



Cray C and C++ Reference Manual

S-2179-83

© 1996-2000, 2002-2014 Cray Inc. All Rights Reserved. This document or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Inc.

U.S. GOVERNMENT RESTRICTED RIGHTS NOTICE

The Computer Software is delivered as "Commercial Computer Software" as defined in DFARS 48 CFR 252.227-7014.

All Computer Software and Computer Software Documentation acquired by or for the U.S. Government is provided with Restricted Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7014, as applicable.

Technical Data acquired by or for the U.S. Government, if any, is provided with Limited Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7013, as applicable.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: Cray and design, Sonexion, Urika, and YarcData. The following are trademarks of Cray Inc.: ACE, Apprentice2, Chapel, Cluster Connect, CrayDoc, CrayPat, CrayPort, ECOPhex, LibSci, NodeKARE, Threadstorm. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark Linux is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

Intel, Aries and Gemini are trademarks of Intel Corporation in the United States and/or other countries. CUDA, Kepler, OpenACC and NVIDIA are trademarks of NVIDIA Corporation. Google is a trademark of Google. ISO is a trademark of International Organization for Standardization (Organisation Internationale de Normalisation). OpenMP is a trademark of OpenMP architecture Review Board. O2 is a trademark of Silicon Graphics, Inc. Optron is a trademark of Advanced Micro Devices, Inc. PGI is a trademark of The Portland Group Compiler Technology, STMicroelectronics, Inc. Platform is a trademark of Platform Computing Corporation. TotalView is a trademark of Rogue Wave Software, Inc. UNIX, the "X device," X Window System, and X/Open are trademarks of The Open Group.

Portions of this document were copied by permission of OpenMP Architecture Review Board from OpenMP C and C++ Application Program Interface, Version 2.0, March 2002, Copyright © 1997-2002, OpenMP Architecture Review Board.

RECORD OF REVISION

S-2179-83 Published June 2014 Supports the Cray C and C++ compilers running on Cray XE, Cray XK, and XC30 compute nodes.

S-2179-82 Published September 2013 Supports the Cray C and C++ compilers running on Cray XE, Cray XK, and XC30 compute nodes.

S-2179-81 Published September 2012 Supports the Cray C and C++ compilers running on Cray XE and Cray XK compute nodes.

S-2179-80 Published December 2011 Supports the Cray C and C++ compilers running on Cray XE and Cray XK compute nodes.

S-2179-74 Published June 2011 Supports the Cray C and C++ compilers running on Cray XT and Cray XE compute nodes.

S-2179-73 Published December 2010 Supports the Cray C and C++ compilers running on Cray XT and Cray XE compute nodes.

7.2 Published February 2010 Supports the Cray C and C++ compilers running on Cray XT compute nodes.

7.1 Published June 2009 Supports the Cray C and C++ compilers running on Cray XT compute nodes.

7.0 Published December 2008 Supports the Cray C and (Deferred implementation) C++ compilers running on Cray XT compute nodes.

6.0 Published September 2007 Supports the Cray C and Cray C++ 6.0 release running on Cray X1 series and Cray X2 systems.

5.6 Published March 2007 Supports Cray C++ 5.6 and Cray C 8.6 releases running on Cray X1 series systems.

5.5 Published December 2005 Supports Cray C++ 5.5 and Cray C 8.5 releases running on UNICOS/mp 3.0 or later operating systems.

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

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

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

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

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

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

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

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

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

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

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

3.3 Published July 1999 This para 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.

3.2 Published January 1999 This para 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.

3.1 Published August 1998 This para 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.

3.0.2 Published March 1998 This para 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.

3.0 Published 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.

2.0 Published 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.

Changes to this Document

S-2179-83

Added information:

- `-h develop` selects compiler optimization levels to balance compile time against application execution time. Use this option during application development, when quick turnaround is desired. It minimizes compile time at the cost of execution time performance. See [-h develop on page 53](#).
- `-h flex_mp=strict` provides a level repeatability of between the conservative and intolerant levels. See [-h flex_mp=level on page 37](#).
- `-h concurrent` is equivalent to adding a `CONCURRENT` directive (`pragma`) before every loop in the file, including loops created from array syntax. See [-h concurrent on page 45](#) and [concurrent Directive on page 85](#).
- The `-Wx, arg` option can be used to pass command line arguments to the PTX assembler for OpenACC applications. See [-W phase, "opt ..." on page 57](#).
- The `-Wc, arg` option can be used to pass command line arguments to the CUDA linker for OpenACC applications. See [-W phase, "opt ..." on page 57](#).
- The new `cray_upc_sheap_info()` call provides symmetric heap usage information to UPC applications. See the `cray_upc_sheap_info(3c)` manpage.
- The new `cray_upc_shared_cast()` call creates a pointer-to-shared from a pointer-to-local, providing the inverse functionality to `upc_cast`. See the `cray_upc_shared_cast(3c)` man page.
- See new [CRAY_OMP_CHECK_AFFINITY Environment Variable on page 106](#).
- Added descriptions for `__sync_val_compare_and_swap`, and `__sync_lock_test_and_set`. See [Local Atomic Memory Operations on page 217](#).

Changed information:

- The `-h accmodel=fast_addr` performance option is now safe for all OpenACC applications and is enabled by default.
- Correction to macro name. See the `_RELEASE_MAJOR` macro in [Macros Based on the Compiler on page 161](#).
- `-h nokeep_frame_pointer` is not an option. See [-h keep_frame_pointer on page 65](#).

S-2179-82

Added information:

- This release supports *OpenACC Application Programming Interface, Version 2.0* standard developed by PGI, Cray Inc., NVIDIA, with support from CAPS enterprise. See [Chapter 5, Using OpenACC on page 115](#).
- This release offers support for the UPC 1.3 draft specification. See [C/C++ Standard Compatibility on page 22](#). Cray extensions to UPC that are not part of the UPC Language Specification 1.3 are listed here. See [Cray Extensions on page 128](#).
- This release introduces Coarray C++, a template library that implements the coarray concept in C++. See [Chapter 8, Using Coarray C++ on page 133](#).
- Intel™ Xeon™ processor E5-2600 v2 product family, code named *Ivy Bridge* cpu target support. See [-h cpu=target_system on page 64](#).
- Object files created in a pre-8.2 environment should be recompiled if they need to be relinked in the CCE 8.2 compiler environment. See [Binary Compatibility on page 22](#).
- New compilation option for C++ applications that link with PGAS code. See [-h \[no\]pgas_runtime on page 62](#).
- Added [\[no\]vector Directive on page 90](#).
- Added [ivdep Directive on page 86](#).
- Added details regarding the printing of 128-bit floating point types. See [128-Bit Floating Point and 256-Bit Complex Predefined Types on page 189](#).
- Added more detailed description of pointer arithmetic. [Arrays and Pointers on page 189](#).
- Cloning directives are recognized at `-Oipa` levels > 1 . Revised description of IPA levels in [Table 4](#).
- Added note. See [-h \[no\]conform\(CC, cc\), -h \[no\]stdc\(cc\) on page 28](#).
- Updated predefined macros for UPC 1.3. See [Predefined Macros on page 126](#).

Changed information:

- DWARF is generated by default. See [-h nodwarf on page 53](#).
- Updated supported platforms. This release supports the Cray XE, Cray XK, or Cray XC30 systems. See [Chapter 1, Introduction](#).
- `-hpgas_runtime` is now default, changed from `-hnopgas_runtime`. See [-h \[no\]pgas_runtime on page 62](#).
- Numeric values replace alphabetic values in the *level* specification of the `-G level` option. `-G0` replaces `-Gn`, `-G1` replaces `-Gp`, `-G2` replaces `-Gf`, and `-G3` replaces `-Gfast`. The usage of alphabetic values results in a WARNING message and support will be removed in a future release. See [Debugging Options and Compiler Debugging Options on page 166](#).
- More specific guidance regarding the use of `-K trap=opt[,opt] ...` on [page 67](#).

Deleted information:

- Directives `shortloop`, `shortloop128` are deprecated. `loop_info min_trips(1) max_trips(64)` replaces `shortloop` and `loop_info min_trips(1) max_trips(128)` replaces `shortloop128`.
- Remove references to compile option `-h matherror=method` deprecated.

Contents

	<i>Page</i>
Introduction [1]	21
1.1 General Compiler Description	21
1.1.1 Cray C Compiler	21
1.1.2 Cray C++ Compiler	21
1.2 C/C++ Standard Compatibility	22
1.3 Binary Compatibility	22
1.4 Related Publications	23
Invoking the C and C++ Compilers [2]	25
2.1 CC Command	26
2.2 cc Command	26
2.3 Command Line Options	27
2.4 Standard Language Conformance Options	28
2.4.1 -h [no]c99 (cc)	28
2.4.2 -h [no]conform (CC, cc), -h [no]stdc (cc)	28
2.4.3 -h cfront (CC)	29
2.4.4 -h [no]parse_templates (CC)	29
2.4.5 -h [no]dep_name (CC)	29
2.4.6 -h [no]exceptions (CC)	29
2.4.7 -h [no]anachronisms (CC)	29
2.4.8 -h [no]new_for_init (CC)	30
2.4.9 -h [no]tolerant (cc)	30
2.4.10 -h [no]const_string_literals (CC)	30
2.4.11 -h [no]gnu	31
2.5 Virtual Function Options	33
2.5.1 -h forcevtbl (CC)	33
2.5.2 -h suppressvtbl (CC)	33
2.6 General Language Options	33
2.6.1 -h keep= <i>file</i> (CC)	34
2.6.2 -h restrict= <i>args</i>	34

	<i>Page</i>
2.6.3 -h [no]calchars	35
2.6.4 -h [no]signedshifts	35
2.7 General Optimization Options	36
2.7.1 -h [no]add_paren	36
2.7.2 -h [no]aggress	36
2.7.3 -h [no]autoprefetch	36
2.7.4 -h [no]autothread	36
2.7.5 -h display_opt	36
2.7.6 -h flex_mp= <i>level</i>	37
2.7.7 -h fusionn	37
2.7.8 -h [no]intrinsic	38
2.7.9 -h list	38
2.7.10 -h [no]msgs	39
2.7.11 -h [no]negmsgs	39
2.7.12 -h [no]omp_trace	39
2.7.13 -h [no]func_trace	40
2.7.14 -h [no]overindex	40
2.7.15 -h [no]pattern	40
2.7.16 -h pl= <i>program_library</i>	40
2.7.17 -h profile_generate	41
2.7.18 -h threadn	41
2.7.19 -h unrolln	42
2.7.20 -h wp	43
2.7.21 -O <i>level</i>	43
2.8 Automatic Cache Management Options	44
2.8.1 -h cachenn	44
2.9 Vector Optimization Options	45
2.9.1 -h concurrent	45
2.9.2 -h vectorn	45
2.10 Interprocedural Analysis (IPA) Optimization Options	46
2.10.1 Inlining	46
2.10.2 Cloning	47
2.10.3 -h ipan	47
2.10.4 -h ipafrom= <i>source[:source]</i>	48
2.11 Scalar Optimization Options	48
2.11.1 -h [no]interchange	49
2.11.2 -h scalarn	49

	<i>Page</i>
2.11.3 -h [no]zeroinc	49
2.12 Math Options	50
2.12.1 -h <i>fprn</i>	50
2.13 Debugging Options	51
2.13.1 -G <i>level</i> and -g	52
2.13.2 -h [no]bounds (cc)	52
2.13.3 -h <i>develop</i>	53
2.13.4 -h <i>dir_check</i>	53
2.13.5 -h <i>nodwarf</i>	53
2.13.6 -h <i>gasp</i> [= <i>opt</i> [: <i>opt</i>]]	53
2.13.7 -h <i>zero</i>	54
2.14 Compiler Message Options	54
2.14.1 -h <i>msglevel_n</i>	54
2.14.2 -h [no]message= <i>n</i> [: <i>n</i> ...]	55
2.14.3 -h <i>report=args</i>	55
2.14.4 -h [no]abort	55
2.14.5 -h <i>errorlimit</i>	56
2.15 Compilation Phase Options	56
2.15.1 -E	56
2.15.2 -P	56
2.15.3 -h <i>feonly</i>	56
2.15.4 -S	57
2.15.5 -c	57
2.15.6 -#, -##, and -###	57
2.15.7 -W <i>phase</i> , " <i>opt</i> ..."	57
2.15.8 -Y <i>phase</i> , <i>dirname</i>	58
2.16 Preprocessing Options	59
2.16.1 -C	59
2.16.2 -D <i>macro</i> [= <i>def</i>]	59
2.16.3 -h [no]pragma= <i>name</i> [: <i>name</i> ...]	59
2.16.4 -I <i>includir</i>	60
2.16.5 -M	61
2.16.6 -nostdinc	61
2.16.7 -U	62
2.17 Linker Options	62
2.17.1 -h [system default]_alloc	62
2.17.2 -h [no]pgas_runtime	62

	<i>Page</i>
2.17.3 -l <i>libname</i>	62
2.17.4 -L <i>ldir</i>	63
2.17.5 -o <i>outfile</i>	63
2.18 Miscellaneous Options	63
2.18.1 -h [no]acc	64
2.18.2 -h cpu= <i>target_system</i>	64
2.18.3 -h [no]fp_trap	64
2.18.4 -h ident= <i>name</i>	65
2.18.5 -h keepfiles	65
2.18.6 -h keep_frame_pointer	65
2.18.7 -h loop_trips=[tiny small medium large huge]	65
2.18.8 -h mpin	65
2.18.9 -h network= <i>nic</i>	65
2.18.10 -h [no]omp	66
2.18.11 -h [no]omp_acc	66
2.18.12 -h pic, -h PIC	66
2.18.13 -h prototype_intrinsics	66
2.18.14 -h [no]threadsafe	66
2.18.15 -h upc(cc)	67
2.18.16 -K trap= <i>opt</i> [, <i>opt</i>]...	67
2.18.17 -V	68
2.18.18 -X <i>npes</i>	68
2.19 Command Line Examples	68
2.20 Compile Time Environment Variables	69
2.21 Run Time Environment Variables	71
2.22 OpenMP Environment Variables	71
Using #pragma Directives [3]	73
3.1 Protecting Directives	74
3.2 Directives in Cray C++	74
3.3 Loop Directives	74
3.4 Alternative Directive Form: <code>_Pragma</code>	75
3.5 General Directives	75
3.5.1 [no]autothread Directive	76
3.5.2 [no]bounds Directive	76
3.5.3 cache Directive	77
3.5.4 cache_nt Directive	77
3.5.5 duplicate Directive	78

	<i>Page</i>
3.5.6 <code>ident</code> Directive	80
3.5.7 <code>message</code> Directive	80
3.5.8 <code>[no]opt</code> Directive	80
3.5.9 <code>prefetch</code> Directive	81
3.5.10 Probability Directives	83
3.5.11 <code>weak</code> Directive	84
3.6 Vectorization Directives	85
3.6.1 <code>concurrent</code> Directive	85
3.6.2 <code>hand_tuned</code> Directive	86
3.6.3 <code>ivdep</code> Directive	86
3.6.4 <code>loop_info</code> Directive	87
3.6.5 <code>loop_info prefer_thread, prefer_nothread</code> Directives	89
3.6.6 <code>nopattern</code> Directive	89
3.6.7 <code>[no]vector</code> Directive	90
3.6.8 <code>permutation</code> Directive	90
3.6.9 <code>[no]pipeline</code> Directive	91
3.6.10 <code>prefervector</code> Directive	92
3.6.11 <code>pgo loop_info</code> Directive	92
3.6.12 <code>safe_address</code> Directive	92
3.6.13 <code>safe_conditional</code> Directive	94
3.7 Scalar Directives	94
3.7.1 <code>blockable</code> Directive	95
3.7.2 <code>blockingsize</code> Directive	95
3.7.3 <code>noblocking</code> Directive	96
3.7.4 <code>collapse</code> and <code>nocollapse</code> Directives	96
3.7.5 <code>[no]interchange</code> Directive	97
3.7.6 <code>suppress</code> Directive	97
3.7.7 <code>[no]unroll</code> Directive	98
3.7.8 <code>nofission</code> Directive	100
3.7.9 <code>[no]fusion</code> Directive	100
3.8 Inlining and Cloning Directives	100
3.8.1 <code>inline_enable, inline_disable, and inline_reset</code> Directives	100
3.8.2 <code>inline_always</code> and <code>inline_never</code> Directives	102
3.8.3 <code>clone_enable, clone_disable, clone_reset</code> Directives	102
3.8.4 <code>clone_always</code> and <code>clone_never</code> Directives	103
3.9 PGAS Directive	103
3.9.1 <code>defer_sync</code> Directive	103

	<i>Page</i>
Using the OpenMP C/C++ API [4]	105
4.1 Standard Support	105
4.2 Cray Enhancements	105
4.2.1 <code>cray_omp_set_wait_policy()</code>	105
4.2.2 <code>CRAY_OMP_CHECK_AFFINITY</code> Environment Variable	106
4.3 Compiling	106
4.4 Executing	106
4.5 Debugging	107
4.6 Cray Implementation Defined Behaviors	107
4.6.1 Directives and Clauses	110
4.6.1.1 <code>atomic</code> Directive	110
4.6.1.2 <code>for</code> Directive	110
4.6.1.3 <code>parallel</code> Directive	111
4.6.1.4 <code>threadprivate</code> Directive	111
4.6.1.5 <code>private</code> Clause	111
4.6.2 Library Routines	111
4.6.2.1 <code>omp_get_max_active_levels()</code>	111
4.6.2.2 <code>omp_set_dynamic()</code>	111
4.6.2.3 <code>omp_set_schedule()</code>	112
4.6.2.4 <code>omp_set_max_active_levels()</code>	112
4.6.2.5 <code>omp_set_nested()</code>	112
4.6.2.6 <code>omp_set_num_threads()</code>	112
4.6.3 OpenMP Environment Variables	112
4.6.3.1 <code>OMP_DYNAMIC</code>	112
4.6.3.2 <code>OMP_MAX_ACTIVE_LEVELS</code>	112
4.6.3.3 <code>OMP_NESTED</code>	112
4.6.3.4 <code>OMP_NUM_THREADS</code>	112
4.6.3.5 <code>OMP_SCHEDULE</code>	113
4.6.3.6 <code>OMP_STACKSIZE</code>	113
4.6.3.7 <code>OMP_THREAD_LIMIT</code>	113
4.6.3.8 <code>OMP_WAIT_POLICY</code>	113
4.7 Limitations	113
4.8 OpenMP Accelerator Support	114
Using OpenACC [5]	115
5.1 OpenACC Execution Model	115
5.2 OpenACC Memory Model	116
5.3 Mapping the OpenACC Programming Model onto Accelerator Components	116

	<i>Page</i>
5.3.1 Streaming Multiprocessors (SM) and Scalar Processor (SP) cores	117
5.3.2 Memory	117
5.4 Mixed Model Support	117
5.5 Compiling	118
5.6 Module Support	119
5.7 Debugging	119
5.8 OpenACC Directives	120
5.9 Runtime Routines	120
5.9.1 Cray Specific Runtime Library Routines	121
5.9.2 CRAY_ACC_DEBUG Output Routines	121
5.10 Environment Variables	122
5.10.1 Cray Specific	122
5.10.2 Standard	122
5.11 OpenACC Examples	123
Using Cray Unified Parallel C (UPC) [6]	125
6.1 Implementation	126
6.1.1 Predefined Macros	126
6.1.2 False Sharing	126
6.2 Compiling and Linking UPC Code	127
6.3 Launching a UPC Application	128
6.4 Cray Extensions	128
6.4.1 Team Collectives	129
6.4.2 Node Affinity	129
6.4.2.1 upc_nodeof ()	129
6.4.2.2 NODES	129
6.4.2.3 MYNODE	130
Using Cray C++ Libraries [7]	131
7.1 Unsupported Standard C++ Library Features	131
Using Coarray C++ [8]	133
8.1 Compiling Coarray C++	134
8.2 Declaring and Accessing Coarrays	135
8.2.1 Basic Types	135
8.2.2 Arrays	137
8.2.3 Pointers	138
8.2.4 Structs, Unions, and Classes	138
8.3 Type System	140

	<i>Page</i>
8.3.1 Coreferences	140
8.3.2 Copointers	141
8.3.3 <code>shape_cast</code>	142
8.4 Control Flow and Synchronization	144
8.4.1 Writing SPMD Code	144
8.4.2 Barriers	144
8.4.3 Function Calls	145
8.4.3.1 <code>coatomic</code>	145
8.4.3.2 <code>coevent</code>	145
8.4.3.3 <code>comutex</code>	146
8.5 Collectives	146
8.5.1 <code>cobroadcast</code>	146
8.5.2 <code>coreduce</code>	147
8.6 Exceptions	148
8.7 Memory Consistency Model	148
8.7.1 <code>atomic_image_fence()</code>	148
8.7.2 Accesses within a Single Image	148
8.7.3 Accesses to Other Images	149
8.7.3.1 Multi-byte Accesses	149
8.7.3.2 From Different Images	149
8.7.3.3 From the Same Image	149
8.8 Blocking Versus Non-blocking Accesses	150
8.8.1 Writes (Puts)	150
8.8.2 Reads (Gets)	151
8.8.3 Cofutures	151
8.9 Code Patterns	152
8.9.1 Coobjects	152
8.9.2 Hoisting a <code>coptr</code>	153
Using Cray C Extensions [9]	155
9.1 Complex Data Extensions	155
9.2 <code>fortran</code> Keyword	155
9.3 Hexadecimal Floating-point Constants	156
Using Predefined Macros [10]	159
10.1 Macros Required by the C and C++ Standards	159
10.2 Macros Based on the Host Machine	160
10.3 Macros Based on the Target Machine	160

	<i>Page</i>
10.4 Macros Based on the Compiler	161
10.5 UPC Predefined Macros	161
Running C and C++ Applications [11]	163
Debugging Cray C and C++ Code [12]	165
12.1 TotalView Debugger	166
12.2 Compiler Debugging Options	166
Using Interlanguage Communication [13]	169
13.1 Calls Between C and C++ Functions	169
13.2 Calling Fortran Functions and Subroutines from C or C++	170
13.2.1 Requirements	171
13.2.2 Argument Passing	171
13.2.3 Array Storage	172
13.2.4 Logical and Character Data	173
13.2.5 Accessing Named Common from C and C++	173
13.2.6 Accessing Blank Common from C or C++	175
13.2.7 Cray C and Fortran Example	177
13.2.8 Calling a Fortran Program from Cray C++	179
13.3 Calling a C or C++ Function from Fortran	180
13.3.1 Portable Interoperability Mechanism	180
13.3.2 Standard Fortran/C Interoperability	183
Implementation-defined Behavior [14]	185
14.1 Messages	185
14.2 Environment	185
14.2.1 Identifiers	186
14.2.2 Types	186
14.2.3 Characters	187
14.2.4 Wide Characters	188
14.2.5 Integers	188
14.2.6 128-Bit Floating Point and 256-Bit Complex Predefined Types	189
14.2.7 Arrays and Pointers	189
14.2.8 Registers	190
14.2.9 Classes, Structures, Unions, Enumerations, and Bit Fields	190
14.2.10 Qualifiers	191
14.2.11 Declarators	191
14.2.12 Statements	191

	<i>Page</i>
14.2.13 Exceptions	191
14.2.14 System Function Calls	191
14.3 Preprocessing	191
Appendix A Using Libraries and the Linker	193
A.1 Cray C and C++ Libraries	193
A.2 Linker	193
Appendix B Using Cray C and C++ Dialects	195
B.1 C++ Language Conformance	195
B.1.1 Supported C++ Language Features	195
B.2 C++ Anachronisms Accepted	198
B.3 Extensions Accepted in Normal C++ Mode	199
B.4 Extensions Accepted in C or C++ Mode	200
B.5 C++ Extensions Accepted in <code>cfront</code> Compatibility Mode	202
Appendix C Using the Compiler Message System	209
C.1 Expanding Messages with the <code>explain</code> Command	209
C.2 Controlling the Use of Messages	209
C.2.1 Command Line Options	210
C.2.2 Environment Options for Messages	210
C.2.3 <code>ORIG_CMD_NAME</code> Environment Variable	211
C.3 Message Severity	211
C.4 Common System Messages	212
Appendix D Using Intrinsic Functions	215
D.1 Atomic Memory Operations	216
D.1.1 Local Atomic Memory Operations	217
D.1.2 Global Atomic Memory Operations	218
D.2 Bit Operations	220
D.3 Mask Operations	221
D.4 Miscellaneous Operations	221
Tables	
Table 1. GCC C Language Extensions	31
Table 2. GCC C++ Language Extensions	33
Table 3. Cache Levels	45
Table 4. IPA Level	47
Table 5. File Types	48
Table 6. Floating-point Optimization Levels	51

	<i>Page</i>
Table 7. <code>-G level</code> Definitions	52
Table 8. <code>-W phase</code> Definitions	57
Table 9. <code>-Y phase</code> Definitions	59
Table 10. <code>-h pragma</code> Directive Processing	60
Table 11. Data Type Mapping	186
Table 12. Packed Characters	187
Table 13. Unrecognizable Escape Sequences	187

Examples

Example 1. <code>CC -X8 -h myprog.C</code>	68
Example 2. <code>CC -h conform myprog.C</code>	68
Example 3. <code>cc -c -h ipa1 myprog.c subprog.c</code>	68
Example 4. <code>cc -I. disc.c vend.c</code>	69
Example 5. <code>cc -P -D DEBUG newprog.c</code>	69
Example 6. <code>cc -c -h report=s mydata1.c</code>	69
Example 7. <code>CC -h ipa5,report=if myfile.C</code>	69
Example 8. Trip counts	88
Example 9. Unrolling outer loops	99
Example 10. Illegal unrolling of outer loops	99
Example 11. Using the <code>inline_enable</code> , <code>inline_disable</code> , and <code>inline_reset</code> directives	101
Example 12. Using <code>inline_reset</code>	101
Example 13. Using <code>defer_sync</code>	104
Example 14. Calling a C function from Fortran	181

The Cray Compiling Environment (CCE) contains both the Cray C and C++ compilers.

