloc`

The synopsis is:

```
#include <upc.h
   upc_lock_t *upc_all_lock_alloc(void);
```

`upc_all_lock_alloc` dynamically allocates a lock and returns a pointer to it. The lock is created in an unlocked state.

The `upc_all_lock_alloc` function is a collective function. The return value on every thread points to the same lock object.

6.7.4.4 `upc_all_lock_free`

The synopsis is:

```
#include <upc.h>

   void upc_all_lock_free(upc_lock_t *ptr);
```

`upc_all_lock_free` frees a lock allocated by the `upc_all_lock_alloc` function.

Note: This is a Cray-specific function.

6.7.4.5 `upc_global_lock_free`

The synopsis is:

```
#include <upc.h>

   void upc_global_lock_free(upc_lock_t *ptr);
```

`upc_global_lock_free` frees a lock allocated by the `upc_global_lock_alloc` function.

The `upc_global_lock_free` function frees all resources associated with lock *ptr*, which was allocated by `upc_global_lock_alloc`. The `upc_global_lock_free` function will free *ptr* whether it is unlocked or locked by any thread. After *ptr* is freed, passing it to any locking functions in any thread will cause undefined behavior.

Only the thread that allocated lock *ptr* should free it. Be cautious when freeing the lock, because there is no implied synchronization with other threads.

If the *ptr* argument is a NULL pointer, the function does nothing. If *ptr* was not allocated by the `upc_global_lock_alloc` function or if it was freed earlier, the behavior of `upc_global_lock_free` will be undefined.

Note: This is a Cray-specific function.

6.7.4.6 `upc_lock_free`

The synopsis is:

```
#include <upc.h>

   void upc_lock_free(upc_lock_t *ptr);
```

`upc_lock_free` frees all resources associated with the dynamically allocated `upc_lock_t` pointed to by *ptr*. If *ptr* is a null pointer, no action occurs. Otherwise, if the argument does not match a pointer earlier returned by the `upc_global_lock_alloc` or `upc_all_lock_alloc` function, or if the lock has been deallocated by a previous call to `upc_lock_free`, the behavior is undefined.

`upc_lock_free` succeeds regardless of whether the referenced lock is currently unlocked or currently locked (by any thread).

Any subsequent calls to locking functions from any threads using *ptr* have undefined effects.

⚠ **Caution:** `upc_lock_free` must be used with MPT 2.3.0.1 and UNICOS/mp 2.4 to work correctly. If `upc_lock_free` is executed with a previous version of MPT or UNICOS/mp, the function will issue a descriptive message and abort the application.

6.7.4.7 `upc_lock`

The synopsis is:

```
#include <upc.h>

   void upc_lock(upc_lock_t *ptr);
```

`upc_lock` locks a shared variable, of type `upc_lock_t`, pointed to by the pointer given as argument.

If the lock is not used by another thread, then the thread making the call gets the lock and the function returns. Otherwise, the function keeps trying to get access to the lock.

A null strict reference is implied after a call to `upc_lock()`.

If the calling thread is already holding the lock referenced by *ptr* (i.e., it has previously locked it using `upc_lock()` or `upc_lock_attempt()`, but not unlocked it), the result is undefined.

6.7.4.8 `upc_lock_attempt`

The synopsis is:

```
#include <upc.h>

    int upc_lock_attempt(upc_lock_t *ptr);
```

`upc_lock_attempt` tries to lock a shared variable, of type `upc_lock_t`, pointed to by the pointer given as argument.

If the lock is not used by another thread, then the thread making the call gets the lock and the function returns 1. Otherwise, the function returns 0.

A null strict reference is implied after a call to `upc_lock_attempt()` that returns 1.

If the calling thread is already holding the lock referenced by *ptr* (i.e., it has previously locked it using `upc_lock()` or `upc_lock_attempt()`, but not unlocked it), the result is undefined.

6.7.4.9 `upc_unlock`

The synopsis is:

```
#include <upc.h>

    void upc_unlock(upc_lock_t *ptr);
```

`upc_unlock` frees the lock and does not return any value.

A null strict reference is implied before a call to `upc_unlock()`.

6.7.5 Shared String Handling Functions

The following sections describe the shared string handling functions.

6.7.5.1 upc_memcpy

The synopsis is:

```
#include <upc.h>

   void upc_memcpy(shared void *dst,
   shared const void *src,
   size_t n);
```

upc_memcpy copies *n* characters from a shared object having affinity with one thread to a shared object having affinity with the same or another thread. If copying takes place between objects that overlap, the behavior is undefined.

The upc_memcpy function treats the *dst* and *src* pointers as if they had type:

shared [] char[*n*]

The effect is equivalent to copying the entire contents from one shared array object with this type (the *src* array) to another shared array object with this type (the *dst* array).

6.7.5.2 upc_memget

The synopsis is:

```
#include <upc.h>

   void upc_memget(void *dst,
   shared const void *src,
   size_t n);
```

upc_memget copies *n* characters from a shared object with affinity to any single thread to a private object on the calling thread. If copying takes place between objects that overlap, the behavior is undefined.

The upc_memget function treats the *src* pointer as if it had type:

shared [] char[*n*]

The effect is equivalent to copying the entire contents from one shared array object with this type (the *src* array) to a private array object (the *dst* array) declared with the type:

char[*n*]

6.7.5.3 upc_memput

The synopsis is:

```
#include <upc.h>

   void upc_memput(shared void *dst,
   const void *src,
   size_t n);
```

upc_memput copies *n* characters from the private object on the calling thread to a shared object with affinity to any single thread. If copying takes place between objects that overlap, the behavior is undefined.

The upc_memput function is equivalent to copying the entire contents from a private array object (the *src* array) declared with the type:

char[*n*]

to a shared array object (the *dst* array) with the type:

shared [] char[*n*]

6.7.5.4 upc_memset

The synopsis is:

```
#include <upc.h>

   void upc_memset(shared void *dst,
   int c, size_t n);
```

upc_memset copies the value of *c*, converted to an unsigned char, to a shared memory object with affinity to any single thread. The number of bytes set is *n*.

The upc_memset function treats the *dst* pointer as if had type:

```
shared [] char[n]
```

The effect is equivalent to setting the entire contents of a shared array object with this type (the *dst* array) to the value *c*.

6.8 Cray Implementation Differences

There is a false sharing hazard when referencing shared `char` and `short` integers on Cray X1 series systems.

On Cray X1 series systems, if two PEs store a `char` or `short` to the same 32-bit word in memory without synchronization, incorrect results can occur. This is because these stores are implemented by reading the entire 32-bit word, inserting the `char` or `short` value and writing the entire word back to memory.

In the situation described above, it is possible for one PE's store to be lost. Imagine two PEs writing two different characters into the same word in memory without synchronization:

```
                   Register        Memory
Initial Value                      0x0000
PE 0 Reads         0x0000          0x0000
PE 1 Reads         0x0000          0x0000
PE 0 Inserts 3     0x3000          0x0000
PE 1 Inserts 7     0x0700          0x0000
PE 0 Writes        0x3000          0x3000
PE 1 Writes        0x0700          0x0700
```

Notice that the value stored by PE 0 has been lost. The final value intended was `0x3700`. This situation is referred to as false sharing. It is the result of supporting data types that are smaller than the smallest type that can be individually read or written by the hardware. On Cray X1 series systems, UPC programmers must take care when storing to shared `char` and `short` data that this situation does not occur.

6.9 Compiling and Executing UPC Code

In order to compile UPC code, you must load the programming environment module (`PrgEnv`) and specify the `-h upc` option on the `cc`, `c89`, or `c99` command line.

The -X *npes* option can optionally be used to define the number of threads to use and statically set the value of the THREADS constant.

This example enables UPC and allows the THREADS symbol to be defined dynamically for the examp1 application:

```
% cc -h upc -o multupc examp1.c
```

This example enables UPC and statically defines the THREADS symbol as 15 for the examp1 application:

```
% cc -h upc -X 15 -o multupc examp1.c
```

The processing elements specified by *npes* are either MSPs or SSPs. To run programs on SSPs, you must specify the -h ssp compiler option. The default is to run on MSPs. See Section 2.10.10, page 27 for more information about using UPC in SSP mode.

After compiling the UPC code, you run the program using the aprun command when the code contains UPC code only, or a mixture of UPC and SHMEM, and/or CAF code. If the code has a mixture of UPC and MPI code, use the mpirun command.

If you use the -X *npes* compiler option, you must specify the same number of threads in the aprun command.

> **Note:** For more information about improving UPC code performance, refer to *Optimizing Applications on Cray X1 Series Systems.*

Cray C++ Libraries [7]

The Cray C++ compiler together with the Dinkum C++ Libraries support the C++ 98 standard (ISO/IEC FDIS 14882) and continues to support existing Cray extensions. Most of the standard C++ features are supported, except for the few mentioned in Section 7.1. The Dinkum C++ Library is described in Section 7.2.

For information about C++ language conformance and exceptions, refer to Appendix D, page 203.

7.1 Unsupported Standard C++ Library Features

The Cray C++ compiler supports the C++ standard except for wide characters and multiple locales as follows:

- String classes using basic string class templates with wide character types or that use the `wstring` standard template class

- I/O streams using wide character objects

- File-based streams using file streams with wide character types (`wfilebuf`, `wifstream`, `wofstream`, and `wfstream`)

- Multiple localization libraries; Cray C++ supports only one locale

  Note: The C++ standard provides a standard naming convention for library routines. Therefore, classes or routines that use wide characters are named appropriately. For example, the `fscanf` and `sprintf` functions do not use wide characters, but the `fwscanf` and `swprintf` function do.

7.2 Dinkum C++ Libraries

The Cray C++ compiler uses the Dinkum C++ libraries, which support standard C++. The Dinkum C++ Library documentation is provided in HTML through CrayDoc. You can also find other references to tutorials and advanced user materials for the standard C++ library in the preface of this document.

Cray C++ Template Instantiation [8]

A *template* describes a class or function that is a model for a family of related classes or functions. The act of generating a class or function from a template is called *template instantiation*.

For example, a template can be created for a stack class, and then a stack of integers, a stack of floats, and a stack of some user-defined type can be used. In source code, these might be written as `Stack<int>`, `Stack<float>`, and `Stack<X>`. From a single source description of the template for a stack, the compiler can create instantiations of the template for each of the types required.

The instantiation of a class template is always done as soon as it is needed during a compilation. However, the instantiations of template functions, member functions of template classes, and static data members of template classes (template entities) are not necessarily done immediately for the following reasons:

- The preferred end result is one copy of each instantiated entity across all object files in a program. This applies to entities with external linkage.

- A specialization of a template entity is allowed. For example, a specific version of `Stack<int>`, or of just `Stack<int>::push` could be written to replace the template-generated version and to provide a more efficient representation for a particular data type.

- If a template function is not referenced, it should not be compiled because such functions could contain semantic errors that would prevent compilation. Therefore, a reference to a template class should not automatically instantiate all the member functions of that class.

The goal of an instantiation mode is to provide trouble-free instantiation. The programmer should be able to compile source files to object code, link them and run the resulting program, without questioning how the necessary instantiations are done.

In practice, this is difficult for a compiler to do, and different compilers use different instantiation schemes with different strengths and weaknesses.

The Cray C++ compiler requires a normal, top-level, explicitly compiled source file that contains the definition of both the template entity and of any types required for the particular instantiation. This requirement is met in one of the following ways:

- Each `.h` file that declares a template entity also contains either the definition of the entity or includes another file containing the definition.

- When the compiler sees a template declaration in a `.h` file and discovers a need to instantiate that entity, implicit inclusion gives the compiler permission to search for an associated definition file having the same base name and a different suffix and implicitly include that file at the end of the compilation (see Section 8.6, page 151).

- The programmer makes sure that the files that define template entities also have the definitions of all the available types and adds code or directives in those files to request instantiation of those entities.

The Cray C++ compiler provides two instantiation mechanisms—simple instantiation and prelinker instantiation. These mechanisms perform template instantiation and provide command line options and `#pragma` directives that give the programmer more explicit control over instantiation.

8.1 Simple Instantiation

The goal of the simple instantiation mode is to provide a method of instantiating templates without the need to create and manage intermediate (`*.ti` and `*.ii)` files.

The Cray C++ compilers accomplishes simple instantiation as follows:

1. When the source files of a program are compiled using the `-h simple_templates` option, each of the `*.o` files contains a copy of all of the template instantiations it uses.

2. When the object files are linked together, the resulting executable file contains multiple copies of the template function.

Unlike in prelinker instantiation, no `*.ti` or `*.ii` files are created. The programmer is not required to manage the naming and location of the intermediate files.

The simple template instantiation process creates slightly larger object files and a slightly larger executable file than is the case for prelinker instantiation.

For example, you have three C++ source files, `x.C`, `y.C`, and `z.C`. The source files reference a template `sortall` that sorts `int`, `float`, and `char` array elements:

```
template <class X> void sortall(X a[])
{
  ... code to sort int, float, char elements ...
}
```

Entering the command `CC -c -h simple_templates x.C y.C z.C` produces object files `x.o`, `y.o`, and `z.o`. Each `*.o` file has three copies of `sortall`, one for `int`s, one for `float`s, and one for `char`s.

Then, entering the command `CC x.o y.o z.o` links the files and any needed library routines, creating `a.out`.

Because the `-h simple_templates` option enables the `-h instantiate=used` option, all needed template entities are instantiated. The programmer can use the `#pragma do_not_instantiate` directive in programs compiled using the `-h simple_templates` option. See Section 3.6, page 79.

8.2 Prelinker Instantiation

In prelinker mode, automatic instantiation is accomplished by the Cray C++ compiler as follows:

1. If the compiler is responsible for doing all instantiations automatically, it can only do so for the entire program. That is, the compiler cannot make decisions about instantiation of template entities until all source files of the complete program have been read.

2. The first time the source files of a program are compiled, no template entities are instantiated. However, the generated object files contain information about things that could have been instantiated in each compilation. For any source file that makes use of a template instantiation, an associated `.ti` file is created, if one does not already exist (for example, the compilation of `abc.C` results in the creation of `abc.ti`).

3. When the object files are linked together, a program called the *prelinker* is run. It examines the object files, looking for references and definitions of template entities and for any additional information about entities that could be instantiated.

⚠️ **Caution:** The prelinker examines the object files in a library (.a) file but, because it does not modify them, is not able to assign template instantiations to them.

4. If the prelinker finds a reference to a template entity for which there is no definition in the set of object files, it looks for a file that indicates that it could instantiate that template entity. Upon discovery of such a file, it assigns the instantiation to that file. The set of instantiations assigned to a given file (for example, abc.C) is recorded in an associated file that has a .ii suffix (for example, abc.ii).

5. The prelinker then executes the compiler to again recompile each file for which the .ii was changed.

6. During compilation, the compiler obeys the instantiation requests contained in the associated .ii file and produces a new object file that contains the requested template entities and the other things that were already in the object file.

7. The prelinker repeats steps 3 through 5 until there are no more instantiations to be adjusted.

8. The object files are linked together.

Once the program has been linked correctly, the .ii files contain a complete set of instantiation assignments. If source files are recompiled, the compiler consults the .ii files and does the indicated instantiations as it does the normal compilations. That means that, except in cases where the set of required instantiations changes, the prelink step from then on will find that all the necessary instantiations are present in the object files and no instantiation assignment adjustments need be done. This is true even if the entire program is recompiled. Because the .ii file contains information on how to recompile when instantiating, it is important that the .o and .ii files are not moved between the first compilation and linkage.