Log in either to a login node or a standalone application development system and use the Cray XE, Cray XK, or Cray XC series Programming Environment, and related products to create an application which executes on compute nodes. For further information about login nodes and the user environment, see the *Cray Programming Environment User's Guide*.

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. All compiler command options apply to Cray C and C++ unless noted.

1.1 General Compiler Description

Both the Cray C and C++ compilers are contained within the Cray Compiling Environment (CCE). If compiling code written in C, use the `cc` command to compile source files. If you are compiling code written in C++, use the `CC` command.

1.1.1 Cray C Compiler

The Cray C compiler consists of a preprocessor, a language parser, an optimizer, and a code generator. Invoke the Cray C compiler with the `cc` compiler driver command. The `cc` command is described in [cc Command on page 26](#). This command and its options are also described in the `craycc(1)` man page. See [Command Line Examples on page 68](#).

1.1.2 Cray C++ Compiler

The Cray C++ compiler consists of a preprocessor, a language parser, an optimizer, and a code generator. Invoke the Cray C++ compiler with the `CC` compiler driver command. The `CC` command is described in [CC Command on page 26](#) and the `crayCC(1)` man page. See [Command Line Examples on page 68](#).

1.2 C/C++ Standard Compatibility

The Cray C compiler conforms to the International Organization of Standards (ISO) standard ISO/IEC 9899:1999 (C99).

The Cray C++ compiler accepts the C++ language as defined by the *ISO/IEC 14882:2003* standard, with some exceptions. The exceptions are noted in [Appendix B, Using Cray C and C++ Dialects on page 195](#).

This release offers preliminary support for the UPC 1.3 draft specification. The proposed UPC 1.3 standard is discussed on the UPC specification website, <http://code.google.com/p/upc-specification>.

1.3 Binary Compatibility

UPC 1.3 requires a non-backwards compatible API change to the UPC runtime library. PGAS files compiled with CCE 8.2 are not binary compatible with those compiled with earlier versions of CCE. Attempting to link a pre-CCE 8.2 object file in the 8.2 CCE runtime environment will result in a **runtime** error message containing the text ". . . source file compiled with compiler expecting PGAS runtime API version 1.1, but linked with PGAS runtime using API version 2.0 . . . ". If you see this error, recompile the indicated source file/s with CCE 8.2 and relink your application.

1.4 Related Publications

The following documents contain additional information that may be helpful:

- `cc(1)` compiler driver man page for all C compilers
- `craycc(1)` man page for the Cray C compiler
- `CC(1)` compiler driver man page for all C++ compilers
- `crayCC(1)` man page for the Cray C++ compiler
- `intro_directives(7)` man page
- `intro_openacc(7)`
- `intro_pgas(7)`
- `ftn(1)` compiler driver man page for Fortran compilers
- `crayftn(1)` man page for the Cray Fortran compiler
- `aprun(1)` man page
- *Cray Fortran Reference Manual*
- *Cray Programming Environments Installation Guide*
- *Cray Programming Environment User's Guide*
- *Using Cray Performance Measurement and Analysis Tools*

Invoking the C and C++ Compilers [2]

This chapter describes the compiler driver commands that used to launch the Cray C and C++ compilers. The following commands invoke the compilers:

- `CC`, which invokes the Cray C++ compiler.
- `cc`, which invokes the Cray C compiler.
- `cpp`, the C language preprocessor, is not part of the Cray Compilation Environment (CCE). The `cpp` command resolves to the GNU `cpp` command and does not predefine any Cray compiler-specific macros ([Chapter 10, Using Predefined Macros on page 159](#)). If the predefinition of the Cray compiler-specific macros is required, then use the `cc` or `CC` command to do the source preprocessing using the `-E` or `-P` option.

A successful compilation creates an executable, named `a.out` by default, that reflects the contents of the source code and any referenced library functions. Use the `aprun` command to run the executable on the compute nodes.

For example, the following command sequence compiles file `mymodule.c` and launches the resulting executable program on 64 compute nodes:

```
% cc mymodule.c
% aprun -n 64 ./a.out
```

With the use of appropriate options, it is possible to direct the compiler to generate 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, it is possible to save the intermediate files and reference them later on a `CC` or `cc` command, with other files or libraries included as necessary.

By default, the `CC` and `cc` commands automatically call the linker, which creates an executable file. If only one source file is specified, the object file (`*.o`) is deleted. If more than one source file is specified, the object files are retained.

For example, the following command creates and retains object files `file1.o`, `file2.o`, and `file3.o`, and creates the executable file `a.out`:

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

The following command creates `file.o` and `a.out`; `file.o` is not retained.

```
% cc file.c
```

2.1 cc Command

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

```
.c  
.C  
.i  
.c++  
.C++  
.cc  
.cxx  
.Cxx  
.CXX  
.CC
```

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 macro[=def]] [-E] [-g] [-G level] [-h arg] [-I inldir]  
[-K trap=opt[,opt]...] [-l libfile] [-L ldir] [-M] [-nostdinc] [-o outfile]  
[-O level] [-P] [-S] [-U macro] [-V] [-Wphase,"opt..."]  
[-X npes] [-Yphase,dirname] [-#] [-##] [-###] files ...
```

For an explanation of the command line options, see [Command Line Options on page 27](#).

2.2 cc Command

The `cc` command invokes the Cray C compiler. The `cc` command accepts 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` command format is as follows:

```
cc [-c] [-C] [-D macro[=def]] [-E] [-g] [-G level] [-h arg] [-I inldir]  
[-K trap=opt[,opt]...] [-l libfile] [-L ldir] [-M] [-nostdinc] [-o outfile]  
[-O level] [-P] [-S] [-U macro] [-V] [-Wphase,"opt..."]  
[-X npes] [-Y phase,dirname] [-#] [-##] [-###] files ...
```

For an explanation of the command line options, see [Command Line Options on page 27](#).

2.3 Command Line Options

The following subsections describe options for the `CC` and `CC` commands. These options are grouped according to the following functions:

- Standard conformance options ([Standard Language Conformance Options on page 28](#))
- Virtual function options ([Virtual Function Options on page 33](#))
- General language options ([General Language Options on page 33](#))
- General optimization options ([General Optimization Options on page 36](#))
- Automatic cache management options (`-h cachem` on page 44)
- Vector optimization options ([Vector Optimization Options on page 45](#))
- Inlining options ([Interprocedural Analysis \(IPA\) Optimization Options on page 46](#))
- Scalar optimization options ([Scalar Optimization Options on page 48](#))
- Math options ([Math Options on page 50](#))
- Debugging options ([Debugging Options on page 51](#))
- Compiler message options ([Compiler Message Options on page 54](#))
- Compilation phase options ([Compilation Phase Options on page 56](#))
- Preprocessing options ([Preprocessing Options on page 59](#))
- Linker options ([Linker Options on page 62](#))
- Miscellaneous options ([Miscellaneous Options on page 63](#))
- Command line examples ([Command Line Examples on page 68](#))

Options other than those described in this manual are passed to the linker.

There are many options that start with `-h`. Specify multiple `-h` options 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 you specify conflicting options, 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. Exceptions to this rule are noted in descriptions of options or `#pragma` directives.

2.4 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.

2.4.1 `-h [no]c99 (cc)`

Default: `-h c99`

This option enables or 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 variable-length arrays (VLAs) and restricted pointers that were available as extensions previously to adoption of the C99 standard remain available to you.

2.4.2 `-h [no]conform (CC, cc), -h [no]stdc (cc)`

Default: `-h noconform, -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: By default, the compiler calls the Cray mathlib versions of intrinsic functions (`abs`, `cos`, `exp`, for example) which do not set `errno` and do not raise IEEE-754 underflow exceptions. If `-hconform` is specified, the compiler calls the `stdc glibc` versions of the runtime intrinsic functions.

2.4.3 `-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.

2.4.4 `-h [no]parse_templates (CC)`

Default: `-h noparse_templates`

This option allows existing code that defines templates using previous versions of the Cray Standard Template Library (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.

2.4.5 `-h [no]dep_name (CC)`

Default: `-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.

2.4.6 `-h [no]exceptions (CC)`

Default: 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. The `-h exceptions` option is enabled by `-h conform`.

2.4.7 `-h [no]anachronisms (CC)`

Default: `-h noanachronisms`

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

2.4.8 -h [no]new_for_init (CC)

Default: `-h new_for_init`

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; the entire `for` statement is wrapped in its own implicitly generated scope. The `-h new_for_init` option 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 for -h new_for_init
```

2.4.9 -h [no]tolerant (cc)

Default: `-h notolerant`

The `-h tolerant` option allows older, less standard C constructs, thereby making it easier to port 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 `-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 a `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.

2.4.10 -h [no]const_string_literals (CC)

Default: `-h noconst_string_literals`

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

2.4.11 -h [no]gnu

Default: `-h nognu`

The `-h gnu` option enables the compiler to recognize the subset of the GCC version 4.4.4 extensions to C listed in [Table 1](#). [Table 2](#) lists the extensions that apply only to C++.

For detailed descriptions of the GCC C and C++ language extensions, see <http://gcc.gnu.org/onlinedocs/>.

Table 1. GCC C Language Extensions

GCC C Language Extension	Description
Typeof	<code>typeof</code> : referring to the type of an expression
Lvalues	Using <code>?:</code> , and casts in lvalues
Conditionals	Omitting the middle operand of a <code>?:</code> expression
Long Long	Double-word integers – <code>long long int</code>
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	<code>'case 1 ... 9'</code> and such
Mixed Declarations	Mixing declarations and code
Attribute Syntax	Formal syntax for attributes
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	<code>\e</code> stands for the character <code><ESC></code>
Alignment	Inquiring about the alignment of a type or variable

GCC C Language Extension	Description
Inline	Defining inline functions (as fast as macros)
Alternate Keywords	<code>__const__</code> , <code>__asm__</code> , and so on, for header files
Incomplete Enums	<code>enum foo;</code> , 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: <code>nothrow</code> ; <code>format</code> , <code>format_arg</code> ; <code>deprecated</code> ; <code>used</code> ; <code>unused</code> ; <code>alias</code> ; <code>weak</code>	Declaring that functions have no side effects, or that they can never return
Variable Attributes: <code>alias</code> ; <code>deprecated</code> ; <code>unused</code> ; <code>used</code> ; <code>transparent_union</code> ; <code>weak</code> ;	Specifying attributes of variables
Type Attributes: <code>deprecated</code> ; <code>unused</code> ; <code>used</code> ; <code>transparent_union</code>	Specifying attributes of types
Asm Labels	Specifying the assembler name to use for a C symbol
Other Builtins: <code>__builtin_types_compatible_p</code> , <code>__builtin_choose_expr</code> , <code>__builtin_constant_p</code> , <code>__builtin_huge_val</code> , <code>__builtin_huge_valf</code> , <code>__builtin_huge_vall</code> , <code>__builtin_inf</code> , <code>__builtin_inff</code> , <code>__builtin_infl</code> , <code>__builtin_nan</code> , <code>__builtin_nanf</code> , <code>__builtin_nanl</code> , <code>__builtin_nans</code> , <code>__builtin_nansf</code> , <code>__builtin_nansl</code>	Other built-in functions

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

The supported subset of the GCC version 4.4.4 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 <code>using</code> directive with <code>__attribute__((strong))</code>
Explicit template specializations	Attributes may be used on explicit template specializations

2.5 Virtual Function Options

This section describes general language options.

2.5.1 `-h forcevtbl` (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 virtual function table for a class is defined in a compilation if the compilation contains a definition of the first non-inline, non-pure 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.

2.5.2 `-h suppressvtbl` (CC)

The `-h suppressvtbl` option suppresses 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.

2.6 General Language Options

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




2.6.1 -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 linker command instead of the `CC` command.

2.6.2 -h restrict=args

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

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

<i>args</i>	Description
a	All pointers to object and incomplete types are considered restricted pointers, regardless of where they appear in the source code. This includes pointers in <code>class</code> , <code>struct</code> , and <code>union</code> declarations, type casts, function prototypes, and so on.
	 Caution: Do not specify <code>restrict=a</code> 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 <code>restrict=f</code> 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 <code>this</code> pointers can be treated as restricted pointers (Cray C++ only).
	 Caution: Do not specify <code>restrict=t</code> if, during execution of any function, an object is modified and that object is referenced through the declared name of the object and a <code>this</code> pointer. Undefined behavior may result.

The *args* arguments tell the compiler to assume that, in the current compilation unit, each pointer (`=a`), 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.



Caution: The arguments make assertions about your program that, if incorrect, can introduce undefined behavior. Do 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."

2.6.3 `-h [no]calchars`

Default: `-h nocalchars`

The `-h calchars` option allows the use of the `$` character in identifier names. This option is useful for porting code containing identifiers that include this character. With `-h nocalchars`, this character is not allowed in identifier names.



Caution: Use this option with extreme care, because identifiers with this character are within CNL name space and are included in many library identifiers, internal compiler labels, objects, and functions. Prevent conflicts between identifiers within CNL name space and your code; any such conflict is an error.

2.6.4 `-h [no]signedshifts`

Default: `-h nosignedshifts`

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, see [Integers on page 188](#) about the effects of this option when shifting integers.

2.7 General Optimization Options

2.7.1 -h [no]add_paren

Default: -h noadd_paren

The -h [no]add_paren option automatically adds parenthesis to select associative operations (+, -, *) to encourage left to right evaluation of floating point and complex expressions. The default is -h noadd_paren. For more information, see the `crayftn(1)` man page.

Left to right evaluation is not required by the language standards, but some applications may expect it.

2.7.2 -h [no]aggress

Default: -h noaggress

The -h aggress option provides greater opportunity to optimize loops that would otherwise be 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.

2.7.3 -h [no]autoprefetch

Default: -h autoprefetch

The -h [no]autoprefetch option controls automatic prefetch optimization. Does not affect the `loop_info [no]prefetch` directive.

2.7.4 -h [no]autothread

Default: -h noautothread

The -h [no]autothread option enables or disables automatic threading.

2.7.5 -h display_opt

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

2.7.6 -h flex_mp=level

Default: `-h flex_mp=default`

The `-h flex_mp=level` option controls the aggressiveness of optimizations which may affect floating point and complex repeatability when application requirements require identical results when varying the number of ranks or threads.

The values for *level* are:

<i>level</i>	Description
<code>intolerant</code>	Has the highest probability of repeatable results, but also the highest performance penalty.
<code>strict</code>	Uses some safe optimizations and yields higher performance than <code>intolerant</code> , with a high probability of repeatable results.
<code>conservative</code>	Uses more aggressive optimization and yields higher performance than <code>strict</code> , but results may not be sufficiently repeatable for some applications.
<code>default</code>	Uses more aggressive optimization and yields higher performance than <code>conservative</code> , but results may not be sufficiently repeatable for some applications.
<code>tolerant</code>	Uses most aggressive optimization and yields highest performance, but results may not be sufficiently repeatable for some applications.

2.7.7 -h fusionn

Default: `-h fusion2`

The `-h fusion n` option controls loop fusion and changes the assertiveness of the `fusion` pragma. Loop fusion can improve the performance of loops, although in rare cases it may degrade performance. The *n* argument allows you to turn loop fusion on or off and determine where fusion should occur.

Note: Loop fusion is disabled when the scalar level is set to 0.

Default: `-h fusion2`

The values for *n* are:

- 0 No fusion (ignore all `fusion` pragmas and do not attempt to fuse other loops)
- 1 Attempt to fuse loops that are marked by the `fusion` pragma.
- 2 Attempt to fuse all loops (includes array syntax implied loops), except those marked with the `nofusion` pragma.

2.7.8 -h [no]intrinsic

Default: -h intrinsic

The `-h intrinsic` 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 D, Using Intrinsic Functions on page 215](#).

2.7.9 -h list

The `-h list=opt` option allows you to create listings and control their formats. The listings are written to `source_file_name_without_suffix.lst`.

The values for *opt* are:

- a Use all list options; `source_file_name_without_suffix.lst` includes a summary report, an options report, and the source listing.

- d 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 source code to improve its performance.

The compiler produces two decompilation listing files with these extensions per specified source file: `.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 quite close to what will be produced as assembly output. This version displays the intermediate text after all 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 source code. The format and contents of the files can be expected to change from release to release.

`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.

`s` Create a complete source listing (include files not expanded).

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

2.7.10 `-h [no]msgs`

Default: `-h nomsgs`

The `-h msgs` option causes the compiler to write optimization messages to `stderr`.

When the `-h msgs` option is in effect, you may request that a listing be produced so that you can see the optimization messages in the listing. For information about obtaining listings, see [-h list on page 38](#).

2.7.11 `-h [no]negmsgs`

Default: `-h nonegmsgs`

The `-h negmsgs` option causes the compiler to generate messages to `stderr` that indicate why optimizations such as vectorization, inlining, or cloning did not occur in a given instance.

The `-h negmsgs` option enables the `-h msgs` option. The `-h list=a` option enables the `-h negmsgs` option.

2.7.12 `-h [no]omp_trace`

Default: `-h noomp_trace` (tracing is off)

The `-h [no]omp_trace` option turns the insertion of the CrayPat OpenMP tracing calls on or off.

2.7.13 -h [no]func_trace

The `-h func_trace` option is for use only with CrayPat. If this option is specified, the compiler inserts CrayPat trace entry points into each function in the compiled source file. The names of the trace entry points are:

- `__pat_tp_func_entry`
- `__pat_tp_func_return`

These are resolved by CrayPat when the program is instrumented using the `pat_build` command. When the instrumented program is executed and it encounters either of these trace entry points, CrayPat captures the address of the current function and its return address.

2.7.14 -h [no]overindex

Default: `-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.

2.7.15 -h [no]pattern

Default: `-h pattern`

The `-h [no]pattern` option globally enables or disables pattern matching.

When the compiler recognizes certain patterns in the source code, it replaces the construct with a call to an optimized library routine. A loop or statement that has been pattern matched and replaced with a call to a library routine is indicated with an `A` in the loopmark listing.

Note: Pattern matching is not always worthwhile. If there is a small amount of work in the pattern-matched construct, the call overhead may outweigh the time saved by using the optimized library routine. When compiling using the default optimization settings, the compiler attempts to determine whether each given candidate for pattern matching will in fact yield improved performance.

2.7.16 -h pl=*program_library*

Create and use a persistent repository of compiler information specified by *program_library*. When used with `-hwp`, this option provides application-wide, cross-file, automatic inlining. See `-h wp` on page 43.

The *program_library* repository is implemented as a directory and the information contained in *program library* is built up with each compiler invocation. Any compilation that does not have the `-hpl` option will not add information to this repository.

Because of the persistence of *program_library*, it is the user's responsibility to manage it. For example, `rm -r program_library` might be added to the `clean` target in an application makefile. Because *program_library* is a directory, use `rm -r` to remove it.

If an application `makefile` works by creating files in multiple directories during a single build, the *program_library* should be an absolute path, otherwise multiple and incomplete program library repositories will be created. For example, avoid `-hpl=./PL.1` and use `-hpl=/fullpath/buildidir/PL.1` instead.

2.7.17 `-h profile_generate`

The `-h profile_generate` option directs that the source code be instrumented for gathering profile information. The compiler inserts calls and data-gathering instructions to allow CrayPat to gather information about the loops in a compilation unit. If you use this option, you must run CrayPat on the resulting executable so the CrayPat data-gathering routines are linked in. For information about CrayPat and profile information, see the *Using Cray Performance Measurement and Analysis Tools* guide.

2.7.18 `-h threadn`

Default: `-h thread2`

The `-h threadn` options control the optimization of both OpenMP and automatic threading.

The values of n are:

- | | |
|---|---|
| 0 | No autothreading or OMP threading. The <code>thread0</code> option is similar to <code>-h noomp</code> , but <code>-h noomp</code> disables OpenMP only and does not affect autothreading. |
| 1 | Specifies strict compliance with the OpenMP standard for directive compilation. Strict compliance is defined as no extra optimizations in or around OpenMP constructs. In other words, the compiler performs only the requested optimizations. |
| 2 | OpenMP parallel regions are subjected to some optimizations; that is, some parallel region expansion. Parallel region expansion is an optimization that merges two adjacent parallel regions in a compilation unit into a single parallel region. |
| 3 | Full optimization: loop restructuring, including modifying iteration space for static schedules (breaking standard compliance). Reduction results may not be repeatable. |

2.7.19 `-h unrolln`

Default: `-h unroll2`

The `-h unrolln` option globally controls loop unrolling and changes the assertiveness of the `unroll` pragma. By default, the compiler attempts to unroll all loops, unless the `nounroll` pragma is specified for a loop. Generally, unrolling loops increases single processor performance at the cost of increased compile time and code size.

The n argument allows you to turn loop unrolling on or off and specify where unrolling should occur. It also affects the assertiveness of the `unroll` pragma.

The values for n are:

- | | |
|---|---|
| 0 | No unrolling (ignore all <code>unroll</code> pragmas and do not attempt to unroll other loops). |
| 1 | Attempt to unroll loops that are marked by the <code>unroll</code> pragma. |
| 2 | Unroll loops when performance is expected to improve. Loops marked with the <code>unroll</code> or <code>nounroll</code> pragma override automatic unrolling. |

Note: Loop unrolling is disabled when the scalar level is set to 0.

2.7.20 `-h wp`

Enables the whole program mode. This option causes the compiler backend (IPA, optimizer, codegenerator) to be invoked at application link time, enabling whole program automatic inlining/cloning and future whole program interprocedural analysis (IPA) optimizations. Since the `-hwp` option provides automatic application-wide inlining, the `-Oipafrom` option is no longer needed for cross-file inlining and using these two options together is not permitted. Requires that `pl=program_library` is also specified. See `-h pl=program_library` on page 40.

The options `-hpl=program_library` and `-hwp` should be specified on all compiler invocations and on the compiler link invocation. Since `-hwp` delays the compiler optimization step until link time, `-c` compiles will take less time and the link step will take longer. Normally, this is just a time shift from one build phase to another with roughly the same overall compile time. In some cases increased inlining may cause an increase in overall compile time. Using `-hwp` allows the compiler backend to be invoked in parallel during a build. Setting the environment variable `NPROC` controls the number of concurrent compiler backend invocations and this parallelism may reduce overall compile time.

2.7.21 `-O level`

The `-O 0`, `-O 1`, `-O 2`, and `-O 3` options allow you to specify a general level of optimization that includes vectorization, scalar optimization, and inlining. Generally, as the optimization level increases, compilation time increases and execution time decreases.

The `-O 1`, `-O 2`, and `-O 3` specifications do not directly correspond to the numeric optimization levels for scalar optimization, vectorization, and inlining. For example, specifying `-O 3` does not necessarily enable `vector3`. Cray reserves the right to alter the specific optimizations performed at these levels from release to release. You can use the `-h display_opt` option to display the optimization options used during compilation.

The `-On` option performs general optimization at these levels: 0 (none), 1 (conservative), 2 (moderate, default), and 3 (aggressive).

- The `-O 0` option disables all optimizations including floating point optimizations. Implies `-h fp0`. This option's characteristics include low compile time, small compile size, and no global scalar optimization.

Most array syntax statements are vectorized, but all other vectorizations are disabled.

- The `-O 1` option specifies conservative optimization. This option's characteristics include moderate compile time and size, global scalar optimizations, and loop nest restructuring. Results may differ from the results obtained when `-O 0` is specified because of operator reassociation. No optimizations will be performed that might create false exceptions.

Only array syntax statements and inner loops are vectorized and the system does not perform some vector reductions. User tasking is enabled, so OpenMP directives are recognized.

- The `-O 2` option specifies moderate optimization. This option's characteristics include moderate compile time and size, global scalar optimizations, pattern matching, and loop nest restructuring.

Results may differ from results obtained when `-O 1` is specified because of vector reductions. The `-O 2` option enables automatic vectorization of array syntax and entire loop nests.

This is the default level of optimization.

- The `-O 3` option specifies aggressive optimization. This option's characteristics include a potentially larger compile time and size, global scalar optimizations, possible loop nest restructuring, and pattern matching. The optimizations performed might create false exceptions in rare instances.

Results may differ from results obtained when `-O 1` is specified because of vector reductions.

[Table 3](#) shows the equivalent level of automatic cache optimization for the `-h` option.

2.8 Automatic Cache Management Options

This section describes the automatic cache management options. Automatic cache management can be overridden by the use of the cache directives (`cache`, `cache_nt`, and `loop_info`).

2.8.1 `-h cachem`

Default: `-h cache2`

The `-h cachem` option specifies the levels of automatic cache management to perform. The default is `-h cache2`.