The prelinker cannot instantiate into and from library files (.a), so if a library is to be shared by many applications its templates should be expanded. You may find that creating a directory of objects with corresponding .ii files and the use of -h prelink_copy_if_nonlocal (see Section 2.7.9, page 21) will work as if you created a library (.a) that is shared.

The -h prelink_local_copy option indicates that only local files (for example, files in the current directory) are candidates for assignment of instantiations. This option is useful when you are sharing some common relocatables but do not want them updated. Another way to ensure that shared

.o files are not updated is to use the −h remove_instantiation_flags
option when compiling the shared .o files. This also makes smaller resulting
shared .o files.

An easy way to create a library that instantiates all references of templates within
the library is to create an empty main function and link it with the library, as
shown in the following example. The prelinker will instantiate those template
references that are within the library to one of the relocatables without generating
duplicates. The empty dummy_main.o file is removed prior to creating the
.a file.

```
% CC a.C b.C c.C dummy_main.C
% ar cr mylib.a a.o b.o c.o
```

Another alternative to creating a library that instantiates all references of
templates is to use the −h one_instantiation_per_object option. This
option directs the prelinker to instantiate each template referenced within a
library in its own object file. The following example shows how to use the option:

```
% CC -h one_instantiation_per_object a.C b.C c.C dummy_main.C
% ar cr mylib.a a.o b.o c.o myInstantiationsDir/*.int.o
```

For more information about this alternative see Section 8.4, page 149 and Section
2.7.3, page 19.

Prelinker instantiation can coexist with partial explicit control of instantiation
by the programmer through the use of #pragma directives or the
−h instantiate=*mode* option.

Prelinker instantiation mode can be disabled by issuing the
−h noautoinstantiate command line option. If prelinker instantiation is
disabled, the information about template entities that could be instantiated in a
file is not included in the object file.

8.3 Instantiation Modes

Normally, during compilation of a source file, no template entities are instantiated (except those assigned to the file by prelinker instantiation). However, the overall instantiation mode can be changed by issuing the -h instantiate=*mode* command line option. The *mode* argument can be specified as follows:

mode	Description
none	Do not automatically create instantiations of any template entities. This is the most appropriate mode when prelinker instantiation is enabled. This is the default instantiation mode.
used	Instantiate those template entities that were used in the compilation. This includes all static data members that have template definitions.
all	Instantiate all template entities declared or referenced in the compilation unit. For each fully instantiated template class, all of its member functions and static data members are instantiated, regardless of whether they were used. Nonmember template functions are instantiated even if the only reference was a declaration.
local	Similar to used mode, except that the functions are given internal linkage. This mode provides a simple mechanism for those who are not familiar with templates. The compiler instantiates the functions used in each compilation unit as local functions, and the program links and runs correctly (barring problems due to multiple copies of local static variables). This mode may generate multiple copies of the instantiated functions and is not suitable for production use. This mode cannot be used in conjunction with prelinker template instantiation. Prelinker instantiation is disabled by this mode.

In the case where the CC(1) command is given a single source file to compile and link, all instantiations are done in the single source file and, by default, the used mode is used and prelinker instantiation is suppressed.

8.4 One Instantiation Per Object File

You can direct the prelinker to instantiate each template referenced in the source into its own object file. This method is preferred over other template instantiation object file generation options because:

- The user of a library pulls in only the instantiations that are needed.

- Multiple libraries with the same template can link. If each instantiation is not placed in its own object file, linking a library with another library that also contains the same instantiations will generate warnings on some platforms.

Use the -h one_instantiation_per_object option to generate one object file per instantiation. For more information about this option, see Section 2.7.3, page 19.

8.5 Instantiation `#pragma` Directives

Instantiation #pragma directives can be used in source code to control the instantiation of specific template entities or sets of template entities. There are three instantiation #pragma directives:

- The #pragma _CRI instantiate directive causes a specified entity to be instantiated.

- The #pragma _CRI do_not_instantiate directive suppresses the instantiation of a specified entity. It is typically used to suppress the instantiation of an entity for which a specific definition is supplied.

- The #pragma _CRI can_instantiate directive indicates that a specified entity can be instantiated in the current compilation, but need not be. It is used in conjunction with prelinker instantiation to indicate potential sites for instantiation if the template entity is deemed to be required by the compiler.

The argument to the #pragma _CRI instantiate directive can be any of the following:

- A template class name. For example: A<int>

- A template class declaration. For example: class A<int>

- A member function name. For example: A<int>::f

- A static data member name. For example: A<int>::i

- A static data declaration. For example: int A<int>::i

- A member function declaration. For example: `void A<int>::f(int, char)`

- A template function declaration. For example: `char* f(int, float)`

A `#pragma` directive in which the argument is a template class name (for example, `A<int>` or `class A<int>`) is equivalent to repeating the directive for each member function and static data member declared in the class. When instantiating an entire class, a given member function or static data member may be excluded using the `#pragma _CRI do_not_instantiate` directive. For example:

```
#pragma _CRI instantiate A<int>
#pragma _CRI do_not_instantiate A<int>::f
```

The template definition of a template entity must be present in the compilation for an instantiation to occur. If an instantiation is explicitly requested by use of the `#pragma _CRI instantiate` directive and no template definition is available or a specific definition is provided, an error is issued.

The following example illustrates the use of the `#pragma _CRI instantiate` directive:

```
template <class T> void f1(T);  // No body provided
template <class T> void g1(T);  // No body provided
void f1(int) {}  // Specific definition
void main()
{
 int    i;
 double  d;
 f1(i);
 f1(d);
 g1(i);
 g1(d);
}
#pragma _CRI instantiate void f1(int) // error-specific definition
#pragma _CRI instantiate void g1(int) // error-no body provided
```

In the preceding example, `f1(double)` and `g1(double)` are not instantiated because no bodies are supplied, but no errors will be produced during the compilation. If no bodies are supplied at link time, a linker error is issued.

A member function name (such as `A<int>::f`) can be used as a `#pragma` directive argument only if it refers to a single, user-defined member function (that is, not an overloaded function). Compiler-generated functions are

not considered, so a name can refer to a user-defined constructor even if a compiler-generated copy constructor of the same name exists. Overloaded member functions can be instantiated by providing the complete member function declaration, as in the following example:

```
#pragma _CRI instantiate char* A<int>::f(int, char*)
```

The argument to an instantiation directive cannot be a compiler-generated function, an inline function, or a pure virtual function.

8.6 Implicit Inclusion

The implicit inclusion feature implies that if the compiler needs a definition to instantiate a template entity declared in a `.h` file, it can implicitly include the corresponding `.C` file to get the source code for the definition. For example, if a template entity `ABC::f` is declared in file `xyz.h`, and an instantiation of `ABC::f` is required in a compilation, but no definition of `ABC::f` appears in the source code processed by the compilation, the compiler will look to see if a file `xyz.C` exists and, if so, process it as if it were included at the end of the main source file.

To find the template definition file for a given template entity, the Cray C++ compiler must know the full path name to the file in which the template was declared and whether the file was included using the system include syntax (such as `#include <file.h>`). This information is not available for preprocessed source code containing `#line` directives. Consequently, the Cray C++ compiler does not attempt implicit inclusion for source code that contains `#line` directives.

The set of definition-file suffixes that are tried by default, is `.c`, `.C`, `.cpp`, `.CPP`, `.cxx`, `.CXX`, and `.cc`.

Implicit inclusion works well with prelinker instantiation; however, they are independent. They can be enabled or disabled independently, and implicit inclusion is still useful without prelinker instantiation.

Cray C Extensions [9]

The Cray C compiler supports these extensions, developed by Cray, to the C standard:

- Complex data extensions (Section 9.1, page 153)

- `fortran` keyword (Section 9.2, page 154)

- Hexadecimal floating-point constants (Section 9.3, page 154)

A program that uses one or more extensions does not strictly conform to the standard. These extensions are not available in strict conformance mode.

9.1 Complex Data Extensions

Cray C extends the complex data facilities defined by standard C with these extensions:

- Imaginary constants

- Incrementing or decrementing `_Complex` data

The Cray C compiler supports the Cray imaginary constant extension and is defined in the `<complex.h>` header file. This imaginary constant has the following form:

Ri

R is either a floating constant or an integer constant; no space or other character can appear between R and i. If you are compiling in strict conformance mode (`-h conform`), the Cray imaginary constants are not available.

The following example illustrates imaginary constants:

```
#include <complex.h>
double complex z1 = 1.2 + 3.4i;
double complex z2 = 5i;
```

The other extension to the complex data facility allows the prefix– and postfix-increment and decrement operators to be applied to the `_Complex` data type. The operations affect only the real portion of a complex number.

9.2 `fortran` Keyword

In extended mode, the identifier `fortran` is treated as a keyword. It specifies a
storage class that can be used to declare a Fortran-coded external function. The
use of the `fortran` keyword when declaring a function causes the compiler to
verify that the arguments used in each call to the function are pass by addresses;
any arguments that are not addresses are converted to addresses.

As in any function declaration, an optional *type-specifier* declares the type
returned, if any. Type `int` is the default; type `void` can be used if no value is
returned (by a Fortran subroutine). The `fortran` storage class causes conversion
of lowercase function names to uppercase, and, if the function name ends with
an underscore character, the trailing underscore character is stripped from the
function name. (Stripping the trailing underscore character is in keeping with
UNIX practice.)

Functions specified with a `fortran` storage class must not be declared elsewhere
in the file with a `static` storage class.

Note: The `fortran` keyword is not allowed in Cray C++.

An example using the `fortran` keyword is shown in Section 13.3.7, page 174.

9.3 Hexadecimal Floating-point Constants

The Cray C compiler supports the standard hexadecimal floating constant
notations and the Cray hexadecimal floating constant notation. The standard
hexadecimal floating constants are portable and have sizes that are dependent
upon the hardware. The remainder of this section discusses the Cray
hexadecimal floating constant.

The Cray hexadecimal floating constant feature is not portable, because identical
hexadecimal floating constants can have different meanings on different systems.
It can be used whenever traditional floating-point constants are allowed.

The hexadecimal constant has the usual syntax: `0x` (or `0X`) followed by
hexadecimal characters. The optional floating suffix has the same form as for
normal floating constants: `f` or `F` (for float), `l` or `L` (for long), optionally followed
by an `i` (imaginary).

The constant must represent the same number of bits as its type, which is
determined by the suffix (or the default of double). The constant's bit length is
four times the number of hexadecimal digits, including leading zeros.

The following example illustrates hexadecimal constant representation:

```
0x7f7fffff.f           32-bit float
0x0123456789012345.    64-bit double
```

The value of a hexadecimal floating constant is interpreted as a value in the specified floating type. This uses an unsigned integral type of the same size as the floating type, regardless of whether an object can be explicitly declared with such a type. No conversion or range checking is performed. The resulting floating value is defined in the same way as the result of accessing a member of floating type in a union after a value has been stored in a different member of integral type.

The following example illustrates hexadecimal floating-point constant representation that use Cray floating-point format:

```
int main(void)
{
  float f1, f2;
  double g1, g2;

  f1 = 0x3ec00000.f;
  f2 = 0x3fc00000.f;
  g1 = 0x40fa400100000000.;
  g2 = 0x40fa400200000000.;

  printf("f1 = %8.8g\n", f1);
  printf("f2 = %8.8g\n", f2);
  printf("g1 = %16.16g\n", g1);
  printf("g2 = %16.16g\n", g2);
  return 1;
}
```

This is the output for the previous example:

```
f1 =    0.375
f2 =    1.5
g1 =    107520.0625
g2 =    107520.125
```

Predefined Macros [10]

Predefined macros can be divided into the following categories:

- Macros required by the C and C++ standards (Section 10.1, page 157)

- Macros based on the host machine (Section 10.2, page 158)

- Macros based on the target machine (Section 10.3, page 158)

- Macros based on the compiler (Section 10.4, page 159)

- UPC macros (Section 10.5, page 159)

Predefined macros provide information about the compilation environment. In this chapter, only those macros that begin with the underscore (_) character are defined when running in strict-conformance mode (see the `-h conform` command line option in Section 2.6.2, page 13).

Note: Any of the predefined macros except those required by the standard (see Section 10.1, page 157) can be undefined by using the `-U` command line option; they can also be redefined by using the `-D` command line option.

A large set of macros is also defined in the standard header files.

10.1 Macros Required by the C and C++ Standards

The following macros are required by the C and C++ standards:

Macro	Description
`__TIME__`	Time of translation of the source file.
`__DATE__`	Date of translation of the source file.
`__LINE__`	Line number of the current line in your source file.
`__FILE__`	Name of the source file being compiled.
`__STDC__`	Defined as the decimal constant 1 if compilation is in strict conformance mode; defined as the decimal constant 2 if the compilation is in extended mode. This macro is defined for Cray C and C++ compilations.

`__cplusplus`	Defined as 1 when compiling Cray C++ code and undefined when compiling Cray C code. The `__cplusplus` macro is required by the ISO C++ standard, but not the ISO C standard.

10.2 Macros Based on the Host Machine

The following macros provide information about the environment running on the host machine:

Macro	Description
`__unix`	Defined as 1 if the machine uses the UNIX OS.
`unix`	Defined as 1 if the machine uses the UNIX OS. This macro is not defined in strict-conformance mode.
`_UNICOSMP`	Defined as 1 if the operating system is UNICOS/mp. This macro is not defined in strict-conformance mode.

10.3 Macros Based on the Target Machine

The following macros provide information about the characteristics of the target machine:

Macro	Description
`_ADDR64`	Defined as 1 if the targeted CPU has 64-bit address registers; if the targeted CPU does not have 64-bit address registers, the macro is not defined.
`__sv`	Defined as 1 on all Cray X1 series systems.
`__sv2`	Defined as 1 and indicates that the current system is a Cray X1 series system.
`_CRAY`	Defined as 1 on UNICOS/mp systems.
`_CRAYIEEE`	Defined as 1 if the targeted CPU type uses IEEE floating-point format.
`_CRAYSV2`	Defined as 1 and indicates that the current system is a Cray X1 series system.
`__crayx1`	Defined as 1 and indicates that the current system is a Cray X1 series system.

_MAXVL	Defined as the maximum hardware vector length, which is 64.
cray	Defined as 1 on UNICOS/mp. This macro is not defined in strict-conformance mode.
CRAY	Defined as 1 on UNICOS/mp systems. This macro is not defined in strict-conformance mode.

10.4 Macros Based on the Compiler

The following macros provide information about compiler features:

Macro	Description
_RELEASE	Defined as the major release level of the compiler.
_RELEASE_MINOR	Defined as the minor release level of the compiler.
_RELEASE_STRING	Defined as a string that describes the version of the compiler.
_CRAYC	Defined as 1 to identify the Cray C and C++ compilers.

10.5 UPC Predefined Macros

The following macros provide information about UPC functions:

Macro	Description
__UPC__	
	The integer constant 1, indicating a conforming implementation.
__UPC_DYNAMIC_THREADS__	
	The integer constant 1 in the dynamic THREADS translation environment.
__UPC_STATIC_THREADS__	
	The integer constant 1 in the static THREADS translation environment.

Running C and C++ Applications [11]

Cray X1 series systems provide the following options for launching applications:

- Launching a single non-MPI application

- Launching a single MPI application

- Launching multiple interrelated applications

11.1 Launching a Single Non-MPI Application

Cray X1 series systems provide two methods of launching single, non-MPI applications. You can use the `aprun` command or the auto `aprun` method.