The values for *n* are:

- | | |
|---|--|
| 0 | Cache blocking (including directive-based blocking) is turned off. This level is compatible with all scalar and vector optimization levels. |
| 1 | Conservative automatic cache management. Characteristics include moderate compile time. Symbols are placed in the cache when the possibility of cache reuse exists and the predicted cache footprint of the symbol in isolation is small enough to experience the reuse. |

2	Moderately aggressive automatic cache management. Characteristics include moderate compile time. Symbols are placed in the cache when the possibility of cache reuse exists and the predicted state of the cache model is such that the symbol will experience the reuse.
3	Aggressive automatic cache management. Characteristics include potentially high compile time. Symbols are placed in the cache when the possibility of cache reuse exists and the allocation of the symbol to the cache is predicted to increase the number of cache hits.

Table 3. Cache Levels

-O Option	Cache Level
-O0	-h cache0
-O1	-h cache1
-O2	-h cache2
-O3	-h cache2

2.9 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.

2.9.1 -h concurrent

Indicates that no data dependence exists between array references of the same loop, for every loop in the file. This can be useful for vectorization optimizations. Equivalent to adding a `CONCURRENT pragma` before every loop in the file, including loops created from array syntax. See [concurrent Directive on page 85](#).

2.9.2 -h vectorn

Default: -h vector2

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

The values of n are:

<u>n</u>	<u>Description</u>
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. Loop nests are restructured. No vectorizations that might create false exceptions are performed. Results may differ slightly from results obtained when <code>-h vector0</code> is specified because of vector reductions. The <code>-h vector1</code> option is compatible with <code>-h scalar1</code> , <code>-h scalar2</code> , and <code>-h scalar3</code> .
2	Specifies moderate vectorization. Characteristics include moderate compile time and size. Loop nests are restructured. The <code>-h vector2</code> option is compatible with <code>-h scalar2</code> and <code>-h scalar3</code> .
3	Specifies aggressive vectorization. Characteristics include potentially high compile time and size. Loop nests are restructured. Vectorizations that might create false exceptions in rare cases may be performed.

For further information, see [Vectorization Directives on page 85](#).

2.10 Interprocedural Analysis (IPA) Optimization Options

Inlining and cloning transform code in ways that increase the opportunity for interprocedural (IPA) optimizations.

The user controls inlining and cloning through the use of command line options alone or command line options in combination with directives placed within the code. By default, the compiler will attempt inline optimizations where appropriate, but not cloning. Inlining and cloning may increase object code size.

Also see [Inlining and Cloning Directives on page 100](#).

2.10.1 Inlining

Inlining is the process of replacing a user procedure call with the procedure definition itself. This can improve performance by saving the expense of the call overhead. It also increases the possibility of additional code optimization of the inlined code. If all calls within a loop are inlined, the loop becomes a candidate for parallelization.

The compiler supports the following inlining modes through the indicated options:

- Automatic inlining allows the compiler to automatically select which functions to inline. This occurs with the `-h ipa2` or greater option. When `-h ipan` is used alone, the candidates for expansion are all those functions that are present in the input file to the compile step. See [-h ipan on page 47](#).
- Explicit inlining allows you to explicitly indicate which procedures the compiler should attempt to inline and occurs with the `-h ipafrom=source[:source]` option alone as described in [-h ipafrom=source\[:source\] . . . on page 48](#).
- Combined inlining allows you to specify potential targets for inline expansion, while applying the selected level of inlining heuristics. If `-h ipan` is used in conjunction with `-h ipafrom=source[:source]`, the candidates for expansion are those functions present in *source* and `-h ipan` selects level of heuristics.

2.10.2 Cloning

Cloning replaces a call to a procedure with a call to a modified version of that same procedure (clone) in which the parameters in the original procedure are replaced with the constant actual parameters present at the call site.

Automatic cloning is enabled at `-Oipa5`. The compiler first attempts to inline a call site. If inlining the call site fails, the compiler attempts to clone the procedure for the specific call site.

2.10.3 -h ipan

Default: `-h ipa3`

The `-h ipan` option controls the level of automatic inlining and cloning. [Table 4](#) explains what type of IPA optimization is performed at each level.

Table 4. IPA Level

IPA level	Description
0	All interprocedural analysis and optimizations disabled. All inlining and cloning compiler directives are ignored.
1	Directive IPA. Inlining/cloning is attempted for call sites and routines that are under the control of a compiler directive. See Inlining and Cloning Directives on page 100 .
2	Inlining. Inline a call site to an arbitrary depth as long as the expansion does not exceed some compiler-determined threshold. The call site must flatten for any expansion to occur. The call site is said to "flatten" when there are no calls present in the expanded code. The call site must reside within the body of a loop and the entire loop body must flatten. A loop body is said to "flatten" when all call sites within the body of the loop are flattened. Includes level 1.

IPA level	Description
3	Constant actual argument inlining and tiny routine inlining. Default level for inlining. This includes levels 1 and 2, plus any call site that contains a constant actual argument. Additionally, any call nest (regardless of location) that is below some small compiler-determined threshold will be inlined provided that call nest completely flattens. Cloning directives are recognized.
4	Aggressive inlining. This includes levels 1, 2, and 3, plus a call site does not have to reside in a loop body to inline nor does the call site have to necessarily flatten.
5	Cloning. Includes levels 1, 2, 3, 4, plus routine cloning is attempted if inlining fails at a given call site.

2.10.4 -h ipafrom=*source*[:*source*] ...

The `-h ipafrom=source[:source]` option allows you to explicitly indicate the procedures to consider for inline expansion or cloning. The *source* arguments identify each file or directory that contains the routines to consider for inlining or cloning.

Note: Spaces are not allowed on either side of the equal sign.

All inlining directives are recognized at `-O ipa levels > 1`. For information about inlining directives, see [Inlining and Cloning Directives on page 100](#).

Note: The routines in *source* are not actually linked with the final program. They are simply templates for the inliner. To have a routine contained in *source* linked with the program, you must include it in an input file to the compilation.

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

Table 5. File Types

C or C++ source files	The routines in C or C++ source files are candidates for expansion and must contain error-free code. Source files that are acceptable are files that have one of the following extensions: <code>.C</code> , <code>.c++</code> , <code>.C++</code> , <code>.cc</code> , <code>.cxx</code> , <code>.Cxx</code> , <code>.CXX</code> , or <code>.CC</code> .
<i>dir</i>	A directory that contains any of the file types described in this table.

2.11 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.

2.11.1 -h [no]interchange

Default: `-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.

2.11.2 -h scalarn

Default: `-h scalar2`

The `-h scalarn` option specifies the level of automatic scalar optimization to be performed. Scalar optimization directives are unaffected by this option (see [Scalar Directives on page 94](#)).

The values for *n* are:

- 0 Minimal automatic scalar optimization. The `-h zeroinc` option is implied by `-h scalar0`.
- 1 Conservative automatic scalar optimization. This level implies `-h nozeroinc`.
- 2 Aggressive automatic scalar optimization. The scalar optimizations that provide the best application performance are used, with some limitations imposed to allow for faster compilation times.
- 3 Very aggressive optimization; compilation times may increase significantly.

2.11.3 -h [no]zeroinc

Default: `-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 constant increment variables (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 scalarn` option (see [-h scalarn on page 49](#)).

2.12 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.

2.12.1 `-h fpn`

Default: `-h fp2`

The `-h fp` option allows you to control the level of floating-point and complex arithmetic optimizations. The n argument controls the level of allowable optimization; 0 gives the compiler minimum freedom to optimize floating-point operations, while 4 gives it maximum freedom. The higher the level, the lesser the floating-point operations conform to the IEEE standard.

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

- The `-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 you specify this level, many identity optimizations are disabled, vectorization of floating-point and complex reductions are disabled, executable code is slower than higher floating-point optimization levels. Use the `-h fp0` option only when your code pushes the limits of IEEE accuracy or requires strong IEEE standard conformance.
- The `-h fp1` option 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, a scaled complex divide mechanism is enabled that increases the range of complex values that can be handled without producing an underflow, and rewrite of division into multiplication by reciprocal is inhibited. You should never use the `-h fp1` option except when your code pushes the limits of IEEE accuracy or requires strong IEEE standard conformance.
- The `-h fp2` option includes optimizations of `-h fp1`.
- The `-h fp3` option includes optimizations of `-h fp2`. You should use the `-h fp3` option when performance is more critical than the level of IEEE standard conformance provided by `-h fp2`. The `-h fp3` option is an acceptable level of optimization for many applications.
- The `-h fp4` option includes optimizations of `-h fp3`. You should only use `-h fp4` if your application uses algorithms which are tolerant of reduced precision.

[Table 6](#) compares the various optimization levels of the `-h fp` option. The table lists some of the optimizations performed; the compiler may perform other optimizations not listed.

Table 6. Floating-point Optimization Levels

Optimization Type	fp0	fp1	fp2 (default)	fp3	fp4
Safety	Maximum	High	High	Moderate	Low
Complex divisions	Accurate and slower	Accurate and slower	Fast ¹	Fast ¹	Fast ¹
Exponentiation rewrite	None	None	When beneficial ²	Always ^{2, 3}	Always ^{2, 3}
Strength reduction	None	None	Fast	Fast	Fast
Rewrite division as reciprocal equivalent ⁴	None	None	Yes	Aggressive	Aggressive
Floating point reductions	Slow	Fast	Fast	Fast	Fast
Expression factoring	None	Yes	Yes	Yes	Yes
Expression tree balancing	None	None	Yes	Yes	Yes
Inline 32-bit operations ⁵	No	No	No	Yes	Yes
Fused multiply-add ⁶	No	Yes	Yes	Yes	Yes

2.13 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.

¹ Algebraically correct but may lack precision in boundary cases.

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 $x * 1.0/y$.

⁵ 32-bit division, square root, and reciprocal square root use very fast but less precise code sequences.

⁶ Uses fused multiply-add instructions on architectures that support it.

2.13.1 `-G level` and `-g`

The `-G level` options enable the generation of debugging information used by symbolic debuggers. These options allow debugging with breakpoints. [Table 7](#) describes the values for the `-G` option.

Table 7. `-G level` Definitions

<i>level</i>	Optimization	Breakpoints Allowed On	Debugging	Execution Speed
<code>-G0</code>	Full	Function entry and exit	Limited	Best
<code>-G1</code>	Partial	Block boundaries	Better	Better
<code>-G2</code>	None	Every executable statement	Best	Limited
<code>-G3</code>	Full	Every executable statement	Best: requires fast-track debugger	Best

Better debugging information comes at the cost of inhibiting certain optimization techniques, so choose the option that best fits the debugging needs of any particular source file in an application.

The `-g` option is equivalent to `-G0`. 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 `-g` and `-G0` options recognize OpenMP directives and disable all optimizations. They imply `-hthread1 -homp -hfp0`.

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, Debugging Cray C and C++ Code on page 165](#).

2.13.2 `-h [no]bounds (cc)`

Default: `-h nobounds`

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

For each dimension, the checks verify that the subscript is greater than or equal to 0 and less than the upper bound. For pointers, the upper bound is computed based on the amount of the memory on the node. This amount is scaled at runtime by the number of UPC threads in the job for UPC pointers-to-shared with definite blocksize. For arrays, the declared (possibly implicit) upper bound of the dimension is used. If the dimension is the THREADS-scaled dimension of a UPC shared array with definite blocksize, the upper bound for the check is computed at runtime based on the number of UPC threads in the job.

2.13.3 -h develop

Default: `off`

Reduce compile time at the expense of optimization. This option is intended to be used when a program is under development and compiled frequently. This option is different from and independent of the `-O` option. For example, `-O0` disables all optimizations, but sometimes can increase compile time because certain optimizations reduce code size, which allow other phases of the compiler to deal with less code.

2.13.4 -h dir_check

Default: `off`

The `-h dir_check` option enables directive checking at run time. Errors detected at compile time are reported during compilation and so are not reported at run time. The following directives are checked: `collapse`, and the `loop_info` clauses `min_trips` and `max_trips`. Violation of a run time check results in an immediate fatal error diagnostic.



Warning: Optimization of enclosing and adjacent loops is degraded when run time directive checking is enabled. This capability, though useful for debugging, is not recommended for production runs.

2.13.5 -h nodwarf

The `-h nodwarf` option disables DWARF generation during compilation. By default, DWARF source line information is generated to support traceback analysis. `-hdwarf` is deprecated. This option has no affect if `-g` or `-G dbg_opt` is specified.

2.13.6 -h gasp[=opt[:opt]]

Default: `disabled`

Request GASP (Global Address Space Performance Analysis) instrumentation. Requests instrumentation of events generated by shared local accesses. Instrumenting these events can add runtime overhead to the application. `#pragma pupc [on|off]` has no effect in the current GASP implementation. Possible values for `opt`:

- `local`

Requests instrumentation of events generated by shared local accesses. Instrumenting these events can add runtime overhead to the application

- `functions`

Enables function instrumentation. Sets `-hipa0`. See [-h ipan on page 47](#).

2.13.7 `-h zero`

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

2.14 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.

2.14.1 `-h msglevel_n`

Default: `-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. Values for `n` are:

0	Comment
1	Note
2	Caution
3	Warning
4	Error

2.14.2 -h [no]message=*n*[:*n*...]

Default: Determined by `-h msglevel_n`

The `-h [no]message=n[:n...]` option enables or disables specified compiler messages, where *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.

2.14.3 -h report=*args*

The `-h report=args` option generates report messages specified in *args* and lets you direct the specified messages to a file. The *args* field can be any combination of the following options:

- `f` Writes specified messages to *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`.
- `i` Generates inlining optimization messages.
- `s` Generates scalar optimization messages.
- `v` Generates vector optimization messages.

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
```

The `-h msgs` option also provides optimization messages.

2.14.4 -h [no]abort

Default: `-h noabort`

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

2.14.5 `-h errorlimit`

Default: `-h errorlimit=100`

The `-h errorlimit[=n]` option specifies the maximum number of error messages the compiler prints before it exits, where *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.

2.15 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.

2.15.1 `-E`

The `-E` option directs the compiler to execute 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 linenumber` preprocessing directives. The `-E` option takes precedence over the `-h feonly`, `-S`, and `-c` options.

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

2.15.2 `-P`

The `-P` option directs the compiler to execute 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 a `.i` suffix substituted for the suffix of the source file. The `-P` option is similar to the `-E` option, except that `#line linenumber` 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.

2.15.3 `-h feonly`

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

2.15.4 -S

The `-S` option compiles the named 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`.

2.15.5 -c

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.

2.15.6 -#, -##, and -###

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.

2.15.7 -w *phase*, "*opt ...*"

The `-w phase` option passes arguments directly to a phase of the compiling system. [Table 8](#) shows the system phases that *phase* can indicate.

Table 8. -w *phase* Definitions

<i>Phase</i>	System Phase	Command
0 (zero)	Compiler	CC and cc
a	Assembler	as
c	CUDA linker	nvlink
l	Linker	ld
x	PTX Assembler	ptxas

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 linker:

```
% 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.

The `-wl, -rpath ldir` option changes the runtime library search algorithm to look for files in directory *ldir*. To request more than one library directory, specify multiple `-rpath` options. Note that a library may be found at link time with a `-L` option, but may not be found at run time if a corresponding `-rpath` option was not supplied. Also note that the compiler driver does not pass the `-rpath` option to the linker. You must explicitly specify `-wl` when using this option.

At link time, all `ldir` arguments are added to the executable. The dynamic linker will search these paths first for shared dynamic libraries at runtime, with one exception. The Linux environment variable `LD_LIBRARY_PATH` precedes all other search paths for shared dynamically linked libraries. The use of `LD_LIBRARY_PATH` is discouraged.



Caution: Caution should be used when setting `LD_LIBRARY_PATH`. Doing so will change the shared dynamically linked library search paths for all executable files in your environment.

The `-wx, arg` option can be used to pass command line arguments to the PTX assembler for OpenACC applications.

The `-wc, arg` option can be used to pass command line arguments to the CUDA linker for OpenACC applications.

2.15.8 `-Y phase,dirname`

The `-Y phase,dirname` option specifies a new directory (*dirname*) from which the designated *phase* should be executed. The values of *phase* are [Table 9](#).

Table 9. -Y phase Definitions

<i>Phase</i>	<i>System Phase</i>	<i>Command</i>
0 (zero)	Compiler	CC, cc
a	Assembler	as
l	Linker	ld

Because there is no separate preprocessor, `-Yp` and `-Y0` are equivalent.

2.16 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.

2.16.1 -C

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 in combination with the `-P` or `-E` option.

2.16.2 -D *macro* [=def]

The `-D macro [=def]` option defines *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, Using Predefined Macros on page 159](#). Any predefined macro except those required by the standard (see [Macros Required by the C and C++ Standards on page 159](#)) 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.

2.16.3 -h [no]pragma=*name*[:*name* . . .]

Default: `-h pragma` (no pragmas disabled)

The `[no]pragma=name[:name . . .]` option enables or disables the processing of specified directives in the source code, where *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

<i>Name</i>	Group	Directives Affected
all	All	All directives
allinline	Inlining	inline_enable, inline_disable, inline_reset, inline_always, inline_never
allscalar	Scalar optimization	blockable, blockingsize, noblocking, nointerchange, suppress, unroll/nounroll
allvector	Vectorization	concurrent, novector, loop_info, hand_tuned, ivdep, nopattern, novector, permutation, pipeline/nopipeline, prefervector, safe_address, safe_conditional
omp	OpenMP	All OpenMP directives
acc	OpenACC	All OpenACC directives

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.

2.16.4 -I *includir*

The `-I includir` 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`.

2.16.5 -M

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`.

2.16.6 -nostdinc

The `-nostdinc` option stops the preprocessor from searching for include files in the standard directories (`/usr/include` and for Cray C++ also `/usr/include/c++`).

2.16.7 -U

The `-U` option removes any initial definition of *macro*. Any predefined macro except those required by the standard (see [Macros Required by the C and C++ Standards on page 159](#)) 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, Using Predefined Macros on page 159](#). Macros defined in the system headers are not predefined macros and are not affected by the `-U` option.

2.17 Linker Options

This section describes compiler options that affect linker tasks.

2.17.1 -h [system|default]_alloc

Default: `-h default_alloc`

By default, the compiler uses a modified `malloc` implementation that offers better support for memory needs. The `-h system_alloc` option directs the compiler to link in the native `malloc` provided by the OS instead of the modified implementation.

2.17.2 -h [no]pgas_runtime

Default: `-h pgas_runtime`

The `-h pgas_runtime` option directs the compiler driver to link with the runtime libraries required when linking programs that use UPC, or coarrays, which is default. In general, `aprun` must be used to launch the resulting executable.

The `-hnopgas_runtime` option prevents this runtime library environment from being added to the link line. Use the `-hnopgas_runtime` option when you have a program, that does not use UPC or coarrays, and you wish to execute it outside of the `aprun/alps` job launch context. For example, you may wish to test a serial program which does not contain any UPC or coarray code on a login or service node, or `fork/exec` an executable on a compute node. Also, compile non-coarray Fortran using the `-hnocaf` option.

2.17.3 -l *libname*

The `-l libname` option directs the compiler driver to search for the specified object library file when linking an executable file. To request more than one library file, specify multiple `-l` options.

When statically linking, the compiler driver searches for libraries by prepending *ldir/lib* on the front of *libname* 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 ([-L *ldir* on page 63](#)).

When dynamically linking, the library search process is similar to the static case, with a few differences. The compiler driver searches for libraries by prepending *ldir/lib* on the front of *libname* and appending *.so* on the end of it, for each *ldir* that has been specified by using the `-L` option. If a matching *.so* is not found, the compiler driver replaces *.so* with *.a* and repeats the process from the beginning. It uses the first file it finds. See also the `-L` option ([-L *ldir* on page 63](#)).

There is no search order dependency for libraries.

If you specify personal libraries by using the `-l` command line option, as in the following example, those libraries are added before the default CCE library list. (The `-l` option is passed to the linker.)

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

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

2.17.4 `-L ldir`

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

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

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

2.17.5 `-o outfile`

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 source file, a relocatable object file named *outfile* is produced.

2.18 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.

2.18.1 -h [no]acc

Default: -h acc

The `-h [no]acc` option enables or disables compiler recognition of OpenACC pragmas.

2.18.2 -h cpu=*target_system*

The `-h cpu=target_system` option specifies the Cray system on which the absolute binary file is to be executed, where *target_system* can be either `x86-64` or `opteron` (single or dual-core), `barcelona` or `shanghai` (quad-core), `istanbul` (6-core), `mc8` (8-core), `mc12` (12-core), `interlagos` (16-core), `interlagos-cu` (8-compute unit), `abudhabi` (16-core), `abudhabi-cu` (8-compute unit), `ivybridge`, `sandybridge`, or `haswell`.

The `interlagos` and `abudhabi` processors contain up to 8 compute units, each of which contains two integer cores and a shared FPU. These targets assume that the user intends to run with one thread per core (up to 16 per processor), while the `cpu-cu` target assumes that the user intends to run with one thread per compute unit (up to 8 per processor or one thread per FPU).

Rather than setting this option directly, users should load one of the targeting modules (`craype-mc12` or `craype-interlagos-cu`, for example). The targeting modules set `CRAY_CPU_TARGET` and define paths to the corresponding libraries. The compiler driver script translates `CRAY_CPU_TARGET` to the corresponding `cpu=target_system` option when calling the compiler.

If the *target_system* is set during compilation of any source file, it must also be set to that same target during linking and loading.

If a user wishes to override the current *target_system* value set by the module environment (via the `CRAY_CPU_TARGET` definition), they should do so by specifying `-hcpu=target_system` on the compiler command line.

See [Compile Time Environment Variables on page 69](#). For more information, see the *Cray Application Developer's Environment User's Guide*.

2.18.3 -h [no]fp_trap

Controls whether the compiler generates code that is compatible with floating-point traps.

Default: `fp_trap`, if traps are enabled using the `-K trap` option, or if `-Ofp[0,1]` is in effect. Otherwise, the default is `nofp_trap`. See [-K trap=*opt*\[, *opt*\] ... on page 67](#).

2.18.4 -h ident=*name*

Default: 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 *name* is specified or the default name is used, the following transformations are performed on *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.

2.18.5 -h keepfiles

The `-h keepfiles` option prevents the removal of the object (`.o`) and temporary assembly (`.s`) 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 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.

2.18.6 -h keep_frame_pointer

Retain call stack information back to main entry point for CrayPat performance sampling. Prevents call stack frame from being optimized out of a function so CrayPat performance sampling is able to trace call stack back to entry point.

Default: `off`

2.18.7 -h loop_trips=[*tiny* | *small* | *medium* | *large* | *huge*]

Specifies runtime loop trip counts for all loops in a compiled source file. This information is used to optimize the runtime characteristics of the application.

2.18.8 -h mpin

Enables or disables optimization of MPI operations. `-h mpi1` enables this option.

Default: `mpi0`

2.18.9 -h network=*nic*

The `-h network=nic` option specifies the target machine's system interconnection network. Currently, supported values for *nic* are `gemini` and `aries`.

2.18.10 `-h [no]omp`

Default: `-h omp` (if `-O1` or higher is implied or specified)

The `-h [no]omp` option enables or disables compiler recognition of OpenMP pragmas. If `-O0` is specified, then `-h noomp` is implied. For details, see [Chapter 4, Using the OpenMP C/C++ API on page 105](#).

2.18.11 `-h [no]omp_acc`

Default: `-h omp_acc`

The `-h [no]omp_acc` option enables or disables compiler recognition of OpenMP acc pragmas.

2.18.12 `-h pic, -h PIC`

Generate position independent code (PIC), which allows a virtual address change from one process to another, as is necessary in the case of shared, dynamically linked objects. The virtual addresses of the instructions and data in PIC code are not known until dynamic link time. For the Cray implementation, the `pic` and `PIC` options have the same effect and should be used to compile codes using more than 2GB of static memory, or for creating dynamically linked libraries.

2.18.13 `-h prototype_intrinsics`

The `-h prototype_intrinsics` option 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. For details, see [Appendix D, Using Intrinsic Functions on page 215](#).

2.18.14 `-h [no]threadsafe`

Default: `-h threadsafe`

The `-h [no]threadsafe` option enables or disables the generation of threadsafe code. Code that is threadsafe can be used with `pthread` and OpenMP. This option is not binary-compatible with code generated by Cray C 8.1 or 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 8.1 or Cray C++ 5.1 and earlier compilers at the expense of not being threadsafe.

C or C++ code compiled with `-h threadsafe` (the default) cannot be linked with C or C++ code compiled with `-h nothreadsafe` or with code compiled with a Cray C 8.1, Cray C++ 5.1, or earlier compiler.

2.18.15 `-h upc (cc)`

Default: `off`

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 implementation of UPC is discussed in [Chapter 6, Using Cray Unified Parallel C \(UPC\) on page 125](#).

2.18.16 `-K trap=opt[,opt] ...`

Enable traps for the specified exceptions. By default, no exceptions are trapped. Enabling traps by using this option also has the effect of setting `-h fp_trap`.

If the specified options contradict each other, the last option predominates. For example, `-K trap=none,fp` is equivalent to `-K trap=fp`.

This option does not affect compile time optimizations; it detects runtime exceptions. This option is processed only at link time and affects the entire program; it is not processed when compiling subprograms. Therefore, traps may be set using this command line option at the beginning of execution of the main program only. The program may subsequently change these settings by calling intrinsic or library procedures. Use of this option may require the specification of `-hfp_trap` when compiling other files of the application.

See [-h \[no\]fp_trap on page 64](#).

<i>opt</i>	Exceptions
<code>denorm</code>	Trap on denormalized operands.
<code>divz</code>	Trap on divide-by-zero.
<code>fp</code>	Trap on <code>divz</code> , <code>inv</code> , or <code>ovf</code> exceptions.
<code>inexact</code>	Trap on inexact result (i.e., rounded result). Enabling traps for inexact results is not recommended.
<code>inv</code>	Trap on invalid operation.
<code>none</code>	Disables all traps (default).
<code>ovf</code>	Trap on overflow (i.e., the result of an operation is too large to be represented).
<code>unf</code>	Trap on underflow (i.e., the result of an operation is too small to be represented).

2.18.17 -v

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
/opt/cray/xt-asyncpe/2.5/bin/CC: INFO: native target is being used
Cray C++ : Version 7.1.0.129 Thu May 21, 2009 12:59:44
```

2.18.18 -x npes

The `-x npes` option specifies the number of processing elements (PEs) that will be specified through `aprun` at job launch. The value for `npes` ranges from 1 through $2^{*}31 - 1$ inclusive.

Ensure that you compile all object files with the same `-x npes` value and run the resulting executable with that number of PEs. If you use mixed `-x npes` values or if the number of PEs provided at run time differs from the `-x npes` value, you will receive a run time error.

You cannot change the number of PEs to use at link or run time. You must recompile the program with a different value for `npes` to change the number of PEs.

For further information about running applications, see the *Cray Application Developer's Environment User's Guide* or the `aprun(1)` man page.

2.19 Command Line Examples

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

Example 1. `CC -x8 -h myprog.C`

This example compiles `myprog.C` and fixes the number of processing elements to 8.

```
% CC -x8 myprog.C
```

Example 2. `CC -h conform myprog.C`

This example compiles `myprog.C`. The `-h conform` option specifies strict conformance to the ISO C++ standard.

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

Example 3. `cc -c -h ipa1 myprog.c subprog.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 linker. Option `-h ipa1` tells the compiler to inline function calls marked with the `inline_always` pragma.

```
% cc -c -h ipa1 myprog.c subprog.c
```

Example 4. `cc -I. disc.c vend.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
```

Example 5. `cc -P -D DEBUG newprog.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
```

Example 6. `cc -c -h report=s mydata1.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
```

Example 7. `CC -h ipa5,report=if myfile.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
```

2.20 Compile Time Environment Variables

The following environment variables are used during compilation.

<u>Variable</u>	<u>Description</u>
-----------------	--------------------

`CRAYOLDCPPLIB`

When set to a nonzero value, enables C++ code to use the following 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. For more information, see [Appendix B, Using Cray C and C++ Dialects on 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`

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.

Identifies your requirements for native language, local customs, and coded character set with regard to compiler messages.

Controls the format in which you receive compiler messages.

Specifies the message system catalogs that should be used.

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.

2.21 Run Time Environment Variables

`CRAY_MALLOPT_OFF`

If set, then the system default `malloc` parameters are used, instead of the compiler default parameters. For most programs, run time performance is improved by using the compiler defaults, but more memory may be used.

`MALLOC_MMAP_MAX`

Specifies the maximum number of memory chunks to allocate with `mmap`. The compiler default value is 0. For most programs, run time performance is improved by using the compiler default, but more memory may be used.

`MALLOC_TRIM_THRESHOLD`

Specifies the minimum size of the unused memory region at the top of the heap before the region is returned to the operating system. The compiler default value is 536870912 bytes. For most programs, run time performance is improved by using the compiler default, but more memory may be used.

`PGAS_ERROR_FILE`

Specifies the location to which `libpgas` (the library which provides an interface to the internal system network) error messages are written. The default is `stderr`. If `stdout` is specified, errors will be written to standard output.

2.22 OpenMP Environment Variables

For Cray-specific information about OpenMP environment variables, see [Chapter 4, Using the OpenMP C/C++ API on page 105](#). For documentation of standard OpenMP environment variables, see the *OpenMP Application Program Interface Version 3.1 July 2011* standard (<http://openmp.org/wp/openmp-specifications/>).

Using #pragma Directives [3]

The #pragma directives are used within the source program to request certain kinds of special processing. The 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; it 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 ([General Directives on page 75](#))
- Vectorization ([Vectorization Directives on page 85](#))
- Scalar ([Scalar Directives on page 94](#))
- Inlining ([Inlining and Cloning Directives on page 100](#))

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 4, Using the OpenMP C/C++ API on page 105](#).

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* is 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 concurrent
```

The following is the same example using the alternative form:

```
_Pragma( "_CRI concurrent" )
```

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

```
#define _str( _X ) # _X
#define COPY( _A, _B, _N

{
    int i;
    _Pragma( "_CRI concurrent" )
    _Pragma( _str( _CRI loop_info cache_nt( _B ) ) )
    for ( i = 0; i < _N; i++ ) {
        _A[i] = _B[i];
    }
}

void
copy_data( int *a, int *b, int n )
{
    COPY( a, b, n );
}
```

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]autothread Directive

Scope: Local

The `autothread` and `noautothread` directives turn autothreading on and off for selected blocks of code.

The format of these directives is as follows:

```
#pragma _CRI autothread
#pragma _CRI noautothread
```

3.5.2 [no]bounds Directive

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

For each dimension, the checks verify that the subscript is greater than or equal to 0 and less than the upper bound. For pointers, the upper bound is computed based on the amount of the memory on the node. This amount is scaled at runtime by the number of UPC threads in the job for UPC pointers-to-shared with definite blocksize. For arrays, the (possibly implicit) declared upper bound of the dimension is used. If the dimension is the THREADS-scaled dimension of a UPC shared array with definite blocksize, the upper bound for the check is computed at runtime based on the number of UPC threads in the job.

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];
    .
    .
    .
}
```

3.5.3 cache Directive

The `cache` directive asserts that all memory operations with the specified symbols as the base are to be allocated in cache. This is an advisory directive. The `cache` directive is meaningful for stores in that it allows the user to override a decision made by the automatic cache management. This directive may be locally overridden by the use of a `#pragma loop_info` directive. This directive overrides automatic cache management decisions (see `-h cachem`).

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

The format of the `cache` directive is:

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

base_name The base name of the object that should be placed into the cache. This can be the base name of any object such as an array, scalar structure, and so on, without member references like `C[10]`. If you specify a pointer in the list, only the references, not the pointer itself, are cached.

3.5.4 cache_nt Directive

The `cache_nt` directive is an advisory directive that specifies objects that should use non-temporal reads and writes. Use this directive to identify objects that should not be placed in cache.

The format of the `cache_nt` directive is:

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

base_name The base name of the object that should use non-temporal reads and writes. This can be the base name of any object such as an array, scalar structure, and so on, without member references like `C[10]`. If you specify a pointer in the list, only the references, not the pointer itself, have the cache non-temporal property.

This directive overrides the automatic cache management level that was specified using the `-h cachem` option on the compiler command line. This directive may be overridden locally by use of a `loop_info` directive.

3.5.5 duplicate Directive

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
```

3.5.6 `ident` Directive

The `ident` pragma 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
```

3.5.7 `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
```

3.5.8 `[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();
}
```

3.5.9 prefetch Directive

The general `prefetch` directive instructs the compiler to generate explicit `prefetch` instructions which load data from memory into cache prior to read or write access. The memory location to be prefetched is defined by `var`, which specifies any valid variable, member, or array element reference. The format of this directive is as follows:

```
#pragma _CRI prefetch [[lines(num)][, level(num)] [,
write][, nt]] var[, var]...
```

The `prefetch` directive supports the following options:

- `lines(num)` Specifies the number of cache lines to be prefetched. *num* is an expression that evaluates to an integer constant at compilation time. By default, the number of cache lines prefetched is 1.
- `level(num)` Specifies the level of cache into which data is loaded. *num* is an expression that evaluates to an integer constant at compilation time. The cache level defaults to 1, the level closest to the processing unit. This level specification has little effect for current x86 targets.
- `write` Specifies that the `prefetch` is for data to be written. When data is to be written, a `prefetch` instruction can move a block into

the cache so that the expected store will be to the cache. Prefetch for write generally brings the data into the cache in an exclusive or modified state. By default, the `prefetch` is for data to be read. If the target architecture does not support `prefetch` for write, the `prefetch` will automatically become a `prefetch` for read.

`nt` Specifies that the `prefetch` is for non-temporal data. By default, the `prefetch` is for temporal data. Data with temporal locality (persistence), is expected to be accessed multiple times.

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

```
void
add( long * restrict a, long * restrict b, const int n )
{
    int i;

#pragma _CRI prefetch (lines(2)) b[0]

    for ( i = 0; i < n; i++ ) {
#pragma _CRI prefetch b[i+16]
        a[i] += b[i];
    }

    return;
}
```

The compiler issues the `prefetch` instruction when it encounters the `prefetch` directive. The directive allows the user to influence almost every aspect of `prefetch` behavior. The default behavior prefetches one cache line, into L1 cache, for read access, and assumes temporal locality.

The `prefetch` directive can be used inside and outside of loops, in a loop preamble, or before a function call to reduce cache-miss memory latency.

The compiler will attempt to avoid multiple prefetches to the same cache line, which can be created as a result of optimization.

All variables specified on the same `prefetch` directive line share the same behavior. If different behavior is needed for different variables, use multiple `prefetch` directive lines.

The general `prefetch` directive supersedes the effects of any relevant `loop_info` `[no]prefetch` directives and the `-h [no]autoprefetch` command line option.

The Cray Fortran compiler command line option `-x prefetch` can be used to disable all general `prefetch` directives in Fortran source code. The Cray C and C++ compiler command line option `-h nopragma=prefetch` can be used to disable all general `prefetch` directives in C and C++ source code.

3.5.10 Probability Directives

The `probability`, `probability_almost_always`, and `probability_almost_never` directives specify information used by interprocedural analysis (IPA) and the optimizer to produce faster code sequences. The specified probability is a hint, rather than a statement of fact. You can also specify `almost_never` and `almost_always` by using the values 0.0 and 1.0, respectively.

These directives have the following format:

```
#pragma probability const
#pragma probability_almost_always
#pragma probability_almost_never
```

const is an expression that evaluates to a floating point constant at compilation time. (0.0 <= *const* <= 1.0.)

These directives can appear anywhere executable code is legal.

Each directive applies to the block of code where it appears. It is important to realize that the directive should not be applied to a conditional test directly; rather, it should be used to indicate the relative probability of a `then` or `else` branch being executed.

Example:

```
    if ( a[i] > b[i] ) {
#pragma probability 0.3
        a[i] = b[i];
    }
```

This example states that the probability of entering the block of code with the assignment statement is 0.3 or 30%. This also means that `a[i]` is expected to be greater than `b[i]` 30% of the time.

Note that the `probability` directive appears within the conditional block of code, rather than before it. This removes some of the ambiguity that has plagued other implementations that tie the directive directly to the conditional code.

This information is used to guide inlining decisions, branch elimination optimizations, branch hint marking, and the choice of the optimal algorithmic approach to the vectorization of conditional code.

The following GCC-style intrinsic is also accepted when it appears in a conditional test:

```
__builtin_expect( expr, const )
```

The following example:

```
if ( __builtin_expect( a[i] > b[i], 0 ) ) {  
    a[i] = b[i];  
}
```

is roughly equivalent to:

```
    if ( a[i] > b[i] ) {  
#pragma _CRI probability_almost_never  
        a[i] = b[i];  
    }
```

3.5.11 weak Directive

Scope: Global

When statically linking, the `weak` directive specifies an external identifier that may remain unresolved throughout the compilation. This directive has no effect when dynamically linking. A weak external reference can be a reference 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 linker 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 as an unsatisfied external. The linker 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 linker 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 linker (using some "strong" reference elsewhere) has actually linked 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 */
}
```

3.6 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.

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

3.6.1 concurrent Directive

Scope: Local

The `concurrent` directive indicates that no data dependence exists between array references in different iterations of the loop. This directive affects the loop that immediately follows it. This can be useful for vectorization 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 integer constant > 0.

If `SAFE_DISTANCE=n` is not specified, the distance is assumed to be infinite, and the compiler ignores all cross-iteration dependencies.

The `concurrent` directive is ignored if the `safe_distance` clause is used and 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 loop iteration *i*.

```
#pragma _CRI concurrent safe_distance=3

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

3.6.2 `hand_tuned` Directive

The format of this directive is:

```
#pragma _CRI hand_tuned
```

This directive asserts that the code in the loop nest has been arranged by hand for maximum performance, and the compiler should restrict some of the more aggressive automatic expression rewrites. The compiler should still fully optimize and vectorize the loop within the constraints of the directive.

The `hand_tuned` directive applies to the next loop in the same manner as the `concurrent` and `safe_address` directives.



Warning: Use of this directive may severely impede performance. Use carefully and evaluate before and after performance.

3.6.3 `ivdep` Directive

When the `ivdep` directive appears before a loop, the compiler ignores vector dependencies, including explicit dependencies, in any attempt to vectorize the loop. `ivdep` applies to the first `for` loop or `while` loop that follows the directive. The directive applies to only the first loop that appears after the directive within the same program unit.

Whether or not `ivdep` is used, conditions other than vector dependencies can inhibit vectorization. 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. That is, no dependency will occur at any vector length.

If no vector length is specified, the vector length used is infinity.

If a loop with an `ivdep` directive is enclosed within another loop with an `ivdep` directive, the `ivdep` directive on the outer loop is ignored.

When the Cray compiler vectorizes a loop, it may reorder the statements in the source code to remove vector dependencies. When `ivdep` is specified, the statements in the loop or array syntax statement are assumed to contain no dependencies as written, and the Cray compiler does not reorder loop statements.

3.6.4 `loop_info` Directive

Scope: Local

The `loop_info` directive allows additional information to be specified about the behavior of a loop, including run time trip count and hints on cache allocation strategy.

`loop_info` directive provides information to the optimizer and can produce faster code sequences. `loop_info` is used immediately before a `for` loop to indicate minimum, maximum, or estimated trip count. The compiler will diagnose misuse at compile time (when able) or when option `-h dir_check` is specified at run time.

For cache allocation hints, the `loop_info` directive can be used to override default settings, `cache` or `cache_nt` directives, or override automatic cache management decisions. The cache hints are local and apply only to the specified loop nest.

The format of this directive is:

```
#pragma _CRI loop_info
[min_trips(c)] [est_trips(c)] [max_trips(c)] [cache(
symbol[, symbol ...] )][cache_nt(symbol[, symbol ...] )
][prefetch] [noprefetch]
```

<i>c</i>	An expression that evaluates to an integer constant at compilation time.
<i>min_trips</i>	Specifies guaranteed minimum number of trips.
<i>est_trips</i>	Specifies estimated or average number of trips.
<i>max_trips</i>	Specifies guaranteed maximum number of trips.
<i>cache</i>	Specifies that <i>symbol</i> is to be allocated in cache; this is the default if no hint is specified and the <i>cache_nt</i> directive is not specified.
<i>cache_nt</i>	Specifies that <i>symbol</i> is to use non-temporal reads and writes.
<i>prefetch</i>	Specifies a preference that prefetches be performed for the following loop.
<i>noprefetch</i>	Specifies a preference that no prefetches be performed for the following loop.
<i>symbol</i>	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 or scalar structure) without member references like <code>C[10]</code> . If you specify a pointer in the list, only the references, not the pointer itself, have the no cache allocate property.

Example 8. Trip counts

In the following example, the minimum trip count is 1 and the maximum trip count is 1000:

```
void
loop_info( double *restrict a, double *restrict b, double s1, int n )
{
    int i;

    #pragma _CRI loop_info min_trips(1) max_trips(1000), cache_nt(b)
    for (i = 0; i < n; i++) {
        if(a[i] != 0.0) {
            a[i] = a[i] + b[i]*s1;
        }
    }
}
```


3.6.5 loop_info prefer_thread, prefer_nothread Directives

Scope: Local

Use these directives to indicate a preference for turning threading on or off for selected loops. Use the `loop_info prefer_thread` directive to indicate your preference that the loop following the directive be threaded. The `loop_info prefer_nothread` indicates your preference that the loop following the directive should not be threaded.

The format of these directives is:

```
#pragma _CRI loop_info prefer_thread
#pragma _CRI loop_info prefer_nothread
```

3.6.6 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];
            }
        }
    }
}
```

3.6.7 [no]vector Directive

[no]vector controls vectorization of for loops. It may affect specific optimizations.

```
#pragma _CRI vector [clause[, clause]... ]
#pragma _CRI novector
```

The novector directive suppresses compiler attempts to vectorize loops and array syntax statements. It overrides any other vectorization-related directives, as well as the -h vector and -O vector*n* command line options. These directives are ignored if vectorization or scalar optimization has been disabled.

In C/C++, the novector directive applies only to the following loop. When applied to an outer loop in a nest, the directive also applies to all inner loops.

After a vector directive is specified, automatic vectorization is enabled for all loop nests.

The vector directive supports the following optional clauses:

- | | |
|-----------|--|
| always | Vectorize the loop that immediately follows the directive. This directive states a vectorization preference and does not guarantee that the loop has no memory-dependence hazard. This directive has the same effect as the prefervector(7) directive. |
| aligned | Directs the compiler to generate aligned data movement instructions for array references when vectorizing. For current INTEL processors, data alignment is necessary for efficient vectorization. Use with care to improve performance. If some of the access patterns are actually unaligned, using the ALIGNED clause may generate incorrect code. This directive also directs the compiler to ignore explicit and implicit vector dependencies. |
| unaligned | Directs the compiler to generate unaligned data movement instructions for all array references when vectorizing. |

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];
}
```

3.6.8 permutation Directive

The permutation directive specifies that an integer array has no repeated values. This directive is useful when the integer array is used as a subscript for another array (vector-valued subscript). This directive may improve code performance.

This directive has the following format:

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

In a sequence of array accesses that read array element values from the specified symbols with no intervening accesses that modify the array element values, each of the accessed elements will have a distinct value.

When an array with a vector-valued subscript appears on the left side of the equal sign in a loop, many-to-one assignment is possible. Many-to-one assignment occurs if any repeated elements exist in the subscripting array. If it is known that the integer array is used merely to permute the elements of the subscripted array, it can often be determined that many-to-one assignment does not exist with that array reference.

Sometimes a vector-valued subscript is used as a means of indirect addressing because the elements of interest in an array are sparsely distributed; in this case, an integer array is used to select only the desired elements, and no repeated elements exist in the integer array, as in the following example:

```
int *ipnt;
#pragma permutation ipnt
...
for ( i = 0; i < N; i++ ) {
    a[ipnt[i]] = b[i] + c[i];
}
```

The permutation directive does not apply to the array `a`. Rather, it applies to the pointer used to index into it, `ipnt`. By knowing that `ipnt` is a permutation, the compiler can safely generate an unordered scatter for the write to `a`.

3.6.9 [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
```

3.6.10 `prefervector` Directive

Scope: Local

The `prefervector` pragma directs the compiler to vectorize the loop immediately following 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:

```
#pragma _CRI prefervector
```

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

```
float a[1000], b[100][1000];

void
f(int m, int n)
{
    int i, j;

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

In this 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 would vectorize the inner loop.

3.6.11 `pgo_loop_info` Directive

Scope: Local

The format of this directive is as follows:

```
#pragma _CRI pgo loop_info
```

The `pgo_loop_info` directive enables profile-guided optimizations by tagging loopmark information as having come from profiling. For information about CrayPat and profile information, see the *Using Cray Performance Measurement and Analysis Tools* guide.

3.6.12 `safe_address` Directive

Scope: Local

The format of this directive is as follows:

```
#pragma _CRI safe_address
```

The `safe_address` directive specifies 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 `-h report=v` or `-h msgs` 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, you 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 safely load `b[i][j]` as a vector, merge `0.0` where the condition is true, and store the resulting vector safely.

3.6.13 `safe_conditional` Directive

The `safe_conditional` directive specifies that it is safe to execute all references and operations 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. This directive specifies that memory and arithmetic operations are safe.

This directive applies to scalar and vector loop nests. It can improve performance by allowing the hoisting of invariant expressions from conditional code and by allowing prefetching of memory references.

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



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

The `safe_conditional` directive has the following format:

```
#pragma _CRI safe_conditional
```

In the following example, without the `safe_conditional` directive, the compiler cannot precompute the invariant expression `s1*s2` because their values are unknown and may cause an arithmetic trap if executed unconditionally. However, if you know that the condition is true at least once, then `s1*s2` is safe to speculatively execute. The `safe_conditional` compiler directive can be used to imply the safety of the operation. With the directive, the compiler evaluates `s1*s2` outside of the loop, rather than under control of the conditional code. In addition, all control flow is removed from the body of the vector loop, because `s1*s2` no longer poses a safety risk.

```
void
safe_cond( double a[restrict 1000], double s1, double s2 )
{
    int i;

    #pragma _CRI safe_conditional
    for (i = 0; i < 1000; i++) {
        if( a[i] != 0.0) {
            a[i] = a[i] + s1*s2;
        }
    }
}
```

3.7 Scalar Directives

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

3.7.1 blockable Directive

The `blockable` directive specifies that it is legal to cache block the subsequent loops.

The format of this directive is as follows:

```
#pragma _CRI blockable (nest-depth)
```

where `nest-depth` specifies the depth of the loop nest to be blocked.

This directive instructs the compiler to perform a cache blocking rewrite in which the following *nest-depth* loops participate. The nest to be transformed must be a fully-permutable, perfect, rectangular nest. A fully-permutable nest is a nest that may be legally interchanged in any order. To be perfect, a nest must have no statements between the participating loops. Rectangular means that the lower and upper bounds (and stride) of each loop must be independent of the value of the loop control variables of all enclosing member loops in the rewrite set. Though the compiler will catch some misuse of this directive, it will not detect all misuses. The `blockingsize` directive may be used to control the size and shape of the cache block, otherwise known as a tile.

3.7.2 blockingsize Directive

When a loop is annotated with a `blockingsize` directive, and this loop is to be cache blocked automatically or by `blockable` directive, the indicated blocking factors are employed. In the absence of a `blockingsize` directive, the compiler will select blocking factors.

The compiler attempts to include this loop within cache, but it cannot guarantee this.

The formats of this directive is as follows:

```
#pragma _CRI blockingsize (n1, [, n2])
```

where *n1* and *n2* are integer constants that indicate the block size, with $0 \leq n1 \leq n2 \leq 2^{*}30$.

Note: The Cray compiler only blocks one level of cache – the secondary cache, which is specified by *n2*. If a single is block size is specified, it is interpreted as the blocking factor for the secondary cache.

A loop with `blockingsize` *n* > 1 is strip mined to length *n* as shown:

```
do ii = 1, trips, n
  do i = ii, min( ii + n - 1, trips )
```

To fully understand the interaction of blockingsizes, it is helpful to examine the two steps employed by the compiler:

1. Strip mine `blockable` nest members according to `blockingsize`
 - if $n > 1$, strip mine by n , creating `loop_outer` and `loop_inner`.
 - if $n = 0$, do not strip mine, treat loop as `loop_inner`. Full tripcount. The entire loop is inside the block.
 - if $n = 1$, do not strip mine, treat loop as `loop_outer`.
2. Interchange the loops resulting from previous step so all "outer" loops are outside the "inner" loops and the relative order within each subset is preserved.

3.7.3 `noblocking` Directive

The `noblocking` directive asserts that the loop following the directive should not be cache blocked for the primary or secondary cache. It is an error to place a `noblocking` directive before a loop that is part of a `blockable` collection.

The format of this directive is as follows:

```
#pragma _CRI noblocking
```

3.7.4 `collapse` and `nocollapse` Directives

Scope: Local

The loop collapse directives control collapse of the immediately following loop nest.

The formats of these directives are as follows:

```
#pragma _CRI collapse(loop-number1,loop-number2[,loop-number3] ...)
#pragma _CRI nocollapse
```

When the `collapse` directive is applied to a loop nest, the loop numbers of the participating loops must be listed in order of increasing access stride. Loop numbers range from 1 to the nesting level of the most deeply nested loop. The directive enables the compiler to assume appropriate conformity between trip counts. The compiler diagnoses misuse at compile time (when able); or, if `-h dir_check` is specified, at run time.

The `nocollapse` directive disqualifies the immediately following loop from collapsing with any other loop. Collapse is almost always desirable, so use this directive sparingly.

Loop collapse is a special form of loop coalesce. Any perfect loop nest may be coalesced into a single loop, with explicit rediscovery of the intermediate values of original loop control variables. The rediscovery cost, which generally involves integer division, is quite high. Therefore, coalesce is rarely suitable for vectorization. It may be beneficial for multithreading.

By definition, loop collapse occurs when loop coalesce may be done without the rediscovery overhead. To meet this requirement, all memory accesses must have uniform stride.

3.7.5 [no]interchange Directive

Scope: Local

The loop interchange control directives specify whether or not the order of the following two or more loops should be interchanged. These directives apply to the loops that they immediately precede.

The formats of these directives are as follows:

```
#pragma _CRI interchange(loop_number1, loop_number2[, loop_number3] ...)
#pragma _CRI nointerchange
```

The first format specifies two or more loop numbers. Loop numbers range from 1 to nesting depth of the most deeply nested loop. They can be specified in any order, and the compiler reorders the loops. The loops must be perfectly nested. If the loops are not perfectly nested, you may receive unexpected results. The compiler reorders the loops such that the loop with *loop-number1* is outermost, then *loop-number2*, then *loop-number3*.

The second format inhibits loop interchange on the loop that immediately follows the directive.

In the following example, the `interchange` directive reorders the loops; the `k` loop becomes the outermost and the `i` loop the innermost:

```
#define N 100

A[N][N][N];

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

#pragma _CRI interchange( 2, 3, 1 )
    for (i=0; i < n; i++) {
        for (k=0; k < n; k++) {
            for (j = 0; j < n; j++) {
                A[k][j][i] = 1.0;
            }
        }
    }
}
```

3.7.6 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.

3.7.7 [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 formats for these compiler directives are:

```
#pragma _CRI unroll n  
#pragma _CRI nounroll
```

The `nounroll` directive disables loop unrolling for the next loop. The `nounroll` directive is functionally equivalent to the `unroll 0` and `unroll 1` 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 ($2 \leq n \leq 63$).

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 10](#)).

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 9. Unrolling outer loops

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

```
#pragma _CRI unroll 2
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 10. 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 unroll 2
for (i = 0; i < 10; i++) {
    for (j = 1; j < 100; j++) {
        a[i][j] = a[i+1][j-1] + 1;
    }
}
```

3.7.8 `nofission` Directive

Scope: Local

The `nofission` directive instructs the compiler not to split statements in a given loop into distinct loops. Fission is prevented only for the loop specified; loops nested within the indicated loop remain fission candidates unless likewise annotated.

3.7.9 `[no]fusion` Directive

Scope: Local

The `nofusion` directive instructs the compiler to not attempt loop fusion on the following loop even when the `-h fusion` option was specified on the compiler command line. The `fusion` directive instructs the compiler to attempt loop fusion on the following loop unless `-h nofusion` was specified on the compiler command line.

The formats for these compiler directives are:

```
#pragma _CRI fusion
#pragma _CRI nofusion
```

3.8 Inlining and Cloning Directives

Inlining and cloning directives can only appear in local scope — inside a function definition.

Inlining directives always take precedence over the command line settings with the exception of `-h ipa0`, which instructs the compiler to ignore inlining directives.

Cloning directives are enabled with `-h ipa5`.

3.8.1 `inline_enable`, `inline_disable`, and `inline_reset` Directives

The `inline_enable` pragma directs 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 `inline_reset` directive returns the inlining state to the state specified on the command line (`-h ipan`). It has the following format:

```
#pragma _CRI inline_reset
```

The following example illustrates the use of these directives.

Example 11. Using the inline_enable, inline_disable, and inline_reset directives

The following code fragment shows how the inline_enable, inline_disable, and inline_reset directives would affect code compiled with the -h ipa4 option:

```
void qux(int x)
{
void bar(void);
int a = 1;
    x = a+a+a+a+a+a+a+a+a+a;
    bar();
}

void foo(void)
{
int j = 1;

/* enable inlining at all call sites here forward */
#pragma _CRI inline_enable
    qux(j);
    qux(j);
/* disable inlining at all call sites here forward */
#pragma _CRI inline_disable
    qux(j);
/* reset control to the command line -hipa4 */
#pragma _CRI inline_reset
    qux(j);
}
```

Example 12. Using inline_reset

The following code fragment shows how the #pragma _CRI inline_reset directive would affect code compiled with the -h ipa3 option:

```
void f1()
{
#pragma _CRI inline_disable /* No inlining will be done in f1;
    ...
}

void f2()
{
/* turn off all inlining to the end of the routine or another directive is encountered. */
#pragma _CRI inline_disable
    ...