To launch an application via `aprun`, you enter the name of the executable and any other desired command line options. Refer to the `aprun`(1) man page for details.

For example, if you want to compile and run programs `prog1`, `prog2`, and `prog3` as application `trio`, you would enter the following command sequence:

```
% CC -c prog1.C prog2.C prog3.C
% CC -o trio prog1.o prog2.o prog3.o
% aprun ./trio
```

You could use the auto `aprun` feature to perform the same functions:

```
% CC -c prog1.C prog2.C prog3.C
% CC -o trio prog1.o prog2.o prog3.o
% ./trio
```

The `CRAY_AUTO_APRUN_OPTIONS` environment variable specifies options for the `aprun` command when the command is called automatically. See Section 2.24, page 59.

11.2 Launching a Single MPI Application

The process for launching a single MPI application is the same as for non-MPI applications except that you use the `mpirun` command instead of `aprun`. The `aprun`(1) man page also describes `mpirun` options.

For example, if you want to compile and run programs `mpiprog1`, `mpiprog2`, and `mpiprog3` as application `mpitrio`, you would enter the following command sequence:

```
% CC -c mpiprog1.C mpiprog2.C mpiprog3.C
% CC -o mpitrio mpiprog1.o mpiprog2.o mpiprog3.o
% mpirun ./trio
```

11.3 Multiple Program, Multiple Data (MPMD) Launch

Cray X1 series enable you to launch multiple interrelated applications with a single `aprun` or `mpirun` command. The applications must have the following characteristics:

- The applications can use MPI, SHMEM, or CAF to perform application-to-application communications. Using UPC for application-to-application communication is not supported.

- Within each application, the supported programming models are MPI, SHMEM, CAF, pthreads, and OpenMP.

- All applications must be of the same mode; that is, they must all be MSP-mode applications or all SSP-mode applications.

- If one or more of the applications in an MPMD job use a shared memory model (OpenMP or pthreads) and need a depth greater than the default of 1, then all of the applications will have the depth specified by the `aprun` or `mpirun -d` option, whether they need it or not.

To launch multiple applications with one command, you use `aprun` or `mpirun`.

For example, suppose you have created three MPI applications which contain CAF statements:

```
% CC -o multiabc a.o b.o c.o
% CC -o multijkl j.o k.o l.o
% CC -o multixyz x.o y.o z.o
```

and the number of processing elements required are 128 for `multiabc`, 16 for `multijkl`, and 4 for `multixyz`.

To launch all three applications simultaneously, you would enter:

```
% mpirun -np 128 multiabc : -np 16 multijkl : -np 4 multixyz
```

Debugging Cray C and C++ Code [12]

The Etnus TotalView symbolic debugger is available to help you debug C and C++ codes (refer to *Etnus TotalView Users Guide*). In addition, the Cray C and C++ compilers provide the following features to help you in debugging codes:

- The `-G` and `-g` compiler options provide symbol information about your source code for use by the Etnus TotalView debugger. For more information on these compiler options, see Section 2.16.1, page 39.

- The `-h [no]bounds` option and the `#pragma _CRI [no]bounds` directive let you check pointer and array references. The `-h [no]bounds` option is described in Section 2.16.2, page 40. The `#pragma _CRI [no]bounds` directive is described in Section 3.5.1, page 70.

- The `#pragma _CRI message` directive lets you add warning messages to sections of code where you suspect problems. The `#pragma _CRI message` directive is described in Section 3.5.3, page 73.

- The `#pragma _CRI [no]opt` directive lets you selectively isolate portions of your code to optimize, or to toggle optimization on and off in selected portions of your code. The `#pragma _CRI [no]opt` directive is described in Section 3.5.6, page 75.

12.1 Etnus TotalView Debugger

Some of the functions available in the TotalView debugger allow you to perform the following actions:

- Set and clear breakpoints, which can be conditional, at both the source code level and the assembly code level

- Examine core files

- Step through a program, including across function calls

- Reattach to the executable file after editing and recompiling

- Edit values of variables and memory locations

- Evaluate code fragments

12.2 Compiler Debugging Options

To use the TotalView debugger in debugging your code, you must first compile your code using one of the debugging options (`-g` or `-G`). These options are specified as follows:

- `-Gf`

  If you specify the `-Gf` debugging option, the TotalView debugger allows you to set breakpoints at function entry and exit and at labels.

- `-Gp`

  If you specify the `-Gp` debugging option, the TotalView debugger allows you to set breakpoints at function entry and exit, labels, and at places where execution control flow changes (for example, loops, `switch`, and `if...else` statements).

- `-Gn` or `-g`

  If you specify the `-Gn` or `-g` debugging option, the TotalView debugger allows you to set breakpoints at function entry and exit, labels, and executable statements. These options force all compiler optimizations to be disabled as if you had specified `-O0`.

Users of the Cray C and C++ compilers do not have to sacrifice run time performance to debug codes. Many compiler optimizations are inhibited by breakpoints generated for debugging. By specifying a higher debugging level, fewer breakpoints are generated and better optimization occurs.

However, consider the following cases in which optimization is affected by the `-Gp` and `-Gf` debugging options:

- Vectorization can be inhibited if a label exists within the vectorizable loop.

- Vectorization can be inhibited if the loop contains a nested block and the `-Gp` option is specified.

- When the `-Gp` option is specified, setting a breakpoint at the first statement in a vectorized loop allows you to stop and display at each vector iteration. However, setting a breakpoint at the first statement in an unrolled loop may not allow you to stop at each vector iteration.

Interlanguage Communication [13]

In some situations, it is necessary or advantageous to make calls to assembly or Fortran functions from C or C++ programs. This chapter describes how to make such calls. It also discusses calls to C and C++ functions from Fortran and assembly language. For additional information on interlanguage communication, see *Interlanguage Programming Conventions*. The calling sequence is described in detail on the `callseq`(3) man page.

The C and C++ compilers provide a mechanism for declaring external functions that are written in other languages. This allows you to write portions of an application in C, C++, Fortran, or assembly language. This can be useful in cases where the other languages provide performance advantages or utilities that are not available in C or C++.

This chapter describes how to call assembly language and Fortran programs from a C or C++ program. It also discusses the issues related to calling C or C++ programs from other languages.

13.1 Calls between C and C++ Functions

The following requirements must be considered when making calls between functions written in C and C++:

- In Cray C++, the `extern "C"` linkage is required when declaring an external function that is written in Cray C or when declaring a Cray C++ function that is to be called from Cray C. Normally the compiler will mangle function names to encode information about the function's prototype in the external name. This prevents direct access to these function names from a C function. The `extern "C"` keyword will prevent the compiler from performing name mangling.

- The program must be linked using the `CC` command.

- The program's main routine must be C or C++ code compiled with the `CC` command.

Objects can be shared between C and C++. There are some Cray C++ objects that are not accessible to Cray C functions (such as classes). The following object types can be shared directly:

- Integral and floating types.

- Structures and unions that are declared identically in C and C++. In order for structures and unions to be shared, they must be declared with identical members in the identical order.

- Arrays and pointers to the above types.

In the following example, a Cray C function (`C_add_func`) is called by the Cray C++ main program:

```
#include <iostream.h>

extern "C" int C_add_func(int, int);
int global_int = 123;

main()
{
  int res, i;

  cout << "Start C++ main" << endl;

   /* Call C function to add two integers and return result. */

  cout << "Call C C_add_func" << endl;
  res = C_add_func(10, 20);
  cout << "Result of C_add_func = " << res << endl;
  cout << "End C++ main << endl;
}
```

The Cray C function (`C_add_func`) is as follows:

```
#include <stdio.h>

extern int global_int;

int C_add_func(int p1, int p2)
{
    printf("\tStart C function C_add_func.\n");
    printf("\t\tp1      = %d\n", p1);
    printf("\t\tp2      = %d\n", p2);
    printf("\t\tglobal_int = %d\n", global_int);
    return p1 + p2;
}
```

The output from the execution of the calling sequence illustrated in the preceding example is as follows:

```
Start C++ main
Call C C_add_func
        Start C function C_add_func.
                p1      = 10
                p2      = 20
                global_int = 123
Result of C_add_func = 30
End C++ main
```

13.2 Calling Assembly Language Functions from a C or C++ Function

You can sometimes avoid bottlenecks in programs by rewriting parts of the program in assembly language, maximizing performance by selecting instructions to reduce machine cycles. When writing assembly language functions that will be called by C or C++ functions, use the standard UNICOS/mp program linkage macros. When using these macros, you do not need to know the specific registers used by the C or C++ program or by the calling sequence of the assembly coded routine.

In Cray C++, use extern "C" to declare the assembly language function.

(Deferred implementation) Support of Cray Assembly Language (CAL) Functions is deferred. The use of Cray Assembly Language (CAL) is described in the *Cray Assembly Language (CAL) for Cray X1 Systems Reference Manual.*

The ALLOC, DEFA, DEFS, ENTER, EXIT, and MXCALLEN macros can be used to define the calling list, A and S register use, temporary storage, and entry and exit points.

13.3 Calling Fortran Functions and Subroutines from a C or C++ Function

This subsection describes the following aspects of calling Fortran from C or C++. Topics include requirements and guidelines, argument passing, array storage, logical and character data, accessing named common, and accessing blank common.

13.3.1 Requirements

Keep the following points in mind when calling Fortran functions from C or C++:

- Fortran uses the call-by-address convention. C and C++ use the call-by-value convention, which means that only pointers should be passed to Fortran subprograms. See Section 13.3.2, page 168.

- Fortran arrays are in column-major order. C and C++ arrays are in row-major order. This indicates which dimension is indicated by the first value in an array element subscript. See Section 13.3.3, page 169.

- Single-dimension arrays of signed 32-bit integers and single dimension arrays of 32-bit floating-point numbers are the only aggregates that can be passed as parameters without changing the arrays.

- Fortran character pointers and character pointers from Cray C and C++ are incompatible. See Section 13.3.4, page 170.

- Fortran logical values and the Boolean values from C and C++ are not fully compatible. See Section 13.3.4, page 170.

- External C and C++ variables are stored in common blocks of the same name, making them readily accessible from Fortran programs if the C or C++ variable is in uppercase.

- When declaring Fortran functions or objects in C or C++, the name must be specified in all uppercase letters, digits, or underscore characters and consist of 31 or fewer characters.

- In Cray C, Fortran functions can be declared using the `fortran` keyword (see Section 9.2, page 154). The `fortran` keyword is not available in Cray C++. Instead, Fortran functions must be declared by specifying `extern "C"`.

13.3.2 Argument Passing

Because Fortran subroutines expect arguments to be passed by pointers rather than by value, C and C++ functions called from Fortran subroutines must pass pointers rather than values.

All argument passing in Cray C is strictly by value. To prepare for a function call between two Cray C functions, a copy is made of each actual argument. A function can change the values of its formal parameters, but these changes cannot affect the values of the actual arguments. It is possible, however, to pass a pointer. (All array arguments are passed by this method.) This capability is analogous to the Fortran method of passing arguments.

In addition to passing by value, Cray C++ also provides passing by reference.

13.3.3 Array Storage

C and C++ arrays are stored in memory in row-major order. Fortran arrays are stored in memory in column-major order. For example, the C or C++ array declaration int A[3][2] is stored in memory as:

A[0][0]	A[0][1]
A[1][0]	A[1][1]
A[2][0]	A[2][1]

The previously defined array is viewed linearly in memory as:

```
A[0][0] A[0][1] A[1][0] A[1][1] A[2][0] A[2][1]
```

The Fortran array declaration INTEGER A(3,2) is stored in memory as:

A(1,1)	A(2,1)	A(3,1)
A(1,2)	A(2,2)	A(3,2)

The previously defined array is viewed linearly in memory as:

```
A(1,1)  A(2,1)  A(3,1)  A(1,2)  A(2,2)  A(3,2)
```

When an array is shared between Cray C, C++, and Fortran, its dimensions are declared and referenced in C and C++ in the opposite order in which they are declared and referenced in Fortran. Arrays are zero-based in C and C++ and are one-based in Fortran, so in C and C++ you should subtract 1 from the array subscripts that you would normally use in Fortran.

For example, using the Fortran declaration of array A in the preceding example, the equivalent declaration in C or C++ is:

```
int a[2][3];
```

The following list shows how to access elements of the array from Fortran and from C or C++:

Fortran	C or C++
A(1,1)	A[0][0]
A(2,1)	A[0][1]
A(3,1)	A[0][2]
A(1,2)	A[1][0]
A(2,2)	A[1][1]
A(3,2)	A[1][2]

13.3.4 Logical and Character Data

Logical and character data need special treatment for calls between C or C++ and Fortran. Fortran has a character descriptor that is incompatible with a character pointer in C and C++. The techniques used to represent logical (Boolean) values also differ between Cray C, C++, and Fortran.

Mechanisms you can use to convert one type to the other are provided by the `fortran.h` header file and conversion macros shown in the following list:

Macro	Description
_btol	Conversion utility that converts a 0 to a Fortran logical `.FALSE.` and a nonzero value to a Fortran logical `.TRUE.`
_ltob	Conversion utility that converts a Fortran logical `.FALSE.` to a 0 and a Fortran logical `.TRUE.` to a 1.

13.3.5 Accessing Named Common from C and C++

The following example demonstrates how external C and C++ variables are accessible in Fortran named common blocks. It shows a C or C++ C function calling a Fortran subprogram, the associated Fortran subprogram, and the associated input and output.

In this example, the C or C++ structure ST is accessed in the Fortran subprogram as common block ST. The name of the structure and the Fortran common block must match. Note that this requires that the structure name be uppercase. The C and C++ C structure member names and the Fortran common block member names do not have to match, as is shown in this example.

The following Cray C main program calls the Fortran subprogram FCTN:

```c
#include <stdio.h>
struct
{
   int i;
   double a[10];
   long double d;
} ST;

main()
{
   int i;

   /* initialize struct ST */
   ST.i = 12345;

   for (i = 0; i < 10; i++)
      ST.a[i] = i;

   ST.d = 1234567890.1234567890L;

   /* print out the members of struct ST */
   printf("In C: ST.i = %d, ST.d = %20.10Lf\n", ST.i, ST.d);
   printf("In C: ST.a = ");
   for (i = 0; i < 10; i++)
      printf("%4.1f", ST.a[i]);
   printf("\n\n");

   /* call the fortran function */
   FCTN();
}
```

The following example is the Fortran subprogram FCTN called by the previous Cray C main program:

```fortran
C *********** Fortran subprogram (f.f): ***********

      SUBROUTINE FCTN

      COMMON /ST/STI, STA(10), STD
      INTEGER STI
      REAL STA
      DOUBLE PRECISION STD
```

```
      INTEGER I

      WRITE(6,100) STI, STD
 100 FORMAT ('IN FORTRAN: STI = ', I5, ', STD = ', D25.20)
      WRITE(6,200) (STA(I), I = 1,10)
 200 FORMAT ('IN FORTRAN: STA =', 10F4.1)
      END
```

The previous Cray C and Fortran examples are executed by the following commands, and they produce the output shown:

```
% cc -c c.c
% ftn -c f.f
% ftn c.o f.o
% ./a.out
 ST.i = 12345, ST.d = 1234567890.1234567890
 In C: ST.a =  0.0 1.0 2.0 3.0 4.0 5.0 6.0 7.0 8.0 9.0

 IN FORTRAN: STI = 12345, STD = .12345678901234567889D+10
 IN FORTRAN: STA = 0.0 1.0 2.0 3.0 4.0 5.0 6.0 7.0 8.0 9.0
```

13.3.6 Accessing Blank Common from C or C++

Fortran includes the concept of a common block. A *common block* is an area of memory that can be referenced by any program unit in a program. A *named common block* has a name specified in names of variables or arrays stored in the block. A *blank common block*, sometimes referred to as blank common, is declared in the same way, but without a name.

There is no way to access blank common from C or C++ similar to accessing a named common block. However, you can write a simple Fortran function to return the address of the first word in blank common to the C or C++ program and then use that as a pointer value to access blank common.

The following example shows how Fortran blank common can be accessed
using C or C++ source code:

```
#include <stdio.h>

struct st
{
  float a;
  float b[10];
} *ST;

#ifdef __cplusplus
  extern "C" struct st *MYCOMMON(void);
  extern "C" void FCTN(void);
#else
  fortran struct st *MYCOMMON(void);
  fortran void FCTN(void);
#endif

main()
{
    int i;

    ST = MYCOMMON();
    ST->a = 1.0;
    for (i = 0; i < 10; i++)
        ST->b[i] = i+2;
    printf("\n In C and C++\n");
    printf("    a = %5.1f\n", ST->a);
    printf("    b = ");
    for (i = 0; i < 10; i++)
        printf("%5.1f ", ST->b[i]);
    printf("\n\n");

    FCTN();
}
```

This Fortran source code accesses blank common and is accessed from the C or C++ source code in the preceding example:

```
SUBROUTINE FCTN
COMMON // STA,STB(10)
PRINT *, "IN FORTRAN"
PRINT *, "    STA = ",STA
PRINT *, "    STB = ",STB
STOP
END

FUNCTION MYCOMMON()
COMMON // A
MYCOMMON = LOC(A)
RETURN
END
```

This is the output of the previous C or C++ source code:

```
 a =  1.0
 b = 2.0   3.0   4.0   5.0   6.0   7.0   8.0   9.0  10.0  11.0
```

This is the output of the previous Fortran source code:

```
 STA = 1.
 STB = 2.,   3.,   4.,   5.,   6.,   7.,   8.,   9.,   10.,   11.
```

13.3.7 Cray C and Fortran Example

Here is an example of a Cray C function that calls a Fortran subprogram. The Fortran subprogram example follows the Cray C function example, and the input and output from this sequence follows the Fortran subprogram example.

Note: This example assumes that the Cray Fortran function is compiled with the `-s default32` option enabled. The examples will not work if the `-s default64` option is enabled.

```
/*                 C program (main.c):                    */

#include <stdio.h>
#include <string.h>
#include <fortran.h>

/* Declare prototype of the Fortran function. Note the last */
/* argument passes the length of the first argument. */
fortran double FTNFCTN (char *, int *, int);

double FLOAT1 = 1.6;
double FLOAT2;  /* Initialized in FTNFCTN */

main()
{
    int clogical, ftnlogical, cstringlen;
    double rtnval;
    char *cstring = "C Character String";


/* Convert clogical to its Fortran equivalent */
    clogical = 1;
    ftnlogical = _btol(clogical);

/* Print values of variables before call to Fortran function */
    printf(" In main: FLOAT1 = %g; FLOAT2 = %g\n",
        FLOAT1, FLOAT2);
    printf(" Calling FTNFCTN with arguments:\n");
    printf(" string = \"%s\"; logical = %d\n\n", cstring, clogical);
    cstringlen = strlen(cstring);
    rtnval = FTNFCTN(cstring, &ftnlogical, cstringlen);

/* Convert ftnlogical to its C equivalent */
    clogical = _ltob(&ftnlogical);

/* Print values of variables after call to Fortran function */
    printf(" Back in main: FTNFCTN returned %g\n", rtnval);
    printf(" and changed the two arguments:\n");
    printf(" string = \"%.*s\"; logical = %d\n",
    cstringlen, cstring, clogical);
}
```

```
C                    Fortran subprogram (ftnfctn.f):


      FUNCTION FTNFCTN(STR, LOG)

      REAL FTNFCTN
      CHARACTER*(*) STR
      LOGICAL LOG

      COMMON /FLOAT1/FLOAT1
      COMMON /FLOAT2/FLOAT2
      REAL FLOAT1, FLOAT2
      DATA FLOAT2/2.4/         ! FLOAT1 INITIALIZED IN MAIN

C     PRINT CURRENT STATE OF VARIABLES
      PRINT*, '     IN FTNFCTN: FLOAT1 = ', FLOAT1,
     1                         ';FLOAT2 = ', FLOAT2
      PRINT*, '     ARGUMENTS:    STR = "', STR, '"; LOG = ', LOG

C     CHANGE THE VALUES FOR STR(ING) AND LOG(ICAL)
      STR = 'New Fortran String'
      LOG = .FALSE.

      FTNFCTN = 123.4
      PRINT*, '    RETURNING FROM FTNFCTN WITH ', FTNFCTN
      PRINT*
      RETURN
      END
```

The previous Cray C function and Fortran subprogram are executed by the following commands and produce the following output:

```
% cc -c main.c
% ftn -c ftnfctn.f
% ftn main.o ftnfctn.o
% ./a.out
In main: FLOAT1 = 1.6;  FLOAT2 = 2.4
Calling FTNFCTN with arguments:
string = "C Character String"; logical = 1

IN FTNFCTN: FLOAT1 = 1.6; FLOAT2 = 2.4
ARGUMENTS:   STR = "C Character String"; LOG = T
RETURNING FROM FTNFCTN WITH 123.4

Back in main: FTNFCTN returned 123.4
and changed the two arguments:
string = "New Fortran String"; logical = 0
```

13.3.8 Calling a Fortran Program from a Cray C++ Program

The following example illustrates how a Fortran program can be called from a Cray C++ program:

```
#include <iostream.h>
extern "C" int FORTRAN_ADD_INTS(int *arg1, int &arg2);

main()
{
    int num1, num2, res;
    cout << "Start C++ main" << endl << endl;

    //Call FORTRAN function to add two integers and return result.
    //Note that the second argument is a reference parameter so
    //it is not necessary to take the address of the
    //variable num2.

    num1 = 10;
    num2 = 20;
```

```
        cout << "Before Call to FORTRAN_ADD_INTS" << endl;
        res = FORTRAN_ADD_INTS(&num1, num2);
        cout << "Result of FORTRAN Add = " << res << endl << endl;
        cout << "End C++ main" << endl;
}
```

The Fortran program that is called from the Cray C++ `main` function in the
preceding example is as follows:

```
INTEGER FUNCTION FORTRAN_ADD_INTS(Arg1, Arg2)
INTEGER Arg1, Arg2

PRINT *," FORTRAN_ADD_INTS, Arg1,Arg2 = ", Arg1, Arg2
FORTRAN_ADD_INTS = Arg1 + Arg2
END
```

The output from the execution of the preceding example is as follows:

```
Start C++ main

Before Call to FORTRAN_ADD_INTS
  FORTRAN_ADD_INTS, Arg1,Arg2 =  10,  20
Result of FORTRAN Add = 30

End C++ main
```

13.4 Calling a C or C++ Function from a Fortran or Assembly Language Program

A C or C++ function can be called from a Fortran or (Deferred implementation)
assembly language program. One of two methods can be used to call C functions
from Fortran: the C interoperability feature provided by the Fortran 2000 facility
or the method documented in this section. C interoperability provides a standard
portable interoperability mechanism for Fortran and C programs. Refer to
Fortran Language Reference Manual, Volume 2 for more information about C
interoperability. If you are using the method documented in this section to call C
functions from Fortran, keep in mind the information in Section 13.3, page 167.

When calling a Cray C++ function from a Fortran or (Deferred implementation)
assembly language program, observe the following rules:

- The Cray C++ function must be declared with `extern "C"` linkage.

- The program must be linked with the `CC`(1) command.

- The program's main routine must be C or C++ code compiled with the CC command.

The example that follows illustrates a Fortran program, main.f, that calls a Cray C function, ctctn.c. The Cray C function being called, the commands required, and the associated input and output are also included.

Note: This example assumes that the Cray Fortran program is compiled with the -s default32 option enabled. The examples will not work if the -s default64 option is enabled.

Example 5: Calling a C Function from a Fortran Program

Fortran program main.f source code:

```
C   Fortran program (main.f):

      PROGRAM MAIN

      REAL CFCTN
      COMMON /FLOAT1/FLOAT1
      COMMON /FLOAT2/FLOAT2
      REAL FLOAT1, FLOAT2
      DATA FLOAT1/1.6/ ! FLOAT2 INITIALIZED IN cfctn.c
      LOGICAL LOG
      CHARACTER*24 STR
      REAL RTNVAL

C  INITIALIZE VARIABLES STR(ING) AND LOG(ICAL)
      STR = 'Fortran Character String'
      LOG = .TRUE.

C  PRINT VALUES OF VARIABLES BEFORE CALL TO C FUNCTION
      PRINT*, 'In main.f: FLOAT1 = ', FLOAT1,
     1                '; FLOAT2 = ', FLOAT2
      PRINT*, 'Calling cfctn.c with these arguments: '
      PRINT*, 'LOG = ', LOG
      PRINT*, 'STR = ', STR

      RTNVAL = CFCTN(STR, LOG)

C  PRINT VALUES OF VARIABLES AFTER CALL TO C FUNCTION
      PRINT*, 'Back in main.f:: cfctn.c returned ', RTNVAL
      PRINT*, 'and changed the two arguments to: '
```

```
        PRINT*, 'LOG = ', LOG
        PRINT*, 'STR = ', STR

        END PROGRAM
```

Compile `main.f`, creating `main.o`:

> **ftn -c main.f**

C function `cfctn.c` source code:

```
/*   C function (cfctn.c)   */
#include <fortran.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

float FLOAT1;    /* Initialized in MAIN */
float FLOAT2 = 2.4;

/* The slen argument passes the length of string in str */
float cfctn_(char * str, int *log, int slen)
{
  int clog;
  float rtnval;
  char *cstring;

/* Convert log passed from Fortran MAIN */
/* into its C equivalent */
  cstring = malloc(slen+1);
  strncpy(cstring, str, slen);
  cstring[slen] = '\0';
  clog = _ltob(log);

/* Print the current state of the variables */
  printf(" In CFCTN: FLOAT1 = %.1f; FLOAT2 = %.1f\n",
         FLOAT1, FLOAT2);
  printf(" Arguments: str = '%s'; log = %d\n",
  cstring, clog);

/* Change the values for str and log */
  strncpy(str, "C Character String ", 24);
  *log = 0;
```

```
  rtnval = 123.4;
  printf(" Returning from CFCTN with %.1f\n\n", rtnval);
  return(rtnval);
}
```

Compile `cfctn.c`, creating `cfctn.o`:

> **cc -c cfctn.c**

Link `main.o` and `cfctn.o`, creating executable `interlang1`:

> **ftn -o interlang1 main.o cfctn.o**

Run program `interlang1`:

> **./interlang1**

Program output:

```
 In main.f: FLOAT1 =  1.60000002 ; FLOAT2 =  2.4000001
 Calling cfctn.c with these arguments:
 LOG =  T
 STR = Fortran Character String
 In CFCTN: FLOAT1 = 1.6; FLOAT2 = 2.4
 Arguments: str = 'Fortran Character String'; log = 1
 Returning from CFCTN with 123.4

 Back in main.f:: cfctn.c returned  123.400002
 and changed the two arguments to:
 LOG =  F
 STR = C Character String
```

Implementation-defined Behavior [14]

This chapter describes compiler behavior that is defined by the implementation according to the C and/or C++ standards. The standards require that the behavior of each particular implementation be documented.

The C and C++ standards define implementation-defined behavior as behavior, for a correct program construct and correct data, that depends on the characteristics of the implementation. The behavior of the Cray C and C++ compilers for these cases is summarized in this section.

14.1 Messages

All diagnostic messages issued by the compilers are reported through the UNICOS/mp message system. For information on messages issued by the compilers and for information about the UNICOS/mp message system, see Appendix E, page 219.

14.2 Environment

When `argc` and `argv` are used as parameters to the `main` function, the array members `argv[0]` through `argv[argc-1]` contain pointers to strings that are set by the command shell. The shell sets these arguments to the list of words on the command line used to invoke the compiler (the argument list). For further information on how the words in the argument list are formed, refer to the documentation on the shell in which you are running. For information on UNICOS/mp shells, see the `sh`(1) or `csh`(1) man page.

A third parameter, `char **envp`, provides access to environment variables. The value of the parameter is a pointer to the first element of an array of null-terminated strings, that matches the output of the `env`(1) command. The array of pointers is terminated by a null pointer.

The compiler does not distinguish between interactive devices and other, noninteractive devices. The library, however, may determine that `stdin`, `stdout`, and `stderr` (`cin`, `cout`, and `cerr` in Cray C++) refer to interactive devices and buffer them accordingly.

14.2.1 Identifiers

The *identifier* (as defined by the standards) is merely a sequence of letters and digits. Specific uses of identifiers are called *names.*

The Cray C compiler treats the first 255 characters of a name as significant, regardless of whether it is an internal or external name. The case of names, including external names, is significant. In Cray C++, all characters of a name are significant.

14.2.2 Types

Table 12 summarizes Cray C and C++ types and the characteristics of each type. *Representation* is the number of bits used to represent an object of that type. *Memory* is the number of storage bits that an object of that type occupies.

In the Cray C and C++ compilers, *size*, in the context of the `sizeof` operator, refers to the size allocated to store the operand in memory; it does not refer to representation, as specified in Table 12. Thus, the `sizeof` operator will return a size that is equal to the value in the *Memory* column of Table 12, page 184 divided by 8 (the number of bits in a byte).

Table 12. Data Type Mapping

Type	Representation Size and Memory Storage Size (bits)
`bool` (C++)	8
`_Bool` (C)	8
`char`	8
`wchar_t`	32
`short`	16
`int`	32
`long`	64
`long long`	64
`float`	32
`double`	64
`long double`	128
`float complex`	64 (each part is 32 bits)

Type	Representation Size and Memory Storage Size (bits)
double complex	128 (each part is 64 bits)
long double complex	256 (each part is 128 bits)
Pointers	64

Variables with 8-bit `char` or 16-bit `short` data types are fully vectorizable when used in one of the following operations within a vector context:

- Reads of 8-bit `char`s and 16-bit `short`s

- Writes to 8-bit `char`s and 16-bit `short`s, except arrays

- Use of 8- and 16-bit variables as targets in a reduction loop. For example, `c` is a 16-bit object in this program fragment:

```
int i;
short c;
int a[100];
c=0;

for (i=0;i<100;i++) {
  c = c + a[i];
}
```

Cray discourages the use of 8-bit `char`s and 16-bit `short`s in contexts other than those listed above because of performance penalties.

14.2.3 Characters

The full 8-bit ASCII code set can be used in source files. Characters not in the character set defined in the standard are permitted only within character constants, string literals, and comments. The `-h [no]calchars` option allows the use of the `@` sign and `$` sign in identifier names. For more information on the `-h [no]calchars` option, see Section 2.9.3, page 23.

A character consists of 8 bits. Up to 8 characters can be packed into a 64-bit word. A plain `char` type, one that is declared without a `signed` or `unsigned` keyword, is treated as an unsigned type.