/* The inlining state is -h ipa3 for the remainder of f2 */
#pragma _CRI inline_reset
    ...
}
```

3.8.2 `inline_always` and `inline_never` Directives

The `inline_always` directive specifies functions that the compiler should always attempt to inline. If the directive is placed in the definition of the function, inlining is attempted at every call site to *name* in the entire input file being compiled. If the directive is placed in a function other than the definition, inlining is attempted at every call site to *name* within the specific function containing the directive.

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. If the directive is placed in the definition of the function, inlining is never attempted at any call site to *name* in the entire input file being compiled. If the directive is placed in a function other than the definition, inlining is never attempted at any call site to *name* within the specific function containing the directive.

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.

When applied to a function template specialization, the inlining directive will apply to all instantiations of the template, not just the specialization. For example, in the following case, both `foo<double>` and `foo<float>` will be affected by the `inline_always` directive:

```
template<typename T>
int foo(T i, int j) {
    return j;
}

template <>
int foo(double f, int i) {
#pragma _CRI inline_always foo
    return i + 5;
}

int bar(float f,int i) {
    return foo(f,i);
}

int main(void) {
    foo(1.0,7);
    bar(1.0,7);
    return 0;
}
```

3.8.3 `clone_enable`, `clone_disable`, `clone_reset` Directives

The `clone_enable` and `clone_disable` directives control whether cloning is attempted over a range of code.

If `clone_enable` is in effect, cloning is attempted at call sites. If `clone_disable` is in effect, cloning is not attempted at call sites.

The `clone_reset` directive resets cloning to the state specified on the compiler command line.

These directives have the following formats:

```
#pragma _CRI clone_enable
#pragma _CRI clone_disable
#pragma _CRI clone_reset
```

One of these directives remains in effect until the opposite directive is encountered, until the end of the program unit, or until the `clone_reset` directive is encountered.

3.8.4 `clone_always` and `clone_never` Directives

The `clone_always` and `clone_never` directives specify functions or procedures that the compiler should always/never attempt to clone. If the directive is placed in the definition of the function, cloning is always/never attempted at every call site to *name* in the entire input file being compiled. If the directive is placed in a function other than the definition, cloning is always/never attempted at every call to *name* within the specific function containing the directive.

The `clone_always`, `clone_never` directives control cloning of a procedure for the compilation of the whole input file.

Use compile options `-hipaN` or `-OipaN` where *N* is equal to 5 to enable cloning directives.

These directives have the following formats:

```
#pragma _CRI clone_always name [,name ...]
#pragma _CRI clone_never name [,name ...]

!DIR$ CLONEALWAYS name [,name ...]
!DIR$ CLONENEVER name [,name ...]
```

The *name* argument is the name of a function.

3.9 PGAS Directive

3.9.1 `defer_sync` Directive

The `defer_sync` directive defers the synchronization of PGAS data until the next synchronization.

Normally the compiler synchronizes the references in a statement as late as possible without violating program semantics. The purpose of the `defer_sync` directive is to synchronize the references even later, beyond where the compiler can determine it is safe. PGAS data references made by the single statement immediately following the `pgas defer_sync` directive are not synchronized until the next fence instruction.

Use this directive to force all references in the next statement to be non-blocking. This helps for cases where the compiler cannot prove that it is safe. For example, if there is a remote-memory access (RMA) `put` near the end of a subroutine, the compiler must guard against the `put` value being read back immediately after the subroutine returns, so the `put` is synchronized just before returning. The programmer, however, may know that the value is not read back and can insert a `pgas defer_sync` directive.

The format is as follows:

```
#pragma pgas defer_sync
```

Example 13. Using `defer_sync`

```
void my_put( shared int* x, int thread, int value ) {  
    #pragma pgas defer_sync  
        x[thread] = value;  
}
```


Using the OpenMP C/C++ API [4]

OpenMP is a parallel programming model that is portable across shared memory architectures from Cray and other vendors.

The OpenMP C/C++ specification is accessible at <http://openmp.org/wp/openmp-specifications/>.

4.1 Standard Support

This compiler supports the OpenMP API, Version 3.1 (*OpenMP Application Program Interface Version 3.1 July 2011 Copyright © 1997-2011 OpenMP Architecture Review Board*), with a few exceptions. See [Limitations on page 113](#).

4.2 Cray Enhancements

This section describes features not described by the API specification which are included in Cray's implementation of OpenMP.

4.2.1 `cray_omp_set_wait_policy()`

This routine allows dynamic modification of the `wait-policy-var` ICV value, which corresponds to the `OMP_WAIT_POLICY` environment variable. The `policy` argument provides a hint to the OpenMP runtime library environment about the desired behavior of waiting threads; acceptable values are `ACTIVE` or `PASSIVE` (case insensitive). It is an error to call this routine in an active parallel region.

The OpenMP runtime library supports a "wait policy" and a "contention policy", both of which can be set with the following environment variables:

```
OMP_WAIT_POLICY=(ACTIVE|PASSIVE)
CRAY_OMP_CONTENTION_POLICY=(Automatic|Standard|MonitorMwait|IntegerDivision)
```

These environment variables allow the policies to be set once at program launch for the entire execution. However, in some circumstances it would be useful for the programmer to explicitly change the policy at various points during a program's execution. This cray-specific routine allows the programmer to dynamically change the wait policy (and potentially the contention policy). This addresses the situation when an application needs OpenMP for the first part of program execution, but there is a clear point after which OpenMP is no longer used. Unfortunately, the idle OpenMP threads still consume resources since they are waiting for more work, resulting in performance degradation for the remainder of the application. A passive-waiting policy might eliminate the performance degradation after OpenMP is no longer needed, but the developer may still want an active-waiting policy for the openmp-intensive region of the application. This routine notifies all threads of the policy change at the same time, regardless of whether they are idle or active (to avoid deadlock from waiting and signaling threads using different policies).

4.2.2 CRAY_OMP_CHECK_AFFINITY Environment Variable

Setting this environment variable to `TRUE` at execution time causes the CCE OpenMP runtime environment to display affinity binding for each OpenMP thread. The messages contain hostname, process identifier, OS thread identifier, OpenMP thread identifier, and affinity binding.

4.3 Compiling

By default, the CCE compiler recognizes OpenMP directives. These Cray C and C++ Compiler options affect OpenMP applications:

- `-h [no]omp` ([-h \[no\]omp on page 66](#))
- `-h threadn` ([-h threadn on page 41](#))

4.4 Executing

For OpenMP applications, use both the `OMP_NUM_THREADS` environment variable to specify the number of threads and the `aprun -d depth` option to specify the number of CPUs hosting the threads. The number of threads specified by `OMP_NUM_THREADS` should not exceed the number of cores in the CPU.

If neither the `OMP_NUM_THREADS` environment variable nor the `omp_set_num_threads()` call is used to set the number of OpenMP threads, the system defaults to 1 thread.

For further information, including example OpenMP programs, see the *Cray Application Developer's Environment User's Guide*.

4.5 Debugging

The `-g` option provides debugging support for OpenMP directives. The `-g` option provides debugging support identical to specifying the `-G0` option. This level of debugging implies `-homp` (most optimizations disabled but OpenMP directives are recognized) and `-h fp0`.

If you want to debug without OpenMP, use `-g -xomp` or `-g -hnoomp`, which will disable OpenMP and turn on debugging.

4.6 Cray Implementation Defined Behaviors

The OpenMP Application Program Interface Specification, Appendix E presents a list of implementation defined behaviors. This section presents a corresponding list of Cray specific behavior:

- Task scheduling points in untied task regions

There are no untied tasks in Cray's implementation of OpenMP.

- Atomicity of memory access by multiple threads

When multiple threads access the same shared memory location and at least one thread is a write, threads should be ordered by explicit synchronization to avoid data race conditions and the potential for non-deterministic results. Always use explicit synchronization for any access smaller than one byte.

- Initial values of internal control variables

<i>nthreads-var</i>	1
<i>dyn-var</i>	TRUE
<i>run-sched-var</i>	static, 0
<i>stacksize-var</i>	128 MB
<i>wait-policy-var</i>	ACTIVE
<i>thread-limit-var</i>	64
<i>max-active-levels-var</i>	1
<i>def-sched-var</i>	static, 0

- Dynamic adjustment of threads

dyn-var is supported and enabled by default. It behaves according to Algorithm 2-1 (See Section 2.4.1 of the OpenMP specifications). Threads may be dynamically created up to an upper limit which is 4 times the number of cores/node. It is up to the programmer to try to limit oversubscription.

- `loop` directive

The integer type or kind used to compute the iteration count of a collapsed loop are signed 64-bit integers, regardless of how the original induction variables and loop bounds are defined.

If the schedule specified by the runtime `schedule` clause is specified and *run-sched-var* is `auto`, then the Cray implementation generates a static schedule.

- `sections` directive

Multiple structured blocks within a single `sections` construct are scheduled in lexical order and an individual block is assigned to the first thread that reaches it. It is possible for a different thread to execute each section block, or for a single thread to execute multiple section blocks. There is not a guaranteed order of execution of structured blocks within a section.

- `single` directive

A single block is assigned to the first thread in the team to reach the block; this thread may or may not be the master thread.

- `atomic` directive

Different optimization levels produce different behavior. At thread-level 1 and lower (`-hthread1`), the atomic construct is transformed into a critical section with a shared lock between all atomic regions. This approach enforces mutual exclusion with respect to all atomic constructs. At thread-level 2 and above (`-hthread2`), the compiler will use hardware atomic instructions when available, which allows atomic constructs that access independent memory locations to execute in parallel. Data types larger than 64-bits are not supported by hardware atomic instructions; in these cases an atomic construct is transformed into a critical section with a shared lock between all atomic regions, regardless of the optimization level.

Note: For all data types that are 64-bits or less, we use hardware atomic instructions. In Fortran we can have 128-bit floats and 64-bit complex (128-bits of total data), and for those types we fall back to a critical section.

- `omp_set_num_threads` routine

Sets *nthreads-var* to a positive integer. If the argument is < 1 , then set *nthreads-var* to 1.

- `omp_set_schedule` routine

Sets the schedule type as defined by the current specification. There are no implementation defined schedule types.

- `omp_set_max_active_levels` routine

Sets the *max-active-levels-var* ICV. Defaults to 1. If argument is < 1 , then set to 1.

- `omp_get_max_active_levels` routine

There is a single *max-active-levels-var* ICV for the entire runtime system. Thus, a call to `omp_get_max_active_levels` will bind to all threads, regardless of which thread calls it.

- `OMP_SCHEDULE` environment variable

If the value of the variable does not conform to the specified format then print a warning and default to `static` with chunk 1 (each thread is a single iteration).

- `OMP_NUM_THREADS` environment variable

`aprun -d depth` has no effect on `OMP_NUM_THREADS`, but the programmer should usually set them to be equal. `OMP_NUM_THREADS` determines the number of threads that the program should create, while `aprun -d depth` determines the number of hardware cores that are available for the program to run on. Setting `OMP_NUM_THREADS` lower than *depth* will likely result in under-utilization of the allocated cores, while setting `OMP_NUM_THREADS` higher than *depth* will likely result in over-utilization of the allocated cores.

- `OMP_DYNAMIC` environment variable

If the value is neither `true` nor `false`, print a warning and set to default (`true`).

- `OMP_NESTED` environment variable

If the value is neither `true` nor `false`, print a warning and set to default (`true`).

- `OMP_STACKSIZE`

- `OMP_WAIT_POLICY` environment variable

`active` causes idle threads to use a spin-wait loop, consuming compute resources. `passive` causes idle threads to enter a blocked wait state, which does not consume compute resources but has a higher wake-up latency.

- `OMP_MAX_ACTIVE_LEVELS`

If the value is less than 1, it is set to 1. There is no explicit upper limit on the maximum number of active levels of parallelism; rather, the limit is determined by the maximum number of threads and the available memory. Additional levels of parallelism may be created after the thread limit is reached, but they will run with only one thread.

- `OMP_THREAD_LIMIT`

Max number of threads is 1024. No warning if try to exceed that. If < 1 , $4 * (\text{cores/node})$ (default)

Fortran only implementation specific behavior:

- `threadprivate` directive

If the conditions for values of data in the `threadprivate` objects of threads (other than the initial thread) to persist between two consecutive active parallel regions do not all hold, the allocation status of an allocatable array in the second region is implementation defined (see Section 2.9.2 on page 81). If a number of threads remain which is less than the allocated number, and if the number of threads increase, then new threads will allocate data.

- `shared` clause

Passing a shared variable to a non-intrinsic procedure may result in the value of the shared variable being copied into temporary storage before the procedure reference, and back out of the temporary storage into the actual argument storage after the procedure reference. Situations where this occurs other than those specified are implementation defined (see Section 2.9.3.2 on page 88). Copies are used only where required by the Fortran spec. Put interface block in place.

- Runtime library definitions

It is implementation defined whether the include file `omp_lib.h` or the module `omp_lib` (or both) is provided. It is implementation defined whether any of the OpenMP runtime library routines that take an argument are extended with a generic interface so arguments of different `KIND` type can be Fortran accommodated (see Section 3.1 on page 108). Both `omp_lib.h` and the module `omp_lib` are provided. Cray Fortran uses generic interfaces for routines.

4.6.1 Directives and Clauses

4.6.1.1 `atomic` Directive

The `atomic` directive is replaced with a `critical` section that encloses the statement.

4.6.1.2 `for` Directive

For the `schedule(guided, chunk)` clause, the size of the initial chunk for the master thread and other team members is approximately equal to the trip count divided by the number of threads.

For the `schedule(runtime)` clause, the schedule type and chunk size can be chosen at run time by setting the `OMP_SCHEDULE` environment variable. If this environment variable is not set, the schedule type and chunk size default to `static` and 0, respectively.

In the absence of the `schedule` clause, the default `schedule` is `static` and the default chunk size is approximately the number of iterations divided by the number of threads.

4.6.1.3 `parallel` Directive

If a parallel region is encountered while dynamic adjustment of the number of threads is disabled, and the number of threads specified for the parallel region exceeds the number that the runtime system can supply, the program terminates.

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 `OMP_NESTED` environment variable and the `omp_set_nested()` call control nested parallelism. To enable nesting, set `OMP_NESTED` to `true` or use the `omp_set_nested()` call. Nesting is disabled by default.

4.6.1.4 `threadprivate` Directive

The `threadprivate` directive specifies that variables are replicated, with each thread having its own copy. If the dynamic threads mechanism is enabled, the definition and association status of a thread's copy of the variable is undefined, and the allocation status of an allocatable array is undefined.

4.6.1.5 `private` Clause

If a variable is declared as `private`, the variable is referenced in the definition of a statement function, and the statement function is used within the lexical extent of the directive construct, then the statement function references the `private` version of the variable.

4.6.2 Library Routines

4.6.2.1 `omp_get_max_active_levels()`

The `omp_get_max_active_levels()` routine returns the maximum number of nested parallel levels currently allowed.

4.6.2.2 `omp_set_dynamic()`

The `omp_set_dynamic()` routine enables or disables dynamic adjustment of the number of threads available for the execution of subsequent parallel regions by setting the value of the *dyn-var* ICV. The default is on.

4.6.2.3 `omp_set_schedule()`

The `omp_set_schedule()` routine affects the schedule that is applied when `runtime` is used as schedule kind, by setting the value of the *run-sched-var* ICV. The default is `on`.

4.6.2.4 `omp_set_max_active_levels()`

The `omp_set_max_active_levels()` routine limits the depth of nested parallelism. The number specified controls the maximum number of nested parallel levels with more than one thread. The default value is 1 (nesting disabled).

4.6.2.5 `omp_set_nested()`

The `omp_set_nested()` routine enables or disables nested parallelism, by setting the *nest-var* internal control variable (ICV). The default is `false`.

4.6.2.6 `omp_set_num_threads()`

If dynamic adjustment of the number of threads is disabled, the `number_of_threads_expr` argument sets the number of threads for all subsequent parallel regions until this procedure is called again with a different value.

4.6.3 OpenMP Environment Variables

4.6.3.1 `OMP_DYNAMIC`

The default value is `true`.

4.6.3.2 `OMP_MAX_ACTIVE_LEVELS`

The default value is 1.

4.6.3.3 `OMP_NESTED`

The default value is `false`.

4.6.3.4 `OMP_NUM_THREADS`

If this environment variable is not set and you do not use the `omp_set_num_threads()` routine to set the number of OpenMP threads, the default is 1 thread.

The maximum number of threads per compute node is 4 times the number of allocated processors. If the requested value of `OMP_NUM_THREADS` is more than the number of threads an implementation 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. If `OMP_DYNAMIC` is `true`, it uses up to 4 times the number of allocated processors. For example, on a 8-core Cray XE system, this means the program can use up to 32 threads per compute node.

4.6.3.5 `OMP_SCHEDULE`

The default values for this environment variable are `static` for *type* and 0 for *chunk*.

4.6.3.6 `OMP_STACKSIZE`

The default value for this environment variable is 128 MB.

4.6.3.7 `OMP_THREAD_LIMIT`

Sets the number of OpenMP threads to use for the entire OpenMP program by setting the *thread-limit-var* ICV. The Cray implementation defaults to 4 times the number of available processors.

4.6.3.8 `OMP_WAIT_POLICY`

Provides a hint to an OpenMP implementation about the desired behavior of waiting threads by setting the *wait-policy-var* ICV. A compliant OpenMP implementation may or may not abide by the setting of the environment variable. The default value for this environment variable is `active`.

4.7 Limitations

The following OpenMP features are not currently supported by the Cray C and Cray C++ compilers.

- Orphaned task constructs may have an implicit `taskwait` directive added to the end of the routine. This is not required by the specification but is currently required by the Cray implementation. This limits the amount of parallelism that may be seen.
- Task switching is not implemented. The thread that starts executing a task is the thread that finishes the task.
- Support for OpenMP Random Access Iterators (RAIs) in the C++ Standard Template Library (STL) is deferred.

4.8 OpenMP Accelerator Support

The OpenMP accelerator directives are deprecated. Please use the OpenACC accelerator directives instead.

Using OpenACC [5]

OpenACC is a parallel programming model which facilitates the use of an accelerator device attached to a host CPU. The OpenACC API allows the programmer to supplement information available to the compilers in order to offload code from a host CPU to an attached accelerator device.

This release supports the *OpenACC Application Programming Interface, Version 2.0* standard developed by PGI, Cray Inc., NVIDIA, with support from CAPS enterprise.

Refer to the OpenACC home page at <http://www.openacc-standard.org>. Under the Downloads link, select the *OpenACC 2.0 Specification*.

For the most current information regarding the Cray implementation of OpenACC, see the `intro_openacc(7)` man page.

5.1 OpenACC Execution Model

The CPU host offloads compute intensive regions to the accelerator device. The accelerator executes parallel regions, which contain work sharing loops executed as *kernels* on the accelerator. The CPU host manages execution on the accelerator by allocating memory on the accelerator, initiating data transfer, sending code, passing arguments to the region, waiting for completion, transferring accelerator results back to the CPU host and releasing memory.

The accelerator supports multiple levels of parallelism. The accelerator executes a *kernel* composed of parallel threads or *vectors*. Vectors (threads) are grouped into sets called *workers*. Threads in a set of workers are scheduled together and execute together. *Workers* are grouped into larger sets called *gangs*. One or more *gangs* may comprise a kernel. To summarize, a *kernel* is executed as a set of *gangs* of *workers* of *vectors*.

The compiler determines the number of gangs/workers/vectors based on the problem and then maps the vectors, workers, and gangs onto the accelerator architecture. Specifying the number of gangs, workers, or vectors is optional but may permit tuning to a particular target architecture. The way that the compiler maps a particular problem onto a constellation of gangs, workers, and vectors which are then mapped onto the accelerator architecture is implementation defined. See [Mapping the OpenACC Programming Model onto Accelerator Components on page 116](#) for more information on this topic.

Note: OpenACC terminology is situated in the context of the PGAS programming model. In the PGAS model, there may be one or more Processing Elements (PEs) per node. Each PE is multi-threaded and each thread can execute vector instructions. The PGAS thread concept is not the same as the OpenACC thread concept.

5.2 OpenACC Memory Model

The memory on the accelerator is separate from host memory. Accelerator device memory is not mapped onto the host's virtual memory space. All data movement between host and accelerator memory is initiated by the host through the library functions that move data. Also, it is not assumed that the accelerator can access host memory, though it is supported by some devices. In this model, data movement between memories is managed by the compiler according to OpenACC directives. The programmer needs to be aware of device memory size, as well as memory bandwidth between host and device in order to effectively accelerate a region of code.

Current accelerators implement a weak memory model; they do not support memory coherence between operations executed by different *execution units* — an execution unit is a hardware abstraction which can execute one or more gangs. If an operation updates a memory location and another reads from the same location, or two operations store a value to the same location, the hardware may not guarantee repeatable results. Some potential errors of this type are prevented by the compiler, but it is possible to write an accelerator parallel region that produces inconsistent results. Memory coherence is guaranteed when memory operations referencing the same location are separated by an explicit barrier.

5.3 Mapping the OpenACC Programming Model onto Accelerator Components

The compiler maps the OpenACC execution model (kernels, gangs, workers, vectors) onto the accelerator architecture as described in the following sections.

5.3.1 Streaming Multiprocessors (SM) and Scalar Processor (SP) cores

The accelerator architecture is comprised of two main components — global memory and some number of streaming multiprocessors (SM). Each SM contains multiple scalar processor (SP) cores, schedulers, special-function units, and memory which is shared among all the SP cores. An SP core contains floating point, integer, logic, branching, and move and compare units. Each thread/vector is executed by a core. The SM manages thread execution.

The OpenACC execution model maps to the NVIDIA GPU hardware as follows (GPU terms are in parenthesis): One or more OpenACC kernels may execute on a GPU. The compiler divides a kernel into one or more gangs (blocks) of vectors (*threads*). Several concurrent gangs (*blocks*) of threads may execute on one SM depending on several factors, including memory requirements, compiler optimizations, or user directives. A single block (gang) does not span SMs and will remain on one SM until completion. When the SM encounters a block (gang), each gang (block) is further broken up into workers (*warps*) which are groups of threads to execute in parallel. Scheduling occurs at the granularity of the worker (warp). Individual threads within a warp start together and execute one common instruction at a time. If conditional branching occurs within a worker (warp), the warp serially executes each branch path taken causing some threads to wait until threads converge back to the same instruction. Data dependent conditional code within a warp usually has negative performance impact. Worker (warp) threads also fetch data from memory together and when accessing global memory, the accesses of the threads within a warp are grouped to minimize transactions. Each thread in a worker (warp) is executed on a different SP core.

There may be up to 32 threads in a worker (warp) — a limit defined by the hardware.

See the `intro_openacc(7)` man page for more detail on Partition Mapping.

5.3.2 Memory

There is a hierarchy of memory spaces used by OpenACC threads. Each thread has its own private local memory. Each gang of workers of threads has shared memory visible to all threads of the gang. All OpenACC threads running on a GPU have access to the same global memory. Global memory on the accelerator is accessible to the host CPU.

5.4 Mixed Model Support

OpenMP directives may appear inside of OpenACC data or host data regions only. OpenMP directives are not allowed inside of any other OpenACC directives.

For example, the following is permitted:

```
#pragma acc data
{
#pragma omp parallel
    {a = 10}
}
```

OpenACC may not appear inside OpenMP directives. If you wish to have OpenACC directives nested inside of OpenMP constructs, place them in calls that are not inlined.

5.5 Compiling

The CCE compiler recognizes OpenACC directives, by default. Use either the `ftn` or `cc` command to compile.

The CCE compiler does not produce CUDA code. It generates PTX (Parallel Thread Execution) instructions which are then translated into assembly.

Note the following interactions between directives and command line options.

- `-h [no]acc`
`-h noacc` disables OpenACC directives.
- `-h [no]pragma`
See `-h [no]pragma=name[:name ...]` on page 59.
- `-h acc_model=option[:option] ...`

Explicitly controls the execution and memory model utilized by the accelerator support system. The option arguments identify the type of behavior desired. There are three option sets. Only one member of a set may be used at a time; however, all three sets may be used together. Default: `auto_async_kernel:fast_addr:no_deep_copy`

option Set 1:

`auto_async_none`

Execute kernels and updates synchronously, unless there is an `async` clause present on the kernels or update directive.

`auto_async_kernel`

Execute all kernels asynchronously ensuring program order is maintained.

`auto_async_all`

Execute all kernels and data transfers asynchronously, ensuring program order is maintained.

option Set 2:`no_fast_addr`

Use default types for addressing.

`fast_addr` (Default) Attempt to use 32 bit integers in all addressing to improve performance. Base addresses remain as 64 bit. The performance is improved by potentially using fewer registers and faster arithmetic for offset calculations. This optimization may result in incorrect behavior for codes that make use within accelerator regions of any of the following: very large arrays (offsets would require greater than 32 bits); very large array lower bounds (max offset plus lower bound is greater than 32 bits); bitfields/other bit operations.

option Set 3:`no_deep_copy`

Do not look inside of an object type to transfer sub-objects. Allocatable members of derived type objects will not be allocated on the device.

`deep_copy` (Fortran only) Look inside of derived type objects and recreate the derived type on the accelerator recursively. A derived type object that contains an allocatable member will have memory allocated on the device for the member.

5.6 Module Support

To compile, ensure that `PrgEnv-cray` module is loaded and that it includes CCE 8.2 or later. Then, either load the `craype-accel-nvidia20` module for Fermi support or the `craype-accel-nvidia35` module for Kepler support.

The `craype-accel-host` module supports compiling and running an OpenACC application on the host X86 processor. This provides source code portability between systems with and without an accelerator. The accelerator directives are automatically converted at compile time to OpenMP equivalent directives.

Use either the `ftn` or `cc` command to compile.

5.7 Debugging

Use either Alinea DDT or Rogue Wave TotalView.

The following apply to all debuggers:

- To enable debugging, compile using the `-g` option.

When compiling with the debug option (`-g`), CCE may require additional memory from the accelerator heap, exceeding the 8MB default. In this case, there will be `malloc` failures during compilation. The environment variable `CRAY_ACC_MALLOC_HEAPSIZE` specifies the accelerator heap size in bytes. It may be necessary to increase the accelerator heap size to 32MB (33554432), 64MB (67108864), or greater by setting `CRAY_ACC_MALLOC_HEAPSIZE` accordingly. The accelerator heap size defaults to 8MB.

- Debug one rank/image/thread/PE per node.
- CCE does not generate CUDA code, but generates PTX code. Debuggers will not display CUDA intermediate code.
- To enter an OpenACC region using a debugger, breakpoints may be set inside the OpenACC region. It is not possible to do a single `step` into the region from the code immediately prior to the start of an OpenACC directive.

5.8 OpenACC Directives

For information on the OpenACC directives, see the *OpenACC 2.0 Specification* available at <http://www.openacc-standard.org>.

For the most current information regarding the Cray implementation of OpenACC, see the `intro_openacc(7)` man page. See the `OpenACC.EXAMPLES(7)` man page for example OpenACC codes.

5.9 Runtime Routines

Runtime routines defined by the standard specification are supported unless otherwise noted in the `intro_openacc(7)` man page.

5.9.1 Cray Specific Runtime Library Routines

The following routines are currently Cray specific. These interfaces are subject to change and their usage may result in non-portable code.

- `void cray_acc_memcpy_to_host_async(void* host_destination, const void* device_source, size_t size, int async_id);`

Asynchronously copies *size* bytes from the accelerator source address to the host destination address; returns destination address. See `async` clause for explanation of *async_id*.

- `void cray_acc_memcpy_to_device_async(void* host_destination, const void* device_source, size_t size, int async_id);`

Asynchronously copies *size* bytes from the accelerator source address to the host destination address; returns destination address. See `async` clause for explanation of *async_id*.

- `bool cray_acc_get_async_info(int async_id, void* async_info);`

Returns true if the *async_id* was found to have any architecture specific `async` information available. The user is responsible for ensuring that the *async_info* pointer points to a `async` structure from the underlying architecture. For an NVIDIA target this would be a CUDA Stream (CUstream).

5.9.2 CRAY_ACC_DEBUG Output Routines

When the runtime environment variable `CRAY_ACC_DEBUG` is set to 1, 2, or 3, CCE writes runtime commentary of accelerator activity to `STDERR` for debugging purposes; every accelerator action on every PE generates output prefixed with "ACC:". This may produce a large volume of output and it may be difficult to associate messages with certain routines and/or certain PEs.

With this set of API calls, the programmer can enable or disable output at certain points in the code, and modify the string that is used as the debug message prefix.

The `cray_acc_set_debug*_prefix(void)` routines define a string that is used as the prefix, with the default being "ACC:". The `cray_acc_get_debug*_prefix(void)` routines are provided so that the previous setting can be restored.

Output from the library is printed with a format string starting with "ACC: %s %s", where the global prefix is printed for the first %s (if not NULL), and the thread prefix is printed for the second %s. The global prefix is shared by all host threads in the application, and the thread prefix is set per-thread. By default, strings used in the %s fields are empty.

The C interface is provided by `omp.h`:

- `char *cray_acc_get_debug_global_prefix(void)`
- `void cray_acc_set_debug_global_prefix(char *)`
- `char *cray_acc_get_debug_thread_prefix(void)`
- `void cray_acc_set_debug_thread_prefix(char *)`

To enable debug output, set level from 1 to 3, with 3 being the most verbose. Setting a level less than or equal to 0 disables the debug output. The get version is provided so the previous setting can be restored. The thread level is an optional override of the global level.

- `int cray_acc_get_debug_global_level(void)`
- `void cray_acc_set_debug_global_level(int level)`
- `int cray_acc_get_debug_thread_level(void)`
- `void cray_acc_set_debug_thread_level(int level)`

5.10 Environment Variables

5.10.1 Cray Specific

- `CRAY_ACC_MALLOC_HEAPSIZE`

Specifies the accelerator heap size in bytes. The accelerator heap size defaults to 8MB. When compiling with the debug option (`-g`), CCE may require additional memory from the accelerator heap, exceeding the 8MB default. In this case, there will be `malloc` failures during compilation. It may be necessary to increase the accelerator heap size to 32MB (33554432), 64MB (67108864), or greater.

- `CRAY_ACC_DEBUG`

When set to 1, 2, or 3 (most verbose), writes runtime commentary of accelerator activity to `STDERR` for debugging purposes. There is also an API which allows the programmer to enable/disable debug output and set the output message prefix from within the application. See `CRAY_ACC_DEBUG` Output Routines.

5.10.2 Standard

The following environment variables are defined by the API specification:

- `ACC_DEVICE_NUM`
- `ACC_DEVICE_TYPE`

5.11 OpenACC Examples

See the `OpenACC.EXAMPLES(7)` man page for examples.

Using Cray Unified Parallel C (UPC) [6]

Unified Parallel C (UPC) is a C language extension for parallel program development. UPC supports a Partitioned Global Address Space (PGAS) programming model.

This release supports the UPC Language Specification, Version 1.3. The proposed UPC 1.3 standard is discussed on the UPC specification website, <http://code.google.com/p/upc-specification>.

This chapter describes the Cray specific UPC functionality available in CCE, and features of the specification which are implementation defined. Also see `intro_pgas(7)`, or refer to the appropriate UPC man page.

You should be 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 1.2* paper.

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.

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 and SHMEM.

Note: UPC 1.3 supports a parallel I/O model which provides control over file synchronization. However, if you continue to use the regular C I/O routines, 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.

Cray UPC supports GASP instrumentation. GASP instrumentation enables the use of external performance tools, such as the Parallel Performance Wizard (PPW) from the University of Florida. For more information on GASP and PPW, see <http://gasp.hcs.ufl.edu><http://ppw.hcs.ufl.edu>. To instrument for GASP, refer to the command line option `-h gasp [=opt [:opt]]` on page 53.

For a description of UPC related command line options, see [Compiling and Linking UPC Code](#) on page 127.

6.1 Implementation

6.1.1 Predefined Macros

The following UPC 1.3 preprocessor macros are supported and defined as follows:

- `__UPC__`: 1
- `__UPC_VERSION__`: 201309 (corresponds to the date that 1.3 spec is published)
- `UPC_MAX_BLOCK_SIZE`: 1073741823
- `__UPC_DYNAMIC_THREADS__`: 1 (if compiling for dynamic threads, otherwise undefined)
- `__UPC_STATIC_THREADS__`: 1 (if compiling for static threads, otherwise undefined)
- `__UPC_COLLECTIVE__`: 1
- `__UPC_TICK__`: 1
- `__UPC_CASTABLE__`: 1
- `__UPC_IO__`: 1
- `__UPC_NB__`: 1

6.1.2 False Sharing

There is a false sharing hazard when referencing shared `char` and `short` integers.

If two PEs store a `char` or `short` to the same 64-bit word in memory without synchronization, incorrect results can occur. It is possible for one PE's store to be lost. This is because these stores are implemented by reading the entire 64-bit word, inserting the `char` or `short` value and writing the entire word back to memory.

The following output is a result of 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. UPC programmers must take care when storing to shared `char` and `short` data that this situation does not occur.

6.2 Compiling and Linking UPC Code

Compiling a PGAS application (UPC, Fortran 2008) requires the `PrgEnv-cray` module to be loaded.

The `-hupc` option is required to enable recognition of UPC syntax because it is not part of the standard C language.

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. See [-X npes on page 68](#) for requirements regarding the use of the `-X npes` option.

The following command creates an executable file:

```
% cc -hupc hello.c -o hello
```

An executable can be created by linking together various object files that were generated from source code written in standard C, UPC, and Fortran. Either `cc` or `ftn` can be used to link the object files:

```
% cc -hupc x.o y.o z.o
% ftn x.o y.o z.o
```

For dynamic linking, add the `-dynamic` option. For information about linking PGAS applications to use huge pages, see the `intro_hugepages(1)` man page.

The Cray implementation of UPC supports adding GASP instrumentation to UPC codes. To instrument for GASP, refer to the command line option `-h gasp [=opt [:opt]]` on page 53.

6.3 Launching a UPC Application

After compiling the UPC code, you run the program using the `aprun` command.

Launch the application using 128 PEs:

```
% aprun -n 128 ./hello
```

If you use the `-X npes` compiler option, you must specify the same number of threads in the `aprun` command. The processing elements specified by `npes` are compute node cores/PEs.

By default, each PE reserves 64 MB of symmetric heap space. To increase or decrease this amount, set the `XT_SYMMETRIC_HEAP_SIZE` environment variable to the desired number of bytes. The suffixes K, M, and G are permitted to simplify requests for large values:

```
% export XT_SYMMETRIC_HEAP_SIZE=512M  
% aprun -n 128 ./hello
```

The UPC run time system uses GNI and DMAPP (low level libraries) to implement a logically shared, distributed memory programming model. The symmetric heap is mapped onto hugepages by DMAPP. It is advisable to also map the static data and/or private heap onto huge pages. See the `intro_hugepages(1)` man page.

6.4 Cray Extensions

Cray extensions to UPC that are not part of the UPC Language Specification 1.3 are listed here.

Note: A number of former extensions to UPC 1.2 have been standardized in UPC 1.3, including non-blocking bulk copies (`upc_nb.h`), privatizability (`upc_castable.h`) and timing (`upc_tick.h`) interfaces. These interfaces have been removed from the `upc_cray.h` header and moved into new headers as required by the UPC 1.3 specification. Additionally, some of the semantics and interfaces have been changed slightly, so existing users of these interfaces may need to update their applications.

6.4.1 Team Collectives

The following interfaces, declared in `upc_collective_cray.h`, provide common collective operations on a subset (team) of threads. These are loosely based on the UPC Collectives Library 2.0 proposal, with changes to argument ordering to better match existing practice in UPC and no explicit initialization.

- `CRAY_UPC_TEAM_ALL`
- `CRAY_UPC_TEAM_NODE`
- `cray_upc_op_create(3c)`
- `cray_upc_op_free(3c)`
- `cray_upc_type_size(3c)`
- `cray_upc_team_rank(3c)`
- `cray_upc_team_size(3c)`
- `cray_upc_team_split(3c)`
- `cray_upc_team_free(3c)`
- `cray_upc_team_barrier(3c)`
- `cray_upc_team_allreduce(3c)`
- `cray_upc_team_reduce(3c)`

6.4.2 Node Affinity

Include `upc_cray.h` to use these extensions.

6.4.2.1 `upc_nodeof()`

Returns the index of the node of the thread that has affinity to the shared object pointed to by *ptr*. Similar to `upc_threadof()`.

6.4.2.2 `NODES`

`NODES` is an expression with a value of type `int`; it specifies the number of nodes and has the same value on every thread in the job.

Similar to `THREADS`, but evaluates to the number of nodes used by the application, equal to the ceiling of the `aprun -n` value divided by the `-N` value.

6.4.2.3 MYNODE

`MYNODE` is an expression with a value of type `int`; it specifies the unique node index associated with the current thread and has the same value on all threads that are located on the same node.

Similar to `MYTHREAD`, but evaluates to a node number in the range 0 to `NODES - 1`, inclusive.

Using Cray C++ Libraries [7]

Most of the standard C++ features are supported, except for the few mentioned in [Unsupported Standard C++ Library Features](#).

For information about C++ language conformance and exceptions, see [Appendix B, Using Cray C and C++ Dialects on page 195](#).

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.

Using Coarray C++ [8]

This release introduces *Coarray C++*, a template library that implements the coarray concept for Partitioned Global Address Space (PGAS) programming in C++. The template library specifications are contained on a set of *.html pages that the CCE installation copies to /opt/cray/cce/version/doc/html/ on the Cray platform; they may be copied to any location which provides HTML web content for your site, or any location that can be accessed by site local web browsers.

The coarray concept used in Coarray C++ is intentionally very similar to Fortran (ISO/IEC 1539-1:2010) coarrays. Users familiar with Fortran coarrays will notice that terminology and even function names are identical, although the syntax follows C++ conventions.

A coarray adds an additional dimension, called a *codimension*, to a normal scalar or array type. The codimension spans instances of a Single-Program Multiple-Data (SPMD) application, called *images*, such that each image contains a slice of the coarray equivalent in shape to the original scalar or array type. Each image has immediate access via processor loads and stores to its own slice of the coarray, which resides in that image's local partition of the global address space. By specifying an image number in the cosubscript of the codimension, each image also has access to the slices residing in other images' partitions.

Images are an orthogonal concept to threads, such as those provided by C++11 or OpenMP. Threads are used for shared memory programming where each thread has immediate access to the address space of a single process and possibly some thread-local storage to which only it has access. Images are a broader concept intended to provide communication among cooperating processes that each have their own address space. The mechanism for this cooperation varies by implementation. Typically it involves network communication between processes that have arranged to have identical virtual memory layouts. This communication is one-sided such that a programmer can have an image read or write data that belongs to a different image without writing any code for the second image. Note that images and threads may coexist in the same application; a large networked system with multicore nodes could use coarrays to communicate among nodes but use threads within each node to exploit the multicore parallelism.

In Coarray C++, a coarray is presented as a class template that collectively allocates an object of a specified type within the address space of each image. The coarray object is responsible for managing storage for the object that it allocates. When used in an expression context, the coarray object automatically converts to its managed object so that an image can access its own slice of the coarray without using special syntax. Accessing a slice that belongs to a different image requires specifying the image number as a cosubscript in parenthesis immediately following the coarray object, before any array subscripts. Therefore, the codimension is the slowest-running array dimension, just like Fortran.

Note: The subscript order is backwards from Fortran because in Fortran the slowest-running dimension is rightmost whereas in C++ it is leftmost.

In addition to providing the fundamental ability to allocate and access a coarray, Coarray C++ provides image synchronization, atomic operations, and collectives.

Although this chapter presents Cray's implementation, Coarray C++ is designed to allow portable applications to be written for a variety of computing platforms in the sense that the template library interface is platform independent and can be compiled by any C++03 (ISO/IEC 14882:2003) or C++11 (ISO/IEC 14882:2011) compliant compiler. The implementation of the template library is likely to differ for each platform due to different transport layers (e.g., shared memory or various networks) for communicating data between images.

8.1 Compiling Coarray C++

The following program is the Coarray C++ equivalent of the classic "Hello World" program. The header file `coarray_cpp.h` provides all Coarray C++ declarations within namespace `coarray_cpp`. Normally a program imports all of the declarations into its namespace with a `using` directive, but having the namespace gives the programmer flexibility to deal with name conflicts.

```
#include <iostream>
#include <coarray_cpp.h>
using namespace coarray_cpp;
int main( int argc, char* argv[] )
{
    std::cout << "Hello from image " << this_image()
              << " of " << num_images() << std::endl;
    return 0;
}
```

The program is compiled with the Cray compiler and executed using four images as follows:

```
> module load PrgEnv-cray
> CC -o hello hello.cpp
> aprun -n4 ./hello
Hello from image 0 of 4
Hello from image 1 of 4
Hello from image 2 of 4
Hello from image 3 of 4
```

8.2 Declaring and Accessing Coarrays

The general form of a coarray declaration is:

```
coarray<T> name;
```

where T is the type of the object that will be allocated in the address space of each image.

A coarray declaration may appear anywhere that a C++ object can be declared. Therefore, a coarray may be declared as a global variable, local variable, static local variable, or as part of a struct or class. It may be allocated statically or dynamically. The only restriction is that a coarray allocation must be executed collectively by all images. The C++ language ensures that this restriction is met for global and static local coarray declarations, but the programmer is responsible for ensuring that local and dynamically-allocated coarrays are declared collectively. For example:

```
coarray<int> x; // global
void
foo( void )
{
    static coarray<int> y; // static local
    coarray<int> z; // local
    coarray<int>* p = new coarray<int>; // dynamically allocated
    ...
    delete p;
} // z is automatically destroyed here
```

8.2.1 Basic Types

A coarray of a basic C++ type is the simplest kind of coarray. Each image has an instance of the basic type that is managed by its coarray object. A coarray of type int is declared as:

```
coarray<int> x;
```

The declaration may pass an initial value to the constructor. Different images may pass different initial values:

```
coarray<int> x(2);
```

The initializer syntax below is not supported. If it were permitted, then automatic conversion from `int` to `coarray<int>` would be allowed, which would loosen type checking and lead to unexpected collective allocations:

```
coarray<int> x = 2;
```

This `coarray` object will behave as if it were the `int` that it manages. Assigning to the `coarray` object will assign a value to the `int` that is managed by the `coarray` object:

```
x = 42;
```

Likewise, using the `coarray` object in any expression where an `int` is expected shall read the value of the managed `int`:

```
int y = x + 1;
```

If the `coarray` object needs to be used in an expression where no particular type is expected, then the managed object can be accessed explicitly via empty parenthesis:

```
// prints the address of the coarray object
std::cout << &x << std::endl;
// prints the address of the int managed by the coarray object
std::cout << &x() << std::endl;
```

Accessing an `int` that is managed by another image requires specifying the image number within the parenthesis:

```
x(5) = 42; // set x = 42 within the address space of image 5
int y = x(2); // obtain the value of x from the address space of image 2
```

Finally, consider an enhanced version of the Hello World program. In this program, all images write their image number to their local object and then call `sync_all()`, which synchronizes control flow across all images. After the `sync_all()`, each image computes the image number of its left and right neighbors in the image space and prints the values that were written by its neighbors.

```
#include <iostream>
#include <coarray_cpp.h>
using namespace coarray_cpp;
int main( int argc, char* argv[] )
{
    coarray<int> x;
    x = this_image();
    sync_all();
    const int left = ( this_image() - 1 ) % num_images();
    const int right = ( this_image() + 1 ) % num_images();
    std::cout << "Hello from image " << x << "
              where x(left) = " << x(left) << " and x(right) = "
              << x(right) << std::endl;
```



```

    return 0;}

> CC -o hello2 hello2.cpp
> aprun -n4 ./hello2
Hello from image 0 where x(left) = 3 and x(right) = 1
Hello from image 3 where x(left) = 2 and x(right) = 0
Hello from image 2 where x(left) = 1 and x(right) = 3
Hello from image 1 where x(left) = 0 and x(right) = 2

```

8.2.2 Arrays

A coarray of an array type gives every image an array of the same shape. An example of a statically-sized coarray is below. The complete array type, including all extents, is provided as the coarray template's type argument:

```
// Declares a coarray of an array of 10 arrays of 20 ints
coarray<int[10][20]> x;
```

Note: The following declaration is very different:

```
// Declares an array of 10 arrays of 20 coarrays
// of type int. Legal, but very inefficient!
coarray<int> bad[10][20];
```

A coarray of a multidimensional array type is not achieved via nested coarray types. Although such declarations are legal, they are strange and not particularly useful:

```
// Declares a coarray of an array of 10 coarrays of arrays of 20 ints
coarray< coarray<int[20]>[10] > weird;
```

In a dynamically-sized coarray declaration, the extent of the leading dimension is left unbounded. The size of this extent cannot be part of the template type because it is not known at compile time. Instead, the size is passed as a constructor argument:

```
coarray<int[][20]> y(n); // each image must pass the same value
```

Later, the extent of the leading dimension can be extracted from the coarray object via the `extent()` member function:

```
size_t y_extent = y.extent();
```

An individual element of the local array managed by the coarray object is accessed by applying subscripts directly to the coarray object. When accessing part of the coarray managed by another image, the cosubscript appears in parenthesis before the subscripts:

```
x[4][5]    = 1; // set x[4][5] = 1 within this image's address space
y(3)[6][7] = 2; // set y[6][7] = 2 within the address space of image 3
```

8.2.3 Pointers

A coarray of pointers is typically used to implement a "ragged array" where different images need to allocate a different amount of memory as part of the same coarray. An example of a coarray of pointers is:

```
coarray<int*> x;
```

Each image allocates additional memory independently from the collective allocation of the coarray object itself:

```
x = new int[n]; // n usually varies per image
```

Due to the independent allocations, the allocated memory might not be located at the same address within every image's address space. Therefore, accessing the data requires an additional read of the pointer from the target image before a normal read or write can occur. This additional read happens automatically as part of the usual syntax for accessing the data:

```
x(i)[3] = 4; // set x[3] = 4 within the address space of image i
```

Note: The address stored within the pointer may be valid only on the allocating image, unless the program is careful to target the pointer at only symmetric virtual addresses. Great care should be taken with the following code pattern:

```
int* p = x(i); // get an address from image i
p[3] = 4;      // and dereference it on this image
```

Finally, the program must ensure prior to performing any accesses that other images have allocated their memory:

```
coarray<int*> x;x = new int[n];
sync_all();
x(i)[3] = 4;
```

8.2.4 Structs, Unions, and Classes

A coarray of a struct, union, or class behaves like a coarray of a basic type when the entire object is accessed, however special syntax is required for member access due to limitations of C++ operator overloading:

```
struct Point { int x, y; };

coarray<Point> pt;
Point p;

pt = p; // set pt = p in this image's address space
pt(2) = p; // set pt = p within the address space of image

2pt->x = 0; // set pt.x = 0 in this image's address space
pt().x = 0; // alternate syntax

// set pt.x = 1 within address space of image i
pt(i).member( &Point::x ) = 1;
```

Calling a member function of an object that resides in the address space of another image (i.e., a remote procedure call) is not supported. By default, when a struct, union, or class is copied between images, it is treated as a Plain Old Data (POD) type such that a bitwise copy occurs. This behavior is not appropriate if the type contains pointers to allocated data. The default behavior can be changed by creating a specialization of `coarray_traits` where `is_trivially_gettable` is false. C++ requires that the specialization be placed in the same namespace as the general template:

```
struct my_string {
    char* data;
    size_t length;
};

namespace coarray_cpp {
    template < >
    struct coarray_traits<my_string> {
        static const bool is_trivially_gettable = false;
        static const bool is_trivially_puttable = false;
    };
}
```

When `is_trivially_gettable` is false for a type, Coarray C++ expects the type to have a special constructor and a special assignment operator to facilitate reading an object from a remote image:

```
struct my_string {
    char* data;
    size_t length;

    // remote constructor
    my_string( const_coref<my_string> ref );
    // remote assignment operator
    my_string& operator = ( const_coref<my_string> ref );
};
```

The role of the remote constructor or remote assignment operator is to read the POD parts of the object from the other image, use that data to calculate how much memory needs to be allocated, allocate the memory, then read the rest of the object into the newly allocated memory. See [Section 8.3.1](#) for an explanation of the parameters of these functions.

Typically, if `is_trivially_gettable` is false for a type, then `is_trivially_puttable` should also be false. When `is_trivially_puttable` is false for a type, a compile time error will occur the program attempts to copy an instance of the type to another image.

8.3 Type System

The Coarray C++ type system is modeled closely on the C++ type system. In addition to the coarray type that extends the C++ array concept across images, there are *coreferences* and *copointers* that extend the C++ concepts of references and pointers to refer to objects on other images.

8.3.1 Coreferences

A coreference is returned when a cosubscript is applied to a coarray. Like a C++ reference, a coreference is always associated with an object, called its referent, can never be rebound to a different object, and can never be null. Typically a coreference is either immediately converted to its referent or subscripted, such that it is not necessary to declare a coreference and its fleeting presence can be ignored. Nevertheless, explicit coreferences are useful in some situations. Suppose that a function needs to have access to an object in another image's address space, but does not need to know anything about the coarray containing the object. For example:

```
void foo( coref<int> );
int main( int argc, char* argv[] ){
    coarray<int> x;
    coarray<int[10]> y;
    ...
    foo( x(2) );
    foo( y(3)[4] );
    ...
return 0;
}
```

In the above code, function `foo` can access an `int` that is part of either `x` or `y` even though `x` and `y` have different shapes. If `foo` were to require a coarray parameter instead, then it could accept either `x` or `y` but not both because the coarrays have different types. Furthermore, `foo`'s coreference parameter makes it clear to someone reading the code that the function's effect is narrow, limited to one object instead of an entire coarray. Two other uses of coreferences are to operate on coarray slices that are larger than a single object and to move data in bulk between images. To make these techniques more useful, coreferences can be created for local objects:

```
int main( int argc, char* argv[] ){
    coarray<int[5][10]> x;
    int local[10];
    coref<int[10]> local_ref( local );
    ...
    // local[0...9] = x(2)[1][0...9]
    local_ref = x(2)[1];
    ...
    // x(3)[4][0...9] = local[0...9]
    x(3)[4] = local_ref;
    ...
return 0;
}
```

For convenience, the `make_coref` and `make_const_coref` functions create coreferences for local objects without requiring the programmer to write the type of the local object:

```
int main( int argc, char* argv[] )
{
    coarray<int[5][10]> x;
    int local[10];
    ...
    // local[0...9] = x(2)[1][0...9]
    make_coref( local ) = x(2)[1];
    ...
    // x(3)[4][0...9] = local[0...9]
    x(3)[4] = make_const_coref( local );
    ...
    return 0;
}
```

A `const_coref` behaves exactly like a `coref` except that it cannot be used to modify its referent.

8.3.2 Copointers

A coreference can be converted to a copointer by calling its address function; the address-of operator is not overloaded. Local pointers are automatically convertible to copointers. Unlike coreferences, a copointer can be reassociated and can be unassociated or null. Arithmetic on a copointer changes the address to which it points but never changes the image to which it points. Comparisons between two copointers are allowed provided that both copointers point to the same image. Copointers can be used as iterators with standard C++ function templates. For example, the following code will not assert:

```
int
main( int argc, char* argv[] )
{
    coarray<int[10]> x;
    const size_t left = ( this_image() - 1 ) % num_images();
    const size_t right = ( this_image() + 1 ) % num_images();
    coptr<int> begin = x(right)[0].address();
    // Apply a standard algorithm, using a coptr as an iterator.
    coptr<int> end = x(right)[10].address();
    std::fill( begin, end, image );
    sync_all();
    for ( int i = 0; i < 10; ++i ) {
        assert( x[i] == left );
    }
    return 0;
}
```

They can be used to form linked lists spanning images. The list even can include links that point to local data:

```
#include <iostream>
#include <coarray_cpp.h>
```

```
using namespace coarray_cpp;

template < typename T >
struct Link {
    T data;
    coptr< Link<T> > next;
};

coarray< Link<int> > global_links;

int main( int argc, char* argv[] )
{
    Link<int> local_link;
    global_links->data = 2 * this_image();
    global_links->next = &local_link;
    local_link.data = 2 * this_image() + 1;
    if ( this_image() < num_images() - 1 ) {
        local_link.next = global_links(this_image() + 1).address();
    }
    else {
        local_link.next = 0;
    }
    sync_all(); // ensure every image has setup the data

    if ( this_image() == 0 ) {
        for ( coptr< Link<int> > p = global_links(0).address();
              p != NULL; p = p->member( &Link<int>::next ) ) {
            std::cout << p->member( &Link<int>::data ) << std::endl;
        }
    }
    // ensure local_link is not destroyed before it's read by image 0
    sync_all();
    return 0;
}
```

Compiling and executing the above program:

```
> CC -o list list.cpp
```

```
> aprun -n4 ./list
```

A `const_coptr` behaves exactly like a `coptr` except that it cannot be used to modify its target.

8.3.3 `shape_cast`

Various different array types have the same number of elements even though they have a different shape. For example, `int[100]`, `int[2][50]`, and `int[2][2][25]` all have 100 elements. A reference or pointer to a coarray of one of these types can be reinterpreted as a coarray of any of the others via a `shape_cast`, which has the same syntax as the standard C++ `static_cast`, `dynamic_cast`, `reinterpret_cast`, and `const_cast`. A `shape_cast` converts between coarray types of the same ultimate type that have different shapes. For example, a `shape_cast` cannot be used to reinterpret a `coarray<int[100]>&` as a `coarray<float[100]>&`; that conversion will throw a `std::bad_cast` exception. A `shape_cast` can be used to convert to a smaller shape but not to a larger shape. For example, a `coarray<int[100]>&`

may be converted to a `coarray<int[50]>&`, in which case the new coarray can access only the first 50 elements of the original, but it may not be converted to a `coarray<int[200]>&` because that requires more storage and will throw a `std::bad_cast` exception. The example code below shows various legal `shape_casts`:

```
#include <cassert>
#include <iostream>
#include <coarray_cpp.h>

using namespace coarray_cpp;

void foo( const coarray<int[]>& y ) { }
void foo10( const coarray<int[10]>& y ) { }
void foo5( const coarray<int[][5]>& y ) { }
void foo10_5( const coarray<int[10][5]>& y ) { }
void foo50( const coarray<int[50]>& y ) { }
int
main( int argc, char* argv[] )
{
    int extent = 10;
    coarray<int[10]>      x_10_s;
    coarray<int[]>      x_10_d(extent);
    coarray<int[10][5]>  x_10_5_s;
    coarray<int[][5]>    x_10_5_d(extent);
    coarray<int>        y;

    // Perform all valid combinations of passing the coarrays to the functions,
    // using shape_cast when necessary.
    foo( x_10_s );
    foo( x_10_d );
    foo( shape_cast<int[]>( x_10_5_s ) );
    foo( shape_cast<int[]>( x_10_5_d ) );
    foo10( x_10_s );
    foo10( x_10_d );
    foo5( shape_cast<int[2][5]>( x_10_s ) );
    foo5( shape_cast<int[][5]>( x_10_d ) );

    foo5( x_10_5_s );
    foo5( x_10_5_d );
    foo10_5( x_10_5_s );
    foo10_5( x_10_5_d );
    foo50( shape_cast<int[50]>( x_10_5_s ) );
    foo50( shape_cast<int[50]>( x_10_5_d ) );

    // Trivial reshape to same shape.
    shape_cast<int>( y );

    // shape_cast from scalar to array.
    shape_cast<int[1]>( y );

    // shape_cast from array to scalar.
    shape_cast<int>( x_10_s );

    // shape_cast to smaller array.
    shape_cast<int[5]>( x_10_s );
```

```
// shape_cast to larger array.
bool passed = false;
try {
    shape_cast<int[25]>( x_10_s );
} catch ( std::bad_cast& e ){
    passed = true;
}
assert( passed );

return 0;
}
```

8.4 Control Flow and Synchronization

8.4.1 Writing SPMD Code

Coarray C++ follows the Single-Program Multiple-Data model where all images begin executing the same main program but may operate on different data. Conditional code is used to restrict execution to certain images:

```
#include <iostream>
#include <coarray_cpp.h>

using namespace coarray_cpp;

int main( int argc, char* argv[] )
{
    if ( this_image() % 2 == 0 ){
        std::cout << "Hello from even image "
                  << this_image() << std::endl;
    }
    else {
        std::cout << "Hello from odd image "
                  << this_image() << std::endl;
    }
    return 0;
}

> aprun -n4 ./a.out
Hello from odd image 3
Hello from even image 0
Hello from even image 2
Hello from odd image 1
```

8.4.2 Barriers

A `sync_all()` ensures that all images must execute a `sync_all()` before any image may proceed beyond the `sync_all()` which it executed. It is not required that all images execute exactly the same `sync_all()` in the source code, just that they must execute some `sync_all()`. Failure of all images to participate will cause deadlock.

8.4.3 Function Calls

A coarray may be passed to a function via a reference or a pointer, but may not be passed by value. If a coarray could be passed by value, the call would have to be collective. There would be a collective allocation of a temporary coarray, the data within the original coarray would need to be copied into the temporary coarray, and eventually the temporary coarray would need to be collectively destroyed. Pass by value is expensive and there are better alternatives, like passing a coarray as a const reference, so it is a compile-time error. No matter how a coarray parameter is declared, the type of the actual argument must agree. Automatic conversions are provided between bounded and unbounded arrays; a conversion from unbounded to bounded performs a run-time check to ensure that the extents match and may throw a `mismatched_extent_error` exception.

8.4.3.1 `coatomic`

The `coatomic` template is similar to the C++11 `std::atomic` template, but provides operations that are atomic with respect to images rather than threads. Specializations exist for all basic types and the same operations are supported as for the C++11 `std::atomic` template. Similar convenience typedefs are provided as well so that, for example, `coatomic_long` can be used in place of `coatomic<long>`.

```
coarray< coatomic<long> > x; // or coarray<coatomic_long>

x(i) ^= 3; // atomic update x = x ^ 3 on image i

long old_value = x(i)++; // atomic increment, saving the old value

long new_value = ++x(i); // atomic increment, saving the new value
```

8.4.3.2 `coevent`

A *coevent* permits point-to-point synchronization between images. It wraps a `coatomic_long` that acts as a counter and provides two operations, `post` and `wait`. `Post` atomically increments the counter and `wait` blocks execution of the calling image until it can atomically decrement the counter to a non-negative value.

```
coarray<coevent> x;

if ( this_image() == 0 ) {
    // do something, then notify image 1
    x(1).post();
}
else if ( this_image() == 1 ) {
    // wait for notification from another image
    x().wait();    // then do something
}
```

8.4.3.3 comutex

A comutex provides mutual exclusion. The lock function blocks until the mutex can be acquired and the unlock function releases the mutex. The try_lock function attempts to acquire the lock and returns a true upon success.

```
coarray<comutex> m;  
  
m(i).lock();  
// critical section, typically guarding access to data on image i  
m(i).unlock();
```

8.5 Collectives

Coarray C++ provides broadcast and reduction collectives.

8.5.1 cobroadcast

cobroadcast replicates the value of a coarray on one image across all other images.

```
#include <cassert>  
#include <iostream>  
#include <coarray_cpp.h>  
  
using namespace coarray_cpp;  
  
int  
main( int argc, char* argv[] )  
{  
    coarray<int> x;  
    size_t image = this_image();  
    size_t n      = num_images();  
    if ( image == 0 ) {  
        x = 42;  
    }  
    sync_all();  
  
    // Make x on every image equal the x on image 0.  
    cobroadcast( x, 0 );  
    sync_all();  
    assert( x == 42 );  
    return 0;  
}
```

8.5.2 coreduce

`coreduce` applies a function across the coarray values of all images. For convenience, template specializations of `coreduce` are provided for the addition, min, and max operations from the C++ functional header. Implementations are likely to provide optimized versions of at least these reductions.

```
#include <cassert>
#include <iostream>
#include <coarray_cpp.h>
using namespace coarray_cpp;
int
main( int argc, char* argv[] )
{
    coarray<int> sum;
    coarray<int> min;
    coarray<int> max;
    size_t image = this_image();
    size_t n      = num_images();
    sum = image;
    min = image;
    max = image;

    sync_all();

    cosum( sum ); // equivalent to coreduce( sum, std::plus<int> )
    comin( min ); // equivalent to coreduce( min, std::less<int> )
    comax( max ); // equivalent to coreduce( max, std::greater<int> )
    sync_all();

    assert( sum == ( n * ( n - 1 ) / 2 ) );
    assert( min == 0 );
    assert( max == ( n - 1 ) );

    return 0;
}
```

8.6 Exceptions

Coarray C++ throws standard C++ exceptions, like `std::bad_cast`, but also throws some special exceptions for coarray-specific errors.

- `invalid_image_error`

This exception is thrown whenever a cosubscript is invalid. For example, given a coarray `x` in a program executed with 4 images, `x(4)` triggers an exception because the only valid image numbers are 0, 1, 2, and 3.

- `invalid_put_error`

This exception is thrown whenever a user-defined type is copied to a different image, but that type has `coarray_traits` that specify that it is not trivially puttable.

- `mismatched_extent_error`

This exception is thrown when two arrays in an array assignment have a different shape.

- `mismatched_image_error`

This exception is thrown when two copointers are compared or subtracted, but the copointers point to objects on different images.

8.7 Memory Consistency Model

8.7.1 `atomic_image_fence()`

The `atomic_image_fence()` function is the Coarray C++ equivalent of the C++11 `std::atomic_thread_fence()` function. It has the same behavior with respect to images as `std::atomic_thread_fence()` has with respect to threads. Typically, it is used to ensure that all memory accesses made by the calling image are visible to all images before performing subsequent memory accesses.

8.7.2 Accesses within a Single Image

The effect of two memory accesses made by an image to its own address space is governed by the C++ memory consistency model. The C++ memory consistency model depends on which version of the C++ standard is implemented by the compiler. In general, a C++03 compiler assumes that an image is single-threaded and offers no memory consistency guarantees if multiple threads perform the accesses, whereas a C++11 compiler provides a detailed memory consistency model that can be used to reason about the effect of memory accesses within a multithreaded image.

8.7.3 Accesses to Other Images

8.7.3.1 Multi-byte Accesses

A memory access of an object of size N bytes shall be treated as if it was performed as N arbitrarily ordered single-byte memory accesses. For example, the target image of a write shall not rely on the N th byte being written last to detect whether the full object has been written.

8.7.3.2 From Different Images

The execution of a program contains a data race if it contains two conflicting actions in different images, at least one of which is not atomic, and neither happens before the other. Any such data race results in undefined behavior. For example, if two images both write to the same object without any synchronization:

```
if ( this_image() == 0 ) {
    x(i) = 0;
} else if ( this_image() == 1 ) {
    x(i) = 1;
}
```

then the final value of the object is undefined. Various forms of synchronization can impose a specific order, such as in this example:

```
if ( this_image() == 0 ) {
    x(i) = 0;
}

sync_all();

if ( this_image() == 1 )
{
    x(i) = 1;
}
```

where the assignment by image 0 happens before the assignment by image 1 because of the `sync_all()`.

Two atomic operations issued by different images to the same coatomic object have the same ordering relationship as two C++11 threads that perform the same atomic operations on the same object.

8.7.3.3 From the Same Image

Two memory accesses issued by the same image to non-conflicting memory addresses are unordered.

Two memory accesses issued by the same image to conflicting memory addresses within the address space of a single, different image shall have the same order as if they were made within the issuing image's address space. For example, in the following code:

```
x(i) = 1; int y = x(i);
```

the value of `y` will be 1 provided that there are no data races. Therefore, a Coarray C++ implementation for a shared memory system could inline `x(i)` as a direct memory access, allowing the compiler to make the following optimization (forward substitution):

```
x(i) = 1; int y = 1;
```

Note: For distributed memory systems, providing this ordering guarantee is unfortunately somewhat onerous, but it is consistent with ordering guarantees of other PGAS languages, namely UPC and Fortran. Two memory accesses issued by an image to the same distant memory location typically will pass through the issuing processor's memory system, a high-speed communication network, and finally the target processor's memory system. Each hardware component is likely to contain multiple data pathways to increase bandwidth and resiliency, such that two memory accesses traveling on different pathways could bypass each other. Providing the ordering guarantee may require constraining two memory accesses to the same target location to always take the same hardware path to prevent bypass. Alternatively, software can track outstanding memory accesses and defer issuing an access if there is a conflict; however, software ordering adds overhead to each memory access to check for conflicts as well as storage overhead to track the accesses

8.8 Blocking Versus Non-blocking Accesses

When an image makes a blocking read or write access, it does not proceed to execute its next operation until the access fully completes. By contrast, a non-blocking read or write access permits an image to proceed to execute its next operation before the access fully completes and provides some mechanism for ensuring that the operation has completed later.

8.8.1 Writes (Puts)

Neither the target image nor any other image besides the issuing image is required to be able to observe the effects of a write until some form of image synchronization occurs. Therefore, an implementation is permitted to issue non-blocking writes for all writes provided that it can ensure that conflicting accesses issued by the same image occur in program order. Whether this guarantee is provided by software or hardware depends on the implementation. To explicitly issue and manage completion of a non-blocking write, see [Cofutures on page 151](#).

8.8.2 Reads (Gets)

A Coarray C++ read access is blocking in order to provide a value for use in an arbitrary expression context:

```
coarray<int> x;
...
int y = x(i) + 1; // read of x(i) shall block
```

A non-blocking read is performed via an explicit `get()` member function of `coref`:

```
    int y;x(i).get( &y );
...           // some code that does not access y
atomic_image_fence();
++y;
```

The `get()` member function issues a non-blocking read that is not guaranteed to complete until the next fence. The `atomic_image_fence()` ensures completion of all previously issued memory accesses. The `get()` plus fence solution is appropriate in many cases, but it may be too broad if the fence would force completion of other accesses on which the issuing image does not yet need to wait. To explicitly issue and manage completion of a non-blocking read, see [Cofutures on page 151](#).

8.8.3 Cofutures

Coarray C++ provides explicit completion management of a non-blocking access via a `cofuture`, which is modeled on C++11's `std::future`. A `coref` plays a similar role to C++11's `std::promise`, providing member functions that create a `cofuture`. Here is an example of a non-blocking read where the storage for the value is contained within the `cofuture`. The value cannot be accidentally used before the operation has completed, but existing storage cannot be used as the target of the read:

```
coarray<int> x;
...
cofuture<int> f = x(i).get_cofuture(); // or just x(i)
...
int z = f + 1; // using f waits then implicitly returns the value
```

For convenience, a `coref` can automatically convert to a `cofuture` so that the `get_cofuture()` call can be omitted. Here is an example of a non-blocking read where the storage for the value is external to the `cofuture`. Care must be taken to not access the storage until `wait()` has been called:

```
coarray<int[100]> x;
int y[100];
...
cofuture<void> f = x(i).get_cofuture( y );
...           // code that does not read or write y
f.wait();
...           // code that reads or writes y
```

Note that the `cofuture`'s parameter type is `void` because it does not store any value.

Here is an example of a non-blocking write. Care must be taken to not overwrite the source of the write until `wait()` has been called.

```
coarray<int> x;  
int y;  
...  
cofuture<void> f = x(i).put_cofuture( y );  
... // code that does not write y  
f.wait(); // ensure that the x(i) = y assignment completed
```

Note that the `cofuture`'s parameter type is `void` because a `cofuture` for a write never stores a value.

8.9 Code Patterns

8.9.1 Coobjects

When a `coarray` is included as a member of a class, it can be allocated with the class object or it can be allocated later:

```
// An X must be allocated and destroyed  
// collectively because it contains a coarray.  
class X {  
    coarray<int> x;  
    ...  
};  
  
// But a Y defers its "collectiveness" until  
// it needs to allocate the coarray.  
class Y {  
    coarray<int>* y;  
    ...  
};
```

These two options provide flexibility for implementing collective objects, or `coobjects`, which can encapsulate `coarray` data movement.

8.9.2 Hoisting a `coptr`

When a coarray of pointer type is accessed within a loop, there may be unnecessary reads of the pointer from the target image if the same image is accessed repeatedly:

```
coarray<int*> x;
...
for ( int i = 0; i < n; ++i ) {
    int y = x(1)[i]; // reads pointer x(1) each time
    ...
}
```

A `coptr` or `const_coptr` can be used to hoist the read of the pointer:

```
coarray<int*> x;
...
const_coptr<int> p = x(1)[0].address(); // reads pointer x(1) once
for ( int i = 0; i < n; ++i ) {
    int y = p[i];
    ...
}
```


Using Cray C Extensions [9]

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 [Cray C and Fortran Example on page 177](#).

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
```


Using Predefined Macros [10]

The macros listed in this chapter are the Cray-specific predefined macros. To see the entire list of predefined macros, add `-Wp, -list_final_macros` to your `cc` command line. For example, if you have the file `c.c`, specify:

```
% cc -Wp,-list_final_macros c.c > out
```

Predefined macros can be divided into the following categories:

- Macros required by the C and C++ standards ([Macros Required by the C and C++ Standards on page 159](#))
- Macros based on the host machine ([Macros Based on the Host Machine on page 160](#))
- Macros based on the target machine ([Macros Based on the Target Machine on page 160](#))
- Macros based on the compiler ([Macros Based on the Compiler on page 161](#))
- UPC macros ([UPC Predefined Macros on page 161](#))

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.

Note: Any of the predefined macros except those required by the standard (see [Macros Required by the C and C++ Standards on page 159](#)) 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
<code>__TIME__</code>	Time of translation of the source file.
<code>__DATE__</code>	Date of translation of the source file.
<code>__LINE__</code>	Line number of the current line in your source file.

Macro	Description
<code>__FILE__</code>	Name of the source file being compiled.
<code>__STDC__</code>	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.
<code>__cplusplus</code>	Defined as 1 when the compiling Cray C++ code and undefined when compiling Cray C code. The <code>__cplusplus</code> 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
<code>__linux</code>	Defined as 1.
<code>__linux__</code>	Defined as 1.
<code>linux</code>	Defined as 1.
<code>__gnu_linux__</code>	Defined as 1.

10.3 Macros Based on the Target Machine

The following macros provide information about the characteristics of the target machine:

Macro	x86	AVX
<code>_ADDR64</code>	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.	
<code>__LITTLE_ENDIAN__</code>	Defined as 1.	
<code>_LITTLE_ENDIAN</code>	Defined as 1.	
<code>_MAXVL_8</code>	Defined as 16, the number of 8-bit elements that fit in an XMM register ("vector length").	Defined as 32, the number of 8-bit elements that fit in a YMM register ("vector length").
<code>_MAXVL_16</code>	Defined as 8.	Defined as 16.

Macro	x86	AVX
<code>_MAXVL_32</code>	Defined as 4.	Defined as 8.
<code>_MAXVL_64</code>	Defined as 2.	Defined as 4.
<code>_MAXVL_128</code>	Defined as 0.	Defined as 2.

10.4 Macros Based on the Compiler

The following macros provide information about compiler features:

Macro	Description
<code>_RELEASE_MAJOR</code>	Defined as the major release level of the compiler.
<code>_RELEASE_MINOR</code>	Defined as the minor release level of the compiler.
<code>_RELEASE_STRING</code>	Defined as a string that describes the version of the compiler.
<code>_CRAYC</code>	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
<code>__UPC__</code>	The integer constant 1, indicating a conforming implementation.
<code>__UPC_DYNAMIC_THREADS__</code>	The integer constant 1 in the dynamic THREADS translation environment.
<code>__UPC_STATIC_THREADS__</code>	The integer constant 1 in the static THREADS translation environment.

Running C and C++ Applications [11]

To run applications, log in to a login node and set up your user environment. See the *Cray Application Developer's Environment User's Guide* for details on setting up your environment. In your working directory, load the appropriate modules, compile your programs, and launch them using the `aprun` command.

To use the Cray C compiler, load the `PrgEnv-cray` module. Use the `module list` command to get a list of currently loaded modules. If another Programming Environment module is loaded, use the `module swap` command. For example, if `PrgEnv-pgi` is loaded, use this command:

```
% module swap PrgEnv-pgi PrgEnv-cray
```

Then use the `cc -V` command to verify that the Cray C compiler is available.