Character constants and string literals can contain any characters defined in the 8-bit ASCII code set. The characters are represented in their full 8-bit form. A character constant can contain up to 8 characters. The integer value of a character

constant is the value of the characters packed into a word from left to right, with the result right-justified, as shown in the following table:

Table 13. Packed Characters

Character constant	Integer value
'a'	0x61
'ab'	0x6162

In a character constant or string literal, if an escape sequence is not recognized, the \ character that initiates the escape sequence is ignored, as shown in the following table:

Table 14. Unrecognizable Escape Sequences

Character constant	Integer value	Explanation
'\a'	0x7	Recognized as the ASCII BEL character
'\8'	0x38	Not recognized; ASCII value for 8
'\['	0x5b	Not recognized; ASCII value for [
'\c'	0x63	Not recognized; ASCII value for c

14.2.4 Wide Characters

Wide characters are treated as signed 64-bit integer types. Wide character constants cannot contain more than one multibyte character. Multibyte characters in wide character constants and wide string literals are converted to wide characters in the compiler by calling the mbtowc(3) function. The current locale in effect at the time of compilation determines the method by which mbtowc(3) converts multibyte characters to wide characters, and the shift states required for the encoding of multibyte characters in the source code. If a wide character, as converted from a multibyte character or as specified by an escape sequence, cannot be represented in the extended execution character set, it is truncated.

14.2.5 Integers

All integral values are represented in a twos complement format. For representation and memory storage requirements for integral types, see Table 12, page 184.

When an integer is converted to a shorter signed integer, and the value cannot be represented, the result is the truncated representation treated as a signed quantity. When an unsigned integer is converted to a signed integer of equal length, and the value cannot be represented, the result is the original representation treated as a signed quantity.

The bitwise operators (unary operator ~ and binary operators <<, >>, &, ^, and |) operate on signed integers in the same manner in which they operate on unsigned integers. The result of E1 >> E2, where E1 is a negative-valued signed integral value, is E1 right-shifted E2 bit positions; vacated bits are filled with 1s. This behavior can be modified by using the -h nosignedshifts option (see Section 2.9.4, page 24). Bits higher than the sixth bit are not ignored. Values higher than 31 cause the result to be 0 or all 1s for right shifts.

The result of the / operator is the largest integer less than or equal to the algebraic quotient when either operand is negative and the result is a nonnegative value. If the result is a negative value, it is the smallest integer greater than or equal to the algebraic quotient. The / operator behaves the same way in C and C++ as in Fortran.

The sign of the result of the percent (%) operator is the sign of the first operand.

Integer overflow is ignored. Because some integer arithmetic uses the floating-point instructions, floating-point overflow can occur during integer operations. Division by 0 and all floating-point exceptions, if not detected as an error by the compiler, can cause a run time abort.

14.2.6 Arrays and Pointers

An unsigned int value can hold the maximum size of an array. The type size_t is defined to be a typedef name for unsigned long in the headers: malloc.h, stddef.h, stdio.h, stdlib.h, string.h, and time.h. If more than one of these headers is included, only the first defines size_t.

A type int can hold the difference between two pointers to elements of the same array. The type ptrdiff_t is defined to be a typedef name for long in the header stddef.h.

If a pointer type's value is cast to a signed or unsigned `long int`, and then cast back to the original type's value, the two pointer values will compare equal.

Pointers on UNICOS/mp systems are byte pointers. Byte pointers use the same internal representation as integers; a byte pointer counts the numbers of bytes from the first address.

A pointer can be explicitly converted to any integral type large enough to hold it. The result will have the same bit pattern as the original pointer. Similarly, any value of integral type can be explicitly converted to a pointer. The resulting pointer will have the same bit pattern as the original integral type.

14.2.7 Registers

Use of the register storage class in the declaration of an object has no effect on whether the object is placed in a register. The compiler performs register assignment aggressively; that is, it automatically attempts to place as many variables as possible into registers.

14.2.8 Classes, Structures, Unions, Enumerations, and Bit Fields

Accessing a member of a union by using a member of a different type results in an attempt to interpret, without conversion, the representation of the value of the member as the representation of a value in the different type.

Members of a class or structure are packed into words from left to right. Padding is appended to a member to correctly align the following member, if necessary. Member alignment is based on the size of the member:

- For a member bit field of any size, alignment is any bit position that allows the member to fit entirely within a 64–bit word.

- For a member with a size less than 64 bits, alignment is the same as the size. For example, a `char` has a size and alignment of 8 bits; a `float` has a size and alignment of 32 bits.

- For a member with a size equal to or greater than 64 bits, alignment is 64 bits.

- For a member with array type, alignment is equal to the alignment of the element type.

A plain `int` type bit field is treated as an `signed int` bit field.

The values of an enumeration type are represented in the type `signed int` in C; they are a separate type in C++.

14.2.9 Qualifiers

When an object that has volatile-qualified type is accessed, it is simply a reference to the value of the object. If the value is not used, the reference need not result in a load of the value from memory.

14.2.10 Declarators

A maximum of 12 pointer, array, and/or function declarators are allowed to modify an arithmetic, structure, or union type.

14.2.11 Statements

The compiler has no fixed limit on the maximum number of case values allowed in a switch statement.

The Cray C++ compiler parses asm statements for correct syntax, but otherwise ignores them.

14.2.12 Exceptions

In Cray C++, when an exception is thrown, the memory for the temporary copy of the exception being thrown is allocated on the stack and a pointer to the allocated space is returned.

14.2.13 System Function Calls

See the exit(3) man page for a description of the form of the unsuccessful termination status that is returned from a call to exit(3).

14.3 Preprocessing

The value of a single-character constant in a constant expression that controls conditional inclusion matches the value of the same character in the execution character set. No such character constant has a negative value. For each, 'a' has the same value in the two contexts:

```
#if 'a' == 97
if ('a' == 97)
```

The -I option and the method for locating included source files is described in Section 2.19.4, page 46.

The source file character sequence in a #include directive must be a valid UNICOS/mp file name or path name. A #include directive may specify a file name by means of a macro, provided the macro expands into a source file character sequence delimited by double quotes or < and > delimiters, as follows:

```
#define myheader "./myheader.h"
#include myheader

#define STDIO <stdio.h>
#include STDIO
```

The macros __DATE__ and __TIME__ contain the date and time of the beginning of translation. For more information, see the description of the predefined macros in Chapter 10, page 157.

The #pragma directives are described in Chapter 3, page 67.

Possible Requirements for non-C99 Code [A]

In order to use C code, developed under previous C compilers of the Cray C++ Programming Environment, with the `c99` command, your code may require one or more of the following modifications:

- Include necessary header files for complete function prototyping.

- Add return statements to all non-`void` functions.

- Ensure that all strings in any macro that begins with an underscore are literals. These macros cannot contain other types of strings.

- Follow C99 conventions

Previous Cray C compilers did not require you to explicitly include header files in many situations because they allowed functions to be implicitly declared. In C99, functions cannot be implicitly declared.

Libraries and Loader [B]

This appendix describes the libraries that are available with the Cray C and C++ compilers and the loader (`ld`).

B.1 Cray C and C++ Libraries Current Programming Environments

Libraries that support Cray C and C++ are automatically available when you use the `CC`, `cc`, `c89`, or `c99` command to compile your programs. These commands automatically issue the appropriate directives to load the program with the appropriate functions. If your program strictly conforms to the C or C++ standards, you do not need to know library names and locations. If your program requires other libraries or if you want direct control over the loading process, more knowledge of the loader and libraries is necessary.

The Standard Template Library (STL) is a C++ library of container classes, algorithms, and iterators; it provides many of the basic algorithms and data structures of computer science. The STL is a generic library, meaning that its components are heavily parameterized: almost every component in the STL is a template. Be sure you have a complete understanding of templates and how they work before using them.

B.2 Loader

When you issue the `cc`(1), `CC`, `c89`, or `c99` command to invoke the compiler, and the program compiles without errors, the loader is called. Specifying the `-c` option on the command line produces relocatable object files without calling the loader. These relocatable object files can then be used as input to the loader command by specifying the file names on the appropriate loader command line.

For example, the following command line compiles a file called `target.c` and produces the relocatable object file called `target.o` in your current working directory:

```
cc -c target.c
```

You can then use file `target.o` as input to the loader or save the file to use with other relocatable object files to compile and create a linked executable file (`a.out` by default).

Because of the special code needed to handle templates, constructors, destructors, and other C++ language features, object files generated by using the `CC`(1) command should be linked using the `CC` command. To link C++ object files using the loader command (`ld`), the `-h keep=`*files* option (see Section 2.9.1, page 22) must be specified on the command line when compiling source files.

The `ld` command can be accessed by using one of the following methods:

- You can access the loader directly by using the `ld` command.

- You can let the `cc`, `CC`, `c89`, or `c99` command choose the loader. This method has the following advantages:

  - You do not need to know the loader command line interface.

  - You do not need to worry about the details of which libraries to load, or the order in which to load them.

  - When using `CC`, you need not worry about template instantiation requirements or about loading the compiler-generated static constructors and destructors.

You can control the operation of the loader with the `ld` command line options. Refer to the `ld`(1) man page.

Compatibility with Older C++ Code [C]

A key feature of the Cray C++ Programming Environment 5.x is the Standard C++ Library. C++ code developed under the C++ Programming Environment 4.2 release or earlier can still be used with Programming Environment release 5.x. If your code uses nonstandard Cray C++ header files, you can continue to use your code without modification by using the `CRAYOLDCPPLIB` environment variable. Another way to use your pre-4.x code with the current Programming Environment release is to make changes to your existing code. The following sections explain how to use either of these methods.

> **Note:** Other changes to your existing C++ code may be required because of differences between the Cray SV1 or Cray T3E systems and the Cray X1 series systems. Refer to the *Cray X1 User Environment Differences.*

C.1 Use of Nonstandard Cray C++ Header Files

The Cray C++ Programming Environment release continues to support some of the nonstandard Cray C++ header files. This allows pre-5.0 code that use these header files to be compiled without modification. These header files are available in the Standard C++ Library at the same location as they were in previous releases.

Here are the Cray nonstandard header files that can be used in Programming Environment 5.x:

* `common.h`

* `complex.h`

* `fstream.h`

* `generic.h`

* `iomanip.h`

* `iostream.h`

* `stdiostream.h`

* `stream.h`

* `strstream.h`

* `vector.h`

The nonstandard header files can be used when you set the CRAYOLDCPPLIB environment variable to a nonzero value. How to set the variable depends on the shell you are using. If you are using ksh or sh, set the variable as this example shows:

```
% export CRAYOLDCPPLIB=1
```

If you are using csh, set the variable as this example shows:

```
% setenv CRAYOLDCPPLIB 1
```

C.2 When to Update Your C++ Code

You are not required to modify your existing C++ codes in order to compile it with the Cray C++ compiler version 5.x, unless you wish to use the Standard C++ Library. One reason for migrating your code to the Standard C++ Library is that the nonstandard Cray C++ header files of Programming Environment 3.5 may not be supported by future versions of the Cray C++ compiler. Another reason for migrating is your C++ code may already contain support for the Standard C++ Library. Often, third-party code contains a configuration script that tests the features of the compiler and system before building a makefile. This script can determine whether the C++ compiler supports the Standard C++ Library.

You can use the following steps to migrate your C++ code:

1. Use the proper header files

2. Add namespace declarations

3. Reconcile header definition differences

4. Recompile all C++ files

C.2.1 Use the Proper Header Files

The first step in migrating your C++ code to use the Standard C++ Library is to ensure that it uses the correct Standard C++ Library header files. The following tables show each header file used by the C++ library version 3.5 and its likely corresponding header file in the current Standard C++ Library. The older header files do not always map directly to the new files. For example, most of the definitions of the Cray C++ version 3.5 STL alloc.h header file are contained in the Standard C++ Library header files memory and xmemory. Anomalies, such as this are noted in the tables.

The tables divide the header files into three groups:

- Run time support library header files

- Stream and class library header files

- Standard Template Library header files

The older header file used by the run time support library originated from Edison Design Group and perform functions such as exception handling and memory allocation and deallocation. Table 15 shows the old and new header files.

Table 15. Run time Support Library Header Files

Cray C++ 3.5 header file	Standard C++ library header file
`exception.h`	`exception`
`new.h`	`new`
`stdexcept.h`	`stdexcept`
`typeinfo.h`	`typeinfo`

The header files in the stream and class library originate from AT&T and define the I/O stream classes along the string, complex, and vector classes. Table 16 shows the old and new header files.

Table 16. Stream and Class Library Header Files

Cray C++ 3.5 header file	Standard C++ Library header file
`common.h`	No equivalent header file
`complex.h`	`complex`
`fstream.h`	`fstream`
`iomanip.h`	`iomanip`
`iostream.h`	`iostream`
`stdiostream.h`	`iosfwd`
`stream.h`	Not available
`strstream.h`	strstream
`vector.h`	`vector`

Note: The use of any of the stream and class library header files from Cray C++ Programming Environment 3.5 requires that you set the `CRAYOLDCPPLIB` environment variable. Refer to Section C.1, page 195.

Table 17 shows the old and new Standard Template Library (STL) header files.

Note: The older STL originated from Silicon Graphics Inc.

Table 17. Standard Template Library Header Files

Cray C++ 3.5 Header File	Standard C++ Header File
`algo.h`	`algorithm`
`algobase.h`	`algorightm`
`alloc.h`	`memory`
`bvector.h`	`vector`
`defalloc.h`[1]	Not available
`deque.h`	`deque`
`function.h`	`functional`
`hash_map.h`	`hash_map`
`hash_set.h`	`hash_set`
`hashtable.h`	`xhash`
`heap.h`	`algorithm`
`iterator.h`	`iterator`
`list.h`	`list`
`map.h`	`map`
`mstring.h`	`string`
`multimap.h`	`map`
`multiset.h`	`set`
`pair.h`	`pair`
`pthread_alloc.h`	No equivalent header file
`rope.h`	`rope`

---

[1] This header file was deprecated in the Cray C++ Programming Environment 3.5 release.

Cray C++ 3.5 Header File	Standard C++ Header File
`ropeimpl.h`	`rope`
`set.h`	`set`
`slist.h`	`slist`
`stack.h`	`stack`
`stl_config.h`	The Standard C++ Library does not need the STL configuration file.
`tempbuf.h`	`memory`
`tree.h`	`xtree`
`vector.h`	`vector`

C.2.2 Add Namespace Declarations

The second step in migrating to the Standard C++ Library is adding namespace declarations. Most classes of the Standard C++ Library are declared under the `std` namespace, so this usually requires that you add this statement to the existing code: `using namespace std`. For example, the following program returns an error when it is compiled with previous versions of the Standard C++ Library:

```
% cat hello.C
#include <iostream>

main() { cout << "hello world\n"; }

% CC hello.C

  CC-20 CC: ERROR File = hello.C, line = 2
  The identifier "cout" is undefined.
  main() { cout <<"hello world\n" ; }
  ^
  Total errors detected in hello.C: 1
```

When you add `using namespace std;` to the example program, it compiles without error:

```
% cat hello.C
    #include <iostream>
    using namespace std;
```

```
                    main() { cout << "hello world\n"; }

          % CC hello.C
          % ./a.out
          hello world
```

C.2.3 Reconcile Header Definition Differences

The most difficult process of migrating to the Standard C++ Library is reconciling the differences between the definitions of the Cray C++ version 3.5 header files and the Standard Cray C++ library header files. For example, the definitions for the complex class differs. In Cray C++ version 3.5, the complex class has real and imaginary components of type double. The Standard C++ Library defines the complex class as a template class, where the user defines the data type of the real and imaginary components.

For example, here is a program written with the Cray C++ version 3.5 header files:

```
% cat complex.C
    #include <iostream.h>
    #include <complex.h>

    main() {
       complex C(1.0, 2.0);
       cout << "C = " << C << endl;
    }

#env CRAYOLDCPPLIB=1 CC complex.C
#a.out
C = ( 1, 2)
```

An equivalent program that uses the Standard C++ Library appears as:

```
% cat complex.C
    #include <iostream>
    #include <complex.h>
    using namespace std;

    main() {
       complex<double> C(1.0, 2.0);

       cout << "C = " << C << endl;
    }
```

```
% CC complex.C
% a.out
C = (1,2)
```

C.2.4 Recompile All C++ Files

Finally, when all of the source files that use the Standard C++ Library header files can be built, you must recompile all C++ source files that belong to the program using only the Standard C++ Library.

Cray C and C++ Dialects [D]

This appendix details the features of the C and C++ languages that are accepted by the Cray C and C++ compilers, including certain language dialects and anachronisms. Users should be aware of these details, especially users who are porting codes from other environments.

D.1 C++ Language Conformance

The Cray C++ compiler accepts the C++ language as defined by the *ISO/IEC 14882:1998* standard, with the exceptions listed in Section D.1.1, page 203.

The Cray C++ compiler also has a `cfront` compatibility mode, which duplicates a number of features and bugs of `cfront`. Complete compatibility is not guaranteed or intended. The mode allows programmers who have used `cfront` features to continue to compile their existing code (see Section 3.5, page 70). Command line options are also available to enable and disable anachronisms (see Section D.2, page 207) and strict standard-conformance checking (see Section D.3, page 208, and Section D.4, page 209). The command line options are described in Chapter 2, page 7.

D.1.1 Unsupported and Supported C++ Language Features

The `export` keyword for templates is not supported. It is defined in the *ISO/IEC 14882:1998* standard, but is not in traditional C++.

The following features, which are in the ISO/IEC 14882:1998 standard but not in traditional C++[1], are supported:

- The dependent statement of an `if`, `while`, `do-while`, or `for` is considered to be a scope, and the restriction on having such a dependent statement be a declaration is removed.

- The expression tested in an `if`, `while`, `do-while`, or `for`, as the first operand of a `?` operator, or as an operand of the `&&`, `||`, or `!` operators may have a pointer-to-member type or a class type that can be converted to a pointer-to-member type in addition to the scalar cases permitted by the ARM.

- Qualified names are allowed in elaborated type specifiers.

---

[1] As defined in *The Annotated C++ Reference Manual* (ARM), by Ellis and Stroustrup, Addison Wesley, 1990.

- A global-scope qualifier is allowed in member references of the form `x.::A::B` and `p->::A::B`.

- The precedence of the third operand of the `?` operator is changed.

- If control reaches the end of the `main()` routine, and the `main()` routine has an integral return type, it is treated as if a `return 0;` statement was executed.

- Pointers to arrays with unknown bounds as parameter types are diagnosed as errors.

- A functional-notation cast of the form `A()` can be used even if `A` is a class without a (nontrivial) constructor. The temporary that is created gets the same default initialization to zero as a static object of the class type.

- A cast can be used to select one out of a set of overloaded functions when taking the address of a function.

- Template friend declarations and definitions are permitted in class definitions and class template definitions.

- Type template parameters are permitted to have default arguments.

- Function templates may have nontype template parameters.

- A reference to `const volatile` cannot be bound to an `rvalue`.

- Qualification conversions such as conversion from `T**` to `T const * const` are allowed.

- Digraphs are recognized.

- Operator keywords (for example, `and` or `bitand`) are recognized.

- Static data member declarations can be used to declare member constants.

- `wchar_t` is recognized as a keyword and a distinct type.

- `bool` is recognized.

- RTTI (run time type identification), including `dynamic_cast` and the `typeid` operator, is implemented.

- Declarations in tested conditions (within `if`, `switch`, `for`, and `while` statements) are supported.

- Array `new` and `delete` are implemented.

- New-style casts (`static_cast`, `reinterpret_cast`, and `const_cast`) are implemented.

- Definition of a nested class outside its enclosing class is allowed.

- `mutable` is accepted on nonstatic data member declarations.

- Namespaces are implemented, including using declarations and directives. Access declarations are broadened to match the corresponding using declarations.

- Explicit instantiation of templates is implemented.

- The `typename` keyword is recognized.

- `explicit` is accepted to declare nonconverting constructors.

- The scope of a variable declared in the `for-init-statement` of a `for` loop is the scope of the loop (not the surrounding scope).

- Member templates are implemented.

- The new specialization syntax (using `template <>`) is implemented.

- Cv qualifiers are retained on `rvalues` (in particular, on function return values).

- The distinction between trivial and nontrivial constructors has been implemented, as has the distinction between process overlay directives (PODs) and non-PODs with trivial constructors.

- The linkage specification is treated as part of the function type (affecting function overloading and implicit conversions).

- A `typedef` name can be used in an explicit destructor call.

- Placement delete is supported.

- An array allocated via a placement new can be deallocated via delete.

- `enum` types are considered to be nonintegral types.

- Partial specification of class templates is implemented.

- Partial ordering of function templates is implemented.

- Function declarations that match a function template are regarded as independent functions, not as "guiding declarations" that are instances of the template.

- It is possible to overload operators using functions that take `enum` types and no class types.

- Explicit specification of function template arguments is supported.

- Unnamed template parameters are supported.

- The new lookup rules for member references of the form `x.A::B` and `p->A::B` are supported.

- The notation `:: template` (and `->template`, etc.) is supported.

- In a reference of the form `f()->g()`, with `g` a static member function, `f()` is evaluated. Likewise for a similar reference to a static data member. The ARM specifies that the left operand is not evaluated in such cases.

- `enum` types can contain values larger than can be contained in an `int`.

- Default arguments of function templates and member functions of class templates are instantiated only when the default argument is used in a call.

- String literals and wide string literals have `const` type.

- Class name injection is implemented.

- Argument-dependent (Koenig) lookup of function names is implemented.

- Class and function names declared only in unqualified friend declarations are not visible except for functions found by argument-dependent lookup.

- A `void` expression can be specified on a return statement in a `void` function.

- `reinterpret_cast` allows casting a pointer to a member of one class to a pointer to a member of another class even when the classes are unrelated.

- Two-phase name binding in templates as described in the *Working Paper* is implemented.

- Putting a `try/catch` around the initializers and body of a constructor is implemented.

- Template `template` parameters are implemented.

- Universal character set escapes (e.g., `\uabcd`) are implemented.

- `extern inline` functions are supported.

- Covariant return types on overriding virtual functions are supported.

D.2 C++ Anachronisms Accepted

C++ anachronisms are enabled by using the `-h anachronisms` command line option (see Section 2.6.7, page 14). When anachronisms are enabled, the following anachronisms are accepted:

- `overload` is allowed in function declarations. It is accepted and ignored.

- Definitions are not required for static data members that can be initialized by using the 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 can 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 can be omitted in a base class initializer if there is only one immediate base class.

- Assignment to the `this` pointer in constructors and destructors is allowed. This is only allowed 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 nonnested class name if 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 can participate in function overloading as though it were prototyped. Default argument promotion is not applied to parameter types of such functions when checking for compatibility, therefore, the following statements declare the overloading of two functions named f:

```
int f(int);

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

Note: In C, this code is legal, but has a different meaning. A tentative declaration of f is followed by its definition.

D.3 Extensions Accepted in Normal C++ Mode

The following C++ extensions are accepted (except when strict standard conformance mode is enabled, in which case a warning or caution message may be issued):

- A friend declaration for a class can omit the class keyword, as shown in the following example:

```
class B;
class A  {
   friend B;   // Should be "friend class B"
};
```

- Constants of scalar type can be defined within classes, as shown in the following example:

```
class A {
   const int size=10;
   int a[size];
};
```

- In the declaration of a class member, a qualified name can be used, as shown in the following example:

```
struct A {
   int A::f();   // Should be int f();
}
```

- An assignment operator declared in a derived class with a parameter type matching one of its base classes is treated as a "default" assignment operator; that is, such a declaration blocks the implicit generation of a copy assignment

operator. This is `cfront` behavior that is known to be relied upon in at least one widely used library. Here is an example:

```
struct A { };
struct B : public A {
   B& operator=(A&);
};
```

By default, as well as in `cfront` compatibility mode, there will be no implicit declaration of `B::operator=(const B&)`, whereas in strict-ANSI mode, `B::operator=(A&)` is not a copy assignment operator and `B::operator=(const B&)` is implicitly declared.

- Implicit type conversion between a pointer to an `extern "C"` function and a pointer to an `extern "C++"` function is permitted. The following is an example:

```
extern "C" void f();  // f's type has extern "C" linkage
void (*pf)()           // pf points to an extern "C++" function
             = &f;    // error unless implicit conversion allowed
```

- The `?` operator, for which the second and third operands are string literals or wide string literals, can be implicitly converted to one of the following:

```
char *
wchar_t *
```

In C++ string literals are `const`. There is a deprecated implicit conversion that allows conversion of a string literal to `char *`, dropping the `const`. That conversion, however, applies only to simple string literals. Allowing it for the result of a `?` operation is an extension:

```
char *p = x ? "abc" : "def";
```

D.4 Extensions Accepted in C or C++ Mode

The following extensions are accepted in C or C++ mode except when strict standard conformance modes is enabled, in which case a warning or caution message may be issued.

- The special `lint` comments `/*ARGSUSED*/`, `/*VARARGS*/` (with or without a count of nonvarying arguments), and `/*NOTREACHED*/` are recognized.

- A translation unit (input file) can contain no declarations.

- Comment text can appear at the ends of preprocessing directives.

- Bit fields can have base types that are `enum` or integral types in addition to `int` and `unsigned int`. This corresponds to A.6.5.8 in the ANSI Common Extensions appendix.

- `enum` tags can be incomplete as long as the tag name is defined and resolved by specifying the brace-enclosed list later.

- An extra comma is allowed at the end of an `enum` list.

- The final semicolon preceding the closing of a `struct` or `union` type specifier can be omitted.

- A label definition can be immediately followed by a right brace (}). (Normally, a statement must follow a label definition.)

- An empty declaration (a semicolon preceded by nothing) is allowed.

- An initializer expression that is a single value and is used to initialize an entire static array, struct, or union does not need to be enclosed in braces. ANSI C requires braces.

- In an initializer, a pointer constant value can be cast to an integral type if the integral type is large enough to contain it.

- The address of a variable with register storage class may be taken.

- In an integral constant expression, an integer constant can be cast to a pointer type and then back to an integral type.

- In duplicate size and sign specifiers (for example, `short short` or `unsigned unsigned`) the redundancy is ignored.

- Benign redeclarations of `typedef` names are allowed. That is, a `typedef` name can be redeclared in the same scope with the same type.

- Dollar sign ($) and at sign (@) characters can be accepted in identifiers by using the `-h calchars` command line option. This is not allowed by default.

- Numbers are scanned according to the syntax for numbers rather than the pp-number syntax. Thus, `0x123e+1` is scanned as three tokens instead of one token that is not valid. If the `-h conform` option is specified, the pp-number syntax is used.

- Assignment and pointer differences are allowed between pointers to types that are interchangeable but not identical, for example, `unsigned char *` and `char *`. This includes pointers to integral types of the same size (for

example, `int *` and `long *`). Assignment of a string constant to a pointer to any kind of character is allowed without a warning.

- Assignment of pointer types is allowed in cases where the destination type has added type qualifiers that are not at the top level (for example, `int **` to `const int **`). Comparisons and pointer difference of such pairs of pointer types are also allowed.

- In operations on pointers, a pointer to `void` is always implicitly converted to another type if necessary, and a null pointer constant is always implicitly converted to a null pointer of the right type if necessary. In ANSI C, these are allowed by some operators, and not by others (generally, where it does not make sense).

- Pointers to different function types may be assigned or compared for equality (`==`) or inequality (`!=`) without an explicit type cast. This extension is not allowed in C++ mode.

- A pointer to `void` can be implicitly converted to or from a pointer to a function type.

- External entities declared in other scopes are visible:

```
void f1(void) { extern void f(); }
void f2() { f(); /* Using out of scope declaration */ }
```

- In C mode, end-of-line comments (`//`) are supported.

- A non-lvalue array expression is converted to a pointer to the first element of the array when it is subscripted or similarly used.

- The `fortran` keyword. For more information, see Section 9.2, page 154.

- Cray hexadecimal floating point constants. For more information, see Section 9.3, page 154.

D.5 C++ Extensions Accepted in `cfront` Compatibility Mode

The cfront compatibility mode is enabled by the `-h cfront` command-line option. The following extensions are accepted in cfront compatibility mode:

- Type qualifiers on the `this` parameter are dropped in contexts such as in the following example:

```
struct A {
   void f() const;
};
void (A::*fp)() = &A::f;
```

This is a safe operation. A pointer to a `const` function can be put into a pointer to non-`const`, because a call using the pointer is permitted to modify the object and the function pointed to will not modify the object. The opposite assignment would not be safe.

- Conversion operators that specify a conversion to `void` are allowed.

- A nonstandard `friend` declaration can introduce a new type. A `friend` declaration that omits the elaborated type specifier is allowed in default mode, however, in cfront mode the declaration can also introduce a new type name. An example follows:

```
struct A {
   friend B;
};
```

- The third operator of the `?` operator is a conditional expression instead of an assignment expression.

- A reference to a pointer type may be initialized from a pointer value without use of a temporary even when the reference pointer type has additional type qualifiers above those present in the pointer value. For example:

```
int *p;
const int *&r = p;   // No temporary used
```

- A reference can be initialized to `NULL`.

- Because cfront does not check the accessibility of types, access errors for types are issued as warnings instead of errors.

- When matching arguments of an overloaded function, a `const` variable with a value of 0 is not considered to be a null pointer constant. In general, in overload resolution, a null pointer constant must be spelled "0" to be

considered a null pointer constant (e.g., '\0' is not considered a null pointer constant).

- An alternate form of declaring pointer-to-member-function variables is supported, as shown in the following example:

```
struct A {
 void f(int);
 static void sf(int);
 typedef void A::T3(int);  // nonstd typedef decl
 typedef void T2(int);     // std typedef
};
typedef void A::T(int);   // nonstd typedef decl
T* pmf = &A::f;           // nonstd ptr-to-member decl
A::T2* pf = A::sf;        // std ptr to static mem decl
A::T3* pmf2 = &A::f;      // nonstd ptr-to-member decl
```

In this example, `T` is construed to name a function type for a nonstatic member function of class `A` that takes an `int` argument and returns `void`; the use of such types is restricted to nonstandard pointer-to-member declarations. The declarations of `T` and `pmf` in combination are equivalent to the following single standard pointer-to-member declaration:

```
void (A::* pmf)(int) = &A::f;
```

A nonstandard pointer-to-member declaration that appears outside of a class declaration, such as the declaration of `T`, is normally not valid and would cause an error to be issued. However, for declarations that appear within a class declaration, such as `A::T3`, this feature changes the meaning of a valid declaration. `cfront` version 2.1 accepts declarations, such as `T`, even when `A` is an incomplete type; so this case is also accepted.

- Protected member access checking is not done when the address of a protected member is taken. For example:

```
class B { protected: int i; };
class D : public B { void mf()};

void D::mf() {
 int B::* pmi1 = &B::i;  // error, OK in cfront mode
 int D::* pmi2 = &D::i;  // OK
}
```

Note: Protected member access checking for other operations (such as everything except taking a pointer-to-member address) is done normally.

- The destructor of a derived class can implicitly call the private destructor of a base class. In default mode, this is an error but in cfront mode it is reduced to a warning. For example:

```
class A {
    ~A();
};
class B : public A {
    ~B();
};
B::~B(){}              // Error except in cfront mode
```

- When disambiguation requires deciding whether something is a parameter declaration or an argument expression, the pattern *type-name-or-keyword*(*identifier* ...) is treated as an argument. For example:

```
class A { A(); };
double d;
A x(int(d));
A(x2);
```

By default, `int(d)` is interpreted as a parameter declaration (with redundant parentheses), and so `x` is a function; but in cfront compatibility mode `int(d)` is an argument and `x` is a variable.

The declaration `A(x2)` is also misinterpreted by cfront. It should be interpreted as the declaration of an object named `x2`, but in cfront mode it is interpreted as a function style cast of `x2` to the type `A`.

Similarly, the following declaration declares a function named `xzy`, that takes a parameter of type function taking no arguments and returning an `int`. In cfront mode, this is interpreted as a declaration of an object that is initialized with the value `int()`, which evaluates to 0.