Compile your application.

```
% cc -o simple simple.c
```

Move your application to a mount point on the Cray system to execute.

```
% aprun -n 4 ./simple | sort
Application 1024906 resources: utime 0, stime 0
hello from pe 0 of 4
hello from pe 1 of 4
hello from pe 2 of 4
hello from pe 3 of 4
```

If you specified the `-X` option on the `cc` command line, then the `aprun -n` option must specify the same number of processing elements (`npes`). Otherwise, you will receive a run time error.

For additional information, see the *Cray Programming Environment User's Guide*.

Debugging Cray C and C++ Code [12]

The TotalView symbolic debugger is available to help you debug C and C++ codes. 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 TotalView debugger. For more information about these compiler options, see [-G level and -g on page 52](#).
- 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 [-h \[no\]bounds \(cc\) on page 52](#). The `#pragma _CRI [no]bounds` directive is described in [\[no\]bounds Directive on page 76](#).
- The `-G3` option optimizes code for use with Cray fast-track debugging and requires use of a debugger that supports fast-track debugging. For more information, see the `lgdb(1)` man page.
- 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 [message Directive on page 80](#).
- 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 [\[no\]opt Directive on page 80](#).

12.1 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

Compiler options control the trade-offs between ease of debugging and compiler optimizations. The compiler produces internal debugger information (DWARF) at all times. The DWARF data provides function and line information to debuggers for tracebacks and breakpoints, as well as type and location information about data variables.

These options are specified as follows:

- `-G3` This option permits both full code optimization and the greatest flexibility in setting breakpoints, but requires use of the Cray fast-track debugger. For more information, see the `lgdb(1)` man page.

- `-G2`

With no DWARF, the executable is optimized and as small as possible, but cannot be easily debugged. Only assembly instructions will be visible and only global symbols will be available.

- `-G1`

With partial DWARF and at least some optimization, tracebacks and limited breakpoints are available in the debugger. The source code will be visible and many more symbols will be available. The executable will be somewhat slower and larger in exchange for increased debugger functionality.

- `-g` or `-G0`

With full DWARF and no optimizations, full debugging will be available, but at the cost of a slower and larger executable.

Note: The `-g` or `-G` options may be specified on a per file basis so that only part of an application incurs the overhead of improved debugging.

However, consider the following cases in which optimization is affected by the `-G1` and `-G2` 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 `-G1` option is specified.
- When the `-G1` 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.

Using Interlanguage Communication [13]

The C and C++ compilers provide mechanisms for declaring external functions written in other languages. This enables you to write portions of an application in C, C++, Fortran, or assembly language, which can be useful in cases where the other languages provide performance advantages or utilities not available in C or C++. The calling sequence is described in detail on the `callseq(3)` man page.

13.1 Calls Between C and C++ Functions

The following requirements apply 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 mangles 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 prevents 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 using 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\t p1      = %d\n", p1);
    printf("\t\t p2      = %d\n", p2);
    printf("\t\t global_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 Fortran Functions and Subroutines from C or C++

The following standard considerations apply when calling Fortran functions from C or C++. In addition, new interoperability features are supported under the more recent Fortran standards. These newer features are described in [Calling a C or C++ Function from Fortran on page 180](#).

13.2.1 Requirements

- 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. For more information, see [Argument Passing on page 171](#).
- 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. For more information, see [Array Storage on page 172](#).
- 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. For more information, see [Logical and Character Data on page 173](#).
- Fortran logical values and the Boolean values from C and C++ are not fully compatible. For more information, see [Logical and Character Data on page 173](#).
- 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 [fortran Keyword on page 155](#)). The `fortran` keyword is not available in Cray C++. Instead, Fortran functions must be declared by specifying `extern "C"`.

13.2.2 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.

13.2.3 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++:

<u>Fortran</u>	<u>C or C++</u>
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]

```

13.2.4 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:

<u>Macro</u>	<u>Description</u>
<code>_btol</code>	Conversion utility that converts a 0 to a Fortran logical <code>.FALSE.</code> and a nonzero value to a Fortran logical <code>.TRUE.</code>
<code>_ltob</code>	Conversion utility that converts a Fortran logical <code>.FALSE.</code> to a 0 and a Fortran logical <code>.TRUE.</code> to a 1.

13.2.5 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++ 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 Fortran common block `ST` will be converted to lower case with a trailing underscore added.

The name of the structure and the converted Fortran common block name must match. The C and 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:

```
#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:

```
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
```

13.2.6 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.
```

13.2.7 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
```

13.2.8 Calling a Fortran Program from Cray C++

The following example illustrates how a Fortran program can be called from a Cray C++ program:

```
#include <iostream>
using namespace std;
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
```

13.3 Calling a C or C++ Function from Fortran

In addition to calling Fortran functions and subroutines from C or C++ programs, C or C++ functions can be called from Fortran. Two methods can be used: either the [Standard Fortran/C Interoperability on page 183](#) or the [Portable Interoperability Mechanism on page 180](#), which provides a standard portable interoperability mechanism for Fortran and C programs.

13.3.1 Portable Interoperability Mechanism

If you use the method documented in this section, keep in mind the information in [Calling Fortran Functions and Subroutines from C or C++ on page 170](#).

When calling a Cray C++ function from a Fortran program, observe the following rules:

- The Cray C++ function must be declared with `extern "C"` linkage.
- The program must be linked using the `CC` command.
- The program's main routine must be C or C++ code compiled using 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, and will not work if the `-s default64` option is enabled.

Example 14. Calling a C function from Fortran

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
```

13.3.2 Standard Fortran/C Interoperability

For more information about C interoperability, see the current Fortran standard.

The ISO_C_BINDING module provides interoperability between Fortran intrinsic types and C types. The ISO_C_BINDING module provides named constants which can be used as KIND type parameters, compatible with C types.

In addition to the named constants required by the Fortran 2003 standard, Cray compiler provides, as an extension, definitions for 128-bit floating, and complex types. C_FLOAT128 and C_FLOAT128_COMPLEX correspond to C types `__float128` and `__float128 complex`.

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, and define implementation-defined behavior as behavior that is dependent on the characteristics of the implementation for a correct program construct and correct data.

14.1 Messages

All diagnostic messages issued by the compilers are reported through the Cray Linux Environment (CLE) message system. For information about messages issued by the compilers and for information about the Cray Linux Environment (CLE) message system, see [Appendix C, Using the Compiler Message System on page 209](#).

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 about how the words in the argument list are formed, refer to the documentation on the shell in which you are running. For information about Cray Linux Environment (CLE) 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` 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 11](#) 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 11](#). Thus, the `sizeof` operator will return a size that is equal to the value in the *Memory* column of [Table 11](#) divided by 8 (the number of bits in a byte).

Table 11. Data Type Mapping

Type	Representation Size and Memory Storage Size (bits)
<code>bool</code> (C++)	8
<code>_Bool</code> (C)	8
<code>char</code>	8
<code>wchar_t</code>	32
<code>short</code>	16
<code>int</code>	32
<code>long</code>	64
<code>long long</code>	64
<code>float</code>	32
<code>double</code>	64
<code>long double</code>	64
<code>float complex</code>	64 (each part is 32 bits)
<code>double complex</code>	128 (each part is 64 bits)
<code>long double complex</code>	128 (each part is 64 bits)
<code>_float128</code>	128

Type	Representation Size and Memory Storage Size (bits)
<code>_float128 complex</code>	256 (each part is 128 bits)
Pointers	64

Note: Vectorization of 8- and 16-bit data types is deferred.

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 in identifier names. For more information about the `-h [no]calchars` option, see [-h \[no\]calchars on page 35](#).

A character consists of 8 bits. Up to 8 characters can be packed into a 64-bit word. A plain `char` type (that is, one that is declared without a `signed` or `unsigned` keyword) is treated as a signed 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 12. Packed Characters

Character Constant	Integer Value
<code>'a'</code>	<code>0x61</code>
<code>'ab'</code>	<code>0x6162</code>

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 13. Unrecognizable Escape Sequences

Character Constant	Integer Value	Explanation
<code>'\a'</code>	<code>0x7</code>	Recognized as the ASCII BEL character
<code>'\8'</code>	<code>0x38</code>	Not recognized; ASCII value for 8
<code>'\['</code>	<code>0x5b</code>	Not recognized; ASCII value for [
<code>'\c'</code>	<code>0x63</code>	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 `mbtowl()` function. The current locale in effect at the time of compilation determines the method by which `mbtowl()` 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 two's complement format. For representation and memory storage requirements for integral types, see [Table 11](#).

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 that `e1` is right-shifted `e2` bit positions; vacated bits are filled with 1s. This behavior can be modified by using the `-h nosignedshifts` option (see [-h \[no\]signedshifts on page 35](#)). Bits higher than the sixth bit are not ignored.

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.

14.2.6 128-Bit Floating Point and 256-Bit Complex Predefined Types

The Cray C and C++ Compilers now support 128-bit floating point and 256-bit complex predefined types using the X86-64 ABI definitions for type names and data layout. These types are sometimes referred to as *quad-precision*. In C and C++, use `__float128`, and `__float128 _complex`. The header file `quadmath.h` defines the 128-bit functions and the header file `complex.h` defines the complex functions.

The base type itself uses 128 bits of storage with a guaranteed minimum alignment on a 128-bit boundary, little endian, has a 15-bit exponent, a 113-bit mantissa, and an exponent bias of 16383, and is compatible with the `gcc` implementation.

In C and C++, `long double` remains identical to `double` – 64-bit IEEE, and not 80-bit extended precision.

C forms of intrinsic math functions offer full support for quad-precision types. See the `intro_quad_precision(3i)` man page for a complete list of intrinsic functions that support quad-precision.

There is no `printf` descriptor for `__float128`. `quadmath.h` defines `gnu` functions `quadmath_snprintf()` and `strtouflt128()` that convert `__float128` to strings:

```
extern __float128 strtouflt128 (const char *, char **);
extern int quadmath_snprintf (char *str, size_t size, const char *format, ...)
```

Alternatively, use a Fortran subroutine to print a 128-bit floating point number:

```
% cat printf128main.c printf128.f
#include <quadmath.h>
#include <math.h>
printf128_(__float128 *x);
void main(){
    __float128 x=1.234567890123456789012345678901234567890Q;
    printf128_(&x);
}

subroutine printf128(qx)
real*16 qx
print "('qx:', f5.1, e45.35)",qx,qx
end

%cc printf128main.c printf128.o && ./a.out
qx: 1.2 0.123456789012345678901234567890123456789012346E+01
```

14.2.7 Arrays and Pointers

An unsigned `long` 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 `long` 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. Type-casting from pointer to `long int` enables pointer arithmetic. For example:

```
static void **Table;
size_t offset = -BlockSize[nr];
Table = (void **) malloc(MAXBLOCKS * sizeof(void *));
Table[i] = (void **) (((long) Table[i]) + offset);
```

Pointers on Cray Linux Environment (CLE) 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.

14.2.8 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.

14.2.9 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 a `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++.

14.2.10 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.

14.2.11 Declarators

A maximum of 12 pointer, array, and/or function declarators are allowed to modify an arithmetic, structure, or union type.

14.2.12 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.

14.2.13 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.

14.2.14 System Function Calls

For a description of the form of the unsuccessful termination status that is returned from a call to `exit(3)`, see the `exit(3)` man page.

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 [-I *includir* on page 60](#).

The source file character sequence in a `#include` directive must be a valid or Cray Linux Environment (CLE) 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, refer to the description of the predefined macros in [Chapter 10, Using Predefined Macros on page 159](#).

The `#pragma` directives are described in [Chapter 3, Using #pragma Directives on page 73](#).

Using Libraries and the Linker [A]

This appendix describes the libraries that are available with the Cray C and C++ compilers and the linker.

A.1 Cray C and C++ Libraries

Libraries that support Cray C and C++ are automatically available when you use the `CC` or `cc` command to compile your programs. These commands automatically issue the appropriate directives to link 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 linking process, more knowledge of the linker 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.

A.2 Linker

When you use the `cc` or `CC` command to invoke the compiler, and the program compiles without errors, the linker is called. Specifying the `-c` option on the command line produces relocatable object files (`*.o`) without calling the linker. These relocatable object files can then be used as input to the linker command by specifying the file names on the appropriate linker 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 linker 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` command should be linked using the `CC` command.

Using Cray C and C++ Dialects [B]

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.

B.1 C++ Language Conformance

The Cray C++ compiler accepts the C++ language as defined by the *ISO/IEC 14882:2003* standard, except for exported templates. C++ supports the ISO 2003 Standard Template Library (STL) headers but abrogates support for pre-standard template headers that have the `.h` extension.

Note: C++ codes that use the pre-standard template headers must be updated to the ISO C++ standard.

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 [General Directives on page 75](#)). Command line options are also available to enable and disable anachronisms (see [C++ Anachronisms Accepted on page 198](#)) and strict standard-conformance checking (see [Extensions Accepted in Normal C++ Mode on page 199](#), and [Extensions Accepted in C or C++ Mode on page 200](#)). The command line options are described in [Chapter 2, Invoking the C and C++ Compilers on page 25](#).

B.1.1 Supported C++ Language Features

The following features, which are in the ISO/IEC 14882:2003 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

¹ As defined in *The Annotated C++ Reference Manual (ARM)*, by Ellis and Stroustrup, Addison Wesley, 1990.

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.
- 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.
- `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.
- 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.
- Type qualifiers, `const` and `volatile` referred to as *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.
- 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` 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.

B.2 C++ Anachronisms Accepted

C++ anachronisms are enabled by using the `-h anachronisms` command line option (see [-h \[no\]anachronisms \(CC\) on page 29](#)). 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 non-nested 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.

B.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";
```

B.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 (\$) 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 [fortran Keyword on page 155](#).
- Cray hexadecimal floating point constants. For more information, see [Hexadecimal Floating-point Constants on page 156](#).

B.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 entered as "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::* pm1 = &B::i; // error, OK in cfront mode
    int D::* pm2 = &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 `xyz`, 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.

Using the Compiler Message System [C]

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.

C.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` 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 indicated octal constant are outside of this range. Change each digit in the octal constant to be within the valid range.

C.2 Controlling the Use of Messages

This section summarizes the command line options that affect the issuing of messages from the compiler.

C.2.1 Command Line Options

Option	Description
-h errorlimit[= <i>n</i>]	Specifies the maximum number of error messages the compiler prints before it exits.
-h [no]message= <i>n</i> [:...]	Enables or disables the specified compiler messages, overriding -h msglevel.
-h msglevel_ <i>n</i>	Specifies the lowest severity level of messages to be issued.
-h [no]msgs	Enables or disables the writing of optimization messages to <code>stderr</code> .
h [no]negmsgs	Enables or disables the writing of messages to <code>stderr</code> that indicate why optimizations such as vectorization, inlining, or cloning did not occur in a given instance.
-h report= <i>args</i>	Generates optimization report messages.

C.2.2 Environment Options for Messages

The following 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.

C.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` 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 linker, 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`.

C.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.

Category	Meaning
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 that a bug report 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 bug report 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.
OPTIMIZATION	Information about general optimizations.
IPA_INFO	Information about interprocedural optimizations.

C.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).

Using Intrinsic Functions [D]

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 most of the intrinsic functions, the Cray C++ compiler requires that either the `intrinsic.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 `intrinsic.h` statement and you cannot modify the code, you can use the `-h prototype_intrinsic` 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 nointrinsic` command line option causes the compiler to treat these calls as regular function calls and not as intrinsic function calls.

There are built-in atomic memory intrinsic functions of the form `__sync_*` that do not require an include file nor any explicit declaration.

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. For details on whether it has a vector version, refer to the appropriate intrinsic function man page.

The following sections group the C and C++ intrinsics according to function and provides a brief description of each intrinsic in that group. For more information, see the corresponding man page.

D.1 Atomic Memory Operations

Atomic memory operations (AMOs), unlike other functions, cannot be interrupted by the system and can allow multiple threads to safely modify the same variable under certain conditions. The AMO intrinsics allow you to add, subtract, AND, NAND, OR, and XOR values together, or compare and swap values.

Local AMOs operate on variables in the processor's local memory (cache domain); they do not use the network interface to access memory. Multiple threads using local atomic memory operations to access the same variable need to be running within the same processor cache domain, which implies that they must be running on the same node. Local AMOs are atomic with respect to each other. The compiler issues an error message if a user tries to apply a local AMO intrinsic to a Unified Parallel C or shared variable or Fortran coarray that is not local to the current thread.

Global AMOs use the network interface to access variables in memory. The variables may or may not be in the processor's local cache domain. Global AMOs are atomic with respect to each other. Global AMOs are used to modify a Unified Parallel C (UPC) shared variable or Fortran coarray and are available only when compiling UPC (`-hupc`) or coarray Fortran (`-hcaf`).

A global AMO uses a different mechanism for achieving atomicity than a local AMO, so local and global AMOs are not atomic with respect to each other. Global and local AMOs should not be used concurrently on the same memory location, without synchronization.

It is possible to safely modify a variable using both atomic and non-atomic operations within a single UPC thread or Fortran image; however, if a thread or image modifies a variable with an atomic operation and a different thread or image concurrently modifies the same variable with a non-atomic operation, the result is indeterminate.

D.1.1 Local Atomic Memory Operations

The following functions, defined in `intrinsics.h`, perform various local atomic memory operations:

`__builtin_ia32_lfence`

(Load fence) Insures that all memory loads issued before this intrinsic are visible in memory before any future loads are executed.

`__builtin_ia32_sfence`

(Store fence) Insures that all memory stores issued before this intrinsic are visible in memory before any future stores are executed.

`__builtin_ia32_mfence`

(Memory fence) Insures that all memory stores and loads issued before this intrinsic are visible in memory before any future stores or loads are executed.

Functions built into the compiler do not require an include file, nor a specific compilation option for use. The following local atomic, built-in functions return the value of the object **before** the named operation occurs:

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. The allowed object and value types are signed and unsigned integer types of 1, 2, 4, or 8 bytes.

- The `__sync_fetch_and_add` function fetches the object pointed to by `ptr`, adds `value`, places the result into the object pointed to by `ptr`, and returns the old value of the object pointed to by `ptr`.
- The `__sync_fetch_and_sub` function fetches the object pointed to by `ptr`, subtracts `value`, places the result into the object pointed to by `ptr`, and returns the old value of the object pointed to by `ptr`.
- The `__sync_fetch_and_or` function fetches the object pointed to by `ptr`, ORs `value`, places the result into the object pointed to by `ptr`, and returns the old value of the object pointed to by `ptr`.
- The `__sync_fetch_and_and` function fetches the object pointed to by `ptr`, ANDs `value`, places the result into the object pointed to by `ptr`, and returns the old value of the object pointed to by `ptr`.
- The `__sync_fetch_and_xor` function fetches the object pointed to by `*ptr`, XORs `value`, places the result into the object pointed to by `ptr`, and returns the old value of the object pointed to by `ptr`.

- The `__sync_fetch_and_nand` function fetches the object pointed to by `ptr`, NANDs `value`, places the result into the object pointed to by `ptr`, and returns the old value of the object pointed to by `ptr`.

The following local atomic, built-in functions return the value of the object **after** the named operation occurs:

- The `__sync_add_and_fetch` function adds `value` to the object pointed to by `ptr` and returns the new value of the object pointed to by `ptr`.
- The `__sync_sub_and_fetch` function subtracts `value` from the object pointed to by `ptr` and returns the new value of the object pointed to by `ptr`.
- The `__sync_or_and_fetch` function ORs `value` with the object pointed to by `ptr` and returns the new value of the object pointed to by `ptr`.
- The `__sync_and_and_fetch` function ANDs `value` with the object pointed to by `ptr` and returns the new value of the object pointed to by `ptr`.
- The `__sync_xor_and_fetch` function XORs `value` with the object pointed to by `ptr` and returns the new value of the object pointed to by `ptr`.
- The `__sync_nand_and_fetch` function NANDs `value` with the current value of `ptr` and returns the new contents of `ptr`.
- The `__sync_val_compare_and_swap` function performs an atomic compare and swap. If the current value of `*ptr` is `compareValue`, then write `replacementValue` into `*ptr` and return the contents of `*ptr` before the operation.
- The `__sync_lock_test_and_set` function writes `value` into `*ptr`, and returns the previous contents of `*ptr`.

D.1.2 Global Atomic Memory Operations

Global atomic memory operations (global AMO) are typically used to atomically modify a Unified Parallel C (UPC) shared variable or Fortran coarray.

The target of a global AMO can be located in a different cache domain, so a global AMO is not atomic with respect to memory operations performed locally within the target's cache domain. Therefore, the application must use synchronization to ensure that global AMOs and local memory operations are not used concurrently on the same memory location.

The following intrinsics are defined in `intrinsics.h`. Functions without the `_upc` suffix accept both shared and non-shared pointers as the first argument. Functions with the `_upc` suffix accept only shared pointers as the first argument.

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.

- The `_amo_aadd` and `_amo_aadd_upc` functions (atomic add) add `value` to the object pointed to by `ptr`.
- The `_amo_aaddf` and `_amo_aaddf` functions (atomic add and fetch) add `value` to the object pointed to by `ptr` and return the new `value`.
- The `_amo_afadd` and `_amo_afadd_upc` functions (atomic fetch and add) add `value` to the object pointed to by `ptr` and return the old value of the object.
- The `_amo_aax` and `_amo_aax_upc` functions (atomic AND and XOR) AND the object pointed to by `ptr` with `andMask`, XOR the result with `xorMask`, and place the result into the object.
- The `_amo_afax` and `_amo_afax_upc` functions (atomic fetch and AND and XOR) AND the object pointed to by `ptr` with `andMask`, XOR the result with `xorMask`, place the result into the object, and return the old value of the object.
- The `_amo_aandf` and `_amo_aandf_upc` functions (atomic AND and fetch) AND the object pointed to by `ptr` with `value`, place the result into the object, and return the new value of the object.
- The `_amo_afand` and `_amo_afand_upc` functions (atomic fetch and AND) AND the object pointed to by `ptr` with `value`, place the result into the object, and return the old value of the object.
- The `_amo_anandf` and `_amo_anandf_upc` functions (atomic NAND and fetch) NAND the object pointed to by `ptr` with `value`, place the result into the object, and return the new value of the object.
- The `_amo_afnand` and `_amo_afnand_upc` functions (atomic fetch and NAND) NAND the object pointed to by `ptr` with `value`, place the result into the object, and return the old value of the object.
- The `_amo_aorf` and `_amo_aorf_upc` functions (atomic OR and fetch) OR the object pointed to by `ptr` with `value`, place the result into the object, and return the new value of the object.
- The `_amo_afor` and `_amo_afor_upc` functions (atomic fetch and OR) OR the object pointed to by `ptr` with `value`, place the result into the object, and return the old value of the object.
- The `_amo_axorf` and `_amo_axorf_upc` functions (atomic XOR and fetch) XOR the object pointed to by `ptr` with `value`, place the result into the object, and return the new value of the object.
- The `_amo_afxor` and `_amo_afxor_upc` functions (atomic fetch and XOR) XOR the object pointed to by `ptr` with `value`, place the result into the object, and returns the old value of the object.
- The `_amo_acswap` and `_amo_acswap_upc` functions (atomic compare and

swap) compare and swap a value by replacing the contents of the object pointed to by `ptr` with `replacementValue` if `compareValue` is equal to the object pointed to by `ptr` and return the old value of the object.

- The `_amo_aswap` and `_amo_aswap_upc` functions (atomic swap) swap a value by replacing the contents of the object pointed to by `ptr` with `replacementValue`. This function always returns the old value.
- The `_amo_aflush` and `_amo_aflush_upc` functions (atomic flush) force `*ptr` to be written to memory.

For more information, see the `amo(3i)` man page.

D.2 Bit Operations

The following intrinsic functions copy, count, or shift bits or computes the parity bit:

<code>_dshiftrl</code>	Move the left most <code>n</code> bits of an integer into the right side of another integer, and return that integer.
<code>_dshiftr</code>	Move the right most <code>n</code> bits of an integer into the left side of another integer and return that integer.
<code>_pbit</code>	Copies the rightmost bit of a word to the <code>nth</code> bit, from the right, of another word.
<code>_pbits</code>	Copies the rightmost <code>m</code> bits of a word to another word beginning at bit <code>n</code> .
<code>_poppar</code>	Computes the parity bit for a variable.
<code>_popcnt</code> <code>_popcnt32</code> <code>_popcnt64</code>	Counts the number of set bits in 32-bit and 64-bit integer words.
<code>_leadz</code> <code>_leadz32</code> <code>_leadz64</code>	Counts the number of leading 0 bits in 32-bit and 64-bit integer words.
<code>_gbit</code>	<code>_gbit</code> returns the value of the <code>nth</code> bit from the right.
<code>_gbits</code>	Returns a value consisting of <code>m</code> bits extracted from a variable, beginning at <code>nth</code> bit from the right.

D.3 Mask Operations

The following intrinsic functions create bit masks:

<code>_mask</code>	Creates a left-justified or right-justified bit mask with all bits set to 1.
<code>_maskl</code>	Returns a left-justified bit mask with <code>i</code> bits set to 1.
<code>_maskr</code>	Returns a right-justified bit mask with <code>i</code> bits set to 1.

D.4 Miscellaneous Operations

The following intrinsic functions perform various functions:

<code>_int_mult_upper</code>	Multiplies integers and returns the uppermost bits. For more information, see the <code>int_mult_upper(3i)</code> man page.
<code>_ranf</code>	Computes a pseudo-random floating-point number ranging from 0.0 through 1.0.
<code>_rtc</code>	Return a real-time clock value expressed in clock ticks.