```
int xyz(int());
```

- A named bit field can have a size of 0. The declaration is treated as though no name had been declared.

- Plain bit fields (such as bit fields declared with a type of `int`) are always signed.

- The name given in an elaborated type specifier can be a `typedef` name that is the synonym for a class name. For example:

```
typedef class A T;
class T *pa;              // No error in cfront mode
```

- No warning is issued on duplicate size and sign specifiers, as shown in the following example:

```
short short int i;  // No warning in cfront mode
```

- Virtual function table pointer-update code is not generated in destructors for base classes of classes without virtual functions, even if the base class virtual functions might be overridden in a further derived class. For example:

```
struct A {
  virtual void f() {}
  A() {}
  ~A() {}
};
struct B : public A {
  B() {}
  ~B() {f();}        // Should call A::f according to ARM 12.7
};
struct C : public B {
  void f() {}
} c;
```

  In cfront compatibility mode, `B::~B` calls `C::f`.

- An extra comma is allowed after the last argument in an argument list. For example:

```
f(1, 2, );
```

- A constant pointer-to-member function can be cast to a pointer-to-function, as in the following example. A warning is issued.

```
struct A {int f();};
main () {
  int (*p)();
  p = (int (*)())A::f;        // Okay, with warning
}
```

- Arguments of class types that allow bitwise copy construction but also have destructors are passed by value like C structures, and the destructor

is not called on the copy. In normal mode, the class object is copied into a temporary, the address of the temporary is passed as the argument, and the destructor is called on the temporary after the call returns. Because the argument is passed by value instead of by address, code like this compiled in `cfront` mode is not calling-sequence compatible with the same code compiled in normal mode. In practice, this is not much of a problem, since classes that allow bitwise copying usually do not have destructors.

- A union member may be declared to have the type of a class for which the user has defined an assignment operator (as long as the class has no constructor or destructor). A warning is issued.

- When an unnamed class appears in a `typedef` declaration, the `typedef` name may appear as the class name in an elaborated type specifier. For example:

```
typedef struct { int i, j; } S;
struct S x;                    // No error in cfront mode
```

- Two member functions may be declared with the same parameter types when one is static and the other is nonstatic with a function qualifier. For example:

```
class A {
  void f(int) const;
  static void f(int); // No error in cfront mode
};
```

- The scope of a variable declared in the `for-init-statement` is the scope to which the `for` statement belongs. For example:

```
int f(int i) {
  for (int j = 0; j < i; ++j) { /* ... */ }
  return j; // No error in cfront mode
}
```

- Function types differing only in that one is declared `extern "C"` and the other `extern "C++"` can be treated as identical:

```
typedef void (*PF)();
extern "C" typedef void (*PCF)();
void f(PF);
void f(PCF);
```

By contrast, in standard C++, PF and PCF are different and incompatible types; PF is a pointer to an `extern "C++"` function whereas PCF is a

pointer to an `extern "C"` function; and the two declarations of `f` create an overload set.

- Functions declared `inline` have internal linkage.

- `enum` types are regarded as integral types.

- An uninitialized const object of non-POD class type is allowed even if its default constructor is implicitly declared as in the following example:

```
struct A { virtual void f(); int i; };
const A a;
```

- A function parameter type is allowed to involve a pointer or reference to array of unknown bounds.

- If the user declares an `operator=` function in a class, but not one that can serve as the default `operator=`, and bitwise assignment could be done on the class, a default `operator=` is not generated. Only the user-written `operator=` functions are considered for assignments, so bitwise assignment is not done.

Compiler Messages [E]

This appendix describes how to use the message system to control and use messages issued by the compiler. Explanatory texts for messages can be displayed online through the use of the `explain` command.

E.1 Expanding Messages with the `explain` Command

You can use the `explain` command to display an explanation of any message issued by the compiler. The command takes as an argument, the message number, including the number's prefix. The prefix for Cray C and C++ is `CC`.

In the following sample dialog, the `cc`(1) command invokes the compiler on source file `bug.c`. Message `CC-24` is displayed. The `explain` command displays the expanded explanation for this message.

```
% cc bug.c
CC-24 cc: ERROR File = bug.c, Line = 1
   An invalid octal constant is used.

   int i = 018;
             ^

1 error detected in the compilation of "bug.c".
% explain CC-24

An invalid octal constant is used.

Each digit of an octal constant must be between 0 and 7,
inclusive. One or more digits in the octal constant on the
indicated line are outside of this range. To avoid issuing
an error for each erroneous digit, the constant will be treated
as a decimal constant. Change each digit in the octal constant
to be within the valid range.
```

E.2 Controlling the Use of Messages

This section summarizes the command line options that affect the issuing of messages from the compiler.

E.2.1 Command Line Options

Option	Description

`-h errorlimit[=`*n*`]`

> Specifies the maximum number of error messages the compiler prints before it exits.

`-h [no]message=`*n*`[:...]`

> Enables or disables the specified compiler messages, overriding `-h msglevel`.

`-h msglevel_`*n*

> Specifies the lowest severity level of messages to be issued.

`-h report=`*args*

> Generates optimization report messages.

E.2.2 Environment Options for Messages

The following environment variables are used by the message system.

Variable	Description
`NLSPATH`	Specifies the default value of the message system search path environment variable.
`LANG`	Identifies your requirements for native language, local customs, and coded character set with regard to the message system.
`MSG_FORMAT`	Controls the format in which you receive error messages.

E.2.3 `ORIG_CMD_NAME` Environment Variable

You can override the command name printed in the message. If the environment variable `ORIG_CMD_NAME` is set, the value of `ORIG_CMD_NAME` is used as the command name in the message. This functionality is provided for use with shell scripts that invoke the compiler. By setting `ORIG_CMD_NAME` to the name of the script, any message printed by the compiler appears as though it was generated by the script. For example, the following C shell script is named `newcc`:

```
#
setenv ORIG_CMD_NAME 'basename $0'
cc $*
```

A message generated by invoking `newcc` resembles the following:

```
 CC-8 newcc: ERROR File = x.c, Line = 1
   A new-line character appears inside a string literal.
```

Because the environment variable `ORIG_CMD_NAME` is set to `newcc`, this appears as the command name instead of `cc(1)` in this message.

⚠ **Caution:** The `ORIG_CMD_NAME` environment variable is not part of the message system. It is supported by the Cray C and C++ compilers as an aid to programmers. Other products, such as the Fortran compiler and the loader, may support this variable. However, you should not rely on support for this variable in any other product.

You must be careful when setting the environment variable `ORIG_CMD_NAME`. If you set `ORIG_CMD_NAME` inadvertently, the compiler may generate messages with an incorrect command name. This may be particularly confusing if, for example, `ORIG_CMD_NAME` is set to `newcc` when the Fortran compiler prints a message. The Fortran message will look as though it came from `newcc`.

E.3 Message Severity

Each message issued by the compiler falls into one of the following categories of messages, depending on the severity of the error condition encountered or the type of information being reported.

Category	Meaning
COMMENT	Inefficient programming practices.
NOTE	Unusual programming style or the use of outmoded statements.

CAUTION	Possible user error. Cautions are issued when the compiler detects a condition that may cause the program to abort or behave unpredictably.
WARNING	Probable user error. Indicates that the program will probably abort or behave unpredictably.
ERROR	Fatal error; that is, a serious error in the source code. No binary output is produced.
INTERNAL	Problems in the compilation process. Please report internal errors immediately to the system support staff, so a Software Problem Report (SPR) can be filed.
LIMIT	Compiler limits have been exceeded. Normally you can modify the source code or environment to avoid these errors. If limit errors cannot be resolved by such modifications, please report these errors to the system support staff, so that an SPR can be filed.
INFO	Useful additional information about the compiled program.
INLINE	Information about inline code expansion performed on the compiled code.
SCALAR	Information about scalar optimizations performed on the compiled code.
VECTOR	Information about vectorization optimizations performed on the compiled code.
STREAM	Information about the MSP optimizations performed on the compiled code.
OPTIMIZATION	Information about general optimizations.

E.4 Common System Messages

The errors in the following list can occur during the execution of a user program. The operating system detects them and issues the appropriate message. These errors are not detected by the compiler and are not unique to C and C++ programs; they may occur in any application program written in any language.

- Operand Range Error

An operand range error occurs when a program attempts to load or store in an area of memory that is not part of the user's area. This usually occurs when an invalid pointer is dereferenced.

- `Program Range Error`

  A program range error occurs when a program attempts to jump into an area of memory that is not part of the user's area. This may occur, for example, when a function in the program mistakenly overwrites the internal program stack. When this happens, the address of the function from which the function was called is lost. When the function attempts to return to the calling function, it jumps elsewhere instead.

- `Error Exit`

  An error exit occurs when a program attempts to execute an invalid instruction. This error usually occurs when the program's code area has been mistakenly overwritten with words of data (for example, when the program stores in a location pointed to by an invalid pointer).

Intrinsic Functions [F]

The C and C++ intrinsic functions either allow for direct access to some hardware instructions or result in generation of inline code to perform some specialized functions. These intrinsic functions are processed completely by the compiler. In many cases, the generated code is one or two instructions. These are called *functions* because they are invoked with the syntax of function calls.

To get access to the intrinsic functions, the Cray C++ compiler requires that either the `intrinsics.h` file be included or that the intrinsic functions that you want to call be explicitly declared. If the source code does not have an `intrinsics.h` statement and you cannot modify the code, you can use the `-h prototype_intrinsics` option instead. If you explicitly declare an intrinsic function, the declaration must agree with the documentation or the compiler treats the call as a call to a normal function, not the intrinsic function. The `-h nointrinsics` command line option causes the compiler to treat these calls as regular function calls and not as intrinsic function calls.

The types of the arguments to intrinsic functions are checked by the compiler, and if any of the arguments do not have the correct type, a warning message is issued and the call is treated as a normal call to an external function. If your intention was to call an external function with the same name as an intrinsic function, you should change the external function name. The names used for the Cray C intrinsic functions are in the name space reserved for the implementation.

> **Note:** Several of these intrinsic functions have both a vector and a scalar version. If a vector version of an intrinsic function exists and the intrinsic is called within a vectorized loop, the compiler uses the vector version of the intrinsic. See the appropriate intrinsic function man page for details on whether it has a vector version.

The following sections groups the C and C++ intrinsics according to function and provides a brief description of each intrinsic in that group. See the corresponding man page for more information.

F.1 Atomic Memory Operations

The following intrinsics perform various atomic memory operations:

> **Note:** In this discussion, an object is an entity that is referred to by a pointer. A value is an actual number, bit mask, etc. that is not referred to by a pointer.

Intrinsic	Description
_amo_aadd	Adds a value to an object that is referred to by a pointer and stores the results in the object.
_amo_aax	ANDs a value and an object that is referred to by a pointer, XORs the result with a third value, and stores the results in the object.
_amo_afadd	Adds a value to an object that is referred to by a pointer and stores the result in the object. The intrinsic returns the original value of the object.
_amo_afax	ANDs a value with an object that is referred to by a pointer, XORs the result with a second value, and stores the result in the object. The intrinsic returns the original value of the object.
_amo_acswap	(Compare and swap) Compares an object that is referenced by a pointer against a value. If equal, a specified value is stored in the object. The intrinsic returns the original value of object.

F.2 BMM Operations

The following intrinsics perform operations on the BMM:

_mtilt	Inverts a bit matrix
_mclr	Logically undefines the BMM unit.
_mld	Loads the BMM functional unit with a matrix vector in transposed form.
_mldmx	Combines the load and multiply functions.
_mmx	Performs a bit matrix multiply.
_mul	Unloads the bit matrix function unit.

F.3 Bit Operations

The following intrinsics copy, count, or shift bits or computes the parity bit:

_dshiftl	Move the left most n bits of an integer into the right side of another integer, and return that integer.

`_dshiftr`	Move the right most n bits of an integer into the left side of another integer and return that integer.
`_pbit`	Copies the rightmost bit of a word to the n[th] bit, from the right, of another word.
`_pbits`	Copies the rightmost m bits of a word to another word beginning at bit n.
`_poppar`	Computes the parity bit for a variable.

`_popcnt`
`_popcnt32`
`_popcnt64`
	Counts the number of set bits in 32-bit and 64-bit integer words.

`_leadz`
`_leadz32`
`_leadz64`
	Counts the number of leading 0 bits in 32-bit and 64-bit integer words.

`_gbit`	`_gbit` returns the value of the n[th] bit from the right.
`_gbits`	Returns a value consisting of `m` bits extracted from a variable, beginning at n[th] bit from the right.

F.4 Function Operations

These intrinsics return information about function arguments:

`_argcount`	Returns the number of arguments explicitly passed to a function, excluding any "hidden" arguments added by the compiler.
`_numargs`	Returns the total number of words in the argument list passed to the function including any "hidden" arguments added by the compiler.

F.5 Mask Operations

These intrinsics create bit masks:

`_mask`	Creates a left-justified or right-justified bit mask with all bits set to 1.
`_maskl`	Returns a left-justified bit mask with `i` bits set to 1.

_maskr   Returns a right-justified bit mask with i bits
                              set to 1.

F.6 Memory Operations

This intrinsic assures that memory references synchronize memory:

_gsync   Performs global synchronization of all memory.

F.7 Miscellaneous Operations

The following intrinsics perform various functions:

_int_mult_upper   Multiplies integers and returns the uppermost bits.
                              Refer to the int_mult_upper(3i) man page.

_ranf   _ranf, compute a pseudo-random floating-point
                              number ranging from 0.0 through 1.0.

_rtc   Return a real-time clock value expressed in clock
                              ticks.

F.8 Streaming Operations

These intrinsics return streaming information:

__sspid   Indicates which SSP is being used by the code.
                              This intrinsic applies to MSP-mode applications,
                              not SSP-mode applications.

__streaming   Indicates whether the code is capable of
                              multistreaming.

Glossary

application node

For UNICOS/mp systems, a node that is used to run user applications. Application nodes are best suited for executing parallel applications and are managed by the strong application placement scheduling and gang scheduling mechanism Psched. See also *node*; *node flavor*.

barrier

An obstacle within a program that provides a mechanism for synchronizing tasks. When a task encounters a barrier, it must wait until all specified tasks reach the barrier.

barrier synchronization

1. An event initiated by software that prevents cooperating tasks from continuing to issue new program instructions until all of the tasks have reached the same point in the program. 2. A feature that uses a barrier to synchronize the processors within a partition. All processors must reach the barrier before they can continue the program.

basic block

A section of a program that does not cross any conditional branches, loop boundaries, or other transfers of control. There is a single entry point and a single exit point. Many compiler optimizations occur within basic blocks.

binding

The way in which one component in a resource specification is related to another component.

blocking

An optimization that involves changing the iteration order of loops that access large arrays so that groups of array elements are processed as many times as possible while they reside in cache.

breakpoint

A point in a program that, when reached, triggers some special behavior useful to the process of debugging; generally, breakpoints are used to either pause

program execution and/or dump the values of some or all of the program variables. Breakpoints may be part of the program itself, or they may be set by the programmer as part of an interactive session with a debugging tool for scrutinizing the execution of the program.

C interoperability

A Fortran 2003 feature that allows Fortran programs to call C functions and access C global objects and also allows C programs to call Fortran procedures and access Fortran global objects.

compute module

For a Cray X1 series mainframe, the physical, configurable, scalable building block. Each compute module contains either one node with 4 MCMs/4MSPs (Cray X1 modules) or two nodes with 4 MCMs/8MSPs (Cray X1E modules). Sometimes referred to as a node module. See also *node*.

Cray Programming Environment Server (CPES)

A server for the Cray X1 series system that runs the Programming Environment software.

Cray streaming directives (CSDs)

Nonadvisory directives that allow you to more closely control multistreaming for key loops.

Cray Workstation (CWS)

For Cray X1 series systems, the system operation, administration, and maintenance workstation.

CrayDoc

Cray's documentation system for accessing and searching Cray books, man pages, and glossary terms from a web browser.

CrayPat

For Cray X1 series systems, the primary high-level tool for identifying opportunities for optimization. CrayPat allows you to perform profiling, sampling, and tracing experiments on an instrumented application and to analyze the results of those experiments; no recompilation is needed to produce

the instrumented program. In addition, the CrayPat tool provides access to all hardware performance counters.

CRInform

An online technical-assistance and problem-reporting service for subscribing Cray customers.

deferred implementation

The label used to introduce information about a feature that will not be implemented until a later release.

distributed memory

The kind of memory in a parallel processor where each processor has fast access to its own local memory and where to access another processor's memory it must send a message via the interprocessor network.

dynamic extent

In OpenMP, an extent that includes both the statements of a lexical extent and the statements of a function that is called from the lexical extent. A dynamic extent is an instance of a region.

entry point

A location in a program or routine at which execution begins. A routine may have several entry points, each serving a different purpose. Linkage between program modules is performed when the linkage editor binds the external references of one group of modules to the entry points of another module.

environment variable

A variable that stores a string of characters for use by your shell and the processes that execute under the shell. Some environment variables are predefined by the shell, and others are defined by an application or user. Shell-level environment variables let you specify the search path that the shell uses to locate executable files, the shell prompt, and many other characteristics of the operation of your shell. Most environment variables are described in the ENVIRONMENT VARIABLES section of the man page for the affected command.

Etnus TotalView

For UNICOS/mp systems, a symbolic source-level debugger designed for debugging the multiple processes of parallel Fortran, C, or C++ programs.

kind

Data representation (for example, single precision, double precision). The kind of a type is referred to as a kind parameter or kind type parameter of the type. The kind type parameter `KIND` indicates the decimal range for the integer type, the decimal precision and exponent range for the real and complex types, and the machine representation method for the character and logical types.

lexical block

The scope within which a C or C++ directive is on or off and is bounded by
the opening curly brace just before the declaration of the directive and the
corresponding closing curly brace. Only applicable executable statements within
the lexical block are affected as indicated by the directive. The lexical block does
not include the statements contained within a procedure that is called from the
lexical block. This example code shows the lexical block for the `upc strict` and
`upc relaxed` directives:

```
void Example(void)
{
  #pragma _CRI upc strict      // UPC strict state is on
  ...

  {
    ...                        // UPC strict state is still on
    #pragma _CRI upc relaxed // UPC strict state is now off
    ...
  }

                               // UPC strict state is back on
  ...
}
```

locale

For UNICOS/mp systems, a collection of culture-dependent information used by
an application to interact with a user.

lock

1. Any device or algorithm that is used to ensure that only one process will
perform some action or use some resource at a time. 2. A synchronization
mechanism that, by convention, forces some data to be accessed by tasks in a
serial fashion. Locks have two states: locked and unlocked. 3. A facility that
monitors critical regions of code.

loop fusion

An optimization that takes the bodies of loops with identical iteration counts and
fuses them into a single loop with the same iteration count.

loop interchange

An optimization that changes the order of loops within a loop nest, to achieve stride minimization or eliminate data dependencies.

loop unrolling

An optimization that increases the step of a loop and duplicates the expressions within a loop to reflect the increase in the step. This can improve instruction scheduling and reduce memory access time.

loopmark listing

A listing that is generated by invoking the Cray Fortran Compiler with the `-rm` option. The loopmark listing displays what optimizations were performed by the compiler and tells you which loops were vectorized, streamed, unrolled, interchanged, and so on.

master thread

The thread that creates a team of threads when an OpenMP parallel region is entered.

MSP mode (multistreaming mode)

One of two types of application modes for UNICOS/mp systems. Programs are compiled either as MSP-mode applications (default) or SSP-mode applications. MSP-mode applications run on one or more MSPs. For MSP-mode applications, each MSP coordinates the interactions of its associated four SSPs. See also *SSP mode.*

multichip module (MCM)

For Cray X1 series systems, the physical packaging that contains processor chips and cache chips.The chips implement either one multistreaming processor (Cray X1 MCM) or two multistreaming processors (Cray X1E MCM). See also *MSP.*

multistreaming processor (MSP)

For UNICOS/mp systems, a basic programmable computational unit. Each MSP is analogous to a traditional processor and is composed of four single-streaming processors (SSPs) and E-cache that is shared by the SSPs. See also *node*; *SSP*; *MSP mode*; *SSP mode.*

nested parallel region

An OpenMP parallel region that appears within a dynamic extent of an OpenMP PARALLEL construct that does not have an if clause or has an if clause that evaluates to true. See also *dynamic extent*.

node

For UNICOS/mp systems, the logical group of four multistreaming processors (MSPs), cache-coherent shared local memory, high-speed interconnections, and system I/O ports. A Cray X1 system has one node with 4 MSPs per compute module. A Cray X1E system has two nodes of 4 MSPs per node, providing a total of 8 MSPs on its compute module. Software controls how a node is used: as an OS node, application node, or support node. See also *compute module*; *MCM*; *MSP*, *node flavor*; *SSP*.

node flavor

For UNICOS/mp systems, software controls how a node is used. A node's software-assigned flavor dictates the kind of processes and threads that can use its resources. The three assignable node flavors are application, OS, and support. See also *application node*; *OS node*; *support node*; *system node*.

OpenMP

An industry-standard, portable model for shared memory parallel programming.

OS node

For UNICOS/mp systems, the node that provides kernel-level services, such as system calls, to all support nodes and application nodes. See also *node*; *node flavor*.

overindexing

The nonstandard practice of referencing an array with a subscript not contained between the declared lower and upper bounds of the corresponding dimension for that array. This practice sometimes, but not always, leads to referencing a storage location outside of the entire array.

page size

The unit of memory addressable through the Translation Lookaside Buffer (TLB). For a UNICOS/mp system, the base page size is 65,536 bytes, but larger page sizes (up to 4,294,967,296 bytes) are also available.

parallel region

See *serial region.*

partitioning

Configuring a UNICOS/mp system into logical systems (partitions). Each partition is independently operated, booted, dumped, and so on without impact on other running partitions. Hardware and software failures in one partition do not affect other partitions.

pointer

A data item that consists of the address of a desired item.

private variable

A variable that is accessible to only one thread in a team of an OpenMP parallel region.

Psched

The UNICOS/mp application placement scheduling tool. The `psched` command can provide job placement, load balancing, and gang scheduling for all applications placed on application nodes.

reduction loop

A loop that contains at least one statement that reduces an array to a scalar value by doing a cumulative operation on many of the array elements. This involves including the result of the previous iteration in the expression of the current iteration.

search loop

An array-processing loop used to perform a table lookup or to find exceptional values within an array.

serial region

An area within a program in which only the master task is executing. Its opposite is a parallel region.

SHMEM

A library of optimized functions and subroutines that take advantage of shared memory to move data between the memories of processors. The routines can either be used by themselves or in conjunction with another programming style such as Message Passing Interface. SHMEM routines can be called from Fortran, C, and C++ programs.

shortloop

A loop that is vectorized but that has been determined by the compiler to have trips less than or equal to the maximum vector length. In this case, the compiler deletes the loop to the top of the loop. If the shortloop directive is used or the trip count is constant, the top test for number of trips is deleted. A shortloop is more efficient than a conventional loop.

single-streaming processor (SSP)

For UNICOS/mp systems, a basic programmable computational unit. See also *node*; *MSP*; *MSP mode*; *SSP mode*.

Software Problem Report (SPR)

A Cray customer service form and process that tracks software problems from first report to resolution. SPR resolution results either from a written reply, the release of software containing the fix to the problem, or the implementation of the requested design change.

SSP mode (single-streaming mode)

One of two types of application modes for UNICOS/mp systems. Programs are compiled either as MSP-mode applications (default) or SSP-mode applications. SSP-mode applications run on one or more SSPs. Each SSP runs independently of the others, executing its own stream of instructions. In contrast, compiler options enable the programmer to develop command-mode programs that run on an SSP on the support node. See also *MSP mode*.

support node

For UNICOS/mp systems, the node that is used to run serial commands, such as shells, editors, and other user commands (`ls`, for example). See also *node*; *node flavor*.

system node

For UNICOS/mp systems, the node that is designated as both an OS node and a support node; this node is often called a system node; however, there is no node flavor of "system." See also *node*; *node flavor*.

thread

The active entity of execution. A sequence of instructions together with machine context (processor registers) and a stack. On a parallel system, multiple threads can be executing parts of a program at the same time.

trigger

A command that a user logged into a Cray X1 series system uses to launch Programming Environment components residing on the CPES. Examples of trigger commands are `ftn`, `CC`, and `pat_build`.

type

A means for categorizing data. Each intrinsic and user-defined data type has four characteristics: a name, a set of values, a set of operators, and a means to represent constant values of the type in a program.

UNICOS/mp

The operating system for Cray X1 series (Cray X1 and Cray X1E) systems.

Unified Parallel C (UPC)

An extension of the C programming language designed for high performance computing on large-scale parallel processing machines. UPC provides a uniform programming model for both shared and distributed memory hardware. Other parallel programming models include Message Passing Interface, SHMEM, Co-array Fortran, and OpenMP.

unrolling

A single-processing-element optimization technique in which the statements within a loop are copied. For example, if a loop has two statements, unrolling might copy those statements four times, resulting in eight statements. The loop control variable would be incremented for each copy, and the stride through the array would also be increased by the number of copies. This technique is often performed directly by the compiler, and the number of copies is usually between two and four.

vector

A series of values on which instructions operate; this can be an array or any subset of an array such as row, column, or diagonal. Applying arithmetic, logical, or memory operations to vectors is called vector processing. See also *vector processing*.

vector length

The number of elements in a vector.

vector processing

A form of instruction-level parallelism in which the vector registers are used to perform iterative operations in parallel on the elements of an array, with each iteration producing up to 64 simultaneous results. See also *vector*.

vectorization

The process, performed by the compiler, of analyzing code to determine whether it contains vectorizable expressions and then producing object code that uses the vector unit to perform vector processing.

Index

-#, 43
-##, 43
-###, 43

A

Advisory directives defined, 74
Affinity, 123
_amo_aadd, 226
_amo_aax, 226
_amo_acswap, 226
_amo_afadd, 226
_amo_afax, 226
Anachronisms
  C++, 207
aprun, 161
_argcount, 227
Argument passing, 168
Arithmetic
  See math
Array storage, 169
Arrays, 187
  dependencies, 91
asm statements, 189
Assembly language
  functions, 167
  output, 43
Assembly source expansions, 7
Auto aprun (see
  CRAY_AUTO_APRUN_OPTIONS.), 59

B

Bit fields, 188
Blank common block, 172
bounds directive, 70
btol conversion utility, 170

C

-c, 193

C extensions, 153
  See also Cray C extensions
C interoperability, 178
C libraries, 193
-c option, 43
-C option, 45
C++
  libraries, 141
  templates, 143
Calls, 165
can_instantiate directive, 80, 149
Cfront, 212
  compatibility mode, 203
  compilers, 13
  option, 13
Character data, 170
Character set, 185
Characters
  wide, 186
CIV
  See Constant increment variables
Classes, 188
Command line options
  -# option, 43
  -## option, 43
  -### option, 43
  -c option, 7, 43
  -C option, 45
  compiler version, 55
  conflicting with directives, 12
  conflicting with other options, 12
  -D *macro*[=*def*], 45
  defaults, 10
  -E option, 7, 42
  examples, 57
  -g option, 39, 163–164
  -G option, 39, 163–164
  -h anachronisms, 207

E

`-E` option,   42
Enumerations,   188
Environment,   183
environment variables
  OpenMP,   62
Environment variables
  compile time,   58
  `CRAYOLDCPPLIB`,   58
  `CRI_c89_OPTIONS`,   59
  `CRI_cc_OPTIONS`,   59
  `CRI_CC_OPTIONS`,   59
  `CRI_cpp_OPTIONS`,   59
  `LANG`,   59, 220
  `MSG_FORMAT`,   59, 220
  `NLSPATH`,   59, 220
  `NPROC`,   59
  `OMP_DYNAMIC`,   64
  `OMP_NESTED`,   64
  `OMP_NUM_THREADS`,   63
  `OMP_SCHEDULE`,   63
  `ORIG_CMD_NAME`,   221
  run time,   59
`Error Exit`,   223
Error messages,   219
Examples
  command line,   57
Exception construct,   14
Exception handling,   14
Exceptions,   189
`explain`,   219
Extensions
  C++ mode,   208
  Cfront compatibility mode,   212
  Cray C,   153
  `_Pragma`,   69
  `#pragma` directives,   67
`extern "C"` keyword,   165
External functions
  declaring,   165

F

Features
  C++,   203
    Cfront compatibility,   203
Files
  `a.out`,   7
  constructor/destructor,   22
  default library,   48
  dependencies,   47
  `.ii` file,   146
  `intrinsics.h`,   225
  library directory,   48
  linking,   22
  output,   49
  personal libraries,   48
Floating constants,   154
Floating-point
  constants,   154
  overflow,   187
`for` directive,   104
Fortran common block,   173
`fortran` keyword,   154
Freeing up memory,   62
`friend` declaration,   212
Functions,   225
  `mbtowc`,   186

G

`-G` *level*,   39
`-g` option,   163–164
`-G` option,   163–164
`_gbit`,   227
`_gbits`,   227
GCC language extensions
  C and C++,   , 16
  C++ only,   , 19
General command functions
  `-h ident=`*name*,   53
  `-V` option,   55
gnu
  GCC language extensions,   16
`_gsync`,   228